



Review Center Guide

October 7, 2026

For the most recent version of this document, visit our [documentation website](#).

Table of Contents

1 Review Center	6
1.1 Review Center overview	7
1.2 Review Center workflow	7
1.3 Understanding the integrative learning classifier	8
1.3.1 Support Vector Machine learning (SVM)	8
1.4 Language support in Review Center	9
1.5 Archiving and restoring workspaces with Review Center	9
1.6 Review Center resources	10
2 Creating a Review Center queue	10
2.1 Installing Review Center	10
2.2 Choosing a queue type	11
2.2.1 Saved search queues	11
2.2.2 Prioritized review queues	11
2.3 How document assignment works	12
2.3.1 Batched document assignment	12
2.3.2 Linear document assignment	12
2.3.3 Keeping document families together	12
2.4 Setting up the reviewer group	13
2.4.1 Checking document permissions	13
2.5 Creating required queue fields	13
2.6 Creating the saved search	14
2.6.1 Reducing document clutter	14
2.6.2 Saved search recommendations for prioritized review queues	14
2.7 Creating a new queue or queue template	15
2.8 Editing premade templates	15
2.9 Queue and template settings list	16
2.9.1 Document delivery options	16
2.9.2 Template settings	17
2.9.3 Admin Settings	17
2.9.4 Reviewer Settings	19
3 Monitoring a Review Center queue	20
3.1 Review Center dashboard	20

3.1.1 Queue tab strip	20
3.1.2 Queue Summary section	21
3.1.3 Review Progress section	22
3.2 Charts and tables	23
3.2.1 General charts and tables	24
3.2.2 Prioritized review charts	26
3.2.3 All Documents table	26
3.2.4 Batches table	27
3.3 Preparing or refreshing the queue	27
3.3.1 Checking refresh status	28
3.3.2 Auto-refreshing the queue	28
3.3.3 Reviewer access during refreshes	29
3.3.4 Auto-refreshing in Coverage Mode	29
3.3.5 How queue refreshes affect batching	30
3.3.6 Caching text in prioritized review queues	30
3.4 Starting the queue	30
3.5 Editing queues and other actions	31
3.5.1 Editing recommendations	31
3.5.2 Turning Coverage Mode on and off	32
3.6 Deleting a queue	32
3.7 Fixing a misconfigured queue	33
3.8 Understanding document ranks	33
3.9 Tracking reviewer decisions	33
3.9.1 Using the Documents tab	33
3.9.2 Using the Field Tree	34
3.9.3 Using the Track Document Field Edits by Reviewer application	35
3.10 Moving Review Center templates and queues	35
3.11 Viewing archived Active Learning projects	35
4 Managing batched queues	36
4.1 Assigning batches from the dashboard	36
4.1.1 Reassigning completed batches	36
4.2 Monitoring queues	36
4.3 Adding or removing documents	37
4.4 Deleting batches	37

5 Checking out batches	37
5.1 Viewing queues and batches	38
5.1.1 The queue and batch list	38
5.2 Selecting and assigning batches	38
5.2.1 Checking out already assigned batches	38
5.2.2 Checking out unassigned batches	39
5.3 Batch options	39
6 Reviewing documents using Review Center	40
6.1 Reviewing documents in the queue	40
6.1.1 Checking in batches	41
6.2 Finding previously viewed documents	41
6.3 Queue statistics	42
6.4 Viewing the dashboard	42
6.5 Best practices for Review Center review	42
6.5.1 Coding according to the "four corners" rule	42
6.5.2 Factors that affect Review Center's predictions	43
7 Review validation	44
7.1 Key definitions	44
7.2 Determining when to validate a Prioritized Review queue	44
7.3 Starting a validation queue	45
7.3.1 Choosing the validation settings	45
7.3.2 Inherited settings	46
7.4 Coding in a validation queue	46
7.5 Monitoring a validation queue	46
7.5.1 Viewing formulas for validation statistics	47
7.5.2 Editing a validation queue	47
7.5.3 Releasing unreviewed documents	48
7.5.4 Tracking sampled documents	48
7.6 Accepting or rejecting validation results	49
7.6.1 Manually rejecting validation results	50
7.7 Reviewing validation results	50
7.7.1 Recalculating validation results	50
7.7.2 Viewing results for previous validation queues	51
7.8 How adding or changing documents affects validation	51

7.8.1 Scenarios that require recalculation	51
7.8.2 Scenarios that require a new validation queue	51
8 Review validation statistics	52
8.1 Defining elusion, recall, richness, and precision	52
8.2 Groups used to calculate validation metrics	52
8.2.1 Variables used in the full calculations	53
8.3 How setting a cutoff affects validation statistics	54
8.3.1 High versus low cutoff	55
8.4 Validation metric calculations	55
8.4.1 Elusion rate	55
8.4.2 Recall	56
8.4.3 Richness	57
8.4.4 Precision	57
8.5 How the validation queue works	58
8.6 How validation handles skipped and neutral documents	58
9 Reusing saved models	59
9.1 Saved model	59
9.2 What a saved model contains	59
9.2.1 How predictions work with multiple models	59
9.2.2 Privacy considerations when reusing saved models	59
9.3 Common use cases for saved models	60
9.4 Creating a saved model	60
9.4.1 Creating from queues with linked models	61
9.5 Linking a saved model to a Review Center queue	61
9.5.1 Removing a linked model from a queue	61
9.5.2 How linked models behave with ARM	62
9.6 Managing saved models	62
9.6.1 Copying saved models to another workspace	62
9.6.2 Editing or deleting saved models	63
9.6.3 Saved models from Active Learning	63
10 Review Center security permissions	63
10.1 Feature permissions	63
10.1.1 Checking document permissions	64
10.2 Granular permissions	65

10.2.1 Creating a Review Center queue or template	65
10.2.2 Editing and controlling Review Center queues	65
10.2.3 Deleting a Review Center queue or template	65
10.2.4 Viewing the Review Center dashboard	65
10.2.5 Assigning batches	66
10.2.6 Managing saved models	66
10.2.7 Tracking reviewer decisions from the Documents tab	66
10.2.8 Reviewer permissions	66
10.2.9 Checking document permissions	67
11 Review Center performance baselines	67
11.1 Queue size recommendations	67
11.2 Reviewer count recommendations	67
11.3 Performance testing definitions	68
11.4 Overall build speed	68
11.5 Build time variance with document coding and caching	68
11.6 Build time variance with more documents	69
12 Active Learning application history	69
12.1 Selecting an Active Learning project	69
12.2 Project Statistics section	69
12.3 Manually Selected section	70
12.4 Prioritized Review section	70
12.5 Coverage Review section	71
12.6 Project Validation History section	71
12.7 Model Updates section	72

1 Review Center

Review Center is a review management tool that helps you build custom queues, use AI to prioritize relevant documents, and leverage a rich reporting dashboard to understand the state of your data and track productivity. With streamlined administrative features and flexible AI algorithms, you can tailor the review process to your needs.

Some of Review Center's key features include:


- **Templatization**—set up best-practice structures ahead of time for easy re-use.
- **Customizable queues**—replace batch administration with queues based on saved searches.
- **Powerful AI classifier**—Review Center uses a new integrative learning classifier that provides even greater efficiency than previous AI classifiers.

- **Clear progress reporting**—a rich dashboard features timeline-based visualizations that show relevance rates and progress.

1.1 Review Center overview

Review Center enables administrators to build review queues from any saved search and choose the order in which the documents will be served up to reviewers. These queues can be ordered using either AI-powered relevance predictions, or custom sort conditions chosen by the admin. After the admin starts the queue, reviewers check out documents from a simple interface. The admin manages all queues, reporting, and progress charts from a central dashboard.

For a guided video showing how to use Review Center, watch the [Review Center: Getting Started](#) on-demand training on the Relativity aiR documentation site.



1.2 Review Center workflow

The basic steps to set up Review Center are:

1. Install the application.
2. Create a saved search containing the documents for review.
3. Create any necessary fields.

After setup, create and manage the Review Center queue:

1. Create a new queue.
2. Assign the reviewer group.
3. Start the queue.
4. Review documents.
5. Monitor the queue using the Review Center dashboard.

After the admin enables the queue, reviewers log into a simple screen showing the queues assigned to them. For more detail on the reviewer's experience, see [Reviewing documents using Review Center on page 40](#).

For detailed instructions on setting up Review Center, see [Creating a Review Center queue on page 10](#).

1.3 Understanding the integrative learning classifier

The integrative learning classifier used by Review Center's AI-powered queues is a scalable, secure, and efficient classification service that can support a variety of use cases and documents. It makes connections among concepts and decisions to serve up relevant documents to reviewers as early as possible.

You do not need to create an Analytics index for Review Center queues. Instead, when you prepare or start an AI-driven queue, the classifier automatically runs in the background to manage documents.

1.3.1 Support Vector Machine learning (SVM)


Review Center's classifier uses a form of classification called a Support Vector Machine, or SVM. With an SVM model, the system learns from reviewers as they code and constantly updates its predictions.

When the SVM runs, it takes reviewer's coding decisions and pulls them into a high-dimensional model. This model is divided into two spaces, with the boundary between them referred to as the hyperplane. The SVM puts documents that are coded Relevant on one side of the hyperplane, and it puts documents coded Not Relevant on the other side.

After the model establishes the hyperplane, it takes all documents without a coding decision and maps them onto either side of the hyperplane. It bases their position on the model's current understanding of the difference between relevant and not relevant. Documents that have strong similarities to relevant or non-relevant documents are positioned farther away from the hyperplane, while documents that might belong in either group are mapped closer to the hyperplane in the middle.

1.3.1.1 Document ranks and the hyperplane

A document's rank measures the confidence the model has in a document being relevant or not relevant. In terms of the SVM model, it measures the document's distance from the hyperplane. Rank is measured on a scale from 100 to 0, with 100 indicating a strongly relevant document, and 0 indicating a strongly non-relevant document. Documents that are mapped close to the hyperplane have ranks close to 50.



1.4 Language support in Review Center

Because the integrative learning classifier is language-agnostic, you can use Review Center for documents written in any language. However, the methods Review Center uses to tokenize text, or break it up into individual words, are primarily based on English, Chinese, and Japanese. If Review Center detects Chinese or Japanese, it uses the tokenization method for those languages. For any other text, it uses the English-based tokenization method that relies on spacing and punctuation. This means that languages with similar spacing and punctuation to English typically have good results.

1.5 Archiving and restoring workspaces with Review Center

Workspaces with Review Center installed can be archived and restored using the ARM application.

When archiving, check **Include Extended Workspace Data** under Extended Workspace Data Options. If this option is not checked during the archive process, the Review Center features in the restored workspace will not be fully functional. If this happens, you will need to manually reinstall Review Center in the restored workspace.

For more information on using ARM, see ARM on the Relativity documentation site.

1.6 Review Center resources

In addition to the help site resources, you can learn more about Review Center through our trainings, knowledge base articles, and white papers.

On-demand training:

- [Review Center Foundations: Managing Prioritized Review Projects](#)

Knowledge base articles:

- [Review Center - Multiple queues powered by the same classifier](#)
- [Review Center workflow for queues exceeding 10M documents](#)
- [Review Center - Legacy Batching vs. Batching in Review Center](#)

2 Creating a Review Center queue

Review Center queues are flexible, customizable, and can be used for any stage of review. You can also create templates for common workflows, which shortens the setup time for a new queue to only a few clicks. These queue templates can be saved as part of workspace templates, making it easy to re-use them for other cases. You can also use the AI training from previous queues to improve the relevance predictions in new queues.

Even after creating a queue, you can still edit the settings or add new documents without interrupting reviewers.

2.1 Installing Review Center

Review Center is available as a secured application from the Application Library.

To install it:

1. Navigate to the **Relativity Applications** tab in your workspace.
2. Select **Install from application library**.
3. Select the **ReviewCenter** application.
4. Click **Install**.

After installation completes, the following tabs will appear in your workspace:

- **Review Library**—manage queue templates.
- **Review Center**—create and manage queues and templates and view the dashboard.
- **Review Queues**—review documents using queues.
- **Saved Models**—manage saved copies of the relevance training from previous queues.

For more information on installing applications, see Relativity Applications in the Admin guide.

2.2 Choosing a queue type

Review Center offers two types of review queues. Based on the needs of your project, you can set up review queues that either focus on custom-sorted sets of documents, or focus on documents that the AI classifier predicts as relevant.

2.2.1 Saved search queues

Saved search queues tie your queue to a saved search. You can use saved searches to group documents based on nearly any criteria, including documents from any existing Review Center queue. With this queue type, documents are served up to reviewers based on the sort method used for the saved search. If the saved search does not have a sort method selected, documents will be served up based on Artifact ID.

2.2.1.1 Batching a saved search queue

When you create a saved search queue, you have two options:

- **Batches**—the document set is divided into groups, then checked out to reviewers group by group. These groups are called batches.
- **Linear**—documents are automatically checked out to reviewers a few at a time, according to the sort order of the saved search.

If you create a batched queue, the administrator can either assign batches of documents to specific reviewers, or let reviewers self-assign batches. You can also group similar documents together into a batch by setting the batch unit field. This takes more setup than a linear review, but it offers more control over how documents are checked in and out.

For more information, see [Managing batched queues on page 36](#).

2.2.2 Prioritized review queues

Prioritized review queues are also based on a saved search, but instead of serving up documents based on their sort order, they use the AI classifier to serve up documents that it predicts as relevant. These relevance rankings are stored in the Rank Output field, and the ranks automatically update every time the queue refreshes.

The AI classifier uses either the extracted text of documents, or another text field you select to make its predictions. Even if other fields are returned in the saved search, they will not affect the results.

If you choose a prioritized review queue, we recommend coding at least two non-empty documents in your data source before preparing or starting the queue: one with the positive choice on your review field, and one with the negative choice. This gives the AI classifier the information it needs to start making its predictions. The more documents are coded, the more accurate the classifier's predictions will be.

If you do not have any coding completed, you can start the prioritized review queue without any coding. The classifier model won't build until at least 50 documents have been coded, with at least one coded positive and one coded negative. After you reach 50 coded documents, your ranks will update upon the next auto-refresh or manual refresh. If you need it to build sooner, you can manually trigger a queue refresh at any point after at least one document has been coded positive and one has been coded negative.

If your prioritized review queue is similar to an older queue, you can copy the relevance training from the older queue and link it to the new queue to jump-start predictions. For more information, see [Reusing saved models on page 59](#).

2.2.2.1 Including random documents in the queue

When you set up a prioritized review queue, you have the option to serve up randomly chosen documents alongside documents that are predicted relevant. This gives the AI classifier a broader variety of coding decisions to learn from, which improves its predictions in the early stages of a review. Having reviewers code a selection of random documents helps the classifier identify a wider range of relevant topics and prevents it from focusing on a limited subject area.

Under the queue setting **Include Random Items**, you can choose to include random documents as up to 20% of the total documents served to reviewers. You can change this setting at any time. We recommend including a high percent of random items during the early stages of review.

Note: If you have Coverage Mode turned on, this overrides the Include Random Items setting.

2.2.2.2 Using Coverage Mode

When Coverage Mode is turned on for a prioritized review queue, the queue switches away from serving up the highest-ranking documents. Instead, it serves up documents that are better for training the model. These are documents with scores near 50, which usually have different content and topics from documents that the model has previously seen. Labeling these helps the model learn from a wider variety of documents and become more confident quickly.

When in Coverage Mode, the AI classifier sorts all documents by their scores' distance from 50, but limits and spreads out the number of exactly 50-ranked documents. This intermixing diversifies the group of documents and lowers the chance of duplicates. The classifier then serves up these sorted documents to reviewers until the next refresh. After each refresh in Coverage Mode, it re-sorts the documents. Coverage Mode also overrides the Include Random Items setting.

You can turn the **Coverage Mode** setting on or off at any time during a review. For instructions, see [Turning Coverage Mode on and off on page 32](#).

Note: Whenever you turn Coverage Mode on or off, manually refresh the queue. This updates the document sorting for reviewers. For more information, see [Turning Coverage Mode on and off on page 32](#).

2.3 How document assignment works

Document assignment depends on whether a queue is batched or linear. Prioritized review queues use linear document assignment.

2.3.1 Batched document assignment

If a saved search queue is set up for batching, all documents in a batch are assigned to the reviewer at a time. The administrator can either assign batches to reviewers from the dashboard, or reviewers can self-assign batches from the Review Queues tab. Reviewers who have the Assign Batches permission can also reassign batches to other reviewers.

After a batch is checked out, the reviewer can check the entire batch in as either Pending, or Completed. The reviewer does not check in individual documents.

2.3.2 Linear document assignment

If a relational field is not set for the queue, five documents are checked out to each active reviewer at a time. As the reviewer saves their progress on those documents, more are checked out as needed.

For example, documents 1 through 5 are assigned to the first reviewer who starts review. If a second reviewer logs in immediately after, documents 6 through 10 are assigned to the second reviewer. As the first reviewer completes their work, documents 11 through 15 are assigned to them, and so on.

If a relational field is set for the queue, then the entire relational group for a document will also be checked out to that document's reviewer. This may mean that a reviewer has more than 5 documents checked out at a time. For more information, see [Keeping document families together below](#). Currently, the number of documents checked out is not customizable.

2.3.3 Keeping document families together

All Review Center queues have the option of setting a relational field. If this is set, the whole relational group of documents present in the queue will be checked out to the same reviewer. This keeps families, email threads, or other

relational groupings together in one queue.

When a relational field is set, it takes priority over the sort method and document rank. For example, if you sort a linear saved search queue by size and set the relational field to Family Group, then the entire family of the largest document will be checked out to the first reviewer, even if it contains small documents. Likewise, if you set the relational field to Family Group for a prioritized review queue, the entire family of the highest ranked document will be checked out to the first reviewer, even if it contains low-ranked documents. Within that family, documents will be served up based on the sort specified in the relational view.

For batched queues, the relational field also takes priority over the maximum batch size. For example, if the Maximum Batch Size is 100 and a document family contains 120 documents, those documents will be put in a batch of 120.

If you plan to code families in the related items pane as part of the reviewer workflow, we recommend that you do not include families in your queue. Otherwise, as you code documents in the related items pane, the coded family documents will still be served to reviewers.

Note: If you set a relational field on a template or queue, set the same field in the Related Items drop-down menu of the saved search Conditions tab. Only relational group members returned by the saved search will be included in the queue. For more information, see [Creating or editing a saved search in the Searching guide](#).

2.4 Setting up the reviewer group

To give reviewers access to a queue, set up a reviewer group. You can either create a brand new group, or modify the permissions for an existing user group. You can assign multiple user groups to the same queue.

To set up a reviewer group:

1. Decide which user group or groups should contain the reviewers for the queue. For information on creating and editing groups, see [Groups in the Admin guide](#).
2. Add each group to the workspace.
3. Assign each reviewer group the following permissions:

Object Security	Tab Visibility
• Document - View, Edit • Review Center Queue - View	• Review Queues

4. Add the reviewers to each group.

For more information about permissions, see [Review Center security permissions on page 63](#).

2.4.1 Checking document permissions

Make sure the reviewer groups have permissions to the documents they need to review. Reviewers will only see documents they have access to. If any documents in the queue are in a secured documents folder, and a reviewer does not have permissions for it, those documents will not be checked out to the reviewer.

If a reviewer sees a message that there are no more documents to review in a queue, but there are uncoded documents left, check the document permissions.

For more information on document security, see [Relativity object security in the Admin guide](#).

2.5 Creating required queue fields

Before creating a prioritized review queue, create the following fields:

- **Review field**—a single-choice field that serves as the coding field for your queue. This field should have at least one positive choice and one negative choice. Any other choices will be considered neutral.
- **Rank Output**—a decimal field that will hold the document ranks. Each prioritized review queue needs a separate Rank Output field on the Document object.

Note: If a field has two colons (::) in the name, this is called a reflected field. Reflected fields typically link two objects, and they cannot be used as the Rank Output field.

If you are creating a saved search queue, you do not need a Rank Output field, and the review field is optional.

For more information about creating new fields, see [Fields in the Admin guide](#).

2.6 Creating the saved search

No matter which queue type you choose, you will need a saved search to serve as the data source. This saved search contains all documents that will be available to reviewers.

When creating the saved search:

- The saved search must be public.
- If you want to make sure documents are not added or removed from the queue, make the saved search static.
- If you set a relational field on a template or queue, set the same field in the Related Items drop-down menu of the saved search Conditions tab. Only relational group members returned by the saved search will be included in the queue.

For full instructions on how to create a saved search, see [Creating or editing a saved search in the Searching guide](#).

2.6.1 Reducing document clutter

To streamline the review process, we recommend taking steps to reduce document clutter in your saved search. This can include any combination of the following:

- Removing out-of-scope documents
- Focusing on inclusive emails
- Running deduplication

2.6.2 Saved search recommendations for prioritized review queues

When setting up a saved search for a prioritized review queue, we recommend excluding the following file types:

- Compressed files, such as ZIP or RAR files
- System files
- Excel files with mostly numbers
- Image files
- CAD drawings
- Maps

We also recommend:

- Remove documents with poor text or no text. By default, the AI model uses the extracted text field to make its predictions. If the documents have text in a different field instead, you can select that field during queue setup.

- Remove files over 30MB when possible. These can slow down queue refresh times.
- Use a saved search that runs quickly. Slow or overly complex searches can slow or stall the queue refresh.

2.6.2.1 Size limits for prioritized review queues

For prioritized review queues, we recommend a maximum of 10 million documents per queue. This assumes an average extracted text size of 30KB. If your documents are larger than 30KB on average, limit the number of documents so that the combined total is less than 300GB of extracted text.


For more information on queue performance, see [Review Center performance baselines on page 67](#).

2.7 Creating a new queue or queue template

From the Review Center dashboard, you can create a new queue. You can also choose to create a queue template when you first create a queue. This makes it easier to re-use the same settings for future queues.

To create a new queue or queue template:

1. Navigate to the **Review Center** tab.
2. Click the **Add Queue** button.



3. Select a queue type, then click **Continue**.
4. Edit the queue settings. For a full list, see [Queue and template settings list on the next page](#).
 1. If you want to create a template out of your current settings, set the **Save as Template** field to **On**. For more information, see [Template settings on page 17](#).
 2. If you want to use an existing template, choose one under **Select Template**. For more information, see [Template settings on page 17](#).
5. Click **Create Queue**.
The new queue appears as a tab in the banner at the top of the dashboard. If you also created a template, the template now appears in the Review Library list.

All queue settings can be edited after creating the queue.

For information on starting, managing, and deleting queues, see [Monitoring a Review Center queue on page 20](#).

Note: We recommend creating queues from the Review Center tab. You can still create some queue types from the Review Library tab, but it does not include options for batched queues.

2.8 Editing premade templates

The Review Center application comes with several premade queue templates to choose from. These are designed for common tasks such as image review, searching for privileged documents, and comparing reviewer coding decisions to the AI model's decisions.

Currently, the premade templates are:

1. **New Blank Queue**—use this template to start a new queue from scratch. Change any settings, including the queue type, to fit your needs.

2. **Prioritized Review**—use this template for relevance review projects. Documents will be served in order based on how strongly the AI classifier predicts that each document is relevant.
3. **Model/Reviewer QC**—use this template to compare reviewer coding decisions against the AI model. Documents will be served in order based on the saved search sort.
4. **Images**—use this template to review images that do not have good extracted text. Documents will be served in order based on the saved search sort.
5. **Spreadsheets**—use this template to review spreadsheets such as Excel files. Documents will be served in order based on the saved search sort.
6. **Potentially Privileged**—use this template to review documents that reviewers marked as being potentially privileged. Documents will be served in order based on the saved search sort.
7. **Needs Further Review**—use this template to review documents that reviewers marked as needing further review. Documents will be served in order based on the saved search sort.
8. **Docs with Ranks Below Zero**—use this template to review documents with ranks below zero. These are documents that could not be classified. Documents will be served in order based on the saved search sort.

You can use these templates as-is. However, we recommend reviewing the settings and tailoring them to your needs.

To edit a queue template:

1. Navigate to the **Review Library** tab.
2. Click the **Edit** icon next to the template you want to use.
3. Review and modify the template settings. For a full list, see [Queue and template settings list below](#).
4. Click **Save**.

2.9 Queue and template settings list

When you create or edit a Review Center queue or template, first choose a queue type, then adjust individual settings. The queue type determines which settings are available, as well as what order those settings appear in.

Note: We recommend creating queues from the Review Center tab. You can still create some queue types from the Review Library tab, but it does not include options for batched queues.

2.9.1 Document delivery options

Each queue type uses a different method of serving documents to the reviewer:

- **Prioritized**—these queues use AI to rank documents by predicted relevance, then serve them to reviewers.
- **Saved Search Queues**—these queues serve up documents either according to batch, or to the sort order of the saved search.
 - **Batches**—the document set is divided into groups, then checked out to reviewers group by group.
 - **Linear**—documents are automatically checked out to reviewers a few at a time, according to the sort order of the saved search.

For more information on each type, see [Choosing a queue type on page 11](#).

2.9.2 Template settings

After you select a queue type, you have the option of selecting a premade template. You can also choose whether to save the current queue as a new template.

At the top of the modal:

- **Select Template**—choose a previously saved template to autofill some of the queue settings. To see the descriptions of each template, click **View Full Details**.

After you select a template to use, review the default values filled in by the template. Change any values that should be different for this queue. In particular, check that the Reviewer Groups match the users you want reviewing the new queue.

Note: After the queue settings have been filled from a template, the two of them are no longer connected. You can edit the queue without affecting the template, or vice versa.

At the bottom of the modal:

- **Save as Template**—whether these queue settings will be used as a template for other queues.
 - If you are creating both a queue and a template, toggle this **On**.
 - If you are creating only a queue, toggle this **Off**.
- **Template Description** (templates only)—enter notes about the template such as its intended use, comments about field settings, and so on.

The Save as Template setting is only available when creating a brand new queue. If you toggle Save as Template on, both a regular queue and a template are created from the settings. Queue templates can be saved as part of your workspace template.

2.9.3 Admin Settings

The admin settings are:

- **Queue Name**—the queue name reviewers will see.
- **Data Source**—select the saved search that contains the documents for your queue. For more information, see [Creating the saved search on page 14](#).
- **Use Case**—select the use case that best describes your queue. If none of them fit, choose **Other**.
- **Rank Output Field** (Prioritized Review only)—select the decimal field you created to hold the document rank scores.
- **Text Precedence** (Prioritized Review only)—select the text field(s) you want Review Center to analyze. By default, this is set to the Extracted Text field.
 - You can choose up to 5 long text fields to analyze.
 - Review Center analyzes text from the first field in the list that is not empty. It does not combine text or analyze multiple fields.
- **Review Field**—select the single choice field you created for review. For a Prioritized Review queue, this field must have two or more choices. For a Saved Search queue, this field and the choices are optional.

- **Positive Choice**—select the choice that represents the positive or responsive designation.
- **Negative Choice**—Select the choice that represents the negative or non-responsive designation.

Note: Any remaining choices are considered neutral.

- **Positive Cutoff** (Prioritized Review only)—on a scale of 0 to 100, enter the document rank that will be the dividing line between positive and negative documents. All documents ranked at or above this number will be predicted positive, and all documents ranked below it will be predicted negative. By default, the cutoff is set at 50.
- **Relational Field**—select a relational field for grouping documents in the queue. This makes reviewers receive related documents together, such as members of the same document family.

Note: If you set a relational field on a template or queue, set the same field in the Related Items drop-down of the saved search Conditions tab. Only relational group members returned by the saved search will be included in the queue. For more information, see [Creating or editing a saved search in the Searching guide](#).

- **Include Random Items** (Prioritized Review only)—select what percentage of random documents to serve to reviewers. For more information, see [Including random documents in the queue on page 11](#).
- **Coverage Mode** (Prioritized Review only)—controls whether the queue serves up the highest-ranking documents first, or documents with middle ranks. For more information, see [Using Coverage Mode on page 12](#).
 - Toggle this **On** to serve up mid-rank documents first. This is useful for training the model as quickly as possible. You can turn this mode off at any time.
 - Toggle this **Off** to serve up highest-rank documents first. This is the default setting for finding relevant documents quickly.
- **Auto Batch** (Batches only)—whether Review Center automatically creates new batches during an auto-refresh. This groups together new documents added to the saved search. It does not re-group documents that were already batched.
 - Toggle this **On** to automatically group newly added documents into batches during an auto-refresh. This also toggles the Auto Refresh setting to **On**.
 - Toggle this **Off** to only create new batches during a manual refresh.
- **Minimum Batch Size** (Batches only)—the minimum number of documents to be batched together during an auto-refresh.
 - This option only appears if Auto Batch is enabled.
 - During a manual refresh, you can choose whether to create a batch smaller than the minimum. For more information, see [How queue refreshes affect batching on page 30](#).
- **Maximum Batch Size** (Batches only)—the maximum number of documents to be batched together.
 - If the Relational Field for the queue is set, some batches may be larger in order to include a whole document family. For example, if the Maximum Batch Size is 100, and a document family contains 120 documents, those documents will be put in a batch of 120.
 - If the Batch Unit Field is set, groupings based on that field will stay within the Maximum Batch Size. For example, if the Batch Unit Field is set to group custodian documents together, but the Maximum Batch Size is set to 100, the batches will not contain more than 100 documents at a time. This stays the same regardless of how many documents each custodian has.
- **Batch Unit Field** (Batches only)—the field Review Center uses to group similar documents together. Documents with different values will not share a batch unless this is overridden by the Relational Field. Currently, this accepts Single Choice and Single Object fields.

2.9.3.1 Additional Admin Settings

The additional admin settings are:

- **Auto Refresh**—controls whether the queue automatically refreshes after coding activity in any queue. This refresh includes re-running the saved search and checking for outside-coded documents. For prioritized review queues, this also re-trains the classifier with the latest coding and re-ranks documents in order of predicted relevance. For more information, see [Auto-refreshing the queue on page 28](#).
 - Toggle this **On** to refresh the queue automatically.
 - Toggle this **Off** to prevent automatic refreshes. You will still be able to manually trigger refreshes using the dashboard.
- **Allow Coded in Review** (Saved Search only)—controls whether documents coded outside of the queue will still be served up in the queue.
 - Toggle this **On** to allow outside-coded documents to be served up.
 - Toggle this **Off** to exclude outside-coded documents from being served up. These are found and removed during queue refreshes and every time a reviewer checks out a document.

Note: Prioritized review queues use outside-coded documents to train their predictions, but they only show them to reviewers if the Relational Field is set. For example, if the relational field is set to Family Group and some members of a document family are already coded, those will be served up to reviewers along with their family.

- **Anonymize Reviewer Names**—controls whether reviewer names appear on the dashboard charts and tables. This setting does not anonymize coding fields added to views or saved searches.
 - Toggle this **On** to hide reviewer names on the dashboard and list them as Anonymous 1, Anonymous 2, and so on. These labels are assigned randomly every session. For example, a reviewer who is listed as Anonymous 1 the first day might be listed as Anonymous 4 another day.
 - Toggle this **Off** to show reviewer names on the dashboard.
- **Saved Model** (Prioritized Review only)—select a saved model that contains relevance training from a previous queue. For more information, see [Reusing saved models on page 59](#).
- **Queue Label**—create and choose organizational labels that will apply to this queue or to queues created from this template. Some label ideas include First Level Review, Second Level Review, or Quality Control. For more information, see [Filtering the queue tab strip on page 21](#).
- **Email Notification Recipients**—enter email addresses to receive notifications about the queue status. These emails tell users when a manually-triggered queue preparation completes, a queue is empty, or a queue encounters an error while populating. To enter multiple email addresses, separate them with a semicolon. Do not include a space.

2.9.4 Reviewer Settings

The reviewer-related settings are:

- **Reviewer Groups**—select the permissions groups that include the reviewers for the queue.
- **Document Panel View**—select a view to control which fields reviewers see in the Documents panel of the Viewer. If you do not choose a view, this defaults to the lowest ordered view the reviewer has permission to access.
 - This panel shows reviewers a list of documents they previously reviewed in their queue. For more information, see [Finding previously viewed documents on page 41](#).

- **Coding Layout**—select the coding layout that you want reviewers to see by default when they enter the queue.
 - You can choose from the first 100 layouts in the workspace.
 - If you do not choose a layout, this defaults to the lowest ordered layout the reviewer has permission to access.

2.9.4.1 Additional Reviewer Settings

Select which statistics you want reviewers to see on the Review Queues tab:

- **Total Remaining Documents in Queue**—the total number of uncoded documents in the queue or batch, across all reviewers. This appears for reviewers as **Total Remaining**.
- **Total Documents in Queue**—the total number of documents in the queue or batch. This appears for reviewers as **Total Docs**.
- **Reviewed Documents Today by Reviewer**—the number of documents the reviewer reviewed in the queue or batch today. This appears for reviewers as **Reviewed Today**.
- **Reviewed Documents Total by Reviewer**—the total number of documents the reviewer has reviewed in the queue or batch. This appears for reviewers as **Reviewed Total**.

3 Monitoring a Review Center queue

The Review Center dashboard provides a centralized location to track, manage, and edit all Review Center queues. In addition, you can track reviewer coding decisions through a variety of methods.


3.1 Review Center dashboard

After creating a queue, navigate to the **Review Center** tab. This tab contains a dashboard showing all queues, their statistics, and controls related to queue progress.

The Review Center dashboard contains the following sections.

3.1.1 Queue tab strip

The queue tab strip contains a tab for each queue that has been created. To make the dashboard show details for a queue, click on its name in the tab strip.



Below the queue name, each queue shows its status. The possible statuses are:

- **Not Started**—the queue has not been prepared or started.
- **Preparing**—the queue is refreshing the saved search for the first time. If it is a prioritized review queue, this also trains the classifier.
- **Prepared**—the queue has finished preparing for the first time, but it has not been started. It may or may not have a reviewer group assigned.
- **Starting**—the admin has started the queue, and the queue is becoming active for reviewers. During this phase, the queue also refreshes the saved search and retrains the classifier if needed.
- **Active**—the queue has started, and reviewers can start reviewing.

- **Paused**—the admin has paused the queue.
- **Canceling**—the admin has canceled the process of starting or refreshing the queue.
- **Complete**—the admin has marked the queue as complete, and it is no longer available to reviewers.
- **Errored**—an error occurred. When this happens, the error details will appear in a banner at the top of the dashboard.
- **Ready for Validation**—a linked validation queue has been created, but not started.
- **Validation Pending**—all documents in the validation queue have been reviewed, and it's ready for you to accept or reject the results.

In addition, if any of the statuses have the word "Validation" added to them (such as "Validation Paused"), this means the status applies to a linked validation queue. For more information, see [Review validation on page 44](#).

At the right of the strip, the Add Queue button lets you quickly create new queues. For instructions, see [Creating a new queue or queue template on page 15](#).

3.1.1.1 Filtering the queue tab strip

If you have a large number of queues, you can filter them according to their assigned labels in the Queue Label field.

To filter the queue tab strip:

1. Click into the search bar above the queue tab strip.
A drop-down list of labels appears.
2. Select the labels you want to filter by. You can also type in the search bar to narrow the list, then press Enter to select or deselect a label.
3. Close the list using the arrow at the right end of the bar.
The queue tab strip now only shows queues whose labels are listed in the search bar. If several labels are listed, queues that match any one of them will appear.


The queue tab filters only apply to the tab strip. They do not affect any of the charts or statistics on the rest of the page.

If you have Active Learning installed in your workspace, an icon appears for the Active Learning History tab. For more information, see [Viewing archived Active Learning projects on page 35](#).

3.1.2 Queue Summary section

The Queue Summary section shows the reviewer group, saved search, coding fields, and controls for actions such as pausing or refreshing the queue. The "<X> Active" statistic under the reviewer group shows how many reviewers currently have documents checked out to them. Additionally, clicking on the saved search name or the coding field name takes you to that saved search or field.

To view all settings for the current queue, click on the arrow symbol on the left side. This expands the Queue Summary panel and shows the detailed setting list.



The Queue Summary section holds controls for common actions. For detailed information on each, see:


- [Preparing or refreshing the queue on page 27](#)
- [Starting the queue on page 30](#)
- [Editing queues and other actions on page 31](#)

3.1.3 Review Progress section

The Review Progress section shows statistics for the current queue's progress.

Review Progress ▾									
Total Docs ?	Docs Coded ?	Responsive ?	Not Responsive ?	Neutral ?	Relevance Rate ?	Uncoded ?	Skipped ?	Predicted Responsi... ?	Predicted Not Respons... ?
15,100	1,484	988	438	23	66.6%	13,616	35	1,424	12,227
100.0%	9.8%	66.6%	29.5%	1.5%	75.4%	90.2%	0.3%	10.5%	89.8%

By default, the section shows a set of statistics that are calculated for all documents in the queue. By clicking the triangle next to the section name, you can select another view.



3.1.3.1 Review Progress view

The default Review Progress view shows statistics for all documents in the queue's data source. If a document has been coded more than once, it counts the most recent coding decision.

The Review Progress statistics are:

- **Total Docs**—the total number of documents currently in the queue's data source. To be counted, the queue must have been prepared or refreshed after the documents were added or removed. The "100%" in smaller print underneath it indicates that this is the total document set.
- **Docs Coded**—the number of documents in the data source that have a value in the review field. This includes documents coded outside the queue. The smaller percentage underneath it shows the percentage of Docs Coded divided by Total Docs.
- **<Positive Choice>**—the number of documents coded with the positive choice on the review field. This includes documents coded outside the queue. The smaller percentage underneath it shows the percentage of <Positive Choice> divided by Docs Coded.
- **<Negative Choice>**—the number of documents coded with the negative choice on the review field. This includes documents coded outside the queue. The smaller percentage underneath it shows the percentage of <Negative Choice> divided by Docs Coded.
- **Neutral**—the number of documents coded with a neutral choice on the review field. This includes documents coded outside the queue. The smaller percentage underneath it shows the percentage of Neutral documents divided by all Docs Coded.
- **Relevance Rate**—the total percentage of documents coded positive. This is calculated by counting the number of documents coded positive, then dividing it by the total number of coded, non-skipped documents. The bold percentage shows the relevance rate including documents coded either inside or outside of the queue, while the smaller percentage underneath it shows the relevance rate only for documents coded inside the queue.
- **Uncoded**—the number of documents in the data source with no value in the review field. This includes documents that were skipped or had their coding decision removed. The smaller percentage underneath it shows the percentage of Uncoded documents divided by Total Docs.

- **Skipped**—the number of documents that were skipped within the queue. The smaller percentage underneath it shows the percentage of Skipped documents divided by all Uncoded documents.
- **Predicted <Positive Choice>** (Prioritized Review only)—the number of documents in the data source with no review field value and a relevance rank greater than or equal to the Positive Cutoff. The smaller percentage underneath it shows the percentage of Predicted <Positive Choice> documents divided by all Uncoded documents.
- **Predicted <Negative Choice>** (Prioritized Review only)—the number of documents in the data source with no review field value and a relevance rank less than the Positive Cutoff. The smaller percentage underneath it shows the percentage of Predicted <Negative Choice> documents divided by all Uncoded documents.

Note: The Predicted <Positive Choice> and Predicted <Negative Choice> fields only show their predictions after 50 or more documents have been coded.

3.1.3.2 Documents Coded Outside Queue view


If you select Documents Coded Outside Queue from the Review Progress drop-down, this shows an alternate view. These statistics count documents that are part of the queue's saved search, but that were coded through some means other than the selected Review Center queue.

The Documents Coded Outside Queue statistics are:

- **Docs Coded**—the number of documents in the data source that were coded outside of the queue. The smaller percentage underneath it shows the percentage of documents coded outside the queue divided by all documents coded.
- **<Positive Choice>**—the number of documents that were coded positive outside of the queue. The smaller percentage underneath it shows the percentage of documents coded positive outside the queue divided by all documents coded.
- **<Negative Choice>**—the number of documents that were coded negative outside of the queue. The smaller percentage underneath it shows the percentage of documents coded negative outside the queue divided by all documents coded.
- **Neutral**—the number of documents that were coded with a neutral choice outside of the queue. The smaller percentage underneath it shows the percentage of documents coded neutral outside the queue divided by all documents coded.


3.2 Charts and tables

The dashboard includes two visualization panels. Both panels have the same options for charts and tables to show, which lets you choose which visualization to show on which panel, in any order.



To navigate the visualization panel:

- To select a different visualization, click the visualization's name. This opens a drop-down menu with all other visualizations.
- To switch from the chart view to the table view, click the **Chart** drop-down in the upper right corner and select **Table**. This shows a table with the same information as the selected chart.
- To zoom in or out on a chart, hover the cursor over it and scroll. All charts reset to their default zoom when you reload the page.
- To download the panel contents, click the **Export** symbol on the upper right. Charts download as .png images, and tables download as .csv files.



Notes:

- The charts and tables only reflect documents coded inside the queue.
- If any documents were coded by reviewers who are not part of this Relativity instance, those reviewers will be listed as Unknown User 1, Unknown User 2, and so on. This can happen if a reviewer was removed from the workspace or if the workspace has been archived and restored into a different instance.

3.2.1 General charts and tables

Some charts and tables are available for any type of queue. These include:

3.2.1.1 Coding Progress

The Coding Progress tab shows the count of documents that have been coded in the queue over time. Coding data is reported in 15-minute increments.

The numbers for Est. Total Docs and Est. Docs Remaining are updated every time the queue refreshes. Because they update at a different time than the coding data, these numbers are estimates.

3.2.1.2 Relevance Rate

The Relevance Rate tab shows the relevance rate over time. This can be shown overall or by user.

Each solid data point represents 100 documents, and a hollow data point represents any remainder. For example, if 201 documents have been coded, there will be 3 points: 2 solid points for each set of 100, and 1 hollow point for the final document.

Other details about the data points include:

- If you have more than one data point in a 15 minute increment, the chart shows them as two points on a vertical line. This can happen if many reviewers are coding quickly.
- The Date field for a data point is the date the last document in the set of 100 was logged.
- If a document has been coded more than once, the overall relevance statistics reflect the most recent decision. However, the per-user statistics show the most recent decision for each user. For example, if User 1 skipped a document and User 2 coded it as relevant, User 1's statistics will still reflect the skip.

For prioritized review queues, the relevancy rate usually declines over time. However, the relevance rate may spike if lots of new documents are added to the queue or if the definition of relevance changes during review. For saved search queues, the shape of the relevancy rate graph varies depending on the saved search being used.

3.2.1.3 Review Speed

The Review Speed tab shows the number of documents coded per hour. Data is reported in 15-minute increments.

Other details about the data points include:

- The Review Speed data can be shown overall or by user. When it's set to show all reviewers, the line chart shows a weighted average of the review speeds of the reviewers. It does not report their aggregate review speed.
- If a reviewer has coded the same document more than once, each time counts as a separate coding decision.

3.2.1.4 Queue History

The Queue History tab shows the state of the queue at every previous refresh. This is shown only as a table, not a chart.

The columns vary depending on the queue type. For saved search queues, it also depends on whether positive and negative choices are selected for the review field.

Possible columns include:

- **Refresh Start Time**
- **Refresh End Time**
- **Total Items**—the total number of documents in the data source.
- **Refresh Type**—this can be either Auto or Manual.
- **Coded <Positive Choice>** (optional for saved search queues)
- **Coded <Negative Choice>** (optional for saved search queues)
- **Uncoded Predicted <Positive Choice>** (prioritized review queues only)


- **Uncoded Predicted <Negative Choice>** (prioritized review queues only)
- **Coverage Mode** (prioritized review queues only)—whether the queue was in Coverage Mode during the refresh.

All document counts show the number of documents in that category at the Refresh End Time.

3.2.2 Prioritized review charts

The **Rank Distribution** chart is available for prioritized review queues. This chart helps you compare the model's predictions to reviewer's actual coding decisions. It shows the number of documents at each rank, from 0 to 100, color-coded by the reviewers' coding decisions on those documents.

A low relevance rank means that the model predicts that the document is more likely to be coded negative, and a high relevance rank means that the model predicts the document is more likely to be coded positive.



If you zoom out on the Rank Distribution chart, you may see documents with ranks below zero. These are documents that could not be classified. For more information, see [Understanding document ranks on page 33](#).

3.2.3 All Documents table

The All Documents table shows all documents in the queue. For reviewed documents, it shows which reviewer coded each document, how long the reviewer took, and how it was coded. This only shows coding that takes place within the Review Center queue.

For saved search queues, the columns depend on whether a review field is set, as well as if positive and negative choices are selected.

Possible columns include:

- **Control Number**—the control number of the document.
- **Reviewer**—the assigned reviewer's name.
- **Reviewed Time**—the check-in time for the document.
- **Reviewed Duration**—how much time passed between the reviewer opening a document in the review queue and saving a coding decision. This includes clicking Save, Save & Next, or Skip. The time is reported in hours, minutes, and seconds (HH:MM:SS).

- **Queue Coding Decision** (optional for saved search queues)—how the document was coded when the reviewer checked it back in. If the document was skipped, this is blank.
- **<Review Field Name>** (optional for saved search queues)—the current coding designation of the document.
- **Batch Number** (batched queues only)—the numerical ID of the batch the document belongs to.
- **Is Reviewed**—whether the document has been reviewed yet.
- **Random** (prioritized review queues only)—whether the document was randomly chosen, or served up according to its rank. For more information, see [Including random documents in the queue on page 11](#).

3.2.4 Batches table

For batched queues, you can click the triangle next to the All Documents title to select the Batches table. This shows a summary of batch-related statistics.

All Documents 

Columns include:

- **Batch Name**—the name of the batch. This defaults to the Queue Name the queue was created with, plus the batch ID.
- **Batch Number**—the numerical ID of the batch.
- **Batch Status**—whether the batch is New, In Progress, Pending, or Complete.
- **Batch Unit Field**—the value in the Batch Unit Field used to group the batch.
- **Assigned To**—the reviewer assigned to the batch. If no reviewer is assigned, this says Unassigned.
- **Reviewed Docs**—the number of documents that have been reviewed in the batch. This includes documents that were skipped.
- **Unreviewed Docs**—the number of documents that have been reviewed in the batch.
- **Batch Size**—the total number of documents in the batch.

From the Batches table, you can also assign or reassign batches to reviewers. For more information, see [Managing batched queues on page 36](#).

3.3 Preparing or refreshing the queue

In order for a queue to function, Review Center has to run the saved search, check for any outside-coded documents, and perform other actions. If it is a prioritized review queue, it also needs to periodically retrain the classifier. This collection of actions is referred to as refreshing the queue.

Depending on your settings, the refresh button may say several things:

- **Prepare Only**—appears when the queue has not been started. This runs the saved search and trains the classifier for the first time, but it does not start the queue. Alternately, you can click Prepare and Start to perform both actions together.

Note: Preparing a new queue in advance makes the Start Review action take only a few seconds. This can be helpful if your data source is very large or if you have a complicated saved search. For example, you might prepare a new queue overnight, then start it in the morning.

- **Refresh Queue**—appears during a review that does not use auto-refresh. Clicking this refreshes the queue.

- **Auto Refresh**—appears during a review that uses auto-refresh. Clicking this starts an immediate refresh of the queue. For more information, see [Auto-refreshing the queue below](#).

After you click Confirm, a **Cancel** option appears. For prioritized review queues, you may also see a confirmation modal with the option to refresh the cache. For more information, see [Caching text in prioritized review queues on page 30](#).

If you edit a queue's settings when the queue is partway through refreshing, the refresh will automatically cancel. Any edits that affect the queue refresh will take effect during the next refresh.

3.3.1 Checking refresh status

When you prepare or refresh a queue, the button displays the refresh progress as a series of steps. Step 1 has several phases, and it typically takes the longest. The system updates the status every 5 minutes.

The steps are:

- **Step 1 of 5**—initiates background processes and collects the documents. This has different phases depending on the queue type.
 - Prioritized review queue phases: **Initiating**, **Ingesting**, **Model training**, and **Ranking docs**. During the Ingesting phase, a banner appears below the Queue Summary section showing how many documents have been ingested.
 - Saved search queue phases: **Initiating** and **Running Search**.
 - Validation queue phases: **Initiating** and **Creating Sample**.
- **Step 2 of 5: Retrieving**—takes the results of Step 1 and prepares them for processing.
- **Step 3 of 5: Processing**—processes the documents. This includes updating ranks for prioritized review queues and retrieving document families. For batched queues, this includes creating batches from newly added documents.
- **Step 4 of 5: Capturing Stats**—records document statistics.
- **Step 5 of 5: Finalizing**—completes the job, finalizes metrics, and sends any requested email notifications.

To see what time the queue began preparing or refreshing, see the Queue History table. You can also use the Queue History table to check how long previous builds took. For more information, see [Queue History on page 25](#).

3.3.2 Auto-refreshing the queue

If Auto Refresh is set to On in the queue settings, the queue will automatically refresh at specific intervals. The interval length depends on the queue type and the coding activity. These refreshes only happen after the queue has been started, and you can change this setting at any time.

Saved search queues refresh every 15 minutes if there is coding activity within the queue.

3.3.2.1 Auto-refresh timing for prioritized review queues

Prioritized review queues refresh when there has been a 20% change in the total number of coded documents in the queue. For example, if 100 documents were previously coded, you would need to code 20 more to trigger a refresh.

As the total number of coded documents grows higher, the auto-refresh triggers less frequently, because 20% change requires more documents. The following table shows an example of how this grows at the beginning of a project.

Total documents coded positive or negative at the last refresh	Number of documents needed to trigger next refresh (rounded up)
50 (coded before queue started)	10
60 (50 plus 10)	12


Total documents coded positive or negative at the last refresh	Number of documents needed to trigger next refresh (rounded up)
72 (60 plus 12)	15
87 (72 plus 15)	18

A prioritized review queue will also auto-refresh if there is coding activity within the queue and it has been 8 hours since the last refresh, regardless of whether 20% more documents have been coded. For example, if there are 1000 coded documents in the queue and a reviewer codes 1 more, the queue will auto-refresh after 8 hours. However, if the queue were to sit completely inactive for 8 hours, with no reviewer coding, the queue would not auto-refresh.

Note: Auto-refresh only triggers after coding within the Review Center queue. It does not trigger if the documents are coded elsewhere in Relativity.

3.3.2.2 Viewing auto-refresh estimates

For prioritized review queues, the Auto Refresh button shows an estimate of how many documents must be coded to trigger the next auto-refresh. When that number have been coded positive or negative, the next auto-refresh will start within about five minutes.



3.3.2.3 Manually triggering or canceling auto-refresh

If you need to trigger an immediate refresh, click on the words **Auto Refresh** to trigger an additional manual refresh. For example, if new documents have been added to the saved search, you can click this to add them to the queue quickly instead of waiting until the next auto-refresh.

While the queue is auto-refreshing, a **Cancel** option appears. If you cancel the current auto-refresh, the queue will still try to auto-refresh again later.

Note: Canceling the queue preparation can take some time. If you need to remove reviewer access immediately while canceling, edit the queue and remove the reviewer group.

3.3.3 Reviewer access during refreshes

Reviewers can still review documents in an active queue while it refreshes. Clicking the refresh button, running an auto-refresh, or canceling a refresh makes no difference to reviewer access.

Similarly, if the queue was paused before the refresh, it will stay unavailable. Active queues stay active, and paused queues stay paused.

3.3.4 Auto-refreshing in Coverage Mode

If your prioritized review queue has automatic refreshes enabled and Coverage Mode turned on, the refreshes trigger at a different time. The queue will automatically refresh each time 100 documents are coded within the queue, or when 5% of the documents have been coded positive or negative, whichever occurs first. The "Next refresh" document count reflects this change whenever you turn on Coverage Mode.

For example, if the queue has 1000 total documents, coding 50 positive or negative would trigger the next auto-refresh. If the queue has 2000 or more documents, coding 100 would trigger the next auto-refresh.

Note: Whenever you turn Coverage Mode on or off, manually refresh the queue. This updates the document sorting for reviewers. For more information, see [Turning Coverage Mode on and off on page 32](#).

3.3.5 How queue refreshes affect batching

For batched queues, refreshing the queue re-runs the saved search and does the following:

- If new documents were added to the saved search, it groups any added documents into new batches.
- If any documents were removed from the saved search, it removes those documents from previously created batches. Removed documents that were already coded will still be listed in the All Documents and Batches tables, but will not be counted in other dashboard statistics.
- If you have a Minimum Batch Size set, batches created during an auto-refresh respect that minimum. When you refresh manually, a confirmation window appears with the option to create a batch smaller than the minimum.

Documents that have already been assigned to a batch remain in their batch, unless they were removed from the saved search. Refreshing the queue does not mix documents from previously created batches with newly added ones.

If the Auto Batch setting is toggled **On** for the queue, the queue will automatically create new batches during each auto-refresh. If it is toggled **Off**, manually refresh the queue to batch new documents.

3.3.6 Caching text in prioritized review queues

The first time you prepare a prioritized review queue, Review Center caches the text of the documents in the queue and stores the documents' data at the workspace level. This significantly speeds up later refreshes, because Review Center references the cache instead of re-analyzing the text. This also speeds up the creation of any other queues in the workspace with the same documents.

By default, Review Center caches the Extracted Text field of each document. If you selected a different field to analyze when you set up your queue, it will cache that text instead.

When you click to manually refresh the queue, a modal appears with an option to refresh the workspace cache:

- If the text of documents in the queue's data source has not changed, leave the box **unchecked**. This makes the refresh process much faster.

Note: You do not need to refresh the cache if you are simply adding or removing documents from the queue.

- If the text of documents in the queue's data source has changed, **check** the box. This tells Review Center to clear the cached fields for all documents in this workspace. It then re-caches the text of documents in this queue. Choosing to re-cache the text may add significant time to the queue refresh, as well as the next refresh for queues with other cached documents.

Note:

If the Text Precedence fields you selected for your queue are Data Grid enabled, preparing the queue for the first time may run up to three times faster than for fields that store data in SQL. After the text has been cached, refresh times are typically equal between both types of fields.

For more information on using Data Grid, see Processing to Data Grid in the Processing guide.

3.4 Starting the queue


The **Start Review** button makes the queue available for review. If the queue has never been prepared before, it will say **Prepare and Start**. This also runs the saved search and trains the classifier for the first time.

After the queue has finished starting, the symbol beside this option changes to a pause button. Clicking this pauses the queue and stops reviewers from checking out more documents.

Before starting a queue, you must have a reviewer group assigned.

3.5 Editing queues and other actions

To edit the queue or perform other less-frequent actions, click on the three-dot menu on the right of the Queue Summary section.



The menu options are:

- **Edit**—opens a modal to edit any of the queue settings.
 - For information on general edits, see [Editing recommendations below](#).
 - For information on Coverage Mode, see [Turning Coverage Mode on and off on the next page](#).
- **Release Checked Out Documents**—releases all documents that are checked out by reviewers. If a reviewer falls inactive and does not review the last few documents in a queue, this frees up those documents for reassignment.
 - To see the number of currently checked out documents, look at the main ribbon for the Queue Summary section.
 - If you release documents while a reviewer is actively reviewing, that person will be able to finish coding, but their documents may get checked out by another reviewer at the same time. To prevent this, ask any active reviewers to exit and re-enter the queue after you click the link.
- **Release Skipped Documents**—adds skipped documents back into the document pool as uncoded documents. They will be available for reviewers, and all metrics will show them as Uncoded instead of Skipped. The coding records tied to the document will still show that it was skipped, as well as any later coding decisions.
- **Save as New Model** (prioritized review queue only)—opens the options to save the relevance training from this queue for re-use with another queue. For more information, see [Reusing saved models on page 59](#).
- **Set up Validation** (prioritized review queue only)—opens the options to create a review validation queue. For more information, see [Review validation on page 44](#).
- **Delete Batches** (batched queues only)—deletes the existing batches for a queue. For more information, see [Deleting batches on page 37](#).
- **Mark as Complete**—sets the queue's status to Complete and moves it to the far right of the queue tab strip. This also removes the queue from the Review Queues tab, and reviewers can no longer access it. After the queue has been marked Complete, this option changes to **Re-enable**. Clicking this sets the queue's status to Not Started and returns it to the Review Queues tab.

3.5.1 Editing recommendations

Many edits are minor, and you can make them without pausing the queue. However, if you make a major change such as changing the data source, we recommend:

1. Pause the queue before editing.
2. Release any checked out documents.
3. Edit the queue.

4. Refresh the queue.
5. Restart the queue.

For descriptions of the editable fields, see [Creating a Review Center queue on page 10](#).


3.5.2 Turning Coverage Mode on and off

When you turn Coverage Mode on or off for a prioritized review queue, this changes the order in which the documents will be served up. Before the new order will take effect, though, you must refresh the queue.


To turn Coverage Mode on or off:

1. On the right side of the Queue Summary section, click the three-dot menu.
2. Select **Edit**.
3. Click the **Coverage Mode** toggle to enable or disable it.
4. Click **Save**.
5. Manually refresh the queue. For more information, see [Preparing or refreshing the queue on page 27](#).

On the right of the Queue Summary section, the start or pause button reflects whether Coverage Mode is turned on. If it is turned **On**, it will refer to the queue as "Coverage Review."



If Coverage Mode is turned **Off**, the button will refer to the queue as "Prioritized Review."



On the Queue History table, you can also see if your queue was in Coverage Mode or not during each queue refresh. For more information, see [Queue History on page 25](#).

3.6 Deleting a queue

Queues can be edited or deleted from the Review Library tab.

To delete a queue:

1. Navigate to the **Review Library** tab.
2. Click on the queue you want to delete.
3. Click **Delete**.
A confirmation pop-up will appear.
4. Click **Delete** again.
After the process completes, you will return to the main Review Library tab.

Deleting a queue does not remove any of the coding decisions or rank values that have been assigned to the documents.

Note: If you delete a main queue that has a validation queue linked to it, it also deletes the validation queue. For more information on validation queues, see [Review validation on page 44](#).

3.7 Fixing a misconfigured queue

If a required field or object that a queue relies on is deleted or moved, this puts the queue into a warning state. Any queue preparation or auto-refresh stops, and a message appears at the top of the Review Center tab directing you to the field or object that needs to be fixed. Your reviewers also see a warning at the top of the Review Queue page telling them which queue is misconfigured and that they should alert their administrator.

When this happens, we recommend pausing the queue and checking its settings. For example, if the saved search was deleted, you may need to link the queue to a new saved search. If a required field was deleted, you may need to create a new one.

If you have checked the queue's settings and still see warnings, contact [Product Support](#).

3.8 Understanding document ranks

During prioritized review, the AI classifier assigns a rank to each document. These ranks are stored in the Rank Output field, and they determine the order in which reviewers will see documents.

Most document ranks range from 0 to 100. The higher the score, the stronger the prediction that the document will be coded on the positive choice. The AI classifier recalculates ranks every time the queue refreshes, and the highest-ranking documents are served up to reviewers.

Note: In order to improve efficiency and performance, Relativity reserves the right to update the prioritized review queue's AI classifier during software upgrades. Although we work hard to minimize disruptions, these upgrades can cause minor differences in document ranking. As a result, administrators may occasionally see minor variations in document ranks after a queue is refreshed, even without any new document coding.

If the classifier cannot classify a document, it will assign the document a value below zero. These values are:

Negative rank	Document error
-1	An error occurred while processing the data through the classifier.
-2	The extracted text field or other selected text field is empty. If you see this rank, consider making a saved search queue to review these documents separately.
-3	The document's extracted text field or other selected text field is larger than the limit of 30MB. If you see this rank, we recommend filtering out large documents from your saved search to improve the performance of the classifier.

3.9 Tracking reviewer decisions

You can view coding decisions made by each reviewer in the All Documents table. For more information, see [All Documents table on page 26](#).

Alternatively, you can also use the following methods.

3.9.1 Using the Documents tab

The Review Center Coding fields track the reviewer names, decisions, and dates. You can add these to views and saved searches from the Documents tab.

Note: We recommend using the fields with double colons (such as Review Center Coding::Queue) for views and saved searches. These are called reflected fields, and they are attached to the Document object.

The field names are:

- **Review Center Coding::Coding Duration**—how much time passed between the reviewer opening a document in the review queue and saving a coding decision. This includes clicking Save, Save & Next, or Skip. The time is reported in hours, minutes, and seconds (HH:MM:SS).
- **Review Center Coding::Field Name**—the name of the Review Field for the queue.
- **Review Center Coding::Queue**—the name of the Review Center queue that contains the document.
- **Review Center Coding::Reviewed By**—the name of the reviewer who made the coding decision.
- **Review Center Coding::Reviewed On**—the date of the coding decision. Dates are based on the UTC time zone.
- **Review Center Coding::Value**—the reviewer's coding decision.
- **Review Center Coding::Was Outside Coded**—whether the document was coded through some means other than the selected Review Center queue.
- **Review Center Coding::RC Batch** (batched queues only)—the name of the batch the document belongs to. This includes the Queue Name the queue was initially created with, plus the batch ID.

These fields are populated for each document when that document is reviewed in a queue. Documents that have not been reviewed yet do not have values in these fields. If a document has been coded multiple times, each coding decision appears as a sub-item within the row.

For more information on creating views and saved searches, see [Creating a view](#) and [Creating or editing a saved search](#) on the Relativity aiR documentation site.

3.9.2 Using the Field Tree

The Field Tree helps you get a quick overview of document coding decisions. Unlike creating a view or saved search, it does not show which reviewer made each decision.

To view coding decisions using the Field Tree:

1. Navigate to the **Documents** tab.
2. In the browser panel, click on the tag symbol to open the Field Tree.
3. Scroll to the folder labeled **Review Center** and expand it.
4. Click on your queue's name. This shows all documents currently in the queue, plus any documents that were coded in the queue but later removed.

Depending on your queue's history, there may also be other tags nested underneath it:

- **<queue name> Validation <#>**—this lists documents in an attached Validation queue. If the queue has several attached Validation queues, each one will have its own tag.
- **Removed**—this lists any documents that were coded in the queue, but later removed from the data source.

If you rename or delete a queue, this renames or deletes the matching Field Tree tags also. The tags are automatically stored in a multi-choice field called Review Center.

3.9.3 Using the Track Document Field Edits by Reviewer application

The Track Document Field Edits by Reviewer application lets you see which reviewer made each coding decision. You can set up the application individually for each of your queues.

Install the application using the instructions from [Track document field edits by reviewer](#) on the Relativity aiR documentation site.

When configuring the application:

1. Put the **Reviewed On** and **Reviewed By** fields you create into a saved search or view for monitoring.

Note: These fields should be created from scratch when setting up the application. Do not use the preexisting Review Center Coding::Reviewed On or Review Center Coding::Reviewed By fields. For full instructions, see [Track document field edits by reviewer](#) on the Relativity aiR documentation site.

2. Set your queue's review field as the **Field to Monitor**.
3. Set **Track Initial Change Only?** as follows:
 1. If this is set to **Yes**, the application tracks only the first reviewer of the document. It does not track any later edits.
 2. If this is set to **No**, the application updates the Reviewed On and Reviewed By fields every time a user edits the document.

If you set up the application after starting your queue, you can still see previous coding decisions by following the steps under Populating Historical Records.

3.10 Moving Review Center templates and queues

Review Center templates and queues are Relativity Dynamic Objects (RDOs), which typically can be moved across workspaces or instances with Relativity Integration Points. However, because of the complexity of an active queue, we do not support moving active queues. Doing so could damage your Review Center environment.


We do support moving queue templates across workspaces or instances using Relativity Integration Points. This process is safe for your environment.

3.11 Viewing archived Active Learning projects

If your workspace includes projects from the older Active Learning application, you can view read-only statistics and results for those projects from the Active Learning History tab.

To access the tab, click the **Active Learning History** icon on the right side of the queue tab strip filter.

Note: Make sure the Active Learning application is still installed in the workspace. Uninstalling Active Learning removes the project data.



4 Managing batched queues

Administrators can use the Review Center dashboard to monitor batch progress, assign batches, and add or remove documents.

4.1 Assigning batches from the dashboard

You can assign batches to reviewers from the Review Center dashboard. If you prefer to let reviewers choose their own from the Review Queues tab, you can leave some or all batches unassigned.

To assign batches from the dashboard, you must have the Assign Batches permission:

Tab Visibility	Other Settings
• Review Center	• Assign Batches

For more information on permissions, see [Review Center security permissions on page 63](#).

To assign batches to reviewers:

1. On the Review Center dashboard, scroll down to the All Documents table.
2. Click the blue arrow next to the table name, then select **Batches**. This switches the panel to show the Batches table.
3. Select the batch or batches you want to assign, then click **Manage Batch Assignment**.
4. Choose the following:
 - Action: **Assign**
 - Reviewer: Select a reviewer.
5. Click **Assign**.

If you do not assign all batches, reviewers can assign themselves available batches from the Review Queues tab. For more information, see [Checking out batches on the next page](#).

By default, batches are offered to reviewers in order of when the batches were created. Reviewers can also select individual batches from the Review Queues tab.

4.1.1 Reassigning completed batches

If a reviewer checks a batch in as Complete, you can still reassign that batch to another reviewer. This is especially helpful if not all documents in the batch were reviewed.

When you reassign a completed batch, the batch status changes to In Progress.

4.2 Monitoring queues

To monitor a batched queue:

- Use the Review Center tab as a dashboard to view progress reports, charts, and queue status. For more information, see [Monitoring a Review Center queue on page 20](#).

- The All Documents table shows both reviewed and unreviewed documents. Within it, the Batch Number column shows which batch each document belongs to. You can filter the table by batch, reviewer, coding decision, and more. For more information, see [All Documents table on page 26](#).
- For assigning, unassigning, and viewing the status of each batch, use the Batches table at the bottom of the Review Center dashboard. This is part of the same panel as the All Documents table, and you can select one or the other by clicking the blue arrow next to the table name. For more information, see [Batches table on page 27](#).

4.3 Adding or removing documents

To add, remove, and batch out new documents:

1. Add any new documents to the saved search that you selected as the queue's data source.
2. Click to refresh the queue. If you have auto refresh turned on, the button will say **Auto Refresh**. If it is turned off, it will say **Refresh Queue**.
3. In the confirmation modal, select or leave blank the following:
 - **Ignore batch minimum**—this creates a small batch of any remaining documents that are smaller than the minimum batch size. For example, if you added 703 documents and have a normal batch size of 100, this would create seven 100-document batches and a 3-document batch.
4. Click **Confirm**.

If new batches were created, you can assign them using the same method as the initial batches. For more information, see [Assigning batches from the dashboard on the previous page](#).

4.4 Deleting batches

If you decide you no longer want to use batching, or if you want to re-batch all documents at the same time, you can delete the existing batches. The documents within the batches are unaffected, and you can create new batches from them by refreshing the queue. This is similar to the Purge Batches action from standard review.

To delete all existing batches for a queue:

1. Pause the queue.
2. On the Review Center dashboard, click the **three-dot menu** at the right side of the Queue Summary section.
3. Select **Delete Batches**. A confirmation modal appears.
4. Click **Delete Batches**.
 - If it succeeds, the Batches table will no longer show any batches.
 - If an error message appears, check that the queue has been paused, then retry.

If you want to create new batches, refresh the queue.

5 Checking out batches


Use the Review Queues tab to check out batches, review documents, and check completed batches back in.

This page is for batches created using Review Center. For other batches, see Batches on the Relativity documentation site.

For instructions on how to assign batches to users as an administrator, see [Managing batched queues on the previous page](#).

5.1 Viewing queues and batches

When you log in as a reviewer, you can see all active queues and batches assigned to you from the Review Queues tab.



Near the top of the page, you can filter which queues and batches to show:

- **My Work**—shows active queues and batches that are assigned to you, as well as any batches that are available to self-assign. This is the default.
- **Available Batches**—shows only batches that are available to self-assign. If you are not assigned to any batched queues, this filter does not appear.

5.1.1 The queue and batch list

Under the My Work filter, you can see queues and batches that are assigned to you.

- Queues that contain batches assigned to you have a Start button on the right.
- If the queue batches were grouped by Batch Unit Field, the name of that field appears underneath the queue name.
- Clicking the arrow on the left side of a batched queue expands the list of batches.

Search...						
		OVERALL PROGRESS			MY PROGRESS	
			Total Remaining	Total Docs	Reviewed Today	Reviewed Total
Queue without Batching		0%	77,674	77,674	0	0
Queue with Batching		0%	77,673	77,674	0	2
> 1 batch assigned to me Custodians						

A search bar appears at the top of the queue and batch list. Currently, you can search by queue name.

5.2 Selecting and assigning batches

Depending on how the queue is set up, your administrator may have assigned you batches in advance. The administrator may also have left some batches unassigned, which you can assign to yourself.


Queues and batches are visible if you have been assigned to the reviewer group for that queue.

5.2.1 Checking out already assigned batches

Queues that contain batches assigned to you have a Start button on the right. Clicking Start without choosing a specific batch will serve up the first batch assigned to you. You can also expand the list of batches and select a specific one to review.

To choose a specific batch:

1. Click the arrow next to the queue name. This expands the list of batches.




2. Locate the batch you want.
 - If there are more than 10 batches within that queue, pagination arrows appear at the upper right.
 - If the batches were grouped by Batch Unit Field, that field and its value appear under the batch name. For example, you may see a custodian's name or a file type.
3. To start coding, click the **Start** button on the row.

5.2.2 Checking out unassigned batches

If the administrator has left some batches unassigned, you can assign them to yourself.

To self-assign a batch:

1. At the top of the Review Queues tab, choose the **Available Batches** filter. All available queues with unassigned batches appear.
2. Locate the queue you want, then click the arrow next to the queue name. This expands the list of batches.




3. Locate the batch you want.
 - If there are more than 10 batches within that queue, pagination arrows appear at the upper right.
 - If the batches were grouped by Batch Unit Field, that field and its value appear under the batch name. For example, you may see a custodian's name or a file type.
4. Click **Check Out**.

The checked-out batch now appears under the My Work filter.

5.3 Batch options

Each batch assigned to you has a three-dot menu on the far right.



The batch options include:

- **Check in as complete**—mark this batch as completed. This removes the completed batch from the list on the Review Queues tab, but it can still be reassigned by an administrator.
- **Check in as pending**—mark this batch as pending.
 - In general, people typically check batches in as pending if there are still documents remaining to be reviewed.
 - Pending batches go to the bottom of the batch list for that queue. You can still check them back out again from the Review Queues tab.
- **Reassign Batch**—reassign the batch to another user or unassign it. This option only appears if you have the Assign Batches permission. For more information, see [Review Center security permissions on page 63](#).

Both pending and complete batches can be reassigned from the Review Center dashboard and then checked back out. For more information, see [Managing batched queues on page 36](#).

6 Reviewing documents using Review Center

The Review Queues tab is the starting point for reviewers. Every Review Center queue that a reviewer is assigned to shows up here.

This topic provides step-by-step instructions for accessing a queue and reviewing documents.

6.1 Reviewing documents in the queue

The Review Queues tab shows all queues assigned to you. Each queue has a row in the grid. If you have any batched queues available, you can expand the queue's row to see smaller rows for each batch.

To review documents in a queue:

1. Navigate to the **Review Queues** tab.
2. Each queue you are assigned to has a separate row. Locate the row with the same name as the queue you want.
3. For batched queues:
 1. Click the arrow beside a queue name to expand the list of batches underneath it.
 2. If there are unassigned batches, you can check them out to claim them. Click **Check Out** on the batch row.
For more detailed instructions, see [Checking out batches on page 37](#).
4. Click **Start**.
This opens the document viewer.
5. Code the current document according to your organization's workflow.
6. Click **Save and Next**.
The next document will appear for review.

If you do not see a Start button, either the queue is paused, or the admin has not started the queue. Talk to your administrator to find out when the queue will be ready.

For more information on using the document viewer, see Viewer in the Admin guide.

Note: If the same document is part of two separate queues, two reviewers can check it out at the same time. If that happens, one reviewer may see a message that the document has been edited elsewhere. To see the other reviewer's edits, reload the page.

6.1.1 Checking in batches

When you reach the last document in a batch, a modal appears with some options. You can:


- **Check in as complete & start next batch**—mark this batch as completed and start reviewing the next assigned batch.
- **Check in as pending & start next batch**—mark this batch as pending and start reviewing the next assigned batch. Pending batches go to the bottom of the batch list for that queue.

You can mark batches as either complete or pending, even if you have not finished reviewing every document. However, completed queues will disappear from the Review Queues tab. For more information, see [Batch options on page 39](#).


6.2 Finding previously viewed documents

As you review, you can click on Documents on the left-hand navigation. This opens the Documents panel.


- For linear queues, this shows you documents you already reviewed in the queue.
- For batched queues, this shows you the other documents in the batch.



To view a document, click on its control number in the panel.



To return to your current document, click on the **Navigate Last** button in the upper right corner of the document viewer.



For linear queues, the Documents panel shows a set number of documents based on your Item List Page Length setting. For information on changing this, see Navigation in the User guide.

Notes:

- When you filter columns in the Documents panel, the filters only apply to documents on the current page of the panel. For a comprehensive list of results, filter within the Documents tab of Relativity or run a search from the saved search browser or field tree.
 - If you have reviewed more than 1000 documents in the queue, the navigation arrows in the Documents panel will only show documents from your current session. This limit applies even if the Item List Page Length is set higher.
-

6.3 Queue statistics

If your admin has enabled it, you may see values in the statistics columns displayed across the grid. These include:

- Overall Progress
 - **Total Remaining**—the total number of uncoded documents in the queue or batch.
 - **Total Docs**—the total number of documents in the queue or batch.
- My Progress
 - **Reviewed Today**—the number of documents you reviewed in the queue or batch today. These are counted starting at midnight in your computer's time zone.
 - **Reviewed Total**—the total number of documents you have reviewed in the queue or batch.

If you look at a document in the queue but skip coding it, it still counts as reviewed. Documents that were coded outside of the queue do not count towards these totals.

6.4 Viewing the dashboard

Your admin may give you access to the Review Center dashboard. The dashboard shows how the review is progressing, including statistics and visualizations.

For more information on the Review Center dashboard, see [Monitoring a Review Center queue on page 20](#).

6.5 Best practices for Review Center review

When reviewing documents in a Review Center queue, we recommend the following guidelines:

- **Double check**—always check the extracted text of a document to be sure it matches the content in other views. Whenever possible, review from the Extracted Text viewer.
- **Stay consistent**—check with fellow reviewers to make sure your team has a consistent definition of relevance. The AI classifier can handle occasional inconsistencies, but you'll get the best results with coordinated, consistent coding.
- **When in doubt, ask**—if something confuses you, don't guess. Ask a system admin or project manager about the right course of action.

6.5.1 Coding according to the "four corners" rule

For queues that use AI, Review Center's AI classifier predicts which documents will be responsive or non-responsive based on the contents of the document itself. Family members, date range, custodian identity, and other outside factors do not affect the rankings. Because of this, the AI classifier learns best when documents are coded based only on text within the four corners of the document.

When you code documents as positive or negative in a Review Center queue, you are both coding the document and teaching the AI classifier what a responsive document looks like. Therefore, your positive or negative coding decisions should follow the "four corners" rule: code only based on text within the body of the document, not based on surrounding factors.


Having one or two documents that fail this rule will not harm the overall accuracy of Review Center's predictions. However, if you want to track large numbers of documents that are responsive for reasons outside of the four corners, we recommend talking to the project manager about either setting up a third, neutral choice on the coding field for those, or adding a secondary coding field. Neutral choices and other coding fields are not used to train the AI classifier.

6.5.1.1 Common scenarios that fail the "four corners" rule

The following scenarios violate the "four corners" rule and do **not** train the AI classifier well:

- The document is a family member of another document which is responsive.
- The document comes from a custodian whose documents are presumptively responsive.
- The document was created within a date range which is presumptively responsive.
- The document comes from a location or repository where documents are presumptively responsive.

For example, the following email has a responsive attachment. However, the body of the email—the text within the four corners—is only a signature and a privacy disclaimer. Because the body of this email is not responsive, this document does not pass the "four corners" rule.



6.5.2 Factors that affect Review Center's predictions

Not all responsive documents inform Review Center equally. The following factors affect how much the AI classifier learns from each document:

- **Sufficient text**—if there are only a few words or short phrases in a document, the engine will not learn very much from it.
- **Images**—text contained only in images, such as a photograph of a contract, cannot be read by Review Center. The system works only with the extracted text or other selected text field of a document.
- **Numbers**—numbers are not considered by Review Center.

7 Review validation

Review validation evaluates the accuracy of a Review Center queue. The goal of validation is to estimate the accuracy and completeness of your relevant document set if you were to stop the queue immediately and not produce any unreviewed documents. The primary statistic, elusion rate, estimates how many uncoded documents are actually relevant documents that you would leave behind if you stopped the queue. The other statistics give further information about the state of the queue.

- For a practical overview of how the validation statistics are calculated, see [Review validation statistics on page 52](#).
- For detailed equations, see [Review Center Validation Deep Dive](#) on the Community site.
- For validating aiR Review Prompt Criteria, see aiR Review Prompt Criteria validation in the aiR Review guide.

Note: Review validation does not check for human error. We recommend that you conduct your own quality checks to make sure reviewers are coding consistently.

7.1 Key definitions

The following definitions are useful for understanding review validation:

- **Discard pile**—the set of documents that are uncoded, skipped, or coded as neutral. This also includes documents that are being reviewed when validation starts, but their coding has not been saved yet.
- **Already-coded documents**—documents that have already been coded as either positive or negative. These are counted as part of the validation process, but they will not be served up to reviewers a second time. Neutral-coded documents are considered part of the discard pile instead, and those may be served up a second time.

7.2 Determining when to validate a Prioritized Review queue

When a Prioritized Review queue is nearing completion, it can become more difficult to find additional relevant documents. As you monitor your queue, the following dashboard charts can help you determine when the queue is ready for validation:

- **Rank Distribution**—look for few or no unreviewed documents with a rank of 50 or higher.
- **Relevance Rate**—you should see a decline in the relevance rate progress line indicating that very few responsive documents are being found.

When you believe you have found most of the relevant documents, run validation to estimate the accuracy and completeness of your relevant document set.

For more information on the dashboard charts, see [Charts and tables on page 23](#).

7.3 Starting a validation queue

When you are ready to validate your progress in a Review Center queue, you can start a linked validation queue that samples documents from the discard pile and serves them to reviewers.

To set up the validation queue:

1. From the **Review Center** tab, click on the queue you want to validate.
2. Pause and stabilize the queue:
 1. If auto-refresh is turned on, turn it off.
 2. If the queue is in the middle of refreshing, wait until the refresh has finished before starting validation.
 3. Click the Pause button.
 4. Release any documents that are currently checked out to reviewers. For more information, see [Editing queues and other actions on page 31](#).
3. On the right side of the Queue Summary section, click on the three-dot menu and select **Set up Validation**. An options modal appears.
4. In the options modal, set the following:
 1. **Validation Reviewer Groups**—the user groups you want reviewing the queue.
 2. **Cutoff**—enable this to set a cutoff for the validation queue. For more information, see [How setting a cutoff affects validation statistics on page 54](#).
 3. **Positive Cutoff**—if Cutoff is enabled, enter a custom cutoff value between 0 and 100 in this field. This rank will be used as the dividing line between documents predicted positive and documents predicted negative. Setting this value also adds a Precision statistic to the validation results and affects the Margin of Error Estimate (Elusion).
 4. Set the sample size using these fields:
 1. **Sample Size**—sets a fixed number of documents for the sample size. By default, this field is set to 1000 documents. The sample size must be larger than 5 and smaller than the size of the discard pile.
 2. **Richness**—the estimated percentage of relevant documents in the data source. By default, this field is set to 10%. We recommend leaving this set to 10% unless you know the richness of your data source. Entering a higher Richness value lowers the Margin of Error Estimate (Recall).
 3. **Margin of Error Estimate (Elusion)**—how accurate the Elusion statistic will be based on the current sample size. Entering a lower margin increases the Sample Size and the Margin of Error Estimate (Recall).
 4. **Margin of Error Estimate (Recall)**—how accurate the Recall statistic will be based on the current sample size. Entering a lower margin increases the Sample Size and the Margin of Error Estimate (Elusion).

Note: Each of these fields affects the others. For more details, see [Choosing the validation settings below](#).

5. Click **Save**.

7.3.1 Choosing the validation settings

When you set the sample size for validation, each of these fields affects the others:

- If you have a Positive Cutoff, raising it increases the number of documents in the sample below the cutoff, which decreases the Margin of Error Estimate (Elusion). Raising this cutoff generally increases the Elusion rate, but the margin of error goes down.
- Increasing the Sample Size lowers both Margin of Error estimates.
- Increasing the Richness percent lowers the Margin of Error Estimate (Recall).
- Lowering the Margin of Error Estimate (Elusion) increases the Sample Size and the Margin of Error Estimate (Recall).
- Lowering the Margin of Error Estimate (Recall) increases the Sample Size and the Margin of Error Estimate (Elusion).

If you know a specific sample size or margin of error is more important than the others for your review process, set that value manually. The other fields will auto-adjust as needed.

The final margin of error estimates may be slightly different from the ones chosen at setup, depending on the documents found during validation. All validation statistics will have 95% confidence intervals alongside the margin of error.

For more information on how validation statistics are calculated, see [Review validation statistics on page 52](#).

7.3.2 Inherited settings

Each validation queue inherits these settings from the main queue:

- Queue Display Options
- Document Panel View
- Coding Layout
- Email Notification Recipients

To change them, edit the validation queue after creating it. For more information, see [Editing a validation queue on the next page](#).

7.4 Coding in a validation queue

Reviewers access the validation queue from the Review Queues tab like all other queues.

During review:


- Have reviewers code documents from the sample until all documents have been served up.
- Sample documents that are coded outside the Review Queues tab will still count towards the validation statistics, but they might not be reflected in dashboard charts and tables.
- For best results, we strongly recommend coding every document in the validation queue as positive or negative. Avoid skipping documents or coding them as neutral. For more information, see [How validation handles skipped and neutral documents on page 58](#).

Have reviewers code documents from the sample until all documents have been served up.

7.5 Monitoring a validation queue

Validation statistics are reported on the Review Center dashboard like any other queue. You can cancel validation from the three-dot menu, and you can pause validation by clicking the **Pause** button. All data in the charts and tables reflect the validation queue.

During validation, the Review Progress section changes to become a Validation Progress section, which shows the progress of the validation queue. To view validation statistics instead, click the arrow next to the section name, then select **Validation Stats**.



7.5.1 Viewing formulas for validation statistics

In the Validation Stats section, you can click on each statistic to show how the statistic is calculated. This includes the formula used, plus a table listing the variables and their values.

Elusion = FORMULA	$\frac{FN_s}{FN_s + TN_s}$	Variable	Meaning	Your Value
		FN _s	False Negatives (Sample) - documents predicted non-responsive but coded responsive in the sample	5
		TN _s	True Negatives (Sample) - documents predicted non-responsive and coded non-responsive in the sample	370

For more information on the validation statistics, see [Reviewing validation results on page 50](#).

7.5.2 Editing a validation queue

You can change some of the queue settings at any time during validation.

To edit the validation queue:

1. On the right side of the Queue Summary section, click on the three-dot menu and select **Edit**.
2. Edit any of the following settings:
 - Additional Admin Settings
 - **Anonymize Reviewer Names**
 - **Queue Label**
 - **Email Notification Recipients**
 - Reviewer Settings
 - **Reviewer Groups**
 - **Document Panel View**
 - **Coding Layout**
 - Additional Reviewer Settings
 - **Total Remaining Documents in Queue**
 - **Total Documents in Queue**
 - **Reviewed Documents Today by Reviewer**
 - **Reviewed Documents Total by Reviewer**

3. Click **Save**.

For descriptions of the queue settings, see [Queue and template settings list on page 16](#).

7.5.3 Releasing unreviewed documents

If a reviewer falls inactive and does not review the last few documents in a validation queue, you can release those documents through the Queue Summary section of the dashboard. For more information, see [Editing queues and other actions on page 31](#).


To see which documents are checked out to a reviewer, filter the All Documents table by the reviewer's name. Any documents that are still checked out will show the Coded Time as blank. For more information, see [All Documents table on page 26](#).

7.5.4 Tracking sampled documents

If you want to run your own calculations or view documents in the validation sample, you can track the sampled documents from the Document list page. This process is optional.


To view sampled documents:

1. From the Documents tab, click on the **Field Tree** icon.



2. Expand the **Review Center** folder.
3. Expand the folder for the queue you're validating.
4. Several subfolders appear.
5. Expand the folder titled **[Queue Name] - Validation [Current Round Number]**. If you have only run validation one time, the round number will be **1**.

Each validation folder contains the documents selected for the sample. It also holds a sub-choice that shows all documents removed from that sample.



To view coding decisions for each document, add **Review Center Coding::Value** to the document view. For other optional fields, see [Tracking reviewer decisions on page 33](#).

7.6 Accepting or rejecting validation results

After all documents in the validation queue have been reviewed, a ribbon appears underneath the Queue Summary section. This ribbon has two buttons: one to accept the validation results, and one to reject them.

If you click Accept:

- The main queue status changes to Validation Complete.
- The model remains frozen. Any future coding decisions will no longer be used to train the model, and the Review Progress statistics will not reflect any new coding.
- The Validation Progress strip on the dashboard displays the final validation statistics.

If you click Reject:

- The validation queue status changes to Rejected, and the main review queue changes to Paused.
- The main review queue re-opens for normal coding, and you can build the model again at any time. Any documents coded since validation began, including those from the validation queue itself, will be included in the model build.
- The Coding Progress strip on the dashboard displays the main queue's statistics.

You can run validation on the queue again at any later time, and you can reject validation rounds as many times as needed. Even if you reject the results, Review Center keeps a record of them. For more information, see [Viewing results for previous validation queues on page 51](#).

7.6.1 Manually rejecting validation results

If you change your mind after accepting the validation results, you can still reject them manually.

To reject the results after accepting them:

1. On the right side of the Queue Summary section, click on the three-dot menu and select **Reject Validation**.
2. Click **Reject**.

After you have rejected the validation results, you can resume normal reviews in the main queue.

7.7 Reviewing validation results

After reviewers code all documents in the sample, the validation results appear in the Validation Progress strip on the Review Center dashboard.

The results include:

- **Relevance Rate**—percentage of sampled documents that were coded relevant by reviewers, out of all coded documents in the sample. If any documents were coded as neutral, this statistic also counts them as relevant.
- **Elusion Rate**—the percentage of unreviewed documents that are predicted as non-relevant, but that are actually relevant. The range listed below it applies the margin of error to the sample elusion rate, which is an estimate of the discard pile elusion rate.

Notes:

- If you do not set a cutoff for your validation queue, this is calculated as the percentage of all unreviewed documents that are actually relevant.
- Documents that are skipped or coded neutral in the validation queue are treated as relevant documents when calculating Elusion Rate. Therefore, coding all documents in the elusion sample as positive or negative guarantees the statistical validity of the calculated elusion rate as an estimate of the entire discard-pile elusion rate.

-
- **Eluded Documents**—the estimated number of relevant documents that have not been found. This is calculated by multiplying the sample elusion rate by the number of documents in the discard pile. The range listed below it applies the margin of error to the document count.
 - **Precision**—the estimated percentage of documents that would be responsive in a production that includes all documents coded positive, plus all uncoded documents at or above the Positive Cutoff. This statistic is only calculated if you set a cutoff when creating the validation queue.
 - **Recall**—percentage of documents that were coded relevant out of the total number of relevant documents, both coded and uncoded. The range listed below it applies the margin of error to the percentage.
 - **Richness**—the percentage of relevant documents across the entire review queue. The range listed below it applies the margin of error to the percentage.

For more information about how these statistics are calculated, see [Review validation statistics on page 52](#).

7.7.1 Recalculating validation results

If you have re-coded any documents from the validation sample, you can recalculate the results without having to re-run validation. For example, if reviewers had initially skipped documents in the sample or coded them as neutral, you can re-code those documents outside the queue, then recalculate the validation results to include the new coding decisions.

When Review Center detects that documents have been coded outside the queue and it's ready for completion, a banner appears with a button to calculate statistics. You can also manually recalculate them at any time.

To recalculate validation results:


1. On the right side of the Queue Summary section, click on the three-dot menu and select **Recalculate Validation**.
2. Click **Recalculate**.

7.7.2 Viewing results for previous validation queues

After you have run validation on a queue, you can switch back and forth between viewing the statistics for the main queue and any linked validation queues that were completed or rejected. Viewing the statistics for linked queues does not affect which queue is active or interrupt reviewers.

To view linked queues:

1. Click the triangle symbol near the right side of the Queue Summary section.



A drop-down menu listing all linked queues appears.

2. Select the queue whose stats you want to view.

When you're done viewing the linked queue's stats, you can use the same drop-down menu to select the main queue or other linked queues.

7.8 How adding or changing documents affects validation

Typically, review validation is linear: The administrator sets up the validation sample, the reviewers code the sample, and the results are calculated from those documents. However, if documents are added or removed, coded documents are re-coded, or other things happen to change the queue being validated, this can affect the validity of the results.

7.8.1 Scenarios that require recalculation

The following scenarios can be fixed by recalculating statistics:

- Changing coding decisions on documents within the validation sample
- Changing already-coded documents outside the sample from positive to negative or negative to positive
- Adding already-coded documents to the queue after validation starts

In these cases, the sample itself is still valid, but the numbers have changed. For these situations, recalculate the validation results to see accurate statistics.

For instructions on how to recalculate results, see [Recalculating validation results on the previous page](#).

7.8.2 Scenarios that require a new validation queue

The following scenarios cannot be fixed by recalculation:

- Adding uncoded or neutral documents to the queue after validation starts
- Changing positive- or negative-coded documents outside the sample to skipped or neutral

In both of these cases, this means that the validation sample is no longer a random sample of all uncoded or neutral documents. For these situations, we recommend starting a new validation queue.

8 Review validation statistics

Review Center provides several metrics for evaluating your review coverage: elusion, richness, recall, and precision. Together, these metrics can help you determine the state of your Review Center project.

Once you have insight into the accuracy and completeness of your relevant document set, you can make an educated decision about whether to stop the Review Center workflow or continue review.

This page gives a practical overview of how validation statistics are calculated. For more detailed equations, see [Review Center Validation Deep Dive](#) on the Community site.

See these related pages:

- For instructions on how to run Project Validation, see [Review validation on page 44](#).
- For aiR Review Prompt Criteria validation statistics, see Prompt Criteria validation statistics in the aiR Review guide.

8.1 Defining elusion, recall, richness, and precision

Validation centers on the following statistics. For all of these, it reports on confidence intervals:

- **Elusion rate**—the percentage of unreviewed documents that are predicted as non-relevant, but that are actually relevant. The rate is rounded to the nearest single digit (tenth of a percent).
- **Recall**—the percentage of truly positive documents that were found by the review.
- **Richness**—the percentage of relevant documents across the entire review.
- **Precision**—the percentage of found documents that were truly positive. This statistic is only calculated if you set a cutoff when creating the validation queue.

In everyday terms, you can think of these as:

- Elusion rate: "How much of what we're leaving behind is relevant?"
- Recall: "How much of the relevant stuff did we find?"
- Richness: "How much of the overall document set is relevant?"
- Precision: "How much junk is mixed in with what we think is relevant?"

The calculations for elusion and recall change depending on whether the validation queue uses a cutoff. For more information, see [How setting a cutoff affects validation statistics on page 54](#).

For each of these metrics, the validation queue assumes that you trust the human coding decisions over machine predictions, and that the prior coding decisions are correct. It does not second-guess human decisions.

Note: Validation does not check for human error. We recommend that you conduct your own quality checks to make sure reviewers are coding consistently.

8.2 Groups used to calculate validation metrics

Validation divides the documents into groups based on two main distinctions:

- Whether or not the document has been coded.
- Whether or not the document is relevant.

Together, these distinctions yield four buckets:

- **Coded, not relevant**—documents that have been coded as not relevant.
- **Coded, relevant**—documents that have been coded as relevant.
- **Uncoded, predicted not relevant**—documents that have not been coded and are predicted not relevant.
- **Uncoded, predicted relevant**—documents that have not been coded and are predicted relevant.



Coded
Not Relevant



Coded
Relevant



Uncoded,
Predicted
Not Relevant



Uncoded,
Predicted
Relevant

Coded and uncoded documents are separated as follows:

- **Coded documents**—have a positive or negative coding label.
- **Uncoded documents**—have not received a positive or negative coding label. This includes any documents that:
 - have not been reviewed yet.
 - are being reviewed at the moment the validation starts, but their coding has not been saved yet.
 - were skipped.
 - received a neutral coding label.

Notes:

- These buckets are determined by a document's status at the start of Project Validation. For the purpose of these calculations, documents do not "switch buckets" during the course of validation.
 - If you choose not to set a cutoff for your queue, buckets 3 and 4 are combined, and all uncoded documents are predicted not relevant. For more information, see [How setting a cutoff affects validation statistics on the next page](#).
-

8.2.1 Variables used in the full calculations

When Review Center calculates the validation metrics, it also subdivides buckets 3 and 4 based on two extra distinctions:

- Whether documents were included in the sample.
- Whether the coding decisions on sample documents matched their predictions.

Together with the first two buckets, this creates a total of ten variables used in the equations:

Variable	Full Name	Definition	Bucket Grouping
CN	Coded Negative (PRQ)	Documents coded non-responsive in the main prioritized review queue (PRQ) population	Bucket 1
CP	Coded Positive (PRQ)	Documents coded responsive in the main PRQ population	Bucket 2
FN _s	False Negatives (Sample)	Documents predicted non-responsive but coded responsive in the sample	Sample documents coded responsive in bucket 3
FP _s	False Positives (Sample)	Documents predicted responsive but coded non-responsive in the sample	Sample documents coded non-responsive in bucket 4
PN	Predicted Negative (PRQ)	Documents predicted non-responsive in the main PRQ population	Bucket 3
PN _s	Predicted Negative (Sample)	Documented predicted non-responsive in the sample (FN _s + TN _s)	Sample documents in bucket 3
PP	Predicted Positive (PRQ)	Documents predicted responsive in the main PRQ population	Bucket 4
PP _s	Predicted Positive (Sample)	Documented predicted responsive in the sample (TP _s + FP _s)	Sample documents in bucket 4
TN _s	True Negatives (Sample)	Documents predicted non-responsive and coded non-responsive in the sample	Sample documents coded non-responsive in bucket 3
TP _s	True Positives (Sample)	Documents predicted responsive and coded responsive in the sample	Sample documents coded responsive in bucket 4

When you click on the metrics in the Validation Stats panel, you can see the equations using these variables, as well as a table of the variables' values for your queue.

Note: Different regions and industries may use "relevant," "responsive," or "positive" to refer to documents that directly relate to a case or project. For the purposes of these calculations, the terms are used interchangeably.

Because this page is geared towards a general overview, the explanations here focus on the main four buckets. For a more granular breakdown using the full equations and variable set, see [Review Center Validation Deep Dive](#) on the Community site.

8.3 How setting a cutoff affects validation statistics

When you set up a validation queue, you have the choice of setting a cutoff. This cutoff rank is used as the dividing line between documents that are predicted positive or predicted negative. Setting a cutoff also enables the Precision statistic.

Depending on your choice, the calculations change as follows:

- If you set a cutoff, Review Center assumes that you still expect some of your uncoded documents to be positive. All statistics are calculated with the assumption that you plan to produce all positive-coded documents, plus all uncoded documents at or above the Positive Cutoff rank (buckets 2 and 4). These documents are considered "found."
- If you do not set a cutoff, Review Center assumes that you do not expect any uncoded documents to be positive. All statistics are calculated with the assumption that you plan to produce only the positive-coded doc-

uments (bucket 2). Only bucket 2 is considered "found." Without a cutoff, there is no bucket 4 - all uncoded documents are treated as bucket 3.

Whether or not you should use a cutoff depends on your review style. TAR1-style reviews, which focus on refining the model until it can be trusted to predict some of the positive documents, usually use a cutoff. TAR2-style reviews, which focus on having humans code all positive documents, usually do not use a cutoff.

8.3.1 High versus low cutoff

If you choose a high cutoff, this generally increases the precision, but lowers recall. If you choose a low cutoff, this generally increases the recall, but lowers precision.

In other words, choosing a high cutoff makes it likely that a high percentage of the documents you produce will be positive. However, it also increases the likelihood that some positive documents will be mistakenly left out. Conversely, if you choose a low cutoff, it's more likely that you may produce a few negative documents. However, you have better odds of finding ("recalling") all of the positive documents.

8.4 Validation metric calculations

8.4.1 Elusion rate

This is the percentage of uncoded, predicted non-relevant documents that are relevant.

The equation for elusion rate is:

Elusion =
FORMULA

$$\frac{FN_s}{FN_s + TN_s}$$

When breaking this down using the bucket method, it can be simplified as:

Relevant documents in



(NC - Not Rel)

Elusion =

All documents in



(NC - Not Rel)

Elusion rate = (Relevant documents in bucket 3's sample) / (All documents in bucket 3's sample)

Elusion measures the "error rate of the discard pile" — meaning, the relevant document rate in bucket 3. Documents that were coded relevant before starting project validation are not included in the calculation, regardless of their rank score. Documents coded outside of the queue during validation count as "eluded" documents.

If you do not set a cutoff for your validation queue, this is calculated as the number of documents in the validation sample that were coded relevant, divided by the entire size of the sample.

For more details, see [Review Center Validation Deep Dive](#) on the Community site.

8.4.1.1 Elusion rate versus eluded documents

In the Review Center dashboard, a metric also appears for eluded documents. This is the number of uncoded, predicted non-relevant documents that are relevant.

Eluded Docs =
FORMULA

Elusion Rate × PN

Eluded documents = Elusion rate * number of documents predicted negative

The elusion rate is a percentage, while the eluded documents statistic is an estimated document count.

8.4.2 Recall

Recall is the number of documents that were either previously coded or correctly predicted relevant, divided by the total number of documents coded relevant in any group.

The equation for recall is:

Recall =
FORMULA

$$\frac{CP \left(\frac{PP_s + PN_s}{PP + PN} \right) + TP_s}{CP \left(\frac{PP_s + PN_s}{PP + PN} \right) + TP_s + FN_s}$$

When breaking this down using the bucket method, it can be simplified as:

2

(C - Rel)

+ Relevant documents in

4

(NC - Rel)

Recall =

2

(C - Rel)

+ Relevant documents in

3

(NC - Not Rel)

and

4

(NC - Rel)

Recall = (Bucket 2 + relevant documents in bucket 4) / (Bucket 2 + relevant documents in buckets 3 and 4)

Recall measures the percentage of truly positive documents that were found by the review. Recall shares a numerator with the precision metric, but the denominators are different. In recall, the denominator is "what is truly relevant;" in

precision, the denominator is "what we are producing." Documents coded outside of the queue during validation count against recall.

If you do not set a cutoff for your validation queue, this is calculated as the number of documents that were previously coded relevant, divided by the total number of documents coded relevant in any group.

For more details, see [Review Center Validation Deep Dive](#) on the Community site.

8.4.3 Richness

This is the percentage of relevant documents across the entire review.

The equation for richness is:

Richness =
FORMULA

$$\frac{CP + (PP + PN) \frac{TP_s + FN_s}{TP_s + FP_s + FN_s + TN_s}}{CP + CN + PP + PN}$$

In the full equation, the numerator uses the sample to estimate the number of relevant documents in buckets 3 and 4. When breaking this down using the bucket method, it can be simplified as:


(C - Rel)

+ Relevant documents in


(NC - Not Rel)

and


(NC - Rel)

Richness = $\frac{\hspace{10em}}{\text{All buckets}}$

Richness = (Bucket 2 + any relevant documents found in buckets 3 and 4) / (All buckets)

Similar to recall, review validation estimates the number of relevant documents in bucket 4 by multiplying the estimated elusion rate by the number of uncoded documents. This is only done for the top half of the formula. For the bottom half, review validation only needs to know the size of the review.

For more details, see [Review Center Validation Deep Dive](#) on the Community site.

8.4.4 Precision

Precision is the percentage of truly positive documents out of all documents that were expected to be positive. Documents that were predicted positive, but coded negative during validation, lower the precision percentage.

The equation for precision is:

Precision =
FORMULA

$$\frac{CP + \left(PP \times \frac{TP_s}{TP_s + FP_s} \right)}{CP + PP}$$

In the full equation, the numerator uses the sample to estimate the number of relevant documents in bucket 4. When breaking this down using the bucket method, it can be simplified as:

$$\text{Precision} = \frac{\begin{array}{c} \text{Bucket 2} \\ (C - Rel) \end{array} + \begin{array}{c} \text{Relevant documents in} \\ \text{Bucket 4} \\ (NC - Rel) \end{array}}{\begin{array}{c} \text{All documents in} \\ \text{Bucket 2 and} \\ \text{Bucket 4} \\ (C - Rel) \quad (NC - Rel) \end{array}}$$

Precision = (Bucket 2 + relevant documents in bucket 4) / (All documents in buckets 2 and 4)

The precision statistic is only calculated if you set a cutoff when creating the validation queue. It assumes that you plan to produce all documents coded positive, plus all uncoded documents at or above the Positive Cutoff.

For more details, see [Review Center Validation Deep Dive](#) on the Community site.

8.5 How the validation queue works

When you start validation, the system puts all sampled documents from buckets 3 and 4 into the queue for reviewers to code.

Documents coded during project validation do not switch buckets during the validation process. Documents that started in buckets 3 and 4 are still considered part of 3 and 4 until validation is complete. This allows the system to keep track of correct or incorrect predictions when calculating metrics, instead of lumping all coded documents in with those which were previously coded.

Review Center reports statistics after all documents in the sample are reviewed. A document is considered reviewed if a reviewer has viewed the document in the Viewer and has clicked Save or Save and Next.

8.6 How validation handles skipped and neutral documents

We strongly recommend coding every document in the validation queue as relevant or non-relevant. Skipping documents or coding them neutral lowers the randomness of the random sampling, which introduces bias into the validation statistics. To counter this, Review Center gives conservative estimates. Each validation statistic counts a skipped or neutral document as an unwanted result.

The following table shows how skipped or neutral documents negatively affect each statistic.

Skipped or Neutral Document	Effect on Elusion	Effect on Recall	Effect on Richness	Effect on Precision
• Low-ranking (validation with cutoff) • Any document in a validation without cutoff	Increases elusion rate (Counts as relevant)	Lowers recall rate (Counts as relevant)	Raises richness estimate (Counts as relevant)	(No effect)
• High-ranking (validation with cutoff)	(No effect)	Lowers recall rate slightly (Counts as if it weren't present)	Raises richness estimate (Counts as relevant)	Lowers precision rate (Counts as non-relevant)

9 Reusing saved models

Saved models in Review Center provide the ability to take the knowledge, or the training, from one Review Center queue and re-use it in another queue. A saved model contains how many times a word occurred and what those words are, essentially remembering what was relevant, what was irrelevant, and how those were defined. With that information, you can use it to find relevant documents in a new queue or workspace.

9.2 What a saved model contains

A saved model contains the training from a previous Review Center queue, which includes:

- What words occurred.
- How many times words occurred.
- Which words generally predicted relevance or irrelevance, and a “relevance score” for each word.

The model does not contain any actual documents from the original queue. Instead, it contains relevance predictors in an encrypted, digested form that cannot be accessed by human users.

9.2.1 How predictions work with multiple models

When a saved model is linked to a queue, Review Center makes its predictions by averaging the relevance scores stored in the linked model and the local model. The local model contains the scores for words based on all coding decisions within the queue.

If there are several linked models, Review Center first takes the mean of the scores within the linked models, then averages that result against the local model's scores.

For example, if linked model 1 assigns a relevance score of 30% to the word "housing," and linked model 2 assigns it a score of 80%, this averages to 55%. If you start a brand new queue with both of these models attached, Review Center scores the word "housing" at 55%.

If you attach those same models to a queue that already contains some coding decisions, the local model may already have an entry for the word "housing." If the local model scores it at 50%, Review Center averages 50% with 55%. This gives the word "housing" a final score of 52.5%.

9.2.2 Privacy considerations when reusing saved models

Saved models can be shared across workspaces within the same instance, regardless of client domain, as long as the person sharing the model has access to both workspaces.

There is some risk that the model will reveal aspects of its training indirectly based on how it classifies. For example, if a document that it predicts as relevant contains the name "Jennifer" and nothing else, users can assume that the

original queue had "Jennifer" in the source data and that it was considered relevant. However, the source documents themselves are not actually revealed. The model does not store any identifying information such as the name of the original workspace, the name of the queue, or the control numbers of the documents that trained the original model.

For more information about permissions related to saved models, see [Review Center security permissions on page 63](#).

9.3 Common use cases for saved models

If you handle cases with similar document types or subject matter, saved models can help you jump-start a new case and start reviewing relevant documents more quickly. Instead of training a new model from scratch, you can link one or more saved models to a new queue and immediately start coding documents that the model predicts as relevant. After the new queue is underway, you can choose either to continue with both the saved models and the local model built from newly coded documents, or you can remove the links to the saved models and continue coding with only the local model.

Some use cases include:

- Anti-trust cases in different jurisdictions.
- Cases with similar subject matter such as medical-based reviews, toxic workplace, or bribery.
- Serial litigations.
- Internal investigations for improper or risky behavior.

You can also use saved models to find particular document types across cases. Examples of these include:

- Culling junk documents such as spam emails.
- Culling office noise such as vacation chatter and party invitations.
- Finding specific categories of documents such as contracts or financial paperwork.


9.4 Creating a saved model

You can save a Review Center model that has at least five positive and five negative documents coded. Queues with more documents coded will have more fully developed models, so we recommend saving models from late-stage or completed Review Center queues.

For a list of required permissions, see [Review Center security permissions on page 63](#).

To create a saved model from an existing Review Center queue:

1. From the Review Center dashboard, select the queue.
2. On the right side of the Queue Summary section, click the three-dot **Menu** icon.



3. Select **Save as New Model**.
4. Fill out the following fields:
 1. **Name**—the name of your saved model. If you plan to save the model from this queue multiple times, consider including a version number or a date.
 2. **Description** (Optional)—identifying features such as the model's purpose or what workspace and queue it originated from.

5. Click **Save**.

When the save completes, a green success banner appears at the top of the dashboard.

Note: The Saved Models feature was released in March 2025. If you want to save a model from a queue older than that, refresh the queue first.

9.4.1 Creating from queues with linked models

If you create a saved model from a queue that already has a saved model linked to it, the newly saved model will contain the training from both the linked model and the local model.


9.5 Linking a saved model to a Review Center queue

After creating a saved model, you can link it to another Review Center queue to jump-start the coding predictions for the new queue. For most situations, we recommend linking models to a newly created or early-stage queue. However, it is possible to link them at any stage. Any documents that are already coded within the destination project will add to the relevance predictions, but they are not required for the model to build.

To link a saved model when creating a new queue, see [Creating a new queue or queue template on page 15](#).

To link a saved model to an existing queue or to switch models:

1. From the Review Center dashboard, select the queue.
2. On the right side of the Queue Summary section, click the three-dot **Menu** icon.




3. Select **Edit**.
4. Next to **Saved Model**, click **Select**.
The model selection options appear.
5. Select the model or models you want to link.
6. Click **Apply**.
7. Click **Save**.
8. Refresh the queue to make the changes take effect.

9.5.1 Removing a linked model from a queue

After you have coded enough documents in the queue for the local model to build, you can remove the linked model at any time. After you remove it, the queue's relevance predictions will be calculated only from coding decisions within the queue.

To remove a linked model from a queue:

1. From the Review Center dashboard, select the queue.
2. On the right side of the Queue Summary section, click the three-dot **Menu** icon.



3. Select **Edit**.
4. Next to **Saved Model**, click **Clear**.
5. Click **Save**.
6. Refresh the queue to make the changes take effect.

9.5.2 How linked models behave with ARM

When using the Archive, Move, Restore (ARM) tool, linked models will be retained if you archive and restore within the same instance. However, if you attempt to restore a workspace with a linked model in another instance, you will receive an error and will not be able to start or refresh the queue. To start or refresh the queue, remove the linked model from it.

9.6 Managing saved models

You can access your saved models on the Saved Models tab in your workspace. This tab shows all models that originate in this workspace, as well as all models that have been copied to this workspace from outside.

Spam Detector Model Model for detecting Spam, exported on 3-22	Created Date/Time 3/14/2025 5:06AM	GUID 112f4caf-0e49-4025-d6b4-08dd62b5d21e	  
Harassment Finder Model for identifying harassment.	Created Date/Time 3/14/2025 5:07AM	GUID 79cacd37-4466-4889-d6ba-08dd62b5d21e	  

The tab shows the following for each model:

- **Name**—the name of the saved model.
- **Description**—any description given to the model when it was saved.
- **Created Date/Time**—the date and time the model was originally saved or copied to this workspace.
- **GUID**—the model's Globally Unique Identifier (GUID). This ID stays the same for a model regardless of which workspace it appears in or whether it has been renamed.

If you have several saved models to manage, you may want to create a workspace to serve as a central model library. Copying all saved models to this library workspace, then re-sharing them to individual workspaces as needed, allows you to manage all of your saved models in one place.

After a model has been copied to the library workspace, it can be safely deleted from its original workspace. Copies of models are not linked to the original copy.

9.6.1 Copying saved models to another workspace

Models can be copied to workspaces within the same instance, regardless of client domain. To copy a model, you must have access to the destination workspace.

For more information on permissions, see [Review Center security permissions on the next page](#). For more information on model sharing and privacy concerns, see [Privacy considerations when reusing saved models on page 59](#).

To copy a model and make it available for use in another workspace:

1. On the Saved Models tab, click the **Copy To** icon.



A workspace list appears. This list includes any workspace you have access to in the instance.

2. Select the workspace or workspaces you want to share to, then click **Copy To**.

After copying, each workspace has its own separate copy of the model.

9.6.2 Editing or deleting saved models

Every saved model is independent. Editing or deleting a saved model in one workspace does not affect shared copies in other workspaces, nor does it affect the Review Center queue the model was saved from.

To edit a saved model:

1. From the Saved Models tab, click the **Edit** icon beside the model.



2. Edit the fields you want to change.
3. Click **Save**.

To delete a saved model:

1. On the Saved Models tab, click the **Delete** icon beside the model.



A confirmation message appears.

2. Click **Delete**.

9.6.3 Saved models from Active Learning

If your workspace contains trained models that were saved from the older Active Learning application, these have been automatically converted to saved models for Review Center. You can view these on the Saved Models tab alongside the other models.

10 Review Center security permissions

This page contains information on the security permissions required for creating and interacting with the Review Center application.

10.1 Feature permissions

Note: This feature supports feature permission actions. Each action grants access to a specific task within the feature. You can grant groups access to individual actions, allowing them to perform only the tasks they need. To grant access to the entire feature, select all of its actions. Manage feature actions on the Features tab in a group's workspace permissions. For instructions, see [Setting workspace permissions](#). For more information about feature permissions, see [Feature permissions](#).

Permissions

Workspace Security

Groups

Features

Object Security

Tab Visibility

Other Settings

0/3 of actions

☐ Imaging

User

2 presets available

☐ OCR

User

2 presets available

Review - Applications

☒ Review Center

Administrator

3/3 of actions

☐ Review Management

☐ Transcripts

0/2 of actions

General

☐ Application Management

User

2 presets available

☐ Search

0/6 of actions

☐ Tab Management

☐ Workspace Details

0/2 of actions

Review Center

Preset

Administrator

Description

Full access to Review Center including queue management and document review.

Actions

☒ Select all

☒ Manage Documents

☒ View Review Center

☒ Manage Review Center Queues

>> [Granular permissions mapping](#)

Review and Save

Revert All

Close

Permissions for Review Center include:

- **Manage Documents**—select to allow users to view and edit documents from the Review Queues tab.
- **View Review Center**—select to allow users to view the Review Center dashboard.
- **Manage Review Center Queues**—select to allow users to create, edit, and delete Review Center queues, manage queues and view their statistics, and review documents.

For more detailed information on what is included in each permission, see Feature permissions in the Admin guide.

10.1.1 Checking document permissions

Make sure the reviewer groups have permissions to the documents they need to review. Reviewers will only see documents they have access to. If any documents in the queue are in a secured documents folder, and a reviewer does not have permissions for it, those documents will not be checked out to the reviewer.

If a reviewer sees a message that there are no more documents to review in a queue, but there are uncoded documents left, check the document permissions.

For more information on document security, see Relativity object security in the Admin guide.

For more information on assigning reviewer groups to a queue, see:

- [Setting up the reviewer group on page 13](#)
- [Creating a new queue or queue template on page 15](#)

10.2 Granular permissions

10.2.1 Creating a Review Center queue or template

To create a Review Center queue or template, you need the following permissions:

Object Security	Tab Visibility
• Queue Refresh Trigger - View, Edit, Add • Review Center Queue - View, Edit, Add	• Review Library • Review Center

10.2.2 Editing and controlling Review Center queues

To edit an existing Review Center queue and use dashboard controls such as Prepare or Start, you need the following permissions:

Object Security	Tab Visibility
• Queue Refresh Trigger - View, Edit, Add • Review Center Queue - View, Edit	• Review Center

10.2.3 Deleting a Review Center queue or template

To delete a Review Center queue or template, you need the following permissions:

Object Security	Tab Visibility	Mass Operation
• Queue Refresh Trigger - View, Edit, Add • Review Center Queue - View, Edit, Delete	• Review Library	• Delete

10.2.4 Viewing the Review Center dashboard

To view the Review Center dashboard, you need the following permissions:

Object Security	Tab Visibility
• Review Center Queue - View	• Review Center

If you want a user group to only see specific queues on the dashboard, you can restrict a queue using item-level security on the Review Library tab. For more information, see [Security and permissions in the Admin guide](#).

10.2.5 Assigning batches

To assign batches from a Review Center queue, you need the following permissions:

Tab Visibility	Other Settings
• Review Center (if assigning from the dashboard) • Review Queues (if assigning from the Review Queues tab)	• Assign Batches

Users do not need this permission to self-assign batches from the Review Queues tab.

10.2.6 Managing saved models

To create a saved model from a Review Center queue, you need the following permissions:

Object Security	Tab Visibility
• Review Center Queue - View, Edit	• Review Center • Saved Models (only needed to view results)

To copy a saved model to another workspace, you must have access to the destination workspace. You also need the following permissions in the origin workspace:

Object Security	Tab Visibility
• Review Center Queue - View, Edit	• Saved Models

To delete a saved model, you need the following permissions:

Object Security	Tab Visibility
• Review Center Queue - View, Edit	• Saved Models

10.2.7 Tracking reviewer decisions from the Documents tab

To track reviewer coding decisions using the Documents tab or the Field Tree, you need the following permissions:

Object Security	Tab Visibility	Browsers
• Review Center Coding - View	• Documents	• Field Tree

Users with access to the Review Center dashboard can also track reviewer decisions using the All Documents table. For more information, see [Tracking reviewer decisions on page 33](#).

10.2.8 Reviewer permissions

A reviewer group accessing a Review Center queue and coding documents must have the following permissions:

Object Security	Tab Visibility
• Document - View, Edit • Review Center Queue - View	• Review Queues

10.2.9 Checking document permissions

Make sure the reviewer groups have permissions to the documents they need to review. Reviewers will only see documents they have access to. If any documents in the queue are in a secured documents folder, and a reviewer does not have permissions for it, those documents will not be checked out to the reviewer.

If a reviewer sees a message that there are no more documents to review in a queue, but there are uncoded documents left, check the document permissions.

For more information on document security, see [Relativity object security](#) in the Admin guide.

For more information on assigning reviewer groups to a queue, see:

- [Setting up the reviewer group on page 13](#)
- [Creating a new queue or queue template on page 15](#)

11 Review Center performance baselines

This page acts as a reference to track the general performance of Review Center in Relativity aiR.

Because of data and configuration differences, do not use this as a benchmark for what to expect in your own organization's environment. The results may not scale linearly.

11.1 Queue size recommendations

For the best user experience, we recommend limiting the queue size as follows. These numbers are based on extensive testing in Relativity aiR.

Volume Type	Limit	Notes
Max documents in saved search	5,027,765 documents	We recommend a maximum of roughly 5 million total documents in the queue's data source.
Max coded documents	1,000,000 documents	We recommend a maximum of roughly 1 million total coded documents in the queue's data source.
Max concurrent reviewers	100 reviewers	We recommend a maximum of roughly 100 people reviewing per queue at the same time.

11.2 Reviewer count recommendations

In general, up to 100 reviewers can review documents simultaneously in a single queue without the Viewer performance degrading. However, this varies by environment. Some users experience issues slightly below this number, while others have successfully supported 200 or more concurrent reviewers without significant impact.

11.3 Performance testing definitions

For these tests, the times listed include the time taken to populate the queue, build the model, and load the results into the Review Center dashboard.

The start and end times are measured as:

- **Start time**—the time the user clicked the button to start the job.
- **End time**—the time the last document became available in Review Center.

11.4 Overall build speed

In a queue with 1 million documents, 10 of which were pre-coded, an initial build took slightly over an hour. It populated data into the queue at a rate of 13.57 gigabytes of data per hour, or 920,000 documents per hour.

The tabulated data is as follows.

Documents in queue	Pre-coded documents	Total run time (h:m-m:ss)	GB per hour	Documents per hour
1,000,000	10	1:05:50	13.57	920,000

11.5 Build time variance with document coding and caching

The following tests used a queue that contained 1 million documents, but each test had a different number of pre-coded documents. These documents were randomly coded 50% responsive and 50% non-responsive.

Each test also recorded the time it took to complete a second build. After the initial run, later builds use cached document tokens, which substantially speeds up the process.

Test Number	Pre-coded documents	Time for initial run (h:mm:ss)	Time for subsequent run (h:mm:ss)
Test 1	10	1:05:50	0:17:18
Test 2	100,000	1:08:50	0:20:15
Test 3	200,000	1:06:05	0:19:28
Test 4	300,000	1:02:19	0:21:19
Test 5	400,000	0:59:42	0:21:58
Test 6	500,000	1:01:23	0:17:42
Test 7	600,000	0:58:06	0:23:51
Test 8	700,000	0:52:38	0:16:17
Test 9	800,000	0:56:36	0:17:05
Test 10	900,000	1:07:43	0:20:07
Test 11	1,000,000	0:58:43	0:21:07

11.6 Build time variance with more documents

The following tests used a queue that contained 5,027,765 documents. Like previous tests, each test had a different number of pre-coded documents. These documents were randomly coded 50% responsive and 50% non-responsive.

Each test also recorded the time it took to complete a second build. After the initial run, later builds use cached document tokens, which substantially speeds up the process.

Test Number	Pre-coded documents	Time for initial run (h:mm:ss)	Time for subsequent run (h:mm:ss)
Test 1	10	4:14:49	0:48:11
Test 2	300,000	4:27:43	0:49:11
Test 3	400,000	4:35:43	0:50:02
Test 4	700,000	4:53:09	0:47:46
Test 5	800,000	4:02:12	0:58:08
Test 6	900,000	4:34:30	0:52:27
Test 7	1,000,000	4:04:45	0:54:06

12 Active Learning application history

If your workspace includes projects from the older Active Learning application, you can review their statistics and results from the Active Learning History tab. This tab shows read-only data for each project.

This tab can be accessed through the Review Center dashboard. For more information, see [Viewing archived Active Learning projects on page 35](#).

To use this tab, make sure that both Active Learning and Review Center are installed in the workspace. Uninstalling Active Learning removes the project data.

12.1 Selecting an Active Learning project

To view the data for a project, select it from the drop-down menu at the top of the page. This menu includes all Active Learning projects from the current workspace.

12.2 Project Statistics section

The Project Statistics section shows overall statistics for the selected project.

It includes:

- **Coded Positive**—the number of documents coded with the positive designation on the review field (excludes additional reviewed by family).
- **Coded Positive Outside**—the number of documents coded outside the Active Learning queue with the positive designation on the review field (excludes additional reviewed by family).
- **Coded Negative**—the number of documents coded with the negative designation on the review field (excludes additional reviewed by family).
- **Coded Negative Outside**—the number of documents coded outside the Active Learning queue with the negative designation on the review field (excludes additional reviewed by family).

- **Coded Neutral**—the number of documents coded with a neutral designation on the review field (excludes additional reviewed by family).
- **Coded Neutral Outside**—the number of documents coded outside the Active Learning queue with a neutral designation on the review field (excludes additional reviewed by family).
- **Skipped**—the number of documents that were saved or had **Save and Next** selected with no coding decision supplied on the review field (excludes additional reviewed by family).
- **Project Size**—the total number of documents included in the Active Learning project.

12.3 Manually Selected section

The Manually Selected section shows the number of document coding decisions made outside the Active Learning queue. These are grouped by date.

The columns include:

- **Manually-selected Documents**—the number of documents coded outside the Active Learning queue.
- **Coded Positive**—the number of these documents coded with the positive designation on the review field.
- **Coded Negative**—the number of these documents coded with the negative designation on the review field.
- **Coded Neutral**—the number of these documents coded with a neutral designation on the review field.
- **Date submitted**—the date in UTC that the statistics were submitted.

12.4 Prioritized Review section

The Prioritized Review section shows the statistics for Prioritized Review queues. The total number of coded documents appears in the section title.

For every 200 documents that are coded (excluding additional reviewed by family), a new row appears in the table. The first row in the table provides a summary for the entire project.

The columns include:

- **Prioritized Review**—the set of documents the statistics apply to (excludes additional reviewed by family). The sum of the count of **Coded Positive**, **Coded Negative**, and **Skipped** documents should equal 200.
- **# of Reviewers**—the number of unique reviewers who reviewed documents in the Prioritized Review queue during this interval.
- **Coded Positive**—the number of documents coded with the positive designation on the review field (excludes additional reviewed by family).
- **Coded Negative**—the number of documents coded with the negative designation on the review field (excludes additional reviewed by family).
- **Coded Neutral**—the number of documents coded with a neutral designation on the review field (excludes additional reviewed by family).
- **Skipped**—the number of documents that were saved or had **Save and Next** selected with no coding decision supplied on the review field (excludes additional reviewed by family).
- **Index Health**—the number of index health documents reviewed in the Prioritized Review queue. These documents are excluded from the relevance rate calculation.
- **Highest Ranked**—the number of highly ranked documents reviewed in the Prioritized Review queue.

- **Highest Ranked Coded Positive**—the number of highly ranked documents that were coded with the positive designation in the Prioritized Review queue.
- **Relevance Rate**—the percentage of documents that were chosen for being highly ranked that were then confirmed as relevant by reviewers' coding decisions. You can calculate the relevance rate manually using the following formula: Highest Ranked Coded Positive / Highest Ranked.
- **Family Group Documents**—the number of family documents reviewed in the Prioritized Review queue.
- **Positive Family Group Documents**—the number of family documents coded with the positive designation in the Prioritized Review queue.

12.5 Coverage Review section

The Coverage Review section shows the statistics for Coverage Review queues. The total number of coded documents appears in the section title.

For every 200 documents that are coded, a new row appears in the table. The first row in the table provides a summary for the entire project.

The columns include:

- **Coverage Review**—the set of documents the statistics apply to.
- **# of Reviewers**—the number of unique reviewers who reviewed documents in the Coverage Review queue during this interval.
- **Coded Positive**—the number of documents coded with the positive designation on the review field.
- **Coded Negative**—the number of documents coded with the negative designation on the review field.
- **Coded Neutral**—the number of documents coded with a neutral designation on the review field.
- **Skipped**—the number of documents that were saved or had **Save and Next** selected with no coding decision supplied on the review field.

12.6 Project Validation History section

The Project Validation History section shows the results for projects that were validated. Each row represents a separate validation round.

The columns include:

- **Validation**—each review is called Elusion with Recall or Elusion Test plus a numeral. For example, the first Elusion with Recall is "Elusion with Recall 1," and the second is "Elusion with Recall 2."
- **Start Date**—the date and time when Project Validation was started. This appears in your computer's time zone.
- **Status**—whether the validation results were accepted or rejected.
- **Rank Cutoff**—the numeric cutoff for positive prediction, fixed before project validation begins.
- **Discard Pile Size**—the number of documents below the rank cutoff that are not coded when Project Validation was started.
- **Elusion Sample Size**—the number of documents sampled for elusion rate. This number is computed when Project Validation is started.
- **Elusion Coded Relevant**—the number of sampled documents from below the cutoff which were coded positive during project validation.

- **Elusion Coded Not Relevant**—the number of sampled documents from below the cutoff which were coded negative during project validation.
- **Elusion Documents Skipped/Coded Neutral**—the number of sampled documents from below the cutoff that were either saved with no coding decision on the review field, or saved with a neutral coding decision.
- **Pending Document Count**—the number of documents whose coding has changed since the last model build (prior to Project Validation starting). This includes documents coded in Project Validation, and documents coded through other means. For instance, if a reviewer manually codes documents after Project Validation is underway, these will be counted as Pending.
- **Elusion Rate (Range)**—the error rate of the discard pile. This is calculated as the percentage of sampled, previously uncoded documents from below the cutoff which are coded positive during Project Validation. The range applies this sampled rate to the entire discard pile, using the confidence level provided by the user and the margin of error calculated from sample size.
- **Confidence Level for Elusion**—user input or calculated when setting up Project Validation.
- **Elusion Margin of Error**—margin of error calculated based on the elusion sample size, the discard pile size, and the elusion rate on the validation sample.
- **Estimated Eluded Documents (Range)**—the projected number of eluded documents. This estimates the number of relevant documents you would miss if you produced all documents marked relevant, as well as those with ranks at or above the cutoff.
- **Recall Rate (Range at CL80%)**—the percentage of truly positive documents which were found by the Active Learning process. A document has been "found" if it was previously coded positive, or if it is uncoded with a rank at or above the cutoff. Recall is calculated on the sample, then estimated for the total population with a confidence level (CL) of 80%.
- **Precision Rate (Range)**—the percentage of found documents which are truly positive. A document has been "found" if it was previously coded positive, or if it is uncoded with a rank at or above the cutoff. Documents which were predicted positive but coded negative during validation will count against precision. Precision is calculated on the sample, then estimated for the total population with a confidence level (CL) of 80%.
- **Precision Margin of Error (CL80%)**—the margin of error for precision as estimated from the sample size, the equivalent portion of the whole project, and the observed precision rate on the validation sample.
- **Richness Rate (Range)**—the percentage of documents which are relevant (positive choice). This is calculated by dividing the number of positive-coded documents in the sample by the total number of documents in the sample. The range predicts the richness for the whole project, subject to a 95% confidence level.
- **Richness Margin of Error (CL95%)**—the margin of error for richness as estimated from the sample size, the whole project size, and the observed richness rate on the sample.
- **Estimated Total Relevant Documents**—the estimated number of relevant documents in the whole project. This is calculated by projecting the richness rate across the whole project.
- **Total Documents in Project**—the number of documents in the project at the time of project validation.

12.7 Model Updates section

The Model Updates section shows each time the Active Learning model built, as well as the document ranks for each build. Each row represents a separate model build.

The columns include:

- **Build Date**—the date and time the model built. This appears in your computer's time zone.
- **0-10**—the number of documents that ranked between 0 and 10 after this build.
- **11-20**—the number of documents that ranked between 11 and 20 after this build.
- **21-30**—the number of documents that ranked between 21 and 30 after this build.
- **31-40**—the number of documents that ranked between 31 and 40 after this build.
- **41-50**—the number of documents that ranked between 41 and 50 after this build.
- **51-60**—the number of documents that ranked between 51 and 60 after this build.
- **61-70**—the number of documents that ranked between 61 and 70 after this build.
- **71-80**—the number of documents that ranked between 71 and 80 after this build.
- **81-90**—the number of documents that ranked between 81 and 90 after this build.
- **91-100**—the number of documents that ranked between 91 and 100 after this build.

Proprietary Rights

This documentation (“**Documentation**”) and the software to which it relates (“**Software**”) belongs to Relativity ODA LLC and/or Relativity’s third party software vendors. Relativity grants written license agreements which contain restrictions. All parties accessing the Documentation or Software must: respect proprietary rights of Relativity and third parties; comply with your organization’s license agreement, including but not limited to license restrictions on use, copying, modifications, reverse engineering, and derivative products; and refrain from any misuse or misappropriation of this Documentation or Software in whole or in part. The Software and Documentation is protected by the **Copyright Act of 1976**, as amended, and the Software code is protected by the **Illinois Trade Secrets Act**. Violations can involve substantial civil liabilities, exemplary damages, and criminal penalties, including fines and possible imprisonment.

©2026. Relativity ODA LLC. All rights reserved. Relativity® is a registered trademark of Relativity ODA LLC.