Relativity one

Review Center Guide

May 16, 2025

For the most recent version of this document, visit our **documentation website**.

Table of Contents

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 RelativityOne documentation site.

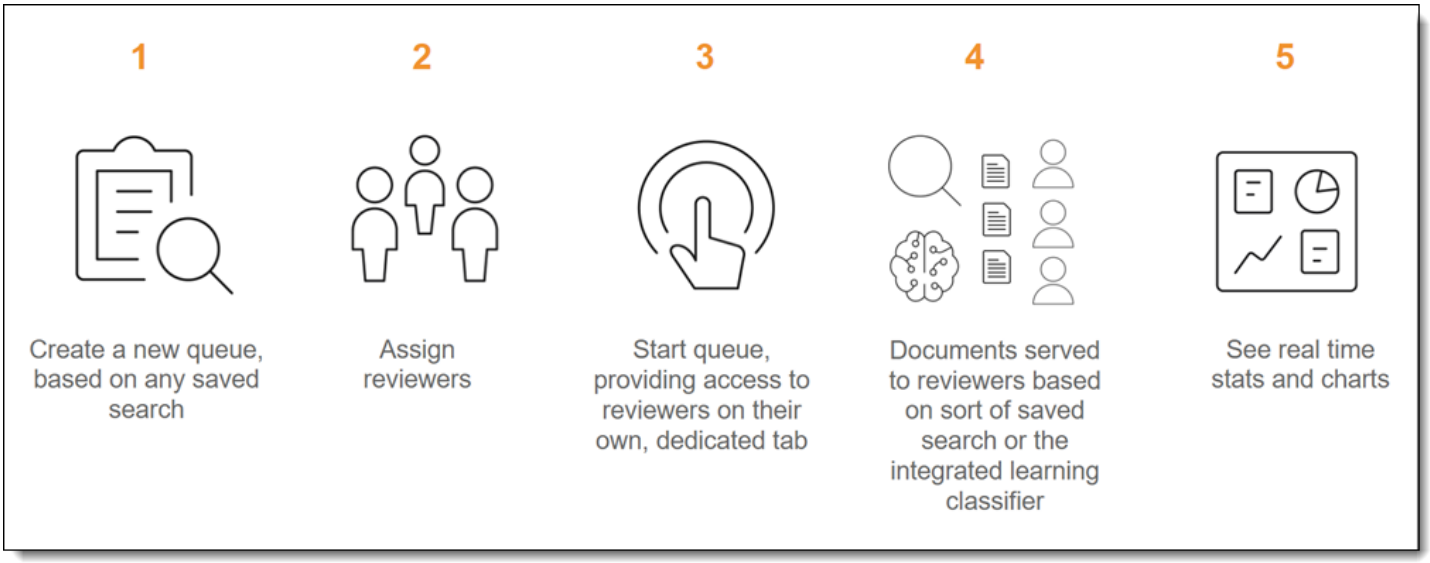

1.2 Review Center workflow

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.
4. Create or customize a review queue template.

After setup, create and manage the Review Center queue:

1. Create a new queue from the template.
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 32.

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.

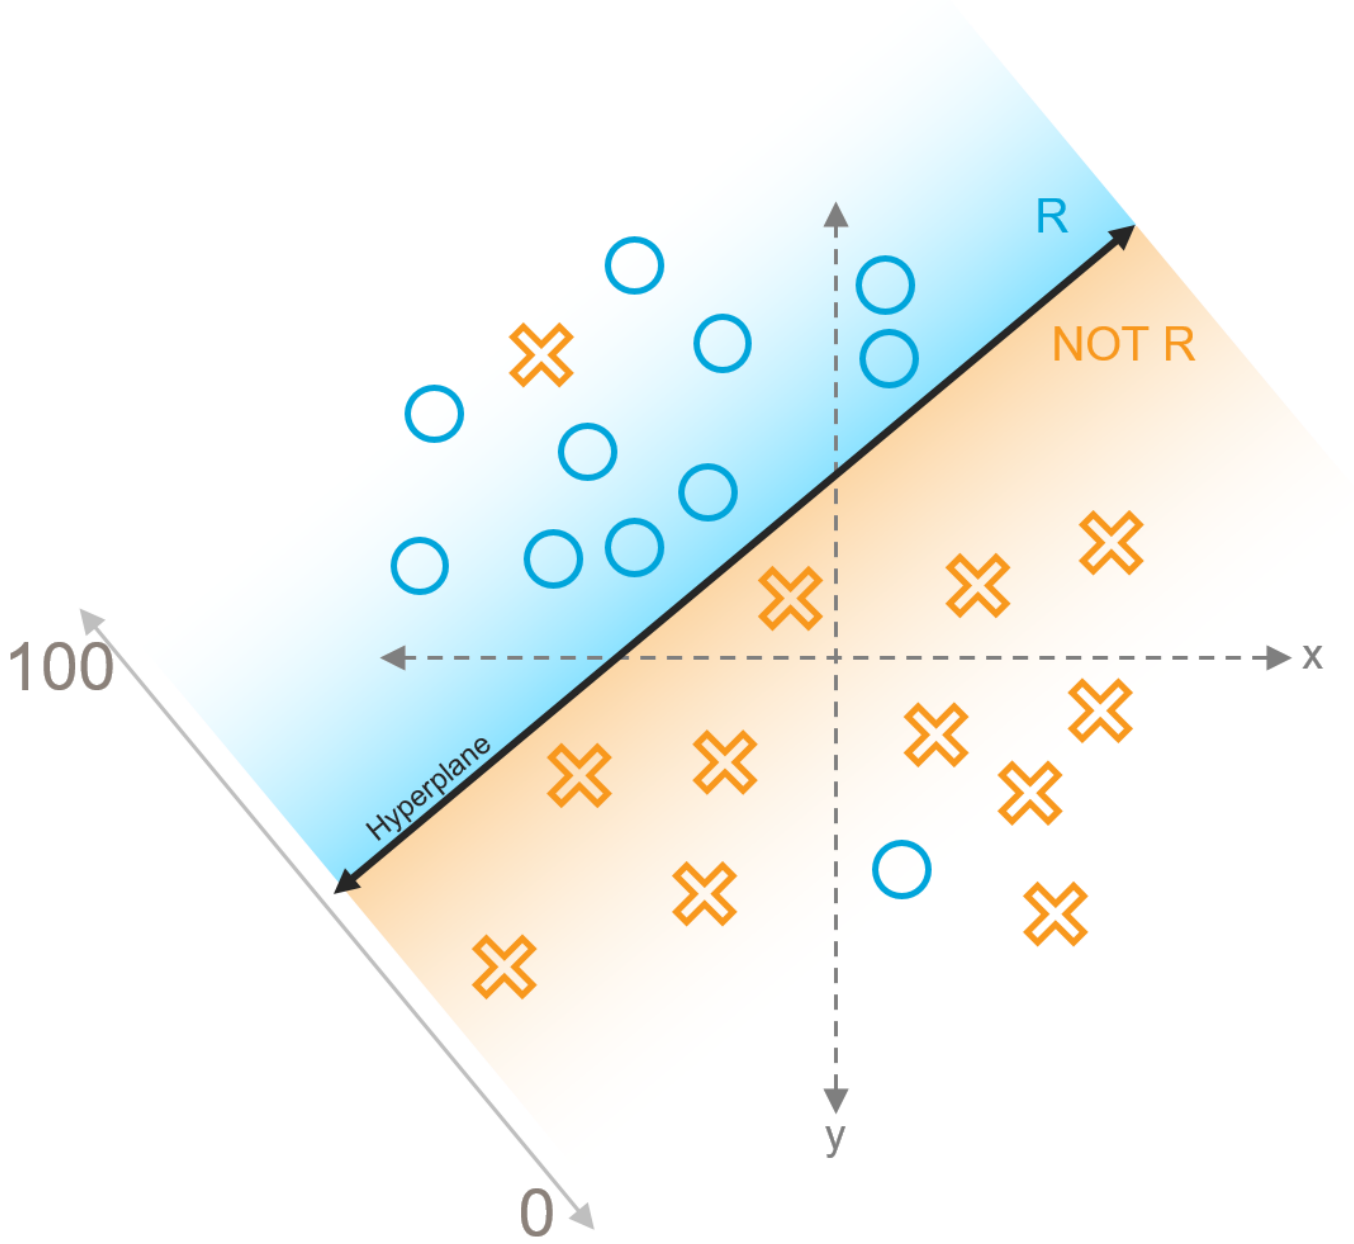

1.4 Language support in Review Center

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 Using Review Center versus batching

Review Center offers many benefits over batched reviews, including:

- **Built-in administrative reporting**—track progress and manage reviews in a single spot.
- **Time-saving templates**—shorten queue creation to a few clicks by creating templates for common workflows.
- **Streamlined assignment**—permissions are simplified, and documents are checked out automatically as each reviewer advances.
- **Simplified entry screen for reviewers**—reviewers enter queues with one click and have fewer distractions than on the standard Documents tab.
- **Easy to change**—you can update queues at any time, whether to add documents or to try out AI-powered review. None of these changes interrupt reviewer access.

If your organization uses custom reporting that requires a specific workflow, you may prefer to continue using batching for now. For other scenarios, though, users often find significant benefit in switching from batching to Review Center.

For more information on the traditional batching workflow, see Batches in the Admin guide.

1.6 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.

1.7 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: Create/Edit Queue Templates and Manage Queues
- Review Center Foundations: Key Functionality & Workflows
- Review Center Foundations: Managing Prioritized Review Projects

Knowledge base articles:

- Review Center - Multiple queues powered by the same classifier
- How to control what fields appear on the Documents card in the Review Center Viewer
- Using Review Center for Batching

White papers:

- Active Learning in RelativityOne 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**—create and manage queue templates.
- **Review Center**—create and manage queues 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.

2.2.2 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 49.

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 22.

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 22.

2.3 How document assignment works

By default, 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. For more information, see Keeping document families together on the next page.

2.3.1 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 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.

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 54.

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 a queue template

Templates are unassigned queues that can be used as the basis for building other queues quickly. Queue templates can also be saved as part of your workspace template.

Most fields which are required for queues, such as the Review Field, are not required for a template. This enables you to create generalized templates ahead of time and leave those decisions to the queue creator.

To create a queue template:

1. Navigate to the **Review Library** tab.
2. Click the **New Review Center Queue** button.
3. Fill out the template settings. For a full list, see Queue and template settings list on the next page.
   - Set the **Is Template** field to **On**. This adds the new template to your template list and makes many fields optional.
   - If a field value needs to be changed for each queue, such as Reviewer Groups, leave that field blank.
4. Click **Save**.
   The template now appears in the Review Library list.

2.7 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**.

2.8 Queue and template settings list

When you create or edit a Review Center queue or template, the following settings appear:

1. **Name**—the queue name reviewers will see.
2. **Is Template**—whether these queue settings will be used as a template for other queues.
   - If you are creating a template, toggle this **On**.
   - If you are creating a regular queue, toggle this **Off**.

   Note: This field exists for all queues. If you toggle the **Is Template** setting to **On** for a regular queue, it disappears from the dashboard and becomes usable as a template for other queues. Toggling it off again returns the queue to the dashboard. The queue keeps all of its statistics and coding decisions, but the queue state resets to Not Started.
3. **Template Description** (templates only)—enter notes about the template such as its intended use, comments about field settings, etc.
4. **Queue Type**—choose either **Saved Search** or **Prioritized Review**.
5. **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 18.
6. **Reviewer Groups**—select the permissions groups that include the reviewers for the queue. This field is not recommended for templates.

7. **Data Source**—select the saved search that contains the documents for your queue.

   Note: If you are using a prioritized review queue:
   - If you want to make sure documents are not added or removed from the queue, used a static saved search as the data source.
   - 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.

8. **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 49.
9. **Rank Output Field** (Prioritized Review only)—select the decimal field you created to hold the document rank scores.
10. **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.
    1. **Positive Choice**—select the choice that represents the positive or responsive designation.
    2. **Negative Choice**—Select the choice that represents the negative or non-responsive designation.

       Note: Any remaining choices are considered neutral.

11. **Positive Cutoff**—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.
12. **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.

13. **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.

14. **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.
15. **Queue Display Options**—select which statistics you want reviewers to see on the queue card in the Review Queues tab.
16. **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 11.

- 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.

17. **Text Precedence 1** (Prioritized Review only)—select the primary text field you want Review Center to analyze. By default, this is set to the Extracted Text field.
18. **Text Precedence 2** (Prioritized Review only)—select the text field you want Review Center to analyze if the primary text field is empty. This field is optional.
    - If you choose a value for Text Precedence 2, another field called Text Precedence 3 appears, and so on. 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.
19. **Queue 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 20.
    - 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.
20. **Reviewer Document 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 32.
    - If there are any conditions applied to the view, those conditions will also limit which documents appear in the panel.
21. **Reviewer Layout**—select the coding layout that you want reviewers to see by default when they enter the queue. If you do not choose a view, this defaults to the lowest ordered layout the reviewer has permission to access.
22. **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 Unknown 1, Unknown 2, and so on. These labels are assigned randomly every session. For example, a reviewer who is listed as Unknown 1 the first day might be listed as Unknown 4 another day.
    - Toggle this **Off** to show reviewer names on the dashboard.
23. **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 Creating a new queue from a template

To create a new queue using a queue template, use the **Add Queue** button on the Review Center dashboard.

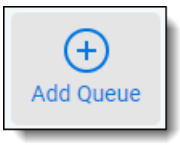

To create a new queue from template using the dashboard:

1. Navigate to the **Review Center** tab.
2. Click the **Add Queue** button.
3. Select the template you want to use, then click **Next**.
4. Edit the queue settings. For a full list, see Queue and template settings list on page 14.
    1. Under **Reviewer Groups**, choose one or more reviewer groups.
    2. In the other fields, check the default values filled in by the template. Change any values that should be different for this queue.
5. Click **Save**.
   The new queue appears as a tab in the banner at the top of the dashboard.

All queue settings can also be edited after creating the queue.

Note: After a queue has been created from a template, the two of them are no longer connected. You can edit the template without affecting the queue.

For information on starting, managing, and deleting queues, see Monitoring a Review Center queue on the next page.

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.

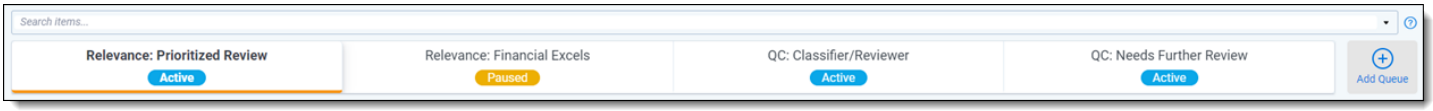

Below the queue name, each queue shows its status. The possible statuses are:

- **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 36.

At the right of the strip, the Add Queue button lets you quickly create new queues. For instructions, see Creating a new queue from a template on page 16.

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 30.

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.

3.1.2.1 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 on the next page.

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 21.

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.

Auto-refreshing the queue

If Queue 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.

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

Prioritized review queues refresh when 20% of documents in the queue have had positive or negative coding changes since the last queue refresh. The 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% of documents have been coded. These refreshes only happen after the queue has been started, and you can change this setting at any time.

For example, if 1000 documents were coded positive or negative at the last refresh, coding another 200 within the queue would trigger the next auto-refresh. If another 10 were coded, the queue would also 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.

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.

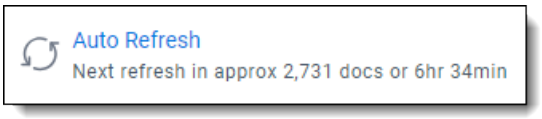

If you need to trigger an immediate refresh, click on the words **Auto Refresh** to trigger an additional manual refresh. For example, if new documents have been added to the saved search, you can click this to add them to the queue quickly instead of waiting until the next auto-refresh.

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.

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.

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 22.

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.

3.1.2.2 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.

3.1.2.3 Editing queues and other actions

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

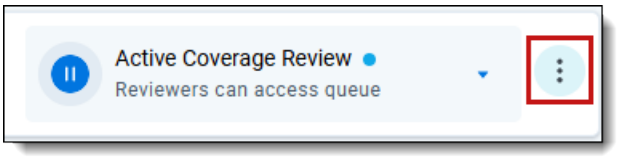

The menu options are:

- **Edit**—opens a modal to edit any of the queue settings.
  - For information on general edits, see Editing recommendations on the next page.
  - 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 49.
- **Set up Validation** (prioritized review queue only)—opens the options to create a review validation queue. For more information, see Review validation on page 36.
- **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.

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.

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 19.

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."

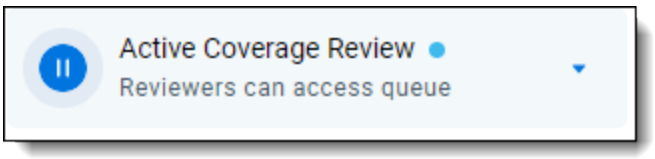

If Coverage Mode is turned **Off**, the button will refer to the queue as "Prioritized Review."

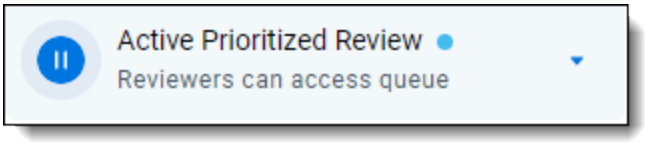

On the Queue History table, you can also see if your queue was in Coverage Mode or not during each queue refresh. For more information, see Queue History on page 26.

3.1.3 Review Progress section

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

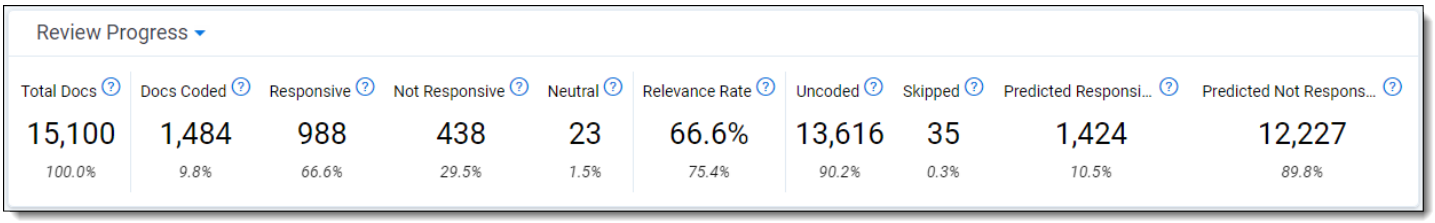

By default, the section shows a set of statistics that are calculated for all documents in the queue. By clicking the triangle next to the section name, you can select another view.

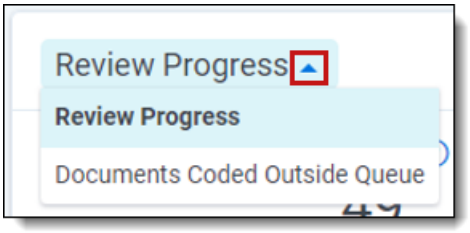

3.1.3.1 Review Progress view

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.

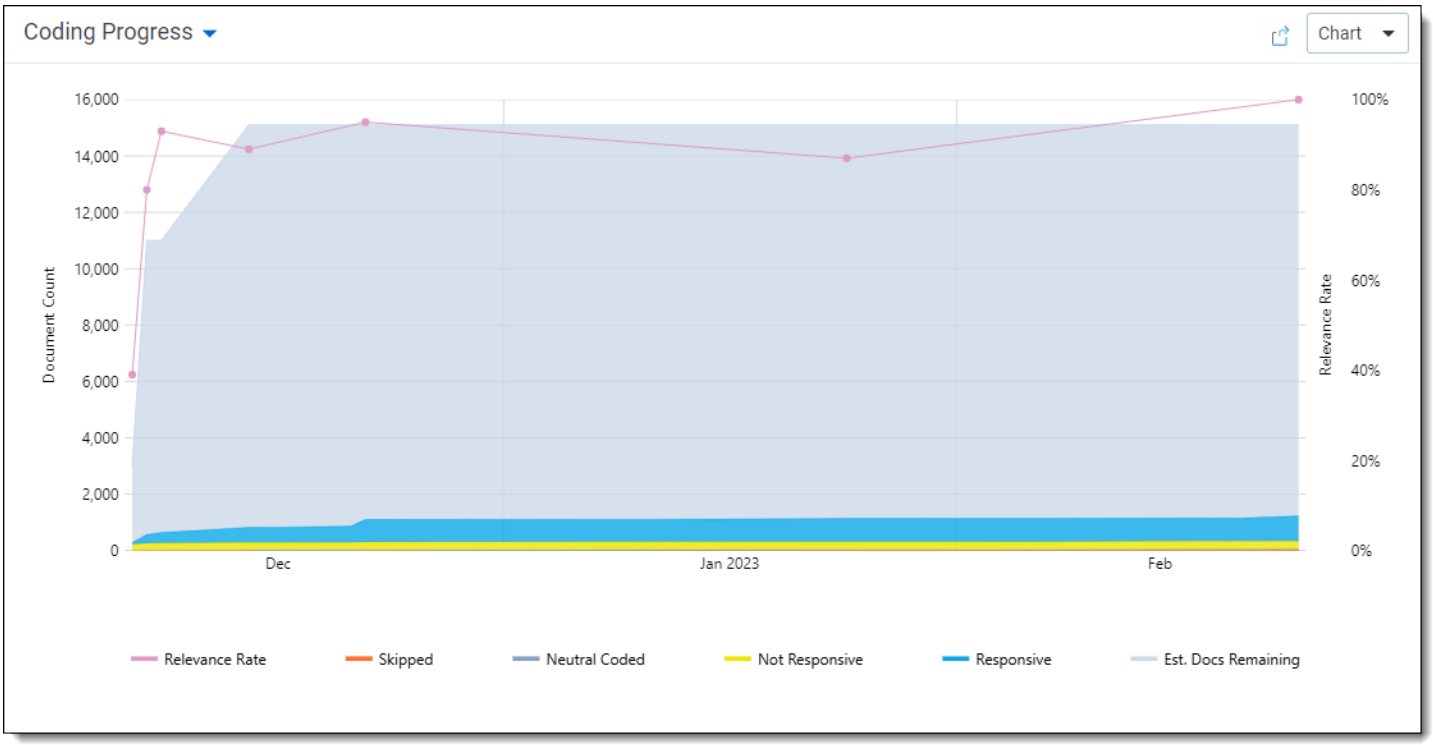

To navigate the visualization panel:

- To select a different visualization, click the **blue arrow** (▾) next to 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 download symbol ( ) on the upper right. Charts download as .png images, and tables download as .csv files.
- 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.

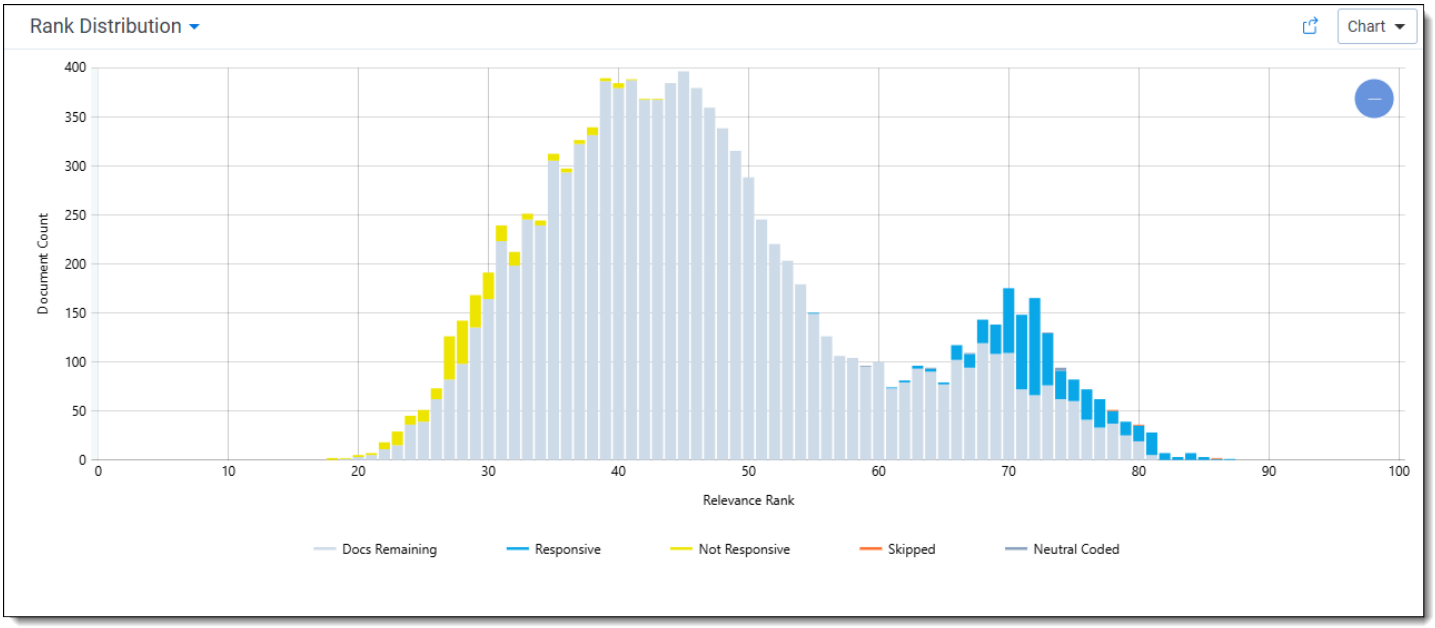

If you zoom out on the Rank Distribution chart, you may see documents with ranks below zero. These are documents that could not be classified. For more information, see Understanding document ranks on the next page.

3.2.3 Reviewed Documents table

The Reviewed Documents table shows which reviewer coded each document, how long the reviewer took, and how it was coded.

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.
- **Coded Time**—the check-in time for the document. If the document is still checked out, this is blank.
- **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).

- **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.

3.3 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 36.

3.4 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.

3.5 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.

3.6 Tracking reviewer decisions

You can view coding decisions made by each reviewer in the Reviewed Documents table. For more information, see Reviewed Documents table on page 27.

Alternatively, you can also use the following methods.

3.6.1 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.

The field names are:

- **Review Center Coding::Reviewed On**—the date of the coding decision. Dates are based on the UTC time zone.
- **Review Center Coding::Reviewed By**—the name of the reviewer who made the coding decision.
- **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::Value**—the reviewer's coding decision.
- **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).

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 RelativityOne documentation site.

3.6.2 Using the Field Tree

The Field Tree helps you get a quick overview of document coding decisions. 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.

3.6.3 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 RelativityOne documentation site.

When configuring the application:

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

3.7 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.

3.8 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.

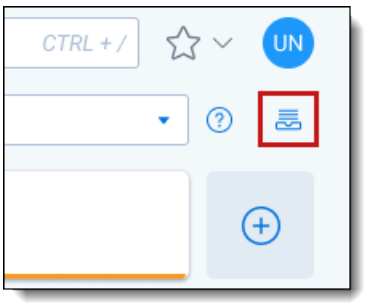

For detailed information on the Active Learning History tab, see Active Learning application history on page 59.

4 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.

4.1 Reviewing documents in the queue

To review documents in a queue:

1. Navigate to the **Review Queues** tab.
2. Each queue you are assigned to has a separate card. Locate the card with the same name as the queue you want.

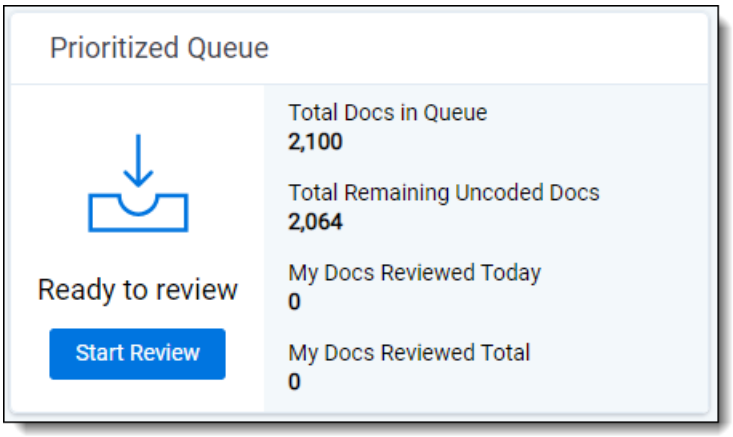

3. Click **Start Review**.
   This opens the document viewer.
4. Review the document as specified by your admin, then enter your coding choice.
5. Click **Save and Next**.
   The next document will appear for review.

If you do not see a Start Review 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.

4.2 Finding previously viewed documents

As you work through the queue, you can see documents you already reviewed in the queue by clicking on **Documents** in the left-hand navigation bar. This opens the Documents panel.

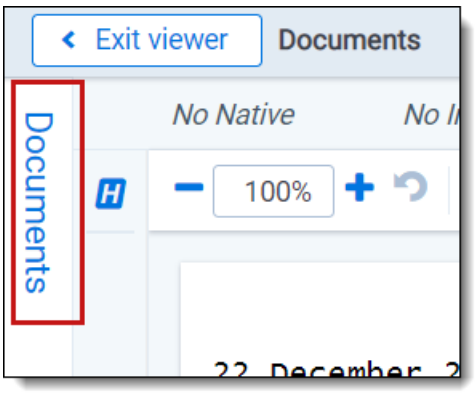

To view a document, click on its control number in the panel.

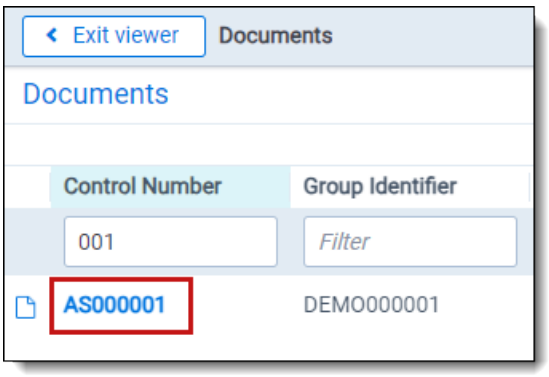

To return to your current document, click on the **Navigate Last** button in the upper right corner of the document viewer.

Note: 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.

4.3 Queue card statistics

If your admin has enabled it, you may see some statistics displayed on the queue cards.

The statistics you may see are:

- **Total docs in queue**—the total number of documents in this queue, across all reviewers.
- **Total remaining uncoded docs**—the total number of uncoded documents in this queue, across all reviewers.
- **My docs reviewed total**—how many documents you have reviewed total in this queue.
- **My docs reviewed today**—how many documents you have reviewed today in this queue. These are counted starting at midnight in your local time.

4.4 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 18.

4.5 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.

4.5.1 Coding according to the "four corners" rule

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.

4.5.1.1 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 typically 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.

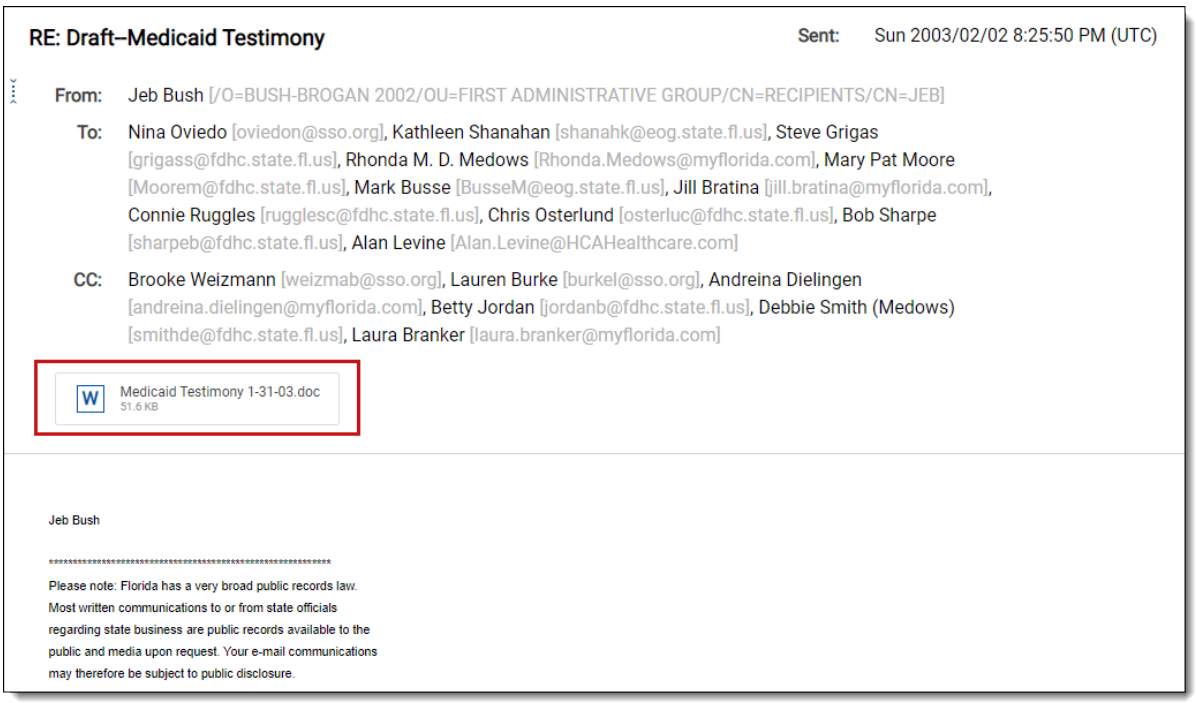
RE: Draft–Medicaid Testimony

Sent: Sun 2003/02/02 8:25:50 PM (UTC)

From: Jeb Bush [/O=BUSH-BROGAN 2002/OU=FIRST ADMINISTRATIVE GROUP/CN=RECIPIENTS/CN=JEB]

To: Nina Oviedo [oviedon@sso.org], Kathleen Shanahan [shanahk@eog.state.fl.us], Steve Grigas [grigass@fdhc.state.fl.us], Rhonda M. D. Medows [Rhonda.Medows@myflorida.com], Mary Pat Moore [Moorem@fdhc.state.fl.us], Mark Busse [BusseM@eog.state.fl.us], Jill Bratina [jill.bratina@myflorida.com], Connie Ruggles [rugglesc@fdhc.state.fl.us], Chris Osterlund [osterluc@fdhc.state.fl.us], Bob Sharpe [sharpeb@fdhc.state.fl.us], Alan Levine [Alan.Levine@HCAHealthcare.com]

CC: Brooke Weizmann [weizmab@sso.org], Lauren Burke [burkel@sso.org], Andreina Dielingen [andreina.dielingen@myflorida.com], Betty Jordan [jordanb@fdhc.state.fl.us], Debbie Smith (Medows) [smithde@fdhc.state.fl.us], Laura Branker [laura.branker@myflorida.com]

Medicaid Testimony 1-31-03.doc
51.6 KB

Jeb Bush

****

Please note: Florida has a very broad public records law. Most written communications to or from state officials regarding state business are public records available to the public and media upon request. Your e-mail communications may therefore be subject to public disclosure.

4.5.2 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.

5 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 43.
- For detailed equations, see Review Center Validation Deep Dive on the Community site.

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.

5.1 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.

5.2 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 25.

5.3 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 the queue.
   - If auto-refresh is turned on, turn it off.
   - If the queue is in the middle of refreshing, wait until the refresh has finished before starting validation.
   - If any documents are currently checked out to reviewers, release them. For more information, see Editing queues and other actions on page 21.
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 44.
   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.
   4. Set the sample size using three interconnected fields:
      1. **Sample Size**—this 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. **Margin of Error Estimate (Elusion)**—this calculates a size for the sample based on how accurate the Elusion statistic will be.
      3. **Margin of Error Estimate (Recall)**—this calculates a size for the sample based on how accurate the Recall statistic will be.

      Note: Each of these fields affects the others. For an explanation of how they work, see Choosing the validation settings below.
5. Click **Save**.

5.3.1 Choosing the validation settings

Validation always samples a specific number of documents, but there are three ways to choose the sample size:

1. You can specify exactly how many documents you want to sample. Review Center automatically calculates the estimated margins of error for both Elusion and Recall based on the sample size you select.
2. You can specify the desired margin of error for the elusion estimate and let Review Center calculate an appropriate sample size. Review Center also automatically calculates the corresponding recall margin of error.
3. You can specify the desired margin of error for the recall estimate and let Review Center calculate an appropriate sample size. Review Center also automatically calculates the corresponding elusion margin of error.

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 aim for a 95% confidence interval alongside the margin of error.

The estimated elusion margin of error depends only on the sample size, and vice versa. Their relationship to the estimated recall margin of error depends on the number of relevant documents that have already been coded and the current size of the discard pile. It may vary among different validation samples, even within the same review.

For more information on how validation statistics are calculated, see Review validation statistics on page 43.

5.3.2 Inherited settings

Each validation queue inherits these settings from the main queue:

- Queue Display Options
- Reviewer Document View
- Reviewer Layout
- Email Notification Recipients

To change them, edit the validation queue after creating it. For more information, see Editing a validation queue below.

5.4 Coding in a validation queue

Reviewers access the validation queue from the Review Queues tab like all other queues. Have reviewers code documents from the sample until all documents have been served up.

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 47.

5.5 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**.

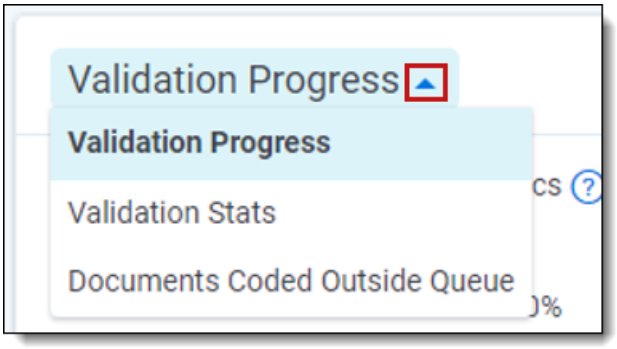

For more information on the validation statistics, see Reviewing validation results on page 41.

5.5.1 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:
    - **Reviewer Groups**
    - **Queue Display Options**
    - **Reviewer Document View**
    - **Reviewer Layout**
    - **Email Notification Recipients**
3. Click **Save**.

For descriptions of the queue settings, see Creating a Review Center queue on page 10.

5.5.2 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 21.

To see which documents are checked out to a reviewer, filter the Reviewed 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 Reviewed Documents table on page 27.

5.5.3 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.

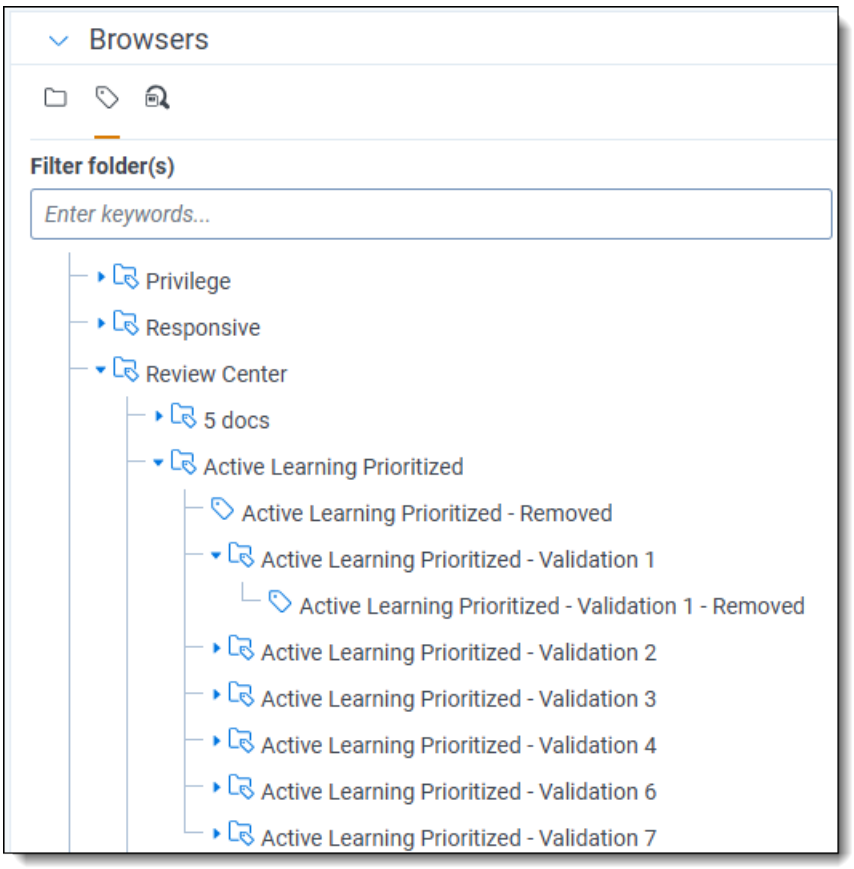

To view coding decisions for each document, add **Review Center Coding::Value** to the document view. For other optional fields, see Tracking reviewer decisions on page 29.

5.6 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 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 42.

5.6.1 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.

5.7 Reviewing validation results

After reviewers code all documents in the sample, the queue status changes to Complete. All 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.
  - 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 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 43.

5.7.1 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.

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**.

5.7.2 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.

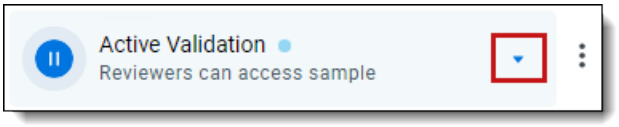

   A drop-down menu listing all linked queues appears.
2. Select the queue whose stats you want to view.

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.

5.8 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.

5.8.1 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.

5.8.2 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.

6 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.

For instructions on how to run Project Validation, see Review validation on page 36.

6.1 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 the next page.

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.

6.2 Groups used to calculate validation metrics

Validation divides the documents into groups based on two 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 but are not relevant.
- **Coded, relevant**—documents that have been coded and are 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.

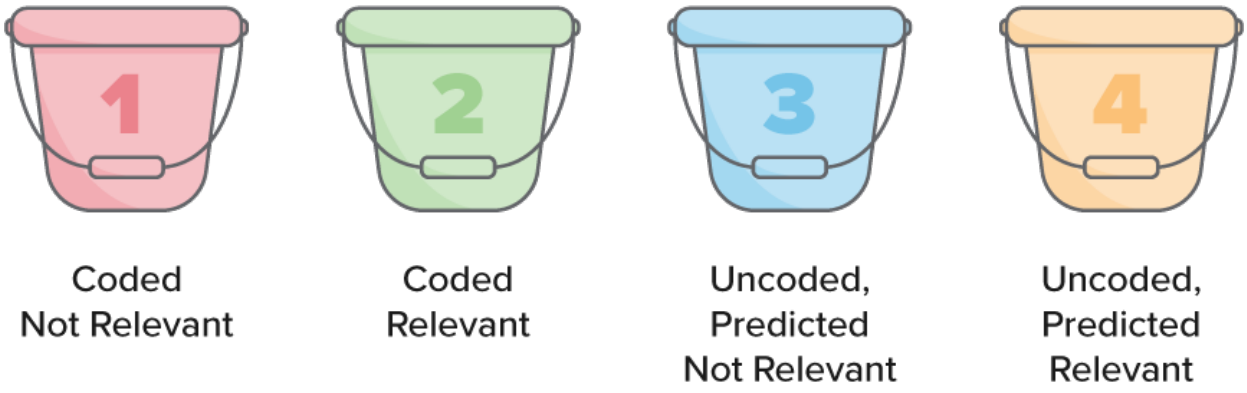

At the start of validation, the system knows exactly how many documents are in buckets 1 and 2.

- **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.

The system also knows how many documents are in buckets 3 and 4 altogether, but not the precise breakdown between the two buckets.

You could find out by exhaustively coding the uncoded documents, but that's time-consuming. Instead, review validation uses statistical estimation to find out approximately how many are in each bucket. This means that any statistics involving bucket 3 or 4 will include a confidence interval that indicates the degree of uncertainty about how close the estimate might be to the true value.

- 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 below.

6.3 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 documents (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 having humans code all positive documents, usually do not use a cutoff. TAR2-style reviews, which focus on refining the model until it can be trusted to predict some of the positive documents, usually use a cutoff.

6.3.1 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.

6.4 Validation metric calculations

6.4.1 Elusion rate

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

$$\text{Elusion} = \frac{\text{Relevant documents in } 3 \text{ (NC - Not Rel)}}{\text{All documents in } 3 \text{ (NC - Not Rel)}}$$

Elusion = (Relevant documents in bucket 3) / (All documents in bucket 3)

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.

6.4.2 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.

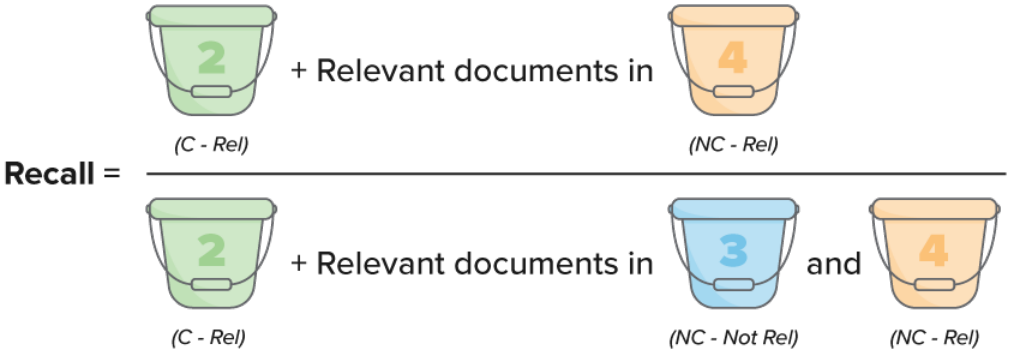

Recall = (Bucket 2 + relevant documents in bucket 4) / (Bucket 2 + relevant documents in buckets 3 and 4)

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.

6.4.3 Richness

This is the percentage of documents in the review that are relevant.

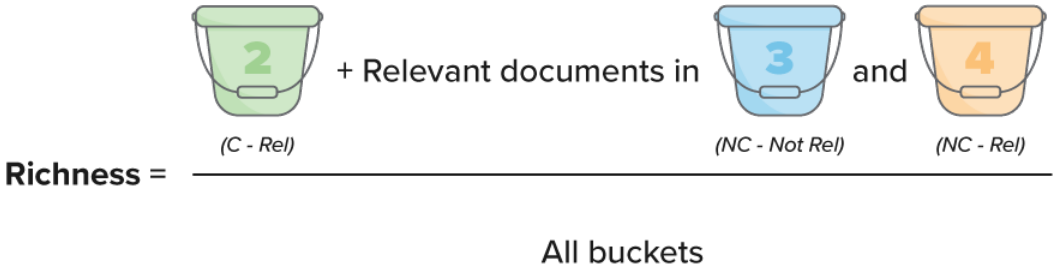

Richness = (Bucket 2 + any relevant documents found in buckets 3 and 4) / (All buckets)

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.

6.4.4 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, lower the precision percentage.

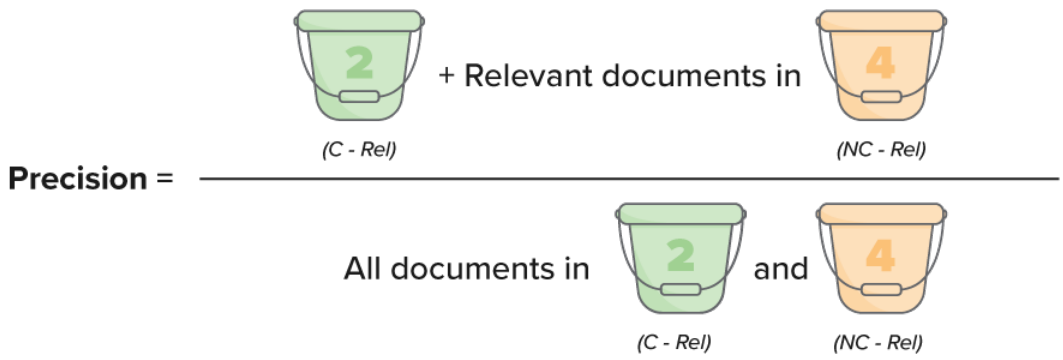

Precision = (Bucket 2 + relevant documents in bucket 4) / (All documents in buckets 2 and 4)

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 documents at or above the Positive Cutoff.

6.5 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.

6.6 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)

Skipped or Neutral Document	Effect on Elusion	Effect on Recall	Effect on Richness	Effect on Precision
■ **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)

7 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.

7.2 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.

7.2.1 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%.

7.2.2 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 54.

7.3 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.

7.4 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 54.

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.

7.4.1 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.

7.5 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 from a template on page 16.

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.

7.5.1 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.

7.5.2 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.

7.6 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.

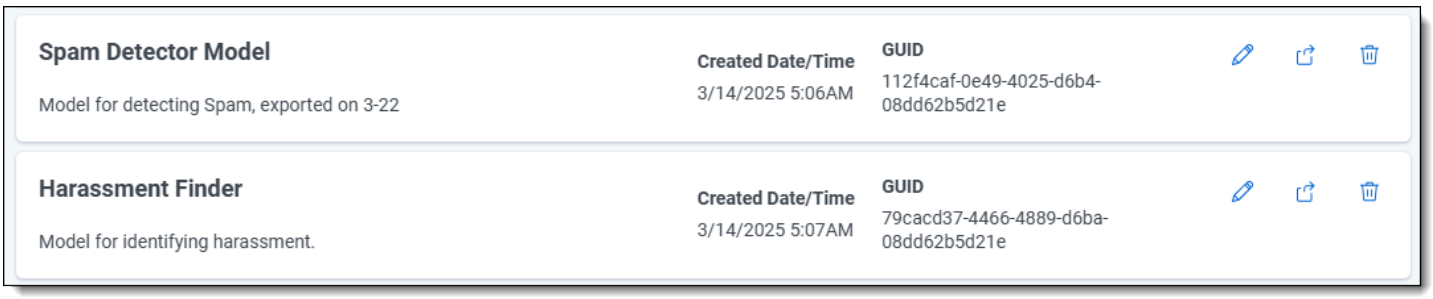

The tab shows the following for each model:

- **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.

7.6.1 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 page 54. For more information on model sharing and privacy concerns, see Privacy considerations when reusing saved models on page 49.

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.

7.6.2 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**.

7.6.3 Saved models from Active Learning

If your workspace contains trained models 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.

8 Review Center security permissions

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

Note: As of February 2025, the new Feature Permissions redefines Relativity's security management by shifting the focus from Object Types and Tab Visibility to feature-based permissions. This new method is simply another option; any feature-specific permissions information already in this topic is still applicable. This new interface enables administrators to manage permissions at the feature level, offering a more intuitive experience. By viewing granular permissions associated with each feature, administrators can ensure comprehensive control, ultimately reducing complexity and minimizing errors. For details see Instance-level permissions and Workspace-level permissions.

8.1 Creating a Review Center template or queue

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

Object Security	Tab Visibility
■ **Queue Refresh Trigger** - View, Edit, Add ■ **Review Center Queue** - View, Edit, Add ■ **Workspace** - Edit Security	■ Review Library ■ Review Center

8.2 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 ■ **Workspace** - Edit Security	■ Review Center

Note: The **Workspace** - Edit Security permission is only required to edit the assigned reviewer group.

8.3 Deleting a Review Center template or queue

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

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

Object Security	Tab Visibility	Mass Operation
Delete		

8.4 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.

8.5 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

8.6 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 Reviewed Documents table. For more information, see Tracking reviewer decisions on page 29.

8.7 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

8.7.1 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 12
- Creating a new queue from a template on page 16

9 Review Center performance baselines

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

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.

9.1 Queue size recommendations

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

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.

9.2 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.

9.3 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:mm:ss)	GB per hour	Documents per hour
1,000,000	10	1:05:50	13.57	920,000

9.4 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

9.5 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

10 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 30.

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.

10.1 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.

10.2 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.

10.3 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.

10.4 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.

10.5 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.

10.6 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.

10.7 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.