Relativity one

Review Center Guide

March 12, 2024

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 Getting Started with Review Center on-demand training on the RelativityOne documentation site.

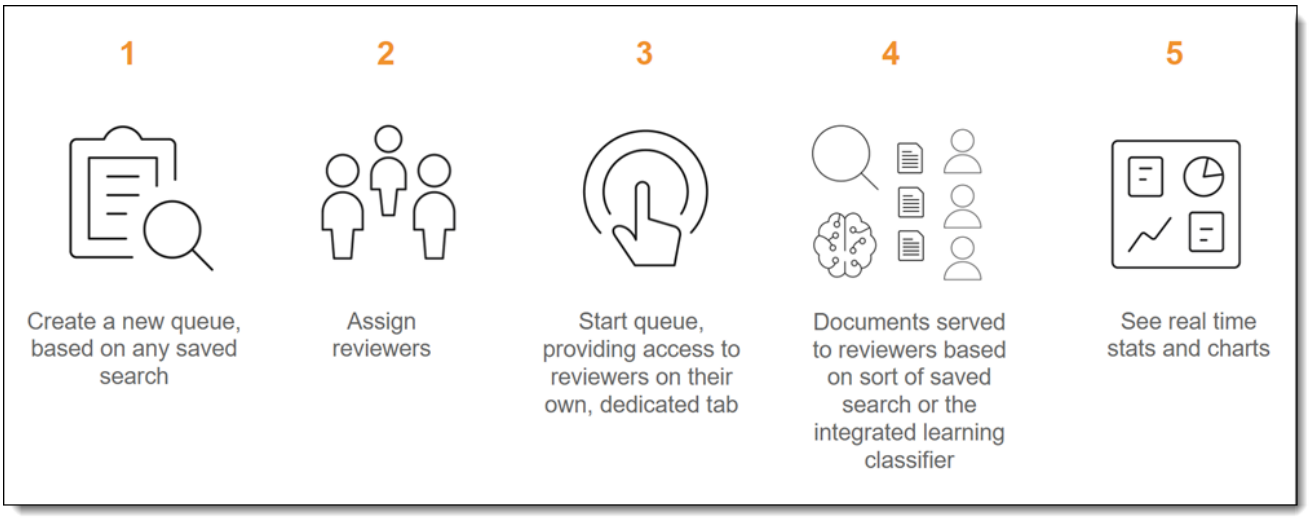

1.2 Review Center workflow

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 25.

For detailed instructions on setting up Review Center, see Creating a Review Center queue on page 8.

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.

Because the classifier is language-agnostic, you can use Review Center for documents written in any language.

1.4 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.

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.

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.

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 Active Learning project or other 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 the extracted text of documents to make its predictions. Even if other fields are returned in the saved search, it will not affect the results.

If you choose a prioritized review queue, you must code 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.

Note: The more documents are coded, the more accurate the classifier's predictions will be. We recommend starting with at least 10 positive-coded and 10 negative-coded documents when possible.

2.3 How document assignment works

By default, two 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 and 2 are assigned to the first reviewer who starts review. If a second reviewer logs in immediately after, documents 3 and 4 are assigned to the second reviewer. As the first reviewer completes their work, documents 5 and 6 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 below.

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.

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 40.

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: You cannot use a reflected field 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. The **Is Template** field should always be toggled to **On** for templates.

The Review Center application comes with several premade queue templates to choose from. However, we recommend tailoring them or creating your own to best suit your needs. These 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. Enter the following information:
    1. **Name**—the queue name reviewers will see.
    2. **Is Template**—toggle this to **On**.

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**—enter notes about the template such as its intended use, comments about field settings, etc.
4. **Queue Label**—create and choose organizational labels that will apply 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 15.
5. **Reviewer Groups**—this is not recommended for templates.
6. **Queue Type**—choose either **Saved Search** or **Prioritized Review**.
7. **Data Source**—select the saved search that contains the documents for your queue.

Note: If you are using a prioritized review queue:

- We recommend a maximum of 2.5 million documents in the data source.
- The classifier ignores documents with an extracted text field larger than 600 KB. We recommend leaving these documents out of the data source.

8. **Rank Output** (Prioritized Review only)—select the decimal field you created to hold the document rank scores.
9. **Review Field**—select the single choice field you created for review. This field must have two or more choices.
    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.
10. **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.

11. **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.

12. **Queue Display Options**—select which statistics you want reviewers to see on the queue card in the Review Queues tab.
13. **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 25.
    - If there are any conditions applied to the view, those conditions will also limit which documents appear in the panel.
14. **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 16.
    - 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.
15. **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.

4. Click **Save**.

   The template now appears in the Review Library list.

2.7 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.

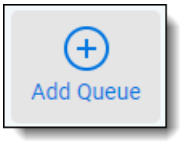

To create a new queue from template using the dashboard:

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. Under **Reviewer Groups**, choose one or more reviewer groups.
5. In the other fields, check the default values filled in by the template. Change any values that should be different for this queue.
6. 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.

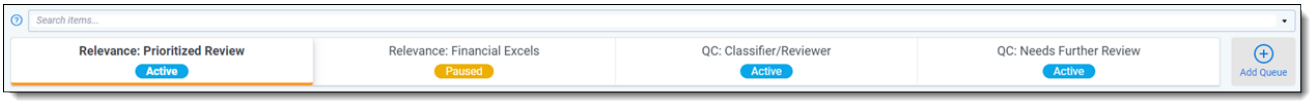

Below the queue name, each queue shows its status. The possible statuses are:

- **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 28.

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 12.

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.

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.

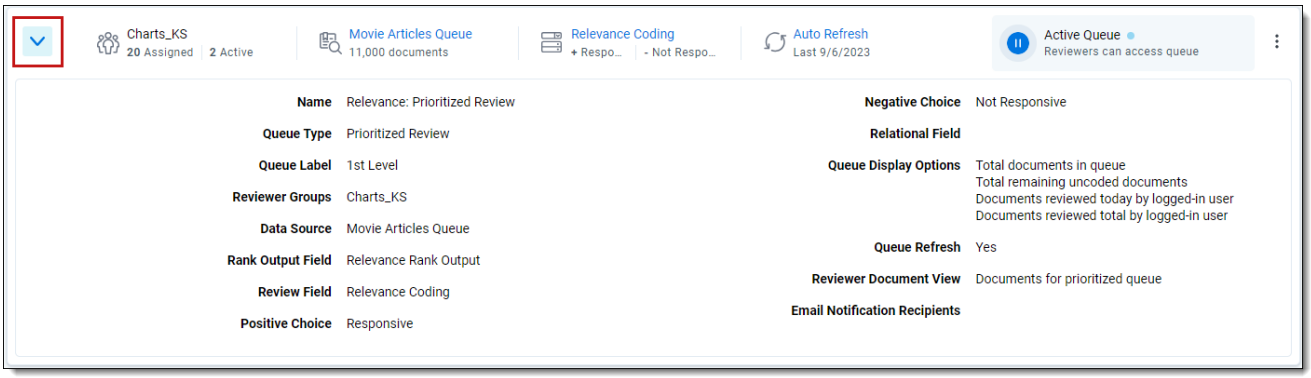

3.1.2.1 Preparing or refreshing the queue

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 below.

After you click the button, a **Cancel** option appears.

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 when a certain percentage of documents have had coding changes. These refreshes only happen after the queue has been started, and you can change this setting at any time.

The amount of document coding that triggers a refresh increases every time. After each refresh, there will need to be slightly more documents coded than the amount that triggered the last refresh.

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

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

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

Reviewer access during refreshes

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

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

3.1.2.2 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.

A few things are required before starting a queue:

- You must have a reviewer group assigned.
- For a prioritized review queue, you must code 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 you code in advance, the better the starting predictions will be.

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.

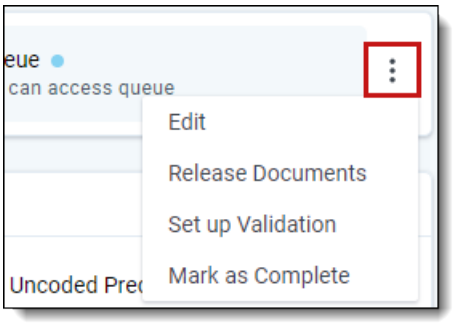

The menu options are:

- **Edit**—opens a modal to edit any of the queue settings. For field descriptions, see Creating a Review Center queue on page 8.
- **Release 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.

  Note: 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.

- **Set up Validation**—opens the options to create a review validation queue. For more information, see Review validation on page 28.
- **Mark as Complete**—sets the queue's status to Complete and moves it to the far right of the queue tab strip. This also removes the queue from the Review Queues tab, and reviewers can no longer access it. After the queue has been marked Complete, this option changes to **Re-enable**. Clicking this sets the queue's status to Not Started and returns it to the Review Queues tab.

3.1.3 Review Progress section

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

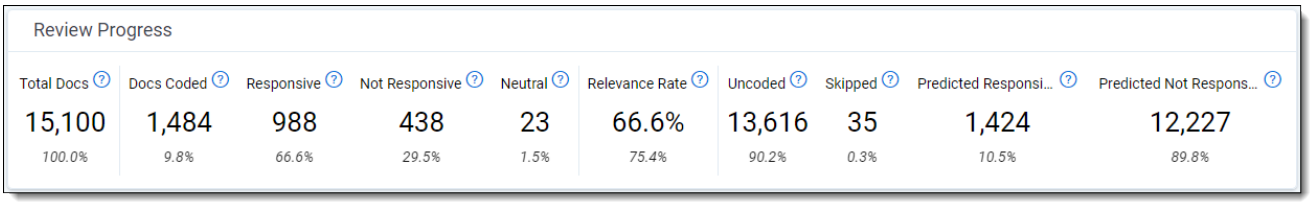

The statistics are:

- **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 either have a value in the review field, or have been skipped. 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 50.00. 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 50.00. 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.

3.2 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.

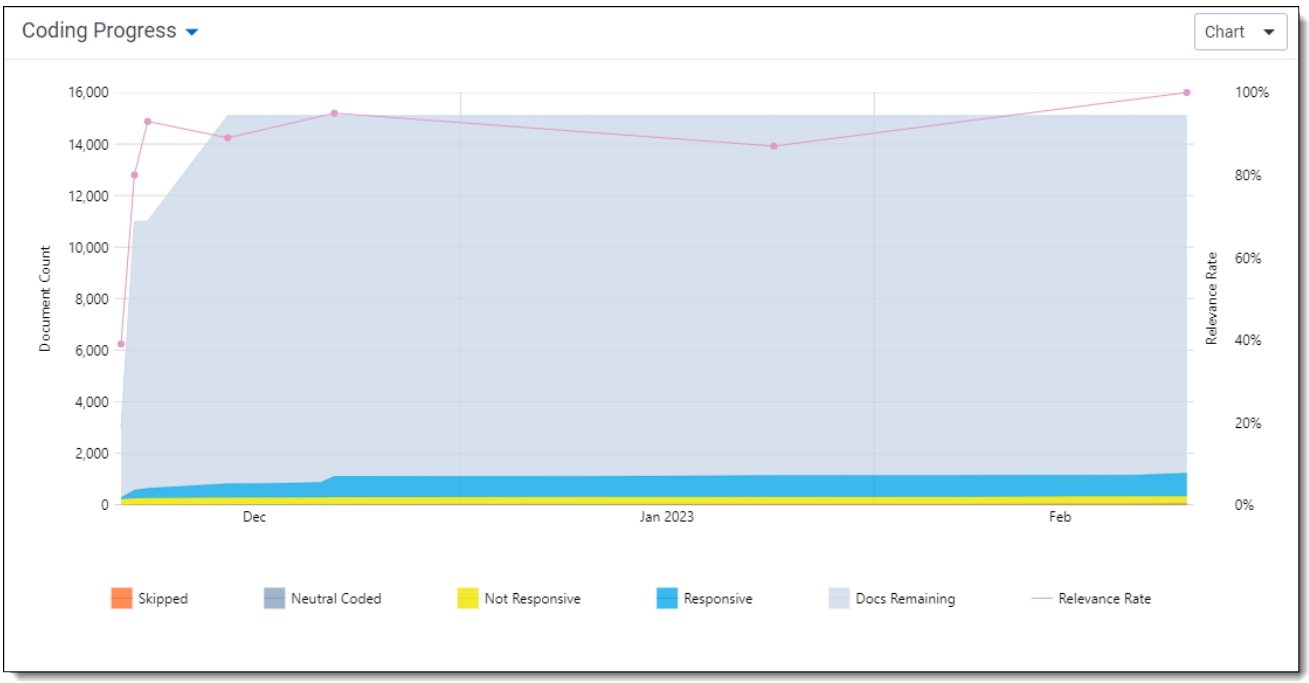

To navigate the visualization panel:

- 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. The table can also be downloaded as a .csv file.
- 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.

Note: 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. Data is reported in 15-minute increments.

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.

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.

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.

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)

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.

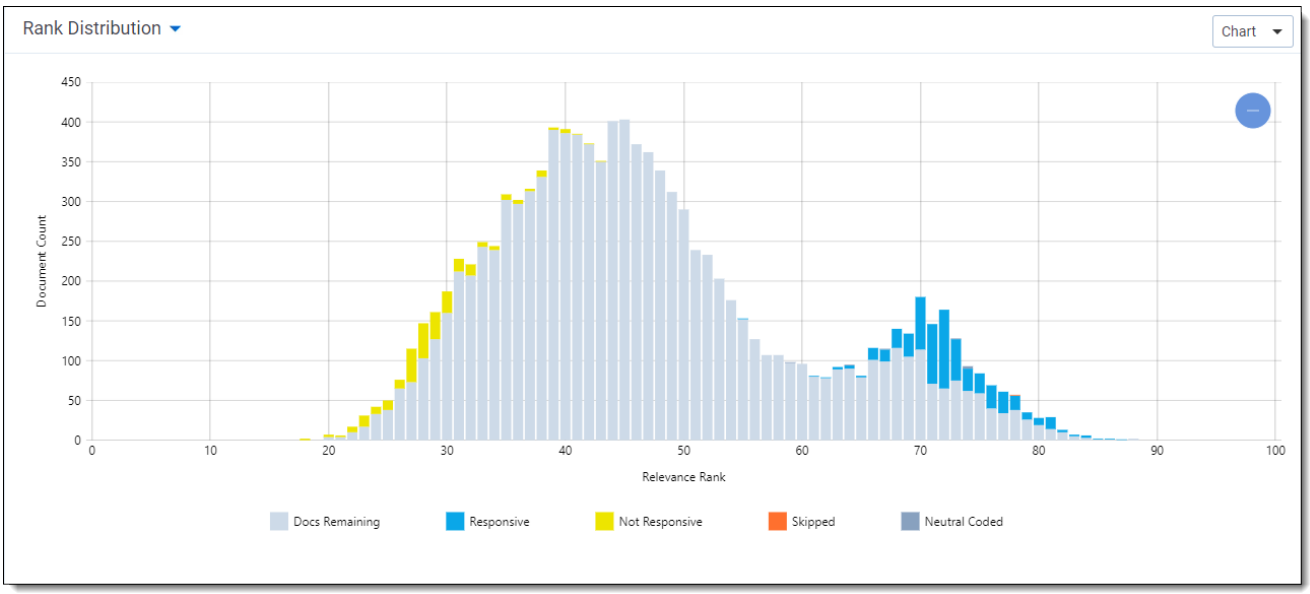

If you zoom out on the Rank Distribution chart, you may see documents with ranks below zero. These are documents that could not be classified. For more information, see Understanding document ranks on the next page.

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 document being checked out to the reviewer and checked back in. This 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 28.

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.

Notes:

- Active Learning and Review Center use similar ranking systems, but the classifiers are not the same. If you use both tools to classify the same document, it will receive separate scores. These scores can be very different depending on circumstances.
- 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 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 is larger than the limit of 600 KB. 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 21.

Alternatively, you can also use the following methods.

3.6.1 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.
5. Click on the **Reviewed** tag to view coded documents or the **Skipped** tag to view skipped documents.

If you rename or delete a queue, this renames or deletes the matching Field Tree tags also.

3.6.2 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. For most use cases, set **Track Initial Change Only?** to **Yes**. This sets it to track the first reviewer of the document, instead of overwriting 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 and Relativity Desktop Client. 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 and Relativity Desktop Client. This process is safe for your environment.

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.

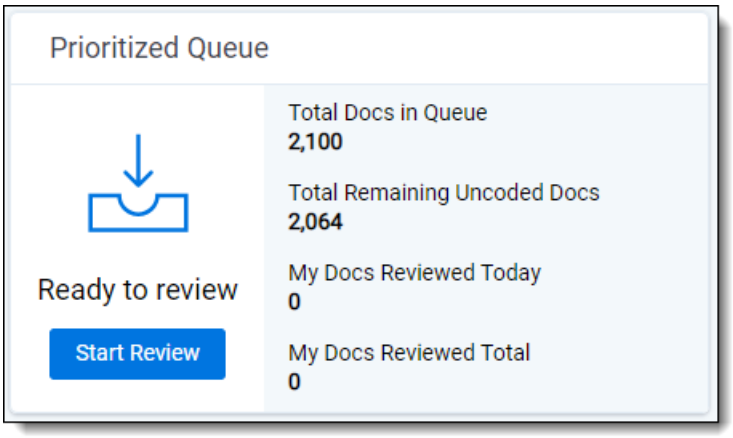

3. Click **Start Review**.

   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.

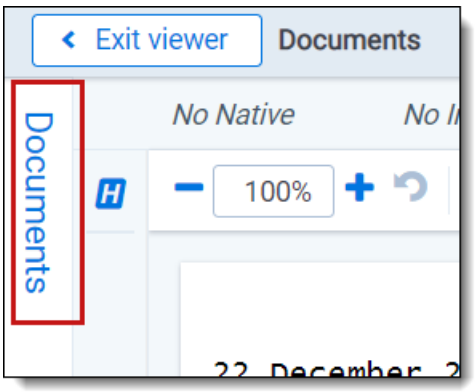

To view a document, click on its control number in the panel.

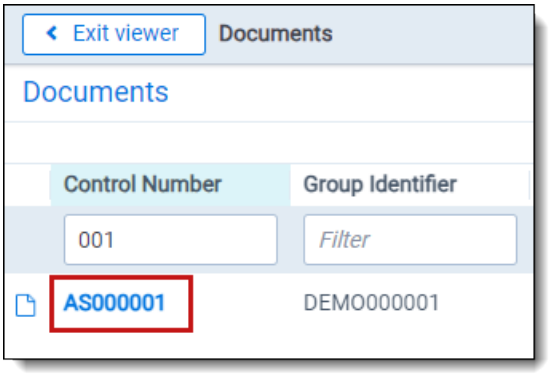

To return to your current document, click on the **Navigate Last** button in the upper right corner of the document viewer.

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 14.

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 detailed explanation of how the validation statistics are calculated, see Review validation statistics on page 36.

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 18.

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.
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. 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.

    Note: This is equivalent to choosing the “fixed” option when configuring an Elusion with Recall test in Active Learning. In contrast with an Active Learning Elusion with Recall test, Review Center only samples the discard pile. This means that the sample size is also the number of documents that will need to be coded.
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.

    Note: This is equivalent to the “statistical” option when configuring an Elusion with Recall test in Active Learning.

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 36.

5.3.2 Inherited settings

Each validation queue inherits these settings from the main queue:

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

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

5.4 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 also pause validation by clicking the **Pause** button. The validation statistics are reported in the Validation Progress section, and the charts and tables reflect the validation queue.

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

Note: 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 39.

5.4.1 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**
- **Email Notification Recipients**

3. Click **Save**.

For descriptions of the queue settings, see Creating a Review Center queue on page 8.

5.4.2 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 17.

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 21.

5.4.3 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.

   Several subfolders appear.
4. 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 two sub-choices: one for documents coded positive or negative, and one for skipped or neutral documents. As documents are coded, they populate under these choices.

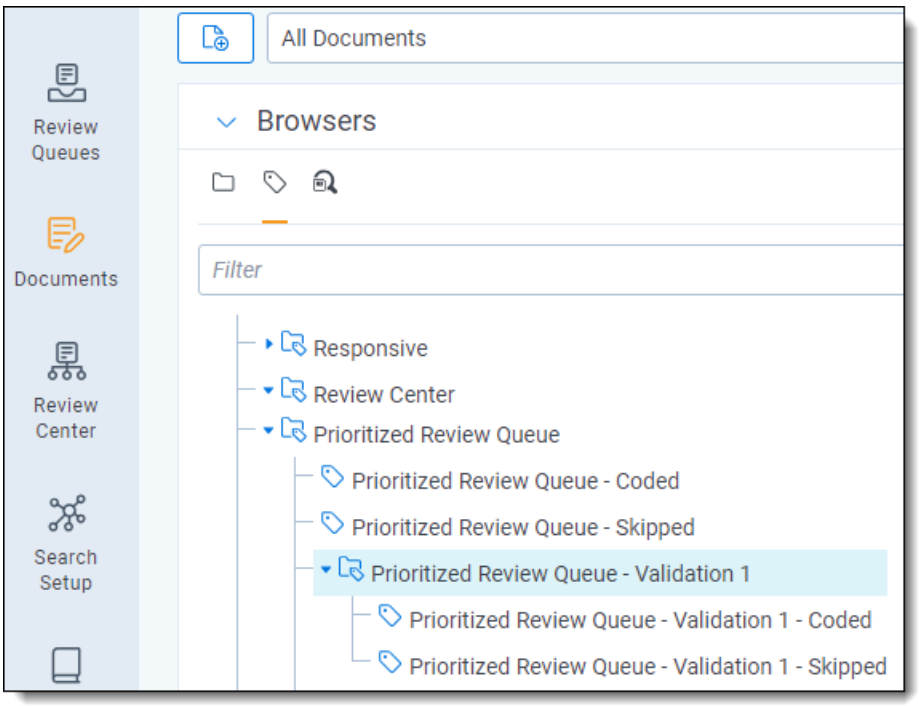

5.5 Accepting or rejecting validation results

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 34.

5.5.1 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.

5.6 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 not-yet-coded documents that are estimated to be 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.

  Note: 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.
- **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 36.

5.6.1 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**.

5.6.2 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.

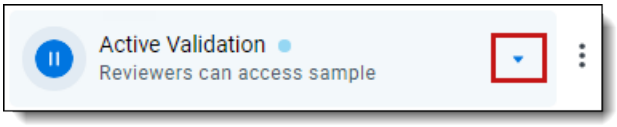

   A drop-down menu listing all linked queues appears.
2. Select the queue whose stats you want to view.

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.

5.7 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.

5.7.1 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.

5.7.2 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 three metrics for evaluating your review coverage: elusion, richness, and recall. 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.

For instructions on how to run Project Validation, see Review validation on page 28.

6.1 Defining elusion, richness, and recall

Validation centers on the following statistics, and for all of these, it reports on confidence intervals:

- **Elusion rate**—the percentage of documents that are relevant but have yet to be coded. The rate is rounded to the nearest single digit (tenth of a percent).
- **Richness**—the percentage of relevant documents across the entire review.
- **Recall**—the percentage of truly positive documents which were found by the review.

For both elusion and recall, a document is considered *found* if it was previously coded positive. Documents coded relevant after the start of the validation are not considered found for the purposes of validation. Documents coded outside of the queue still count toward elusion and against recall, and so do any responsive documents in the validation sample.

In pattern recognition, information retrieval, and machine learning, recall is a performance metric that applies to data retrieved from a collection.

Recall is based on relevance. Specifically, it is the fraction of relevant instances that were retrieved, as displayed in the following equation.

$$\frac{relevant\_retrieved\_instances}{all\_\mathbf{relevant}\_instances}$$

For each of these metrics, the validation queue assumes that you will 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 Validation metric calculations

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

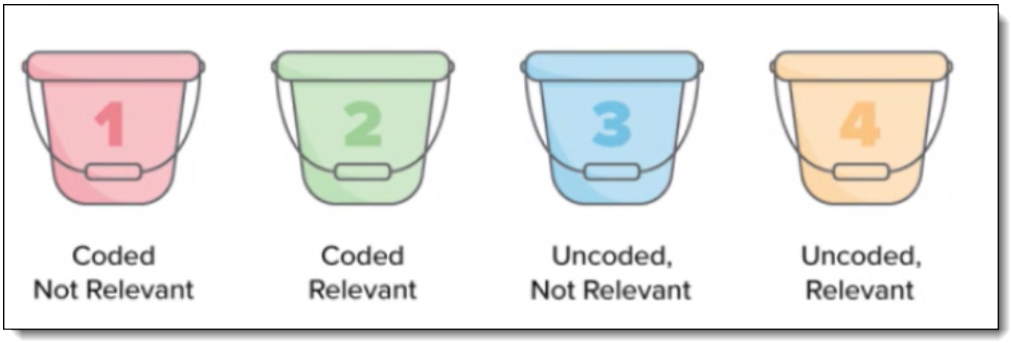

At the start of validation, the system knows exactly how many documents are in buckets 1 and 2.

- **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.

6.2.1 Elusion rate

This is the percentage of uncoded documents that are relevant.

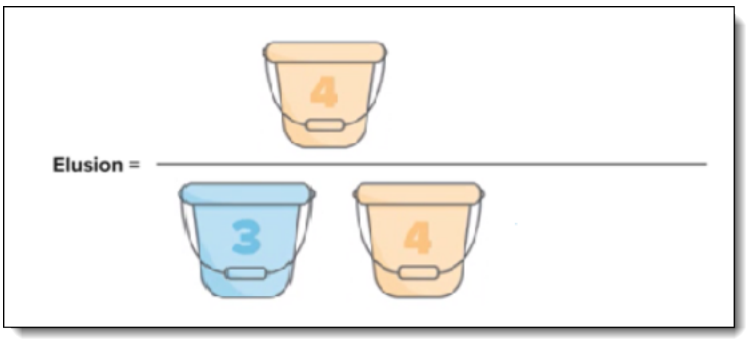

Review validation estimates this directly from the sample. It's calculated as the number of documents in the validation sample that were coded relevant, divided by the size of the sample.

6.2.2 Recall

Recall is the number of documents that were coded relevant, divided by the total number of relevant documents, both coded and uncoded.

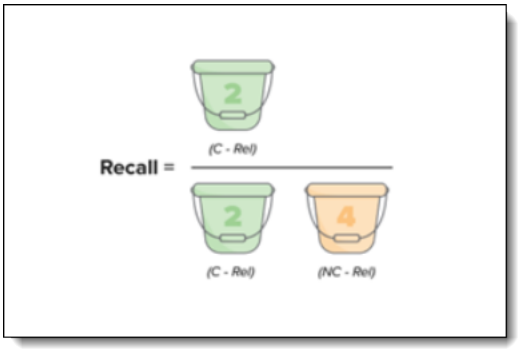

To estimate this, review validation needs an estimate of how many documents are in bucket 4. It finds this by multiplying the elusion rate by the total number of uncoded documents.

6.2.3 Richness

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

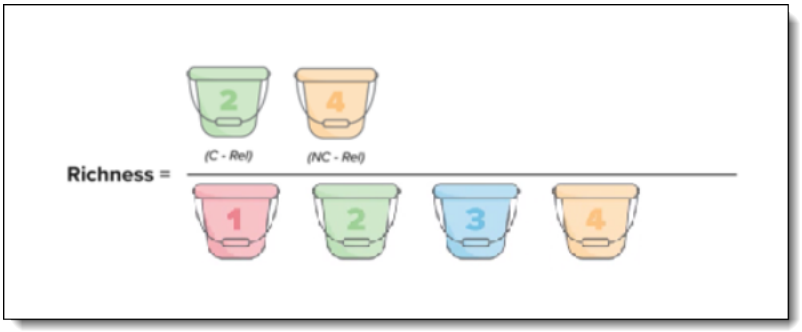

Similar to recall, review validation estimates the number of documents in bucket 4 by multiplying the estimated elusion rate by the number of uncoded documents. This is only done for the top half of the formula. For the bottom half, review validation only needs to know the size of the review. The buckets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

6.3 How the validation queue works

Once 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 once 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.

6.4 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.

For statistics, review validation assumes that skipped and neutral documents in the validation sample are responsive. Doing this guards against systematically overestimating recall.

In practice, this tends to raise the elusion and richness estimates, and lower the recall estimate.

7 Review Center security permissions

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

7.1 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

7.2 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.

7.3 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, Delete	■ Review Library	■ Delete

7.4 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.

7.5 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

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

- Setting up the reviewer group on page 9
- Creating a new queue from a template on page 12

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.