

aiR for Review Guide

May 8, 2024

Table of Contents

1 aiR for Review	4
1.1 aiR for Review overview	4
1.2 aiR for Review workflow	4
1.3 How aiR for Review works	5
1.3.1 Using aiR for Review with non-English languages	5
1.3.2 Using aiR for Review with emojis	6
1.4 Archiving and restoring workspaces with aiR for Review	6
2 Running aiR for Review	7
2.1 Installing aiR for Review	7
2.2 Setting up permissions	7
2.3 Choosing an analysis type	8
2.4 Choosing a processing mode	8
2.5 Best practices for running aiR for Review	9
2.6 Running the analysis	9
2.6.1 Step 1: Selecting documents and setting up the review	9
2.6.2 Step 2: Writing the Prompt Criteria	10
2.6.3 Step 3: Submitting the job for analysis	13
2.7 Editing and collaboration	14
2.8 How document units are calculated	14
2.9 Job capacity and size limitations	15
2.9.1 Size limits	15
2.9.2 Volume limits	15
2.9.3 Speed	15
3 Monitoring aiR for Review jobs	16
3.1 aiR for Review Jobs tab	16
3.1.1 Managing jobs and document linking	16
3.1.2 Viewing job details	17
3.1.3 Jobs tab fields	17
3.2 Fast Track banner statuses	18
4 aiR for Review results	19

4.1 How aiR for Review results work	19
4.1.1 Predictions versus document coding	19
4.1.2 Variability of results	19
4.2 Understanding document ranks	19
4.3 Viewing results for individual documents	20
4.3.1 Citations and highlighting	20
4.3.2 Adding aiR for Review fields to layouts	21
4.4 Viewing results for groups of documents	21
4.4.1 Creating an aiR for Review results view	21
4.4.2 Filtering and sorting aiR for Review results	21
4.5 aiR for Review results fields	22
4.5.1 aiR Relevance Analysis fields	22
4.5.2 aiR Issues Analysis fields	23
4.5.3 aiR Key Analysis fields	24
4.6 Using aiR for Review with Review Center	25
4.7 How document errors are handled	25
4.8 Hallucinations and confluations	26
5 aiR for Review security permissions	27
5.1 Running the aiR for Review mass action	27
5.2 Viewing the aiR for Review Jobs tab	27
5.2.1 Instance-level permissions	27
5.2.2 Workspace-level permissions	28
5.3 Clearing and restoring job results	28
5.4 Viewing highlights in the Viewer	28

1 aiR for Review

aiR for Review harnesses the power of large language models, or LLMs, to review documents. aiR for Review goes far beyond existing classifiers by using generative AI to both predict coding decisions and to support those predictions with descriptive text and document excerpts which explain the decisions.

Some benefits of aiR for Review include:

- Highly efficient, low-cost document analysis
- Quick discovery of important issues and criteria
- Consistent, cohesive analysis across all documents

Note: aiR for Review is currently in limited release. For information about the general release, contact your account representative.

1.1 aiR for Review overview

aiR for Review uses generative AI to simulate the actions of a human reviewer, finding and describing relevant documents according to the review instructions that you provide. It identifies the documents, describes why they are relevant using natural language, and demonstrates relevance using citations from the document.

aiR for Review has three different analysis types:

- **Relevance review**—predict documents responsive to a request for production.
- **Issues review**—locate material relating to different legal issues.
- **Key documents**—find hot documents important to a case or investigation.

Some use cases for aiR for Review include:

- **Kickstarting the review process**—prioritize the most important documents to give to reviewers.
- **First-pass review**—determine what you need to produce and discover essential insights.
- **Gaining early case insights**—learn more about your matter right from the start.
- **Internal investigations**—find documents and insights that help you understand the story hidden in your data.
- **Analyzing productions from other parties**—reduce the effort to find important material and get it into the hands of decision makers.
- **Quality control for traditional review**—compare aiR for Review's coding predictions to decisions made by reviewers to accelerate QC and improve results.

1.2 aiR for Review workflow

aiR for Review's process is similar to training a human reviewer: explain the case and its relevance criteria, hand over the documents, and check the results. If aiR misunderstood any part of the relevance criteria, explain that part in more detail, then try again.

Within Relativity, the main steps are:

1. Select the documents to review
2. Choose the aiR for Review mass action
3. Write and submit the review instructions, called Prompt Criteria
4. Review the results

When setting up the first analysis, we recommend running it on a sample set of documents that was already coded by human reviewers. If aiR's predictions are different from the human coding, revise the Prompt Criteria and try again. This could include rewriting unclear instructions, defining an acronym or a code word, or adding more detail to an issue definition.

Overall, the workflow has three phases:

1. **Develop**—write the Prompt Criteria, test, and tweak until the results match human review.
2. **Verify**—run the Prompt Criteria on a slightly larger set of documents and compare to results from senior reviewers.
3. **Run**—use the verified Prompt Criteria on a much larger set of documents.

For more details, see [Running aiR for Review on page 7](#). For additional workflow help and examples, see [Workflows for Applying aiR for Review](#) on the Community site.

1.3 How aiR for Review works

aiR for Review's analysis is powered by Azure OpenAI's GPT-4 large language model, or LLM. The LLM is designed to understand and generate human language, and it is trained on billions of documents from open datasets and the web.

When you submit Prompt Criteria and a set of documents to aiR for Review, Relativity sends the first document to Azure OpenAI and asks it to review the document according to the Prompt Criteria. After Azure OpenAI returns its results, Relativity sends the next document. The LLM reviews each document independently, and it does not learn from previous documents. Unlike Review Center, which makes its predictions based on learning from the document set, the LLM makes its predictions based on the Prompt Criteria and its preexisting training.

Azure OpenAI does not retain any data from the documents being analyzed. Data you submit for processing by Azure OpenAI is not retained beyond your organization's instance, nor is it used to train any other generative AI models from Relativity, Microsoft, or any other third party. For more information, see the white paper [A Focus on Security and Privacy in Relativity's Approach to Generative AI](#) on the Relativity website.

For more information on using generative AI for document review, we recommend:

- [Relativity Webinar - AI Advantage: How to Accelerate Review with Generative AI](#)
- [MIT's Generative AI for Law resources](#)
- [The State Bar of California's drafted recommendations for the use of generative AI](#)

1.3.1 Using aiR for Review with non-English languages

aiR for Review currently supports English-language Prompt Criteria and analyzing English text. It has not been tested for use with other languages, and doing so may result in unexpected or unverifiable results. However, Relativity will continue to do additional testing and analysis on non-English language scenarios.

1.3.2 Using aiR for Review with emojis

aiR for Review has not been specifically tested for analyzing emojis. However, the underlying LLM does understand Unicode emojis. It also understands other formats that could normally be understood by a human reviewer. For example, an emoji that is extracted to text as `:smile:` would be understood as smiling.

1.4 Archiving and restoring workspaces with aiR for Review

Workspaces with aiR for Review 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 was not checked during the archive process, the aiR for Review features in the restored workspace will not be fully functional. If this happens, you will need to manually reinstall aiR for Review in the restored workspace.

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

2 Running aiR for Review

aiR for Review uses generative AI to simulate the actions of a human reviewer, finding and describing relevant documents using the review instructions that you provide. These analyses can be customized to search for relevance, key documents, or specific issues as needed.

The instructions you give aiR for Review are called Prompt Criteria. For best results, we recommend analyzing a small set of documents, tweaking the Prompt Criteria as needed, then finally analyzing a larger set of documents. This lets you see immediately how aiR's coding compares to a human reviewer's coding and adjust the prompts accordingly.

Note: aiR for Review is currently in limited release. For information about the general release, contact your account representative.

2.1 Installing aiR for Review

aiR for Review is available as a secured application from the Application Library. You must have an active aiR for Review contract to use it, and it is not available for repository workspaces.

To install it:

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

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

- **aiR Relevance Analysis**—records the Relevance results of aiR for Review analysis runs.
- **aiR Issue Analysis**—records the Issue results of aiR for Review analysis runs.
- **aiR Key Analysis**—records the Key results of aiR for Review analysis runs.
- **aiR for Review Prompt Criteria**—records the Prompt Criteria settings and contents for each analysis run. This also records Prompt Criteria drafts for each user.

Installing aiR for Review creates two versions of the aiR for Review Jobs tab: one at the instance level, and one at the workspace level. Some users also find it helpful to manually create tabs for the three Analysis objects, but this is optional.

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

2.2 Setting up permissions

For detailed information on aiR for Review user permissions, see [aiR for Review security permissions on page 27](#).

2.3 Choosing an analysis type

aiR for Review supports three types of analysis. Each one is geared towards a different phase of a review or investigation.

For each aiR for Review job, choose one analysis type:

- **Relevance**—analyzes whether documents are relevant to a case or situation that you describe, such as documents responsive to a production request.
- **Relevance and Key Documents**—analyzes documents for both relevance and whether they are “hot” or key to a case.
- **Issues**—analyzes documents for whether they include content that falls under specific categories. For example, you might use this to check whether documents involve coercion, retaliation, or a combination of both.

Based on the analysis type you choose, you will need the following fields:

- **Relevance**—one single-choice results field. The field must have at least one choice.
- **Relevance and Key Documents**—two single-choice results fields. These should have distinct names such as “Relevant” and “Key,” and each field should have at least one choice.
- **Issues**—one multi-choice results field. Each choice should represent one of the issues you want to analyze.

Note: Currently, aiR for Review analyzes a maximum of five issues per run. You can have as many choices for the field as you want, but you can only analyze five at a time.

aiR for Review does not actually write to these fields. Instead, it uses them as reference when making and reporting on its predictions.

2.4 Choosing a processing mode

aiR for Review supports two processing modes: Fast Track and Batch. These are designed for the workflow of fine-tuning Prompt Criteria on a small set of documents, then using the Prompt Criteria on a large set of documents.

Because the large language model (LLM) has limited capacity, we dedicate some bandwidth exclusively to Fast Track jobs. This returns speedy results for smaller jobs without them having to wait for larger jobs in the queue.

For each aiR for Review job, choose one mode:

- **Fast Track**—processes up to 50 documents quickly. Use this to test and refine your Prompt Criteria on a small set of test documents.
 - Each user can have only one Fast Track job running at a time.
 - These jobs use dedicated bandwidth that is not available to larger batch jobs. They typically return results within 5-10 minutes.
- **Batch**—processes up to 50,000 documents. Use this to run your previously refined Prompt Criteria on a larger set of documents.

- Each instance can have up to three Batch jobs running at a time.
- Processing time varies based on document load and total load on the LLM.

For more detailed information about each mode's capacity, see [Job capacity and size limitations on page 15](#).

2.5 Best practices for running aiR for Review

aiR for Review works best after fine-tuning the Prompt Criteria. Analyzing just a few documents at first, comparing the results to human coding, and then adjusting the Prompt Criteria as needed yields more accurate results than diving in with a full document set.

We recommend the following workflow:

1. For your first analysis, run the Prompt Criteria on a set of 10 test documents that are a mix of relevant and not relevant.
2. Compare the results to human coding. In particular, look for documents that aiR coded differently than the humans did and investigate possible reasons. This could include unclear instructions, needing to define an acronym or code word, or other blind spots in the Prompt Criteria.
3. Tweak the Prompt Criteria to adjust for blind spots.
4. Repeat steps 1 through 3 until aiR predicts coding decisions accurately for the test documents.
5. Test the Prompt Criteria on 50 documents and compare results. Continue tweaking as needed.
6. Finally, run the Prompt Criteria on a larger set of documents.

aiR only sees the extracted text of a document. It does not see any non-text elements like advanced formatting, embedded images, or videos. We do not recommend using aiR for Review on documents such as images, videos, or spreadsheets with heavy formulas. Instead, use it on documents whose extracted text accurately represents their content and meaning.

2.6 Running the analysis

aiR for Review works as a mass action found on the Documents tab. Running the analysis has three basic parts:

1. Selecting documents and setting up the review
2. Writing the Prompt Criteria
3. Submitting the job for analysis

At any point in this process, you can click **Save and Close** in the mass action modal. This saves your progress so that you can keep working on it at a later time.

When you reopen the mass action modal, the last Prompt Criteria that you saved will display. For more information, see [Running aiR for Review on page 7](#).

2.6.1 Step 1: Selecting documents and setting up the review

To start an aiR for Review analysis job:

1. From the **Documents** tab, select the documents you want to analyze.
2. Under Mass Actions, select **aiR for Review**. A modal with several tabs appears.
3. On the Setup tab of the modal, set the following:
 1. **Prompt Criteria Name**—give the Prompt Criteria a unique name. You can also click **Load Prompt Criteria** to select Prompt Criteria that you or another user previously wrote. For more information, see [Editing and collaboration on page 14](#).
 2. **Review Type**—select one of the following. For more information, see [Choosing an analysis type on page 8](#).
 - **Relevance**—analyzes whether documents are relevant to a case or situation that you describe, such as documents responsive to a production request.
 - **Relevance and Key Documents**—analyzes documents for both relevance and whether they are “hot” or key to a case.
 - **Issues**—analyzes documents for whether they include content that falls under specific categories.
 3. **Run in Fast Track**—toggle this **On** for 50 documents or fewer, and **Off** for more than 50 documents. For more information, see [Choosing a processing mode on page 8](#).

The first time you use the mass action, all the fields will be blank. When you click **Save and Next** or **Save and Close**, your progress saves and persists for the next analysis.

If any required fields on any of the tabs are empty or misconfigured, the Save and Next button will be unavailable. Click on the title of each tab to fill out its fields.

2.6.2 Step 2: Writing the Prompt Criteria

The Prompt Criteria are a set of inputs that give aiR the context it needs to understand the matter and evaluate each document. Writing the Prompt Criteria is a way of training your “reviewer,” similar to training a human reviewer.

Depending which type of analysis you chose, you will see a different set of tabs. All Prompt Criteria include the Case Summary tab.

2.6.2.1 General writing guidelines

For all of the setup tabs, we recommend:

- Write as if “less is more.” Instead of pasting in a long review protocol as-is, summarize where possible and include only key passages. The Prompt Criteria have an overall length limit of 10,000 characters.
- Phrase things in a positive way when possible. Avoid negatives (“not” statements) and double negatives.
- Do not include explanations of the law.
- Do not give the LLM commands, such as “you will review XX.” Instead, simply describe the case.
- Use whatever writing format makes the most sense to a human reader. For example, bullet points might be useful for the People and Aliases section, but paragraphs might make sense in another section.

- The LLM has essentially “read the whole Internet.” It understands widely used slang and abbreviations, but it does not necessarily know jargon or phrases that are specific to an organization.

When you start to write your first Prompt Criteria, the fields contain grayed-out helper text that shows examples of what to enter. Use this as a guideline for crafting your own entries.

Note:

For more guidance on prompt writing, see the following resources on the Community site:

- [aiR for Review Prompt Writing Best Practices](#)—downloadable PDF of writing guidelines
 - [aiR for Review example project](#)—detailed example of adapting a review protocol into Prompt Criteria
-

2.6.2.2 Filling out the Case Summary tab

The Case Summary gives the broad context surrounding a matter. It includes an overview of the matter, people and entities involved, and any jargon or terms that are needed to understand the document set.

Limit the Case Summary content to roughly 20 sentences overall, and 20 each of People and Aliases, Noteworthy Organizations, and Noteworthy Terms.

To fill out the Case Summary tab:

1. Within the setup modal, click on the **Case Summary** tab.
2. Fill out the following:
 1. **Matter Overview**—provide a concise overview of the case. Include the names of the plaintiff and defendant, the nature of the dispute, and other important case characteristics.
 2. **People and Aliases**—list the names and aliases of key custodians who authored or received the documents. Include their role and any other affiliations.
 3. **Noteworthy Organizations**—list the organizations and other relevant entities involved in the case. Highlight any key relationships or other notable characteristics.
 4. **Noteworthy Terms**—list and define any relevant words, phrases, acronyms, jargon, or slang that might be important to the analysis.
 5. **Additional Context**—list any additional information that does not fit the other fields. This section is typically left blank.

Depending on which Review Type you chose, the remaining tabs will be called Relevance, Key Documents, or Issues. Fill out those tabs according to the guide sections below.

2.6.2.3 Filling out the Relevance tab

If you chose either Relevance or Relevance and Key Documents as the Review Type, you will see the Relevance tab. This defines the fields and criteria used for determining if a document is relevant to the case.

To fill out the Relevance tab:

1. Within the setup modal, click on the **Relevance** tab.
2. Fill out the following:
 1. **Relevance Field**—select a single-choice field that represents whether a document is relevant or non-relevant.
 2. **Relevant Choice**—select the field choice you use to mark a document as relevant.

3. **Relevance Criteria**—summarize the criteria that determine whether a document is relevant. Include:
 - Keywords, phrases, legal concepts, parties, entities, and legal claims
 - Any criteria that would make a document non-relevant, such as relating to a project that is not under dispute
4. **Issues Field (Optional)**—select a single-choice or multi-choice field that represents the issues in the case.
 - **Choice Criteria**—select each of the field choices one by one. For each choice, write a summary in the text box listing the criteria that determine whether that issue applies to a document. For more information, see [Filling out the Issues tab below](#).

Note: aiR does not make Issue predictions during Relevance review, but you can use this field for reference when writing the Relevance Criteria. For example, you could tell aiR that any documents related to these issues are relevant.

For best results when writing the Relevance Criteria:

- Limit the Relevance Criteria to 5-10 sentences.
- Do not paste in the original request for production (RFP); those are often too long and complex to give good results. Instead, summarize it and include key excerpts.
- Group similar criteria together when you can. For example, if an RFP asks for “emails pertaining to X” and “documents pertaining to X,” write “emails or documents pertaining to X.”

2.6.2.4 Filling out the Key Documents tab

If you chose Relevance and Key as the Review Type, you will see the Key Documents tab. This defines the fields and criteria used for determining if a document is “hot” or key to the case.

To fill out the Key Documents tab:

1. Within the setup modal, click on the **Key Documents** tab.
2. Fill out the following:
 1. **Key Document Field**—select a single-choice field that represents whether a document is key to the case.
 2. **Key Document Choice**—select the field choice you use to mark a document as key.
 3. **Key Document Criteria**—summarize the criteria that determine whether a document is key. Include:
 - Keywords, phrases, legal concepts, parties, entities, and legal claims
 - Any criteria that would exclude a document from being key, such as falling outside a certain date range

For best results, limit the Key Document Criteria to 5-10 sentences.

2.6.2.5 Filling out the Issues tab

If you chose Issues as the Review Type, you will see the Issues tab. This defines the fields and criteria used for determining whether a document relates to a set of specific topics or issues.

To fill out the Issues tab:

1. Within the setup modal, click on the **Issues** tab.
2. Fill out the following:
 1. **Field**—select a multi-choice field that represents the issues in the case.
 2. **Choice Criteria**—select each of the field choices one by one. For each choice, write a summary in the text box listing the criteria that determine whether that issue applies to a document. Include:
 - Keywords, phrases, legal concepts, parties, entities, and legal claims
 - Any criteria that would exclude a document from being key, such as falling outside a certain date range

For best results when writing the Choice Criteria:

- Limit the criteria description for each choice to 5-10 sentences.
- Each of the choices must have its own criteria. If a choice has no criteria, either fill it in or remove the choice.

Removing issue choices

aiR analyzes a maximum of 5 choices. If the issue field has more than 5 choices:

1. Select the choice you want to remove.
2. Click the **Remove Choice** button on the right.
3. Repeat with any other unwanted choices.

2.6.3 Step 3: Submitting the job for analysis

After filling out the Setup, Case Summary, and other tabs, review the job and submit it for analysis.

To submit a job:

1. Click **Save and Next**.
2. Review the confirmation summary. This includes:
 - **Total Docs**—number of documents to be analyzed.
 - **Est. Total Doc Units**—number of document units counted for billing purposes. For more information, see [How document units are calculated on the next page](#).
 - **Est. Time to Start**—estimated wait time from when you submit the job, to when aiR can begin analyzing your job. Longer wait times appear when aiR already has other work queued up.
 - **Est. Run Time**—estimated time aiR will take to analyze and return the results of the documents selected. This does not include time waiting in the queue.
3. Click **Start Classification**.

A banner appears showing that the job was successfully submitted for analysis.

Note: If you try to run a job that is too large or when too many jobs are already running, an error will appear. You can still save and edit the Prompt Criteria, but you will not be able to start the job. For more information, see [Job capacity and size limitations on the next page](#).

For information on monitoring jobs in progress, see [Monitoring aiR for Review jobs on page 16](#).

2.7 Editing and collaboration

By default, when you select the mass action, the last Prompt Criteria you saved will display. This makes it easy to edit the Prompt Criteria without re-entering information.

If you want to edit a different set of Prompt Criteria or collaborate with another user, you can load previous Prompt Criteria into the mass action modal. From there, any edits you make will be saved as a new set of Prompt Criteria.

To load previous Prompt Criteria:

1. From the **Documents** tab, check the documents you want to analyze.
2. Under Mass Actions, select **aiR for Review**.
3. On the Setup tab of the modal, click **Load Prompt Criteria**. A pop-up opens with two tabs:
 - **Prompt Criteria**—Prompt Criteria for jobs that already ran in the workspace.
 - **Drafts**—each user's most recently saved Prompt Criteria. These may or may not have been run yet.
4. Select a row from either tab.

The right-hand panel shows a preview of the Prompt Criteria.

5. Click **Load**.

The Prompt Criteria Name, Analysis Type, and all criteria load into the aiR for Review modal.

If you run the previous Prompt Criteria without making any changes, the Prompt Criteria Name stays the same. If you edit the prompt before running it, a number such as (1) or (2) will be automatically added to the end of the Prompt Criteria Name. You can also manually enter your own name for the Prompt Criteria as you edit it.

If you and another user both edit the same Prompt Criteria at the same time, your edits are saved as separate drafts. To collaborate on the same draft, we recommend having the first user finish their edits, then pass the draft off to the second user.

Note: Only one Prompt Criteria draft is saved for each user. If you save a draft, then load in different Prompt Criteria, that draft will be overwritten. To save Prompt Criteria long-term, run them with one or more documents.

2.8 How document units are calculated

A document unit is a document with between 1 and 15,000 characters of text. If a document has more than 15,000 characters in it, it is counted as two or more document units. For example, a 1-character document counts as one document unit, but a 16,000-character document counts as two document units.

Note: Any Unicode character counts as one character, regardless of storage size.

Because the document unit estimation calculates white space slightly differently than actual billing, small discrepancies may appear for document sizes that are right on the border between two document units. To find the actual document units that are billed, see the Cost Explorer.

2.9 Job capacity and size limitations

Based on the limits of the underlying large language model (LLM), aiR has size limits for the documents and prompts you submit, as well as volume limits for the overall jobs.

2.9.1 Size limits

The documents and Prompt Criteria have the following size limits:

- The Prompt Criteria have an overall length limit of 10,000 characters.
- Each document's extracted text should be under 150KB if possible. aiR has a hard limit of 300KB extracted text per document, but due to how the LLM processes document size, it sometimes rejects documents that are smaller than 300KB. To avoid this, we recommend treating 150KB as the upper limit.
- Each document's extracted text, when combined with the Prompt Criteria, must be less than 32,000 “tokens” (roughly equivalent to words, symbols, or whitespace). This is usually not a problem for documents under 150KB.

2.9.2 Volume limits

The volume limits for aiR for Review jobs are as follows:

Volume Type	Limit	Notes
Max job size for Fast Track mode	50 documents	For over 50 documents, use Batch mode.
Max job size for Batch mode	50,000 documents	A single Batch job can include up to 50,000 documents.
Concurrent Fast Track jobs per user	1 job	Each user can only have one Fast Track job running at a time. However, multiple users in one instance can run Fast Track jobs at the same time.
Concurrent Batch jobs per instance	3 jobs	Only 3 Batch jobs can be queued or running at the same time within an instance.

2.9.3 Speed

After a job is submitted, aiR analyzes roughly 25-50 documents per minute. Fast Track jobs typically take under 10 minutes. Batch job speeds vary widely depending on the number of documents, the overall load on the LLM, and other factors.

3 Monitoring aiR for Review jobs

After an aiR for Review job has started, you can use the aiR for Review jobs tab to monitor its progress, view prompt details, or cancel it. You can also view completed jobs and choose which analysis results are connected to the documents.

For Fast Track jobs, banner statuses appear on the Documents tab to let you know when a job is complete. You can also monitor these from the aiR for Review jobs tab like any other job.

Note: aiR for Review is currently in limited release. For information about the general release, contact your account representative.

3.1 aiR for Review Jobs tab

There are two versions of the aiR for Review Jobs tab: one at the instance level, and one at the workspace level. The instance-level tab shows all jobs across all workspaces, and it includes several extra columns to identify the workspace, matter, and client connected to each job. Most users only need access to the workspace-level tab. However, because some of aiR's volume limits are instance-wide, the instance-level tab makes it easy to see exactly how much capacity is being used.

Both versions of the tab show aiR for Review jobs that have been submitted for analysis. You can use the tab to view prompt details, cancel queued or in-progress jobs, and manage the job results.




For information on managing tab permissions, see [aiR for Review security permissions on page 27](#).

Note: If the aiR for Review Jobs tab says that aiR for Review is not currently available, check with your administrator. Your organization might not have an active contract for aiR for Review.

3.1.1 Managing jobs and document linking

You can use the aiR for Review Jobs tab to cancel jobs, clear job results out of the document fields, and restore previous job results.

To manage jobs, use the following icons:

- **Cancel symbol** ()—cancels a queued or in-progress job. Any results that were already received from the large language model (LLM) will stay in the fields, and those results will still be billed.
- **Clear symbol** ()—clears job results from the documents in this run. This empties the aiR for Review fields and removes highlighting from the Viewer, but it does not permanently delete the results. The results can be restored and re-linked at any time.
- **Restore symbol** ()—re-links the results of the selected job to the documents in the run. This replaces the results of any other job with the same result type.

For example, if you realize your current Prompt Criteria gives you less helpful results than a previous Prompt Criteria did, you can restore the previous job's results. This immediately gives reviewers access to the old predictions without needing to re-run the old Prompt Criteria.

Notes:

- If you run a new job on documents that were part of a previous job, you may temporarily see both sets of results linked to those documents. The old results will be unlinked after the new job is complete.
 - To avoid seeing doubled results, clear the previous result set using the aiR for Review Jobs tab.
-

3.1.2 Viewing job details

To see an aiR for Review job's Prompt Criteria, click on its row. A detail panel opens showing the setup details, case summary, fields, and criteria for analysis.

You can control a user's access to the detail panel using both item-level and workspace-level permissions. For more information, see [aiR for Review security permissions on page 27](#).

3.1.3 Jobs tab fields

The following fields appear on the aiR for Review Jobs tab:

- **Job ID**—the unique ID assigned to a job.
- **Prompt Criteria Name**—the name of the Prompt Criteria used by the job.
 - If several jobs ran using the same Prompt Criteria, this name will be the same for those jobs.
 - If a user edited the Prompt Criteria before running the job but did not change the name, the Prompt Criteria name will have a version number such as (1) or (2) appended after it.
- **Job Status**—the current state of the job. The possible statuses are:
 - Not Started
 - Queued
 - In Progress
 - Completed
 - Cancelling
 - Errored
- **Job Type**—the job's processing mode. For more information, see [Choosing a processing mode on page 8](#).
- **Client Name** (workspace-level only)—the client associated with the job's workspace.
- **Matter Name** (workspace-level only)—the matter name associated with the job's workspace.
- **Matter Number** (workspace-level only)—the matter number associated with the job's workspace.
- **Workspace ID** (workspace-level only)—the ID of the job's workspace.
- **Workspace Name** (workspace-level only)—the name of the job's workspace.
- **Doc Count**—the number of documents submitted for analysis.
- **Docs Successful**—the number of documents that were successfully analyzed.

- **Docs Pending**—the number of documents that are waiting to be analyzed.
- **Docs Errored**—the number of documents that encountered an error during analysis.
- **Docs Skipped**—the number of documents that aiR did not return results for. This can happen for reasons such as cancelling a job, network errors, and partial or complete job failures.
- **User Name**—the user who submitted the job.
- **Submitted Time**—the time the user submitted the job.
- **Completed Time**—the time the job successfully completed. If the job failed or was cancelled early, this field is blank.
- **Terminated Time**—the time the job stopped running, regardless of whether it was cancelled, failed, or completed successfully.
- **Job Failure Reason**—if the job failed, the reason is listed here. If the job completed successfully, this field is blank.
- **Estimated Wait Time**—the initial estimate for how long the job will wait between when the user submits the job and when the job can start running.
- **Estimated Run Time**—the initial estimate for how long the job will take to run.
- **Document Units**—the number of documents counted for billing purposes. For more information, see [How document units are calculated on page 14](#).

3.2 Fast Track banner statuses

When you submit a Fast Track job, the confirmation banner on the Documents tab updates automatically to show you the status of the job. This makes it easier to keep track of a small job's status without needing to check the dedicated Jobs tab. This banner only appears for the user who submitted the job.

The banner updates when the job is queued, in progress, and complete.

4 aiR for Review results

When aiR for Review analyzes documents, it makes predictions about the relevance of documents to different topics or issues. If it predicts that a document is relevant or relates to an issue, it includes a written justification of that prediction, as well as a counterargument and in-text citations. You can view these predictions, citations, and justifications either from the Viewer, or from a custom document view.

Note: aiR for Review is currently in limited release. For information about the general release, contact your account representative.

4.1 How aiR for Review results work

When aiR for Review finishes its analysis, it returns a prediction about how each document should be categorized, as well as its reasons for that prediction. This analysis has several parts:

- **Rank**—a numerical rank that indicates how strongly relevant the document is or how well it matches the predicted issue.
- **Prediction**—the relevance, key, or issue label that aiR predicts should apply to the document.
- **Rationale**—an explanation of why aiR chose this rank and prediction.
- **Considerations**—a counterargument explaining why the rationale might possibly be wrong.
- **Citations**—excerpts from the document that support the prediction and rationale.

In general, citations are left empty for non-relevant documents and documents that don't match an issue. However, aiR occasionally provides a citation for low-ranking documents if it helps to clarify why it was marked non-relevant. For example, if aiR is searching for changes of venue, it might cite an email that ends with "Hang on, gotta run, more later" as worth noting, even though it understands that's not a true change of venue request.

4.1.1 Predictions versus document coding

Even though aiR refers to the relevance, key, and issue fields during its analysis, it does not actually write to these fields. All of aiR's results are stored in aiR-specific fields such as the Prediction field. We recommend using these aiR fields for reference and reserving the actual relevance, key, and issue fields for human coding.

For ideas on how to integrate aiR for Review results into a larger review workflow, see [Using aiR for Review with Review Center on page 25](#).

4.1.2 Variability of results

Because of how large language models work, results can vary slightly from run to run. aiR's results for an individual document can potentially change even when given the same set of inputs. However, this is relatively rare; from our testing, it happens about 4% of the time.

4.2 Understanding document ranks


aiR ranks documents from 0 to 4 according to how relevant they are or how well they match an issue. The higher the number, the more relevant the document is predicted to be. In addition, aiR assigns a rank of -1 to any errored documents. Because these were not properly analyzed, they cannot receive a normal rank.

The aiR for Review ranks are:

Rank	Description
-1	The document either encountered an error or could not be analyzed. For more information, see How document errors are handled on page 25 .
0	The document is “junk” data such as system files or sets of random characters.
1	The document is predicted not relevant. aiR did not find any evidence that it relates to the case or issue.
2	The document is predicted borderline relevant. aiR found some content that might relate to the case or issue. It usually has citations.
3	The document is predicted relevant to the issue. Citations show the relevant text.
4	The document is predicted very relevant to the issue. aiR found direct, strong evidence that the content relates to the case or issue. Citations show the relevant text.

4.3 Viewing results for individual documents

From the Viewer, you can see the aiR for Review results for each individual document. Predictions show up in the left-hand pane, and all citations are automatically highlighted.

To view a document's aiR for Review results, click on the **aiR for Review Analysis** icon () to expand the pane. The aiR for Review Analysis pane displays the following:

1. Analysis Name
2. Prediction
3. Rationale and Considerations
4. Citation

Notes:

- If you run a new job on documents that were part of a previous job, you may temporarily see both sets of results linked to those documents. The old results will be unlinked after the new job is complete.
 - To avoid seeing doubled results, clear the previous result set using the aiR for Review Jobs tab.
-

4.3.1 Citations and highlighting

To jump to a specific citation, click the citation card. You can also toggle highlighting on or off by clicking the toggle at the top of the aiR for Review Analysis pane.

The highlight colors depend on the type of citation:

- **Relevance citation**—orange.
- **Key Document citation**—purple.
- **Issue citation**—color set chosen in the Color Map application. For more information, see [Changing the color associated with a coding choice in the Analytics guide](#).

If the same passage is cited by two types of results, the highlight blends their colors.

4.3.2 Adding aiR for Review fields to layouts

Because of how aiR for Review results fields are structured, you cannot add them directly to layouts. If the highlighting is not enough, you can add an object list to the layout that shows all linked results. For more information, see [Adding and editing an object list in the Admin guide](#).

4.4 Viewing results for groups of documents

You can view and compare aiR for Review results for large groups of documents by adding their fields to document views and saved searches.

Each field name is formatted as `aiR <review type> Analysis::<fieldname>`. For example, the Prediction field for a Relevance analysis is called `aiR Relevance Analysis::Prediction`.

For a full field list, see [aiR for Review results fields on the next page](#).

Notes:

- If you run a new job on documents that were part of a previous job, you may temporarily see both sets of results linked to those documents. The old results will be unlinked after the new job is complete.
 - To avoid seeing doubled results, clear the previous result set using the aiR for Review Jobs tab.
-

4.4.1 Creating an aiR for Review results view

When creating a view for aiR for Review results, we recommend including these fields:

- Edit
- Control Number
- <Review Field>
- aiR <Review Type> Analysis::Rank
- aiR <Review Type> Analysis::Prediction

Because the Rationale, Citation, and Considerations fields have larger blocks of text, those tend to be less helpful for comparing many documents. However, you can also add those if desired.

For a full field list, see [aiR for Review results fields on the next page](#).

4.4.2 Filtering and sorting aiR for Review results

Documents have a one-to-many relationship with the aiR for Review's results fields. For example, a single document might be linked to five Issue results. This creates some limitations when sorting and filtering results:

- Filter one column at a time in the Document list. Combining filters may include more results than you expect.
- If you need to filter by more than one field at a time, we recommend using search conditions instead.
- You can add these fields to views and widgets, but you cannot sort the view or the widget by these fields.

4.5 aiR for Review results fields

The results of every aiR for Review analysis are stored as part of an analysis object. Each of the three result types has its own object type to match:

- aiR Relevance Analysis
- aiR Key Analysis
- aiR Issue Analysis

aiR also links the results to each of the documents that were analyzed. These linked fields, called reflected fields, update to link to the newest results every time the document is analyzed. However, aiR keeps a record of all previous job results, and you can link the documents to a different job at any time. For more information, see [Managing jobs and document linking on page 16](#).

The reflected fields are the most useful for reviewing analysis results. These are formatted as `aiR <review type> Analysis::<fieldname>`. For example, the Prediction field for a Relevance analysis is called `aiR Relevance Analysis::Prediction`.

4.5.1 aiR Relevance Analysis fields

The fields for aiR Relevance Analysis are:

Field name	Field type	Description
Name	Fixed-length Text	The name of this specific result. This formatted as <code><Document Artifact ID>_<Job ID></code> .
Job ID	Fixed-length Text	The unique ID of the job this result came from.
Rank	Whole Number	Numerical rank indicating how strongly relevant the document is. For more information, see Understanding document ranks on page 19 .
Document	Multiple Object	The Control Number of the document this result is linked to. If the result is not currently linked to any documents, this field is blank.
Prediction	Fixed-length Text	aiR's prediction of whether this qualifies as a relevant document.
Rationale	Fixed-length Text	An explanation of why aiR chose this rank and prediction.
Considerations	Fixed-length Text	A counterargument explaining why the rationale might possibly be wrong.
Citation 1	Fixed-length	Excerpt from the document that supports the prediction and rationale. This may be blank for some documents.

Field name	Field type	Description
	Text	
Citation 2	Fixed-length Text	Second excerpt from the document that supports the prediction and rationale. This may be blank for some documents.
Citation 3	Fixed-length Text	Third excerpt from the document that supports the prediction and rationale. This may be blank for some documents.
Citation 4	Fixed-length Text	Fourth excerpt from the document that supports the prediction and rationale. This may be blank for some documents.
Citation 5	Fixed-length Text	Fifth excerpt from the document that supports the prediction and rationale. This may be blank for some documents.
Error Details	Fixed-length Text	If the document encountered an error, the error message displays here. For an error list, see How document errors are handled on page 25 .

4.5.2 aiR Issues Analysis fields

The fields for aiR Issues Analysis are:

Field name	Field type	Description
Name	Fixed-length Text	The name of this specific result. This formatted as <Document ID>_<Job ID>.
Job ID	Fixed-length Text	The unique ID of the job this result came from.
Choice Analyzed	Fixed Text	The name of the issue choice being analyzed for this result.
Choice Analyzed ID	Whole Number	The Artifact ID of the issue choice being analyzed for this result.
Document	Multiple Object	The Control Number of the document this result is linked to. If the result is not currently linked to any documents, this field is blank.
Rank	Whole Number	Numerical rank indicating how well the document matches an issue. For more information, see Understanding document ranks on page 19 .
Prediction	Fixed-length Text	aiR's predicted issue choice for this document.
Rationale	Fixed-	An explanation of why aiR chose this rank and prediction.

Field name	Field type	Description
	length Text	
Considerations	Fixed-length Text	A counterargument explaining why the rationale might possibly be wrong.
Citation	Fixed-length Text	Excerpt from the document that supports the prediction and rationale. This may be blank for some documents.
Error Details	Fixed-length Text	If the document encountered an error, the error message displays here. For an error list, see How document errors are handled on the next page .

4.5.3 aiR Key Analysis fields

The fields for aiR Key Analysis are:

Field name	Field type	Description
Name	Fixed-length Text	The name of this specific result. This formatted as <Document ID>_<Job ID>.
Job ID	Fixed-length Text	The unique ID of the job this result came from.
Document	Multiple Object	The Control Number of the document this result is linked to. If the result is not currently linked to any documents, this field is blank.
Rank	Whole Number	Numerical rank indicating how strongly relevant the document is. For more information, see Understanding document ranks on page 19 .
Prediction	Fixed-length Text	aiR's prediction of whether this qualifies as a key document.
Rationale	Fixed-length Text	An explanation of why aiR chose this rank and prediction.
Considerations	Fixed-length Text	A counterargument explaining why the rationale might possibly be wrong.
Citation 1	Fixed-length Text	Excerpt from the document that supports the prediction and rationale. This may be blank for some documents.
Citation 2	Fixed-	Second excerpt from the document that supports the prediction and

Field name	Field type	Description
	length Text	rationale. This may be blank for some documents.
Citation 3	Fixed-length Text	Third excerpt from the document that supports the prediction and rationale. This may be blank for some documents.
Citation 4	Fixed-length Text	Fourth excerpt from the document that supports the prediction and rationale. This may be blank for some documents.
Citation 5	Fixed-length Text	Fifth excerpt from the document that supports the prediction and rationale. This may be blank for some documents.
Error Details	Fixed-length Text	If the document encountered an error, the error message displays here. For an error list, see How document errors are handled below .

4.6 Using aiR for Review with Review Center

One option for integrating aiR for Review into a larger review workflow is to combine it with Review Center. After analyzing the documents with aiR for Review, you can use aiR's predictions to prioritize which documents to include in a Review Center queue.

For example, you may want to review all documents that aiR for Review ranked as borderline or above for relevance. To do that:

1. Set up a saved search for documents where **aiR Relevance Analysis::Rank** is greater than **1**. This returns all documents ranked 2 or higher.
2. Create a Review Center queue using that saved search as the data source.

Because of how the aiR for Review fields are structured, you cannot sort by them. However, you can either sort by another field, or use a prioritized review queue to dynamically serve up documents that may be most relevant.

For more information, see the Review Center guide.

4.7 How document errors are handled

If aiR encounters a problem when analyzing a document, it will not return results for that document. Instead, it ranks the document as -1 and returns an error message in the Error Details column. Your organization is not charged for any errored documents.

The possible error messages are:

Error message	Description
Failed to parse completion	The large language model (LLM) encountered an error.

Error message	Description
Completion is not valid JSON	The large language model (LLM) encountered an error.
Hallucination detected in completion	The results for this document may include a hallucination. For more information, see Hallucinations and conflations below .
Conflation detected in completion	The results for this document may include a conflation. For more information, see Hallucinations and conflations below .
Document text is empty	The extracted text of the document was empty.
Document text is too short	There was not enough extracted text to analyze in the document.
Document text is too long	The document's extracted text was too long to analyze.
Model API error occurred	A communication error occurred between the large language model (LLM) and Relativity.
Uncategorized error occurred	An unknown error occurred.

4.8 Hallucinations and conflations

Two types of errors deserve special mention:

- **Hallucinations**—these occur when the aiR results citation cannot be found anywhere in the prompt text. The large language model (LLM) appears to be citing sentences that don't exist in the prompt.
- **Conflations**—these occur when the aiR results citation comes from something other than the document itself, but which is still part of the full prompt. For example, it might cite text that was part of the Prompt Criteria instead of the document's extracted text.

When aiR receives the analysis results from the LLM, it checks all citations against the prompt text. Any possible hallucinations or conflations are marked as errors, and they receive a rank of -1 instead of whatever rank they were originally assigned. We recommend manually reviewing errored documents.

Hallucinations are typically rare. However, highly structured documents such as Excel spreadsheets and PDF forms have a higher hallucination rate than other document types.

Note: Due to the way that columns of text are scanned in OCR, OCR'd documents are occasionally marked as hallucinations when the citations do actually exist in the original document.

5 aiR for Review security permissions

This page contains information on the security permissions required for interacting with aiR for Review.

Note: aiR for Review is currently in limited release. For information about the general release, contact your account representative.

5.1 Running the aiR for Review mass action

To run the aiR for Review mass action, you need the following permissions:

Object Security	Other Settings
<ul style="list-style-type: none">■ aiR for Review Profile - View, Edit, Add	<ul style="list-style-type: none">■ Mass Operations - aiR for Review

You must also belong to at least one user group other than the Workspace Admin Group.

5.2 Viewing the aiR for Review Jobs tab

There are two versions of the aiR for Review Jobs tab: one at the instance level, and one at the workspace level. The instance-level tab shows all jobs across all workspaces, and it includes several extra columns to identify the workspace, matter, and client connected to each job.

The following permissions allow users to see the job list and click on each job to view Prompt Criteria details. Users with access to this tab can also cancel in-progress jobs.

5.2.1 Instance-level permissions

To view the instance-level aiR for Review Jobs tab, you need the following permissions:

Object Security	Tab Visibility
<ul style="list-style-type: none">■ Admin Repository - View■ aiR for Review Prompt Criteria - View	<ul style="list-style-type: none">■ aiR for Review Jobs

Assign these permissions under the **Instance Details** tab.

5.2.1.1 Viewing Prompt Criteria at the instance level

To view Prompt Criteria details for a job, you also need some permissions within that job's workspace:

- You must belong to more than just the Workspace Admin Group within the workspace.
- You must have **aiR for Review Prompt Criteria - View** rights within that job's workspace.

If you do not have these, you will be able to see jobs from that workspace, but you will not be able click on those jobs to view their Prompt Criteria.

You can also use item-level permissions to restrict access to a specific job's aiR for Review Prompt Criteria. For more information, see [Security and permissions in the Admin guide](#).

5.2.2 Workspace-level permissions

To view the workspace-level aiR for Review Jobs tab, you need the following permissions:

Object Security	Tab Visibility
<ul style="list-style-type: none">■ aiR for Review Prompt Criteria - View	<ul style="list-style-type: none">■ aiR for Review Jobs

Assign these permissions under the **Workspace Details** tab within the chosen workspace.

You can also use item-level permissions to restrict access to a specific job's aiR for Review Prompt Criteria. For more information, see [Security and permissions in the Admin guide](#).

5.3 Clearing and restoring job results

To clear or restore job results using the aiR for Review Jobs tab, you need the following permissions:

Object Security	Tab Visibility
<ul style="list-style-type: none">■ Document - View, Edit■ aiR Relevance Analysis - View, Edit■ aiR Issue Analysis - View, Edit■ aiR Key Analysis - View, Edit	<ul style="list-style-type: none">■ aiR for Review Jobs

If you have Edit permissions for only one of the analysis types, you will only be able to clear or restore results of that type.

For more information on clearing and restoring results, see [Managing jobs and document linking on page 16](#).

5.4 Viewing highlights in the Viewer

To see aiR for Review results highlighted in the Viewer, you need the following permissions:

Object Security
<ul style="list-style-type: none">■ aiR Relevance Analysis - View■ aiR Issue Analysis - View■ aiR Key Analysis - View

If you only have access to some of these, you will only see highlighting for those analysis types.

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.

©2024. Relativity ODA LLC. All rights reserved. Relativity® is a registered trademark of Relativity ODA LLC.