

---

# **Rater Scoring Modeling Tool**

***Release 12.0.0***

**Nitin Madnani, Jeremy Biggs, and Anastassia Loukina**


**Mar 11, 2024**



# CONTENTS

<b>1</b>	<b>Documentation</b>	<b>3</b>
1.1	Who is RSMTool for? . . . . .	3
1.2	Overview of RSMTool Pipeline . . . . .	4
1.3	Evaluation Metrics . . . . .	7
1.4	Installation . . . . .	14
1.5	Tutorial . . . . .	15
1.6	Using RSMTool . . . . .	18
1.7	Advanced Uses of RSMTool . . . . .	41
1.8	Writing custom RSMTool sections . . . . .	89
1.9	Auto-generating configuration files . . . . .	92
1.10	API Documentation . . . . .	98
1.11	Utility Scripts . . . . .	170
1.12	Frequently Asked Questions . . . . .	170
1.13	Contributing to RSMTool . . . . .	172
1.14	Internal Documentation . . . . .	177
<b>2</b>	<b>Indices and tables</b>	<b>181</b>
	<b>Python Module Index</b>	<b>183</b>
	<b>Index</b>	<b>185</b>





Automated scoring of written and spoken responses is a growing field in educational natural language processing. Automated scoring engines employ machine learning models to predict scores for such responses based on features extracted from the text/audio of these responses. Examples of automated scoring engines include [MI Write](#) for written responses and [SpeechRater](#) for spoken responses.

RSMTTool is a python package which automates and combines in a *single pipeline* multiple analyses that are commonly conducted when building and evaluating automated scoring models. The output of RSMTTool is a comprehensive, customizable HTML statistical report that contains the outputs of these multiple analyses. While RSMTTool does make it really simple to run this set of standard analyses using a single command, it is also fully customizable and allows users to easily exclude unneeded analyses, modify the standard analyses, and even include custom analyses in the report.

We expect the *primary users* of RSMTTool to be researchers working on developing new automated scoring engines or on improving existing ones. Note that RSMTTool is not a scoring engine by itself but rather a tool for building and evaluating machine learning models that may be used in such engines.

The primary means of using RSMTTool is via the *command-line*.



## DOCUMENTATION

---

**Note:** If you use the [Dash](#) app on macOS, you can also download the complete RSMTTool documentation for offline use. Go to the Dash preferences, click on “Downloads”, then “User Contributed”, and search for “RSMTTool”.

---

### 1.1 Who is RSMTTool for?

We expect the primary users of RSMTTool to be researchers working on developing new automated scoring engines or on improving existing ones. Here’s the most common scenario.

A group of researchers already *has* a set of responses such as essays or recorded spoken responses which have already been assigned numeric scores by human graders. They have also processed these responses and extracted a set of (numeric) features using systems such as [Coh-Metrix](#), [TextEvaluator](#), [OpenSmile](#), or using their own custom text/speech processing pipeline. They wish to understand how well the set of chosen features can predict the human score.

They can then run an RSMTTool “experiment” to build a regression-based scoring model (using one of many available regressors) and produce a report. The report includes descriptive statistics for all their features, diagnostic information about the trained regression model, and a comprehensive evaluation of model performance on a held-out set of responses.

While they could use R, PASW (SPSS) or other tools to perform each of the RSMTTool analyses individually and compile a report himself, RSMTTool does all of this work for them with just a single command. Furthermore, the analyses included into the tool highlight educational measurement criteria important to building automated scoring models. If they wish, they can conduct further exploratory analysis using their preferred tools for data analysis by using the output of RSMTTool as a starting point.

RSMTTool is designed to be customizable:

1. The users can choose to either run all of the default analyses or select *only* the subset applicable to their particular study by changing the appropriate settings in a configuration file.
2. RSMTTool provides explicit support for adding *custom analyses* to the report if the user has some analysis in mind that is not already provided by RSMTTool. These analyses can then be automatically included in all subsequent experiments.

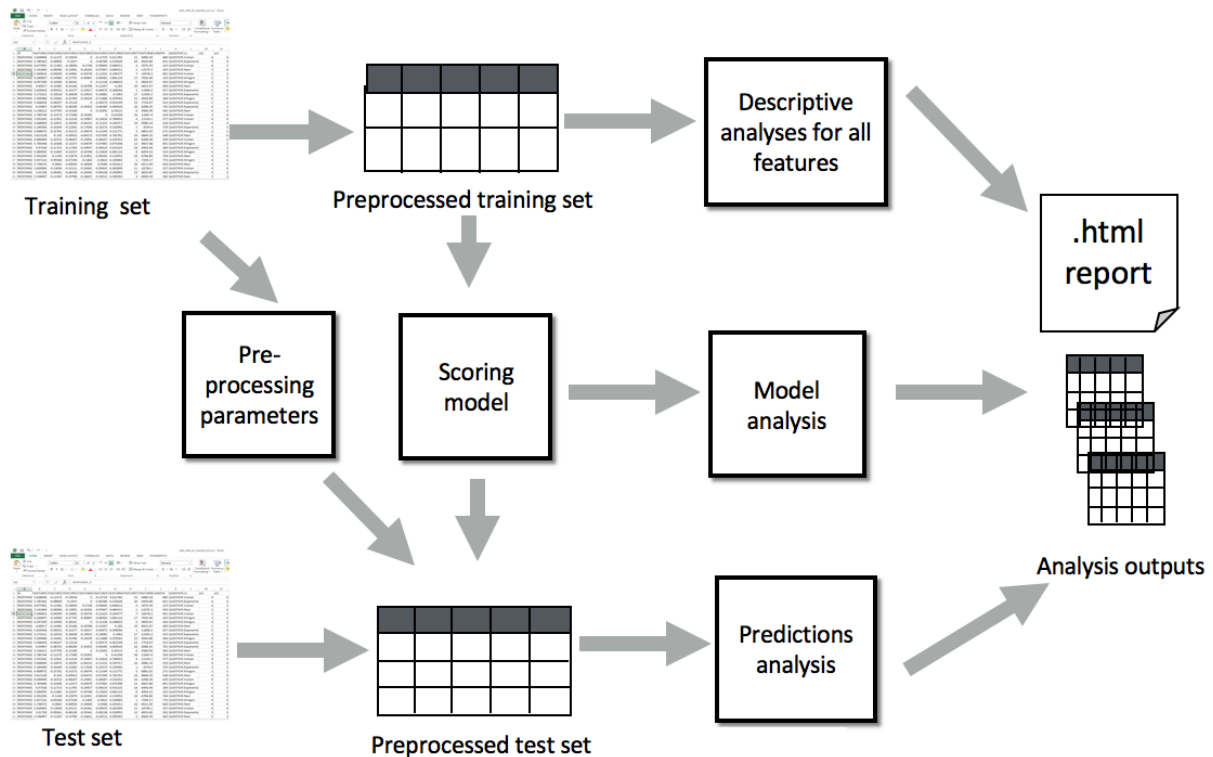
While training and evaluating a scoring model represents the most common use case for RSMTTool, it can do a lot more for advanced users such as:

- evaluating predictions obtained from *any* ML model using [rsmeval](#),
- generating predictions for new data from an `rsmttool` model using [rsmpredict](#),
- generating a detailed comparison between two different `rsmttool` models using [rsmcompare](#),
- generating a summary report for multiple `rsmttool` models using [rmssummarize](#),

- running cross-validation to estimate the performance of `rsmttool` models using `rsmxval`, and
- explaining the predictions of non-linear `rsmttool` models using `rsmexplain`.

## 1.2 Overview of RSMTTool Pipeline

The following figure gives an overview of the RSMTTool pipeline:



As its primary input, RSMTTool takes a *data file* containing a table with numeric, non-sparse features and a human scores for all responses, *pre-processes* them and lets you train a regression-based **Scoring Model** to predict the human score from the features. Available regression models include Ridge, SVR, AdaBoost, and Random Forests, among many others.

This trained model can then be used to generate scores for a held-out evaluation data whose feature values are pre-processed using the same *Pre-processing Parameters*. In addition to the raw scores predicted by the model, the **Prediction Analysis** component of the pipeline generates several additional *post-processed scores* that are commonly used in automated scoring.

The primary output of RSMTTool is a comprehensive, customizable HTML statistical report that contains the multiple analyses required for a comprehensive evaluation of an automated scoring model including descriptive analyses for all features, model analyses, subgroup comparisons, as well as several different *evaluation measures* illustrating model efficacy. More Details about these analyses can be found in this documentaion and in a separate *technical paper*.

In addition to the HTML report, RSMTTool also saves the intermediate outputs of all of the performed analyses as *CSV files*.



### 1.2.1 Input file format

The input files for the training and evaluation data should either be in a tabular format with responses as rows and features and score(s) in the columns in the `jsonlines` format with a JSON object per responses on each line. See below for a more detailed description of the `jsonlines` format.

RSMTTool supports input files in `.csv`, `.tsv`, `.sas7bdat`, `.xlsx`, or `.jsonlines` format. The format of the file is determined based on the extension. In all cases the *output files* will be saved in `.csv` format by default (see *file format* for other intermediate file output options).

For Excel spreadsheets, all data must be stored in the first sheet.

---

**Note:** RSMTTool 8.1 and higher no longer support `.xls` files. All Excel files must be in `.xlsx` format.

---

In a `jsonlines` format file, each line corresponds to a response and is represented as a dictionary with column names as the keys and column values as the values. For example:

```
{
  "id": "RESPONSE_1",
  "FEATURE1": 4.9380646013,
  "FEATURE2": -0.0584019308,
  "FEATURE3": -0.0846667513,
  "FEATURE4": -0.3167939755,
  "FEATURE5": -0.059868434,
  "FEATURE6": 4.6559139785,
  "FEATURE7": 16,
  "FEATURE8": -7940.9812058284,
  "LENGTH": 279,
  "QUESTION": "QUESTION_1",
  "L1": "Esperanto",
  "y": 3,
  "score2": 3
}
{
  "id": "RESPONSE_2",
  "FEATURE1": 5.3798973535,
  "FEATURE2": -0.0781006576,
  "FEATURE3": -0.1592030845,
  "FEATURE4": -0.2715376933,
  "FEATURE5": -0.1073346977,
  "FEATURE6": 4.7626728111,
  "FEATURE7": 10,
  "FEATURE8": -6683.4324801506,
  "LENGTH": 434,
  "QUESTION": "QUESTION_1",
  "L1": "Klingon",
  "y": 4,
  "score2": 4
}
```

Although RSMTTool does allow for nesting in the JSON objects on each line of a `jsonlines` format file, the top-level keys will be ignored when parsing the files and processing the data. Therefore, in the example below, the keys `x` and `metadata` will be ignored.

```
{
  "id": "RESPONSE_1",
  "y": 3,
  "x": {
    "FEATURE1": 4.9380646013,
    "FEATURE2": -0.0584019308,
    "FEATURE3": -0.0846667513,
    "FEATURE4": -0.3167939755,
    "FEATURE5": -0.059868434,
    "FEATURE6": 4.6559139785,
    "FEATURE7": 16.0,
    "FEATURE8": -7940.9812058284
  },
  "metadata": {
    "LENGTH": 279,
    "QUESTION": "QUESTION_1",
    "L1": "Esperanto",
    "score2": 3
  }
}
{
  "id": "RESPONSE_2",
  "y": 4,
  "x": {
    "FEATURE1": 5.3798973535,
    "FEATURE2": -0.0781006576,
    "FEATURE3": -0.1592030845,
    "FEATURE4": -0.2715376933,
    "FEATURE5": -0.1073346977,
    "FEATURE6": 4.7626728111,
    "FEATURE7": 10.0,
    "FEATURE8": -6683.4324801506
  },
  "metadata": {
    "LENGTH": 434,
    "QUESTION": "QUESTION_1",
    "L1": "Klingon",
    "score2": 4
  }
}
```

If the file contains nesting of more than two levels, the column names for nested records beyond the top level will be generated using `.` to separate levels: For example, given the JSON object `{'foo': {'bar': {'foo2': 0, 'foo3': 0}}}`, `foo` will be ignored and the columns will be named `bar.foo2` and `bar.foo3`.

### 1.2.2 Feature pre-processing

#### Data filtering

1. Remove all training and evaluation responses that have non-numeric for any of the features (see *column selection methods* for different ways to select features).
2. Remove all training and evaluation responses with non-numeric values for human scores.
3. Optionally remove all training and evaluation responses with zero values for human scores. Zero scored responses are usually removed since in many scoring rubrics, zero scores usually indicate non-scorable responses.

4. Remove all features with values that do not change across responses (i.e., those with a standard deviation close to 0).

### Data preprocessing

1. Truncate/clamp any outlier feature values, where outliers are defined as  $\mu \pm 4 * \sigma$ , where  $\mu$  is the mean and  $\sigma$  is the standard deviation.
2. Apply pre-specified *transformations* to feature values.
3. Standardize all transformed feature values into z-scores.
4. Flip the signs for feature values if necessary.

### Pre-processing parameters

Any held-out evaluation data on which the model is to be evaluated needs to be pre-processed in the same way as the training data. Therefore, the following parameters are computed on the training set, saved to disk, and re-used when pre-processing the evaluation set:

- Mean and standard deviation of raw feature values. These are used to compute floor and ceiling for truncating any outliers in the evaluation set;
- Any transformation and sign changes that were applied;
- Mean and standard deviation of transformed feature values. These are used to convert feature values in the evaluation set to z-scores.

### 1.2.3 Score post-processing

RSMTTool computes six different versions of scores commonly used in different applications of automated scoring:

#### raw

The raw predictions generated by the model.

#### raw\_trim

The raw predictions “trimmed” to be in the score range acceptable for the item. The scores are trimmed to be within the following range:  $score_{min} - tolerance$  and  $score_{max} + tolerance$ . Unless specified otherwise,  $score_{min}$  and  $score_{max}$  are set to the lowest and highest points on the scoring scale respectively and *tolerance* is set to 0.4998.

This approach represents a compromise: it provides scores that are real-valued and, therefore, provide more information than human scores that are likely to be integer-valued. However, it also ensures that the scores fall within the expected scale.

---

**Note:** In RSMTTool v6.x and earlier, the default value for *tolerance* was 0.49998 (note the extra “9”). Therefore, the *raw\_trim* values computed for outliers by RSMTTool v7.0 and onwards will be *different* from those computed by previous versions. If you wish to replicate results obtained with older versions, set the new *trim\_tolerance* field in the experiment configuration file to 0.49998.

---

## raw\_trim\_round

The raw\_trim predictions rounded to the nearest integer.

---

**Note:** The rounding is done using rint function from numpy. See [numpy documentation](#) for treatment of values such as 1.5.

---

## scale

The raw predictions rescaled to match the human score distribution on the training set. The raw scores are first converted to z-scores using the mean and standard deviation of the *machine scores* predicted for the training set. The z-scores are then converted back to “scaled” scores using the mean and standard deviation of the *human scores*, also computed on the training set.

## scale\_trim

The scaled scores trimmed in the same way as raw\_trim scores.

## scale\_trim\_round

The scale\_trim scores rounded to the nearest integer.

# 1.3 Evaluation Metrics

This section documents the exact mathematical definitions of the primary metrics used in RSMTTool for evaluating the performance of automated scoring engines. RSMTTool reports also include many secondary evaluations as described in *intermediary files* and the *report sections*.

The following conventions are used in the formulas in this section:

$N$  \- total number of responses in the *evaluation set* with numeric human scores and numeric system scores. Zero human scores are, by default, excluded from evaluations unless *exclude\_zero\_scores* was set to false.

$M$  \- system score. The primary evaluation metrics in the RSMTTool report are computed for *all* six types of *scores*. For some secondary evaluations, the user can choose between raw and scaled scores using the *use\_scaled\_predictions* configuration field for RSMTTool or the *scale\_with* field for RSMEval.

$H$  \- human score. The score values in *test\_label\_column* for RSMTTool or *human\_score\_column* for RSMEval.

$H2$  \- second human score (if available). The score values in *second\_human\_score\_column*.

$N_2$  \- total number of responses in the evaluation set where both  $H$  and  $H2$  are available and are numeric and non-zero (unless *exclude\_zero\_scores* was set to false).

$$\bar{M} \text{ \- Mean of } M ; \bar{M} = \sum_{n=1}^N \frac{M_i}{N}$$

$$\bar{H} \text{ \- Mean of } H ; \bar{H} = \sum_{n=1}^N \frac{H_i}{N}$$

$$\sigma_M \text{ \- Standard deviation of } M ; \sigma_M = \sqrt{\frac{\sum_{i=1}^N (M_i - \bar{M})^2}{N - 1}}$$

$$\sigma_H \text{ \textbackslash- \textbackslash- Standard deviation of } H ; \sigma_H = \sqrt{\frac{\sum_{i=1}^N (H_i - \bar{H})^2}{N - 1}}$$

$$\sigma_{H2} \text{ \textbackslash- \textbackslash- Standard deviation of } H2 ; \sigma_{H2} = \sqrt{\frac{\sum_{i=1}^{N_2} (H2_i - \bar{H2})^2}{N_2 - 1}}$$

### 1.3.1 Accuracy Metrics (Observed score)

These metrics show how well system scores  $M$  predict observed human scores  $H$ . The computed metrics are available in the [intermediate file](#) `eval`, with a subset of the metrics also available in the intermediate file `eval_short`.

#### Percent exact agreement (rounded scores only)

Percentage responses where human and system scores match exactly.

$$A = \sum_{i=1}^N \frac{w_i}{N} \times 100$$

where  $w_i = 1$  if  $M_i = H_i$  and  $w_i = 0$  if  $M_i \neq H_i$

The percent exact agreement is computed using [rsmttool.utils.agreement](#) with `tolerance` set to 0.

#### Percent exact + adjacent agreement

Percentage responses where the absolute difference between human and system scores is 1 or less.

$$A_{adj} = \sum_{i=1}^N \frac{w_i}{N} \times 100$$

where  $w_i = 1$  if  $|M_i - H_i| \leq 1$  and  $w_i = 0$  if  $|M_i - H_i| > 1$ .

The percent exact + adjacent agreement is computed using [rsmttool.utils.agreement](#) with `tolerance` set to 1.

#### Cohen's kappa (rounded scores only)

$$\kappa = 1 - \frac{\sum_{k=0}^{K-1} \sum_{j=1}^K w_{jk} X_{jk}}{\sum_{k=0}^{K-1} \sum_{j=1}^K w_{jk} m_{jk}}$$

when  $k = j$  then  $w_{jk} = 0$  and when  $k \neq j$  then  $w_{jk} = 1$

where:

- $K$  is the number of scale score categories (maximum observed rating - minimum observed rating + 1). Note that for  $\kappa$  computation the values of  $H$  and  $M$  are shifted to  $H\text{-minimum\_rating}$  and  $M\text{-minimum\_rating}$  so that the lowest value is 0. This is done to support negative labels.
- $X_{jk}$  is the number times where  $H = j$  and  $M = k$ .
- $m_{jk}$  is the percent chance agreement:

$$m_{jk} = \sum_{k=1}^K \frac{n_{k+}}{N} \frac{n_{+k}}{N}, \text{ where}$$

–  $n_{k+}$  - total number of responses where  $H_i = k$

- $n_{+k}$  - total number of responses where  $M_i = k$

Kappa is computed using `skl.metrics.kappa` with `weights` set to `None` and `allow_off_by_one` set to `False` (default).

**Note:** See [this discussion](#) for the explanation of how the [SKLL implementation](#) differs from the [scikit-learn implementation](#). The two implementations might produce different results if the matrix contains missing labels. For example, consider the hypothetical scenario where our predictions only contain the labels 1, 2, and 4. In the SKLL implementation, the missing 3 will be automatically added to the list of labels whereas in the scikit-learn implementation, the 3 would only be added if a complete list of labels was passed to the function via the optional `labels` keyword argument.

### Quadratic weighted kappa (QWK)

Unlike [Cohen's kappa](#) which is only computed for rounded scores, quadratic weighted kappa is computed for continuous scores using the following formula:

$$QWK = \frac{2 * Cov(M, H)}{Var(H) + Var(M) + (\bar{M} - \bar{H})^2}$$

Note that in this case the variances and covariance are computed by dividing by  $N$  and not by  $N-1$ , as in other cases.

QWK is computed using `rsmttool.utils.quadratic_weighted_kappa` with `ddof` set to `0`.

See [Haberman \(2019\)](#) for the full derivation of this formula. The discrete case is simply treated as a special case of the continuous one.

**Note:** In RSMTTool v6.x and earlier QWK was computed using `skl.metrics.kappa` with `weights` set to `"quadratic"`. Continuous scores were rounded for computation. Both formulas produce the same scores for discrete (rounded scores) but QWK values for continuous scores computed by RSMTTool starting with v7.0 will be *different* from those computed by earlier versions.

### Pearson Correlation coefficient (r)

$$r = \frac{\sum_{i=1}^N (H_i - \bar{H})(M_i - \bar{M})}{\sqrt{\sum_{i=1}^N (H_i - \bar{H})^2 \sum_{i=1}^N (M_i - \bar{M})^2}}$$

Pearson correlation coefficients is computed using `scipy.stats.pearsonr`.

If the variance of human or system scores is `0` (all scores are the same) or only one response is available, RSMTTool returns `None`.

**Note:** In `scipy` v1.4.1 and later, the implementation uses the following formula:

$$r = \frac{H - \bar{H}}{\|H - \bar{H}\|_2} \cdot \frac{M - \bar{M}}{\|M - \bar{M}\|_2}$$

This implementation is more robust to very large values but is more likely to return a value slightly smaller than 1 (for example, 0.9999999999999998) for perfect correlation when  $n$  is small. See [this comment](#) for further detail.

### Standardized mean difference (SMD)

This metrics ensures that the distribution of system scores is centered on a point close to what is observed with human scoring.

$$SMD = \frac{\bar{M} - \bar{H}}{\sigma_H}$$

SMD between system and human scores is computed using `rsmttool.utils.standardized_mean_difference` with the `method` argument set to "unpooled".

---

**Note:** In RSMTool v6.x and earlier SMD was computed with the `method` argument set to "williamson" as described in Williamson et al. (2012). The values computed by RSMTool starting with v7.0 will be *different* from those computed by earlier versions.

---

### Mean squared error (MSE)

The mean squared error of a machine score as a predictor of observed human score  $H$ :

$$MSE(H|M) = \frac{1}{N} \sum_{i=1}^N (H_i - M_i)^2$$

MSE is computed using `sklearn.metrics.mean_squared_error`.

### Proportional reduction in mean squared error for observed score (R2)

$$R2 = 1 - \frac{MSE(H|M)}{\sigma_H^2}$$

R2 is computed using `sklearn.metrics.r2_score`. If only one response is available, RSMTool returns None.

## 1.3.2 Accuracy Metrics (True score)

According to test theory, an observed score is a combination of the true score  $T$  and a measurement error. The true score cannot be observed, but its distribution parameters can be estimated from observed scores. Such an estimation requires that two human scores be available for *at least a* subset of responses in the evaluation set since these are necessary to estimate the measurement error component.

Evaluating system against true score produces performance estimates that are robust to errors in human scores and remain stable even when human-human agreement varies (see Loukina et al. (2020)).

The true score evaluations computed by RSMTool are available in the *intermediate file* `true_score_eval`.

### Proportional reduction in mean squared error for true scores (PRMSE)

PRMSE shows how well system scores can predict true scores. This metric generally varies between 0 (random prediction) and 1 (perfect prediction), although in some cases it can take negative values (suggesting a very bad fit) or exceed 1 (suggesting that the sample size is too small to reliably estimate rater error variance).

PRMSE for true scores is defined similarly to *PRMSE for observed scores*, but with the true score  $T$  used instead of the observed score  $H$ , that is, as the percentage of variance in the true scores explained by the system scores.

$$PRMSE = 1 - \frac{MSE(T|M)}{\sigma_T^2}$$

In the simple case where all responses have two human scores,  $MSE(T|M)$  (**mean squared error when predicting true score with system score**) and  $\sigma_T^2$  (**variance of true score**) are estimated from their observed score counterparts  $MSE(H|M)$  and  $\sigma_H^2$  as follows:

- $\hat{H}$  is used instead of  $H$  to compute  $MSE(\hat{H}|M)$  and  $\sigma_{\hat{H}}^2$ .  $\hat{H}$  is the average of two human scores for each response ( $\hat{H}_i = \frac{H_i + H_{2i}}{2}$ ). These evaluations use  $\hat{H}$  rather than  $H$  because the measurement errors for each rater are assumed to be random and, thus, can partially cancel out making the average  $\hat{H}$  closer to true score  $T$  than  $H$  or  $H_2$ .
- To compute estimates for true scores, the values for observed scores are adjusted for **variance of measurement errors** ( $\sigma_e^2$ ) in human scores defined as:

$$\sigma_e^2 = \frac{1}{2 \times N_2} \sum_{i=1}^{N_2} (H_i - H_{2i})^2$$

In the simple case, where **all responses are double-scored**,  $MSE(T|M)$  is estimated as:

$$MSE(T|M) = MSE(\hat{H}|M) - \frac{1}{2}\sigma_e^2$$

and  $\sigma_T^2$  is estimated as:

$$\sigma_T^2 = \sigma_{\hat{H}}^2 - \frac{1}{2}\sigma_e^2$$

The PRMSE formula implemented in RSMTool is more general and can also handle the case where the number of available ratings varies across the responses (e.g. **only a subset of responses is double-scored**). While `rsmtool` and `rsmeval` only support evaluations with two raters, the implementation of the PRMSE formula available via the *API* supports cases where some of the responses have **more than two** ratings available. The formula was derived by Matt S. Johnson and is explained in more detail in [Loukina et al. \(2020\)](#).

In this case, the variance of rater errors is computed as a pooled variance estimator.

We first calculate the within-subject variance of human ratings for each response,  $V_i$ , using denominator  $c_i - 1$ :

$$V_i = \frac{\sum_{j=1}^{c_i} (H_{i,j} - \bar{H}_i)^2}{c_i - 1}$$

where

- $H_{i,j}$  is the human score assigned by rater  $j$  to response  $i$
- $c_i$  is the total number of human scores available for response  $i$ . For double-scored responses this equals 2.
- $\bar{H}_i$  is the average human rating for response  $i$ .

We then take a weighted average of those within-responses variances:

$$\sigma_e^2 = \frac{\sum_{i=1}^N V_i * (c_i - 1)}{\sum_{i=1}^N (c_i - 1)}$$

The **true score variance**  $\sigma_T^2$  is then estimated as

$$\sigma_T^2 = \frac{\sum_{i=1}^N c_i (\bar{H}_i - \bar{H})^2 - (N-1)\sigma_e^2}{c. - \frac{\sum_{i=1}^N c_i^2}{c.}}$$

where

- $c. = \sum_{i=1}^N c_i$  is the total number of observed human scores.
- $\bar{H}_i$  is the average human rating for response  $i$ . For responses with only one rating this will be the single human score  $H$ .

**Mean squared error**  $MSE(T|M)$  is estimated as:

$$MSE(T|M) = \frac{1}{c.} \left( \sum_{i=1}^N c_i (\bar{H}_i - M_i)^2 - N\sigma_e^2 \right)$$

The formulas are derived to ensure consistent results regardless of the number of raters and of the number of ratings available for each response.

PRMSE is computed using the `rsmtool.utils.prmse_true` function.

In some cases, it may be appropriate to compute variance of human errors using a different sample than the one used for main evaluations. This can be accomplished using `rsmtool.utils.variance_of_errors` and using an optional configuration field `rater_error_variance` in `rsmtool` or `rsmtool`

---

**Note:** The PRMSE formula assigns higher weight to discrepancies between system scores and human scores when human score is the average of two or more human scores than when the human score is based on a single score.

---

### 1.3.3 Fairness

Fairness of automated scores is an important component of RSMTTool evaluations (see [Madnani et al, 2017](#)).

When defining an experiment, the RSMTTool user has the option of specifying which subgroups should be considered for such evaluations using `subgroups` field. These subgroups are then used in all fairness evaluations.

All fairness evaluations are conducted on the evaluation set. The metrics are only computed for either `raw_trim` or `scale_trim` scores (see [score postprocessing](#) for further details) depending on the value of `use_scaled_predictions` in RSMTTool or the value of `scale_with` in RSMEval.

#### Differences between standardized means for subgroups (DSM)

This is a standard evaluation used for evaluating subgroup differences. The metrics are available in the `intermediate files eval_by_<SUBGROUP>`.

DSM is computed as follows:

1. For each group, get the  $z$ -score for each response  $i$ , using the  $\bar{H}$ ,  $\bar{M}$ ,  $\sigma_H$ , and  $\sigma_S$  for system and human scores for the whole evaluation set:

$$z_{H_i} = \frac{H_i - \bar{H}}{\sigma_H}$$

$$z_{M_i} = \frac{M_i - \bar{M}}{\sigma_M}$$

2. For each response  $i$ , calculate the difference between machine and human scores:  $z_{M_i} - z_{H_i}$
3. Calculate the mean of the difference  $z_{M_i} - z_{H_i}$  by subgroup of interest.

DSM is computed using `rsmtool.utils.difference_of_standardized_means` with:



$\text{population\_y\_true\_observe\_mn} = \bar{H}$  for the whole evaluation set

$\text{population\_y\_pred\_mn} = \bar{M}$  for the whole evaluation set

$\text{population\_y\_true\_observed\_sd} = \sigma_H$  for the whole evaluation set

$\text{population\_y\_pred\_sd} = \sigma_M$  for the whole evaluation set

---

**Note:** In RSMTTool v6.x and earlier, subgroup differences were computed using *standardized mean difference* with the `method` argument set to "williamson". Since the differences computed in this manner were very sensitive to score distributions, RSMTTool no longer uses this function to compute subgroup differences starting with v7.0.

---

### Additional fairness evaluations

Starting with v7.0, RSMTTool includes additional fairness analyses suggested in Loukina, Madnani, & Zechner, 2019. The computed metrics from these analyses are available in *intermediate files* `fairness_metrics_by_<SUBGROUP>`.

These include:

- Overall score accuracy: percentage of variance in squared error  $(M - H)^2$  explained by subgroup membership
- Overall score difference: percentage of variance in absolute error  $(M - H)$  explained by subgroup membership
- Conditional score difference: percentage of variance in absolute error  $(M - H)$  explained by subgroup membership when controlling for human score

Please refer to the paper for full descriptions of these metrics.

The fairness metrics are computed using `rsmttool.fairness_utils.get_fairness_analyses`.

### 1.3.4 Human-human agreement

If scores from a second human ( $H2$ ) are available, RSMTTool computes the following additional metrics for human-human agreement using only the  $N_2$  responses, including only responses that contain numeric values for both the  $H$  and  $H2$  columns.

The computed metrics are available in the *intermediate file* `consistency`.

#### Percent exact agreement

Same as *percent exact agreement for observed scores* but substituting  $H2$  for  $M$ .

#### Percent exact + adjacent agreement

Same as *percent adjacent agreement for observed scores* but substituting  $H2$  for  $M$  and  $N_2$  for  $N$ .

### Cohen's kappa

Same as *Cohen's kappa for observed scores* but substituting  $H_2$  for  $M$  and  $N_2$  for  $N$ .

### Quadratic weighted kappa (QWK)

Same as *QWK for observed scores* but substituting  $H_2$  for  $M$  and  $N_2$  for  $N$ .

### Pearson Correlation coefficient (r)

Same as *r for observed scores* but substituting  $H_2$  for  $M$  and  $N_2$  for  $N$ .

### Standardized mean difference (SMD)

$$SMD = \frac{\bar{H}_2 - \bar{H}_1}{\sqrt{\frac{\sigma_H^2 + \sigma_{H_2}^2}{2}}}$$

Unlike *SMD for human-system scores*, the denominator in this case is the “pooled” standard deviation of  $H_1$  and  $H_2$ . Therefore, SMD between two human scores is computed using `rsmttool.utils.standardized_mean_difference` with the `method` argument set to “pooled”.

---

**Note:** In RSMTTool v6.x and earlier, SMD was computed with the `method` argument set to “williamson” as described in [Williamson et al. \(2012\)](#). Starting with v7.0, the values computed by RSMTTool will be *different* from those computed by earlier versions.

---

## 1.4 Installation

Note that RSMTTool currently works with Python 3.8, 3.9, and 3.10.

### 1.4.1 Installing with conda

Currently, the recommended way to install RSMTTool is by using the conda package manager. If you have already installed conda, you can skip straight to Step 2.

1. To install conda, follow the instructions on [this page](#).
2. Create a new conda environment (say, `rsmttool`) and install the RSMTTool conda package for your preferred Python version. For example, for Python 3.8, run:

```
conda create -n rsmttool -c conda-forge -c ets python=3.8 rsmttool
```

3. Activate this conda environment by running `conda activate rsmttool`. You should now have all of the RSMTTool command-line utilities in your path.
4. From now on, you will need to activate this conda environment whenever you want to use RSMTTool. This will ensure that the packages required by RSMTTool will not affect other projects.

RSMTTool can also be downloaded directly from [GitHub](#).

### 1.4.2 Installing with pip

You can also use `pip` to install RSMTool instead of `conda`. To do so, simply run:

```
pip install rsmtool
```

However, note that `pip` may be quite slow since its dependency resolution is not as good as `conda`'s. Also, note that if you are on macOS, you will need to have the following line in your `.bashrc` for RSMTool to work properly:

```
export MPLBACKEND=Agg
```

## 1.5 Tutorial

First you'll want to *install RSMTool* and make sure you the `conda` environment in which you installed it is activated.

### 1.5.1 Workflow

---

**Important:** Although this tutorial provides feature values for the purpose of illustration, RSMTool does *not* include any functionality for feature extraction; the tool is *designed for researchers* who use their own NLP/Speech processing pipeline to extract features for their data.

---

Once you have the features extracted from your data, using RSMTool is fairly straightforward:

1. Make sure your data is split into a training set and a held-out evaluation set. You will need two files containing the feature values: one for each set. The files should be in one of the *supported formats*
2. Create an *experiment configuration file* describing the modeling experiment you would like to run.
3. Run that configuration file with *rsmtool* and generate the experiment HTML report as well as the *intermediate CSV files*.
4. Examine the HTML report to check various aspects of model performance.

Note that the above workflow does not use the customization features of RSMTool, e.g., *choosing which sections to include in the report* or *adding custom analyses sections* etc. However, we will stick with this workflow for our tutorial since it is likely to be the most common use case.

### 1.5.2 ASAP Example

Let's see how we can apply this basic RSMTool workflow to a simple example based on a 2012 Kaggle competition on automated essay scoring called the *Automated Student Assessment Prize (ASAP) contest*. As part of that contest, responses to 8 different essay questions written by students in grades 6-10 were provided. The responses were scored by humans and given a holistic score indicating the English proficiency of the student. The goal of the contest was to build an automated scoring model with the highest accuracy on held-out evaluation data.

For our tutorial, we will use one of the questions from this data to illustrate how to use RSMTool to train a scoring model and evaluate its performance. All of the data we refer to below can be found in the `examples/rsmtool` folder in the [github repository](#).

### Extract features

Kaggle provides the actual text of the responses along with the human scores in a `.tsv` file. The first step, as with any automated scoring approaches, is to extract features from the written (or spoken) responses that might be useful in predicting the score we are interested in. For example, some features that might be useful for predicting the English proficiency of a response are:

- Does the response contain any grammatical errors and if so, how many and of what kind since some grammatical errors are more serious than others.
- Is the response well organized, e.g., whether there is a clear thesis statement and a conclusion, etc.
- Is the response coherent, i.e., do the ideas expressed in the various paragraphs well connected?
- Does the response contain any spelling errors?

For more examples of what features might look like, we refer the reader to this paper by Attali & Burstein<sup>1</sup>. For our ASAP2 data, we extracted the following four features:

- **GRAMMAR**: indicates how grammatically fluent the responses is.
- **MECHANICS**: indicates how free of mechanical errors (e.g., spelling errors) the response is.
- **DISCOURSE**: indicates whether the discourse transitions in the response make sense.
- **ORGANIZATION**: indicates whether the response is organized well.

---

**Note:** These features were extracted using proprietary NLP software. To reiterate, RSMTTool does *not* include any NLP/speech components for feature extraction.

---

### Create a configuration file

The next step is to create an *RSMTTool experiment configuration file* in `.json` format.

```
1 {
2   "experiment_id": "ASAP2",
3   "train_file": "train.csv",
4   "test_file": "test.csv",
5   "model": "LinearRegression",
6   "description": "A model which uses all features and a LinearRegression.",
7   "test_label_column": "score",
8   "train_label_column": "score",
9   "use_scaled_predictions": true,
10  "trim_min": 1,
11  "trim_max": 6,
12  "id_column": "ID",
13  "second_human_score_column": "score2",
14  "length_column": "LENGTH"
15 }
```

Let's take a look at the options in our configuration file.

- **Line 2:** We define an experiment ID

---

<sup>1</sup> Attali, Y., & Burstein, J. (2006). Automated essay scoring with e-rater® V.2. *Journal of Technology, Learning, and Assessment*, 4(3). <https://ejournals.bc.edu/index.php/jtla/article/download/1650/1492>

- **Lines 3-4:** We list the paths to our training and evaluation files with the feature values. For this tutorial we used .csv format, but RSMTool also supports several other *input file formats*.
- **Line 5:** We choose to use a linear regression model to combine those features into a score.
- **Line 6:** We also provide a description which will be included in the experiment report.
- **Lines 7-8:** These two fields indicate that the human scores in the two .csv files are located in columns named score.
- **Lines 9:** Next, we indicate that we would like to use the *scaled scores* for our evaluation analyses.
- **Lines 10-11:** These fields indicate that the lowest score on the scoring scale is a 1 and the highest score is a 6. This information is usually part of the rubric used by human graders.
- **Line 12:** This field indicates that the unique IDs for the responses in the two .csv files are located in columns named ID.
- **Line 13:** This field indicates that scores from a second set of human graders are also available (useful for comparing the agreement between human-machine scores to the agreement between two sets of humans) and are located in the score2 column in the test set .csv file.
- **Line 14:** This field indicates that response lengths are also available for the training data in a column named LENGTH.

Documentation for all of the available configuration options is available [here](#).

---

#### Note:

1. For this example, we are using *all* of the non-metadata columns in the training and evaluation .csv files as features in the model. However, it is also possible to *choose specific columns* to be used for training the model.
  2. You can also use our nifty capability to *automatically generate* rsmtool configuration files rather than creating them manually.
  3. To view the tutorial output on a [Weights & Biases](#) project, use the file *config\_rsmtool\_with\_wandb.json* in the same tutorial folder. Make sure to edit the *W&B entity* field to work with your Weights & Biases credentials.
- 

## Run the experiment

Now that we have our features in the right format and our configuration file in .json format, we can use the *rsmtool* command-line script to run our modeling experiment.

```
$ cd examples/rsmtool
$ rsmtool config_rsmtool.json
```

This should produce output like:

```
Output directory: /Users/nmadnani/work/rsmtool/examples/rsmtool
Loading experiment data
Reading configuration file: /Users/nmadnani/work/rsmtool/examples/rsmtool/config_rsmtool.
↪ json
Reading training data: /Users/nmadnani/work/rsmtool/examples/rsmtool/train.csv
Reading evaluation data: /Users/nmadnani/work/rsmtool/examples/rsmtool/test.csv
Pre-processing training and test set features
Saving training and test set data to disk
Training LinearRegression model
```

(continues on next page)

(continued from previous page)

```
The exact solution is x = 0
Running analyses on training set
Generating training and test set predictions
Processing test set predictions
Saving training and test set predictions to disk
Scaling the coefficients and saving them to disk
Running analyses on test set predictions
Starting report generation
Merging sections
Exporting HTML
Executing notebook with kernel: python3
```

Once the run finishes, you will see the output, figure, report, and feature sub-directories in the current directory. Each of these directories contain *useful information* but we are specifically interested in the report/ASAP2\_report.html file, which is the final RSMTTool experiment report.

## Examine the report

Our experiment report contains all the information we would need to determine how well our linear regression model with four features is doing. It includes:

1. Descriptive feature statistics.
2. Inter-feature correlations.
3. Model fit and diagnostics.
4. Several different metrics indicating how well the machine's scores agree with the humans'.

... and *much more*.

If we are not satisfied with the performance of the model, we can then take specific action, e.g., either add new features or perhaps try a more sophisticated model. And then run `rsmttool` again and examine the report. Building a scoring model is generally an iterative process.

### 1.5.3 What's next?

Next, you should read the detailed documentation on *rsmttool* since the tutorial only covered a very basic example. If you are interested in RSMTTool but your use case is more nuanced, we have probably got you *covered*.

## References

## 1.6 Using RSMTTool

For most users, the primary means of using RSMTTool will be via the command-line utility `rsmttool`. We refer to each run of `rsmttool` as an "experiment".

### 1.6.1 Input

rsmtool requires a single argument to run an experiment: the path to *a configuration file*. It can also take an output directory as an optional second argument. If the latter is not specified, rsmtool will use the current directory as the output directory.

Here are all the arguments to the rsmtool command-line script.

**config\_file**

The *JSON configuration file* for this experiment.

**output\_dir** (optional)

The output directory where all the files for this experiment will be stored.

**-f, --force**

If specified, the contents of the output directory will be overwritten even if it already contains the output of another rsmtool experiment.

**-h, --help**

Show help message and exist.

**-V, --version**

Show version number and exit.

### Experiment configuration file

This is a file in *.json* format that provides overall configuration options for an rsmtool experiment. Here's an *example configuration file* for rsmtool.

---

**Note:** To make it easy to get started with rsmtool, we provide a way to **automatically generate** configurations file both interactively as well as non-interactively. Novice users will find interactive generation more helpful while more advanced users will prefer non-interactive generation. See *this page* for more details.

---

Next, we describe all of the rsmtool configuration fields in detail. There are four required fields and the rest are all optional. We first describe the required fields and then the optional ones (sorted alphabetically).

#### experiment\_id

An identifier for the experiment that will be used to name the report and all *intermediate files*. It can be any combination of alphanumeric values, must *not* contain spaces, and must *not* be any longer than 200 characters.

#### model

The machine learner you want to use to build the scoring model. Possible values include *built-in linear regression models* as well as all of the learners available via *SKLL*. With SKLL learners, you can customize the *tuning objective* and also *compute expected scores as predictions*.

### train\_file

The path to the training data feature file in one of the *supported formats*. Each row should correspond to a single response and contain numeric feature values extracted for this response. In addition, there should be a column with a unique identifier (ID) for each response and a column with the human score for each response. The path can be absolute or relative to the location of config file.

### test\_file

The path to the evaluation data feature file in one of the *supported formats*. Each row should correspond to a single response and contain numeric feature values extracted for this response. In addition, there should be a column with a unique identifier (ID) for each response and a column with the human score for each response. The path can be absolute or relative to the location of config file.

---

#### Note:

1. For both the training and evaluation files, the default behavior of `rsmtool` is to look for a column named `spkitemid` in order to get the unique IDs for each response and a column named `sc1` to get train/test labels. The optional fields *id\_column* and *train\_label\_column* can be used to specify different names for these columns.
  2. `rsmtool` also assumes that all other columns present in these files (other than those containing IDs and labels) contain feature values. If this is not the case, one can use other configuration file fields to identify columns containing non-feature information useful for various analyses, e.g., *second\_human\_score\_column*, *flag\_column*, *subgroups* et cetera below.
  3. Any columns not explicitly identified in (1) and (2) will be considered feature columns and used by `rsmtool` in the model. To use only a *subset* of these remaining columns as features, one can employ the four optional fields *features*, *feature\_subset\_file*, *feature\_subset*, and *sign*. See *selecting feature columns* for more details on how to achieve this.
- 

### candidate\_column (Optional)

The name for an optional column in the training and test data containing unique candidate IDs. Candidate IDs are different from response IDs since the same candidate (test-taker) might have responded to multiple questions.

### custom\_sections (Optional)

A list of custom, user-defined sections to be included into the final report. These are IPython notebooks (`.ipynb` files) created by the user. The list must contain paths to the notebook files, either absolute or relative to the configuration file. All custom notebooks have access to some *pre-defined variables*.



**description (Optional)**

A brief description of the experiment. This will be included in the report. The description can contain spaces and punctuation. It's blank by default.

**exclude\_zero\_scores (Optional)**

By default, responses with human scores of 0 will be excluded from both training and evaluation set. Set this field to `false` if you want to keep responses with scores of 0. Defaults to `true`.

**feature\_subset (Optional)**

Name of the pre-defined feature subset to be used if using *subset-based column selection*.

**feature\_subset\_file (Optional)**

Path to the feature subset file if using *subset-based column selection*.

**features (Optional)**

Path to the file with list of features if using *fine-grained column selection*. Alternatively, you can pass a list of feature names to include in the experiment.

**file\_format (Optional)**

The format of the *intermediate files*. Options are `csv`, `tsv`, or `xlsx`. Defaults to `csv` if this is not specified.

**flag\_column (Optional)**

This field makes it possible to only use responses with particular values in a given column (e.g. only responses with a value of 0 in a column called ADVISORY). The field takes a dictionary in Python format where the keys are the names of the columns and the values are lists of values for responses that will be used to train the model. For example, a value of `{"ADVISORY": 0}` will mean that `rsmtool` will *only* use responses for which the ADVISORY column has the value 0. If this field is used without `flag_column_test`, the conditions will be applied to *both* training and evaluation set and the specified columns must be present in both sets. When this field is used in conjunction with `flag_column_test`, the conditions will be applied to *training set only* and the specified columns must be present in the training set. Defaults to `None`.

---

**Note:** If several conditions are specified (e.g., `{"ADVISORY": 0, "ERROR": 0}`) only those responses which satisfy *all* the conditions will be selected for further analysis (in this example, these will be the responses where the ADVISORY column has a value of 0 *and* the ERROR column has a value of 0).

---

---

**Note:** When reading the values in the supplied dictionary, `rsmtool` treats numeric strings, floats and integers as the same value. Thus 1, 1.0, "1" and "1.0" are all treated as the 1.0.

---

### flag\_column\_test (Optional)

This field makes it possible to only use a separate Python flag dictionary for the evaluation set. If this field is not passed, and `flag_column` is passed, then the same advisories will be used for both training and evaluation sets.

When this field is used, the specified columns must be present in the evaluation set. Defaults to `None` or `flag_column`, if `flag_column` is present. Use `flag_column_test` only if you want filtering of the test set.

---

**Note:** When used, `flag_column_test` field determines *all* filtering conditions for the evaluation set. If it is used in conjunction with `flag_column` field, the filtering conditions defined in `flag_column` will *only* be applied to the training set. If you want to apply a subset of conditions to both partitions with additional conditions applied to the evaluation set only, you will need to specify the overlapping conditions separately for each partition.

---

### general\_sections (Optional)

RSMTTool provides pre-defined sections for `rsmttool` (listed below) and, by default, all of them are included in the report. However, you can choose a subset of these pre-defined sections by specifying a list as the value for this field.

- `data_description`: Shows the total number of responses in training and evaluation set, along with any responses have been excluded due to non-numeric features/scores or *flag columns*.
- `data_description_by_group`: Shows the total number of responses in training and evaluation set for each of the *subgroups* specified in the configuration file. This section only covers the responses used to train/evaluate the model.
- `feature_descriptives`: Shows the descriptive statistics for all raw feature values included in the model:
  - a table showing mean, standard deviation, min, max, correlation with human score etc.;
  - a table with percentiles and outliers; and
  - a barplot showing the number of truncated outliers for each feature.
- `features_by_group`: Shows boxplots with distributions of raw feature values by each of the *subgroups* specified in the configuration file.
- `preprocessed_features`: Shows analyses of preprocessed features:
  - histograms showing the distributions of preprocessed features values;
  - the correlation matrix between all features and the human score;
  - a barplot showing marginal and partial correlations between all features and the human score, and, optionally, response length if *length\_column* is specified in the config file.
- `dff_by_group`: Differential feature functioning by group. The plots in this section show average feature values for each of the *subgroups* conditioned on human score.
- `consistency`: Shows metrics for *human-human agreement*, the difference (“degradation”) between the human-human and human-system agreement, and the disattenuated human-machine correlations. This notebook is only generated if the config file specifies *second\_human\_score\_column*.
- `model`: Shows the parameters of the learned regression model. For linear models, it also includes the standardized and relative coefficients as well as model diagnostic plots.
- `evaluation`: Shows the *standard set of evaluations* recommended for scoring models on the evaluation data:
  - a table showing human-system association metrics;

- the confusion matrix; and
- a barplot showing the distributions for both human and machine scores.
- **evaluation\_by\_group**: Shows barplots with the main evaluation metrics by each of the subgroups specified in the configuration file.
- **fairness\_analyses**: Additional *fairness analyses* suggested in Loukina, Madnani, & Zechner, 2019. The notebook shows:
  - percentage of variance in squared error explained by subgroup membership
  - percentage of variance in raw (signed) error explained by subgroup membership
  - percentage of variance in raw (signed) error explained by subgroup membership when controlling for human score
  - plots showing estimates for each subgroup for each model
- **true\_score\_evaluation**: evaluation of system scores against the *true scores* estimated according to test theory. The notebook shows:
  - Number of single and double-scored responses.
  - Variance of human rater errors and estimated variance of true scores
  - Mean squared error (MSE) and proportional reduction in mean squared error (PRMSE) when predicting true score with system score.
- **pca**: Shows the results of principal components analysis on the processed feature values:
  - the principal components themselves;
  - the variances; and
  - a Scree plot.

The analysis keeps all components. The total number of components usually equals the total number of features. In cases where the total number of responses is smaller than the number of features, the number of components is the same as the number of responses.
- **intermediate\_file\_paths**: Shows links to all of the intermediate files that were generated while running the experiment.
- **sysinfo**: Shows all Python packages along with versions installed in the current environment while generating the report.

### **id\_column** (*Optional*)

The name of the column containing the response IDs. Defaults to `spkitemid`, i.e., if this is not specified, `rsmttool` will look for a column called `spkitemid` in the training and evaluation files.

### length\_column (Optional)

The name for the optional column in the training and evaluation data containing response length. If specified, length is included in the inter-feature and partial correlation analyses. Note that this field *should not* be specified if you want to use the length column as an actual feature in the model. In the latter scenario, the length column will automatically be included in the analyses, like any other feature. If you specify `length_column` and include the same column name as a feature in the *feature file*, `rsmttool` will ignore the `length_column` setting. In addition, if `length_column` has missing values or if its standard deviation is 0 (both somewhat unlikely scenarios), `rsmttool` will *not* include any length-based analyses in the report.

### min\_items\_per\_candidate (Optional)

An integer value for the minimum number of responses expected from each candidate. If any candidates have fewer responses than the specified value, all responses from those candidates will be excluded from further analysis. Defaults to `None`.

### min\_n\_per\_group (Optional)

A single numeric value or a dictionary with keys as the group names listed in the *subgroups* field and values as the thresholds for the groups. When specified, only groups with *at least* this number of instances will be displayed in the tables and plots contained **in the report**. Note that this parameter *only* affects the HTML report and the figures. For all analyses – including the computation of the population parameters – data from *all* groups will be used. In addition, the *intermediate files* will still show the results for *all* groups.

---

**Note:** If you supply a dictionary, it *must* contain a key for *every* subgroup listed in *subgroups* field. If no threshold is to be applied for some of the groups, set the threshold value for this group to 0 in the dictionary.

---

---

**Note:** Any provided thresholds will be applied when displaying the feature descriptive analyses conducted on the training set *and* the results of the performance analyses computed on the evaluation set.

---

### predict\_expected\_scores (Optional)

If a probabilistic SKLL classifier is chosen to build the scoring model, then *expected scores* — probability-weighted averages over contiguous, numeric score points — can be generated as the machine predictions instead of the most likely score point, which would be the default for a classifier. Set this field to `true` to compute expected scores as predictions. Defaults to `false`.

---

**Note:** You may see slight differences in expected score predictions if you run the experiment on different machines or on different operating systems most likely due to very small probability values for certain score points which can affect floating point computations.

---

**rater\_error\_variance (Optional)**

*True score evaluations* require an estimate of rater error variance. By default, `rsmtool` will compute this variance from double-scored responses in the data. However, in some cases, one may wish to compute the variance on a different sample of responses. In such cases, this field can be used to set the rater error variance to a precomputed value which is then used as-is by `rsmtool`. You can use the `rsmtool.utils.variance_of_errors` function to compute rater error variance outside the main evaluation pipeline.

**second\_human\_score\_column (Optional)**

The name for an optional column in the test data containing a second human score for each response. If specified, additional information about human-human agreement and degradation will be computed and included in the report. Note that this column must contain either numbers or be empty. Non-numeric values are *not* accepted. Note also that the *exclude\_zero\_scores (Optional)* option below will apply to this column too.

---

**Note:** You do not need to have second human scores for *all* responses to use this option. The human-human agreement statistics will be computed as long as there is at least one response with numeric value in this column. For responses that do not have a second human score, the value in this column should be blank.

---

**section\_order (Optional)**

A list containing the order in which the sections in the report should be generated. Any specified order must explicitly list:

1. Either *all* pre-defined sections if a value for the *general\_sections* field is not specified OR the sections specified using *general\_sections*, and
2. All custom section names specified using *custom\_sections*, i.e., file prefixes only, without the path and without the *.ipynb* extension, and

**select\_transformations (Optional)**

If this option is set to `true` the system will try apply feature transformations to each of the features and then choose the transformation for each feature that yields the highest correlation with human score. The possible transformations are:

- `raw`: no transformation, use original feature value
- `org`: same as raw
- `inv`:  $1/x$
- `sqrt`: square root
- `addOneInv`:  $1/(x+1)$
- `addOneLn`:  $\ln(x+1)$

We only consider transformations that produce numeric results for *all* values for a given feature column. For example, if a feature column contains a single negative value, the `sqrt` transformation will be ignored even if it would have resulted in the highest correlation with human score for the remaining values. In addition, the `inv` and `addOneInv` transformations are never used for feature columns that contain both positive and negative values.

Defaults to `false`.

### See also:

It is also possible to manually apply transformations to any feature as part of the *feature column selection* process.

### **sign (Optional)**

Name of the column containing expected correlation sign between each feature and human score if using *subset-based column selection*.

### **skll\_fixed\_parameters (Optional)**

Any fixed hyperparameters to be used if a SKLL model is chosen to build the scoring model. This should be a dictionary with the names of the hyperparameters as the keys. To determine what hyperparameters are available for the SKLL learner you chose, consult the scikit-learn documentation for the learner with the same name as well as the [SKLL documentation](#). Any values you specify here will override both the scikit-learn and SKLL defaults. The values for a key can be string, integer, float, or boolean depending on what the hyperparameter expects. Note that if this option is specified with the *built-in linear regression models*, it will simply be ignored.

### **skll\_grid\_search\_jobs (Optional)**

Number of folds to run in parallel when using SKLL grid search. Defaults to 1 (no parallelization). For very large data sets, setting this option to a higher number (e.g., number of CPU cores) may significantly improve running time.

### **skll\_objective (Optional)**

The tuning objective to use if a SKLL model is chosen to build the scoring model. Possible values are the objectives available via [SKLL](#). Defaults to `neg_mean_squared_error` for SKLL regressors and `f1_score_micro` for SKLL classifiers. Note that if this option is specified with the *built-in linear regression models*, it will simply be ignored.

### **standardize\_features (Optional)**

If this option is set to `false` features will not be standardized by subtracting the mean and dividing by the standard deviation. Defaults to `true`.

### **subgroups (Optional)**

A list of column names indicating grouping variables used for generating analyses specific to each of those defined subgroups. For example, `["prompt", "gender", "native_language", "test_country"]`. These subgroup columns need to be present in both training *and* evaluation data. If subgroups are specified, `rsmttool` will generate:

- description of the data by each subgroup;
- boxplots showing the feature distribution for each subgroup on the training set; and
- tables and barplots showing human-system agreement for each subgroup on the evaluation set.

**test\_label\_column (Optional)**

The name for the column containing the human scores in the test data. If set to `fake`, fake scores will be generated using randomly sampled integers. This option may be useful if you only need descriptive statistics for the data and do not care about the other analyses. Defaults to `sc1`.

**train\_label\_column (Optional)**

The name for the column containing the human scores in the training data. If set to `fake`, fake scores will be generated using randomly sampled integers. This option may be useful if you only need descriptive statistics for the data and do not care about the other analyses. Defaults to `sc1`.

---

**Note:** All responses with non-numeric values in either `train_label_column` or `test_label_column` and/or those with non-numeric values for relevant features will be automatically excluded from model training and evaluation. By default, zero scores in either `train_label_column` or `test_label_column` will also be excluded. See [\*exclude\\_zero\\_scores \(Optional\)\*](#) if you want to keep responses with zero scores.

---

**trim\_max (Optional)**

The single numeric value for the highest possible integer score that the machine should predict. This value will be used to compute the ceiling value for *trimmed (bound)* machine scores as `trim_max + trim_tolerance`. Defaults to the highest observed human score in the training data or 10 if there are no numeric human scores available.

**trim\_min (Optional)**

The single numeric value for the lowest possible integer score that the machine should predict. This value will be used to compute the floor value for *trimmed (bound)* machine scores as `trim_min - trim_tolerance`. Defaults to the lowest observed human score in the training data or 1 if there are no numeric human scores available.

**trim\_tolerance (Optional)**

The single numeric value that will be used to pad the trimming range specified in `trim_min` and `trim_max`. This value will be used to compute the ceiling and floor values for *trimmed (bound)* machine scores as `trim_max + trim_tolerance` for ceiling value and `trim_min - trim_tolerance` for floor value. Defaults to 0.4998.

---

**Note:** For more fine-grained control over the trimming range, you can set `trim_tolerance` to 0 and use `trim_min` and `trim_max` to specify the exact floor and ceiling values.

---

### `truncate_outliers` (Optional)

If this option is set to `false`, outliers (values more than 4 standard deviations away from the mean) in feature columns will `_not_` be truncated. Defaults to `true`.

### `use_scaled_predictions` (Optional)

If set to `true`, certain evaluations (confusion matrices, score distributions, subgroup analyses) will use the scaled machine scores. If set to `false`, these evaluations will use the raw machine scores. Defaults to `false`.

---

**Note:** All evaluation metrics (e.g., kappa and pearson correlation) are automatically computed for *both* scaled and raw scores.

---

### `use_thumbnails` (Optional)

If set to `true`, the images in the HTML will be set to clickable thumbnails rather than full-sized images. Upon clicking the thumbnail, the full-sized images will be displayed in a separate tab in the browser. If set to `false`, full-sized images will be displayed as usual. Defaults to `false`.

### `use_truncation_thresholds` (Optional)

If set to `true`, use the `min` and `max` columns specified in the `features` file to clamp outlier feature values. This is useful if users would like to clamp feature values based on some pre-defined boundaries, rather than having these boundaries calculated based on the training set. Defaults to `false`.

---

**Note:** If `use_truncation_thresholds` is set, a `features` file `_must_` be specified, and this file `_must_` include `min` and `max` columns. If no `feature` file is specified or these columns are missing, an error will be raised.

---

### `use_wandb` (Optional)

If set to `true`, the generated reports and all intermediate tables will be logged to [Weights & Biases](#). The [Weights & Biases](#) *entity* and *project name* can be specified in the appropriate configuration fields.

The tables and plots will be logged in a section named *rsmtool* in a new run under the given project, and the report will be added to a *reports* section in that run. In addition, some evaluation metrics will be logged to the run's history for easy comparison between different runs in the project. With each metric name representing the context and the table from which the metric is taken, for example: *rsmeval/consistency.SMD* is the *SMD* metric from the *consistency* table calculated by *rsmtool*. Defaults to `false`.



**wandb\_project (Optional)**

The Weights & Biases project name if logging to Weights & Biases is enabled. If a project by this name does not already exist, it will be created.

---

**Important:**

1. Before using Weights & Biases for the first time, users should log in and provide their API key as described in [W&B Quickstart guidelines](#).
  2. Note that when using W&B logging, the rsmttool run may take significantly longer due to the network traffic being sent to W&B.
- 

**wandb\_entity (Optional)**

The Weights & Biases entity name if logging to Weights & Biases is enabled. Entity can be a user name or the name of a team or organization.

## 1.6.2 Output

rsmttool produces a set of folders in the experiment output directory. This is either the current directory in which rsmttool is run or the directory specified as the second optional command-line argument. If logging to [Weights & Biases](#) is *enabled*, the reports and all intermediate files are also logged to the specified Weights & Biases project.

**report**

This folder contains the final RSMTool report in HTML format as well in the form of a Jupyter notebook (a `.ipynb` file).

**output**

This folder contains all of the *intermediate files* produced as part of the various analyses performed, saved as `.csv` files. rsmttool will also save in this folder a copy of the *configuration file*. Fields not specified in the original configuration file will be pre-populated with default values.

**figure**

This folder contains all of the figures generated as part of the various analyses performed, saved as `.svg` files.

## feature

This folder contains a `.csv` file that lists all features, signs and transformation as used in the *final* model, taking into account any manual or automatic feature selection. See [feature column selection](#) for more details.

### 1.6.3 Selecting Feature Columns

By default, `rsmttool` will use all columns included in the training and evaluation data files as features. The only exception are any columns explicitly identified in the configuration file as containing non-feature information (e.g., `id_column`, `train_label_column`, `test_label_column`, etc.)

However, there are certain scenarios in which it is useful to choose specific columns in the data to be used as features. For example, let's say that you have a large number of very different features and you want to use a different subset of features to score different types of questions on a test. In this case, the ability to easily choose the desired features for any `rsmttool` experiment becomes quite important. The alternative of manually pre-processing the data to remove the features you don't need is quite cumbersome.

There are two ways to select specific columns in the data as features:

1. **Fine-grained column selection:** In this method, you manually create a list of the columns that you wish to use as features for an `rsmttool` experiment. See [fine-grained selection](#) for more details.
2. **Subset-based column selection:** In this method, you can pre-define subsets of features and then select entire subsets at a time for any `rsmttool` experiment. See [subset-based selection](#) for more details.

While fine-grained column selection is better for a single experiment, subset-based selection is more convenient when you need to run several experiments with somewhat different subsets of features.

**Warning:** `rsmttool` will filter the training and evaluation data to remove any responses with non-numeric values in any of the feature columns before training the model. If your data includes a column containing string values and you do *not* use any of these methods of feature selection *nor* specify this column as the `id_column` or the `candidate_column` or a subgroup column, `rsmttool` will filter out *all* the responses in the data.

#### Fine-grained column selection

To manually select columns to be used as features, you can provide a data file in one of the [supported formats](#). The file must contain a column named `feature` which specifies the names of the feature columns that should be used for scoring model building. For additional flexibility, the same file also allows you to describe transformations to be applied to the values in these feature columns before being used in the model. The path to this file should be set as an argument to `features` in the experiment configuration file. (Note: If you do not wish to perform any feature transformations, but would simply like to select certain feature columns to include, you can also pass a list of feature names as an argument to `features`.)

Here's an example of what such a file might look like.

```
feature,transform,sign
feature1,raw,1
feature2,inv,-1
```

There is one required column and two optional columns.

## feature

The exact name of the column in the training and evaluation data files, including capitalization. Column names cannot contain hyphens. The following strings are reserved and cannot not be used as feature column names: `spkitemid`, `spkitemlab`, `itemType`, `r1`, `r2`, `score`, `sc`, `sc1`, and `adj`. In addition, any column names provided as values for `id_column`, `train_label_column`, `test_label_column`, `length_column`, `candidate_column`, and `subgroups` may also not be used as feature column names.

## transform (optional)

A transformation that should be applied to the column values before using it in the model. Possible values are:

- `raw`: no transformation, use original value
- `org`: same as raw
- `inv`:  $1/x$
- `sqrt`: square root
- `addOneInv`:  $1/(x+1)$
- `addOneLn`:  $\ln(x+1)$

Note that `rsmttool` will raise an exception if the values in the data do not allow the supplied transformation (for example, if `inv` is applied to a column which has 0 values). If you really want to use the transformation, you must pre-process your training and evaluation data files to remove the problematic cases.

If the feature file contains no `transform` column, `rsmttool` will use the original values for all features (`raw` transform).

## sign (optional)

After transformation, the column values will be multiplied by this number, which can be either 1 or -1 depending on the expected sign of the correlation between transformed feature and human score. This mechanism is provided to ensure that all features in the final models have a positive correlation with the score, if that is so desired by the user.

If the feature file contains no `sign` column, `rsmttool` will multiply all values by 1.

When determining the sign, you should take into account the correlation between the original feature and the score as well as any applied transformations. For example, if you use feature which has a negative correlation with the human score and apply `sqrt` transformation, `sign` should be set to -1. However, if you use the same feature but apply the `inv` transformation, `sign` should now be set to 1.

To ensure that this is working as expected, you can check the sign of correlations for both raw and processed features in the final report.

---

**Note:** You can use the fine-grained method of column selection in combination with a *model with automatic feature selection*. In this case, the features that end up being used in the final model can be found in the `.csv` file in the `feature` folder in the experiment output directory.

---

## Subset-based column selection

For more advanced users, `rsmttool` offers the ability to assign columns to named subsets in a data file in one of the *supported formats* and then select a set of columns by simply specifying the name of that pre-defined subset.

If you want to run multiple `rsmttool` experiments, each choosing from a large number of features, generating a separate *feature file* for each experiment listing columns to use can quickly become tedious.

Instead you can define feature subsets by providing a subset definition file in one of the *supported formats* which lists *all* feature names under a column named `feature`. Each subset is an additional column with a value of either 0 (denoting that the feature does *not* belong to the subset named by that column) or 1 (denoting that the feature does belong to the subset named by that column).

Here's an example of a subset definition file, say `subset.csv`.

```
feature,A,B
feature1,0,1
feature2,1,1
feature3,1,0
```

In this example, `feature2` and `feature3` belong to a subset called "A" and `feature1` and `feature1` and `feature2` belong to a subset called "B".

This feature subset file can be provided to `rsmttool` using the *feature\_subset\_file* field in the configuration file. Then, to select a particular pre-defined subset of features, you simply set the *feature\_subset* field in the configuration file to the name of the subset that you wish to use.

Then, in order to use feature subset "A" (`feature2` and `feature3`) in an experiment, we need to set the following two fields in our experiment configuration file:

```
{
  ...
  "feature_subset_file": "subset.csv",
  "feature_subset": "A",
  ...
}
```

## Transformations

Unlike in *fine-grained selection*, the feature subset file does not list any transformations to be applied to the feature columns. However, you can automatically select transformation for each feature *in the selected subset* by applying all possible transforms and identifying the one which gives the highest correlation with the human score. To use this functionality set the *select\_transformations* field in the configuration file to `true`.

## Signs

Some guidelines for building scoring models require all coefficients in the model to be positive and all features to have a positive correlation with human score. `rsmttool` can automatically flip the sign for any pre-defined feature subset. To use this functionality, the feature subset file should provide the expected correlation sign between each feature and human score under a column called `sign_<SUBSET>` where `<SUBSET>` is the name of the feature subset. Then, to tell `rsmttool` to flip the sign for this subset, you need to set the *sign* field in the configuration file to `<SUBSET>`.

To understand this, let's re-examine our earlier example of a subset definition file `subset.csv`, but with an additional column.

```
feature,A,B,sign_A
feature1,0,1,+
feature2,1,1,-
feature3,1,0,+
```

Then, in order to use feature subset “A” (feature2 and feature3) in an experiment with the sign of feature3 flipped appropriately (multiplied by -1) to ensure positive correlations with score, we need to set the following three fields in our experiment configuration file:

```
{
  ...
  "feature_subset_file": "subset.csv",
  "feature_subset": "A",
  "sign": "A"
  ...
}
```

**Note:** If *select\_transformations* is set to `true`, `rsmttool` is intelligent enough to take it into account when flipping the signs. For example, if the expected correlation sign for a given feature is negative, `rsmttool` will multiply the feature values by `-1` if the `sqrt` transform has the highest correlation with score. However, if the best transformation turns out to be `inv` – which already changes the polarity of the feature – no such multiplication will take place.

### 1.6.4 Intermediate files

Although the primary output of RSMTool is an HTML report, we also want the user to be able to conduct additional analyses outside of RSMTool. To this end, all of the tables produced in the experiment report are saved as files in the format as specified by `file_format` parameter in the output directory. The following sections describe all of the intermediate files that are produced.

**Note:** The names of all files begin with the `experiment_id` provided by the user in the experiment configuration file. In addition, the names for certain columns are set to default values in these files irrespective of what they were named in the original data files. This is because RSMTool standardizes these column names internally for convenience. These values are:

- `spkitemid` for the column containing response IDs.
- `sc1` for the column containing the human scores used as training labels.
- `sc2` for the column containing the second human scores, if this column was specified in the configuration file.
- `length` for the column containing response length, if this column was specified in the configuration file.
- `candidate` for the column containing candidate IDs, if this column was specified in the configuration file.

### Feature values

filenames: `train_features`, `test_features`, `train_preprocessed_features`,  
`test_preprocessed_features`

These files contain the raw and pre-processed feature values for the training and evaluation sets. They include *only* includes the rows that were used for training/evaluating the models after filtering. For models with feature selection, these files *only* include the features that ended up being included in the model.

---

**Note:** By default RSMTTool filters out non-numeric feature values and non-numeric/zero human scores from both the training and evaluation sets. Zero scores can be kept by setting the `exclude_zero_scores` to `false`.

---

### Flagged responses

filenames: `train_responses_with_excluded_flags`, `test_responses_with_excluded_flags`

These files contain all of the rows in the training and evaluation sets that were filtered out based on conditions specified in `flag_column`.

---

**Note:** If the training/evaluation files contained columns with internal names such as `sc1` or `length` but these columns were not actually used by `rsmttool`, these columns will also be included into these files but their names will be changed to `##name##` (e.g. `##sc1##`).

---

### Excluded responses

filenames: `train_excluded_responses`, `test_excluded_responses`

These files contain all of the rows in the training and evaluation sets that were filtered out because of feature values or scores. For models with feature selection, these files *only* include the features that ended up being included in the model.

### Response metadata

filenames: `train_metadata`, `test_metadata`

These files contain the metadata columns (`id_column`, subgroups if provided) for the rows in the training and evaluation sets that were *not* excluded for some reason.

### Unused columns

filenames: `train_other_columns`, `test_other_columns`

These files contain all of the columns from the original features files that are not present in the `*_feature` and `*_metadata` files. They only include the rows from the training and evaluation sets that were not filtered out.

---

**Note:** If the training/evaluation files contained columns with internal names such as `sc1` or `length` but these columns were not actually used by `rsmttool`, these columns will also be included into these files but their names will be changed to `##name##` (e.g. `##sc1##`).

---

## Response length

filename: `train_response_lengths`

If *length\_column* is specified, then this file contains the values from that column for the training data under a column called `length` with the response IDs under the `spkitemid` column.

## Human scores

filename: `test_human_scores`

This file contains the human scores for the evaluation data under a column called `sc1` with the response IDs under the `spkitemid` column. If *second\_human\_score\_column* was specified, then it also contains the values from that column under a column called `sc2`. Only the rows that were not filtered out are included.

---

**Note:** If *exclude\_zero\_scores* was set to `true` (the default value), all zero scores in the *second\_human\_score\_column* will be replaced by `nan`.

---

## Data composition

filename: `data_composition`

This file contains the total number of responses in training and evaluation set and the number of overlapping responses. If applicable, the table will also include the number of different subgroups for each set.

## Excluded data composition

filenames: `train_excluded_composition`, `test_excluded_composition`

These files contain the composition of the set of excluded responses for the training and evaluation sets, e.g., why were they excluded and how many for each such exclusion.

## Missing features

filename: `train_missing_feature_values`

This file contains the total number of non-numeric values for each feature. The counts in this table are based only on those responses that have a numeric human score in the training data.

## Subgroup composition

filename: `data_composition_by_<SUBGROUP>`

There will be one such file for each of the specified subgroups and it contains the total number of responses in that subgroup in both the training and evaluation sets.

### Feature descriptives

filenames: `feature_descriptives`, `feature_descriptivesExtra`

The first file contains the main descriptive statistics (mean, std. dev., correlation with human score etc.) for all features included in the final model. The second file contains percentiles, mild, and extreme outliers for the same set of features. The values in both files are computed on raw feature values before pre-processing.

### Feature outliers

filename: `feature_outliers`

This file contains the number and percentage of outlier values truncated to  $[\text{MEAN}-4\text{SD}, \text{MEAN}+4\text{SD}]$  during feature pre-processing for each feature included in the final model.

### Inter-feature and score correlations

filenames: `cors_orig`, `cors_processed`

The first file contains the pearson correlations between each pair of (raw) features and between each (raw) feature and the human score. The second file is the same but with the pre-processed feature values instead of the raw values.

### Marginal and partial correlations with score

filenames: `margcor_score_all_data`, `pcor_score_all_data`, ``pcor_score_no_length_all_data`

The first file contains the marginal correlations between each pre-processed feature and human score. The second file contains the partial correlation between each pre-processed feature and human score after controlling for all other features. The third file contains the partial correlations between each pre-processed feature and human score after controlling for response length, if `length_column` was specified in the configuration file.

### Marginal and partial correlations with length

filenames: `margcor_length_all_data`, `pcor_length_all_data`

The first file contains the marginal correlations between each pre-processed feature and response length, if `length_column` was specified. The second file contains the partial correlations between each pre-processed feature and response length after controlling for all other features, if `length_column` was specified in the configuration file.

### Principal components analyses

filenames: `pca`, `pcavar`

The first file contains the results of a Principal Components Analysis (PCA) using pre-processed feature values from the training set and its singular value decomposition. The second file contains the eigenvalues and variance explained by each component.



## Various correlations by subgroups

Each of following files may be produced for every subgroup, assuming all other information was also available.

- `margcor_score_by_<SUBGROUP>`: the marginal correlations between each pre-processed feature and human score, computed separately for the subgroup.
- `pcor_score_by_<SUBGROUP>`: the partial correlations between pre-processed features and human score after controlling for all other features, computed separately for the subgroup.
- `pcor_score_no_length_by_<SUBGROUP>`: the partial correlations between each pre-processed feature and human score after controlling for response length (if available), computed separately for the subgroup.
- `margcor_length_by_<SUBGROUP>`: the marginal correlations between each feature and response length (if available), computed separately for each subgroup.
- `pcor_length_by_<SUBGROUP>`: partial correlations between each feature and response length (if available) after controlling for all other features, computed separately for each subgroup.

---

**Note:** All of the feature descriptive statistics, correlations (including those for subgroups), and PCA are computed *only* on the training set.

---

## Model information

- `feature`: *pre-processing parameters* for all features used in the model.
- `coefficients`: model coefficients and intercept (for *built-in models* only).
- `coefficients_scaled`: scaled model coefficients and intercept (linear models only). Although RSMTTool generates scaled scores by scaling the predictions of the model, it is also possible to achieve the same result by scaling the coefficients instead. This file shows those scaled coefficients.
- `betas`: standardized and relative coefficients (for built-in models only).
- `model_fit`: R squared and adjusted R squared computed on the training set. Note that these values are always computed on raw predictions without any trimming or rounding.
- `.model`: the serialized RSMTTool `Modeler` object containing the fitted SKLL `Learner` object (before scaling the coefficients).
- `.ols`: a serialized object of type `pandas.stats.ols.OLS` containing the fitted model (for built-in models excluding `LassoFixedLambda` and `PositiveLassoCV`).
- `ols_summary.txt`: a text file containing a summary of the above model (for built-in models excluding `LassoFixedLabmda` and `PositiveLassoCV`)
- `postprocessing_params`: the parameters for trimming and scaling predicted scores. Useful for generating predictions on new data.

## Predictions

filenames: `pred_processed`, `pred_train`

The first file contains the predicted scores for the evaluation set and the second file contains the predicted scores for the responses in the training set. Both of them contain the raw scores as well as different types of post-processed scores.

## Evaluation metrics

- **eval**: This file contains the descriptives for predicted and human scores (mean, std.dev etc.) as well as the association metrics (correlation, quadratic weighted kappa, SMD etc.) for the raw as well as the post-processed scores.
- **eval\_by\_<SUBGROUP>**: the same information as in `*_eval.csv` computed separately for each subgroup. However, rather than SMD, a difference of standardized means (DSM) will be calculated using z-scores.
- **eval\_short**: a shortened version of **eval** that contains specific descriptives for predicted and human scores (mean, std.dev etc.) and association metrics (correlation, quadratic weighted kappa, SMD etc.) for specific score types chosen based on recommendations by Williamson (2012). Specifically, the following columns are included (the raw or scale version is chosen depending on the value of the `use_scaled_predictions` in the configuration file).
  - `h_mean`
  - `h_sd`
  - `corr`
  - `sys_mean [raw/scale_trim]`
  - `sys_sd [raw/scale_trim]`
  - `SMD [raw/scale_trim]`
  - `adj_agr [raw/scale_trim_round]`
  - `exact_agr [raw/scale_trim_round]`
  - `kappa [raw/scale_trim_round]`
  - `wtkappa [raw/scale_trim]`
  - `sys_mean [raw/scale_trim_round]`
  - `sys_sd [raw/scale_trim_round]`
  - `SMD [raw/scale_trim_round]`
  - `R2 [raw/scale_trim]`
  - `RMSE [raw/scale_trim]`
- **score\_dist**: the distributions of the human scores and the rounded raw/scaled predicted scores, depending on the value of `use_scaled_predictions`.
- **confMatrix**: the confusion matrix between the human scores and the rounded raw/scaled predicted scores, depending on the value of `use_scaled_predictions`.

---

**Note:** Please note that for raw scores, SMD values are likely to be affected by possible differences in scale.

---

- **true\_score\_eval**: evaluation of how well system scores can predict true scores.

## Human-human Consistency

These files are created only if a second human score has been made available via the `second_human_score_column` option in the configuration file.

- `consistency`: contains descriptives for both human raters as well as the agreement metrics between their ratings.
- `consistency_by_<SUBGROUP>`: contains the same metrics as in `consistency` file computed separately for each group. However, rather than SMD, a difference of standardized means (DSM) will be calculated using z-scores.
- `degradation`: shows the differences between human-human agreement and machine-human agreement for all association metrics and all forms of predicted scores.
- `confMatrix_h1h2`: the confusion matrix between the human scores for double-scored responses.

## Evaluations based on test theory

- `disattenuated_correlations`: shows the correlation between human-machine scores, human-human scores, and the disattenuated human-machine correlation computed as human-machine correlation divided by the square root of human-human correlation.
- `disattenuated_correlations_by_<SUBGROUP>`: contains the same metrics as in `disattenuated_correlations` file computed separately for each group.
- `true_score_eval`: evaluations of system scores against estimated true score. Contains total counts of single and double-scored response, variance of human rater error, estimated true score variance, and mean squared error (MSE) and proportional reduction in mean squared error (PRMSE) when predicting true score using system score.

## Additional fairness analyses

These files contain the results of additional fairness analyses suggested in suggested in Loukina, Madnani, & Zechner, 2019.

- `<METRICS>_by_<SUBGROUP>.ols`: a serialized object of type `pandas.stats.ols.OLS` containing the fitted model for estimating the variance attributed to a given subgroup membership for a given metric. The subgroups are defined by the [configuration file](#). The metrics are `osa` (overall score accuracy), `osd` (overall score difference), and `csd` (conditional score difference).
- `<METRICS>_by_<SUBGROUP>.ols_summary.txt`: a text file containing a summary of the above model
- `estimates_<METRICS>_by_<SUBGROUP>``: coefficients, confidence intervals and *p*-values estimated by the model for each subgroup.
- `fairness_metrics_by_<SUBGROUP>`: the  $R^2$  (percentage of variance) and *p*-values for all models.

## 1.6.5 Built-in RSMTTool Linear Regression Models

### Models which use the full feature set

- **LinearRegression**: A model that learns empirical regression weights using ordinary least squares regression (OLS).
- **EqualWeightsLR**: A model with all feature weights set to 1.0; a naive model.
- **ScoreWeightedLR**: a model that learns empirical regression weights using weighted least squares. The weights are determined based on the number of responses with different score levels. Score levels with lower number of responses are assigned higher weight.
- **RebalancedLR** - empirical regression weights are rebalanced by using a small portion of positive weights to replace negative beta values. This model has no negative coefficients.

### Models with automatic feature selection

- **LassoFixedLambdaThenLR**: A model that learns empirical OLS regression weights with feature selection using Lasso regression with all coefficients set to positive. The hyperparameter `lambda` is set to  $\sqrt{n - \lg(p)}$  where  $n$  is the number of responses and  $p$  is the number of features. This approach was chosen to balance the penalties for error vs. penalty for too many coefficients to force Lasso perform more aggressive feature selection, so it may not necessarily achieve the best possible performance. The feature set selected by LASSO is then used to fit an OLS linear regression. Note that while the original Lasso model is constrained to positive coefficients only, small negative coefficients may appear when the coefficients are re-estimated using OLS regression.
- **PositiveLassoCVThenLR**: A model that learns empirical OLS regression weights with feature selection using Lasso regression with all coefficients set to positive. The hyperparameter `lambda` is optimized using cross-validation for loglikelihood. The feature set selected by LASSO is then used to fit an OLS linear regression. Note that this approach will likely produce a model with a large  $N$  features and any advantages of running Lasso would be effectively negated by later adding those features to OLS regression.
- **NNLR**: A model that learns empirical OLS regression weights with feature selection using non-negative least squares regression. Note that only the coefficients are constrained to be positive: the intercept can be either positive or negative.
- **NNLRIterative**: A model that learns empirical OLS regression weights with feature selection using an iterative implementation of non-negative least squares regression. Under this implementation, an initial OLS model is fit. Then, any variables whose coefficients are negative are dropped and the model is re-fit. Any coefficients that are still negative after re-fitting are set to zero.
- **LassoFixedLambdaThenNNLR**: A model that learns empirical OLS regression weights with feature selection using Lasso regression as above followed by non-negative least squares regression. The latter ensures that no feature has negative coefficients even when the coefficients are estimated using least squares without penalization.
- **LassoFixedLambda**: same as **LassoFixedLambdaThenLR** but the model uses the original Lasso weights. Note that the coefficients in Lasso model are estimated using an optimization routine which may produce slightly different results on different systems.
- **PositiveLassoCV**: same as *PositiveLassoCVThenLR* but using the original Lasso weights. Please note: the coefficients in Lasso model are estimated using an optimization routine which may produce slightly different results on different systems.

---

#### Note:

1. NNLR, NNLRIterative, LassoFixedLambdaThenNNLR, LassoFixedLambda and PositiveLassoCV all have no negative coefficients.

2. For all feature selection models, the final set of features will be saved in the `feature` folder in the experiment output directory.
- 

## 1.7 Advanced Uses of RSMTTool

In addition to providing the `rsmttool` utility training and evaluating regression-based scoring models, the RSMTTool package also provides six other command-line utilities for more advanced users.

### 1.7.1 `rsmeval` - Evaluate external predictions

RSMTTool provides the `rsmeval` command-line utility to evaluate existing predictions and generate a report with all the built-in analyses. This can be useful in scenarios where the user wants to use more sophisticated machine learning algorithms not available in RSMTTool to build the scoring model but still wants to be able to evaluate that model's predictions using the standard analyses.

For example, say a researcher *has* an existing automated scoring engine for grading short responses that extracts the features and computes the predicted score. This engine uses a large number of binary, sparse features. She cannot use `rsmttool` to train her model since it requires numeric features. So, she uses `scikit-learn` to train her model.

Once the model is trained, the researcher wants to evaluate her engine's performance using the analyses recommended by the educational measurement community as well as conduct additional investigations for specific subgroups of test-takers. However, these kinds of analyses are not available in `scikit-learn`. She can use `rsmeval` to set up a customized report using a combination of existing and custom sections and quickly produce the evaluation that is useful to her.

#### Tutorial

For this tutorial, you first need to *install RSMTTool* and make sure the conda environment in which you installed it is activated.

#### Workflow

`rsmeval` is designed for evaluating existing machine scores. Once you have the scores computed for all the responses in your data, the next steps are fairly straightforward:

1. Create a data file in one of the *supported formats* containing the computed system scores and the human scores you want to compare against.
2. Create an *experiment configuration file* describing the evaluation experiment you would like to run.
3. Run that configuration file with `rsmeval` and generate the experiment HTML report as well as the *intermediate CSV files*.
4. Examine the HTML report to check various aspects of model performance.

Note that the above workflow does not use any customization features, e.g., *choosing which sections to include in the report* or *adding custom analyses sections* etc. However, we will stick with this workflow for our tutorial since it is likely to be the most common use case.

### ASAP Example

We are going to use the same example from 2012 Kaggle competition on automated essay scoring that we used for the *rsmtool tutorial*.

### Generate scores

rsmeval is designed for researchers who have developed their own scoring engine for generating scores and would like to produce an evaluation report for those scores. For this tutorial, we will use the scores we generated for the ASAP2 evaluation set using *rsmtool tutorial*.

### Create a configuration file

The next step is to create an *experiment configuration file* in .json format.

```
1 {  
2   "experiment_id": "ASAP2_evaluation",  
3   "description": "Evaluation of the scores generated using rsmtool.",  
4   "predictions_file": "ASAP2_scores.csv",  
5   "system_score_column": "system",  
6   "human_score_column": "human",  
7   "id_column": "ID",  
8   "trim_min": 1,  
9   "trim_max": 6,  
10  "second_human_score_column": "human2",  
11  "scale_with": "asis"  
12 }
```

Let's take a look at the options in our configuration file.

- **Line 2:** We define an experiment ID.
- **Line 3:** We also provide a description which will be included in the experiment report.
- **Line 4:** We list the path to the file with the predicted and human scores. For this tutorial we used .csv format, but RSMTool also supports several other *input file formats*.
- **Line 5:** This field indicates that the system scores in our .csv file are located in a column named `system`.
- **Line 6:** This field indicates that the human (reference) scores in our .csv file are located in a column named `human`.
- **Line 7:** This field indicates that the unique IDs for the responses in the .csv file are located in columns named `ID`.
- **Lines 8-9:** These fields indicate that the lowest score on the scoring scale is a 1 and the highest score is a 6. This information is usually part of the rubric used by human graders.
- **Line 10:** This field indicates that scores from a second set of human graders are also available (useful for comparing the agreement between human-machine scores to the agreement between two sets of humans) and are located in the `human2` column in the .csv file.
- **Line 11:** This field indicates that the provided machine scores are already re-scaled to match the distribution of human scores. rsmeval itself will not perform any scaling and the report will refer to these as `scaled` scores.

Documentation for all of the available configuration options is available [here](#).

**Note:** You can also use our nifty capability to *automatically generate* `rsmeval` configuration files rather than creating them manually.

## Run the experiment

Now that we have our scores in the right format and our configuration file in `.json` format, we can use the *rsmeval* command-line script to run our evaluation experiment.

```
$ cd examples/rsmeval
$ rsmeval config_rsmeval.json
```

This should produce output like:

```
Output directory: /Users/nmadnani/work/rsmtool/examples/rsmeval
Assuming given system predictions are already scaled and will be used as such.
  predictions: /Users/nmadnani/work/rsmtool/examples/rsmeval/ASAP2_scores.csv
Processing predictions
Saving pre-processed predictions and the metadata to disk
Running analyses on predictions
Starting report generation
Merging sections
Exporting HTML
Executing notebook with kernel: python3
```

Once the run finishes, you will see the `output`, `figure`, and `report` sub-directories in the current directory. Each of these directories contain *useful information* but we are specifically interested in the `report/ASAP2_evaluation_report.html` file, which is the final evaluation report.

## Examine the report

Our experiment report contains all the information we would need to evaluate the provided system scores against the human scores. It includes:

1. The distributions for the human versus the system scores.
2. Several different metrics indicating how well the machine's scores agree with the humans'.
3. Information about human-human agreement and the difference between human-human and human-system agreement.

... and *much more*.

## Input

`rsmeval` requires a single argument to run an experiment: the path to *a configuration file*. It can also take an output directory as an optional second argument. If the latter is not specified, `rsmeval` will use the current directory as the output directory.

Here are all the arguments to the `rsmeval` command-line script.

### `config_file`

The *JSON configuration file* for this experiment.

### **output\_dir** (optional)

The output directory where all the files for this experiment will be stored.

### **-f, --force**

If specified, the contents of the output directory will be overwritten even if it already contains the output of another rsmeval experiment.

### **-h, --help**

Show help message and exist.

### **-V, --version**

Show version number and exit.

## Experiment configuration file

This is a file in `.json` format that provides overall configuration options for an `rsmeval` experiment. Here's an [example configuration file](#) for `rsmeval`.

---

**Note:** To make it easy to get started with `rsmeval`, we provide a way to **automatically generate** configurations file both interactively as well as non-interactively. Novice users will find interactive generation more helpful while more advanced users will prefer non-interactive generation. See [this page](#) for more details.

---

Next, we describe all of the `rsmeval` configuration fields in detail. There are four required fields and the rest are all optional. We first describe the required fields and then the optional ones (sorted alphabetically).

### **experiment\_id**

An identifier for the experiment that will be used to name the report and all *intermediate files*. It can be any combination of alphanumeric values, must *not* contain spaces, and must *not* be any longer than 200 characters.

### **predictions\_file**

The path to the file with predictions to evaluate. The file should be in one of the *supported formats*. Each row should correspond to a single response and contain the predicted and observed scores for this response. In addition, there should be a column with a unique identifier (ID) for each response. The path can be absolute or relative to the location of the configuration file.

### **system\_score\_column**

The name for the column containing the scores predicted by the system. These scores will be used for evaluation.



**trim\_min**

The single numeric value for the lowest possible integer score that the machine should predict. This value will be used to compute the floor value for *trimmed (bound)* machine scores as `trim_min - trim_tolerance`.

**trim\_max**

The single numeric value for the highest possible integer score that the machine should predict. This value will be used to compute the ceiling value for *trimmed (bound)* machine scores as `trim_max + trim_tolerance`.

---

**Note:** Although the `trim_min` and `trim_max` fields are optional for `rsmtool`, they are *required* for `rsmeval`.

---

**candidate\_column (Optional)**

The name for an optional column in prediction file containing unique candidate IDs. Candidate IDs are different from response IDs since the same candidate (test-taker) might have responded to multiple questions.

**custom\_sections (Optional)**

A list of custom, user-defined sections to be included into the final report. These are IPython notebooks (`.ipynb` files) created by the user. The list must contain paths to the notebook files, either absolute or relative to the configuration file. All custom notebooks have access to some *pre-defined variables*.

**description (Optional)**

A brief description of the experiment. This will be included in the report. The description can contain spaces and punctuation. It's blank by default.

**exclude\_zero\_scores (Optional)**

By default, responses with human scores of 0 will be excluded from evaluations. Set this field to `false` if you want to keep responses with scores of 0. Defaults to `true`.

**file\_format (Optional)**

The format of the *intermediate files*. Options are `csv`, `tsv`, or `xlsx`. Defaults to `csv` if this is not specified.

### flag\_column (Optional)

This field makes it possible to only use responses with particular values in a given column (e.g. only responses with a value of 0 in a column called ADVISORY). The field takes a dictionary in Python format where the keys are the names of the columns and the values are lists of values for responses that will be evaluated. For example, a value of {"ADVISORY": 0} will mean that rsmeval will *only* use responses for which the ADVISORY column has the value 0. Defaults to None.

---

**Note:** If several conditions are specified (e.g., {"ADVISORY": 0, "ERROR": 0}) only those responses which satisfy *all* the conditions will be selected for further analysis (in this example, these will be the responses where the ADVISORY column has a value of 0 *and* the ERROR column has a value of 0).

---



---

**Note:** When reading the values in the supplied dictionary, rsmeval treats numeric strings, floats and integers as the same value. Thus 1, 1.0, "1" and "1.0" are all treated as the 1.0.

---

### general\_sections (Optional)

RSMTool provides pre-defined sections for rsmeval (listed below) and, by default, all of them are included in the report. However, you can choose a subset of these pre-defined sections by specifying a list as the value for this field.

- **data\_description:** Shows the total number of responses, along with any responses have been excluded due to non-numeric/zero scores or *flag columns*.
- **data\_description\_by\_group:** Shows the total number of responses for each of the *subgroups* specified in the configuration file. This section only covers the responses used to evaluate the model.
- **consistency:** shows metrics for *human-human agreement*, the difference (“degradation”) between the human-human and human-system agreement, and the disattenuated human-machine correlations. This notebook is only generated if the config file specifies *second\_human\_score\_column*.
- **evaluation:** Shows the *standard set of evaluations* recommended for scoring models on the evaluation data:
  - a table showing human-system association metrics;
  - the confusion matrix; and
  - a barplot showing the distributions for both human and machine scores.
- **evaluation by group:** Shows barplots with the main evaluation metrics by each of the subgroups specified in the configuration file.
- **fairness\_analyses:** Additional *fairness analyses* suggested in Loukina, Madnani, & Zechner, 2019. The notebook shows:
  - percentage of variance in squared error explained by subgroup membership
  - percentage of variance in raw (signed) error explained by subgroup membership
  - percentage of variance in raw (signed) error explained by subgroup membership when controlling for human score
  - plots showing estimates for each subgroup for each model
- **true\_score\_evaluation:** evaluation of system scores against the true scores estimated according to test theory. The notebook shows:
  - Number of single and double-scored responses.

- Variance of human rater errors and estimated variance of true scores
- Mean squared error (MSE) and proportional reduction in mean squared error (PRMSE) when predicting true score with system score.
- `intermediate_file_paths`: Shows links to all of the intermediate files that were generated while running the evaluation.
- `sysinfo`: Shows all Python packages along with versions installed in the current environment while generating the report.

### `human_score_column` (*Optional*)

The name for the column containing the human scores for each response. The values in this column will be used as observed scores. Defaults to `sc1`.

---

**Note:** All responses with non-numeric values or zeros in either `human_score_column` or `system_score_column` will be automatically excluded from evaluation. You can use [\*`exclude\_zero\_scores`\*](#) (*Optional*) to keep responses with zero scores.

---

### `id_column` (*Optional*)

The name of the column containing the response IDs. Defaults to `spkitemid`, i.e., if this is not specified, `rsmeval` will look for a column called `spkitemid` in the prediction file.

### `min_items_per_candidate` (*Optional*)

An integer value for the minimum number of responses expected from each candidate. If any candidates have fewer responses than the specified value, all responses from those candidates will be excluded from further analysis. Defaults to `None`.

### `min_n_per_group` (*Optional*)

A single numeric value or a dictionary with keys as the group names listed in the `subgroups` field and values as the thresholds for the groups. When specified, only groups with *at least* this number of instances will be displayed in the tables and plots contained **in the report**. Note that this parameter *only* affects the HTML report and the figures. For all analyses – including the computation of the population parameters – data from *all* groups will be used. In addition, the *intermediate files* will still show the results for *all* groups.

---

**Note:** If you supply a dictionary, it *must* contain a key for *every* subgroup listed in `subgroups` field. If no threshold is to be applied for some of the groups, set the threshold value for this group to 0 in the dictionary.

---

### **rater\_error\_variance** (*Optional*)

*True score evaluations* require an estimate of rater error variance. By default, `rsmeval` will compute this variance from double-scored responses in the data. However, in some cases, one may wish to compute the variance on a different sample of responses. In such cases, this field can be used to set the rater error variance to a precomputed value which is then used as-is by `rsmeval`. You can use the `rsmttool.utils.variance_of_errors` function to compute rater error variance outside the main evaluation pipeline.

### **scale\_with** (*Optional*)

In many scoring applications, system scores are *re-scaled* so that their mean and standard deviation match those of the human scores for the training data.

If you want `rsmeval` to re-scale the supplied predictions, you need to provide – as the value for this field – the path to a second file in one of the *supported formats* containing the human scores and predictions of the same system on its training data. This file *must* have two columns: the human scores under the `sc1` column and the predicted score under the `prediction`.

This field can also be set to "asis" if the scores are already scaled. In this case, no additional scaling will be performed by `rsmeval` but the report will refer to the scores as "scaled".

Defaults to "raw" which means that no-rescaling is performed and the report refers to the scores as "raw".

### **second\_human\_score\_column** (*Optional*)

The name for an optional column in the test data containing a second human score for each response. If specified, additional information about human-human agreement and degradation will be computed and included in the report. Note that this column must contain either numbers or be empty. Non-numeric values are *not* accepted. Note also that the *exclude\_zero\_scores* (*Optional*) option below will apply to this column too.

---

**Note:** You do not need to have second human scores for *all* responses to use this option. The human-human agreement statistics will be computed as long as there is at least one response with numeric value in this column. For responses that do not have a second human score, the value in this column should be blank.

---

### **section\_order** (*Optional*)

A list containing the order in which the sections in the report should be generated. Any specified order must explicitly list:

1. Either *all* pre-defined sections if a value for the *general\_sections* field is not specified OR the sections specified using *general\_sections*, and
2. *All* custom section names specified using *custom\_sections*, i.e., file prefixes only, without the path and without the *.ipynb* extension, and

**subgroups (Optional)**

A list of column names indicating grouping variables used for generating analyses specific to each of those defined subgroups. For example, ["prompt", "gender", "native\_language", "test\_country"]. These subgroup columns need to be present in the input predictions file. If subgroups are specified, rsmeval will generate:

- tables and barplots showing human-system agreement for each subgroup on the evaluation set.

**trim\_tolerance (Optional)**

The single numeric value that will be used to pad the trimming range specified in `trim_min` and `trim_max`. This value will be used to compute the ceiling and floor values for *trimmed (bound)* machine scores as `trim_max + trim_tolerance` for ceiling value and `trim_min - trim_tolerance` for floor value. Defaults to 0.4998.

---

**Note:** For more fine-grained control over the trimming range, you can set `trim_tolerance` to 0 and use `trim_min` and `trim_max` to specify the exact floor and ceiling values.

---

**use\_thumbnails (Optional)**

If set to `true`, the images in the HTML will be set to clickable thumbnails rather than full-sized images. Upon clicking the thumbnail, the full-sized images will be displayed in a separate tab in the browser. If set to `false`, full-sized images will be displayed as usual. Defaults to `false`.

**use\_wandb (Optional)**

If set to `true`, the generated reports and all intermediate tables will be logged to [Weights & Biases](#). The [Weights & Biases](#) *entity* and *project name* should be specified in the appropriate configuration fields. The tables and plots will be logged in a section named *rsmeval* in a new run under the given project, and the report will be added to a *reports* section in that run. In addition, some evaluation metrics will be logged to the run's history, see more details [here](#). Defaults to `false`.

**wandb\_project (Optional)**

The [Weights & Biases](#) project name if logging to [Weights & Biases](#) is enabled. If a project by this name does not already exist, it will be created.

---

**Important:**

1. Before using [Weights & Biases](#) for the first time, users should log in and provide their API key as described in [W&B Quickstart guidelines](#).
  2. Note that when using W&B logging, the rsmttool run may take significantly longer due to the network traffic being sent to W&B.
-

### wandb\_entity (Optional)

The Weights & Biases entity name if logging to Weights & Biases is enabled. Entity can be a user name or the name of a team or organization.

### Output

rsmeval produces a set of folders in the output directory. If logging to [Weights & Biases](#) is *enabled*, the reports and all intermediate files are also logged to the specified Weights & Biases project.

#### report

This folder contains the final RSMEval report in HTML format as well as in the form of a Jupyter notebook (a `.ipynb` file).

#### output

This folder contains all of the *intermediate files* produced as part of the various analyses performed, saved as `.csv` files. rsmeval will also save in this folder a copy of the *configuration file*. Fields not specified in the original configuration file will be pre-populated with default values.

#### figure

This folder contains all of the figures generated as part of the various analyses performed, saved as `.svg` files.

### Intermediate files

Although the primary output of rsmeval is an HTML report, we also want the user to be able to conduct additional analyses outside of rsmeval. To this end, all of the tables produced in the experiment report are saved as files in the format as specified by `file_format` parameter in the output directory. The following sections describe all of the intermediate files that are produced.

---

**Note:** The names of all files begin with the `experiment_id` provided by the user in the experiment configuration file. In addition, the names for certain columns are set to default values in these files irrespective of what they were named in the original data files. This is because RSMEval standardizes these column names internally for convenience. These values are:

- `spkitemid` for the column containing response IDs.
  - `sc1` for the column containing the human scores used as observed scores
  - `sc2` for the column containing the second human scores, if this column was specified in the configuration file.
  - `candidate` for the column containing candidate IDs, if this column was specified in the configuration file.
-

## Predictions

filename: `pred_processed`

This file contains the post-processed predicted scores: the predictions from the model are truncated, rounded, and re-scaled (if requested).

## Flagged responses

filename: `test_responses_with_excluded_flags`

This file contains all of the rows in input predictions file that were filtered out based on conditions specified in *flag\_column*.

---

**Note:** If the predictions file contained columns with internal names such as `sc1` that were not actually used by `rsmeval`, they will still be included in these files but their names will be changed to `##name##` (e.g. `##sc1##`).

---

## Excluded responses

filename: `test_excluded_responses`

This file contains all of the rows in the predictions file that were filtered out because of non-numeric or zero scores.

## Response metadata

filename: `test_metadata`

This file contains the metadata columns (`id_column`, subgroups if provided) for all rows in the predictions file that used in the evaluation.

## Unused columns

filename: `test_other_columns`

This file contains all of the the columns from the input predictions file that are not present in the `*_pred_processed` and `*_metadata` files. They only include the rows that were not filtered out.

---

**Note:** If the predictions file contained columns with internal names such as `sc1` but these columns were not actually used by `rsmeval`, these columns will also be included into these files but their names will be changed to `##name##` (e.g. `##sc1##`).

---

### Human scores

filename: `test_human_scores`

This file contains the human scores, if available in the input predictions file, under a column called `sc1` with the response IDs under the `spkitemid` column.

If `second_human_score_column` was specified, then it also contains the values in the predictions file from that column under a column called `sc2`. Only the rows that were not filtered out are included.

---

**Note:** If `exclude_zero_scores` was set to `true` (the default value), all zero scores in the `second_human_score_column` will be replaced by `nan`.

---

### Data composition

filename: `data_composition`

This file contains the total number of responses in the input predictions file. If applicable, the table will also include the number of different subgroups.

### Excluded data composition

filenames: `test_excluded_composition`

This file contains the composition of the set of excluded responses, e.g., why were they excluded and how many for each such exclusion.

### Subgroup composition

filename: `data_composition_by_<SUBGROUP>`

There will be one such file for each of the specified subgroups and it contains the total number of responses in that subgroup.

### Evaluation metrics

- **eval:** This file contains the descriptives for predicted and human scores (mean, std.dev etc.) as well as the association metrics (correlation, quadratic weighted kappa, SMD etc.) for the raw as well as the post-processed scores.
- **eval\_by\_<SUBGROUP>:** the same information as in `*_eval.csv` computed separately for each subgroup. However, rather than SMD, a difference of standardized means (DSM) will be calculated using z-scores.
- **eval\_short:** a shortened version of **eval** that contains specific descriptives for predicted and human scores (mean, std.dev etc.) and association metrics (correlation, quadratic weighted kappa, SMD etc.) for specific score types chosen based on recommendations by Williamson (2012). Specifically, the following columns are included (the raw or scale version is chosen depending on the value of the `use_scaled_predictions` in the configuration file).
  - `h_mean`
  - `h_sd`



- corr
- sys\_mean [raw/scale trim]
- sys\_sd [raw/scale trim]
- SMD [raw/scale trim]
- adj\_agr [raw/scale trim\_round]
- exact\_agr [raw/scale trim\_round]
- kappa [raw/scale trim\_round]
- wtkappa [raw/scale trim]
- sys\_mean [raw/scale trim\_round]
- sys\_sd [raw/scale trim\_round]
- SMD [raw/scale trim\_round]
- R2 [raw/scale trim]
- RMSE [raw/scale trim]
- **score\_dist**: the distributions of the human scores and the rounded raw/scaled predicted scores, depending on the value of **use\_scaled\_predictions**.
- **confMatrix**: the confusion matrix between the the human scores and the rounded raw/scaled predicted scores, depending on the value of **use\_scaled\_predictions**.

---

**Note:** Please note that for raw scores, SMD values are likely to be affected by possible differences in scale.

---

## Human-human Consistency

These files are created only if a second human score has been made available via the **second\_human\_score\_column** option in the configuration file.

- **consistency**: contains descriptives for both human raters as well as the agreement metrics between their ratings.
- **consistency\_by\_<SUBGROUP>**: contains the same metrics as in **consistency** file computed separately for each group. However, rather than SMD, a difference of standardized means (DSM) will be calculated using z-scores.
- **degradation**: shows the differences between human-human agreement and machine-human agreement for all association metrics and all forms of predicted scores.

## Evaluations based on test theory

- **disattenuated\_correlations**: shows the correlation between human-machine scores, human-human scores, and the disattenuated human-machine correlation computed as human-machine correlation divided by the square root of human-human correlation.
- **disattenuated\_correlations\_by\_<SUBGROUP>**: contains the same metrics as in **disattenuated\_correlations** file computed separately for each group.
- **true\_score\_eval**: evaluations of system scores against estimated true score. Contains total counts of single and double-scored response, variance of human rater error, estimated true score variance, and mean squared

error (MSE) and proportional reduction in mean squared error (PRMSE) when predicting true score using system score.

### Additional fairness analyses

These files contain the results of additional fairness analyses suggested in suggested in Loukina, Madnani, & Zechner, 2019.

- `<METRICS>_by_<SUBGROUP>.ols`: a serialized object of type `pandas.stats.ols.OLS` containing the fitted model for estimating the variance attributed to a given subgroup membership for a given metric. The subgroups are defined by the [configuration file](#). The metrics are `osa` (overall score accuracy), `osd` (overall score difference), and `csd` (conditional score difference).
- `<METRICS>_by_<SUBGROUP>.ols_summary.txt`: a text file containing a summary of the above model
- `estimates_<METRICS>_by_<SUBGROUP>``: coefficients, confidence intervals and  $p$ -values estimated by the model for each subgroup.
- `fairness_metrics_by_<SUBGROUP>`: the  $R^2$  (percentage of variance) and  $p$ -values for all models.

### 1.7.2 rsmpredict - Generate new predictions

RSMTool provides the `rsmpredict` command-line utility to generate predictions for new data using a model already trained using the `rsmtool` utility. This can be useful when processing a new set of responses to the same task without needing to retrain the model.

`rsmpredict` pre-processes the feature values according to user specifications before using them to generate the predicted scores. The generated scores are post-processed in the same manner as they are in `rsmtool` output.

---

**Note:** No score is generated for responses with non-numeric values for any of the features included into the model.

If the original model specified transformations for some of the features and these transformations led to NaN or Inf values when applied to the new data, `rsmpredict` will raise a warning. No score will be generated for such responses.

---

### Tutorial

For this tutorial, you first need to [install RSMTool](#) and make sure the conda environment in which you installed it is activated.

### Workflow

---

**Important:** Although this tutorial provides feature values for the purpose of illustration, `rsmpredict` does *not* include any functionality for feature extraction; the tool is [designed for researchers](#) who use their own NLP/Speech processing pipeline to extract features for their data.

---

`rsmpredict` allows you to generate the scores for new data using an existing model trained using RSMTool. Therefore, before starting this tutorial, you first need to complete [rsmtool tutorial](#) which will produce a train RSMTool model. You will also need to process the new data to extract the same features as the ones used in the model.

Once you have the features for the new data and the RSMTool model, using `rsmpredict` is fairly straightforward:

1. Create a file containing the features for the new data. The file should be in one of the *supported formats*.
2. Create an *experiment configuration file* describing the experiment you would like to run.
3. Run that configuration file with *rsmpredict* to generate the predicted scores.

---

**Note:** You do not *need* human scores to run *rsmpredict* since it does not produce any evaluation analyses. If you do have human scores for the new data and you would like to evaluate the system on this new data, you can first run *rsmpredict* to generate the predictions and then run *rsmeval* on the output of *rsmpredict* to generate an evaluation report.

---

## ASAP Example

We are going to use the same example from 2012 Kaggle competition on automated essay scoring that we used for the *rsmttool tutorial*. Specifically, We are going to use the linear regression model we trained in that tutorial to generate scores for new data.

---

**Note:** If you have not already completed that tutorial, please do so now. You may need to complete it again if you deleted the output files.

---

## Extract features

We will first need to generate features for the new set of responses for which we want to predict scores. For this experiment, we will simply re-use the test set from the *rsmttool* tutorial.

---

**Note:** The features used with *rsmpredict* should be generated using the *same* NLP/Speech processing pipeline that generated the features used in the *rsmttool* modeling experiment.

---

## Create a configuration file

The next step is to create an *rsmpredict experiment configuration file* in *.json* format.

```

1 {
2   "experiment_dir": "../rsmttool",
3   "experiment_id": "ASAP2",
4   "input_features_file": "../rsmttool/test.csv",
5   "id_column": "ID",
6   "human_score_column": "score",
7   "second_human_score_column": "score2"
8 }
```

Let's take a look at the options in our configuration file.

- **Line 2:** We give the path to the directory containing the output of the *rsmttool* experiment.
- **Line 3:** We provide the *experiment\_id* of the *rsmttool* experiment used to train the model. This can usually be read off the output/*<experiment\_id>.model* file in the *rsmttool* experiment output directory.

- **Line 4:** We list the path to the data file with the feature values for the new data. For this tutorial we used .csv format, but RSMTTool also supports several other *input file formats*.
- **Line 5:** This field indicates that the unique IDs for the responses in the .csv file are located in a column named ID.
- **Lines 6-7:** These fields indicate that there are two sets of human scores in our .csv file located in the columns named score and score2. The values from these columns will be added to the output file containing the predictions which can be useful if we want to evaluate the predictions using rsmeval.

Documentation for all of the available configuration options is available [here](#).

---

**Note:** You can also use our nifty capability to *automatically generate* rsmpredict configuration files rather than creating them manually.

---

### Run the experiment

Now that we have the model, the features in the right format, and our configuration file in .json format, we can use the *rsmpredict* command-line script to generate the predictions and to save them in predictions.csv.

```
$ cd examples/rsmpredict
$ rsmpredict config_rsmpredict.json predictions.csv
```

This should produce output like:

```
WARNING: The following extraneous features will be ignored: {'spkitemid', 'sc1', 'sc2',
↳ 'LENGTH'}
Pre-processing input features
Generating predictions
Rescaling predictions
Trimming and rounding predictions
Saving predictions to /Users/nmadnani/work/rsmttool/examples/rsmpredict/predictions.csv
```

You should now see a file named predictions.csv in the current directory which contains the predicted scores for the new data in the predictions column.

### Input

rsmpredict requires two arguments to generate predictions: the path to *a configuration file* and the path to the output file where the generated predictions are saved in .csv format.

If you also want to save the pre-processed feature values, ``rsmpredict`` can take a third optional argument `--features` to specify the path to a .csv file to save these values.

Here are all the arguments to the rsmpredict command-line script.

#### config\_file

The *JSON configuration file* for this experiment.

#### output\_file

The output .csv file where predictions will be saved.

`--features <preproc_feats_file>`

If specified, the pre-processed values for the input features will also be saved in this .csv file.

**-h, --help**

Show help message and exist.

**-V, --version**

Show version number and exit.

## Experiment configuration file

This is a file in `.json` format that provides overall configuration options for an `rsmpredict` experiment. Here's an [example configuration file](#) for `rsmpredict`.

---

**Note:** To make it easy to get started with `rsmpredict`, we provide a way to **automatically generate** configurations file both interactively as well as non-interactively. Novice users will find interactive generation more helpful while more advanced users will prefer non-interactive generation. See [this page](#) for more details.

---

Next, we describe all of the `rsmpredict` configuration fields in detail. There are three required fields and the rest are all optional. We first describe the required fields and then the optional ones (sorted alphabetically).

### `experiment_dir`

The path to the directory containing `rsmtool` model to use for generating predictions. This directory must contain a sub-directory called `output` with the model files, feature pre-processing parameters, and score post-processing parameters. The path can be absolute or relative to the location of configuration file.

### `experiment_id`

The `experiment_id` used to create the `rsmtool` model files being used for generating predictions. If you do not know the `experiment_id`, you can find it by looking at the prefix of the `.model` file under the `output` directory.

### `input_feature_file`

The path to the file with the raw feature values that will be used for generating predictions. The file should be in one of the [supported formats](#). Each row should correspond to a single response and contain feature values for this response. In addition, there should be a column with a unique identifier (ID) for each response. The path can be absolute or relative to the location of config file. Note that the feature names *must* be the same as used in the original `rsmtool` experiment.

---

**Note:** `rsmpredict` will only generate predictions for responses in this file that have numeric values for the features included in the `rsmtool` model.

---

### See also:

`rsmpredict` does not require human scores for the new data since it does not evaluate the generated predictions. If you do have the human scores and want to evaluate the new predictions, you can use the [rsmeval](#) command-line utility.

### **candidate\_column** (*Optional*)

The name for the column containing unique candidate IDs. This column will be named `candidate` in the output file with predictions.

### **file\_format** (*Optional*)

The format of the *intermediate files*. Options are `csv`, `tsv`, or `xlsx`. Defaults to `csv` if this is not specified.

### **flag\_column** (*Optional*)

See description in the *rsmtool configuration file* for further information. No filtering will be done by `rsmpredict`, but the contents of all specified columns will be added to the predictions file using the original column names.

### **human\_score\_column** (*Optional*)

The name for the column containing human scores. This column will be renamed to `sc1`.

### **id\_column** (*Optional*)

The name of the column containing the response IDs. Defaults to `spkitemid`, i.e., if this is not specified, `rsmpredict` will look for a column called `spkitemid` in the prediction file.

There are several other options in the configuration file that, while not directly used by `rsmpredict`, can simply be passed through from the input features file to the output predictions file. This can be particularly useful if you want to subsequently run *rsmeval* to evaluate the generated predictions.

### **predict\_expected\_scores** (*Optional*)

If the original model was a probabilistic SKLL classifier, then *expected scores* — probability-weighted averages over a contiguous, numeric score points — can be generated as the machine predictions instead of the most likely score point, which would be the default. Set this field to `true` to compute expected scores as predictions. Defaults to `false`.

---

#### **Note:**

1. If the model in the original `rsmtool` experiment is an SVC, that original experiment *must* have been run with `predict_expected_scores` set to `true`. This is because SVC classifiers are fit differently if probabilistic output is desired, in contrast to other probabilistic SKLL classifiers.
  2. You may see slight differences in expected score predictions if you run the experiment on different machines or on different operating systems most likely due to very small probability values for certain score points which can affect floating point computations.
-

**second\_human\_score\_column (Optional)**

The name for the column containing the second human score. This column will be renamed to `sc2`.

**standardize\_features (Optional)**

If this option is set to `false` features will not be standardized by dividing by the mean and multiplying by the standard deviation. Defaults to `true`.

**subgroups (Optional)**

A list of column names indicating grouping variables used for generating analyses specific to each of those defined subgroups. For example, `["prompt", "gender", "native_language", "test_country"]`. All these columns will be included into the predictions file with the original names.

**truncate\_outliers (Optional)**

If this option is set to `false`, outliers (values more than 4 standard deviations away from the mean) in feature columns will `_not_` be truncated. Defaults to `true`.

**use\_wandb (Optional)**

If set to `true`, the generated report and the predictions table will be logged to [Weights & Biases](#). The [Weights & Biases](#) *entity* and *project name* should be specified in the appropriate configuration fields. The predictions table will be logged in a section named *rsmpredict* in a new run under the given project, and the report will be added to a *reports* section in that run. Defaults to `false`.

**wandb\_project (Optional)**

The [Weights & Biases](#) project name if logging to [Weights & Biases](#) is enabled. If a project by this name does not already exist, it will be created.

---

**Important:**

1. Before using [Weights & Biases](#) for the first time, users should log in and provide their API key as described in [W&B Quickstart guidelines](#).
  2. Note that when using [W&B](#) logging, the `rsmttool` run may take significantly longer due to the network traffic being sent to [W&B](#).
-

### wandb\_entity (Optional)

The Weights & Biases entity name if logging to Weights & Biases is enabled. Entity can be a user name or the name of a team or organization.

### Output

rsmpredict produces a .csv file with predictions for all responses in new data set, and, optionally, a .csv file with pre-processed feature values. If any of the responses had non-numeric feature values in the original data or after applying transformations, these are saved in a file name PREDICTIONS\_NAME\_excluded\_responses.csv where PREDICTIONS\_NAME is the name of the predictions file supplied by the user without the extension.

The predictions .csv file contains the following columns:

- spkitemid : the unique response IDs from the original feature file.
- sc1 and sc2 : the human scores for each response from the original feature file (human\_score\_column and second\_human\_score\_column, respectively).
- raw : raw predictions generated by the model.
- raw\_trim, raw\_trim\_round, scale, scale\_trim, scale\_trim\_round : raw scores *post-processed* in different ways.

If logging to [Weights & Biases](#) is *enabled*, these csv files are also logged to the specified Weights & Biases project.

## 1.7.3 rsmcompare - Create a detailed comparison of two scoring models

RSMTool provides the rsmcompare command-line utility to compare two models and to generate a detailed comparison report including differences between the two models. This can be useful in many scenarios, e.g., say the user wants to compare the changes in model performance after adding a new feature into the model. To use rsmcompare, the user must first run two experiments using either *rsmtool* or *rsmeval*. rsmcompare can then be used to compare the outputs of these two experiments to each other.

---

**Note:** Currently rsmcompare takes the outputs of the analyses generated during the original experiments and creates comparison tables. These comparison tables were designed with a specific comparison scenario in mind: comparing a baseline model with a model which includes new feature(s). The tool can certainly be used for other comparison scenarios if the researcher feels that the generated comparison output is appropriate.

---

rsmcompare can be used to compare:

1. Two rsmtool experiments, or
2. Two rsmeval experiments, or
3. An rsmtool experiment with an rsmeval experiment (in this case, only the evaluation analyses will be compared).

---

**Note:** It is strongly recommended that the original experiments as well as the comparison experiment are all done using the same version of RSMTool.

---



## Tutorial

For this tutorial, you first need to *install RSMTool* and make sure the conda environment in which you installed it is activated.

## Workflow

rsmcompare is designed to compare two existing rsmtool or rsmeval experiments. To use rsmcompare you need:

1. Two experiments that were run using *rsmtool* or *rsmeval*.
2. Create an *experiment configuration file* describing the comparison experiment you would like to run.
3. Run that configuration file with *rsmcompare* and generate the comparison experiment HTML report.
4. Examine HTML report to compare the two models.

Note that the above workflow does not use the customization features of rsmcompare, e.g., *choosing which sections to include in the report* or *adding custom analyses sections* etc. However, we will stick with this workflow for our tutorial since it is likely to be the most common use case.

## ASAP Example

We are going to use the same example from 2012 Kaggle competition on automated essay scoring that we used for the *rsmtool tutorial*.

## Run rsmtool (or rsmeval) experiments

rsmcompare compares the results of the two existing rsmtool (or rsmeval) experiments. For this tutorial, we will compare model trained in the *rsmtool tutorial* to itself.

---

**Note:** If you have not already completed that tutorial, please do so now. You may need to complete it again if you deleted the output files.

---

## Create a configuration file

The next step is to create an *experiment configuration file* in .json format.

```

1 {
2   "comparison_id": "ASAP2_vs_ASAP2",
3   "experiment_id_old": "ASAP2",
4   "experiment_dir_old": "../rsmtool/",
5   "description_old": "RSMTool experiment.",
6   "use_scaled_predictions_old": true,
7   "experiment_id_new": "ASAP2",
8   "experiment_dir_new": "../rsmtool",
9   "description_new": "RSMTool experiment (copy).",
10  "use_scaled_predictions_new": true
11 }
```

Let's take a look at the options in our configuration file.

- **Line 2:** We provide an ID for the comparison experiment.
- **Line 3:** We provide the `experiment_id` for the experiment we want to use as a baseline.
- **Line 4:** We also give the path to the directory containing the output of the original baseline experiment.
- **Line 5:** We give a short description of this baseline experiment. This will be shown in the report.
- **Line 6:** This field indicates that the baseline experiment used scaled scores for some evaluation analyses.
- **Line 7:** We provide the `experiment_id` for the new experiment. We use the same experiment ID for both experiments since we are comparing the experiment to itself.
- **Line 8:** We also give the path to the directory containing the output of the new experiment. As above, we use the same path because we are comparing the experiment to itself.
- **Line 9:** We give a short description of the new experiment. This will also be shown in the report.
- **Line 10:** This field indicates that the new experiment also used scaled scores for some evaluation analyses.

Documentation for all of the available configuration options is available [here](#).

---

**Note:** You can also use our nifty capability to *automatically generate* `rsmcompare` configuration files rather than creating them manually.

---

### Run the experiment

Now that we have the two experiments we want to compare and our configuration file in `.json` format, we can use the `rsmcompare` command-line script to run our comparison experiment.

```
$ cd examples/rsmcompare
$ rsmcompare config_rsmcompare.json
```

This should produce output like:

```
Output directory: /Users/nmadnani/work/rsmtool/examples/rsmcompare
Starting report generation
Merging sections
Exporting HTML
Executing notebook with kernel: python3
```

Once the run finishes, you will see an HTML file named `ASAP2_vs_ASAP2_report.html`. This is the final `rsmcompare` comparison report.

### Examine the report

Our experiment report contains all the information we would need to compare the new model to the baseline model. It includes:

1. Comparison of feature distributions between the two experiments.
2. Comparison of model coefficients between the two experiments.
3. Comparison of model performance between the two experiments.

**Note:** Since we are comparing the experiment to itself, the comparison is not very interesting, e.g., the differences between various values will always be 0.

---

## Input

`rsmcompare` requires a single argument to run an experiment: the path to *a configuration file*. It can also take an output directory as an optional second argument. If the latter is not specified, `rsmcompare` will use the current directory as the output directory.

Here are all the arguments to the `rsmcompare` command-line script.

### `config_file`

The *JSON configuration file* for this experiment.

### `output_dir` (optional)

The output directory where the report files for this comparison will be stored.

### `-h, --help`

Show help message and exist.

### `-V, --version`

Show version number and exit.

## Experiment configuration file

This is a file in `.json` format that provides overall configuration options for an `rsmcompare` experiment. Here's an *example configuration file* for `rsmcompare`.

---

**Note:** To make it easy to get started with `rsmcompare`, we provide a way to **automatically generate** configurations file both interactively as well as non-interactively. Novice users will find interactive generation more helpful while more advanced users will prefer non-interactive generation. See *this page* for more details.

---

Next, we describe all of the `rsmcompare` configuration fields in detail. There are seven required fields and the rest are all optional. We first describe the required fields and then the optional ones (sorted alphabetically).

### `comparison_id`

An identifier for the comparison experiment that will be used to name the report. It can be any combination of alphanumeric values, must *not* contain spaces, and must *not* be any longer than 200 characters.

### **experiment\_id\_old**

An identifier for the “baseline” experiment. This ID should be identical to the `experiment_id` used when the baseline experiment was run, whether `rsmttool` or `rsmeval`. The results for this experiment will be listed first in the comparison report.

### **experiment\_id\_new**

An identifier for the experiment with the “new” model (e.g., the model with new feature(s)). This ID should be identical to the `experiment_id` used when the experiment was run, whether `rsmttool` or `rsmeval`. The results for this experiment will be listed first in the comparison report.

### **experiment\_dir\_old**

The directory with the results for the “baseline” experiment. This directory is the output directory that was used for the experiment and should contain subdirectories `output` and `figure` generated by `rsmttool` or `rsmeval`.

### **experiment\_dir\_new**

The directory with the results for the experiment with the new model. This directory is the output directory that was used for the experiment and should contain subdirectories `output` and `figure` generated by `rsmttool` or `rsmeval`.

### **description\_old**

A brief description of the “baseline” experiment. The description can contain spaces and punctuation.

### **description\_new**

A brief description of the experiment with the new model. The description can contain spaces and punctuation.

### **custom\_sections (Optional)**

A list of custom, user-defined sections to be included into the final report. These are IPython notebooks (`.ipynb` files) created by the user. The list must contain paths to the notebook files, either absolute or relative to the configuration file. All custom notebooks have access to some *pre-defined variables*.

### **general\_sections (Optional)**

RSMTool provides pre-defined sections for `rsmcompare` (listed below) and, by default, all of them are included in the report. However, you can choose a subset of these pre-defined sections by specifying a list as the value for this field.

- **feature\_descriptives**: Compares the descriptive statistics for all raw feature values included in the model:
  - a table showing mean, standard deviation, skewness and kurtosis;
  - a table showing the number of truncated outliers for each feature; and
  - a table with percentiles and outliers;

- a table with correlations between raw feature values and human score in each model and the correlation between the values of the same feature in these two models. Note that this table only includes features and responses which occur in both training sets.
- **features\_by\_group**: Shows boxplots for both experiments with distributions of raw feature values by each of the *subgroups* specified in the configuration file.
- **preprocessed\_features**: Compares analyses of preprocessed features:
  - histograms showing the distributions of preprocessed features values;
  - the correlation matrix between all features and the human score;
  - a table showing marginal correlations between all features and the human score; and
  - a table showing partial correlations between all features and the human score.
- **preprocessed\_features\_by\_group**: Compares analyses of preprocessed features by subgroups: marginal and partial correlations between each feature and human score for each subgroup.
- **consistency**: Compares metrics for human-human agreement, the difference ('degradation') between the human-human and human-system agreement, and the disattenuated correlations for the whole dataset and by each of the *subgroups* specified in the configuration file.
- **score\_distributions**:
  - tables showing the distributions for both human and machine scores; and
  - confusion matrices for human and machine scores.
- **model**: Compares the parameters of the two regression models. For linear models, it also includes the standardized and relative coefficients.
- **evaluation**: Compares the standard set of evaluations recommended for scoring models on the evaluation data.
- **true\_score\_evaluation**: compares the evaluation of system scores against the true scores estimated according to test theory. The notebook shows:
  - Number of single and double-scored responses.
  - Variance of human rater errors and estimated variance of true scores
  - Mean squared error (MSE) and proportional reduction in mean squared error (PRMSE) when predicting true score with system score.
- **pca**: Shows the results of principal components analysis on the processed feature values for the new model only:
  - the principal components themselves;
  - the variances; and
  - a Scree plot.
- **notes**: Notes explaining the terminology used in comparison reports.
- **sysinfo**: Shows all Python packages along with versions installed in the current environment while generating the report.

### **section\_order** (*Optional*)

A list containing the order in which the sections in the report should be generated. Any specified order must explicitly list:

1. Either *all* pre-defined sections if a value for the *general\_sections* field is not specified OR the sections specified using *general\_sections*, and
2. All custom section names specified using *custom\_sections*, i.e., file prefixes only, without the path and without the *.ipynb* extension, and

### **subgroups** (*Optional*)

A list of column names indicating grouping variables used for generating analyses specific to each of those defined subgroups. For example, ["prompt", "gender", "native\_language", "test\_country"].

---

**Note:** In order to include subgroups analyses in the comparison report, both experiments must have been run with the same set of subgroups.

---

### **use\_scaled\_predictions\_old** (*Optional*)

Set to `true` if the “baseline” experiment used scaled machine scores for confusion matrices, score distributions, subgroup analyses, etc. Defaults to `false`.

### **use\_scaled\_predictions\_new** (*Optional*)

Set to `true` if the experiment with the new model used scaled machine scores for confusion matrices, score distributions, subgroup analyses, etc. Defaults to `false`.

**Warning:** For `rsmtool` and `rsmeval`, primary evaluation analyses are computed on both raw and scaled scores, but some analyses (e.g., the confusion matrix) are only computed for either raw or re-scaled scores based on the value of `use_scaled_predictions`. `rsmcompare` uses the existing outputs and does not perform any additional evaluations. Therefore if this field was set to `true` in the original experiment but is set to `false` for `rsmcompare`, the report will be internally inconsistent: some evaluations use raw scores whereas others will use scaled scores.

### **use\_thumbnails** (*Optional*)

If set to `true`, the images in the HTML will be set to clickable thumbnails rather than full-sized images. Upon clicking the thumbnail, the full-sized images will be displayed in a separate tab in the browser. If set to `false`, full-sized images will be displayed as usual. Defaults to `false`.

### **use\_wandb (Optional)**

If set to `true`, the generated reports and all intermediate tables will be logged to [Weights & Biases](#). The Weights & Biases *entity* and *project name* should be specified in the appropriate configuration fields. The report will be added to a *reports* section in a new run under the given project. Defaults to `false`.

### **wandb\_project (Optional)**

The Weights & Biases project name if logging to Weights & Biases is enabled. If a project by this name does not already exist, it will be created.

---

#### **Important:**

1. Before using Weights & Biases for the first time, users should log in and provide their API key as described in [W&B Quickstart guidelines](#).
  2. Note that when using W&B logging, the `rsmtool` run may take significantly longer due to the network traffic being sent to W&B.
- 

### **wandb\_entity (Optional)**

The Weights & Biases entity name if logging to Weights & Biases is enabled. Entity can be a user name or the name of a team or organization.

## **Output**

`rsmtcompare` produces the comparison report in HTML format as well as in the form of a Jupyter notebook (a `.ipynb` file) in the output directory. If logging to [Weights & Biases](#) is *enabled*, the report will be also logged to the specified Weights & Biases project.

## **1.7.4 rsmsummarize - Compare multiple scoring models**

RSMTTool provides the `rsmsummarize` command-line utility to compare multiple models and to generate a comparison report. Unlike `rsmtcompare` which creates a detailed comparison report between the two models, `rsmsummarize` can be used to create a more general overview of multiple models.

`rsmsummarize` can be used to compare:

1. Multiple `rsmttool` experiments, or
2. Multiple `rsmeval` experiments, or
3. A mix of `rsmttool` and `rsmeval` experiments (in this case, only the evaluation analyses will be compared).

---

**Note:** It is strongly recommend that the original experiments as well as the summary experiment are all done using the same version of RSMTTool.

---

### Tutorial

For this tutorial, you first need to *install RSMTool* and make sure the conda environment in which you installed it is activated.

### Workflow

`rsmsummarize` is designed to compare several existing `rsmttool` or `rsmeval` experiments. To use `rsmsummarize` you need:

1. Two or more experiments that were run using *rsmttool* or *rsmeval*.
2. Create an *experiment configuration file* describing the comparison experiment you would like to run.
3. Run that configuration file with *rsmsummarize* and generate the comparison experiment HTML report.
4. Examine HTML report to compare the models.

Note that the above workflow does not use the customization features of `rsmsummarize`, e.g., *choosing which sections to include in the report* or *adding custom analyses sections* etc. However, we will stick with this workflow for our tutorial since it is likely to be the most common use case.

### ASAP Example

We are going to use the same example from 2012 Kaggle competition on automated essay scoring that we used for the *rsmttool tutorial*.

### Run `rsmttool` and `rsmeval` experiments

`rsmsummarize` compares the results of the two or more existing `rsmttool` (or `rsmeval`) experiments. For this tutorial, we will compare model trained in the *rsmttool tutorial* to the evaluations we obtained in the *rsmeval tutorial*.

---

**Note:** If you have not already completed these tutorials, please do so now. You may need to complete them again if you deleted the output files.

---

### Create a configuration file

The next step is to create an *experiment configuration file* in `.json` format.

```
1 {
2   "summary_id": "model_comparison",
3   "description": "a comparison of the results of the rsmttool sample experiment, rsmeval_
  ↳ sample experiment and once again the rsmttool sample experiment",
4   "experiment_dirs": ["../rsmttool", "../rsmeval", "../rsmttool"],
5   "experiment_names": ["RSMTool experiment 1", "RSMEval experiment", "RSMTool experiment 2
  ↳ "]
6 }
```

Let's take a look at the options in our configuration file.

- **Line 2:** We provide the `summary_id` for the comparison. This will be used to generate the name of the final report.



- **Line 3:** We give a short description of this comparison experiment. This will be shown in the report.
- **Line 4:** We also give the list of paths to the directories containing the outputs of the experiments we want to compare.
- **Line 5:** Since we want to compare experiments that all used the same experiment id (ASAP2), we instead list the names that we want to use for each experiment in the summary report.

Documentation for all of the available configuration options is available [here](#).

---

**Note:** You can also use our nifty capability to *automatically generate* `rsmsummarize` configuration files rather than creating them manually.

---

## Run the experiment

Now that we have the list of the experiments we want to compare and our configuration file in `.json` format, we can use the `rsmsummarize` command-line script to run our comparison experiment.

```
$ cd examples/rsmsummarize
$ rsmsummarize config_rsmsummarize.json
```

This should produce output like:

```
Output directory: /Users/nmadnani/work/rsmtool/examples/rsmsummarize
Starting report generation
Merging sections
Exporting HTML
Executing notebook with kernel: python3
```

Once the run finishes, you will see a new folder `report` containing an HTML file named `model_comparison_report.html`. This is the final `rsmsummarize` summary report.

## Examine the report

Our experiment report contains the overview of main aspects of model performance. It includes:

1. Brief description of all experiments.
2. Information about model parameters and model fit for all `rsmtool` experiments.
3. Model performance for all experiments.

---

**Note:** Some of the information such as model fit and model parameters are only available for `rsmtool` experiments.

---

### Input

`rsmsummarize` requires a single argument to run an experiment: the path to *a configuration file*. You can specify which models you want to compare and the name of the report by supplying the path to *a configuration file*. It can also take an output directory as an optional second argument. If the latter is not specified, `rsmsummarize` will use the current directory as the output directory.

Here are all the arguments to the `rsmsummarize` command-line script.

#### `config_file`

The *JSON configuration file* for this experiment.

#### `output_dir` (optional)

The output directory where the report and intermediate `.csv` files for this comparison will be stored.

#### `-f, --force`

If specified, the contents of the output directory will be overwritten even if it already contains the output of another `rsmsummarize` experiment.

#### `-h, --help`

Show help message and exist.

#### `-V, --version`

Show version number and exit.

### Experiment configuration file

This is a file in `.json` format that provides overall configuration options for an `rsmsummarize` experiment. Here's an [example configuration file](#) for `rsmsummarize`.

---

**Note:** To make it easy to get started with `rsmsummarize`, we provide a way to **automatically generate** configurations file both interactively as well as non-interactively. Novice users will find interactive generation more helpful while more advanced users will prefer non-interactive generation. See [this page](#) for more details.

---

Next, we describe all of the `rsmsummarize` configuration fields in detail. There are two required fields and the rest are all optional. We first describe the required fields and then the optional ones (sorted alphabetically).

#### `summary_id`

An identifier for the `rsmsummarize` experiment. This will be used to name the report. It can be any combination of alphanumeric values, must *not* contain spaces, and must *not* be any longer than 200 characters.

#### `experiment_dirs`

The list of the directories with the results of the experiment. These directories should be the output directories used for each experiment and should contain subdirectories `output` and `figure` generated by `rsmtool` or `rsmeval`.

### custom\_sections (Optional)

A list of custom, user-defined sections to be included into the final report. These are IPython notebooks (.ipynb files) created by the user. The list must contain paths to the notebook files, either absolute or relative to the configuration file. All custom notebooks have access to some *pre-defined variables*.

### description (Optional)

A brief description of the summary. The description can contain spaces and punctuation.

### experiment\_names (Optional)

The list of experiment names to use in the summary report and *intermediate files*. The names should be listed in the same order as the experiments in *experiment\_dirs*. When this field is not specified, the report will show the original experiment\_id for each experiment.

### file\_format (Optional)

The format of the *intermediate files* generated by `rsmsummarize`. Options are `csv`, `tsv`, or `xlsx`. Defaults to `csv` if this is not specified.

---

**Note:** In the `rsmsummarize` context, the `file_format` parameter refers to the format of the intermediate files generated by `rsmsummarize`, not the intermediate files generated by the original experiment(s) being summarized. The format of these files does not have to match the format of the files generated by the original experiment(s).

---

### general\_sections (Optional)

RSMTTool provides pre-defined sections for `rsmsummarize` (listed below) and, by default, all of them are included in the report. However, you can choose a subset of these pre-defined sections by specifying a list as the value for this field.

- **preprocessed\_features** : compares marginal and partial correlations between all features and the human score, and optionally response length if this was computed for any of the models.
- **model**: Compares the parameters of the two regression models. For linear models, it also includes the standardized and relative coefficients.
- **evaluation**: Compares the standard set of evaluations recommended for scoring models on the evaluation data.
- **true\_score\_evaluation**: compares the evaluation of system scores against the true scores estimated according to test theory. The notebook shows:
  - Number of single and double-scored responses.
  - Variance of human rater errors and estimated variance of true scores
  - Mean squared error (MSE) and proportional reduction in mean squared error (PRMSE) when predicting true score with system score.
- **intermediate\_file\_paths**: Shows links to all of the intermediate files that were generated while running the summary.
- **sysinfo**: Shows all Python packages along with versions installed in the current environment while generating the report.

### **section\_order** (*Optional*)

A list containing the order in which the sections in the report should be generated. Any specified order must explicitly list:

1. Either *all* pre-defined sections if a value for the *general\_sections* field is not specified OR the sections specified using *general\_sections*, and
2. *All* custom section names specified using *custom\_sections*, i.e., file prefixes only, without the path and without the *.ipynb* extension, and

### **use\_thumbnails** (*Optional*)

If set to `true`, the images in the HTML will be set to clickable thumbnails rather than full-sized images. Upon clicking the thumbnail, the full-sized images will be displayed in a separate tab in the browser. If set to `false`, full-sized images will be displayed as usual. Defaults to `false`.

### **use\_wandb** (*Optional*)

If set to `true`, the generated report will be logged to [Weights & Biases](#). The Weights & Biases *entity* and *project name* should be specified in the appropriate configuration fields. The report will be added to a *reports* section in a new run under the given project. Defaults to `false`.

### **wandb\_project** (*Optional*)

The Weights & Biases project name if logging to Weights & Biases is enabled. If a project by this name does not already exist, it will be created.

---

#### **Important:**

1. Before using Weights & Biases for the first time, users should log in and provide their API key as described in [W&B Quickstart guidelines](#).
  2. Note that when using W&B logging, the `rsmtool` run may take significantly longer due to the network traffic being sent to W&B.
- 

### **wandb\_entity** (*Optional*)

The Weights & Biases entity name if logging to Weights & Biases is enabled. Entity can be a user name or the name of a team or organization.

## Output

`rsmsummarize` produces a set of folders in the output directory. If logging to [Weights & Biases](#) is *enabled*, the reports and all intermediate files are also logged to the specified Weights & Biases project.

### report

This folder contains the final `rsmsummarize` report in HTML format as well as in the form of a Jupyter notebook (a `.ipynb` file).

### output

This folder contains all of the *intermediate files* produced as part of the various analyses performed, saved as `.csv` files. `rsmsummarize` will also save in this folder a copy of the *configuration file*. Fields not specified in the original configuration file will be pre-populated with default values.

### figure

This folder contains all of the figures that may be generated as part of the various analyses performed, saved as `.svg` files. Note that no figures are generated by the existing `rsmsummarize` notebooks.

## Intermediate files

Although the primary output of RSMSummarize is an HTML report, we also want the user to be able to conduct additional analyses outside of RSMTTool. To this end, all of the tables produced in the experiment report are saved as files in the format as specified by `file_format` parameter in the output directory. The following sections describe all of the intermediate files that are produced.

---

**Note:** The names of all files begin with the `summary_id` provided by the user in the experiment configuration file.

---

## Marginal and partial correlations with score

filenames: `margcor_score_all_data`, `pcor_score_all_data`, ``pcor_score_no_length_all_data`

The first file contains the marginal correlations between each pre-processed feature and human score. The second file contains the partial correlation between each pre-processed feature and human score after controlling for all other features. The third file contains the partial correlations between each pre-processed feature and human score after controlling for response length, if `length_column` was specified in the configuration file.

### Model information

- `model_summary`

This file contains the main information about the models included into the report including:

- Total number of features
- Total number of features with non-negative coefficients
- The learner
- The label used to train the model
- `betas`: standardized coefficients (for built-in models only).
- `model_fit`: R squared and adjusted R squared computed on the training set. Note that these values are always computed on raw predictions without any trimming or rounding.

---

**Note:** If the report includes a combination of `rsmttool` and `rsmeval` experiments, the summary tables with model information will only include `rsmttool` experiments since no model information is available for `rsmeval` experiments.

---

### Evaluation metrics

- `eval_short` - descriptives for predicted and human scores (mean, std.dev etc.) and association metrics (correlation, quadratic weighted kappa, SMD etc.) for specific score types chosen based on recommendations by Williamson (2012). Specifically, the following columns are included (the raw or scale version is chosen depending on the value of the `use_scaled_predictions` in the configuration file).
  - `h_mean`
  - `h_sd`
  - `corr`
  - `sys_mean` [raw/scale trim]
  - `sys_sd` [raw/scale trim]
  - `SMD` [raw/scale trim]
  - `adj_agr` [raw/scale trim\_round]
  - `exact_agr` [raw/scale trim\_round]
  - `kappa` [raw/scale trim\_round]
  - `wtkappa` [raw/scale trim\_round]
  - `sys_mean` [raw/scale trim\_round]
  - `sys_sd` [raw/scale trim\_round]
  - `SMD` [raw/scale trim\_round]
  - `R2` [raw/scale trim]
  - `RMSE` [raw/scale trim]

---

**Note:** Please note that for raw scores, SMD values are likely to be affected by possible differences in scale.

---

## Evaluations based on test theory

- `true_score_eval`: evaluations of system scores against estimated true score. Contains total counts of single and double-scored response, variance of human rater error, estimated true score variance, and mean squared error (MSE) and proportional reduction in mean squared error (PRMSE) when predicting true score using system score.

### 1.7.5 `rsmxval` - Run cross-validation experiments

RSMTTool provides the `rsmxval` command-line utility to run cross-validation experiments with scoring models. Why would one want to use cross-validation rather than just using the simple train-and-evaluate loop provided by the `rsmtool` utility? Using cross-validation can provide more accurate estimates of scoring model performance since those estimates are averaged over *multiple* train-test splits that are randomly selected based on the data. Using a single train-test split may lead to biased estimates of performance since those estimates will depend on the specific characteristics of that split. Using cross-validation is more likely to provide estimates of how well the scoring model will generalize to unseen test data, and more easily flag problems with overfitting and selection bias, if any.

Cross-validation experiments in RSMTTool consist of the following steps:

1. The given training data file is first shuffled randomly (with a fixed seed for reproducibility) and then split into the requested number of folds. It is also possible for the user to provide a CSV file containing a pre-determined set of folds, e.g., from another part of the data pipeline.
2. For each fold (or train-test split), `rsmtool` is run to train a scoring model on the training split and evaluate on the test split. All of the outputs for each of the `rsmtool` runs are saved on disk and represent the per-fold performance.
3. The predictions generated by `rsmtool` for each of the folds are all combined into a single file, which is then used as input for `rsmeval`. The output of this evaluation run is saved to disk and provides a more accurate estimate of the predictive performance of a scoring model trained on the given data.
4. A summary report comparing all of the folds is generated by running `rsmsummarize` on all of the fold directories created in the Step 1 and its output is also saved to disk. This summary output can be useful to see if the performance for any of the folds stands out for any reason, which could point to a potential problem.
5. Finally, a scoring model is trained on the complete training data file using `rsmtool`, which also generates a report that contains only the feature and model descriptives. The model is what will most likely be deployed for inference assuming the analyses produced in this step and Steps 1–4 meet the stakeholders' requirements.

## Tutorial

For this tutorial, you first need to *install RSMTTool* and make sure the conda environment in which you installed it is activated.

## Workflow

`rsmxval` is designed to run cross-validation experiments using a single file containing human scores and features. Just like `rsmtool`, `rsmxval` does not provide any functionality for feature extraction and assumes that users will extract features on their own. The workflow steps are as follows:

1. Create a data file in one of the *supported formats* containing the extracted features for each response in the data along with human score(s) assigned to it.
2. Create an *experiment configuration file* describing the cross-validation experiment you would like to run.
3. Run that configuration file with `rsmxval` and generate its *outputs*.

4. Examine the various HTML reports to check various aspects of model performance.

Note that unlike `rsmttool` and `rsmeval`, `rsmxval` currently does not support customization of the HTML reports generated in each step. This functionality may be added in future versions.

### ASAP Example

We are going to use the same example from the 2012 Kaggle competition on automated essay scoring that we used for the [rsmttool tutorial](#).

### Extract features

We are using the same features for this data as described in the [rsmttool tutorial](#).

### Create a configuration file

The next step is to create an *experiment configuration file* in `.json` format.

```
1 {  
2   "experiment_id": "ASAP2_xval",  
3   "description": "Cross-validation with two human scores using a LinearRegression  
↪model.",  
4   "train_file": "train.csv",  
5   "folds": 3,  
6   "train_label_column": "score",  
7   "id_column": "ID",  
8   "model": "LinearRegression",  
9   "trim_min": 1,  
10  "trim_max": 6,  
11  "second_human_score_column": "score2",  
12  "use_scaled_predictions": true  
13 }
```

Let's take a look at the options in our configuration file.

- **Line 2:** We define an experiment ID used to identify the files produced as part of this experiment.
- **Line 3:** We provide a description which will be included in the various reports.
- **Line 4:** We list the path to our training file with the feature values and human scores. For this tutorial, we used `.csv` format, but several other *input file formats* are also supported.
- **Line 5:** This field indicates the number of cross-validation folds we want to use. If this field is not specified, `rsmxval` uses 5-fold cross-validation by default.
- **Line 6:** This field indicates that the human (reference) scores in our `.csv` file are located in a column named `score`.
- **Line 7:** This field indicates that the unique IDs for the responses in the `.csv` file are located in a column named `ID`.
- **Line 8:** We choose to use a linear regression model to combine the feature values into a score.
- **Lines 9-10:** These fields indicate that the lowest score on the scoring scale is a 1 and the highest score is a 6. This information is usually part of the rubric used by human graders.



- **Line 11:** This field indicates that scores from a second set of human graders are also available (useful for comparing the agreement between human-machine scores to the agreement between two sets of humans) and are located in the `score2` column in the training `.csv` file.
- **Line 12:** Next, we indicate that we would like to use the *scaled scores* for all our evaluation analyses at each step.

Documentation for all of the available configuration options is available [here](#).

---

**Note:** You can also use our nifty capability to *automatically generate* `rsmxval` configuration files rather than creating them manually.

---

## Run the experiment

Now that we have our input file and our configuration file, we can use the *rsmxval* command-line script to run our evaluation experiment.

```
$ cd examples/rsmxval
$ rsmxval config_rsmxval.json output
```

This should produce output like:

```
Output directory: output
Saving configuration file.
Generating 3 folds after shuffling
Running RSMTTool on each fold in parallel
Progress: 100%|| 3/3 [00:08<00:00, 2.76s/it]
Creating fold summary
Evaluating combined fold predictions
Training model on full data
```

Once the run finishes, you will see an output sub-directory in the current directory. Under this directory you will see multiple sub-directories, each corresponding to a different cross-validation step, as described [here](#).

## Examine the reports

The cross-validation experiment produces multiple HTML reports – an `rsmtool` report for each of the 3 folds (`output/folds/{01,02,03}/report/ASAP2_xval_fold{01,02,03}.html`), the evaluation report for the cross-validated predictions (`output/evaluation/report/ASAP2_xval_evaluation_report.html`), a report summarizing the salient characteristics of the 3 folds (`output/fold-summary/report/ASAP2_xval_fold_summary_report.html`), and a report showing the feature and model descriptives (`output/final-model/report/ASAP2_xval_model_report.html`). Examining these reports will provide a relatively complete picture of how well the predictive performance of the scoring model will generalize to unseen data.

### Input

`rsmxval` requires a single argument to run an experiment: the path to *a configuration file*. It can also take an output directory as an optional second argument. If the latter is not specified, `rsmxval` will use the current directory as the output directory.

Here are all the arguments to the `rsmxval` command-line script.

#### **config\_file**

The *JSON configuration file* for this cross-validation experiment.

#### **output\_dir** (optional)

The output directory where all the sub-directories and files for this cross-validation experiment will be stored. If a non-empty directory with the same name already exists, an error will be raised.

#### **-h, --help**

Show help message and exit.

#### **-V, --version**

Show version number and exit.

### Experiment configuration file

This is a file in `.json` format that provides overall configuration options for an `rsmxval` experiment. Here's an *example configuration file* for `rsmxval`.

---

**Note:** To make it easy to get started with `rsmxval`, we provide a way to **automatically generate** configuration files both interactively as well as non-interactively. Novice users will find interactive generation more helpful while more advanced users will prefer non-interactive generation. See *this page* for more details.

---

Configuration files for `rsmxval` are almost identical to `rsmttool` configuration files with only a few differences. Next, we describe the three required `rsmxval` configuration fields in detail.

#### **experiment\_id**

An identifier for the experiment that will be used as part of the names of the reports and intermediate files produced in each of the steps. It can be any combination of alphanumeric values, must *not* contain spaces, and must *not* be any longer than 200 characters. Suffixes are added to this experiment ID by each of the steps for the reports and files they produce, i.e., `_fold<N>` in the per-fold `rsmttool` step where `<N>` is a two digit number, `_evaluation` by the `rsmeval` evaluation step, `_fold_summary` by the `rsmsummarize` step, and `_model` by the final full-data `rsmttool` step.

#### **model**

The machine learner you want to use to build the scoring model. Possible values include *built-in linear regression models* as well as all of the learners available via *SKLL*. With SKLL learners, you can customize the *tuning objective* and also *compute expected scores as predictions*.

## **train\_file**

The path to the training data feature file in one of the *supported formats*. Each row should correspond to a single response and contain numeric feature values extracted for this response. In addition, there should be a column with a unique identifier (ID) for each response and a column with the human score for each response. The path can be absolute or relative to config file's location.

---

**Important:** Unlike `rsmtool`, `rsmxval` does not accept an evaluation set and will raise an error if the `test_file` field is specified.

---

Next, we will describe the optional fields that are unique to `rsmxval`.

## **folds (Optional)**

The number of folds to use for cross-validation. This should be an integer and defaults to 5.

## **folds\_file (Optional)**

The path to a file containing custom, pre-specified folds to be used for cross-validation. This should be a `.csv` file (no other formats are accepted) and should contain only two columns: `id` and `fold`. The `id` column should contain the same IDs of the responses that are contained in `train_file` above. The `fold` column should contain an integer representing which fold the response with the `id` belongs to. IDs not specified in this file will be skipped and not included in the cross-validation at all. Just like `train_file`, this path can be absolute or relative to the config file's location. Here's an [example of a folds file containing 2 folds](#).

---

**Note:** If *both* `folds_file` and `folds` are specified, then the former will take precedence unless it contains a non-existent path.

---

## **use\_wandb (Optional)**

If set to `true`, the generated reports and all intermediate tables will be logged to [Weights & Biases](#). The [Weights & Biases](#) *entity* and *project name* should be specified in the appropriate configuration fields. A new run will be created in the specified project and it will include different sections for the different steps of the experiment:

- `rsmxval` section will include the combined predictions file generated by the folds.
- `rsmeval` section will include the tables and plots from the evaluation run on the predictions file.
- `rsmtool` section will include the output of the final `rsmtool` run that creates the final model.
- All the reports generated in the process will be logged to the *reports* section.

In addition, some evaluation metrics will be logged to the run's history with a name representing the context in which each score is calculated (see more details [here](#)).

Note that the output of the individual `rsmtool` runs for each fold are not logged to W&B.

Defaults to `false`.

In addition to the fields described so far, an `rsmxval` configuration file also accepts the following optional fields used by `rsmtool`:

- `candidate_column`

- description
- exclude\_zero\_scores
- feature\_subset
- feature\_subset\_file
- features
- file\_format
- flag\_column
- flag\_column\_test
- id\_column
- length\_column
- min\_items\_per\_candidate
- min\_n\_per\_group
- predict\_expected\_scores
- rater\_error\_variance
- second\_human\_score\_column
- select\_transformations
- sign
- skll\_fixed\_parameters
- skll\_objective
- standardize\_features
- subgroups
- train\_label\_column
- trim\_max
- trim\_min
- trim\_tolerance
- truncate\_outliers
- use\_scaled\_predictions
- use\_thumbnails
- use\_truncation\_thresholds
- skll\_grid\_search\_jobs
- use\_wandb
- wandb\_entity
- wandb\_project

Please refer to these fields' descriptions on the page describing the *rsmttool configuration file*.

## Output

rsmxval produces a set of folders in the output directory. If logging to [Weights & Biases](#) is *enabled*, the reports generated in this run are also logged to the specified Weights & Biases project.

### folders

This folder contains the output of each of the per-fold `rsmttool` experiments. It contains as many sub-folders as the number of specified folds, named 01, 02, 03, etc. Each of these numbered sub-folders contains the output of one `rsmttool` experiment conducted using the training split of that fold as the training data and the test split as the evaluation data. Each of the sub-folders contains the *output directories produced by rsmttool*. The report for each fold lives in the `report` sub-directory, e.g., the report for the first fold is found at `folders/01/report/<experiment_id>_fold01_report.html`, and so on. The messages that are usually printed out by `rsmttool` to the screen are instead logged to a file and saved to disk as, e.g., `folders/01/rsmttool.log`.

### evaluation

This folder contains the output of the `rsmeval` evaluation experiment that uses the cross-validated predictions from each fold. This folder contains the *output directories produced by rsmeval*. The evaluation report can be found at `evaluation/report/<experiment_id>_evaluation_report.html`. The messages that are usually printed out by `rsmeval` to the screen are instead logged to a file and saved to disk as `evaluation/rsmeval.log`.

### fold-summary

This folder contains the output of the `rmsummarize` experiment that provides a quick summary of all of the folds in a single, easily-scanned report. The folder contains the *output directories produced by rmsummarize*. The summary report can be found at `fold-summary/report/<experiment_id>_fold_summary_report.html`. The messages that are usually printed out by `rmsummarize` to the screen are instead logged to a file and saved to disk as `fold-summary/rmsummarize.log`.

### final-model

This folder contains the output of the `rsmttool` experiment that trains a model on the full training data and provides a report showing the feature and model descriptives. It contains the *output directories produced by rsmttool*. The primary artifacts of this experiment are the report (`final-model/report/<experiment_id>_model_report.html`) and the final trained model (`final-model/output/<experiment_id>_model.model`). The messages that are usually printed out by `rsmttool` to the screen are instead logged to a file and saved to disk as `final-model/rsmttool.log`.

---

**Note:** Every `rsmttool` experiment requires both a training and an evaluation set. However, in this step, we are using the full training data to train the model and `rsmxval` does not use a separate test set. Therefore, we simply randomly sample 10% of the full training data as a dummy test set to make sure that `rsmttool` runs successfully. The report in this step *only* contains the model and feature descriptives and, therefore, does not use this dummy test set at all. Users should ignore any intermediate files under the `final-model/output` and `final-model/figure` sub-directories that are derived from this dummy test set. If needed, the data used as the dummy test set can be found at `final-model/dummy_test.csv` (or in the *chosen format*).

---

In addition to these folders, `rsmxval` will also save a copy of the *configuration file* in the output directory at the same-level as the above folders. Fields not specified in the original configuration file will be pre-populated with default values.

## 1.7.6 `rsmexplain` - Explain non-linear models

RSMTTool provides the `rsmexplain` command-line utility to generate a report explaining the predictions made by a model trained using the `rsmttool` utility. These explanations contain useful information about the contribution of each feature to the final score, even if the model is non-linear or black-box in nature. The `rsmexplain` command-line utility uses the [SHAP library](#) to compute the explanations.

---

**Note:** `rsmexplain` uses the sampling explainer which is model agnostic and should, in principle, work for any type of model. However, `rsmexplain` currently only supports regressors since they are the most popular model type used for automated scoring.

---

### Tutorial

For this tutorial, you first need to *install RSMTTool* and make sure the conda environment in which you installed it is activated.

### Workflow

`rsmexplain` is designed to explain the predictions from a model trained as part of an existing `rsmttool` experiment. The steps to do this are as follows:

1. Successfully run an *`rsmttool`* experiment so that the model we would like to explain is trained and available.
2. Create an *experiment configuration file* describing the explanation experiment you would like to run.
3. Run the created configuration file with *`rsmexplain`* to generate the explanation HTML report.
4. Examine the HTML report to see the explanations for the `rsmttool` model on the selected responses.

### ASAP Example

We are going to use the same example from 2012 Kaggle competition on automated essay scoring that we used for the *`rsmttool` tutorial*.

### Run `rsmttool` experiments with chosen model

`rsmexplain` requires an existing `rsmttool` experiment with a trained model. For this tutorial, we will explain the model trained as part of the *`rsmttool` tutorial*.

---

**Note:** If you have not already completed the `rsmttool` tutorial, please do so now. You may need to complete it again if you deleted the output files.

---

## Create a configuration file

The next step is to create an *experiment configuration file* for `rsmexplain` in `.json` format.

```

1 {
2   "description": "Explaining a linear regression model trained on all features.",
3   "experiment_dir": "../rsmttool",
4   "experiment_id": "ASAP2",
5   "background_data": "../rsmttool/train.csv",
6   "explain_data": "../rsmttool/test.csv",
7   "id_column": "ID",
8   "sample_size": 1,
9   "num_features_to_display": 15
10 }
```

Let's take a look at the options in our configuration file.

- **Line 2:** We give a short description of this experiment. This will be shown in the report.
- **Line 3:** We give the path to the directory containing the output of the original `rsmttool` experiment. Note that this is the top-level directory that contains the output directory produced by `rsmttool`.
- **Line 4:** We provide the `experiment_id` of the `rsmttool` experiment used to train the model. This can usually be read off the output/`<experiment_id>.model` file in the `rsmttool` experiment output directory.
- **Line 5:** We provide the path to the data file that will be used as the background distribution.
- **Line 6:** We provide the path to the data file that will be used to generate the explanations.
- **Line 7:** *This field* indicates that the unique IDs for the responses in both data files are located in a column named `ID`.
- **Line 8:** *This field* indicates that we wish to explain one *randomly chosen* example from the second data file. If we wish to explain a *specific* example from that file, we would use the *sample\_ids* option instead.
- **Line 9:** *This field* indicates the number of top features that should be displayed in the plots in the `rsmexplain` report.

Documentation for all of the available configuration options is available [here](#).

---

**Note:** You can also use our nifty capability to *automatically generate* `rsmexplain` configuration files rather than creating them manually.

---

## Run explanation experiment

Now that we have the `rsmttool` experiment, the data files, and our configuration file, we can use the *rsmexplain* command-line script to run our explanation experiment.

```

$ cd examples/rsmexplain
$ rsmexplain config_rsmexplain.json
```

This should produce output like:

```

Output directory: /Users/nmadnani/work/rsmttool/examples/rsmexplain
Saving configuration file.
```

(continues on next page)

(continued from previous page)

```

WARNING: The following extraneous features will be ignored: {'LENGTH', 'score2', 'score'}
Pre-processing input features
WARNING: The following extraneous features will be ignored: {'LENGTH', 'score2', 'score'}
Pre-processing input features
Generating SHAP explanations for 1 examples from ../rsmtool/test.csv
100%| 1/1 [00:00<00:00, 21.99it/s]
Merging sections
Exporting HTML
Success

```

Once the run finishes, you will see the `output`, `figure`, and `report` sub-directories in the current directory. Each of these directories contain *useful information* but we are specifically interested in the `report/ASAP2_explain_report.html` file, which is the final evaluation report.

## Examine the report

Our experiment report contains all the information we would need to explain the trained model. It includes:

1. The various absolute value variants of the SHAP values.
2. Several SHAP plots indicating how different features contribute to the predicted score. Since we chose to explain a single example in this tutorial, the following plots will be displayed in the report: [global bar plot](#), [beeswarm plot](#), [decision plot](#), and [waterfall plot](#).

**Note:** We encourage you to re-run the tutorial by modifying the configuration file to explain multiple examples instead of a single one. You can do so either by setting *sample\_size* to a value larger than 1, by explicitly specifying multiple example indices via *sample\_ids*, or by setting *sample\_range* to an appropriate range of example indices. For a multiple-example explanation run, the following plots will be displayed in the report: [global bar plot](#), [beeswarm plot](#), and [heatmap plots](#).

## Input

`rsmexplain` requires only one argument to generate the explanation report: the path to *a configuration file*.

Here are all the arguments to the `rsmexplain` command-line script.

### **config\_file**

The *JSON configuration file* for this experiment.

### **output\_dir**

The output directory where all the files for this experiment will be stored.

### **-f, --force**

If specified, the contents of the output directory will be overwritten even if it already contains the output of another `rsmeval` experiment.

### **-h, --help**

Show help message and exist.

### **-V, --version**

Show version number and exit.



## Experiment configuration file

This is a file in `.json` format that provides overall configuration options for an `rsmexplain` experiment. Here's an [example configuration file](#) for `rsmexplain`.

---

**Note:** To make it easy to get started with `rsmexplain`, we provide a way to **automatically generate** configurations file both interactively as well as non-interactively. Novice users will find interactive generation more helpful while more advanced users will prefer non-interactive generation. See [this page](#) for more details.

---

Next, we describe all of the `rsmexplain` configuration fields in detail. There are four required fields and the rest are all optional. We first describe the required fields and then the optional ones (sorted alphabetically).

### background\_data

The path to the background data feature file in one of the [supported formats](#). Each row should correspond to a single response and contain the numeric feature values extracted for this response. In addition, there should be a column containing a unique identifier (ID) for each response. This path can be absolute or relative to the location of the config file. It must contain at least 300 responses to ensure meaningful explanations.

### explain\_data

The path to the file containing the data that we want to explain. The file should be in one of the [supported formats](#). Each row should correspond to a single response and contain numeric feature values extracted for this response. In addition, there should be a column containing a unique identifier (ID) for each response. The path can be absolute or relative to the location of config file.

### experiment\_id

An identifier for the `rsmexplain` experiment. This will be used to name the report. It can be any combination of alphanumeric values, must *not* contain spaces, and must *not* be any longer than 200 characters.

### experiment\_dir

The directory containing the `rsmttool` models we want to explain. This directory should contain an `output` sub-directory and that sub-directory should contain two files: the `<experiment_id>.model` and `<experiment_id>_feature.csv`. Note that `<experiment_id>` refers to the one defined in this same configuration file. As an example of this directory structure, take a look at the `existing_experiment` directory [here](#).

### background\_kmeans\_size (Optional)

The size of the k-means sample for background sampling. Defaults to 500. We summarize the dataset specified in `background_data` with this many k-means clusters (each cluster is weighted by the number of data points it represents) and then use the summarized data set for sampling *instead* of the original. The k-means clustering allows us to speed up the explanation process but may sacrifice some accuracy. The default value of 500 has been shown to provide a good balance between speed and accuracy in our experiments. You may use a higher value if you have a very large or very diverse background dataset and you want to ensure that it's accurately summarized.

**Warning:** `background_kmeans_size` *must* be smaller than the size of the original background data. If not, you may see errors like this: `ValueError: n_samples=500 should be >= n_clusters=750`.

### **custom\_sections** (*Optional*)

A list of custom, user-defined sections to be included into the final report. These are IPython notebooks (.ipynb files) created by the user. The list must contain paths to the notebook files, either absolute or relative to the configuration file. All custom notebooks have access to some *pre-defined variables*.

### **description** (*Optional*)

A brief description of the `rsmexplain` experiment that will be shown at the top of the report. The description can contain spaces and punctuation.

### **general\_sections** (*Optional*)

RSMTTool provides pre-defined sections for `rsmexplain` (listed below) and, by default, all of them are included in the report. However, you can choose a subset of these pre-defined sections by specifying a list as the value for this field.

- `data_description`: Shows the number/identity of responses that are being explained.
- `shap_values`: Shows different types of SHAP values for the features.
- `shap_plots`: Shows various SHAP explanation plots for the features.
- `sysinfo`: Shows all Python packages along with versions installed in the current environment while generating the report.

### **id\_column** (*Optional*)

The name of the column containing the response IDs. Defaults to `spkitemid`, i.e., if this is not specified, `rsmexplain` will look for a column called `spkitemid` in both `background_data` and `explain_data` files. Note that the name of the `id_column` must be the same in these two files.

### **num\_features\_to\_display** (*Optional*)

Number of top features that should be displayed in `rsmexplain` plots. Defaults to 15.

### **sample\_ids** (*Optional*)

If we want to explain a specific set of responses from the `explain_data`, we can specify their IDs here as a comma-separated string. Note that the IDs must be values from the *id\_column*. For example, if `explain_data` has IDs of the form "EXAMPLE\_1", "EXAMPLE\_2", etc., and we want to explain the fifth, tenth, and twelfth example, the value of this field must be set to "EXAMPLE\_5, EXAMPLE\_10, EXAMPLE\_12". Defaults to None.

### **sample\_range** (*Optional*)

If we want to explain a specific range of responses from the `explain_data`, we can specify that range here. Note that the range is specified in terms of the *location* of the responses in the `explain_data` file and that the locations are zero-indexed. So, for example, to explain only the first 50 responses in the file, we should set a value of "0-49" for this option. Defaults to `None`.

### **sample\_size** (*Optional*)

If we want to explain a random sample of the responses in `explain_data`, we can specify the size of that random sample here. For example, to explain a random sample of 10 responses, we would set this to 10. Defaults to `None`.

---

**Note:** Only one of `sample_ids`, `sample_range` or `sample_size` must be specified. If none of them are specified, explanations will be generated for the entire set of responses in `explain_data` which could be very slow, depending on its size.

---

### **section\_order** (*Optional*)

A list containing the order in which the sections in the report should be generated. Any specified order must explicitly list:

1. Either *all* pre-defined sections if a value for the `general_sections` field is not specified OR the sections specified using `general_sections`, and
2. All custom section names specified using `custom_sections`, i.e., file prefixes only, without the path and without the `.ipynb` extension

### **show\_auto\_cohorts** (*Optional*)

If this option is set to `true`, auto cohort bar plots will be displayed. These plots can be useful for detecting interaction effects between cohorts and features. If a cohort shows a high feature value, then there may be an interaction between that cohort and the feature. Defaults to `false`. These plots are not shown by default because these plots may be unstable or provide misleading information if `explain_data` is not large enough. For smaller datasets, SHAP may not be able to detect strong feature interactions and compute clear cohorts. If that happens, the plots will be too specific to be useful. If you have a large enough dataset, you can set this option to `true` and see if the plots are useful.

---

**Important:** By default, the auto cohort bar plots are treated as a custom section and added at the *end* of the report, *after* the system information section. The `section_order` option can be used to move this section to a different place in the report. Use "auto\_cohorts" as the name for this section when specifying an order.

---

### **standardize\_features** (*Optional*)

If this option is set to `false`, the feature values for the responses in `background_data` and `explain_data` will *not* be standardized using the mean and standard deviation parameters for the `rsmttool` experiment. These parameters are expected to be part of the feature information contained in `<experiment_dir>/output/<experiment_id>_feature.csv`. Defaults to `true`.

---

**Important:** If `experiment_dir` contains the `rsmttool` configuration file, that file's value for `standardize_features` will override the value specified by the user. The reason is that if `rsmttool` trained the model with (or without) standardized features, then `rsmexplain` must do the same for the explanations to be meaningful.

---

### **truncate\_outliers** (*Optional*)

If this option is set to `false`, outliers (values more than 4 standard deviations away from the mean) in feature columns will *not* be truncated. Defaults to `true`.

### **use\_wandb** (*Optional*)

If set to `true`, the generated report will be logged to [Weights & Biases](#). The Weights & Biases *entity* and *project name* should be specified in the appropriate configuration fields. The report will be added to a *reports* section in a new run under the given project. Defaults to `false`.

### **wandb\_project** (*Optional*)

The Weights & Biases project name if logging to Weights & Biases is enabled. If a project by this name does not already exist, it will be created.

---

#### **Important:**

1. Before using Weights & Biases for the first time, users should log in and provide their API key as described in [W&B Quickstart guidelines](#).
  2. Note that when using W&B logging, the `rsmttool` run may take significantly longer due to the network traffic being sent to W&B.
- 

### **wandb\_entity** (*Optional*)

The Weights & Biases entity name if logging to Weights & Biases is enabled. Entity can be a user name or the name of a team or organization.

## Output

`rsmexplain` produces a set of folders in the output directory. If logging to [Weights & Biases](#) is *enabled*, the reports are also logged to the specified Weights & Biases project.

### report

This folder contains the final explanation report in HTML format as well as in the form of a Jupyter notebook (a `.ipynb` file).

### output

This folder contains various SHAP values and their absolute value variants. `rsmexplain` also saves a copy of the *configuration file* in this folder. Fields not specified in the original configuration file will be pre-populated with default values. The SHAP explanation object is saved as `<experiment_id>_explanation.pkl` and a mapping between the position of each explained response in the data file and its unique ID is saved in `<experiment_id>_ids.pkl`.

### figure

This folder contains all of the figures containing the various SHAP plots, saved as `.svg` files.

## 1.8 Writing custom RSMTTool sections

RSMTTool allows users to include custom Jupyter notebooks in the reports generated by *rsmttool*, *rsmeval*, and *rsmcompare* using the `custom_sections` fields in their configuration files. This can be particularly useful if a researcher wants to include custom analyses specific to their own scoring engine; she can do so while still using the familiar RSMTTool pipeline.

### 1.8.1 Available variables

When writing such notebooks, some or all of the python variables below will be available for use in the notebooks.

#### `experiment_id`

The experiment ID from the respective configuration file.

#### `description`

The description string from the respective configuration file.

### **model\_name**

The name of the model from the respective configuration file.

### **context**

Two possible values: `rsmtool` or `rsmeval` depending on which command-line tool is being used to generate the report (not available in `rsmcompare` notebooks).

### **use\_scaled\_predictions**

A boolean Python variable containing the value of the setting with the same name from the respective configuration file.

### **exclude\_zero\_scores**

A boolean Python variable containing the value of the setting with the same name from the respective configuration file.

### **length\_column** (*rsmtool only*)

The name of the column in the training and/or evaluation data containing response length. `None` if not specified in the configuration file.

### **second\_human\_score\_column**

The name of the column in the evaluation data containing the second human score. `None` if not specified in the respective configuration file.

### **groups\_eval**

A list containing the names of metadata or subgroup columns as specified in the respective configuration file.

### **min\_items**

The minimal number of items expected from each candidate. The value is set to 0 if the user did not specify a minimal number in the respective configuration file.

### **features\_used** (*rsmtool only*)

A list containing the names of all the features that are used for training the model. [*rsmtool only*]

In addition, several pandas data frames are also available. Many of these contain the same information produced in the intermediate CSV files produced by *rsmtool* or *rsmeval*. We have made these available to authors of custom notebooks to avoid the need for reading them from disk.

**df\_features (rsmtool only)**

A data frame containing information about the feature columns that were included in the final model training. Same information as in *feature.csv*.

**df\_betas (rsmtool only)**

Relative and standardized coefficients (*betas.csv*)

**df\_train\_orig (rsmtool only)****df\_test\_orig (rsmtool only)**

Data frames containing the original training and testing data as specified in the config file, without any changes.

**df\_train (rsmtool only)****df\_train\_preproc (rsmtool only)****df\_test (rsmtool only)****df\_test\_preproc (rsmtool only)**

Data frames containing the *raw* and *pre-processed* feature values.

**df\_train\_other\_columns (rsmtool only)****df\_test\_other\_columns**

Data frames containing the *unused columns* from the training and evaluation data.

**df\_train\_responses\_with\_excluded\_flags (rsmtool only)****df\_test\_responses\_with\_excluded flags**

Data frames containing the *flagged responses*.

**df\_train\_length (rsmtool only)**

A data frame containing response lengths under the `length` column for the training data, along with the response IDs under the `spkitemid` column. These are *only* available (a) if *length\_column* was specified in the configuration file, and (b) if no values in that column are missing, and (c) if the values in that column are not distributed with a standard deviation  $\leq 0$ .

### **df\_test\_human\_scores**

A data frame containing the two human scores for the responses in the evaluation data under the `sc1` and `sc2` columns, along with the response IDs under the `spkitemid` column. This frame is *only* available if `second_human_score_column` was specified in the config file.

---

**Note:** This data frame will contain NaN for the responses for which no numeric second human score was available or for which the second score was 0 and `exclude_zero_scores` was set to `true`.

---

### **df\_pred\_preproc**

A data frame containing the *raw and post-processed predictions for the evaluation data*.

### **df\_feature\_subset\_specs (rsmtool only)**

A data frame containing the contents of *feature\_subset\_file* if it was specified in configuration file. None if not specified.

Finally, the following variables are also available but you are strong encouraged *not* to re-read the files under these directories which are already available as data frames.

### **output\_dir**

The output sub-directory under the experiment output directory that contains all the intermediate CSV files.

### **figure\_dir**

The figure sub-directory under the experiment output directory that contains all of the generated SVG and PNG figures.

---

**Note:** All dataframes apart from `df_train_orig` and `df_test_orig` contain an `spkitemid` column which contains the unique response IDs.

All data frames except the `df_*_other_columns` contain an `sc1` column which contains the human score for the responses.

`df_train_orig` and `df_test_orig` will contain the response IDs and human scores under columns with the original names, *not* `spkitemid` and `sc1`.

---

## 1.9 Auto-generating configuration files

Configuration files for *rsmtool*, *rsmeval*, *rsmexplain*, *rsmcompare*, *rsmpredict*, *rsmsummarize* and *rsmxval* can be difficult to create manually due to the large number of configuration options supported by these tools. To make this easier for users, all of these tools support *automatic* creation of configuration files, both interactively and non-interactively.



### 1.9.1 Interactive generation

For novice users, it is easiest to use a guided, interactive mode to generate a configuration file. For example, to generate an `rsmtool` configuration file interactively, run the following command:

```
rsmtool generate --interactive --output example_rsmtool.json
```

The following screencast shows an example interactive session after the above command is run (click to play):

The configuration file `example_rsmtool.json` generated via the session is shown below:

```
{
  // Reference: https://rsmtool.readthedocs.io/en/stable/usage_rsmtool.html#experiment-
  ↪ configuration-file
  "experiment_id": "test_rsmtool",
  "model": "LinearRegression",
  "train_file": "/Users/nmadnani/train.csv",
  "test_file": "/Users/nmadnani/test.csv",
  // OPTIONAL: replace default values below based on your data.
  "candidate_column": null,
  "custom_sections": null,
  "description": "Test experiment.",
  "exclude_zero_scores": true,
  "feature_subset": null,
  "feature_subset_file": null,
  "features": null,
  "file_format": "xlsx",
  "flag_column": null,
  "flag_column_test": null,
  "general_sections": [
    "data_description",
    "feature_descriptives",
    "preprocessed_features",
    "consistency",
    "model",
    "evaluation",
    "true_score_evaluation",
    "pca",
    "intermediate_file_paths",
    "sysinfo"
  ],
  "id_column": "ID",
  "length_column": "length",
  "min_items_per_candidate": null,
  "min_n_per_group": null,
  "predict_expected_scores": false,
  "second_human_score_column": null,
  "section_order": null,
  "select_transformations": false,
  "sign": null,
  "skll_objective": null,
  "standardize_features": true,
  "subgroups": [],
  "test_label_column": "score",
```

(continues on next page)

(continued from previous page)

```

    "train_label_column": "score",
    "trim_max": 1,
    "trim_min": 6,
    "trim_tolerance": 0.4998,
    "use_scaled_predictions": true,
    "use_thumbnails": false,
    "use_truncation_thresholds": false
}

```

**Note:** Although we use `rsmttool` in the example above, the same instructions apply to all 6 tools; simply replace `rsmttool` with `rsmeval`, `rsmcompare`, etc.

There are some configuration options that can accept multiple inputs. For example, the `experiment_dirs` option for `rsmsummarize` takes a list of `rsmttool` experiment directories for a summary report. These options are handled differently in interactive mode. To illustrate this, let's generate a configuration file for `rsmsummarize` by using the following command:

```
rsmsummarize generate --interactive --output example_rsmsummarize.json
```

The following screencast shows the interactive session (click to play):

And here is the generated configuration file for `rsmsummarize`:

```

{
  // Reference: https://rsmttool.readthedocs.io/en/stable/advanced_usage.html#config-
  ↪file-rsmsummarize
  "summary_id": "test_summary",
  "experiment_dirs": [
    "/Users/nmadnani/work",
    "/Users/nmadnani/work/rsmttool",
    "/Users/nmadnani/work/rsmttool/tests"
  ],
  // OPTIONAL: replace default values below based on your data.
  "custom_sections": null,
  "description": "This is a test.",
  "experiment_names": null,
  "file_format": "tsv",
  "general_sections": [
    "preprocessed_features",
    "model",
    "evaluation",
    "true_score_evaluation",
    "intermediate_file_paths",
    "sysinfo"
  ],
  "section_order": null,
  "subgroups": [],
  "use_thumbnails": true
}

```

**Important:** If you want to include subgroup information in the reports for `rsmttool`, `rsmeval`, `rsmcompare`,

and `rsmxval`, you should add `--subgroups` to the command. For example, when you run `rsmeval generate --interactive --subgroups` you would be prompted to enter the subgroup column names and the `general_sections` list (if shown<sup>1</sup>) will also include subgroup-based sections. Since the `subgroups` option can accept multiple inputs, it is handled in the same way as the `experiment_dirs` option for `rsmsummarize` above.

We end with a list of important things to note about interactive generation:

- Carefully read the instructions and notes displayed at the top when you first enter interactive mode.
- If you do not specify an output file using `--output`, the generated configuration file will simply be printed out.
- You may see messages like “invalid option” and “invalid file” on the bottom left while you are entering the value for a field. This is an artifact of real-time validation. For example, when choosing a training file for `rsmttool`, the message “invalid file” may be displayed while you navigate to the actual file. Once you get to a valid file, this message should disappear.
- Required fields will *not* accept a blank input (just pressing enter) and will show an error message in the bottom left until a valid input is provided.
- Optional fields will accept blank inputs since they have default values that will be used if no user input is provided. In some cases, default values are shown underlined in parentheses.
- You can also use `-i` as an alias for `--interactive` and `-g` as an alias for `--subgroups`. So, for example, if you want to interactively generate a configuration file with subgroups for `rsmttool`, just run `rsmttool generate -ig` instead of `rsmttool generate --interactive --subgroups`.
- The configuration files generated interactively contain comments (as indicated by `// ...`). While `RSMTool` handles JSON files with comments just fine, you may need to remove the comments manually if you wish to use these files outside of `RSMTool`.

## 1.9.2 Non-interactive Generation

For more advanced or experienced users who want to quickly get started with a dummy configuration file that they feel comfortable editing manually, `RSMTool` also provides the capability to generate configuration files non-interactively. To do so, simply omit the `--interactive` switch in the commands above. For example, to generate a dummy configuration file for `rsmttool`, the command to run would be:

```
rsmttool generate --output dummy_rsmttool.json
```

When running this command, the following warning would be printed out to `stderr`:

```
WARNING: Automatically generated configuration files MUST be edited to add values
for required fields and even for optional ones depending on your data
```

This warning explains that the generated file *cannot* be used directly as input to `rsmttool` since the required fields are filled with dummy values. This can be confirmed by looking at the configuration file the command generates:

```
{
  // Reference: https://rsmttool.readthedocs.io/en/stable/usage_rsmttool.html#experiment-
  ↳ configuration-file
  // REQUIRED: replace "ENTER_VALUE_HERE" with the appropriate value!
  "experiment_id": "ENTER_VALUE_HERE",
  "model": "ENTER_VALUE_HERE",
  "train_file": "ENTER_VALUE_HERE",
```

(continues on next page)

<sup>1</sup> Recall that `rsmxval` does not support customizing the section list and, therefore, will not display the `general_sections` field in the auto-generated configuration file.

(continued from previous page)

```

"test_file": "ENTER_VALUE_HERE",
// OPTIONAL: replace default values below based on your data.
"candidate_column": null,
"custom_sections": null,
"description": "",
"exclude_zero_scores": true,
"feature_subset": null,
"feature_subset_file": null,
"features": null,
"file_format": "csv",
"flag_column": null,
"flag_column_test": null,
"general_sections": [
    "data_description",
    "feature_descriptives",
    "preprocessed_features",
    "consistency",
    "model",
    "evaluation",
    "true_score_evaluation",
    "pca",
    "intermediate_file_paths",
    "sysinfo"
],
"id_column": "spkitemid",
"length_column": null,
"min_items_per_candidate": null,
"min_n_per_group": null,
"predict_expected_scores": false,
"second_human_score_column": null,
"section_order": null,
"select_transformations": false,
"sign": null,
"skill_objective": null,
"standardize_features": true,
"subgroups": [],
"test_label_column": "sc1",
"train_label_column": "sc1",
"trim_max": null,
"trim_min": null,
"trim_tolerance": 0.4998,
"use_scaled_predictions": false,
"use_thumbnails": false,
"use_truncation_thresholds": false
}

```

Note the two comments demarcating the locations of the required and optional fields. Note also that the required fields are filled with the dummy value “ENTER\_VALUE\_HERE” that *must* be manually edited by the user. The optional fields are filled with default values that may also need to be further edited depending on the data being used.

Just like interactive generation, non-interactive generation is supported by all 6 tools: `rsmtool`, `rsmeval`, `rsmexplain`, `rsmcompare`, `rsmpredict`, `rsmsummarize`, and `rsmxval`.

Similarly, to include subgroup information in the reports for `rsmtool`, `rsmeval`, and `rsmcompare`, just add

--subgroups (or -g) to the command. Note that unlike in interactive mode, this would *only* add subgroup-based sections to the `general_sections` list in the output file. You will need to manually edit the `subgroups` option in the configuration file to enter the subgroup column names.

### 1.9.3 Generation API

Interactive generation is only meant for end users and can only be used via the 6 command-line tools `rsmttool`, `rsmeval`, `rsmexplain`, `rsmcompare`, `rsmpredict`, `rsmsummarize`, and `rsmxval`. It cannot be used via the RSM-Tool API.

However, the non-interactive generation *can* be used via the API which can be useful for more advanced RSMTTool users. To illustrate, here's some example Python code to generate a configuration for `rsmttool` in the form of a dictionary:

```
# import the ConfigurationGenerator class
from rsmttool.utils.commandline import ConfigurationGenerator

# instantiate it with the options as needed
# we want a dictionary, not a string
# we do not want to see any warnings
# we want to include subgroup-based sections in the report
generator = ConfigurationGenerator('rsmttool',
                                   as_string=False,
                                   suppress_warnings=True,
                                   use_subgroups=True)

# generate the configuration dictionary
configdict = generator.generate()

# remember we still need to replace the dummy values
# for the required fields
configdict["experiment_id"] = "test_experiment"
configdict["model"] = "LinearRegression"
configdict["train_file"] = "train.csv"
configdict["test_file"] = "test.csv"

# and don't forget about adding the subgroups
configdict["subgroups"] = ["GROUP1", "GROUP2"]

# make other changes to optional fields based on your data
...

# now we can use this dictionary to run an rsmttool experiment via the API
from rsmttool import run_experiment
run_experiment(configdict, "/tmp/output")
```

For more details, refer to the API documentation for the *ConfigurationGenerator* class.

## 1.10 API Documentation

The primary method of using RSMTTool is via the command-line scripts *rsmttool*, *rsmeval*, *rsmpredict*, *rsmcompare*, and *rsmsummarize*. However, there are certain functions in the *rsmttool* API that may also be useful to advanced users for use directly in their Python code. We document these functions below.

### 1.10.1 rsmttool Package

```
rsmttool.run_experiment(config_file_or_obj_or_dict: str | Configuration | Dict[str, Any] | Path, output_dir: str,
                        overwrite_output: bool = False, logger: Logger | None = None, wandb_run: Run |
                        RunDisabled | None = None) → None
```

Run an rsmttool experiment using the given configuration.

Run rsmttool experiment using the given configuration file, object, or dictionary. All outputs are generated under `output_dir`. If `overwrite_output` is `True`, any existing output in `output_dir` is overwritten.

#### Parameters

- **config\_file\_or\_obj\_or\_dict** (`Union[str, Configuration, Dict[str, Any], Path]`) – Path to the experiment configuration file as either a string or as a `pathlib.Path` object. Users can also pass a `Configuration` object that is in memory or a Python dictionary with keys corresponding to fields in the configuration file. Given a configuration file, any relative paths in the configuration file will be interpreted relative to the location of the file. Given a `Configuration` object, relative paths will be interpreted relative to the `configdir` attribute, that `_must_` be set. Given a dictionary, the reference path is set to the current directory.
- **output\_dir** (`str`) – Path to the experiment output directory.
- **overwrite\_output** (`bool`) – If `True`, overwrite any existing output under `output_dir`. Defaults to `False`.
- **logger** (`Optional[logging.Logger]`) – A `Logger` object. If `None` is passed, get logger from `__name__`. Defaults to `None`.
- **wandb\_run** (`Union[wandb.wandb_run.Run, wandb.sdk.lib.RunDisabled, None]`) – A `wandb` run object that will be used to log artifacts and tables. If `None` is passed, a new `wandb` run will be initialized if `wandb` is enabled in the configuration. Defaults to `None`.

#### Raises

- **FileNotFoundError** – If any of the files contained in `config_file_or_obj_or_dict` cannot be located.
- **IOError** – If `output_dir` already contains the output of a previous experiment and `overwrite_output` is `False`.
- **ValueError** – If the current configuration specifies a non-linear model but `output_dir` already contains the output of a previous experiment that used a linear model with the same experiment ID.

#### Return type

`None`

```
rsmttool.run_evaluation(config_file_or_obj_or_dict: str | Configuration | Dict[str, Any] | Path, output_dir: str,
                        overwrite_output: bool = False, logger: Logger | None = None, wandb_run: Run |
                        RunDisabled | None = None) → None
```

Run an rsmeval experiment using the given configuration.

All outputs are generated under `output_dir`. If `overwrite_output` is `True`, any existing output in `output_dir` is overwritten.

#### Parameters

- **config\_file\_or\_obj\_or\_dict** (`Union[str, Configuration, Dict[str, Any], Path]`) – Path to the experiment configuration file either a string or as a `pathlib.Path` object. Users can also pass a `Configuration` object that is in memory or a Python dictionary with keys corresponding to fields in the configuration file. Given a configuration file, any relative paths in the configuration file will be interpreted relative to the location of the file. Given a `Configuration` object, relative paths will be interpreted relative to the `configdir` attribute, that `_must_` be set. Given a dictionary, the reference path is set to the current directory.
- **output\_dir** (`str`) – Path to the experiment output directory.
- **overwrite\_output** (`bool`) – If `True`, overwrite any existing output under `output_dir`. Defaults to `False`.
- **logger** (`Optional[logging.Logger]`) – A `Logger` object. If `None` is passed, get logger from `__name__`. Defaults to `None`.
- **wandb\_run** (`Union[wandb.wandb_run.Run, wandb.sdk.lib.RunDisabled, None]`) – A `wandb run` object that will be used to log artifacts and tables. If `None` is passed, a new `wandb run` will be initialized if `wandb` is enabled in the configuration. Defaults to `None`.

#### Raises

- **FileNotFoundError** – If any of the files contained in `config_file_or_obj_or_dict` cannot be located.
- **IOError** – If `output_dir` already contains the output of a previous experiment and `overwrite_output` is `False`.

#### Return type

`None`

```
rsmtool.run_comparison(config_file_or_obj_or_dict: str | Configuration | Dict[str, Any] | Path, output_dir: str)
    → None
```

Run an rsmcompare experiment using the given configuration.

Use the given configuration file, object, or dictionary and generate the report in the given directory.

#### Parameters

- **config\_file\_or\_obj\_or\_dict** (`Union[str, Configuration, Dict[str, Any], Path]`) – Path to the experiment configuration file either a string or as a `pathlib.Path` object. Users can also pass a `Configuration` object that is in memory or a Python dictionary with keys corresponding to fields in the configuration file. Given a configuration file, any relative paths in the configuration file will be interpreted relative to the location of the file. Given a `Configuration` object, relative paths will be interpreted relative to the `configdir` attribute, that `_must_` be set. Given a dictionary, the reference path is set to the current directory.
- **output\_dir** (`str`) – Path to the experiment output directory.

#### Raises

- **FileNotFoundError** – If either of the two input directories in `config_file_or_obj_or_dict` do not exist.
- **FileNotFoundError** – If the directories do not contain rsmttool outputs at all.

#### Return type

None

```
rsmttool.run_summary(config_file_or_obj_or_dict: str | Configuration | Dict[str, Any] | Path, output_dir: str,
                    overwrite_output: bool = False, logger: Logger | None = None, wandb_run: Run |
                    RunDisabled | None = None) → None
```

Run rsmssummarize experiment using the given configuration.

Summarize several rsmttool experiments using the given configuration file, object, or dictionary. All outputs are generated under `output_dir`. If `overwrite_output` is True, any existing output in `output_dir` is overwritten.

#### Parameters

- **config\_file\_or\_obj\_or\_dict** (`Union[str, Configuration, Dict[str, Any], Path]`) – Path to the experiment configuration file either a string or as a `pathlib.Path` object. Users can also pass a `Configuration` object that is in memory or a Python dictionary with keys corresponding to fields in the configuration file. Given a configuration file, any relative paths in the configuration file will be interpreted relative to the location of the file. Given a `Configuration` object, relative paths will be interpreted relative to the `configdir` attribute, that `_must_` be set. Given a dictionary, the reference path is set to the current directory.
- **output\_dir** (`str`) – Path to the experiment output directory.
- **overwrite\_output** (`bool`) – If True, overwrite any existing output under `output_dir`. Defaults to False.
- **logger** (`Optional[logging.Logger]`) – A Logger object. If None is passed, get logger from `__name__`. Defaults to None.
- **wandb\_run** (`Union[wandb.wandb_run.Run, wandb.sdk.lib.RunDisabled, None]`) – A wandb run object that will be used to log artifacts and tables. If None is passed, a new wandb run will be initialized if wandb is enabled in the configuration. Defaults to None.

#### Raises

**IOError** – If `output_dir` already contains the output of a previous experiment and `overwrite_output` is False.

#### Return type

None

```
rsmttool.compute_and_save_predictions(config_file_or_obj_or_dict: str | Configuration | Dict[str, Any] |
                                     Path, output_file: str, feats_file: str | None = None, logger: Logger |
                                     None = None, wandb_run: Run | RunDisabled | None = None) →
                                     None
```

Run rsmpredict using the given configuration.

Generate predictions using given configuration file, object, or dictionary. Predictions are saved in `output_file`. Optionally, pre-processed feature values are saved in `feats_file`, if specified.

#### Parameters

- **config\_file\_or\_obj\_or\_dict** (`Union[str, Configuration, Dict[str, Any], Path]`) – Path to the experiment configuration file either a string or as a `pathlib.Path`



object. Users can also pass a `Configuration` object that is in memory or a Python dictionary with keys corresponding to fields in the configuration file. Given a configuration file, any relative paths in the configuration file will be interpreted relative to the location of the file. Given a `Configuration` object, relative paths will be interpreted relative to the `configdir` attribute, that `_must_` be set. Given a dictionary, the reference path is set to the current directory.

- **output\_file** (*str*) – The path to the output file.
- **feats\_file** (*Optional[str]*) – Path to the output file for saving preprocessed feature values.
- **logger** (*Optional[logging.Logger]*) – A Logger object. If `None` is passed, get logger from `__name__`. Defaults to `None`.
- **wandb\_run** (*Union[wandb.wandb\_run.Run, wandb.sdk.lib.RunDisabled, None]*) – A wandb run object that will be used to log artifacts and tables. If `None` is passed, a new wandb run will be initialized if wandb is enabled in the configuration. Defaults to `None`.

#### Raises

- **FileNotFoundError** – If any of the files contained in `config_file_or_obj_or_dict` cannot be located.
- **FileNotFoundError** – If `experiment_dir` does not exist.
- **FileNotFoundError** – If `experiment_dir` does not contain the required output needed from an rsmttool experiment.
- **RuntimeError** – If the name of the output file does not end in “.csv”, “.tsv”, or “.xlsx”.

#### Return type

`None`

```
rsmttool.fast_predict(input_features: Dict[str, float], modeler: Modeler, df_feature_info: DataFrame | None =
    None, trim: bool = False, trim_min: float | None = None, trim_max: float | None = None,
    trim_tolerance: float | None = None, scale: bool = False, train_predictions_mean: float |
    None = None, train_predictions_sd: float | None = None, h1_mean: float | None = None,
    h1_sd: float | None = None, logger: Logger | None = None) → Dict[str, float]
```

Compute predictions for a single instance against given model.

The main difference between this function and the `compute_and_save_predictions()` function is that the former is meant for batch prediction and reads all its inputs from disk and writes its outputs to disk. This function, however, is meant for real-time inference rather than batch. To this end, it operates *entirely* in memory. Note that there is still a bit of overlap between the two computation paths since we want to use the RSMTTool API as much as possible.

This function should only be used when the goal is to generate predictions using RSMTTool models in production. The user should read everything from disk in a separate thread/function and pass the inputs to this function.

Note that this function only computes regular predictions, not expected scores.

#### Parameters

- **input\_features** (*Dict[str, float]*) – A dictionary containing the features for the instance for which to generate the model predictions. The keys should be names of the features on which the model was trained and the values should be the *raw* feature values.
- **modeler** (*Modeler*) – The RSMTTool `Modeler` object from which the predictions are to be generated. This object should be created from the already existing `.model` file in the “output” directory of the previously run RSMTTool experiment.

- **df\_feature\_info** (*Optional* [*pandas.DataFrame*]) – If *None*, this function will try to extract this information from *modeler*.

A *DataFrame* containing the information regarding the model features. The index of the dataframe should be the names of the features and the columns should be:

- “sign” : 1 or -1. Indicates whether the feature value needs to be multiplied by -1.
- “transform” : *transformation* that needs to be applied to this feature.
- “train\_mean”, “train\_sd” : mean and standard deviation for outlier truncation.
- “train\_transformed\_mean”, “train\_transformed\_sd” : mean and standard deviation for computing z-scores.

This dataframe should be read from the “feature.csv” file under the “output” directory of the previously run RSMTool experiment.

Defaults to *None*.

- **trim** (*bool*) – Whether to trim the predictions. If *True*, *trim\_min* and *trim\_max* must be specified or be available as attributes of the *modeler*. Defaults to *False*.
- **trim\_min** (*Optional* [*float*]) – The lowest possible score that the machine should predict. If *None*, this function will try to extract this value from *modeler*. If *None*, no such attribute exists, and *trim=True*, a *ValueError* will be raised. Defaults to *None*.
- **trim\_max** (*Optional* [*float*]) – The highest possible score that the machine should predict. If *None*, this function will try to extract this value from *modeler*. If *None*, no such attribute exists, and *trim=True*, a *ValueError* will be raised. Defaults to *None*.
- **trim\_tolerance** (*Optional* [*float*]) – The single numeric value that will be used to pad the trimming range specified in *trim\_min* and *trim\_max*. If *None*, this function will try to extract this value from *modeler*. If no such attribute can be found, the value will default to 0.4998. Defaults to *None*.
- **scale** (*bool*) – Whether to scale predictions. If *True*, all of *train\_predictions\_mean*, *train\_predictions\_sd*, *h1\_mean*, and *h1\_sd* must be specified or be available as attributes of *modeler*. Defaults to *False*.
- **train\_predictions\_mean** (*Optional* [*float*]) – The mean of the predictions on the training set used to re-scale the predictions. May be read from the “postprocessing\_params.csv” file under the “output” directory of the RSMTool experiment used to train the model. If *None*, this function will try to extract this value from *modeler*. If *None*, no such attribute exists, and *scale=True*, a *ValueError* will be raised. Defaults to *None*.
- **train\_predictions\_sd** (*Optional* [*float*]) – The standard deviation of the predictions on the training set used to re-scale the predictions. May be read from the “postprocessing\_params.csv” file under the “output” directory of the RSMTool experiment used to train the model. If *None*, this function will try to extract this value from *modeler*. If *None* and no such attribute exists, predictions will not be scaled. Defaults to *None*.
- **h1\_mean** (*Optional* [*float*]) – The mean of the human scores in the training set also used to re-scale the predictions. May be read from the “postprocessing\_params.csv” file under the “output” directory of the RSMTool experiment used to train the model. If *None*, this function will try to extract this value from *modeler*. If *None*, no such attribute exists, and *scale=True*, a *ValueError* will be raised. Defaults to *None*.
- **h1\_sd** (*Optional* [*float*]) – The standard deviation of the human scores in the training set used to re-scale the predictions. May be read from the “postprocessing\_params.csv” file under the “output” directory of the RSMTool experiment used to train the model. If *None*,

this function will try to extract this value from `modeler`. If `None`, no such attribute exists, and `scale=True`, a `ValueError` will be raised. Defaults to `None`.

- **logger** (*Optional*[`logging.Logger`]) – A Logger object. If `None` is passed, get logger from `__name__`. Defaults to `None`.

### Returns

A dictionary containing the raw, scaled, trimmed, and rounded predictions for the input features. It always contains the “raw” key and may contain the following additional keys depending on the availability of the various optional arguments: “raw\_trim”, “raw\_trim\_round”, “scale”, “scale\_trim”, and “scale\_trim\_round”.

### Return type

`Dict[str, float]`

### Raises

- **ValueError** – If `input_features` contains any non-numeric features
- **ValueError** – If trimming/scaling is turned on but related parameters are either not specified or cannot be found as attributes in `modeler`/have a value of `None`
- **ValueError** – If trimming/scaling-related parameters are specified but trimming/scaling is turned off
- **ValueError** – If feature information is either not specified or cannot be found as an attribute in `modeler`/has a value of `None`

`rsmtool.generate_explanation(config_file_or_obj_or_dict: str | Configuration | Dict[str, Any] | Path, output_dir: str, overwrite_output: bool = False, logger: Logger | None = None, wandb_run: Run | RunDisabled | None = None)`

Generate a `shap.Explanation` object.

This function does all the heavy lifting. It loads the model, creates an explainer, and generates an explanation object. It then calls `generate_report()` in order to generate a SHAP report.

### Parameters

- **config\_file\_or\_obj\_or\_dict** (*Union*[`str`, `Configuration`, `Dict[str, Any]`, `Path`]) – Path to the experiment configuration file either as a string or as a `pathlib.Path` object. Users can also pass a `Configuration` object that is in memory or a Python dictionary with keys corresponding to fields in the configuration file. Given a configuration file, any relative paths in the configuration file will be interpreted relative to the location of the file. Given a `Configuration` object, relative paths will be interpreted relative to the `configdir` attribute, that `_must_` be set. Given a dictionary, the reference path is set to the current directory.
- **output\_dir** (`str`) – Path to the experiment output directory.
- **overwrite\_output** (`bool`) – If `True`, overwrite any existing output under `output_dir`. Defaults to `False`.
- **logger** (*Optional*[`logging.Logger`]) – A Logger object. If `None` is passed, get logger from `__name__`. Defaults to `None`.
- **wandb\_run** (*Union*[`wandb.wandb_run.Run`, `wandb.sdk.lib.RunDisabled`, `None`]) – A wandb run object that will be used to log artifacts and tables. If `None` is passed, a new wandb run will be initialized if wandb is enabled in the configuration. Defaults to `None`.

### Raises

- **FileNotFoundError** – If any file contained in `config_file_or_obj_or_dict` cannot be located.
- **ValueError** – If both `sample_range` and `sample_size` are defined in the configuration file.

## From analyzer Module

Classes for analyzing RSMTool predictions, metrics, etc.

### author

Jeremy Biggs ([jbiggs@ets.org](mailto:jbiggs@ets.org))

### author

Anastassia Loukina ([aloukina@ets.org](mailto:aloukina@ets.org))

### author

Nitin Madnani ([nmadnani@ets.org](mailto:nmadnani@ets.org))

### organization

ETS

**class** `rsmttool.analyzer.Analyzer`(*logger: [Logger](#) | [None](#) = [None](#)*)

Bases: [object](#)

Class to perform analysis on all metrics, predictions, etc.

Initialize the Analyzer object.

**static** `analyze_excluded_responses`(*df: [DataFrame](#), features: [List\[str\]](#), header: [str](#), exclude\_zero\_scores: [bool](#) = [True](#), exclude\_listwise: [bool](#) = [False](#)*) → [DataFrame](#)

Compute statistics for responses excluded from analyses.

This method computes various statistics for the responses that were excluded from analyses, either in the training set or in the test set.

### Parameters

- **df** ([pandas.DataFrame](#)) – Data frame containing the excluded responses
- **features** ([List\[str\]](#)) – List of column names containing the features to which we want to restrict the analyses.
- **header** ([str](#)) – String to be used as the table header for the output data frame.
- **exclude\_zero\_scores** ([bool](#)) – Whether or not the zero-score responses should be counted in the exclusion statistics. Defaults to [True](#).
- **exclude\_listwise** ([bool](#)) – Whether or not the candidates were excluded based on minimal number of responses. Defaults to [False](#).

### Returns

**df\_full\_crosstab** – Two-dimensional data frame containing the exclusion statistics.

### Return type

[pandas.DataFrame](#)

**static** `analyze_used_predictions`(*df\_test: [DataFrame](#), subgroups: [List\[str\]](#), candidate\_column: [str](#)*) → [DataFrame](#)

Compute various statistics for predictions used in analyses.

### Parameters

- **df\_test** (*pandas.DataFrame*) – Data frame containing the test set predictions.
- **subgroups** (*List[str]*) – List of column names that contain grouping information.
- **candidate\_column** (*str*) – Column name that contains candidate identification information.

### Returns

**df\_analysis** – Data frame containing information about the used predictions.

### Return type

*pandas.DataFrame*

**static analyze\_used\_responses**(*df\_train: DataFrame, df\_test: DataFrame, subgroups: List[str], candidate\_column: str*) → *DataFrame*

Compute statistics for responses used in analyses.

This method computes various statistics on the responses that were used in analyses, either in the training set or in the test set.

### Parameters

- **df\_train** (*pandas.DataFrame*) – Data frame containing the response information for the training set.
- **df\_test** (*pandas.DataFrame*) – Data frame containing the response information for the test set.
- **subgroups** (*List[str]*) – List of column names that contain grouping information.
- **candidate\_column** (*str*) – Column name that contains candidate identification information.

### Returns

**df\_analysis** – Data frame containing information about the used responses.

### Return type

*pandas.DataFrame*

**static check\_frame\_names**(*data\_container: DataContainer, dataframe\_names: List[str]*) → *None*

Check that all specified dataframes are available.

This method checks to make sure all specified DataFrames are in the given data container object.

### Parameters

- **data\_container** (*DataContainer*) – A DataContainer object
- **dataframe\_names** (*List[str]*) – The names of the DataFrames expected in the DataContainer object.

### Raises

**KeyError** – If a given dataframe\_name is not in the DataContainer object.

### Return type

*None*

**static check\_param\_names**(*configuration\_obj: Configuration, parameter\_names: List[str]*) → *None*

Check that all specified parameters are available.

This method checks to make sure all specified parameters are in the given configuration object.

### Parameters

- **configuration\_obj** (*Configuration*) – A configuration object
- **parameter\_names** (*List[str]*) – The names of the parameters (keys) expected in the Configuration object.

#### Raises

**KeyError** – If a given parameter\_name is not in the Configuration object.

#### Return type

*None*

**static compute\_basic\_descriptives**(*df: DataFrame, selected\_features: List[str]*) → *DataFrame*

Compute basic descriptive statistics for columns in the given data frame.

#### Parameters

- **df** (*pandas.DataFrame*) – Input data frame containing the feature values.
- **selected\_features** (*List[str]*) – List of feature names for which to compute the descriptives.

#### Returns

**df\_desc** – DataFrame containing the descriptives for each of the features.

#### Return type

*pandas.DataFrame*

**compute\_correlations\_by\_group**(*df: DataFrame, selected\_features: List[str], target\_variable: str, grouping\_variable: str, include\_length: bool = False*) → *Tuple[DataFrame, DataFrame, DataFrame]*

Compute marginal and partial correlations against target variable.

This method computes various marginal and partial correlations of the given columns in the given data frame against the target variable for all data and for each level of the grouping variable.

#### Parameters

- **df** (*pandas.DataFrame*) – Input data frame.
- **selected\_features** (*List[str]*) – List of feature names for which to compute the correlations.
- **target\_variable** (*str*) – Feature name indicating the target variable i.e., the dependent variable
- **grouping\_variable** (*str*) – Feature name that contain the grouping information
- **include\_length** (*bool*) – Whether or not to include the length when computing the partial correlations. Defaults to False.

#### Return type

*Tuple[DataFrame, DataFrame, DataFrame]*

#### Returns

- **df\_target\_cors** (*pandas.DataFrame*) – Data frame containing Pearson’s correlation coefficients for marginal correlations between features and *target\_variable*.
- **df\_target\_partcors** (*pandas.DataFrame*) – Data frame containing Pearson’s correlation coefficients for partial correlations between each feature and *target\_variable* after controlling for all other features. If *include\_length* is set to True, the “length” column will not be included in the partial correlation computation.

- **df\_target\_partcors\_no\_length** (*pandas.DataFrame*) – If `include_length` is set to `True`: Data frame containing Pearson’s correlation coefficients for partial correlations between each feature and `target_variable` after controlling for “length”. Otherwise, it will be an empty data frame.

**compute\_degradation\_and\_disattenuated\_correlations**(*df: DataFrame, use\_all\_responses: bool = True*) → *Tuple[DataFrame, DataFrame]*

Compute the degradation in performance when using system score.

This method computes the degradation in performance when using the system to predict the score instead of a second human and also the disattenuated correlations between human and system scores. These are computed as the Pearson’s correlation between the human score and the system score divided by the square root of correlation between two human raters.

For this, we can compute the system performance either only on the double scored data or on the full dataset. Both options have their pros and cons. The default is to use the full dataset. This function also assumes that the `sc2` column exists in the given data frame, in addition to `sc1` and the various types of predictions.

#### Parameters

- **df** (*pandas.DataFrame*) – Input data frame.
- **use\_all\_responses** (*bool*) – Use the full data set instead of only using the double-scored subset. Defaults to `True`.

#### Return type

*Tuple[DataFrame, DataFrame]*

#### Returns

- **df\_degradation** (*pandas.DataFrame*) – Data frame containing the degradation statistics.
- **df\_correlations** (*pandas.DataFrame*) – Data frame containing the human-system correlations, human-human correlations and disattenuated correlation.

**static compute\_disattenuated\_correlations**(*human\_system\_corr: Series, human\_human\_corr: Series*) → *DataFrame*

Compute disattenuated correlations between human and system scores.

These are computed as the Pearson’s correlation between the human score and the system score divided by the square root of correlation between two human raters.

#### Parameters

- **human\_system\_corr** (*pandas.Series*) – Series containing of pearson’s correlation coefficients human-system correlations.
- **human\_human\_corr** (*pandas.Series*) – Series containing of pearson’s correlation coefficients for human-human correlations. This can contain a single value or have the index matching that of human-system correlations.

#### Returns

**df\_correlations** – Data frame containing the human-system correlations, human-human correlations, and disattenuated correlations.

#### Return type

*pandas.DataFrame*

**compute\_metrics**(*df: DataFrame, compute\_shortened: bool = False, use\_scaled\_predictions: bool = False, include\_second\_score: bool = False, population\_sd\_dict: Dict[str, float | None] | None = None, population\_mn\_dict: Dict[str, float | None] | None = None, smd\_method: str = 'unpooled', use\_diff\_std\_means: bool = False*) → *Tuple[DataFrame, DataFrame, DataFrame]*



Compute association metrics for scores in the given data frame.

This function compute association metrics for all score types. If `include_second_score` is `True`, then it is assumed that a column called `sc2` containing a second human score is available and it should be used to compute the human-human evaluation stats and the performance degradation statistics.

If `compute_shortened` is `True`, then this function also computes a shortened version of the full human-system metrics data frame. See `filter_metrics()` for the description of the default columns included in the shortened data frame.

#### Parameters

- **df** (*pandas.DataFrame*) – Input data frame
- **compute\_shortened** (*bool*) – Also compute a shortened version of the full metrics data frame. Defaults to `False`.
- **use\_scaled\_predictions** (*bool*) – Use evaluations based on scaled predictions in the shortened version of the metrics data frame. Defaults to `False`.
- **include\_second\_score** (*bool*) – Second human score available. Defaults to `False`.
- **population\_sd\_dict** (*Optional[Dict[str, Optional[float, None]]]*) – Dictionary containing population standard deviation for each column containing human or system scores. This is used to compute SMD for subgroups. If `None`, a dummy dictionary is created that sets the standard deviation for all columns to `None`. Defaults to `None`.
- **population\_mn\_dict** (*Optional[Dict[str, Optional[float]]]*) – Dictionary containing population mean for each column containing human or system scores. This is used to compute SMD for subgroups. If `None`, a dummy dictionary is created that sets the standard deviation for all columns to `None`. Defaults to `None`.
- **smd\_method** (*str*) – The SMD method to use, only used if `use_diff_std_means` is `False`. All methods have the same numerator  $\text{mean}(y_{\text{pred}}) - \text{mean}(y_{\text{true\_observed}})$  and the following denominators:
  - “williamson”: pooled population standard deviation of human and system scores computed based on values in `population_sd_dict`.
  - “johnson”: population standard deviation of human scores computed based on values in `population_sd_dict`.
  - “pooled”: pooled standard deviation of `y_true_observed` and `y_pred` for this group.
  - “unpooled”: standard deviation of `y_true_observed` for this group.
 Defaults to “unpooled”.
- **use\_diff\_std\_means** (*bool*) – Whether to use the difference of standardized means, rather than the standardized mean difference. This is most useful with subgroup analysis. Defaults to `False`.

#### Return type

`Tuple[DataFrame, DataFrame, DataFrame]`

#### Returns

- **df\_human\_system\_eval** (*pandas.DataFrame*) – Data frame containing the full set of evaluation metrics.
- **df\_human\_system\_eval\_filtered** (*pandas.DataFrame*) – Data frame containing the human-human statistics but is empty if `include_second_score` is `False`.



- **df\_human\_human\_eval** (*pandas.DataFrame*) – A shortened version of the first data frame but is empty if `compute_shortened` is `False`.

**compute\_metrics\_by\_group**(*df\_test: DataFrame*, *grouping\_variable: str*, *use\_scaled\_predictions: bool = False*, *include\_second\_score: bool = False*) → *Tuple[DataFrame, DataFrame]*

Compute a subset of evaluation metrics by subgroups.

This method computes a subset of evaluation metrics for the scores in the given data frame by group specified in `grouping_variable`. See `filter_metrics()` above for a description of the subset that is selected.

#### Parameters

- **df\_test** (*pandas.DataFrame*) – Input data frame.
- **grouping\_variable** (*str*) – Feature name indicating the column that contains grouping information.
- **use\_scaled\_predictions** (*bool*) – Include scaled predictions when computing the evaluation metrics. Defaults to `False`.
- **include\_second\_score** (*bool*) – Include human-human association statistics. Defaults to `False`.

#### Return type

*Tuple[DataFrame, DataFrame]*

#### Returns

- **df\_human\_system\_by\_group** (*pandas.DataFrame*) – Data frame containing the correlation human-system association statistics.
- **df\_human\_human\_by\_group** (*pandas.DataFrame*) – Data frame that either contains the human-human statistics or is an empty data frame, depending on whether `include_second_score` is `True`.

**static compute\_outliers**(*df: DataFrame*, *selected\_features: List[str]*) → *DataFrame*

Compute number and percentage of outliers for given columns.

This method computes the number and percentage of outliers that lie outside the range mean +/- 4 SD for each of the given columns in the given data frame.

#### Parameters

- **df** (*pandas.DataFrame*) – Input data frame containing the feature values.
- **selected\_features** (*List[str]*) – List of feature names for which to compute outlier information.

#### Returns

**df\_output** – Data frame containing outlier information for each of the features.

#### Return type

*pandas.DataFrame*

**static compute\_pca**(*df: DataFrame*, *selected\_features: List[str]*) → *Tuple[DataFrame, DataFrame]*

Compute PCA decomposition of the given features.

This method computes the PCA decomposition of features in the data frame, restricted to the given columns. The number of components is set to be `min(n_features, n_samples)`.

#### Parameters

- **df** (*pandas.DataFrame*) – Input data frame containing feature values.

- **selected\_features** (*List[str]*) – List of feature names to be used in the PCA decomposition.

#### Return type

*Tuple[DataFrame, DataFrame]*

#### Returns

- **df\_components** (*pandas.DataFrame*) – Data frame containing the PCA components.
- **df\_variance** (*pandas.DataFrame*) – Data frame containing the variance information.

**static compute\_percentiles**(*df: DataFrame, selected\_features: List[str], percentiles: List[int] | None = None*) → *DataFrame*

Compute percentiles and outliers for columns in the given data frame.

#### Parameters

- **df** (*pandas.DataFrame*) – Input data frame containing the feature values.
- **selected\_features** (*List[str]*) – List of feature names for which to compute the percentile descriptives.
- **percentiles** (*Optional[List[int]]*) – The percentiles to calculate. If *None*, use the percentiles {1, 5, 25, 50, 75, 95, 99}. Defaults to *None*.

#### Returns

**df\_output** – Data frame containing the percentile information for each of the features.

#### Return type

*pandas.DataFrame*

**static correlation\_helper**(*df: DataFrame, target\_variable: str, grouping\_variable: str, include\_length: bool = False*) → *Tuple[DataFrame, DataFrame, DataFrame]*

Compute marginal and partial correlations for all columns.

This helper method computes marginal and partial correlations of all the columns in the given data frame against the target variable separately for each level in the the grouping variable. If *include\_length* is *True*, it additionally computes partial correlations of each column in the data frame against the target variable after controlling for the “length” column.

#### Parameters

- **df** (*pandas.DataFrame*) – Input data frame containing numeric feature values, the numeric *target variable* and the *grouping variable*.
- **target\_variable** (*str*) – The name of the column used as a reference for computing correlations.
- **grouping\_variable** (*str*) – The name of the column defining groups in the data
- **include\_length** (*bool*) – If *True* compute additional partial correlations of each column in the data frame against *target variable* only partialling out “length” column.

#### Return type

*Tuple[DataFrame, DataFrame, DataFrame]*

#### Returns

- **df\_target\_cors** (*pandas.DataFrame*) – Data frame containing Pearson’s correlation coefficients for marginal correlations between features and *target\_variable*.

- **df\_target\_partcors** (*pandas.DataFrame*) – Data frame containing Pearson’s correlation coefficients for partial correlations between each feature and *target\_variable* after controlling for all other features. If *include\_length* is set to *True*, the “length” column will not be included in the partial correlation computation.
- **df\_target\_partcors\_no\_length** (*pandas.DataFrame*) – If *include\_length* is set to *True*: Data frame containing Pearson’s correlation coefficients for partial correlations between each feature and *target\_variable* after controlling for “length”. Otherwise, it will be an empty data frame.

**filter\_metrics**(*df\_metrics: DataFrame*, *use\_scaled\_predictions: bool = False*, *chosen\_metric\_dict: Dict[str, List[str]] | None = None*) → *DataFrame*

Filter data frame to retain only the given metrics.

This method filters the data frame *df\_metrics* – containing all of the metric values by all score types (raw, raw\_trim etc.) – to retain only the metrics as defined in the given dictionary *chosen\_metric\_dict*. This dictionary maps score types (“raw”, “scale”, “raw\_trim” etc.) to metric names. The available metric names are:

- “corr”
- “kappa”
- “wtkappa”
- “exact\_agr”
- “adj\_agr”
- “SMD” or “DSM”, depending on what is in *df\_metrics*.
- “RMSE”
- “R2”
- “sys\_min”
- “sys\_max”
- “sys\_mean”
- “sys\_sd”
- “h\_min”
- “h\_max”
- “h\_mean”
- “h\_sd”
- “N”

#### Parameters

- **df\_metrics** (*pandas.DataFrame*) – The *DataFrame* to filter.
- **use\_scaled\_predictions** (*bool*) – Whether to use scaled predictions. Defaults to *False*.
- **chosen\_metric\_dict** (*Optional[Dict[str, List[str]]]*) – The dictionary mapping each score type to the metrics that should be computed for it. Defaults to *None*.

#### Returns

**df\_filtered\_metrics** – The filtered *DataFrame*.

### Return type

`pandas.DataFrame`

**Note:** The last five metrics will be the *same* for all score types. If `chosen_metric_dict` is not specified, the following default dictionary with the recommended metrics is used:

```
{
    "X_trim": ["N", "h_mean", "h_sd", "sys_mean", "sys_sd", "wtkappa",
               "corr", "RMSE", "R2", "SMD"],
    "X_trim_round": ["sys_mean", "sys_sd", "kappa",
                     "exact_agr", "adj_agr", "SMD"]}

```

where X = “raw” or “scale” depending on whether `use_scaled_predictions` is `False` or `True`, respectively.

```
static metrics_helper(human_scores: Series, system_scores: Series, population_human_score_sd: float
    | None = None, population_system_score_sd: float | None = None,
    population_human_score_mn: float | None = None,
    population_system_score_mn: float | None = None, smd_method: str =
    'unpooled', use_diff_std_means: bool = False) → Series

```

Compute basic association metrics between system and human scores.

### Parameters

- **human\_scores** (`pandas.Series`) – Series containing numeric human (reference) scores.
- **system\_scores** (`pandas.Series`) – Series containing numeric scores predicted by the model.
- **population\_human\_score\_sd** (`Optional[float]`) – Reference standard deviation for human scores. This must be specified when the function is used to compute association metrics for a subset of responses, for example, responses from a particular demographic subgroup. If `smd_method` is set to “williamson” or “johnson”, this should be the standard deviation for the whole population (in most cases, the standard deviation for the whole test set). If `use_diff_std_means` is `True`, this must be the standard deviation for the whole population and `population_human_score_mn` must also be specified. Otherwise, it is ignored. Defaults to `None`.
- **population\_system\_score\_sd** (`Optional[float]`) – Reference standard deviation for system scores. This must be specified when the function is used to compute association metrics for a subset of responses, for example, responses from a particular demographic subgroup. If `smd_method` is set to “williamson”, this should be the standard deviation for the whole population (in most cases, the standard deviation for the whole test set). If `use_diff_std_means` is `True`, this must be the standard deviation for the whole population and `population_system_score_mn` must also be specified. Otherwise, it is ignored. Defaults to `None`.
- **population\_human\_score\_mn** (`Optional[float]`) – Reference mean for human scores. This must be specified when the function is used to compute association metrics for a subset of responses, for example, responses from a particular demographic subgroup. If `use_diff_std_means` is `True`, this must be the mean for the whole population (in most cases, the full test set) and `population_human_score_sd` must also be specified. Otherwise, it is ignored. Defaults to `None`.
- **population\_system\_score\_mn** (`Optional[float]`) – Reference mean for system scores. This must be specified when the function is used to compute association metrics for a subset of responses, for example, responses from a particular demographic subgroup.

If `use_diff_std_means` is `True`, this must be the mean for the whole population (in most cases, the full test set) and `population_system_score_sd` must also be specified. Otherwise, it is ignored. Defaults to `None`.

- **`smd_method`** (*str*) – The SMD method to use, only used if `use_diff_std_means` is `False`. All methods have the same numerator  $\text{mean}(y_{\text{pred}}) - \text{mean}(y_{\text{true\_observed}})$  and the following denominators :
  - “williamson”: pooled population standard deviation of `y_true_observed` and `y_pred` computed using `population_human_score_sd` and `population_system_score_sd`.
  - “johnson”: `population_human_score_sd`.
  - “pooled”: pooled standard deviation of `y_true_observed` and `y_pred` for this group.
  - “unpooled”: standard deviation of `y_true_observed` for this group.
 Defaults to “unpooled”.
- **`use_diff_std_means`** (*bool*) – Whether to use the difference of standardized means, rather than the standardized mean difference. This is most useful with subgroup analysis. Defaults to `False`.

## Returns

**`metrics`** – Series containing different evaluation metrics comparing human and system scores. The following metrics are included:

- *kappa*: unweighted Cohen’s kappa
- *wtkappa*: quadratic weighted kappa
- *exact\_agr*: exact agreement
- *adj\_agr*: adjacent agreement with tolerance set to 1
- One of the following :
  - *SMD*: standardized mean difference, if `use_diff_std_means` is `False`.
  - *DSM*: difference of standardized means, if `use_diff_std_means` is `True`.
- *corr*: Pearson’s r
- *R2*: r squared
- *RMSE*: root mean square error
- *sys\_min*: min system score
- *sys\_max*: max system score
- *sys\_mean*: mean system score (ddof=1)
- *sys\_sd*: standard deviation of system scores (ddof=1)
- *h\_min*: min human score
- *h\_max*: max human score
- *h\_mean*: mean human score (ddof=1)
- *h\_sd*: standard deviation of human scores (ddof=1)
- *N*: total number of responses

### Return type

`pandas.Series`

**run\_data\_composition\_analyses\_for\_rsmeval**(*data\_container*: `DataContainer`, *configuration*: `Configuration`) → `Tuple[Configuration, DataContainer]`

Run all data composition analyses for RSMEval.

### Parameters

- **data\_container** (`DataContainer`) – The `DataContainer` object which must include the following `DataFrames`: {“test\_metadata”, “test\_excluded”}.
- **configuration** (`Configuration`) – The `Configuration` object which must include the following parameters (keys): {“subgroups”, “candidate\_column”, “exclude\_zero\_scores”, “exclude\_listwise”}.

### Return type

`Tuple[Configuration, DataContainer]`

### Returns

- **configuration** (`Configuration`) – The input `Configuration` object that is passed through unmodified.
- **data\_container** (`DataContainer`) – A new `DataContainer` object with the following `DataFrames`:
  - test\_excluded\_composition
  - data\_composition
  - data\_composition\_by\_\*

**run\_data\_composition\_analyses\_for\_rsmtool**(*data\_container*: `DataContainer`, *configuration*: `Configuration`) → `Tuple[Configuration, DataContainer]`

Run all data composition analyses for RSMTool.

### Parameters

- **data\_container** (`DataContainer`) – The `DataContainer` object which must include the following `DataFrames`: {“test\_metadata”, “train\_metadata”, “train\_excluded”, “test\_excluded”, “train\_features”}.
- **configuration** (`Configuration`) – The `Configuration` object which must include the following parameters (keys): {“subgroups”, “candidate\_column”, “exclude\_zero\_scores”, “exclude\_listwise”}.

### Return type

`Tuple[Configuration, DataContainer]`

### Returns

- **configuration** (`Configuration`) – The input `Configuration` object that is passed through unmodified.
- **data\_container** (`DataContainer`) – A new `DataContainer` object with the following `DataFrames`:
  - test\_excluded\_composition
  - train\_excluded\_composition

- data\_composition
- data\_composition\_by\_\*

**run\_prediction\_analyses**(data\_container: [DataContainer](#), configuration: [Configuration](#), wandb\_run: *Run | RunDisabled | None = None*) → [Tuple](#)[[Configuration](#), [DataContainer](#)]

Run all analyses on the system scores (predictions).

#### Parameters

- **data\_container** ([DataContainer](#)) – The [DataContainer](#) object which must include the following DataFrames: {“train\_features”, “train\_metadata”, “train\_preprocessed\_features”, “train\_length”, “train\_features”}.
- **configuration** ([Configuration](#)) – The [Configuration](#) object which must include the following parameters (keys): {“subgroups”, “second\_human\_score\_column”, “use\_scaled\_predictions”}.
- **wandb\_run** (*Union[[wandb.wandb\\_run.Run](#), [wandb.sdk.lib.RunDisabled](#), [None](#)]*) – The wandb run object if wandb is enabled, [None](#) otherwise. If enabled, all the output data frames will be logged to this run as tables. Defaults to [None](#).

#### Return type

[Tuple](#)[[Configuration](#), [DataContainer](#)]

#### Returns

- **configuration** ([Configuration](#)) – The input [Configuration](#) object that is passed through unmodified.
- **data\_container** ([DataContainer](#)) – A new [DataContainer](#) object with the following DataFrames:
  - eval
  - eval\_short
  - consistency
  - degradation
  - disattenuated\_correlations
  - confMatrix
  - confMatrix\_h1h2
  - score\_dist
  - eval\_by\_\*
  - consistency\_by\_\*
  - disattenuated\_correlations\_by\_\*
  - true\_score\_eval

**run\_training\_analyses**(data\_container: [DataContainer](#), configuration: [Configuration](#)) → [Tuple](#)[[Configuration](#), [DataContainer](#)]

Run all analyses on the training data.

#### Parameters

- **data\_container** ([DataContainer](#)) – The [DataContainer](#) object which must include the following DataFrames: {“train\_features”, “train\_metadata”, “train\_preprocessed\_features”, “train\_length”, “train\_features”}.

- **configuration** (*Configuration*) – The Configuration object which must include the following parameters (keys): {“length\_column”, “subgroups”, “selected\_features”}.

**Return type**

`Tuple[Configuration, DataContainer]`

**Returns**

- **configuration** (*Configuration*) – The input Configuration object that is passed through unmodified.
- **data\_container** (*DataContainer*) – A new DataContainer object with the following DataFrames:
  - feature\_descriptives
  - feature\_descriptivesExtra
  - feature\_outliers
  - cors\_orig
  - cors\_processed
  - margcor\_score\_all\_data
  - pcor\_score\_all\_data
  - pcor\_score\_no\_length\_all\_data
  - margcor\_length\_all\_data
  - pcor\_length\_all\_data
  - pca
  - pcavar
  - margcor\_length\_by\_\*
  - pcor\_length\_by\_\*
  - margcor\_score\_by\_\*
  - pcor\_score\_by\_\*
  - pcor\_score\_no\_length\_by\_\*

**From `comparer` Module**

Classes for comparing outputs of two RSMTTool experiments.

**author**

Jeremy Biggs ([jbiggs@ets.org](mailto:jbiggs@ets.org))

**author**

Anastassia Loukina ([aloukina@ets.org](mailto:aloukina@ets.org))

**author**

Nitin Madnani ([nmadnani@ets.org](mailto:nmadnani@ets.org))

**organization**

ETS



**class** `rsmttool.comparer.Comparer`

Bases: `object`

Class to perform comparisons between two RSMTool experiments.

**static** `compute_correlations_between_versions`(*df\_old*: `DataFrame`, *df\_new*: `DataFrame`,  
*human\_score*: *str* = 'sc1', *id\_column*: *str* =  
'spkitemid') → `DataFrame`

Compute correlations between old and new feature values.

This method computes correlations between old and new feature values in the two given frames as well as the correlations between each feature value and the human score.

#### Parameters

- **df\_old** (`pandas.DataFrame`) – Data frame with feature values for the ‘old’ model.
- **df\_new** (`pandas.DataFrame`) – Data frame with feature values for the ‘new’ model.
- **human\_score** (*str*) – Name of the column containing human score. Must be the same for both data sets. Defaults to "sc1".
- **id\_column** (*str*) – Name of the column containing id for each response. Must be the same for both data sets. Defaults to "spkitemid".

#### Returns

**df\_correlations** –

**Data frame with a row for each feature and the following columns:**

- "N": total number of responses
- "human\_old": correlation with human score in the old frame
- "human\_new": correlation with human score in the new frame
- "old\_new": correlation between old and new frames

#### Return type

`pandas.DataFrame`

#### Raises

- **ValueError** – If there are no shared features between the two sets.
- **ValueError** – If there are no shared responses between the two sets.

**load\_rsmttool\_output**(*filedir*: *str*, *figdir*: *str*, *experiment\_id*: *str*, *prefix*: *str*, *groups\_eval*: `List[str]`) →  
`Tuple[Dict[str, DataFrame], Dict[str, str], str]`

Load all of the outputs of an rsmttool experiment.

For each type of output, we first check whether the file exists to allow comparing experiments with different sets of outputs.

#### Parameters

- **filedir** (*str*) – Path to the directory containing output files.
- **figdir** (*str*) – Path to the directory containing output figures.
- **experiment\_id** (*str*) – Original `experiment_id` used to generate the output files.
- **prefix** (*str*) – Must be set to "scale" or "raw". Indicates whether the score is scaled or not.
- **groups\_eval** (`List[str]`) – List of subgroup names used for subgroup evaluation.

**Return type**

`Tuple[Dict[str, DataFrame], Dict[str, str], str]`

**Returns**

- **files** (*Dict[str, pd.DataFrame]*) – A dictionary mapping data frame names converted to the actual pandas data frames. If a particular type of output did not exist for the experiment, its value will be an empty data frame.
- **figs** (*Dict[str, str]*) – A dictionary mapping figure names to the paths of the files containing the figures.
- **file\_format** (*str*) – The file format used for the output files, e.g., csv, tsv, or xlsx.

**static** `make_summary_stat_df(df: DataFrame) → DataFrame`

Compute summary statistics for the data in the given frame.

**Parameters**

**df** (*pandas.DataFrame*) – Data frame containing numeric data.

**Returns**

**res** – Data frame containing summary statistics for data in the input frame.

**Return type**

`pandas.DataFrame`

**static** `process_confusion_matrix(conf_matrix: DataFrame) → DataFrame`

Add “human” and “machine” to column names in the confusion matrix.

**Parameters**

**conf\_matrix** (*pandas.DataFrame*) – data frame containing the confusion matrix.

**Returns**

**conf\_matrix\_renamed** – pandas Data Frame containing the confusion matrix with the columns renamed.

**Return type**

`pandas.DataFrame`

## From configuration\_parser Module

Configuration parser.

Classes related to parsing configuration files and creating configuration objects.

**author**

Jeremy Biggs ([jbiggs@ets.org](mailto:jbiggs@ets.org))

**author**

Anastassia Loukina ([aloukina@ets.org](mailto:aloukina@ets.org))

**author**

Nitin Madnani ([nmadnani@ets.org](mailto:nmadnani@ets.org))

**organization**

ETS

**class** `rsmttool.configuration_parser.Configuration`(*configdict: Dict[str, Any], \*, configdir: str | Path | None = None, context: str = 'rsmttool', logger: Logger | None = None*)

Bases: `object`

Configuration class.

Encapsulates all of the configuration parameters and methods to access these parameters.

Create an object of the *Configuration* class.

This method can be used to directly instantiate a Configuration object.

#### Parameters

- **configdict** (*Dict*[*str*, *Any*]) – A dictionary of configuration parameters. The dictionary must be a valid configuration dictionary with default values filled as necessary.
- **configdir** (*Optional*[*Union*[*str*, *Path*]]) – The reference path used to resolve any relative paths in the configuration object. When *None*, will be set during initialization to the current working directory. Defaults to *None*,
- **context** (*str*) – The context of the tool. One of {"rsmtool", "rsmeval", "rsmcompare", "rsmpredict", "rsmsummarize"}. Defaults to "rsmtool".
- **logger** (*Optional*[*logging.Logger*]) – A Logger object. If *None* is passed, get logger from `__name__`. Defaults to *None*.

**check\_exclude\_listwise()** → *bool*

Check for candidate exclusion.

Check if we are excluding candidates based on number of responses, and add this to the configuration file.

#### Returns

**exclude\_listwise** – Whether to exclude list-wise.

#### Return type

*bool*

**check\_flag\_column**(*flag\_column*: *str* = 'flag\_column', *partition*: *str* = 'unknown') → *Dict*[*str*, *List*[*str*]]

Make sure the column name in *flag\_column* is correctly specified.

Get flag columns and values for filtering, if any, and convert single values to lists. Raises an exception if the column name in *flag\_column* is not correctly specified in the configuration file.

#### Parameters

- **flag\_column** (*str*) – The flag column name to check. Currently used names are "flag\_column" or "flag\_column\_test". Defaults to "flag\_column".
- **partition** (*str*) – The data partition which is filtered based on the flag column name. One of {"train", "test", "both", "unknown"}. Defaults to "unknown".

#### Returns

**new\_filtering\_dict** – Properly formatted dictionary for the column name in *flag\_column*.

#### Return type

*Dict*[*str*, *List*[*str*]]

#### Raises

- **ValueError** – If the specified value of the column name in *flag\_column* is not a dictionary.
- **ValueError** – If the value of *partition* is not in the expected list.
- **ValueError** – If the value of *partition* does not match the *flag\_column*.

**property configdir:** `str`

Get the path to configuration directory.

Get the path to the configuration reference directory that will be used to resolve any relative paths in the configuration.

**Returns**

**configdir** – The path to the configuration reference directory.

**Return type**

`str`

**property context:** `str`

Get the context.

**copy**(*deep*: `bool = True`) → *Configuration*

Return a copy of the object.

**Parameters**

**deep** (`bool`) – Whether to perform a deep copy. Defaults to True.

**Returns**

**copy** – A new configuration object.

**Return type**

*Configuration*

**get**(*key*: `str`, *default*: `Any = None`) → `Any`

Get value or default for the given key.

**Parameters**

- **key** (`str`) – Key to check in the Configuration object.
- **default** (`Any`) – The default value to return, if no key exists. Defaults to None.

**Returns**

The value in the configuration object dictionary.

**Return type**

`value`

**get\_default\_converter**() → `Dict[str, Any]`

Get default converter dictionary for data reader.

**Returns**

**default\_converter** – The default converter for a train or test file.

**Return type**

`Dict[str, Any]`

**get\_names\_and\_paths**(*keys*: `List[str]`, *names*: `List[str]`) → `Tuple[List[str], List[str]]`

Get values (paths) for the given keys and names.

This method is mainly used to retrieve values that are paths and it skips any values that are None.

**Parameters**

- **keys** (`List[str]`) – A list of keys whose values to retrieve.
- **names** (`List[str]`) – The names corresponding to the keys.

**Return type**

`Tuple[List[str], List[str]]`

**Returns**

- **existing\_names** (*List[str]*) – The names for values that were not None.
- **existing\_paths** (*List[str]*) – The paths (values) for the given keys (that were not None).

**Raises**

**ValueError** – If there are any duplicate keys or names.

**get\_rater\_error\_variance()** → *float*

Get specified rater error variance, if any, and make sure it's numeric.

**Returns**

**rater\_error\_variance** – Specified rater error variance.

**Return type**

*float*

**get\_trim\_min\_max\_tolerance()** → *Tuple[float | None, float | None, float | None]*

Get trim min, trim max, and tolerance values.

Get the specified trim min and max, and trim\_tolerance if any, and make sure they are numeric.

**Return type**

*Tuple[Optional[float], Optional[float], Optional[float]]*

**Returns**

- **spec\_trim\_min** (*Optional[float]*) – Specified trim min value.
- **spec\_trim\_max** (*Optional[float]*) – Specified trim max value.
- **spec\_trim\_tolerance** (*Optional[float]*) – Specified trim tolerance value.

**items()** → *List[Tuple[str, Any]]*

Return configuration items as a list of tuples.

**Returns**

**items** – A list of (key, value) tuples in the configuration object.

**Return type**

*List[Tuple[str, Any]]*

**keys()** → *List[str]*

Return keys as a list.

**Returns**

**keys** – A list of keys in the configuration object.

**Return type**

*List[str]*

**pop(key: str, default: Any = None)** → *Any*

Remove and return an element from the object having the given key.

**Parameters**

- **key** (*str*) – Key to pop in the configuration object.
- **default** (*Any*) – The default value to return, if no key exists. Defaults to None.

**Returns**

**value** – The value removed from the object.

**Return type**

Any

**save**(*output\_dir*: *str* | *None* = *None*) → *None*

Save the configuration file to the output directory specified.

**Parameters**
**output\_dir** (*Optional*[*str*]) – The path to the output directory. If *None*, the current directory is used. Defaults to *None*.

**Return type**
*None*
**to\_dict**() → *Dict*[*str*, *Any*]

Get a dictionary representation of the configuration object.

**Returns**
**config** – The configuration dictionary.

**Return type**
*Dict*[*str*, *Any*]

**values**() → *List*[*Any*]

Return configuration values as a list.

**Returns**
**values** – A list of values in the configuration object.

**Return type**
*List*[*Any*]

**class** *rsmttool.configuration\_parser.ConfigurationParser*(*pathlike*: *str* | *Path*, *logger*: *Logger* | *None* = *None*)

Bases: *object*
*ConfigurationParser* class to create *Configuration* objects.

Instantiate a *ConfigurationParser* for a given config file path.

**Parameters**

- **pathlike** (*Union*[*str*, *Path*]) – A string containing the path to the configuration file that is to be parsed. A *pathlib.Path* instance is also acceptable.
- **logger** (*Optional*[*logging.Logger*]) – Custom logger object to use, if not *None*. Otherwise a new logger is created. Defaults to *None*.

**Raises**

- **FileNotFoundError** – If the given path does not exist.
- **OSError** – If the given path is a directory, not a file.
- **ValueError** – If the file at the given path does not have a valid extension (“.json”).

**logger** = *None*
**parse**(*context*: *str* = ‘rsmtool’) → *Configuration*

Parse configuration file.

Parse the configuration file for which this parser was instantiated.

### Parameters

**context** (*str*) – Context of the tool in which we are validating. One of: {"rsmtool", "rsmeval", "rsmpredict", "rsmcompare", "rsmsummarize", "rsmxval", "rsmexplain"}. Defaults to "rsmtool".

### Returns

**configuration** – A configuration object containing the parameters in the file that we instantiated the parser for.

### Return type

*Configuration*

**classmethod process\_config**(*config: Dict[str, Any]*) → *Dict[str, Any]*

Process the given configuration dictionary.

Converts fields which are read in as string to the appropriate format. Fields which can take multiple string values are converted to lists if they have not been already formatted as such.

### Parameters

**config** (*Dict[str, Any]*) – Given Configuration dictionary to be processed.

### Returns

**new\_config** – A copy of the given configuration dictionary with all fields converted to the appropriate format.

### Return type

*Dict[str, Any]*

### Raises

- **NameError** – If config does not exist, or config could not be read.
- **ValueError** – If boolean configuration fields contain a value other than "true" or "false" (in JSON).

**classmethod validate\_config**(*config: Dict[str, Any], context: str = 'rsmtool'*) → *Dict[str, Any]*

Validate the given configuration dictionary.

Ensure that all required fields are specified, add default values values for all unspecified fields, and ensure that all specified fields are valid.

### Parameters

- **config** (*Dict[str, Any]*) – Given configuration dictionary to be validated.
- **context** (*str*) – Context of the tool in which we are validating. One of {"rsmtool", "rsmeval", "rsmpredict", "rsmcompare", "rsmsummarize"}. Defaults to "rsmtool".

### Returns

**new\_config** – A copy of the given configuration dictionary with all required fields specified and with the default values for unspecified fields.

### Return type

*Dict[str, Any]*

### Raises

**ValueError** – If the Configuration object does not contain all required fields, has any unrecognized fields, or has any fields with invalid values.

`rsmtool.configuration_parser.config`(*context: str, config\_file\_or\_obj\_or\_dict: str | Configuration | Dict[str, Any] | Path*) → *Configuration*

Create a Configuration object.

Get the configuration for context from the input `config_file_or_obj_or_dict`.

#### Parameters

- **context** (*str*) – The context that is being configured. Must be one of “rsmtool”, “rsmeval”, “rsmcompare”, “rsmsummarize”, “rsmpredict”, or “rsmxval”.
- **config\_file\_or\_obj\_or\_dict** (*Union[str, Configuration, Dict[str, Any], Path]*) – Path to the experiment configuration file either a string or as a `pathlib.Path` object. Users can also pass a `Configuration` object that is in memory or a Python dictionary with keys corresponding to fields in the configuration file. Given a configuration file, any relative paths in the configuration file will be interpreted relative to the location of the file. Given a `Configuration` object, relative paths will be interpreted relative to the `configdir` attribute, that `_must_` be set. Given a dictionary, the reference path is set to the current directory.

#### Returns

**configuration** – The `Configuration` object for the tool.

#### Return type

*Configuration*

#### Raises

- **AttributeError** – If the `configdir` attribute for the `Configuration` input is not set.
- **ValueError** – If `config_file_or_obj_or_dict` contains anything except a string, a path, a dictionary, or a `Configuration` object.

## From container Module

Class to encapsulate data contained in multiple pandas DataFrames.

It represents each of the multiple data sources as a “dataset”. Each dataset is represented by three properties: - “name” : the name of the data set - “frame” : the pandas `DataFrame` that contains the actual data - “path” : the path to the file on disk from which the data was read

#### author

Jeremy Biggs ([jbiggs@ets.org](mailto:jbiggs@ets.org))

#### author

Anastassia Loukina ([aloukina@ets.org](mailto:aloukina@ets.org))

#### author

Nitin Madnani ([nmadnani@ets.org](mailto:nmadnani@ets.org))

#### organization

ETS

**class** `rsmtool.container.DataContainer`(*datasets: List[DatasetDict] | None = None*)

Bases: `object`

Class to encapsulate datasets.

Initialize a `DataContainer` object.

#### Parameters

**datasets** (*Optional[List[DatasetDict]]*) – A list of dataset dictionaries. Each dict should have the following keys: “name” containing the name of the dataset, “frame” containing the dataframe object representing the dataset, and “path” containing the path to the file from which the frame was read.



**add\_dataset**(*dataset\_dict*: DatasetDict, *update*: bool = False) → None

Add a new dataset (or update an existing one).

#### Parameters

- **dataset\_dict** (DatasetDict) – The dataset dictionary to add or update with the “name”, “frame”, and “path” keys.
- **update** (bool) – Update an existing DataFrame, if True. Defaults to False.

#### Return type

None

**copy**(*deep*: bool = True) → DataContainer

Return a copy of the container object.

#### Parameters

- **deep** (bool) – If True, create a deep copy of the underlying data frames. Defaults to True.

#### Returns

**data\_container** – A copy of the input container object.

#### Return type

DataContainer

**drop**(*name*: str) → DataContainer

Drop a given dataset from the container and return instance.

#### Parameters

- **name** (str) – The name of the dataset to drop.

#### Returns

**data\_container** – The input container object with the dataset dropped.

#### Return type

DataContainer

**get\_frame**(*name*: str, *default*: DataFrame | None = None) → DataFrame | None

Get the data frame given the dataset name.

#### Parameters

- **name** (str) – The name for the dataset.
- **default** (Optional[pandas.DataFrame]) – The default value to return if the named dataset does not exist. Defaults to None.

#### Returns

**frame** – The data frame for the named dataset.

#### Return type

Optional[pandas.DataFrame]

**get\_frames**(*prefix*: str | None = None, *suffix*: str | None = None) → Dict[str, DataFrame]

Get all data frames with a given prefix or suffix in their name.

Note that the selection by prefix or suffix is case-insensitive.

#### Parameters

- **prefix** (Optional[str]) – Only return frames with the given prefix. If None, then do not exclude any frames based on their prefix. Defaults to None.

- **suffix** (*Optional[str]*) – Only return frames with the given suffix. If None, then do not exclude any frames based on their suffix. Defaults to None.

#### Returns

**frames** – A dictionary with the data frames that contain the specified prefix and/or suffix in their corresponding names. The names are the keys and the frames are the values.

#### Return type

Dict[str, pandas.DataFrame]

**get\_path**(name: str, default: str | None = None) → str | None

Get the path for the dataset given the name.

#### Parameters

- **name** (str) – The name for the dataset.
- **default** (*Optional[str]*) – The default path to return if the named dataset does not exist. Defaults to None.

#### Returns

**path** – The path for the named dataset.

#### Return type

Optional[str]

**items**() → List[Tuple[str, DataFrame]]

Return the container items as a list of (name, frame) tuples.

#### Returns

**items** – A list of (name, frame) tuples in the container object.

#### Return type

List[Tuple[str, pandas.DataFrame]]

**keys**() → List[str]

Return the container keys (dataset names) as a list.

#### Returns

**keys** – A list of keys (names) in the container object.

#### Return type

List[str]

**rename**(name: str, new\_name: str) → DataContainer

Rename a given dataset in the container and return instance.

#### Parameters

- **name** (str) – The name of the current dataset in the container object.
- **new\_name** (str) – The new name for the dataset in the container object.

#### Returns

**data\_container** – The input container object with the dataset renamed.

#### Return type

DataContainer

**static to\_datasets**(data\_container: DataContainer) → List[DatasetDict]

Convert container object to a list of dataset dictionaries.

Each dictionary will contain the “name”, “frame”, and “path” keys.

#### Parameters

**data\_container** ([DataContainer](#)) – The container object to convert.

#### Returns

**dataset\_dicts** – A list of dataset dictionaries.

#### Return type

List[[DatasetDict](#)]

**values()** → List[[DataFrame](#)]

Return all data frames as a list.

#### Returns

**values** – A list of all data frames in the container object.

#### Return type

List[[pandas.DataFrame](#)]

**class** `rsmttool.container.DatasetDict`

Bases: [TypedDict](#)

Type definition for a dataset dictionary.

**frame:** [DataFrame](#)

**name:** [str](#)

**path:** [Optional\[str\]](#)

### From `convert_feature_json` Module

`rsmttool.convert_feature_json.convert_feature_json_file(json_file: str, output_file: str, delete=False)`  
→ [None](#)

Convert given feature JSON file into tabular format.

The specific format is inferred by the extension of the output file.

#### Parameters

- **json\_file** ([str](#)) – Path to feature JSON file to be converted.
- **output\_file** ([str](#)) – Path to CSV/TSV/XLSX output file.
- **delete** ([bool](#)) – Whether to delete the original file after conversion. Defaults to `False`.

#### Raises

- [RuntimeError](#) – If the given input file is not a valid feature JSON file.
- [RuntimeError](#) – If the output file has an unsupported extension.

#### Return type

[None](#)

## From fairness\_utils Module

```
rsmtool.fairness_utils.get_fairness_analyses(df: DataFrame, group: str, system_score_column: str,
                                             human_score_column: str = 'sc1', base_group: str | None
                                             = None) → Tuple[Dict[str, RegressionResults],
                                             DataContainer]
```

Compute analyses from Loukina et al. 2019.

The function computes how much variance group membership explains in overall score accuracy (osa), overall score difference (osd), and conditional score difference (csd). See the paper for more details.

### Parameters

- **df** (*pandas.DataFrame*) – A dataframe containing columns with numeric human scores, columns with numeric system scores and a column with group membership.
- **group** (*str*) – Name of the column containing group membership.
- **system\_score\_column** (*str*) – Name of the column containing system scores.
- **human\_score\_column** (*str*) – Name of the column containing human scores. Defaults to "sc1".
- **base\_group** (*Optional[str]*) – Name of the group to use as the reference category. If None, the group with the largest number of cases will be used as the reference category. Ties are broken alphabetically. Defaults to None.

### Return type

*Tuple[Dict[str, RegressionResults], DataContainer]*

### Returns

- **model\_dict** (*Dict[str, RegressionResults]*) – A dictionary with different proposed metrics as keys and fitted models as values.
- **fairness\_container** (*DataContainer*) –  
A datacontainer with the following datasets:
  - "estimates\_<METRIC>\_by\_<GROUP>" where <GROUP> corresponds to the given group and <METRIC> can be osa, osd and csd estimates for each group computed by the respective models.
  - "fairness\_metrics\_by\_<GROUP>" - a summary of model fits ( $R^2$  and p values).

## From modeler Module

Class for training and predicting with built-in or SKLL models.

### author

Jeremy Biggs ([jbiggs@ets.org](mailto:jbiggs@ets.org))

### author

Anastassia Loukina ([aloukina@ets.org](mailto:aloukina@ets.org))

### author

Nitin Madnani ([nmadnani@ets.org](mailto:nmadnani@ets.org))

### organization

ETS

```
class rsmtool.modeler.Modeler(logger: Logger | None = None) → None
```

Bases: `object`

Class to train model and generate predictions with built-in or SKLL models.

---

**Note:** The learner and scaling-/trimming-related attributes are set to `None`.

---

Initialize empty `Modeler` object.

**Parameters**

**logger** (*Optional*[`logging.Logger`]) – Logger object to use for logging messages. If `None`, a new logger instance will be created. Defaults to `None`.

```
create_fake_skll_learner(df_coefficients: DataFrame) → Learner
```

Create a fake SKLL linear regression learner from given coefficients.

**Parameters**

**df\_coefficients** (`pandas.DataFrame`) – The data frame containing the linear coefficients we want to create the fake SKLL model with.

**Returns**

**learner** – SKLL `Learner` object representing a `LinearRegression` model with the specified coefficients.

**Return type**

`skll.learner.Learner`

```
get_coefficients() → ndarray | None
```

Get the coefficients of the model, if available.

**Returns**

**coefficients** – The coefficients of the model, if available.

**Return type**

`Optional`[`np.ndarray`]

```
get_feature_names() → List[str] | None
```

Get the feature names, if available.

**Returns**

**feature\_names** – A list of feature names, or `None` if no learner was trained.

**Return type**

`Optional`[`List`[`str`]]

```
get_intercept() → float | None
```

Get the intercept of the model, if available.

**Returns**

**intercept** – The intercept of the model.

**Return type**

`Optional`[`float`]

```
classmethod load_from_file(path: str) → Modeler
```

Load a `Modeler` object from a file on disk.

The file must contain either a `Modeler` or a `SKLL Learner`, in which case a `Modeler` object will be created from the `Learner`.

**Parameters**

**path** (*str*) – File path from which to load the modeler object.

**Returns**

**model** – A Modeler instance.

**Return type**

*Modeler*

**Raises**

**ValueError** – If path does not end with “.model”.

**classmethod load\_from\_learner**(*learner: Learner*) → *Modeler*

Create a new Modeler object with a pre-populated learner.

**Parameters**

**learner** (*skll.learner.Learner*) – A SKLL Learner object.

**Returns**

**modeler** – The newly created Modeler object.

**Return type**

*Modeler*

**Raises**

**TypeError** – If learner is not a SKLL Learner instance.

**static model\_fit\_to\_dataframe**(*fit: RegressionResults*) → *DataFrame*

Extract fit metrics from a statsmodels fit object into a data frame.

**Parameters**

**fit** (*statsmodels.regression.linear\_model.RegressionResults*) – Model fit object obtained from a linear model trained using statsmodels.OLS.

**Returns**

**df\_fit** – The output data frame with the main model fit metrics.

**Return type**

*pandas.DataFrame*

**static ols\_coefficients\_to\_dataframe**(*coefs: Series*) → *DataFrame*

Convert series containing OLS coefficients to a data frame.

**Parameters**

**coefs** (*pandas.Series*) – Series with feature names in the index and the coefficient values as the data, obtained from a linear model trained using statsmodels.OLS.

**Returns**

**df\_coef** – Data frame with two columns: the feature name and the coefficient value.

**Return type**

*pandas.DataFrame*

---

**Note:** The first row in the output data frame is always for the intercept and the rest are sorted by feature name.

---

**predict**(*df: DataFrame, min\_score: float | None = None, max\_score: float | None = None, predict\_expected: bool = False*) → *DataFrame*

Get raw predictions from given SKLL model on data in given data frame.

### Parameters

- **df** (*pandas.DataFrame*) – Data frame containing features on which to make the predictions. The data must contain pre-processed feature values, an ID column named “spkitemid”, and a label column named “sc1”.
- **min\_score** (*Optional[float]*) – Minimum score level to be used if computing expected scores. If None, trying to compute expected scores will raise an exception. Defaults to None.
- **max\_score** (*Optional[float]*) – Maximum score level to be used if computing expected scores. If None, trying to compute expected scores will raise an exception. Defaults to None.
- **predict\_expected** (*bool*) – Predict expected scores for classifiers that return probability distributions over score. This will be ignored with a warning if the specified model does not support probability distributions. Note also that this assumes that the score range consists of contiguous integers - starting at min\_score and ending at max\_score. Defaults to False.

### Returns

**df\_predictions** – Data frame containing the raw predictions, the IDs, and the human scores.

### Return type

*pandas.DataFrame*

### Raises

- **ValueError** – If no model has been trained yet.
- **ValueError** – If the model cannot predict probability distributions and predict\_expected is set to True.
- **ValueError** – If the score range specified by min\_score and max\_score does not match what the model predicts in its probability distribution.
- **ValueError** – If predict\_expected is True but min\_score and max\_score are None.

**predict\_train\_and\_test**(*df\_train: DataFrame, df\_test: DataFrame, configuration: Configuration*) → *Tuple[Configuration, DataContainer]*

Generate raw, scaled, and trimmed predictions on given data.

### Parameters

- **df\_train** (*pandas.DataFrame*) – Data frame containing the pre-processed training set features.
- **df\_test** (*pandas.DataFrame*) – Data frame containing the pre-processed test set features.
- **configuration** (*Configuration*) – A configuration object containing “trim\_max” and “trim\_min” keys.

### Return type

*Tuple[Configuration, DataContainer]*

### Returns

- **configuration** (*Configuration*) – A copy of the given configuration object also containing the “train\_predictions\_mean”, “train\_predictions\_sd”, “human\_labels\_mean”, “human\_labels\_sd”, “trim\_min”, and “trim\_max” parameters.
- **data\_container** (*DataContainer*) – A data container object containing the “pred\_train”, “pred\_test”, and “postprocessing\_params” data sets.

**save**(*model\_path*: *str*) → *None*

Save an instance of this class to disk.

**Parameters**

**model\_path** (*str*) – Destination path for model file

**Return type**

*None*

**scale\_coefficients**(*configuration*: *Configuration*) → *DataContainer*

Scale coefficients using human scores & training set predictions.

This procedure approximates what is done in operational setting but does not apply trimming to predictions.

**Parameters**

**configuration** (*Configuration*) – A configuration object containing the “train\_predictions\_mean”, “train\_predictions\_sd”, and “human\_labels\_sd” parameters.

**Returns**

**data\_container** – A container object containing the “coefficients\_scaled” dataset. The frame for this dataset contains the scaled coefficients and the feature names, along with the intercept.

**Return type**

*DataContainer*

**Raises**

**RuntimeError** – If the model is non-linear and no coefficients are available.

**static skill\_learner\_params\_to\_dataframe**(*learner*: *Learner*) → *DataFrame*

Extract parameters from the given SKLL learner into a data frame.

**Parameters**

**learner** (*skill.learner.Learner*) – A SKLL learner object.

**Returns**

**df\_coef** – The data frame containing the model parameters from the given SKLL Learner object.

**Return type**

*pandas.DataFrame*

---

**Note:**

1. We use the `coef_` attribute of the scikit-learn model underlying the SKLL learner instead of the latter’s `model_params` attribute. This is because `model_params` ignores zero coefficients, which we do not want.
  2. The first row in the output data frame is always for the intercept and the rest are sorted by feature name.
- 

**train**(*configuration*: *Configuration*, *data\_container*: *DataContainer*, *filedir*: *str*, *file\_format*: *str* = ‘csv’)

Train the given model on the given data and save the results.

The main driver function to train the given model on the given data and save the results in the given directories using the given experiment ID as the prefix.

**Parameters**

- **configuration** (*Configuration*) – A configuration object containing “experiment\_id” and “model\_name” parameters.



- **data\_container** (*DataContainer*) – A data container object containing “train\_preprocessed\_features” data set.
- **filedir** (*str*) – Path to the “output” experiment output directory.
- **file\_format** (*str*) – The format in which to save files. One of {"csv", "tsv", "xlsx"}. Defaults to "csv".

**Returns**

**model** – The trained SKLL Learner object.

**Return type**

*skll.learner.Learner*

**train\_builtin\_model**(*model\_name: str, df\_train: DataFrame, experiment\_id: str, fildir: str, file\_format: str = 'csv'*) → *Learner*

Train one of the *built-in linear regression models*.

**Parameters**

- **model\_name** (*str*) – Name of the built-in model to train.
- **df\_train** (*pandas.DataFrame*) – Data frame containing the features on which to train the model. The data frame must contain the ID column named “spkitemid” and the numeric label column named “sc1”.
- **experiment\_id** (*str*) – The experiment ID.
- **fildir** (*str*) – Path to the *output* experiment output directory.
- **file\_format** (*str*) – The format in which to save files. One of {"csv", "tsv", "xlsx"}. Defaults to "csv".

**Returns**

**learner** – SKLL LinearRegression Learner object containing the coefficients learned by training the built-in model.

**Return type**

*skll.learner.Learner*

**train\_equal\_weights\_lr**(*df\_train: DataFrame, feature\_columns: List[str]*) → *Tuple[Learner, RegressionResults, DataFrame, List[str]]*

Train an “EqualWeightsLR” model.

This model assigns the same weight to all features.

**Parameters**

- **df\_train** (*pandas.DataFrame*) – Data frame containing the features on which to train the model.
- **feature\_columns** (*List[str]*) – A list of feature columns to use in training the model.

**Return type**

*Tuple[Learner, RegressionResults, DataFrame, List[str]]*

**Returns**

- **learner** (*skll.learner.Learner*) – The SKLL learner object.
- **fit** (*statsmodels.regression.linear\_model.RegressionResults*) – A statsmodels regression results object.
- **df\_coef** (*pandas.DataFrame*) – Data frame containing the model coefficients.

- **used\_features** (*List[str]*) – A list of features used in the final model.

**train\_lasso\_fixed\_lambda**(*df\_train: DataFrame, feature\_columns: List[str]*) → *Tuple[Learner, None, DataFrame, List[str]]*

Train a “LassoFixedLambda” model.

This is a Lasso model with a fixed lambda.

#### Parameters

- **df\_train** (*pandas.DataFrame*) – Data frame containing the features on which to train the model.
- **feature\_columns** (*List[str]*) – A list of feature columns to use in training the model.

#### Return type

*Tuple[Learner, None, DataFrame, List[str]]*

#### Returns

- **learner** (*skll.learner.Learner*) – The SKLL learner object
- **fit** (*None*) – This is always *None* since there is no OLS model fitted in this case.
- **df\_coef** (*pandas.DataFrame*) – Data frame containing the model coefficients.
- **used\_features** (*List[str]*) – A list of features used in the final model.

**train\_lasso\_fixed\_lambda\_then\_lr**(*df\_train: DataFrame, feature\_columns: List[str]*) → *Tuple[Learner, RegressionResults, DataFrame, List[str]]*

Train a “LassoFixedLambdaThenLR” model.

First do feature selection using lasso regression with a fixed lambda and then use only those features to train a second linear regression

#### Parameters

- **df\_train** (*pandas.DataFrame*) – Data frame containing the features on which to train the model.
- **feature\_columns** (*List[str]*) – A list of feature columns to use in training the model.

#### Return type

*Tuple[Learner, RegressionResults, DataFrame, List[str]]*

#### Returns

- **learner** (*skll.learner.Learner*) – The SKLL learner object
- **fit** (*statsmodels.regression.linear\_model.RegressionResults*) – A statsmodels regression results object.
- **df\_coef** (*pandas.DataFrame*) – The model coefficients in a data\_frame
- **used\_features** (*List[str]*) – A list of features used in the final model.

**train\_lasso\_fixed\_lambda\_then\_non\_negative\_lr**(*df\_train: DataFrame, feature\_columns: List[str]*) → *Tuple[Learner, RegressionResults, DataFrame, List[str]]*

Train an “LassoFixedLambdaThenNNLR” model.

First do feature selection using lasso regression and positive only weights. Then fit an NNLR (see above) on those features.

#### Parameters

- **df\_train** (*pandas.DataFrame*) – Data frame containing the features on which to train the model.
- **feature\_columns** (*List[str]*) – A list of feature columns to use in training the model.

#### Return type

*Tuple[Learner, RegressionResults, DataFrame, List[str]]*

#### Returns

- **learner** (*skll.learner.Learner*) – The SKLL learner object.
- **fit** (*statsmodels.regression.linear\_model.RegressionResults*) – A statsmodels regression results object.
- **df\_coef** (*pandas.DataFrame*) – Data frame containing the model coefficients.
- **used\_features** (*List[str]*) – A list of features used in the final model.

**train\_linear\_regression**(*df\_train: DataFrame, feature\_columns: List[str]*) → *Tuple[Learner, RegressionResults, DataFrame, List[str]]*

Train a “LinearRegression” model.

This model is a simple linear regression model.

#### Parameters

- **df\_train** (*pandas.DataFrame*) – Data frame containing the features on which to train the model.
- **feature\_columns** (*List[str]*) – A list of feature columns to use in training the model.

#### Return type

*Tuple[Learner, RegressionResults, DataFrame, List[str]]*

#### Returns

- **learner** (*skll.learner.Learner*) – The SKLL learner object.
- **fit** (*statsmodels.regression.linear\_model.RegressionResults*) – A statsmodels regression results object.
- **df\_coef** (*pandas.DataFrame*) – Data frame containing the model coefficients.
- **used\_features** (*List[str]*) – A list of features used in the final model.

**train\_non\_negative\_lr**(*df\_train: DataFrame, feature\_columns: List[str]*) → *Tuple[Learner, RegressionResults, DataFrame, List[str]]*

Train an “NNLR” model.

To do this, we first do feature selection using non-negative least squares (NNLS) and then use only its non-zero features to train another linear regression (LR) model. We do the regular LR at the end since we want an LR object so that we have access to R<sup>2</sup> and other useful statistics. There should be no difference between the non-zero coefficients from NNLS and the coefficients that end up coming out of the subsequent LR.

#### Parameters

- **df\_train** (*pandas.DataFrame*) – Data frame containing the features on which to train the model.
- **feature\_columns** (*List[str]*) – A list of feature columns to use in training the model.

#### Return type

*Tuple[Learner, RegressionResults, DataFrame, List[str]]*

### Returns

- **learner** (*skll.learner.Learner*) – The SKLL learner object.
- **fit** (*statsmodels.regression.linear\_model.RegressionResults*) – A statsmodels regression results object.
- **df\_coef** (*pandas.DataFrame*) – Data frame containing the model coefficients.
- **used\_features** (*List[str]*) – A list of features used in the final model.

**train\_non\_negative\_lr\_iterative**(*df\_train: DataFrame, feature\_columns: List[str]*) → *Tuple[Learner, RegressionResults, DataFrame, List[str]]*

Train an “NNLR\_iterative” model.

For applications where there is a concern that standard NNLS may not converge, an alternate method of training NNLR by iteratively fitting OLS models, checking the coefficients, and dropping negative coefficients. First, fit an OLS model. Then, identify any variables whose coefficients are negative. Drop these variables from the model. Finally, refit the model. If any coefficients are still negative, set these to zero.

### Parameters

- **df\_train** (*pandas.DataFrame*) – Data frame containing the features on which to train the model.
- **feature\_columns** (*List[str]*) – A list of feature columns to use in training the model.

### Return type

*Tuple[Learner, RegressionResults, DataFrame, List[str]]*

### Returns

- **learner** (*skll.learner.Learner*) – The SKLL learner object.
- **fit** (*statsmodels.regression.linear\_model.RegressionResults*) – A statsmodels regression results object.
- **df\_coef** (*pandas.DataFrame*) – Data frame containing the model coefficients.
- **used\_features** (*List[str]*) – A list of features used in the final model.

**train\_positive\_lasso\_cv**(*df\_train: DataFrame, feature\_columns: List[str]*) → *Tuple[Learner, None, DataFrame, List[str]]*

Train a “PositiveLassoCV” model.

Do feature selection using lasso regression optimized for log likelihood using cross validation. All coefficients are constrained to have positive values.

### Parameters

- **df\_train** (*pandas.DataFrame*) – Data frame containing the features on which to train the model.
- **feature\_columns** (*List[str]*) – A list of feature columns to use in training the model.

### Return type

*Tuple[Learner, None, DataFrame, List[str]]*

### Returns

- **learner** (*skll.learner.Learner*) – The SKLL learner object
- **fit** (*None*) – This is always *None* since there is no OLS model fitted in this case.
- **df\_coef** (*pandas.DataFrame*) – Data frame containing the model coefficients.

- **used\_features** (*List[str]*) – A list of features used in the final model.

**train\_positive\_lasso\_cv\_then\_lr**(*df\_train: DataFrame, feature\_columns: List[str]*) → *Tuple[Learner, RegressionResults, DataFrame, List[str]]*

Train a “PositiveLassoCVThenLR” model.

First do feature selection using lasso regression optimized for log likelihood using cross validation and then use only those features to train a second linear regression.

#### Parameters

- **df\_train** (*pandas.DataFrame*) – Data frame containing the features on which to train the model.
- **feature\_columns** (*List[str]*) – A list of feature columns to use in training the model.

#### Return type

*Tuple[Learner, RegressionResults, DataFrame, List[str]]*

#### Returns

- **learner** (*skll.learner.Learner*) – The SKLL learner object.
- **fit** (*statsmodels.regression.linear\_model.RegressionResults*) – A statsmodels regression results object.
- **df\_coef** (*pandas.DataFrame*) – Data frame containing the model coefficients.
- **used\_features** (*List[str]*) – A list of features used in the final model.

**train\_rebalanced\_lr**(*df\_train: DataFrame, feature\_columns: List[str]*) → *Tuple[Learner, RegressionResults, DataFrame, List[str]]*

Train a “RebalancedLR” model.

This model balances empirical weights by changing betas (adapted from [here](#)).

#### Parameters

- **df\_train** (*pandas.DataFrame*) – Data frame containing the features on which to train the model.
- **feature\_columns** (*List[str]*) – A list of feature columns to use in training the model.

#### Return type

*Tuple[Learner, RegressionResults, DataFrame, List[str]]*

#### Returns

- **learner** (*skll.learner.Learner*) – The SKLL learner object.
- **fit** (*statsmodels.regression.linear\_model.RegressionResults*) – A statsmodels regression results object.
- **df\_coef** (*pandas.DataFrame*) – Data frame containing the model coefficients.
- **used\_features** (*List[str]*) – A list of features used in the final model.

**train\_score\_weighted\_lr**(*df\_train: DataFrame, feature\_columns: List[str]*) → *Tuple[Learner, RegressionResults, DataFrame, List[str]]*

Train a “ScoreWeightedLR” model.

This is a linear regression model weighted by score.

#### Parameters

- **df\_train** (*pandas.DataFrame*) – Data frame containing the features on which to train the model.
- **feature\_columns** (*List[str]*) – A list of feature columns to use in training the model.

#### Return type

*Tuple[Learner, RegressionResults, DataFrame, List[str]]*

#### Returns

- **learner** (*skll.learner.Learner*) – The SKLL learner object
- **fit** (*statsmodels.regression.linear\_model.RegressionResults*) – A statsmodels regression results object.
- **df\_coef** (*pandas.DataFrame*) – Data frame containing the model coefficients.
- **used\_features** (*List[str]*) – A list of features used in the final model.

**train\_skll\_model** (*model\_name: str, df\_train: DataFrame, custom\_fixed\_parameters: Dict[str, Any] | None = None, custom\_objective: str | None = None, predict\_expected\_scores: bool = False, skll\_grid\_search\_jobs: int = 1*) → *Tuple[Learner, str]*

Train a SKLL classification or regression model.

#### Parameters

- **model\_name** (*str*) – Name of the SKLL model to train.
- **df\_train** (*pandas.DataFrame*) – Data frame containing the features on which to train the model.
- **custom\_fixed\_parameters** (*Optional[Dict[str, Any]]*) – A dictionary containing any fixed parameters for the SKLL model. Defaults to None.
- **custom\_objective** (*Optional[str]*) – Name of custom user-specified objective. If not specified or None, “neg\_mean\_squared\_error” is used as the objective. Defaults to None.
- **predict\_expected\_scores** (*bool*) – Whether we want the trained classifiers to predict expected scores. Defaults to False.
- **skll\_grid\_search\_jobs** (*int*) – Number of folds to run in parallel when using SKLL grid search. Defaults to 1.

#### Return type

*Tuple[Learner, str]*

#### Returns

- **learner** (*skll.learner.Learner*) – A SKLL Learner object of the appropriate type.
- **objective** (*str*) – The chosen tuning objective.

## From preprocessor Module

Classes for preprocessing input data in various contexts.

#### author

Jeremy Biggs ([jbiggs@ets.org](mailto:jbiggs@ets.org))

#### author

Anastassia Loukina ([aloukina@ets.org](mailto:aloukina@ets.org))

#### author

Nitin Madnani ([nmadnani@ets.org](mailto:nmadnani@ets.org))

**organization**  
ETS

**class** `rsmtool.preprocessor.FeaturePreprocessor`(*logger: [Logger](#) | [None](#) = [None](#)*)

Bases: [object](#)

Class to preprocess features in training and testing sets.

Initialize the FeaturePreprocessor object.

**check\_model\_name**(*model\_name: [str](#)*) → [str](#)

Check that the given model name is valid and determine its type.

**Parameters**

**model\_name** ([str](#)) – Name of the model.

**Returns**

**model\_type** – One of “BUILTIN” or “SKLL”.

**Return type**

[str](#)

**Raises**

**ValueError** – If the model is not supported.

**check\_subgroups**(*df: [DataFrame](#), subgroups: [List](#)[[str](#)]*) → [DataFrame](#)

Validate subgroup names in the given data.

Check that all subgroups, if specified, correspond to columns in the provided data frame, and replace all NaNs in subgroup values with ‘No info’ for later convenience.

Raises an exception if any specified subgroup columns are missing.

**Parameters**

- **df** ([pandas.DataFrame](#)) – Input data frame with subgroups to check.
- **subgroups** ([List](#)[[str](#)]) – List of column names that contain grouping information.

**Returns**

**df** – Modified input data frame with NaNs replaced.

**Return type**

[pandas.DataFrame](#)

**Raises**

**KeyError** – If the data does not contain columns for all specified subgroups.

**filter\_data**(*df: [DataFrame](#), label\_column: [str](#), id\_column: [str](#), length\_column: [str](#) | [None](#), second\_human\_score\_column: [str](#) | [None](#), candidate\_column: [str](#), requested\_feature\_names: [List](#)[[str](#)], reserved\_column\_names: [List](#)[[str](#)], given\_trim\_min: [float](#) | [None](#), given\_trim\_max: [float](#) | [None](#), flag\_column\_dict: [Dict](#)[[str](#), [Any](#)], subgroups: [List](#)[[str](#)], exclude\_zero\_scores: [bool](#) = [True](#), exclude\_zero\_sd: [bool](#) = [False](#), feature\_subset\_specs: [DataFrame](#) | [None](#) = [None](#), feature\_subset: [str](#) | [None](#) = [None](#), min\_candidate\_items: [int](#) | [None](#) = [None](#), use\_fake\_labels: [bool](#) = [False](#)) → [Tuple](#)[[DataFrame](#), [DataFrame](#), [DataFrame](#), [DataFrame](#), [DataFrame](#), [DataFrame](#), [DataFrame](#), [float](#), [float](#), [List](#)[[str](#)]]*

Filter rows with zero/non-numeric values for `label_column`.

Check whether any features that are specifically requested in `requested_feature_names` are missing from the data. If no feature names are requested, the feature list is generated based on column names and subset information, if available. The function then excludes non-numeric values for any feature. It will also

exclude zero scores if `exclude_zero_scores` is `True`. If the user requested to exclude candidates with less than `min_candidate_items`, such candidates are also excluded.

It also generates fake labels between 1 and 10 if `use_fake_parameters` is `True`. Finally, it renames the ID and label columns and splits the data into: (a) data frame with feature values and scores (b) data frame with information about subgroup and candidate (metadata) and (c) the data frame with all other columns.

### Parameters

- `df` (*pandas.DataFrame*) – The data frame to filter.
- `label_column` (*str*) – The label column in the data.
- `id_column` (*str*) – The ID column in the data.
- `length_column` (*Optional[str]*) – The length column in the data.
- `second_human_score_column` (*Optional[str]*) – The second human score column in the data.
- `candidate_column` (*str*) – The candidate column in the data.
- `requested_feature_names` (*List[str]*) – A list of requested feature names.
- `reserved_column_names` (*List[str]*) – A list of reserved column names.
- `given_trim_min` (*Optional[float]*) – The minimum trim value.
- `given_trim_max` (*Optional[float]*) – The maximum trim value.
- `flag_column_dict` (*Dict[str, Any]*) – A dictionary of flag columns.
- `subgroups` (*List[str]*) – List containing subgroup names.
- `exclude_zero_scores` (*bool*) – Whether to exclude zero scores. Defaults to `True`.
- `exclude_zero_sd` (*bool*) – Whether to exclude zero standard deviation. Defaults to `False`.
- `feature_subset_specs` (*Optional[pandas.DataFrame]*) – The data frame containing the feature subset specifications. Defaults to `None`.
- `feature_subset` (*Optional[str]*) – The feature subset group (e.g. 'A'). Defaults to `None`.
- `min_candidate_items` (*Optional[int]*) – The minimum number of items needed to include candidate. Defaults to `None`.
- `use_fake_labels` (*bool*) – Whether to use fake labels. Defaults to `False`.

### Return type

`Tuple[DataFrame, DataFrame, DataFrame, DataFrame, DataFrame, DataFrame, DataFrame, float, float, List[str]]`

### Returns

- `df_filtered_features` (*pandas.DataFrame*) – Data frame with filtered features.
- `df_filtered_metadata` (*pandas.DataFrame*) – Data frame with filtered metadata.
- `df_filtered_other_columns` (*pandas.DataFrame*) – Data frame with other columns filtered.
- `df_excluded` (*pandas.DataFrame*) – Data frame with excluded records.
- `df_filtered_length` (*pandas.DataFrame*) – Data frame with length column(s) filtered.
- `df_filtered_human_scores` (*pandas.DataFrame*) – Data frame with human scores filtered.



- **df\_responses\_with\_excluded\_flags** (*pandas.DataFrame*) – Data frame containing responses with excluded flags.
- **trim\_min** (*float*) – The maximum trim value.
- **trim\_max** (*float*) – The minimum trim value.
- **feature\_names** (*List[str]*) – A list of feature names.

**filter\_on\_column**(*df: DataFrame, column: str, exclude\_zeros: bool = False, exclude\_zero\_sd: bool = False*) → *Tuple[DataFrame, DataFrame]*

Filter out rows containing non-numeric values.

Filter out the rows in the given data frame that contain non-numeric (or zero, if specified) values in the specified column. Additionally, it may exclude any columns if they have a standard deviation (*sigma*) of 0.

#### Parameters

- **df** (*pandas.DataFrame*) – The data frame containing the data to be filtered.
- **column** (*str*) – Name of the column from which to filter out values.
- **exclude\_zeros** (*bool*) – Whether to exclude responses containing zeros in the specified column. Defaults to *False*.
- **exclude\_zero\_sd** (*bool*) – Whether to perform the additional filtering step of removing columns that have *sigma* = 0. Defaults to *False*.

#### Return type

*Tuple[DataFrame, DataFrame]*

#### Returns

- **df\_filtered** (*pandas.DataFrame*) – Data frame containing the responses that were *not* filtered out.
- **df\_excluded** (*pandas.DataFrame*) – Data frame containing the non-numeric or zero responses that were filtered out.

---

**Note:** The columns with *sigma* = 0 are removed from both output data frames, assuming *exclude\_zero\_scores* is *True*.

---

**filter\_on\_flag\_columns**(*df: DataFrame, flag\_column\_dict: Dict[str, Any]*) → *Tuple[DataFrame, DataFrame]*

Filter based on specific flag columns.

Check that all *flag\_columns* are present in the given data frame, convert these columns to strings and filter out the values which do not match the condition in *flag\_column\_dict*.

#### Parameters

- **df** (*pandas.DataFrame*) – The *DataFrame* to filter on.
- **flag\_column\_dict** (*Dict[str, Any]*) – Dictionary containing the flag column information.

#### Return type

*Tuple[DataFrame, DataFrame]*

#### Returns

- **df\_responses\_with\_requested\_flags** (*pandas.DataFrame*) – Data frame containing the responses remaining after filtering using the specified flag columns.
- **df\_responses\_with\_excluded\_flags** (*pandas.DataFrame*) – Data frame containing the responses filtered out using the specified flag columns.

#### Raises

- **KeyError** – If the columns listed in the dictionary are not actually present in the data frame.
- **ValueError** – If no responses remain after filtering based on the flag column information.

**generate\_feature\_names**(*df: DataFrame, reserved\_column\_names: List[str], feature\_subset\_specs: DataFrame, feature\_subset: str | None*) → *List[str]*

Generate feature names from column names in data frame.

This method also selects the specified subset of features.

#### Parameters

- **df** (*pandas.DataFrame*) – The data frame from which to generate feature names.
- **reserved\_column\_names** (*List[str]*) – Names of reserved columns.
- **feature\_subset\_specs** (*pandas.DataFrame*) – Feature subset specifications.
- **feature\_subset** (*Optional[str]*) – Feature subset column.

#### Returns

**feature\_names** – List of generated features names.

#### Return type

*List[str]*

**preprocess\_feature**(*values: ndarray, feature\_name: str, feature\_transform: str, feature\_mean: float, feature\_sd: float, exclude\_zero\_sd: bool = False, raise\_error: bool = True, truncations: DataFrame | None = None, truncate\_outliers: bool = True*) → *ndarray*

Remove outliers and transform the values in given numpy array.

Use the given outlier and transformation parameters.

#### Parameters

- **values** (*numpy.ndarray*) – The feature values to preprocess.
- **feature\_name** (*str*) – Name of the feature being pre-processed.
- **feature\_transform** (*str*) – Name of the transformation function to apply.
- **feature\_mean** (*float*) – Mean value to use for outlier detection instead of the mean of the given feature values.
- **feature\_sd** (*float*) – Std. dev. value to use for outlier detection instead of the std. dev. of the given feature values.
- **exclude\_zero\_sd** (*bool*) – Exclude the feature if it has zero standard deviation. Defaults to False.
- **raise\_error** (*bool*) – Raise an error if any of the transformations lead to “inf” values or may change the ranking of feature values. Defaults to True.
- **truncations** (*Optional[pandas.DataFrame]*) – Set of pre-defined truncation values. Defaults to None.
- **truncate\_outliers** (*bool*) – Whether to truncate outlier values. Defaults to True.

**Returns**

**transformed\_feature** – Numpy array containing the transformed and clamped feature values.

**Return type**

`numpy.ndarray`

**Raises**

**ValueError** – If the preprocessed feature values have zero standard deviation and `exclude_zero_sd` is set to `True`.

**preprocess\_features**(*df\_train*: `DataFrame`, *df\_test*: `DataFrame`, *df\_feature\_specs*: `DataFrame`, *standardize\_features*: `bool = True`, *use\_truncations*: `bool = False`, *truncate\_outliers*: `bool = True`) → `Tuple[DataFrame, DataFrame, DataFrame]`

Preprocess features in given data using corresponding specifications.

Preprocess the feature values in the training and testing data frames whose specifications are contained in `df_feature_specs`. Also returns a third data frame containing the feature specifications and other information.

**Parameters**

- **df\_train** (`pandas.DataFrame`) – Data frame containing the raw feature values for the training set.
- **df\_test** (`pandas.DataFrame`) – Data frame containing the raw feature values for the test set.
- **df\_feature\_specs** (`pandas.DataFrame`) – Data frame containing the various specifications from the feature file.
- **standardize\_features** (`bool`) – Whether to standardize the features. Defaults to `True`.
- **truncate\_outliers** (`bool`) – Truncate outlier values if set in the configuration file. Defaults to `True`.
- **use\_truncations** (`bool`) – Whether we should use the truncation set for removing outliers. Defaults to `False`.

**Return type**

`Tuple[DataFrame, DataFrame, DataFrame]`

**Returns**

- **df\_train\_preprocessed** (`pandas.DataFrame`) – Data frame with preprocessed training data.
- **df\_test\_preprocessed** (`pandas.DataFrame`) – Data frame with preprocessed test data.
- **df\_feature\_info** (`pandas.DataFrame`) – Data frame with feature information.

**preprocess\_new\_data**(*df\_input*: `DataFrame`, *df\_feature\_info*: `DataFrame`, *standardize\_features*: `bool = True`, *truncate\_outliers*: `bool = True`) → `Tuple[DataFrame, DataFrame]`

Preprocess feature values using the parameters in `df_feature_info`.

For more details on what these preprocessing parameters are, see [documentation](#).

**Parameters**

- **df\_input** (`pandas.DataFrame`) – Data frame with raw feature values that will be used to generate the scores. Each feature is stored in a separate column. Each row corresponds to one response. There should also be a column named “spkitemid” containing a unique ID for each response.

- **df\_feature\_info** (*pandas.DataFrame*) – Data frame with preprocessing parameters in the following columns:
  - "feature" : the name of the feature; should match the feature names in `df_input`.
  - "sign" : 1 or -1. Indicates whether the feature value needs to be multiplied by -1.
  - "transform" : *transformation* that needs to be applied to this feature.
  - "train\_mean", "train\_sd" : mean and standard deviation for outlier truncation.
  - "train\_transformed\_mean", "train\_transformed\_sd" : mean and standard deviation for computing z-scores.
- **standardize\_features** (*bool*) – Whether the features should be standardized prior to prediction. Defaults to True.
- **truncate\_outliers** (*bool*) – Whether outlier should be truncated prior to prediction. Defaults to True.

#### Return type

`Tuple[DataFrame, DataFrame]`

#### Returns

- **df\_features\_preprocessed** (*pandas.DataFrame*) – Data frame with processed feature values.
- **df\_excluded** (*pandas.DataFrame*) – Data frame with responses excluded from further analysis due to non-numeric feature values in the original file or after applying transformations. This data frame always contains the original feature values.

#### Raises

- **KeyError** – if some of the features specified in `df_feature_info` are not present in `df_input`.
- **ValueError** – If all responses have at least one non-numeric feature value and, therefore, no score can be generated for any of the responses.

**process\_data**(*config\_obj*: *Configuration*, *data\_container\_obj*: *DataContainer*, *context*: *str* = 'rsmtool') → `Tuple[Configuration, DataContainer]`

Process and setup the data for an experiment in the given context.

#### Parameters

- **config\_obj** (*Configuration*) – The configuration object.
- **data\_container\_obj** (*DataContainer*) – The data container object.
- **context** (*str*) – The tool context: one of {"rsmtool", "rsmeval", "rsmpredict"}. Defaults to "rsmtool".

#### Return type

`Tuple[Configuration, DataContainer]`

#### Returns

- **config\_obj** (*Configuration*) – New configuration object containing the updated configuration.
- **data\_container** (*DataContainer*) – new data container object containing the preprocessed data.

#### Raises

**ValueError** – If the context is not one of {"rsmtool", "rsmeval", "rsmpredict"}.

**process\_data\_rsmeval**(*config\_obj*: *Configuration*, *data\_container\_obj*: *DataContainer*) →  
*Tuple[Configuration, DataContainer]*

Set up rsmeval experiment by loading & preprocessing evaluation data.

This function takes a configuration object and a container object as input and returns the same types of objects as output after the loading, normalizing, and preprocessing.

#### Parameters

- **config\_obj** (*Configuration*) – The configuration object.
- **data\_container\_obj** (*DataContainer*) – The data container object.

#### Return type

*Tuple[Configuration, DataContainer]*

#### Returns

- **config\_obj** (*Configuration*) – New configuration object containing the updated rsmeval configuration.
- **data\_container** (*DataContainer*) – New data container object containing the preprocessed data.

#### Raises

- **KeyError** – If columns specified in the configuration do not exist in the predictions file.
- **ValueError** – If the columns containing the human scores and the system scores in the predictions file have the same name.
- **ValueError** – If the columns containing the first set of human scores and the second set of human scores in the predictions file have the same name.
- **ValueError** – If the predictions file contains the same response ID more than once.
- **ValueError** – No responses were left after filtering out zero or non-numeric values for the various columns.

**process\_data\_rsmexplain**(*config\_obj*: *Configuration*, *data\_container\_obj*: *DataContainer*) →  
*Tuple[Configuration, DataContainer]*

Process data for rsmexplain experiments.

This function takes a configuration object and a container object as input and returns the same types of objects as output after the loading, normalizing, and preprocessing.

#### Parameters

- **config\_obj** (*Configuration*) – The configuration object.
- **data\_container\_obj** (*DataContainer*) – The data container object.

#### Return type

*Tuple[Configuration, DataContainer]*

#### Returns

- **config\_obj** (*Configuration*) – New configuration object containing the updated rsmexplain configuration.
- **data\_congtainer** (*DataContainer*) – New data container object containing the preprocessed data.

#### Raises

- **ValueError** – If data contains duplicate response IDs.

**process\_data\_rsmpredict**(*config\_obj*: Configuration, *data\_container\_obj*: DataContainer) → Tuple[Configuration, DataContainer]

Process data for rsmpredict experiments.

This function takes a configuration object and a container object as input and returns the same types of objects as output after the loading, normalizing, and preprocessing.

#### Parameters

- **config\_obj** (Configuration) – The configuration object.
- **data\_container\_obj** (DataContainer) – The data container object.

#### Return type

Tuple[Configuration, DataContainer]

#### Returns

- **config\_obj** (Configuration) – New configuration object containing the updated rsmpredict configuration.
- **data\_container** (DataContainer) – New data container object containing the preprocessed data.

#### Raises

- **KeyError** – If columns specified in the configuration do not exist in the data.
- **ValueError** – If data contains duplicate response IDs.

**process\_data\_rsmtool**(*config\_obj*: Configuration, *data\_container\_obj*: DataContainer) → Tuple[Configuration, DataContainer]

Set up rsmtool experiment by loading & preprocessing train/test data.

This function takes a configuration object and a container object as input and returns the same types of objects as output after the loading, normalizing, and preprocessing.

#### Parameters

- **config\_obj** (Configuration) – The configuration object.
- **data\_container\_obj** (DataContainer) – The data container object.

#### Return type

Tuple[Configuration, DataContainer]

#### Returns

- **config\_obj** (Configuration) – New configuration object containing the updated rsmtool configuration.
- **data\_container** (DataContainer) – New data container object containing the preprocessed data.

#### Raises

- **ValueError** – If columns specified in the configuration do not exist in the data.
- **ValueError** – If the test label column and second human score columns have the same name.
- **ValueError** – If the length column is requested as a feature.
- **ValueError** – If the second human score column is requested as a feature.

- **ValueError** – If “use\_truncations” was specified in the configuration, but no feature CSV file was found.

**process\_predictions**(*df\_test\_predictions: DataFrame, train\_predictions\_mean: float, train\_predictions\_sd: float, human\_labels\_mean: float, human\_labels\_sd: float, trim\_min: float, trim\_max: float, trim\_tolerance: float = 0.4998*) → *DataFrame*

Process predictions to create scaled, trimmed and rounded predictions.

#### Parameters

- **df\_test\_predictions** (*pandas.DataFrame*) – Data frame containing the test set predictions.
- **train\_predictions\_mean** (*float*) – The mean of the predictions on the training set.
- **train\_predictions\_sd** (*float*) – The std. dev. of the predictions on the training set.
- **human\_labels\_mean** (*float*) – The mean of the human scores used to train the model.
- **human\_labels\_sd** (*float*) – The std. dev. of the human scores used to train the model.
- **trim\_min** (*float*) – The lowest score on the score point, used for trimming the raw regression predictions.
- **trim\_max** (*float*) – The highest score on the score point, used for trimming the raw regression predictions.
- **trim\_tolerance** (*float*) – Tolerance to be added to trim\_max and subtracted from trim\_min. Defaults to 0.4998.

#### Returns

**df\_pred\_processed** – Data frame containing the various trimmed and rounded predictions.

#### Return type

*pandas.DataFrame*

**static remove\_outliers**(*values: ndarray, mean: float | None = None, sd: float | None = None, sd\_multiplier: int = 4*) → *ndarray*

Remove outliers from given array of values by clamping them.

Clamp any given values that are  $\pm$  *sd\_multiplier* (*m*) standard deviations (*sigma*) away from the mean (*mu*). Use given *mean* and *sd* instead of computing *sigma* and *mu*, if specified. The values are clamped to the interval:

$$\left[ \begin{array}{l} \mu - m * \\ \sigma, \\ \mu + m * \\ \sigma \end{array} \right]$$

#### Parameters

- **values** (*numpy.ndarray*) – The values from which to remove outliers, usually corresponding to a given feature.
- **mean** (*Optional[float]*) – Use the given mean value when computing outliers instead of the mean from the data. Defaults to *None*.
- **sd** (*Optional[float]*) – Use the given std. dev. value when computing outliers instead of the std. dev. from the data. Defaults to *None*.

- **sd\_multiplier** (*int*) – Use the given multiplier for the std. dev. when computing the outliers. Defaults to 4.

#### Returns

**new\_values** – Numpy array with the outliers clamped.

#### Return type

*numpy.ndarray*

**remove\_outliers\_using\_truncations**(*values: ndarray, feature\_name: str, truncations: DataFrame*) → *ndarray*

Remove outliers using pre-specified truncation groups.

This is different from `remove_outliers()` which calculates the outliers based on the training set rather than looking up the truncation values from a pre-specified data frame.

#### Parameters

- **values** (*numpy.ndarray*) – The values from which to remove outliers, usually corresponding to a given feature.
- **feature\_name** (*str*) – Name of the feature whose outliers are being clamped.
- **truncations** (*pandas.DataFrame*) – A data frame with truncation values. The features should be set as the index.

#### Returns

**new\_values** – Numpy array with the outliers clamped.

#### Return type

*numpy.ndarray*

**rename\_default\_columns**(*df: DataFrame, requested\_feature\_names: List[str], id\_column: str, first\_human\_score\_column: str | None, second\_human\_score\_column: str | None, length\_column: str | None, system\_score\_column: str | None, candidate\_column: str | None*) → *DataFrame*

Standardize column names and rename columns with reserved column names.

RSMTool reserves some column names for internal use, e.g., “sc1”, “spkitemid” etc. If the given data already contains columns with these names, then they must be renamed to prevent conflict. This method renames such columns to “##NAME##”, e.g., an existing column named “sc1” will be renamed to “##sc1##”.

#### Parameters

- **df** (*pandas.DataFrame*) – The data frame containing the columns to rename.
- **requested\_feature\_names** (*List[str]*) – List of feature column names that we want to include in the scoring model.
- **id\_column** (*str*) – Column name containing the response IDs.
- **first\_human\_score\_column** (*Union[str, None]*) – Column name containing the H1 scores. Should be None if no H1 scores are available.
- **second\_human\_score\_column** (*Union[str, None]*) – Column name containing the H2 scores. Should be None if no H2 scores are available.
- **length\_column** (*Union[str, None]*) – Column name containing response lengths. Should be None if lengths are not available.
- **system\_score\_column** (*Union[str, None]*) – Column name containing the score predicted by the system. This is only used for `rsmeval`.



- **candidate\_column** (*Union[str, None]*) – Column name containing identifying information at the candidate level. Should be *None* if such information is not available.

#### Returns

**df** – Modified input data frame with all the approximate re-namings.

#### Return type

*pandas.DataFrame*

**select\_candidates**(*df: DataFrame, N: int, candidate\_col: str = 'candidate'*) → *Tuple[DataFrame, DataFrame]*

Select candidates which have responses to N or more items.

#### Parameters

- **df** (*pandas.DataFrame*) – The data frame from which to select candidates with N or more items.
- **N** (*int*) – Minimal number of items per candidate
- **candidate\_col** (*str*) – Name of the column which contains candidate ids. Defaults to “candidate”.

#### Return type

*Tuple[DataFrame, DataFrame]*

#### Returns

- **df\_included** (*pandas.DataFrame*) – Data frame with responses from candidates with responses to N or more items.
- **df\_excluded** (*pandas.DataFrame*) – Data frame with responses from candidates with responses to less than N items.

**trim**(*values: List[float] | ndarray, trim\_min: float, trim\_max: float, tolerance: float = 0.4998*) → *ndarray*

Trim values in given numpy array.

The trimming uses *trim\_min* - *tolerance* as the floor and *trim\_max* + *tolerance* as the ceiling.

#### Parameters

- **values** (*Union[List[float], numpy.ndarray]*) – The values to trim.
- **trim\_min** (*float*) – The lowest score on the score point, used for trimming the raw regression predictions.
- **trim\_max** (*float*) – The highest score on the score point, used for trimming the raw regression predictions.
- **tolerance** (*float*) – The tolerance that will be used to compute the trim interval. Defaults to 0.4998.

#### Returns

**trimmed\_values** – Numpy array containing the trimmed values.”

#### Return type

*numpy.ndarray*

#### Raises

**ValueError** – If *trim\_min*, *trim\_max*, or *tolerance* are *None*.

**class** `rsmttool.preprocessor.FeatureSpecsProcessor`(*logger: [Logger](#) | [None](#) = [None](#)*)

Bases: [object](#)

Encapsulate feature file processing methods.

Initialize the FeatureSpecsProcessor object.

**find\_feature\_sign**(*feature: [str](#), sign\_dict: [Dict](#)[[str](#), [str](#)]*) → [float](#)

Get the feature sign from the feature CSV file.

**Parameters**

- **feature** ([str](#)) – The name of the feature.
- **sign\_dict** ([dict](#)) – A dictionary of feature signs.

**Returns**

**feature\_sign\_numeric** – The signed feature.

**Return type**

[float](#)

**generate\_default\_specs**(*feature\_names: [List](#)[[str](#)]*) → [DataFrame](#)

Generate default feature “specifications” for given feature names.

The specifications are stored as a data frame with three columns “feature”, “transform”, and “sign”.

**Parameters**

**feature\_names** ([List](#)[[str](#)]) – List of feature names for which to generate specifications.

**Returns**

**feature\_specs** – A dataframe with feature specifications that can be saved as a *feature list file*.

**Return type**

[pandas.DataFrame](#)

---

**Note:** Since these are default specifications, the values for the “transform” column for each feature will be “raw” and the value for the “sign” column will be 1.

---

**generate\_specs**(*df: [DataFrame](#), feature\_names: [List](#)[[str](#)], train\_label: [str](#), feature\_subset: [DataFrame](#) | [None](#) = [None](#), feature\_sign: [int](#) | [None](#) = [None](#)*) → [DataFrame](#)

Generate feature specifications using the feature CSV file.

Compute the specifications for “sign” and the correlation with score to identify the best transformation.

**Parameters**

- **df** ([pandas.DataFrame](#)) – The input data frame from which to generate the specifications.
- **feature\_names** ([List](#)[[str](#)]) – A list of feature names.
- **train\_label** ([str](#)) – The label column for the training data
- **feature\_subset** ([Optional](#)[[pandas.DataFrame](#)]) – A data frame containing the feature subset specifications, if any. Defaults to [None](#).
- **feature\_sign** ([Optional](#)[[int](#)]) – The sign of the feature. Defaults to [None](#).

**Returns**

**df\_feature\_specs** – The output data frame containing the feature specifications.

**Return type**

[pandas.DataFrame](#)

**validate\_feature\_specs**(*df*: *DataFrame*, *use\_truncations*: *bool* = *False*) → *DataFrame*

Validate given feature specifications.

Check given feature specifications to make sure that there are no duplicate feature names and that all columns are in the right format. Add the default values for “transform” and “sign” if none are given.

#### Parameters

- **df** (*pandas.DataFrame*) – The feature specification DataFrame to validate.
- **use\_truncations** (*bool*) – Whether to use truncation values. If this is *True* and truncation values are not specified, an exception is raised. Defaults to *False*.

#### Returns

**df\_specs\_new** – The output data frame with normalized values.

#### Return type

*pandas.DataFrame*

#### Raises

- **KeyError** – If the input data frame does not have a “feature” column.
- **ValueError** – If there are duplicate values in the “feature” column.
- **ValueError** – if the “sign” column contains invalid values.
- **ValueError** – If *use\_truncations* is set to *True*, and no “min” and “max” columns exist in the data set.

**class** *rsmttool.preprocessor.FeatureSubsetProcessor*(*logger*: *Logger* | *None* = *None*)

Bases: *object*

Class to encapsulate feature sub-setting methods.

Initialize the FeatureSubsetProcessor object.

**check\_feature\_subset\_file**(*df*: *DataFrame*, *subset*: *str* | *None* = *None*, *sign*: *str* | *None* = *None*) → *None*

Check that feature subset file is complete and in the correct format.

Raises an exception if it finds any errors but otherwise returns nothing.

#### Parameters

- **df** (*pandas.DataFrame*) – The data frame containing the feature subset file.
- **subset** (*Optional[str]*) – Name of a pre-defined feature subset. Defaults to *None*.
- **sign** (*Optional[str]*) – Value of the sign. Defaults to *None*.

#### Raises

- **ValueError** – If any columns are missing from the subset file.
- **ValueError** – If any of the columns contain invalid values.

#### Return type

*None*

**select\_by\_subset**(*feature\_columns*: *List[str]*, *feature\_subset\_specs*: *DataFrame*, *subset*: *str*) → *List[str]*

Select feature columns using feature subset specifications.

#### Parameters

- **feature\_columns** (*List[str]*) – A list of feature columns

- **feature\_subset\_specs** (*pandas.DataFrame*) – The feature subset specification data frame.
- **subset** (*str*) – The column to subset.

#### Returns

**feature\_names** – A list of feature names to include.

#### Return type

List[*str*]

### From `prmse` Module

`rsmttool.utils.prmse.prmse_true(system: ndarray, human_scores: ndarray, variance_errors_human: float | None = None) → float | None`

Compute PRMSE when predicting true score from system scores.

PRMSE = Proportional Reduction in Mean Squared Error. The formula to compute PRMSE implemented in RSMTTool was derived at ETS by Matthew S. Johnson. See [Loukina et al. \(2020\)](#) for further information about PRMSE.

#### Parameters

- **system** (*numpy.ndarray*) – System scores for each response of shape (n\_samples,).
- **human\_scores** (*numpy.ndarray*) – Human ratings for each response of shape (n\_samples, n\_ratings).
- **variance\_errors\_human** (*Optional[float]*) – Estimated variance of errors in human scores. If *None*, the variance will be estimated from the data. In this case at least some responses *must* have more than one human score. Defaults to *None*.

#### Returns

**prmse** – Proportional reduction in mean squared error. If the variance of errors in human scores is not available and cannot be estimated from the data, returns *None*.

#### Return type

*Optional[float]*

#### Raises

**ValueError** – If variance of true scores or MSE could not be computed.

`rsmttool.utils.prmse.variance_of_errors(human_scores: ndarray) → float | None`

Estimate the variance of errors in human scores.

#### Parameters

**human\_scores** (*numpy.ndarray*) – Human ratings for each response of shape (n\_samples, n\_ratings).

#### Returns

**variance\_of\_errors** – Estimated variance of errors in human scores. If the variance of errors cannot be estimated from the data, returns *None*.

#### Return type

*Optional[float]*

## From reader Module

Classes for reading data files (or dictionaries) into DataContainer objects.

**author**

Jeremy Biggs (jbiggs@ets.org)

**author**

Anastassia Loukina (aloukina@ets.org)

**author**

Nitin Madnani (nmadnani@ets.org)

**organization**

ETS

```
class rsmtool.reader.DataReader(filepaths: List[str], framenames: List[str], file_converters: Dict[str,
Dict[str, Any]] | None = None)
```

Bases: `object`

Class to generate DataContainer objects.

Initialize a DataReader object.

**Parameters**

- **filepaths** (*List[str]*) – A list of paths to files that are to be read in. Some of the paths can be empty strings.
- **framenames** (*List[str]*) – A list of names for the data sets to be included in the container.
- **file\_converters** (*Optional[Dict[str, Dict[str, Any]]]*) – A dictionary of file converter dicts. The keys are the data set names and the values are the converter dictionaries to be applied to the corresponding data set. Defaults to None.

**Raises**

- **AssertionError** – If `len(filepaths)` does not equal `len(framenames)`.
- **ValueError** – If `file_converters` is not a dictionary or if any of its values is not a dictionary.
- **NameError** – If a key in `file_converters` does not exist in `framenames`.
- **ValueError** – If any of the specified file paths is None.

```
static locate_files(filepaths: str | List[str], configdir: str) → List[str]
```

Locate an experiment file, or a list of experiment files.

If the given path doesn't exist, then maybe the path is relative to the path of the config file. If neither exists, then return an empty string.

**Parameters**

- **filepaths** (*Union[str, List[str]]*) – Name(s) of the experiment file we want to locate.
- **configdir** (*str*) – Path to the reference configuration directory (usually the directory of the config file)

**Returns**

**retval** – List of absolute paths to the located files. If a file does not exist, the corresponding element in the list is an empty string.

**Return type**

List[str]

**Raises**
**ValueError** – If filepaths is not a string or a list.

`read(kwarg_dict: Dict[str, Dict[str, Any]] | None = None) → DataContainer`

Read all files contained in self.dataset\_paths.

**Parameters**
**kwarg\_dict** (Optional[Dict[str, Dict[str, Any]]]) – Dictionary with the names of the datasets as keys and dictionaries of keyword arguments to pass to the pandas reader for each dataset as values. The keys in those dictionaries are the names of the keyword arguments and the values are the values of the keyword arguments. Defaults to None.

**Returns**
**datacontainer** – A data container object.

**Return type**

DataContainer

**Raises**
**FileNotFoundError** – If any of the files in self.dataset\_paths does not exist.

`static read_from_file(filename: str, converters: Dict[str, Any] | None = None, **kwargs)`

Read a CSV/TSV/XLSX/JSONLINES/SAS7BDAT file and return a data frame.

**Parameters**

- **filename** (str) – Name of file to read.
- **converters** (Optional[Dict[str, Any]]) – A dictionary specifying how the types of the columns in the file should be converted. Specified in the same format as for `pandas.read_csv()`. Defaults to None.

**Returns**
**df** – Data frame containing the data in the given file.

**Return type**

pandas.DataFrame

**Raises**

- **ValueError** – If the file has an unsupported extension.
- **pandas.errors.ParserError** – If the file is badly formatted or corrupt.

---

**Note:** Any additional keyword arguments are passed to the underlying pandas IO reader function.

---

`rsmttool.reader.read_jsonlines(filename: str, converters: Dict[str, Any] | None = None) → DataFrame`

Read a data file in .jsonlines format into a data frame.

Normalize nested jsons with up to one level of nesting.

**Parameters**

- **filename** (str) – Name of file to read.
- **converters** (Optional[Dict[str, Any]]) – A dictionary specifying how the types of the columns in the file should be converted. Specified in the same format as for `pandas.read_csv()`. Defaults to None.

### Returns

**df** – Data frame containing the data in the given file.

### Return type

`pandas.DataFrame`

`rsmtool.reader.try_to_load_file(filename: str, converters: Dict[str, Any] | None = None, raise_error: bool = False, raise_warning: bool = False, **kwargs) → DataFrame | None`

Read a single file, if it exists.

Optionally raises an error or warning if the file cannot be found. Otherwise, returns *None*.

### Parameters

- **filename** (*str*) – Name of file to read.
- **converters** (*Optional[Dict[str, Any]]*) – A dictionary specifying how the types of the columns in the file should be converted. Specified in the same format as for `pandas.read_csv()`. Defaults to *None*.
- **raise\_error** (*bool*) – Raise an error if the file cannot be located. Defaults to *False*.
- **raise\_warning** (*bool*) – Raise a warning if the file cannot be located. Defaults to *False*.

### Returns

**df** – *DataFrame* containing the data in the given file, or *None* if the file does not exist.

### Return type

*Optional[pandas.DataFrame]*

### Raises

**FileNotFoundError** – If `raise_error` is *True* and the file cannot be located.

## From reporter Module

Classes for dealing with report generation.

### author

Jeremy Biggs ([jbiggs@ets.org](mailto:jbiggs@ets.org))

### author

Anastassia Loukina ([aloukina@ets.org](mailto:aloukina@ets.org))

### author

Nitin Madnani ([nmadnani@ets.org](mailto:nmadnani@ets.org))

### organization

ETS

`class rsmtool.reporter.Reporter(logger: Logger | None = None, wandb_run: Run | RunDisabled | None = None)`

Bases: `object`

Class to generate Jupyter notebook reports and convert them to HTML.

Initialize the Reporter object.

### Parameters

- **logger** (*Optional[logging.Logger]*) – A *Logger* object. If *None* is passed, get logger from `__name__`. Defaults to *None*.

- **wandb\_run** (*Union*[*wandb.wandb\_run.Run*, *wandb.sdk.lib.RunDisabled*, *None*]) – A wandb run object that will be used to log artifacts and tables. If *None* is passed, a new wandb run will be initialized if wandb is enabled in the configuration. Defaults to *None*.

**static check\_section\_names**(*specified\_sections: List[str]*, *section\_type: str*, *context: str = 'rsmtool'*) → *None*

Validate the specified section names.

This function checks whether the specified section names are valid and raises an exception if they are not.

#### Parameters

- **specified\_sections** (*List[str]*) – List of report section names.
- **section\_type** (*str*) – One of “general” or “special”.
- **context** (*str*) – Context in which we are validating the section names. One of {"rsmtool", "rsmeval", "rsmcompare"}. Defaults to "rsmtool".

#### Raises

**ValueError** – If any of the section names of the given type are not valid in the context of the given tool.

#### Return type

*None*

**static check\_section\_order**(*chosen\_sections: List[str]*, *section\_order: List[str]*) → *None*

Check the order of the specified sections.

#### Parameters

- **chosen\_sections** (*List[str]*) – List of chosen section names.
- **section\_order** (*List[str]*) – An ordered list of the chosen section names.

#### Raises

**ValueError** – If any sections specified in the order are missing from the list of chosen sections or vice versa.

#### Return type

*None*

**static convert\_ipynb\_to\_html**(*notebook\_file: str*, *html\_file: str*)

Convert given Jupyter notebook (.ipynb) to HTML file.

#### Parameters

- **notebook\_file** (*str*) – Path to input Jupyter notebook file.
- **html\_file** (*str*) – Path to output HTML file.

---

**Note:** This function is also exposed as the *render\_notebook* command-line utility.

---

**create\_comparison\_report**(*config, csvdir\_old, figdir\_old, csvdir\_new, figdir\_new, output\_dir*)

Generate an HTML report for comparing two rsmtool experiments.

#### Parameters

- **config** (*configuration\_parser.Configuration*) – A configuration object
- **csvdir\_old** (*str*) – The old experiment CSV output directory.



- **figdir\_old** (*str*) – The old figure output directory
- **csvdir\_new** (*str*) – The new experiment CSV output directory.
- **figdir\_new** (*str*) – The old figure output directory
- **output\_dir** (*str*) – The output dir for the new report.

**create\_explanation\_report**(*config*, *csv\_dir*, *output\_dir*)

Generate a html report for rsmexplain.

#### Parameters

- **config** (*configuration\_parser.Configuration*) – A configuration object
- **csv\_dir** (*str*) – The experiment output directory containing CSV files with SHAP values
- **output\_dir** (*str*) – The directory for the html report

**create\_report**(*config*: *Configuration*, *csvdir*: *str*, *figdir*: *str*, *context*: *str* = 'rsmttool') → *None*

Generate HTML report for an rsmttool/rsmeval experiment.

#### Parameters

- **config** (*Configuration*) – A configuration object
- **csvdir** (*str*) – The CSV output directory.
- **figdir** (*str*) – The figure output directory
- **context** (*str*) – Context of the tool in which we are validating. One of {"rsmttool", "rsmeval"}. Defaults to "rsmttool".

#### Raises

**KeyError** – If the *test\_file\_location* or *pred\_file\_location* fields are not specified in the given configuration.

#### Return type

*None*

**create\_summary\_report**(*config*, *all\_experiments*, *csvdir*)

Generate an HTML report for summarizing the given rsmttool experiments.

#### Parameters

- **config** (*configuration\_parser.Configuration*) – A configuration object
- **all\_experiments** (*list of str*) – A list of experiment configuration files to summarize.
- **csvdir** (*str*) – The experiment CSV output directory.

**determine\_chosen\_sections**(*general\_sections*: *List[str]*, *custom\_sections*: *List[str]*, *subgroups*: *List[str]*, *context*: *str* = 'rsmttool') → *List[str]*

Compile a combined list of section names to be included in the report.

#### Parameters

- **general\_sections** (*List[str]*) – List of specified general section names.
- **custom\_sections** (*List[str]*) – List of specified custom sections, if any.
- **subgroups** (*List[str]*) – List of column names that contain grouping information.
- **context** (*str*) – Context of the tool in which we are validating. One of {"rsmttool", "rsmeval", "rsmcompare"} Defaults to "rsmttool".

#### Returns

**chosen\_sections** – Final list of chosen sections that are to be included in the HTML report.

#### Return type

List[str]

#### Raises

**ValueError** – If a subgroup report section is requested but no subgroups were specified in the configuration file.

**get\_ordered\_notebook\_files**(*general\_sections*: List[str], *custom\_sections*: List[str] = [], *section\_order*: List[str] | None = None, *subgroups*: List[str] = [], *model\_type*: str | None = None, *context*: str = 'rsmtool') → List[str]

Check all section names and the order of the sections.

Combine all section names with the appropriate file mapping, and generate an ordered list of notebook files that are needed to generate the final report.

#### Parameters

- **general\_sections** (List[str]) – List of specified general sections.
- **custom\_sections** (List[str]) – List of specified custom sections, if any. Defaults to [].
- **section\_order** (Optional[List[str]]) – Ordered list in which the user wants the specified sections. Defaults to None.
- **subgroups** (List[str]) – List of column names that contain grouping information. Defaults to [].
- **model\_type** (Optional[List[str]]) – Type of the model. Possible values are {"BUILTIN", "SKLL", None.}. We allow None here so that rsmeval can use the same function. Defaults to None.
- **context** (str) – Context of the tool in which we are validating. One of {"rsmtool", "rsmeval", "rsmcompare"}. Defaults to "rsmtool".

#### Returns

**chosen\_notebook\_files** – List of the IPython notebook files that have to be rendered into the HTML report.

#### Return type

List[str]

**get\_section\_file\_map**(*custom\_sections*: List[str], *model\_type*: str | None = None, *context*: str = 'rsmtool') → Dict[str, str]

Map section names to IPython notebook filenames.

#### Parameters

- **custom\_sections** (List[str]) – List of custom sections.
- **model\_type** (Optional[str]) – Type of the model. One of {"BUILTIN", "SKLL", None.}. We allow None here so that rsmeval can use the same function. Defaults to None.
- **context** (str) – Context of the tool in which we are validating. One of {"rsmtool", "rsmeval", "rsmcompare"}. Defaults to "rsmtool".

#### Returns

**section\_file\_map** – Dictionary mapping each section name to the corresponding IPython notebook filename.

**Return type**

dict

**static** `locate_custom_sections`(*custom\_report\_section\_paths*: List[str], *configdir*: str) → List[str]

Locate custom report section files.

Get the absolute paths for custom report sections and check that the files exist. If a file does not exist, raise an exception.

**Parameters**

- **custom\_report\_section\_paths** (List[str]) – List of paths to IPython notebook files representing the custom sections.
- **configdir** (str) – Path to the experiment configuration directory.

**Returns**

**custom\_report\_sections** – List of absolute paths to the custom section notebooks.

**Return type**

List[str]

**Raises**

**FileNotFoundError** – If any of the files cannot be found.

**static** `merge_notebooks`(*notebook\_files*: List[str], *output\_file*: str) → None

Merge the given Jupyter notebooks into a single Jupyter notebook.

**Parameters**

- **notebook\_files** (List[str]) – List of paths to the input Jupyter notebook files.
- **output\_file** (str) – Path to output Jupyter notebook file

**Return type**

None

**From transformer Module**

Class for transforming features.

**author**

Jeremy Biggs (jbiggs@ets.org)

**author**

Anastassia Loukina (aloukina@ets.org)

**author**

Nitin Madnani (nmadnani@ets.org)

**organization**

ETS

**class** `rsmttool.transformer.FeatureTransformer`(*logger*: Logger | None = None)

Bases: `object`

Encapsulate feature transformation methods.

Initialize the FeatureTransformer object.

**Parameters**

**logger** (Optional[logging.Logger]) – Logger object to use in the transformer. If not provided, a logger will be created with the name of this class.

**apply\_add\_one\_inverse\_transform**(*name: str, values: ndarray, raise\_error: bool = True*) → ndarray

Apply the “addOneInv” (add one and invert) transform to *values*.

**Parameters**

- **name** (*str*) – Name of the feature to transform.
- **values** (*numpy.ndarray*) – Numpy array containing the feature values.
- **raise\_error** (*bool*) – If True, raises an error if the transform is applied to a feature that has zero or negative values. Defaults to True.

**Returns**

**new\_data** – Numpy array containing the transformed feature values.

**Return type**

*numpy.ndarray*

**Raises**

**ValueError** – If the transform is applied to a feature that has negative values and *raise\_error* is True.

**apply\_add\_one\_log\_transform**(*name: str, values: ndarray, raise\_error: bool = True*) → ndarray

Apply the “addOneLn” (add one and log) transform to *values*.

**Parameters**

- **name** (*str*) – Name of the feature to transform.
- **values** (*numpy.ndarray*) – Numpy array containing the feature values.
- **raise\_error** (*bool*) – If True, raises an error if the transform is applied to a feature that has zero or negative values. Defaults to True.

**Returns**

**new\_data** – Numpy array that contains the transformed feature values.

**Return type**

*numpy.ndarray*

**Raises**

**ValueError** – If the transform is applied to a feature that has negative values and *raise\_error* is True.

**apply\_inverse\_transform**(*name: str, values: ndarray, raise\_error: bool = True, sd\_multiplier: int = 4*) → ndarray

Apply the “inv” (inverse) transform to *values*.

**Parameters**

- **name** (*str*) – Name of the feature to transform.
- **values** (*numpy.ndarray*) – Numpy array containing the feature values.
- **raise\_error** (*bool*) – If True, raises an error if the transform is applied to a feature that has zero values or to a feature that has both positive and negative values. Defaults to True.
- **sd\_multiplier** (*int*) – Use this std. dev. multiplier to compute the ceiling and floor for outlier removal and check that these are not equal to zero. Defaults to 4.

**Returns**

**new\_data** – Numpy array containing the transformed feature values.

**Return type**

*numpy.ndarray*

**Raises**

**ValueError** – If the transform is applied to a feature that is zero or to a feature that can have different signs, and `raise_error` is `True`.

**apply\_log\_transform**(*name*: *str*, *values*: *ndarray*, *raise\_error*: *bool* = *True*) → *ndarray*

Apply the “log” transform to *values*.

**Parameters**

- **name** (*str*) – Name of the feature to transform.
- **values** (*numpy.ndarray*) – Numpy array containing the feature values.
- **raise\_error** (*bool*) – If `True`, raises an error if the transform is applied to a feature that has zero or negative values. Defaults to `True`.

**Returns**

**new\_data** – Numpy array containing the transformed feature values.

**Return type**

*numpy.ndarray*

**Raises**

**ValueError** – If the transform is applied to a feature that has zero or negative values and `raise_error` is `True`.

**apply\_sqrt\_transform**(*name*: *str*, *values*: *ndarray*, *raise\_error*: *bool* = *True*) → *ndarray*

Apply the “sqrt” transform to *values*.

**Parameters**

- **name** (*str*) – Name of the feature to transform.
- **values** (*numpy.ndarray*) – Numpy array containing the feature values.
- **raise\_error** (*bool*) – If `True`, raises an error if the transform is applied to a feature that has negative values. Defaults to `True`.

**Returns**

**new\_data** – Numpy array containing the transformed feature values.

**Return type**

*numpy.ndarray*

**Raises**

**ValueError** – If the transform is applied to a feature that has negative values and `raise_error` is `True`.

**find\_feature\_transform**(*feature\_name*: *str*, *feature\_value*: *Series*, *scores*: *Series*) → *str*

Identify best transformation for feature given correlation with score.

The best transformation is chosen based on the absolute Pearson correlation with human score.

**Parameters**

- **feature\_name** (*str*) – Name of feature for which to find the transformation.
- **feature\_value** (*pandas.Series*) – Series containing feature values.
- **scores** (*pandas.Series*) – Numeric human scores.

**Returns**

**best\_transformation** – The name of the transformation which gives the highest correlation between the feature values and the human scores. See [documentation](#) for the full list of transformations.

### Return type

`str`

**transform\_feature**(values: `ndarray`, column\_name: `str`, transform: `str`, raise\_error: `bool` = `True`) → `ndarray`

Apply given transform to all values in the given numpy array.

The values are assumed to be for the feature with the given name.

### Parameters

- **values** (`numpy.ndarray`) – Numpy array containing the feature values.
- **column\_name** (`str`) – Name of the feature to transform.
- **transform** (`str`) – Name of the transform to apply. One of {"inv", "sqrt", "log", "addOneInv", "addOneLn", "raw", "org"}.
- **raise\_error** (`bool`) – If True, raise a ValueError if a transformation leads to invalid values or may change the ranking of the responses. Defaults to True.

### Returns

**new\_data** – Numpy array containing the transformed feature values.

### Return type

`numpy.ndarray`

### Raises

**ValueError** – If the given transform is not recognized.

---

**Note:** Many of these transformations may be meaningless for features which span both negative and positive values. Some transformations may throw errors for negative feature values.

---

## From utils Module

**class** `rsmttool.utils.commandline.ConfigurationGenerator`(context: `str`, as\_string: `bool` = `False`, suppress\_warnings: `bool` = `False`, use\_subgroups: `bool` = `False`)

Class to encapsulate automated batch-mode and interactive generation.

### context

Name of the command-line tool for which we are generating the configuration file.

### Type

`str`

### as\_string

If True, return a formatted and indented string representation of the configuration, rather than a dictionary. Note that this only affects the batch-mode generation. Interactive generation always returns a string. Defaults to False.

### Type

`bool`

### suppress\_warnings

If True, do not generate any warnings for batch-mode generation. Defaults to False.

### Type

`bool`

### use\_subgroups

If True, include subgroup-related sections in the list of general sections in the configuration file. Defaults to False.

#### Type

bool

Create a new ConfigurationGenerator instance.

See attributes above.

ConfigurationGenerator.generate() → str | Dict[str, Any]

Automatically generate an example configuration in batch mode.

#### Returns

**configuration** – The generated configuration either as a dictionary or a formatted string, depending on the value of the as\_string attribute.

#### Return type

Union[str, Dict[str, Any]]

rsmttool.utils.metrics.agreement(score1: List[int], score2: List[int], tolerance: int = 0) → float

Compute the agreement between two raters, under given tolerance.

#### Parameters

- **score1** (List[int]) – List of rater 1 scores
- **score2** (List[int]) – List of rater 2 scores
- **tolerance** (int) – Difference in scores that is acceptable. Defaults to 0.

#### Returns

**agreement\_value** – The percentage agreement between the two scores.

#### Return type

float

rsmttool.utils.metrics.difference\_of\_standardized\_means(y\_true\_observed: ndarray, y\_pred: ndarray, population\_y\_true\_observed\_mn: float | None = None, population\_y\_pred\_mn: float | None = None, population\_y\_true\_observed\_sd: float | None = None, population\_y\_pred\_sd: float | None = None, ddof: int = 1) → float | None

Calculate the difference between standardized means.

First, standardize both observed and predicted scores to z-scores using mean and standard deviation for the whole population. Then calculate differences between standardized means for each subgroup.

#### Parameters

- **y\_true\_observed** (numpy.ndarray) – The observed scores for the group or subgroup.
- **y\_pred** (numpy.ndarray) – The predicted scores for the group or subgroup.
- **population\_y\_true\_observed\_mn** (Optional[float]) – The population true score mean. When the DSM is being calculated for a subgroup, this should be the mean for the whole population. Defaults to None.
- **population\_y\_pred\_mn** (Optional[float]) – The predicted score mean. When the DSM is being calculated for a subgroup, this should be the mean for the whole population. Defaults to None.

- **population\_y\_true\_observed\_sd** (*Optional[float]*) – The population true score standard deviation. When the DSM is being calculated for a subgroup, this should be the standard deviation for the whole population. Defaults to `None`.
- **population\_y\_pred\_sd** (*Optional[float]*) – The predicted score standard deviation. When the DSM is being calculated for a subgroup, this should be the standard deviation for the whole population. Defaults to `None`.
- **ddof** (*int*) – The delta degrees of freedom. The divisor used in calculations is  $N - \text{ddof}$  where  $N$  represents the number of elements. Defaults to 1.

#### Returns

**difference\_of\_std\_means** – The difference of standardized means.

#### Return type

*Optional[float]*

#### Raises

- **ValueError** – If only one of `population_y_true_observed_mn` and `population_y_true_observed_sd` is not `None`.
- **ValueError** – If only one of `population_y_pred_mn` and `population_y_pred_sd` is not `None`.

`rsmttool.utils.metrics.partial_correlations(df: DataFrame) → DataFrame`

Implement the R `pcor` function from `ppcor` package in Python.

This computes partial correlations of each pair of variables in the given data frame `df`, excluding all other variables.

#### Parameters

**df** (*pandas.DataFrame*) – Data frame containing the feature values.

#### Returns

**df\_pcor** – Data frame containing the partial correlations of each pair of variables in the given data frame `df`, excluding all other variables.

#### Return type

*pandas.DataFrame*

`rsmttool.utils.metrics.quadratic_weighted_kappa(y_true_observed: ndarray, y_pred: ndarray, ddof: int = 0) → float`

Calculate quadratic-weighted kappa for both discrete and continuous values.

The formula to compute quadratic-weighted kappa for continuous values was developed at ETS by Shelby Haberman. See [Haberman \(2019\)](#) for the full derivation. The discrete case is simply treated as a special case of the continuous one.

The formula is as follows:

$QWK =$

*displaystyle*

$\frac{2 * Cov(M, H)Var(H) + Var(M) + (\bar{M} - \bar{H})^2}{2 * Cov(M, H)Var(H) + Var(M) + (\bar{M} - \bar{H})^2}$ , where

- $Cov$  - covariance with normalization by  $N$  (total number of observations)
- $H$  - the human score
- $M$  - the system score
- $\bar{H}$  - mean of  $H$



- $\bar{M}$  - mean of  $M$
- $Var(X)$  - variance of  $X$

#### Parameters

- **y\_true\_observed** (*numpy.ndarray*) – The observed scores.
- **y\_pred** (*numpy.ndarray*) – The predicted scores.
- **ddof** (*int*) – Means Delta Degrees of Freedom. The divisor used in calculations is  $N - \text{ddof}$ , where  $N$  represents the number of elements. When **ddof** is set to zero, the results for discrete case match those from the standard implementations. Defaults to 0.

#### Returns

**kappa** – The quadratic weighted kappa

#### Return type

float

#### Raises

**AssertionError** – If the number of elements in **y\_true\_observed** is not equal to the number of elements in **y\_pred**.

```
rsmtool.utils.metrics.standardized_mean_difference(y_true_observed: ndarray, y_pred: ndarray,
                                                    population_y_true_observed_sd: float | None =
                                                    None, population_y_pred_sd: float | None =
                                                    None, method: str = 'unpooled', ddof: int = 1) →
                                                    float
```

Compute the standardized mean difference between system and human scores.

The numerator is calculated as  $\text{mean}(y\_pred) - \text{mean}(y\_true\_observed)$  for all of the available methods.

#### Parameters

- **y\_true\_observed** (*numpy.ndarray*) – The observed scores for the group or subgroup.
- **y\_pred** (*numpy.ndarray*) – The predicted scores for the group or subgroup.
- **population\_y\_true\_observed\_sd** (*Optional[float]*) – The population true score standard deviation. When the SMD is being calculated for a subgroup, this should be the standard deviation for the whole population. Defaults to None.
- **population\_y\_pred\_sd** (*Optional[float]*) – The predicted score standard deviation. When the SMD is being calculated for a subgroup, this should be the standard deviation for the whole population. Defaults to None.
- **method** (*str*) – The SMD method to use. Possible options are:
  - "williamson": Denominator is the pooled population standard deviation of **y\_true\_observed** and **y\_pred** computed using **population\_y\_true\_observed\_sd** and **population\_y\_pred\_sd**.
  - "johnson": Denominator is **population\_y\_true\_observed\_sd**.
  - "pooled": Denominator is the pooled standard deviation of **y\_true\_observed** and **y\_pred** for this group.
  - "unpooled": Denominator is the standard deviation of **y\_true\_observed** for this group.
 Defaults to "unpooled".

- **ddof** (*int*) – The delta degrees of freedom. The divisor used in calculations is  $N - \text{ddof}$ , where  $N$  represents the number of elements. Defaults to 1.

#### Returns

**smd** – The SMD for the given group or subgroup.

#### Return type

*float*

#### Raises

- **ValueError** – If method is “williamson” and either `population_y_true_observed_sd` or `population_y_pred_sd` is `None`.
- **ValueError** – If method is “johnson” and `population_y_true_observed_sd` is `None`.
- **ValueError** – If method is not one of {“unpooled”, “pooled”, “williamson”, “johnson”}.

---

#### Note:

- The “williamson” implementation was recommended by Williamson, et al. (2012).
  - The metric is only applicable when both sets of scores are on the same scale.
- 

`rsmttool.utils.metrics.compute_expected_scores_from_model`(*model*: *Learner*, *featureset*: *FeatureSet*, *min\_score*: *float*, *max\_score*: *float*) → *ndarray*

Compute expected scores using probability distributions over labels.

This function only works with SKLL models.

#### Parameters

- **model** (*skll.learner.Learner*) – The SKLL learner object to use for computing the expected scores.
- **featureset** (*skll.data.FeatureSet*) – The SKLL featureset object for which predictions are to be made.
- **min\_score** (*float*) – Minimum score level to be used for computing expected scores.
- **max\_score** (*float*) – Maximum score level to be used for computing expected scores.

#### Returns

**expected\_scores** – A numpy array containing the expected scores.

#### Return type

*numpy.ndarray*

#### Raises

- **ValueError** – If the given model cannot predict probability distributions.
- **ValueError** – If the score range specified by `min_score` and `max_score` does not match what the model predicts in its probability distribution.

`rsmttool.utils.notebook.get_thumbnail_as_html`(*path\_to\_image*: *str*, *image\_id*: *int*, *path\_to\_thumbnail*: *str* | *None* = *None*) → *str*

Generate HTML for a clickable thumbnail of given image.

Given the path to an image file, generate the HTML for a clickable thumbnail version of the image. When clicked, this HTML will open the full-sized version of the image in a new window.

### Parameters

- **path\_to\_image** (*str*) – The absolute or relative path to the image. If an absolute path is provided, it will be converted to a relative path.
- **image\_id** (*int*) – The id of the <img> tag in the HTML. This must be unique for each <img> tag.
- **path\_to\_thumbnail** (*Optional[str]*) – If you would like to use a different thumbnail image, specify the path to this thumbnail. Defaults to None.

### Returns

**image** – The HTML string generated for the image.

### Return type

*str*

### Raises

**FileNotFoundError** – If the image file cannot be located.

`rsmttool.utils.notebook.show_thumbnail(path_to_image: str, image_id: int, path_to_thumbnail: str | None = None) → None`

Display the HTML for an image thumbnail in a Jupyter notebook.

Given the path to an image file, generate the HTML for its thumbnail and display it in the notebook.

### Parameters

- **path\_to\_image** (*str*) – The absolute or relative path to the image. If an absolute path is provided, it will be converted to a relative path.
- **image\_id** (*int*) – The id of the <img> tag in the HTML. This must be unique for each <img> tag.
- **path\_to\_thumbnail** (*Optional[str]*) – If you would like to use a different thumbnail image, specify the path to the thumbnail. Defaults to None.

### Return type

*None*

`rsmttool.utils.files.parse_json_with_comments(pathlike: str | Path) → Dict[str, Any]`

Parse a JSON file after removing any comments.

Comments can use either `//` for single-line comments or `/* ... */` for multi-line comments. The input filepath can be a string or `pathlib.Path`.

### Parameters

**filename** (*Union[str, Path]*) – Path to the input JSON file either as a string or as a `pathlib.Path` object.

### Returns

**obj** – JSON object representing the input file.

### Return type

`Dict[str, Any]`

---

**Note:** This code was adapted from: <https://web.archive.org/web/20150520154859/http://www.lifl.fr/~riquetd/parse-a-json-file-with-comments.html>

---

## From writer Module

Class for writing DataContainer frames to disk.

### author

Jeremy Biggs ([jbiggs@ets.org](mailto:jbiggs@ets.org))

### author

Anastassia Loukina ([aloukina@ets.org](mailto:aloukina@ets.org))

### author

Nitin Madnani ([nmadnani@ets.org](mailto:nmadnani@ets.org))

### organization

ETS

```
class rsmtool.writer.DataWriter(experiment_id: str | None = None, context: str | None = None, wandb_run:
    Run | RunDisabled | None = None)
```

Bases: `object`

Class to write out DataContainer objects.

Initialize the DataWriter object.

### Parameters

- **experiment\_id** (*Optional*[*str*]) – The experiment name to be used in the output file names. Defaults to *None*.
- **context** (*Optional*[*str*]) – The context in which this writer is used. Defaults to *None*.
- **wandb\_run** (*Union*[*wandb.wandb\_run.Run*, *wandb.sdk.lib.RunDisabled*, *None*]) – The wandb run object if wandb is enabled, *None* otherwise. If enabled, all the output data frames will be logged to this run as tables. Defaults to *None*.

```
write_experiment_output(csvdir: str, container_or_dict: DataContainer | Dict[str, DataFrame],
    dataframe_names: List[str] | None = None, new_names_dict: Dict[str, str] |
    None = None, include_experiment_id: bool = True, reset_index: bool = False,
    file_format: str = 'csv', index: bool = False, **kwargs) → None
```

Write out each of the named frames to disk.

This function writes out each of the given list of data frames as a “.csv”, “.tsv”, or .xlsx file in the given directory. Each data frame was generated as part of running an RSMTTool experiment. All files are prefixed with the given experiment ID and suffixed with either the name of the data frame in the DataContainer (or dict) object, or a new name if *new\_names\_dict* is specified. Additionally, the indexes in the data frames are reset if so specified.

### Parameters

- **csvdir** (*str*) – Path to the output experiment sub-directory that will contain the CSV files corresponding to each of the data frames.
- **container\_or\_dict** (*Union*[*container.DataContainer*, Dict[*str*, *pd.DataFrame*]]) – A DataContainer object or dict, where keys are data frame names and values are pandas.DataFrame objects.
- **dataframe\_names** (*Optional*[List[*str*]]) – List of data frame names, one for each of the data frames. Defaults to *None*.
- **new\_names\_dict** (*Optional*[Dict[*str*, *str*]]) – New dictionary with new names for the data frames, if desired. Defaults to *None*.

- **include\_experiment\_id** (*bool*) – Whether to include the experiment ID in the file name. Defaults to True.
- **reset\_index** (*bool*) – Whether to reset the index of each data frame before writing to disk. Defaults to False.
- **file\_format** (*str*) – The file format in which to output the data. One of {"csv", "xlsx", "tsv"}. Defaults to "csv".
- **index** (*bool*) – Whether to include the index in the output file. Defaults to False.

#### Raises

**KeyError** – If `file_format` is not valid, or a named data frame is not present in `container_or_dict`.

#### Return type

None

**write\_feature\_csv**(*featuringdir: str, data\_container: DataContainer, selected\_features: List[str], include\_experiment\_id: bool = True, file\_format: str = 'csv'*) → None

Write out the selected features to disk.

#### Parameters

- **featuringdir** (*str*) – Path to the experiment output directory where the feature JSON file will be saved.
- **data\_container** (*DataContainer*) – A data container object.
- **selected\_features** (*List[str]*) – List of features that were selected for model building.
- **include\_experiment\_id** (*bool*) – Whether to include the experiment ID in the file name. Defaults to True.
- **file\_format** (*str*) – The file format in which to output the data. One of {"csv", "tsv", "xlsx"}. Defaults to "csv".

#### Return type

None

**static write\_frame\_to\_file**(*df: DataFrame, name\_prefix: str, file\_format: str = 'csv', index: bool = False, \*\*kwargs*) → None

Write given data frame to disk with given name and file format.

#### Parameters

- **df** (*pandas.DataFrame*) – Data frame to write to disk
- **name\_prefix** (*str*) – The complete prefix for the file to be written to disk. This includes everything except the extension.
- **file\_format** (*str*) – The file format (extension) for the file to be written to disk. One of {"csv", "xlsx", "tsv"}. Defaults to "csv".
- **index** (*bool*) – Whether to include the index in the output file. Defaults to False.

#### Raises

**KeyError** – If `file_format` is not valid.

#### Return type

None

## 1.11 Utility Scripts

In addition to the *rsmtool*, *rsmeval*, *rsmpredict*, *rscompare*, and *rsmsummarize* scripts, RSMTool also comes with a number of helpful utility scripts.

### 1.11.1 render\_notebook

Convert a given Jupyter notebook file (.ipynb) to HTML (.html) format.

**ipynb\_file**

Path to input Jupyter notebook file.

**html\_file**

Path to output HTML file.

**-h, --help**

Show help message and exit.

### 1.11.2 convert\_feature\_json

Convert an older feature JSON file to a new file in tabular format.

**--json**

Path to input feature JSON file that is to be converted.

**--output**

Path to output CSV/TSV/XLSX file containing the features in tabular format.

**--delete**

Delete original feature JSON file after conversion.

**-h, --help**

Show help message and exit.

## 1.12 Frequently Asked Questions

### . Can I train a deep learning model using RSMTool? [#]

*rsmtool* supports a large number of *learners* but not deep learning algorithms. However, you can still use *rsmeval* to obtain a comprehensive evaluation report from any model.

### . How can I customize *rsmtool* and *rsmeval* reports? [#]

The reports generated by *rsmtool* and *rsmeval* are designed to be fully customizable. You can:

- Omit sections using the *general\_sections* configuration option;
- Change the order of the sections using the *section\_order* option;
- Create your own “*custom*” sections.

. **I get the following error:** No responses remaining after filtering out non-numeric feature values. No further analysis can be run. **What happened?** [#]

rsmttool is designed to work only with numeric features. Non-numeric values *including* missing values are filtered out. Some of the common reasons for this error are:

- The human score column or one of the features contains only non-numeric values. You can either exclude this feature or convert it to [one-hot encoding](#)
- You have features with missing values. The solution is to replace missing values with zeros. Note that this applies *even* if you use the `.jsonlines` format. We have an open [issue](#) to provide an option to automatically convert missing values to zeros.
- You have a lot of missing feature values and none of the responses has numeric features for every single feature. Inspect the `*_excluded_responses` files in the output directory to see what responses have been excluded.

. **Can I pass a set of learners to the model option in the configuration file or do I need a separate run for each learner I want to try?** [#]

rsmttool cannot train multiple models via a configuration file. If you wish to use multiple learners/models, you should use the [RSMTool API](#) instead of the command line.

. **Can I compute predictions on new data using a model trained with rsmttool ?** [#]

Yes! We have built [rsmpredict](#) to do exactly this!

. **Why did rsmttool change the sign of some features in feature.csv ? I thought rsmttool assumes a positive sign for all features?** [#]

rsmttool indeed assumes a default positive sign for all raw features. However, if you set `select_transformations` to `true` in your experiment configuration file, RSMTool will automatically apply [transformations](#) to some of the features. Some transformations such as `inv` (inverse transform) change the polarity of the feature. In such cases, RSMTool takes this into account and changes the sign accordingly. See [note here](#).

. **Why does rsmpredict exclude a different number of responses compared to rsmttool ?** [#]

This mismatch may happen due to the following:

1. First, rsmttool is run with one of the built-in [feature selection models](#) and excludes responses where *any* of the candidate features has missing/non-numeric values in either the training or the evaluation set. Next, it performs the feature selection, trains a model with those features using the filtered training set and generates predictions on the filtered test set.
2. Next, rsmpredict is run to generate predictions on the evaluation set using the trained model and excludes *only* those responses where any of the *selected* features (a subset of the original features) has missing/non-numeric values.

As a result, there is a mismatch in the number of predictions generated by rsmttool and rsmpredict.

. **The relative betas for the linear regression model in my rsmttool run do not sum to 1. Is that a problem?** [#]

Please check if your model has any negative coefficients. Relative coefficients *only* make sense when all model coefficients are positive. Their sum is expected to be less than 1 if there are negative coefficients. Note that if this is the case, the relative coefficients will *not* be included into the report.

. **Can I use Weights & Biases to log artifacts and reports?** [#]

Yes! W&B logging is natively supported by the following command-line tools: `rsmtool`, `rsmeval`, `rsmxval`, `rsmexplain`, `rsmcompare`, and `rsmsummarize`. To enable it, set the `use_wandb` configuration option to `true` and set the `wandb_entity` and `wandb_project` options appropriately.

. **It's so difficult to remember all the configuration options.** [#]

We agree that it can be quite overwhelming to remember all the configuration options. This is why all command-line tools support *automatic generation of configuration files*!

## 1.13 Contributing to RSMTool

Contributions to RSMTool are very welcome. You can use the instructions below to get started on developing new features or functionality for RSMTool. When contributing to RSMTool, all of your contributions must come with updates to documentation as well as tests.

### 1.13.1 Setting up

To set up a local development environment, follow the steps below:

1. Clone the [Github repository](#) for RSMTool. If you have already have a local version of the repository, pull the latest version from GitHub and switch to the main branch.
2. If you already have the conda package manager installed, skip to the next step. If you do not, follow the instructions on [this page](#) to install conda.
3. Create a new conda environment (say, `rsmdev`) and install the packages specified in the `requirements.dev` file by running:

```
conda create -n rsmdev -c conda-forge -c ets --file requirements.dev
```

4. Activate the environment using `conda activate rsmdev`.
5. Run `pip install -e .` to install `rsmtool` into the environment in editable mode which is what we need for development.
6. Run `pre-commit install --hook-type pre-commit --hook-type commit-msg --overwrite --install-hooks` to install the git hooks for pre-commit checks.
7. Create a new git branch with a useful and descriptive name.
8. Make your changes and add tests. See the next section for more on writing new tests.
9. Run `nose2 --quiet -s tests` to run the tests. We use the `--quiet` switch, since otherwise test failures for some tests tend to produce very long Jupyter notebook traces.



### 1.13.2 Documentation

Note that the file `doc/requirements.txt` is meant specifically for the ReadTheDocs documentation build process and should not be used locally. To build the documentation locally, you *must* use the same conda environment created above.

### 1.13.3 Code style

The RSMTTool codebase enforces a certain code style via pre-commit checks and this style is automatically applied to any new contributions. This code style consists of:

1. All Python code is formatted via the `black` pre-commit check.
2. The f-string specification is used for all format strings in the Python code as enforced by the `flynt` pre-commit check.
3. Any code violations for `PEP 8`, `PEP 257`, and import statements are checked using the `ruff` pre-commit check.
4. Commit messages must follow the `conventional commit specification`. This is enforced by the `conventional-pre-commit` pre-commit check.
5. The imports at the top of any Python files are grouped and sorted as follows: `STDLIB`, `THIRDPARTY`, `FIRST-PARTY`, `LOCALFOLDER`. As an example, consider the imports at the top of `reporter.py` which look like this:

```
import argparse
import asyncio
import json
import logging
import os
import sys
from os.path import abspath, basename, dirname, join, splitext

from nbconvert.exporters import HTMLExporter
from nbconvert.exporters.templateexporter import default_filters
from traitlets.config import Config

from .reader import DataReader
```

Rather than doing this grouping and sorting manually, we use the `isort` pre-commit hook to achieve this.

6. All classes, functions, and methods in the main code files have `numpy-formatted docstrings` that comply with `PEP 257`. This is enforced via the `pydocstyle` pre-commit check. Additionally, when writing docstrings, make sure to use the appropriate quotes when referring to argument names vs. argument values. As an example, consider the docstring for the `train_skill_model` method of the `rsmttool.modeler.Modeler` class. Note that string argument values are enclosed in double quotes (e.g., `"csv"`, `"neg_mean_squared_error"`) whereas values of other built-in types are written as literals (e.g., `True`, `False`, `None`). Note also that if one had to refer to an argument name in the docstring, this referent should be written as a literal. In general, we strongly encourage looking at the docstrings in the existing code to make sure that new docstrings follow the same practices.

### 1.13.4 RSMTTool tests

Existing tests for RSMTTool are spread across the various `test_*.py` files under the `tests` directory after you check out the RSMTTool source code from GitHub.

There are two kinds of existing tests in RSMTTool:

1. The first type of tests are **unit tests**, i.e., very specific tests for which you have a single example (usually embedded in the test itself) and you compare the generated output with known or expected output. These tests should have a very narrow and well defined scope. To see examples of such unit tests, see the test functions in the file `tests/test_utils.py`.
2. The second type of tests are **functional tests** which are generally written from the users' perspective to test that RSMTTool is doing things that users would expect it to. In RSMTTool, most (if not all) functional tests are written in the form of "experiment tests", i.e., we first define an experimental configuration using an `rsmttool` (or `rsmeval/rsmpredict/rsmcompare/rsmsummarize`) configuration file, then we run the experiment, and then compare the generated output files to expected output files to make sure that RSMTTool components are operating as expected. To see examples of such tests, you can look at any of the `tests/test_experiment_*.py` files.

---

**Note:** RSMTTool functional tests are *parameterized*, i.e., since most are identical other than the configuration file that needs to be run, the basic functionality of the test has been factored out into utility functions. Each line starting with *param* in any of the `test_experiment_*.py` files represents a specific functional test.

---

Any new contributions to RSMTTool, no matter how small or trivial, *must* be accompanied by updates to documentations as well as new unit and/or functional tests. Adding new unit tests is fairly straightforward. However, adding new functional tests is a little more involved.

### 1.13.5 Writing new functional tests

To write a new experiment test for RSMTTool (or any of the other tools):

- (a) Create a new directory under `tests/data/experiments` using a descriptive name.
- (b) Create a JSON configuration file under that directory with the various fields appropriately set for what you want to test. Feel free to use multiple words separated by hyphens to come up with a name that describes the testing condition. The name of the configuration file should be the same as the value of the `experiment_id` field in your JSON file. By convention, that's usually the same as the name of the directory you created but with underscores instead of hyphens. If you are creating a new test for `rsmcompare` or `rsmsummarize`, copy over one or more of the existing `rsmttool` or `rsmeval` test experiments as input(s) and keep the same name. This will ensure that these inputs will be regularly updated and remain consistent with the current outputs generated by these tools. If you must create a test for a scenario not covered by a current tool, create a new `rsmttool/rsmeval` test first following the instructions on this page.
- (c) Next, you need to add the test to the list of parameterized tests in the appropriate test file based on the tool for which you are adding the test, e.g., `rsmeval` tests should be added to `tests/test_experiment_rsmeval.py`, `rsmpredict` tests to `tests/test_experiment_rsmpredict.py`, and so on. Tests for `rsmttool` can be added to any of the four files. The arguments for the `param()` call can be found in the [Table 1](#) below.
- (d) In some rare cases, you might want to use a non-parameterized experiment test if you are doing something very different. These should be few and far between. Examples of these can also be seen in various `tests/test_experiment_*.py` files.

- (e) Another rare scenario is the need to create an entirely new `tests/test_experiment_X.py` file instead of using one of the existing ones. This should *not* be necessary unless you are trying to test a newly added tool or component.

Table 1: Table 1: Arguments for `param()` when adding new parameterized functional tests

<p>Writing test(s) for <code>rsmtool</code></p> <ul style="list-style-type: none"> <li>• First positional argument is the name of the test directory you created.</li> <li>• Second positional argument is the experiment ID from the JSON configuration file.</li> <li>• Use <code>consistency=True</code> if you have set <code>second_human_score_column</code> in the configuration file.</li> <li>• Use <code>skill=True</code> if you are writing a test for a SKLL model.</li> <li>• Set <code>subgroups</code> keyword argument to the same list that you specified in the configuration file.</li> <li>• Set <code>file_format="tsv"</code> (or <code>"xlsx"</code>) if you specified the same field in the configuration file.</li> </ul>
<p>Writing test(s) for <code>rsmeval</code></p> <ul style="list-style-type: none"> <li>• Same arguments as <code>RSMTool</code> except the <code>skill</code> keyword argument is not applicable.</li> </ul>
<p>Writing test(s) for <code>rsmpredict</code></p> <ul style="list-style-type: none"> <li>• The only positional argument is the name of the test directory you created.</li> <li>• Use <code>excluded=True</code> if you want to check the excluded responses file as part of the test.</li> <li>• Set <code>file_format="tsv"</code> (or <code>"xlsx"</code>) if you specified the same field in the configuration file.</li> </ul>
<p>Writing test(s) for <code>rsmcompare</code></p> <ul style="list-style-type: none"> <li>• First positional argument is the name of the test directory you created.</li> <li>• Second positional argument is the comparison ID from the JSON configuration file.</li> </ul>
<p>Writing test(s) for <code>rsmsummarize</code></p> <ul style="list-style-type: none"> <li>• The only positional argument is the name of the test directory you created.</li> <li>• Set <code>file_format="tsv"</code> (or <code>"xlsx"</code>) if you specified the same field in the configuration file.</li> </ul>
<p>Writing test(s) for <code>rsmexplain</code></p> <ul style="list-style-type: none"> <li>• First positional argument is the name of the test directory you created.</li> <li>• Second positional argument is the experiment ID from the JSON configuration file.</li> </ul>

Once you have added all new functional tests, commit all of your changes. Next, you should run `nose2` to run all the tests. Obviously, the newly added tests will fail since you have not yet generated the expected output for that test.

To do this, you should now run the following:

```
python tests/update_files.py --tests tests --outputs test_outputs
```

This will copy over the generated outputs for the newly added tests and show you a report of the files that it added. It will also update the input files for tests for `rsmcompare` and `rsmsummarize`. If run correctly, the report should *only* refer to the files affected by the functionality you implemented. If you run `nose2` again, your newly added tests should now pass.

At this point, you should inspect all of the new test files added by the above command to make sure that the outputs are as expected. You can find these files under `tests/data/experiments/<test>/output` where `<test>` refers to the test(s) that you added.

However, if your changes resulted in updates to the inputs to `rsmsummarize` or `rsmcompare` tests, you will first need to re-run the tests for these two tools and then re-run the `update_files.py` to update the outputs.

Once you are satisfied that the outputs are as expected, you can commit them.

The two examples below might help make this process easier to understand:

**Example 1: You made a code change to better handle an edge case that only affects one test.**

1. Run `nose2 --quiet -s tests`. The affected test failed.
2. Run `python tests/update_files.py --tests tests --outputs test_outputs` to update test outputs. You will see the total number of deleted, updated and missing files. There should be no deleted files and no missing files. Only the files for your new test should be updated. There are no warnings in the output.
3. If this is the case, you are now ready to commit your change and the updated test outputs.

**Example 2: You made a code change that changes the output of many tests. For example, you renamed one of the evaluation metrics.**

1. Run `nose2 --quiet -s tests`. Many tests will now fail since the output produced by the tool(s) has changed.
2. Run `python tests/update_files.py --tests tests --outputs test_outputs` to update test outputs. The files affected by your change are shown as added/deleted. You also see the following warning:

WARNING: X input files for rsmcompare/rsmsummarize tests have been updated. You need to re-run these tests and update test outputs

3. This means that the changes you made to the code changed the outputs for one or more `rsmtool/rsmeval` tests that served as inputs to one or more `rsmcompare/rsmsummarize` tests. Therefore, it is likely that the current test outputs no longer match the expected output and the tests for those two tools must be re-run.
4. Run `nose2 --quiet -s tests $(find tests -name 'test_*rsmsummarize*.py' | cut -d '/' -f2 | sed 's/.py//')` and `nose2 --quiet -s tests $(find tests -name 'test_*rsmcompare*.py' | cut -d '/' -f2 | sed 's/.py//')`. If you see any failures, make sure they are related to the changes you made since those are expected.
5. Next, re-run `python tests/update_files.py --tests tests --outputs test_outputs` which should only update the outputs for the `rsmcompare/rsmsummarize` tests.
6. If this is the case, you are now ready to commit your changes.

### 1.13.6 Advanced tips and tricks

Here are some advanced tips and tricks when working with RSMTool tests.

1. To run a specific test function in a specific test file, simply use `nose2 --quiet -s tests test_X.Y.Z` where `test_X` is the name of the test file, `Y` is the enclosing `unittest.TestCase` subclass, and `Z` is the desired test function. Note that this will not work for parameterized tests. If you want to run a specific parameterized test, you can comment out all of the other parameters in the `params` and run the `test_run_experiment_parameterized()` function as above.
2. If you make any changes to the code that can change the output that the tests are expected to produce, you *must* re-run all of the tests and then update the *expected* test outputs using the `update_files.py` command as shown [above](#).
3. In the rare case that you *do* need to create an entirely new `tests/test_experiment_X.py` file instead of using one of the existing ones, you can choose whether to exclude the tests contained in this file from updating their

expected outputs when `update_files.py` is run by setting `_AUTO_UPDATE=False` at the top of the file. This should *only* be necessary if you are absolutely sure that your tests will never need updating.

4. The `--debugger/-D` option for nose2 is your friend. If you encounter test errors or test failures where the cause may not be immediately clear, re-run the nose2 command with this option. Doing so will drop you into an interactive PDB session as soon as an error (or failure) is encountered and then you inspect the variables at that point (or use “u” and “d” to go up and down the call stack). This may be particularly useful for tests in `tests/test_cli.py` that use `subprocess.run()`. If these tests are erroring out, use `-D` and inspect the “`stderr`” variable in the resulting PDB session to see what the error is.
5. In RSMTool 8.0.1 and later, the tests will pass even if any of the reports contain warnings. To catch any warnings that may appear in the reports, run the tests in strict mode (`STRICT=1 nose2 --quiet -s tests`).

## 1.14 Internal Documentation

This section of the documentation is meant only for the project administrators, not users and developers.

### 1.14.1 RSMTool Release Process

This process is only meant for the project administrators, not users and developers.

1. Recreate the development environment so all unpinned packages are updated to their latest versions. See instructions for this [here](#).
2. Make sure any and all tests are passing in `main`. Make sure you have also run tests locally in strict mode (`STRICT=1 nose2 -s tests`) to catch any warnings in the HTML report that can be fixed before the release.
3. Run the `tests/update_files.py` script with the appropriate arguments to make sure that all test data in the new release have correct experiment ids and filenames. If any (non-model) files need to be changed this should be investigated before the branch is released. Please see more details about running this [here](#).

---

**Note:** Several files have been excluded from the repository due to their non-deterministic nature so please do not add them back to the repository. The following files are currently excluded:

- Fairness test files for *lr-eval-system-score-constant* test
- Predictions and all evaluation files for *linearsvr* test.

Note that the full set of outputs from these test files are also used as input for *rsmcompare* and *rsmsummarize* tests. These *input* files need to be updated following the process under **Example 2** in [Writing new functional tests](#). You can also see [this pull request](#) for more information.

---

4. Create a release branch `release/XX` on GitHub.
5. In the release branch:
  - Update the version number in `version.py`.
  - Make sure that `requirements.txt` only has the actual dependencies that are needed to run RSMTool. Any dependencies needed only for development/testing (e.g., *sphinx*, *nose2* etc.) should be moved to `requirements.dev`. This means that `requirements.txt` must be a strict subset of `requirements.dev`.
  - Make sure the versions in `doc/requirements.txt` are up to date with `requirements.txt` and only contains the dependencies needed to build the documentation.
  - Make sure `.readthedocs.yml` is still accurate.

- Update the conda recipe.
  - Update the documentation with any new features or details about changes.
  - Run `make linkcheck` on the documentation (i.e. from the `doc/` folder) and fix any redirected/broken links.
  - Update the README and this release documentation, if necessary. Make sure to update the documentation badge in the README to use the latest release tag. That documentation version will be created in the steps below.
6. Build the PyPI source and wheel distributions using `python setup.py sdist build` and `python setup.py bdist_wheel build` respectively. Note that you should delete the `build` directory after running the `sdist` command and before running the `bdist_wheel` command.
  7. Upload the source and wheel distributions to TestPyPI using `twine upload --repository testpypi dist/*`. You will need to have the `twine` package installed and set up your `$HOME/.pypirc` correctly. See details [here](#). You will need to have the appropriate permissions for the `ets` organization on TestPyPI.
  8. Install the TestPyPI package as follows:

```
pip install --index-url https://test.pypi.org/simple/ --extra-index-url https://  
↪pypi.org/simple rsmtool
```

9. Then run some tests from a RSMTool working copy. If the TestPyPI package works, then move on to the next step. If it doesn't, figure out why and rebuild and re-upload the package.
10. Build the new conda package by running the following command in the `conda-recipe` directory (note that this assumes that you have cloned RSMTool in a directory named `rsmtool`). You should delete the `build` directory before running the command below. Note that you may need to comment out lines in your `$HOME/.condarc` file if you are using ETS Artifactory and you get conflicts:

```
conda build -c conda-forge .
```

11. This will create a noarch package with the path to the package printed out to the screen.
12. Upload the package file to anaconda.org using `anaconda upload --user ets <path_to_file>`. You will need to have the appropriate permissions for the `ets` organization.
13. Create pull requests on the [rsmtool-conda-tester](#) and [rsmtool-pip-tester](#) repositories to test the conda and TestPyPI packages on Linux and Windows.
14. Draft a release on GitHub while the Linux and Windows package tester builds are running.
15. Once both builds have passed, make a pull request with the release branch to be merged into `main` and request code review.
16. Once the build for the PR passes and the reviewers approve, merge the release branch into `main`.
17. Upload the already-built source and wheel packages to PyPI using `twine upload dist/*`. You will need to have the `twine` package installed and set up your `$HOME/.pypirc` correctly. You will need to have the appropriate permissions for the `ets` organization on PyPI.
18. Make sure that the ReadTheDocs build for `main` passes by examining the badge at this [URL](#) - this should say "passing" in green.
19. Tag the latest commit in `main` with the appropriate release tag.
20. Create a new release in Github based on this new tag and auto-generate the release notes. Edit the auto-generated release notes based on the format from the previous releases.
21. Publish the release on GitHub.

22. Go to ReadTheDocs and re-run the build for the *main* branch. This will pull in the new tag. Once the new tag is visible in ReadTheDocs, create a new documentation version based on this tag and make it the default version.





## INDICES AND TABLES

- `genindex`
- `modindex`
- `search`



## PYTHON MODULE INDEX

### r

- `rsmttool.analyzer`, [104](#)
- `rsmttool.comparer`, [116](#)
- `rsmttool.configuration_parser`, [118](#)
- `rsmttool.container`, [124](#)
- `rsmttool.modeler`, [128](#)
- `rsmttool.preprocessor`, [138](#)
- `rsmttool.reader`, [153](#)
- `rsmttool.reporter`, [155](#)
- `rsmttool.transformer`, [159](#)
- `rsmttool.writer`, [168](#)



## Symbols

-v  
 rsmcompare command line option, 63  
 rsmeval command line option, 44  
 rsmpredict command line option, 57, 84  
 rsmsummarize command line option, 70  
 rsmtool command line option, 19  
 rsmxval command line option, 78  
 --delete  
 convert\_feature\_json command line option, 170  
 --features  
 rsmpredict command line option, 56  
 --force  
 rsmeval command line option, 44  
 rsmpredict command line option, 84  
 rsmsummarize command line option, 70  
 rsmtool command line option, 19  
 --help  
 convert\_feature\_json command line option, 170  
 render\_notebook command line option, 170  
 rsmcompare command line option, 63  
 rsmeval command line option, 44  
 rsmpredict command line option, 56, 84  
 rsmsummarize command line option, 70  
 rsmtool command line option, 19  
 rsmxval command line option, 78  
 --json  
 convert\_feature\_json command line option, 170  
 --output  
 convert\_feature\_json command line option, 170  
 --version  
 rsmcompare command line option, 63  
 rsmeval command line option, 44  
 rsmpredict command line option, 57, 84  
 rsmsummarize command line option, 70  
 rsmtool command line option, 19  
 rsmxval command line option, 78  
 -f

rsmeval command line option, 44  
 rsmpredict command line option, 84  
 rsmsummarize command line option, 70  
 rsmtool command line option, 19  
 -h  
 convert\_feature\_json command line option, 170  
 render\_notebook command line option, 170  
 rsmcompare command line option, 63  
 rsmeval command line option, 44  
 rsmpredict command line option, 56, 84  
 rsmsummarize command line option, 70  
 rsmtool command line option, 19  
 rsmxval command line option, 78

## A

add\_dataset() (*rsmtool.container.DataContainer* method), 124  
 agreement() (*in module rsmtool.utils.metrics*), 163  
 analyze\_excluded\_responses() (*rsmtool.analyzer.Analyzer* static method), 104  
 analyze\_used\_predictions() (*rsmtool.analyzer.Analyzer* static method), 104  
 analyze\_used\_responses() (*rsmtool.analyzer.Analyzer* static method), 105  
 Analyzer (*class in rsmtool.analyzer*), 104  
 apply\_add\_one\_inverse\_transform() (*rsmtool.transformer.FeatureTransformer* method), 159  
 apply\_add\_one\_log\_transform() (*rsmtool.transformer.FeatureTransformer* method), 160  
 apply\_inverse\_transform() (*rsmtool.transformer.FeatureTransformer* method), 160  
 apply\_log\_transform() (*rsmtool.transformer.FeatureTransformer* method), 161  
 apply\_sqrt\_transform() (*rsmtool.transformer.FeatureTransformer* method), 161

as\_string(*rsmtool.utils.commandline.ConfigurationGenerator* attribute), 162

## C

check\_exclude\_listwise() (*rsmtool.configuration\_parser.Configuration* method), 119

check\_feature\_subset\_file() (*rsmtool.preprocessor.FeatureSubsetProcessor* method), 151

check\_flag\_column() (*rsmtool.configuration\_parser.Configuration* method), 119

check\_frame\_names() (*rsmtool.analyzer.Analyzer* static method), 105

check\_model\_name() (*rsmtool.preprocessor.FeaturePreprocessor* method), 139

check\_param\_names() (*rsmtool.analyzer.Analyzer* static method), 105

check\_section\_names() (*rsmtool.reporter.Reporter* static method), 156

check\_section\_order() (*rsmtool.reporter.Reporter* static method), 156

check\_subgroups() (*rsmtool.preprocessor.FeaturePreprocessor* method), 139

Comparer (class in *rsmtool.comparer*), 116

compute\_and\_save\_predictions() (in module *rsmtool*), 100

compute\_basic\_descriptives() (*rsmtool.analyzer.Analyzer* static method), 106

compute\_correlations\_between\_versions() (*rsmtool.comparer.Comparer* static method), 117

compute\_correlations\_by\_group() (*rsmtool.analyzer.Analyzer* method), 106

compute\_degradation\_and\_disattenuated\_correlations() (*rsmtool.analyzer.Analyzer* method), 107

compute\_disattenuated\_correlations() (*rsmtool.analyzer.Analyzer* static method), 107

compute\_expected\_scores\_from\_model() (in module *rsmtool.utils.metrics*), 166

compute\_metrics() (*rsmtool.analyzer.Analyzer* method), 107

compute\_metrics\_by\_group() (*rsmtool.analyzer.Analyzer* method), 109

compute\_outliers() (*rsmtool.analyzer.Analyzer* static method), 109

compute\_pca() (*rsmtool.analyzer.Analyzer* static method), 109

compute\_percentiles() (*rsmtool.analyzer.Analyzer* static method), 110

config\_file  
rsmcompare command line option, 63

rsmeval command line option, 43

rsmpredict command line option, 56, 84

rsmsummarize command line option, 70

rsmtool command line option, 19

rsmxval command line option, 78

configdir (*rsmtool.configuration\_parser.Configuration* property), 119

Configuration (class in *rsmtool.configuration\_parser*), 118

ConfigurationGenerator (class in *rsmtool.utils.commandline*), 162

ConfigurationParser (class in *rsmtool.configuration\_parser*), 122

configure() (in module *rsmtool.configuration\_parser*), 123

context (*rsmtool.configuration\_parser.Configuration* property), 120

context (*rsmtool.utils.commandline.ConfigurationGenerator* attribute), 162

convert\_feature\_json command line option  
--delete, 170  
--help, 170  
--json, 170  
--output, 170  
-h, 170

convert\_feature\_json\_file() (in module *rsmtool.convert\_feature\_json*), 127

convert\_ipynb\_to\_html() (*rsmtool.reporter.Reporter* static method), 156

copy() (*rsmtool.configuration\_parser.Configuration* method), 120

copy() (*rsmtool.container.DataContainer* method), 125

correlation\_helper() (*rsmtool.analyzer.Analyzer* static method), 110

create\_comparison\_report() (*rsmtool.reporter.Reporter* method), 156

create\_explanation\_report() (*rsmtool.reporter.Reporter* method), 157

create\_fake\_skill\_learner() (*rsmtool.modeler.Modeler* method), 129

create\_report() (*rsmtool.reporter.Reporter* method), 157

create\_summary\_report() (*rsmtool.reporter.Reporter* method), 157

## D

DataContainer (class in *rsmtool.container*), 124

DataReader (class in *rsmtool.reader*), 153

DatasetDict (class in *rsmtool.container*), 127

DataWriter (class in *rsmtool.writer*), 168

determine\_chosen\_sections() (*rsmtool.reporter.Reporter* method), 157

difference\_of\_standardized\_means() (in module *rsmtool.utils.metrics*), 163

`drop()` (*rsmtool.container.DataContainer* method), 125

## F

`fast_predict()` (in module *rsmtool*), 101

`FeaturePreprocessor` (class in *rsmtool.preprocessor*), 139

`FeatureSpecsProcessor` (class in *rsmtool.preprocessor*), 149

`FeatureSubsetProcessor` (class in *rsmtool.preprocessor*), 151

`FeatureTransformer` (class in *rsmtool.transformer*), 159

`filter_data()` (*rsmtool.preprocessor.FeaturePreprocessor* method), 139

`filter_metrics()` (*rsmtool.analyzer.Analyzer* method), 111

`filter_on_column()` (*rsmtool.preprocessor.FeaturePreprocessor* method), 141

`filter_on_flag_columns()` (*rsmtool.preprocessor.FeaturePreprocessor* method), 141

`find_feature_sign()` (*rsmtool.preprocessor.FeatureSpecsProcessor* method), 150

`find_feature_transform()` (*rsmtool.transformer.FeatureTransformer* method), 161

`frame` (*rsmtool.container.DatasetDict* attribute), 127

## G

`generate()` (*rsmtool.utils.commandline.ConfigurationGenerator* method), 163

`generate_default_specs()` (*rsmtool.preprocessor.FeatureSpecsProcessor* method), 150

`generate_explanation()` (in module *rsmtool*), 103

`generate_feature_names()` (*rsmtool.preprocessor.FeaturePreprocessor* method), 142

`generate_specs()` (*rsmtool.preprocessor.FeatureSpecsProcessor* method), 150

`get()` (*rsmtool.configuration\_parser.Configuration* method), 120

`get_coefficients()` (*rsmtool.modeler.Modeler* method), 129

`get_default_converter()` (*rsmtool.configuration\_parser.Configuration* method), 120

`get_fairness_analyses()` (in module *rsmtool.fairness\_utils*), 128

`get_feature_names()` (*rsmtool.modeler.Modeler* method), 129

`get_frame()` (*rsmtool.container.DataContainer* method), 125

`get_frames()` (*rsmtool.container.DataContainer* method), 125

`get_intercept()` (*rsmtool.modeler.Modeler* method), 129

`get_names_and_paths()` (*rsmtool.configuration\_parser.Configuration* method), 120

`get_ordered_notebook_files()` (*rsmtool.reporter.Reporter* method), 158

`get_path()` (*rsmtool.container.DataContainer* method), 126

`get_rater_error_variance()` (*rsmtool.configuration\_parser.Configuration* method), 121

`get_section_file_map()` (*rsmtool.reporter.Reporter* method), 158

`get_thumbnail_as_html()` (in module *rsmtool.utils.notebook*), 166

`get_trim_min_max_tolerance()` (*rsmtool.configuration\_parser.Configuration* method), 121

## H

`html_file`  
`render_notebook` command line option, 170

## I

`ipynb_file`  
`render_notebook` command line option, 170

`items()` (*rsmtool.configuration\_parser.Configuration* method), 121

`items()` (*rsmtool.container.DataContainer* method), 126

## K

`keys()` (*rsmtool.configuration\_parser.Configuration* method), 121

`keys()` (*rsmtool.container.DataContainer* method), 126

## L

`load_from_file()` (*rsmtool.modeler.Modeler* class method), 129

`load_from_learner()` (*rsmtool.modeler.Modeler* class method), 130

`load_rsmtool_output()` (*rsmtool.comparer.Comparer* method), 117

`locate_custom_sections()` (*rsmtool.reporter.Reporter* static method), 159

`locate_files()` (*rsmtool.reader.DataReader* static method), 153

`logger` (*rsmtool.configuration\_parser.ConfigurationParser* attribute), 122

## M

`make_summary_stat_df()` (*rsm-tool.comparer.Comparer static method*), 118  
`merge_notebooks()` (*rsmtool.reporter.Reporter static method*), 159  
`metrics_helper()` (*rsmtool.analyzer.Analyzer static method*), 112  
`model_fit_to_dataframe()` (*rsm-tool.modeler.Modeler static method*), 130  
`Modeler` (*class in rsmtool.modeler*), 128  
`module`  
     *rsmtool.analyzer*, 104  
     *rsmtool.comparer*, 116  
     *rsmtool.configuration\_parser*, 118  
     *rsmtool.container*, 124  
     *rsmtool.modeler*, 128  
     *rsmtool.preprocessor*, 138  
     *rsmtool.reader*, 153  
     *rsmtool.reporter*, 155  
     *rsmtool.transformer*, 159  
     *rsmtool.writer*, 168

## N

`name` (*rsmtool.container.DatasetDict attribute*), 127

## O

`ols_coefficients_to_dataframe()` (*rsm-tool.modeler.Modeler static method*), 130  
`output_dir`  
     *rsmcompare* command line option, 63  
     *rsmeval* command line option, 43  
     *rsmpredict* command line option, 84  
     *rsmsummarize* command line option, 70  
     *rsmtool* command line option, 19  
     *rsmxval* command line option, 78  
`output_file`  
     *rsmpredict* command line option, 56

## P

`parse()` (*rsmtool.configuration\_parser.ConfigurationParser method*), 122  
`parse_json_with_comments()` (*in module rsm-tool.utils.files*), 167  
`partial_correlations()` (*in module rsm-tool.utils.metrics*), 164  
`path` (*rsmtool.container.DatasetDict attribute*), 127  
`pop()` (*rsmtool.configuration\_parser.Configuration method*), 121  
`predict()` (*rsmtool.modeler.Modeler method*), 130  
`predict_train_and_test()` (*rsm-tool.modeler.Modeler method*), 131

`preprocess_feature()` (*rsm-tool.preprocessor.FeaturePreprocessor method*), 142  
`preprocess_features()` (*rsm-tool.preprocessor.FeaturePreprocessor method*), 143  
`preprocess_new_data()` (*rsm-tool.preprocessor.FeaturePreprocessor method*), 143  
`prmse_true()` (*in module rsmtool.utils.prmse*), 152  
`process_config()` (*rsm-tool.configuration\_parser.ConfigurationParser class method*), 123  
`process_confusion_matrix()` (*rsm-tool.comparer.Comparer static method*), 118  
`process_data()` (*rsm-tool.preprocessor.FeaturePreprocessor method*), 144  
`process_data_rsmeval()` (*rsm-tool.preprocessor.FeaturePreprocessor method*), 145  
`process_data_rsmexplain()` (*rsm-tool.preprocessor.FeaturePreprocessor method*), 145  
`process_data_rsmpredict()` (*rsm-tool.preprocessor.FeaturePreprocessor method*), 145  
`process_data_rsmtool()` (*rsm-tool.preprocessor.FeaturePreprocessor method*), 146  
`process_predictions()` (*rsm-tool.preprocessor.FeaturePreprocessor method*), 147

## Q

`quadratic_weighted_kappa()` (*in module rsm-tool.utils.metrics*), 164

## R

`read()` (*rsmtool.reader.DataReader method*), 154  
`read_from_file()` (*rsmtool.reader.DataReader static method*), 154  
`read_jsonlines()` (*in module rsmtool.reader*), 154  
`remove_outliers()` (*rsm-tool.preprocessor.FeaturePreprocessor static method*), 147  
`remove_outliers_using_truncations()` (*rsm-tool.preprocessor.FeaturePreprocessor method*), 148  
`rename()` (*rsmtool.container.DataContainer method*), 126  
`rename_default_columns()` (*rsm-tool.preprocessor.FeaturePreprocessor*



`method`), 148  
`render_notebook` command line option  
  `--help`, 170  
  `-h`, 170  
  `html_file`, 170  
  `ipynb_file`, 170  
`Reporter` (*class in `rsmtool.reporter`*), 155  
`rsmcompare` command line option  
  `-V`, 63  
  `--help`, 63  
  `--version`, 63  
  `-h`, 63  
  `config_file`, 63  
  `output_dir`, 63  
`rsmeval` command line option  
  `-V`, 44  
  `--force`, 44  
  `--help`, 44  
  `--version`, 44  
  `-f`, 44  
  `-h`, 44  
  `config_file`, 43  
  `output_dir`, 43  
`rsmpredict` command line option  
  `-V`, 57, 84  
  `--features`, 56  
  `--force`, 84  
  `--help`, 56, 84  
  `--version`, 57, 84  
  `-f`, 84  
  `-h`, 56, 84  
  `config_file`, 56, 84  
  `output_dir`, 84  
  `output_file`, 56  
`rsmsummarize` command line option  
  `-V`, 70  
  `--force`, 70  
  `--help`, 70  
  `--version`, 70  
  `-f`, 70  
  `-h`, 70  
  `config_file`, 70  
  `output_dir`, 70  
`rsmtool` command line option  
  `-V`, 19  
  `--force`, 19  
  `--help`, 19  
  `--version`, 19  
  `-f`, 19  
  `-h`, 19  
  `config_file`, 19  
  `output_dir`, 19  
`rsmtool.analyzer`  
  module, 104  
`rsmtool.comparer`  
  module, 116  
`rsmtool.configuration_parser`  
  module, 118  
`rsmtool.container`  
  module, 124  
`rsmtool.modeler`  
  module, 128  
`rsmtool.preprocessor`  
  module, 138  
`rsmtool.reader`  
  module, 153  
`rsmtool.reporter`  
  module, 155  
`rsmtool.transformer`  
  module, 159  
`rsmtool.writer`  
  module, 168  
`rsmxval` command line option  
  `-V`, 78  
  `--help`, 78  
  `--version`, 78  
  `-h`, 78  
  `config_file`, 78  
  `output_dir`, 78  
`run_comparison()` (*in module `rsmtool`*), 99  
`run_data_composition_analyses_for_rsmeval()`  
  (*`rsmtool.analyzer.Analyzer` method*), 114  
`run_data_composition_analyses_for_rsmtool()`  
  (*`rsmtool.analyzer.Analyzer` method*), 114  
`run_evaluation()` (*in module `rsmtool`*), 98  
`run_experiment()` (*in module `rsmtool`*), 98  
`run_prediction_analyses()` (*`rsm-`*  
  *`tool.analyzer.Analyzer` method*), 115  
`run_summary()` (*in module `rsmtool`*), 100  
`run_training_analyses()` (*`rsm-`*  
  *`tool.analyzer.Analyzer` method*), 115  

## S

`save()` (*`rsmtool.configuration_parser.Configuration`  
  *method**), 122  
`save()` (*`rsmtool.modeler.Modeler` method*), 132  
`scale_coefficients()` (*`rsmtool.modeler.Modeler`  
  *method**), 132  
`select_by_subset()` (*`rsm-`*  
  *`tool.preprocessor.FeatureSubsetProcessor`*  
  *method*), 151  
`select_candidates()` (*`rsm-`*  
  *`tool.preprocessor.FeaturePreprocessor`*  
  *method*), 149  
`show_thumbnail()` (*in module `rsmtool.utils.notebook`*),  
  167  
`skill_learner_params_to_dataframe()` (*`rsm-`*  
  *`tool.modeler.Modeler` static method*), 132

- `standardized_mean_difference()` (in module *rsm-tool.utils.metrics*), 165  
`suppress_warnings` (*rsm-tool.utils.commandline.ConfigurationGenerator* attribute), 162
- ## T
- `to_datasets()` (*rsmtool.container.DataContainer* static method), 126  
`to_dict()` (*rsmtool.configuration\_parser.Configuration* method), 122  
`train()` (*rsmtool.modeler.Modeler* method), 132  
`train_built_in_model()` (*rsmtool.modeler.Modeler* method), 133  
`train_equal_weights_lr()` (*rsm-tool.modeler.Modeler* method), 133  
`train_lasso_fixed_lambda()` (*rsm-tool.modeler.Modeler* method), 134  
`train_lasso_fixed_lambda_then_lr()` (*rsm-tool.modeler.Modeler* method), 134  
`train_lasso_fixed_lambda_then_non_negative_lr()` (*rsmtool.modeler.Modeler* method), 134  
`train_linear_regression()` (*rsm-tool.modeler.Modeler* method), 135  
`train_non_negative_lr()` (*rsmtool.modeler.Modeler* method), 135  
`train_non_negative_lr_iterative()` (*rsm-tool.modeler.Modeler* method), 136  
`train_positive_lasso_cv()` (*rsm-tool.modeler.Modeler* method), 136  
`train_positive_lasso_cv_then_lr()` (*rsm-tool.modeler.Modeler* method), 137  
`train_rebalanced_lr()` (*rsmtool.modeler.Modeler* method), 137  
`train_score_weighted_lr()` (*rsm-tool.modeler.Modeler* method), 137  
`train_skill_model()` (*rsmtool.modeler.Modeler* method), 138  
`transform_feature()` (*rsm-tool.transformer.FeatureTransformer* method), 162  
`trim()` (*rsmtool.preprocessor.FeaturePreprocessor* method), 149  
`try_to_load_file()` (in module *rsmtool.reader*), 155
- ## U
- `use_subgroups` (*rsmtool.utils.commandline.ConfigurationGenerator* attribute), 162
- ## V
- `validate_config()` (*rsm-tool.configuration\_parser.ConfigurationParser* class method), 123
- `validate_feature_specs()` (*rsm-tool.preprocessor.FeatureSpecsProcessor* method), 150  
`values()` (*rsmtool.configuration\_parser.Configuration* method), 122  
`values()` (*rsmtool.container.DataContainer* method), 127  
`variance_of_errors()` (in module *rsm-tool.utils.prmse*), 152
- ## W
- `write_experiment_output()` (*rsm-tool.writer.DataWriter* method), 168  
`write_feature_csv()` (*rsmtool.writer.DataWriter* method), 169  
`write_frame_to_file()` (*rsmtool.writer.DataWriter* static method), 169