aiR for Review Guide

July 8, 2025

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

Table of Contents

1 aiR for Review	5
1.1 aiR for Review workflow	5
1.2 How it works	6
1.3 Regional availability of aiR for Review	7
1.4 Language support	7
1.5 Analyzing emojis	. 8
1.6 Archiving and restoring workspaces	8
2 Installing aiR for Review	9
2.1 Object types available	9
2.2 Tabs available	10
3 Permissions	11
3.1 Viewing the aiR for Review dashboard	11
3.2 Creating and running an aiR for Review project	11
3.3 Editing and running an existing aiR for Review project	11
3.4 Viewing highlights in the Viewer	12
3.5 Viewing the aiR for Review Jobs tab	12
3.5.1 Instance-level permissions	12
3.5.2 Workspace-level permissions	13
3.6 Clearing and restoring job results	13
3.7 Running the aiR for Review mass action	13
4 Best practices	14
4.1 Tips for writing Prompt Criteria	14
4.2 Prompt criteria iteration sample documents	15
4.3 Prompt criteria iteration workflow	15
5 Creating an aiR for Review project	16
5.1 Choosing an analysis type	16
5.2 Choosing the data source	16
5.3 Overview of setting up an aiR for Review project	17
5.4 Step 1: Creating an aiR for Review project	17
5.5 Step 2: Writing the Prompt Criteria	18

5.5.1 Case Summary tab	
5.5.2 Relevance tab	
5.5.3 Key Documents tab	20
5.5.4 Issues tab	20
5.6 Step 3: Running the first analysis	21
5.7 Running aiR for Review as a mass operation	22
5.8 Editing and collaboration	
5.9 Job capacity and size limitations	
5.9.1 Size limits	23
5.9.2 Volume limits	23
5.9.3 Speed	24
5.10 Understanding documents and billing	24
6 Iterating on the prompt criteria	
6.1 Navigating the aiR for Review dashboard	25
6.1.1 Project details strip	25
6.1.2 Prompt Criteria panel	25
6.1.3 Aspect selector bar	
6.1.4 Project metrics section	27
6.1.5 Analysis Results section	29
6.1.6 Clearing selected documents	
6.1.7 Saving selected documents as list	
6.1.8 Exporting Prompt Criteria	
6.2 How Prompt Criteria versioning works	
6.2.1 How version controls affect the Viewer	
6.3 Revising the Prompt Criteria	
6.4 Increasing the job size	
7 Managing aiR for Review jobs	
7.1 aiR for Review Jobs tab	
7.2 How aiR for Review document linking works	
7.3 Managing jobs and document linking	35
7.4 Viewing job details	
7.5 Jobs tab fields	

8 aiR for Review results	
8.1 How aiR for Review analysis results work	
8.1.1 Predictions versus document coding	
8.1.2 Variability of results	
8.2 Understanding document scores	
8.3 Viewing results from the dashboard	
8.4 Viewing results for individual documents	
8.4.1 Citations and highlighting	
8.4.2 Adding aiR for Review fields to layouts	
8.5 Creating document views and saved searches	
8.5.1 Creating an aiR for Review results view	
8.5.2 Filtering and sorting aiR for Review results	
8.6 aiR for Review results fields	
8.6.1 aiR Relevance Analysis fields	43
8.6.2 aiR Issues Analysis fields	
8.6.3 aiR Key Analysis fields	45
8.7 Using aiR for Review with Review Center	
8.8 How document errors are handled	
8.9 Ungrounded citations	
9 aiR for Review Analysis	
9.1 Permissions	
9.2 Reviewing aiR for Review Analysis results	
9.3 aiR for Review Analysis highlighting	
9.4 aiR for Review Analysis results order	

1 aiR for Review

aiR for Review harnesses the power of large language models (LLM) to review documents. It uses generative artificial intelligence (AI) to simulate the actions of a human reviewer by finding and describing relevant documents according to the review instructions (prompt criteria) that you provide. It identifies the documents, describes why they are relevant using natural language, and demonstrates relevance using citations from the document.

Three different analysis types are available:

- Relevance review—use to predict documents responsive to a request for production.
- Issues review—use to locate material relating to different legal issues.
- **Key documents**—use to find key documents important to a case or investigation, including those that might be critical or embarrassing to one party or another.

A few benefits of application include:

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

Below are some common use cases for it:

- 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.

1.1 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 the application misunderstood any part of the relevance criteria, simply explain that part in more detail, then try again.

Within Relativity, the main steps are:

- 1. Select the documents to review.
- 2. Create an aiR for Review project.
- 3. Write and submit the review instructions, called Prompt Criteria.
- 4. Review the results (citations, rationale, considerations, recommendation).

When setting up the first analysis, we recommend running it on a sample set of documents that was already coded by human reviewers. If the resulting 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 (review instructions), test on a small document set, and tweak until the results align with human review.
- 2. **Validate**—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 much larger sets of documents.

For more details, see <u>Creating an aiR for Review project on page 16</u>. For additional workflow help and examples, see <u>Workflows for Applying aiR for Review</u> on the Community site.

1.2 How it works

aiR for Review's analysis is powered by Azure OpenAI's GPT-4 Omni large language model. The LLM is designed to understand and generate human language. 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 built-in 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.

Note: For European Economic Area (EEA) customers, aiR for Review data may be processed elsewhere in the EU, but it will always be processed in compliance with applicable laws. For more information, please contact your account manager.

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

Prelativity one

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

1.3 Regional availability of aiR for Review

aiR for Review's availability may vary by region, as well as the availability of the LLM model used. Once OpenAI releases an LLM model to a region, Relativity tests it and notifies clients before upgrading aiR for Review.

The following table lists the current LLM model available and date it was deployed to aiR for Review per region. Also listed is the current version of aiR for Review, which may vary by region.

Region	Current LLM Model	aiR for Review Model Deployment Date	Current aiR for Review Version
United States	GPT-4 Omni - November	2024-06-16	2025.06.1
United Kingdom	GPT-4 Omni - November	2024-06-16	2025.06.1
Australia	GPT-4 Omni - November	2024-06-16	2025.06.1
Canada	GPT-4 Omni - November	2024-06-16	2025.06.1
France	GPT-4 Omni - November	2024-06-16	2025.06.1
Germany	GPT-4 Omni - November	2024-06-16	2025.06.1
Ireland	GPT-4 Omni - November	2024-06-16	2025.06.1
Netherlands	GPT-4 Omni - November	2024-06-16	2025.06.1
Switzerland	GPT-4 Omni - November	2024-06-16	2025.06.1

For more details about availability in your region, contact your account representative.

For technical specifications of your region's current LLM model, see documentation on the <u>Azure</u> website.

1.4 Language support

The underlying LLM used by aiR for Review has been evaluated for use with 83 languages. While aiR for Review itself has been primarily tested on English-language documents, unofficial testing with non-English datasets shows encouraging results.

If you use the application with non-English data sets, we recommend the following:

- Rigorously follow best practices for writing and iterating on the Prompt Criteria. For more information, see Step 2: Writing the Prompt Criteria on page 18 and Iterating on the prompt criteria on page 25.
- Analyze the extracted text as-is. You do not need to translate it into English.
- When possible, write the Prompt Criteria in the same language as the documents being analyzed. This should also be the subject matter expert's native language. If that is not possible, write the Prompt Criteria in English.

When you view the results of the analysis, all citations stay in the same language as the document they cite. By default, the rationales and considerations are in English. If you want the rationales and considerations to be in a different language, type "Write rationales and considerations in [desired language]" in the Additional Context field of the Prompt Criteria.

For the study used to evaluate Azure OpenAI's GPT-4 model across languages, see <u>MEGAVERSE</u>: <u>Benchmarking Large Language Models Across Languages</u>, <u>Modalities</u>, <u>Models and Tasks</u> on the arXiv website.

1.5 Analyzing 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.

1.6 Archiving and restoring workspaces

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

When archiving in ARM, check **Include Extended Workspace Data** under Extended Workspace Data Options. If this option is 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.

Note: If you restore a workspace that includes previous aiR for Review jobs, the pre-restoration jobs will not appear on the instance-level aiR for Review Jobs tab. The jobs and their results will still be visible at the workspace level.

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

2 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. This application is not available for repository workspaces.

To install it in the workspace:

- 1. Navigate to a workspace where you want to install the application.
- 2. Click the Workspace Admin tab and the Relativity Applications tab.
- 3. Click New Relativity Application to display an application form.
- 4. Click the Select from Application Library radio button in the Application Type section.

Note: Global applications are not listed in the **Select from Application Library** option when attempting to add an application to a workspace.

5. Click

ck **——** in the **Choose from Application Library** field.

- 6. Select the **aiR for Review** application on the Select Library Application dialog. This dialog displays only applications added to the Application Library.
- 7. Click Ok. The application form displays the following fields:
 - Schema Version—displays the version of the application that you are installing.
 - User-friendly URL—displays a user-friendly version of the application's URL. This field may be blank.
 - Application Artifacts—displays object types and other application components.
- 8. (Optional) Click Clear to remove the application from the form.
- 9. Click **Import** to save your mappings and import the application.
- 10. Review the import status of the application. Verify that the install was successful or resolve errors.

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

See <u>Permissions on page 11</u> for the list of permissions required for using aiR for Review.

2.1 Object types available

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.
- aiR for Review Project—records the details of each aiR for Review project.

2.2 Tabs available

The following tabs will appear:

- **aiR for Review Projects** (workspace level)—create and manage aiR for Review projects and view the project dashboard.
- **aiR for Review Jobs** (workspace level)—view and manage jobs created by the aiR for Review application within the workspace.
- **aiR for Review Jobs** (instance level)—view and manage jobs created by the aiR for Review application across all workspaces in the instance.

3 Permissions

This information below outlines the permissions required for using aiR for Review.

3.1 Viewing the aiR for Review dashboard

To view the aiR for Review dashboard, you need the following permissions:

Object Security	Tab Visibility
• aiR for Review Project - View	• aiR for Review
 aiR for Review Prompt Criteria View 	Projects
• aiR Relevance Analysis - View	
• aiR Issue Analysis - View	
• aiR Key Analysis - View	

You can view results only for the analysis types you have permission to view.

3.2 Creating and running an aiR for Review project

To creating an aiR for Review project and run the analysis, you need the following permissions:

Object Security	Tab Visibility
• aiR for Review Project - View, Edit, Add	• aiR for Review
 aiR for Review Prompt Criteria - View, Edit, Add 	Projects
 aiR Relevance Analysis - View 	
• aiR Issue Analysis - View	
• aiR Key Analysis - View	

You can run the job without permissions for the analysis types, but you won't be able to see the results.

3.3 Editing and running an existing aiR for Review project

To edit an existing aiR for Review project and run the analysis, you need the following permissions:

Object Security	Tab Visibility
• aiR for Review Project - View, Edit	• aiR for Review
 aiR for Review Prompt Criteria - View, Edit 	Projects

Object Security	Tab Visibility
• aiR Relevance Analysis - View	
• aiR Issue Analysis - View	
• aiR Key Analysis - View	

You can run the job without permissions for the analysis types, but you won't be able to see the results.

3.4 Viewing highlights in the Viewer

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

Object Security

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

You will see highlighting only for the analysis types you have permission to view.

3.5 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. 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.

3.5.1 Instance-level permissions

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

Tab Visibility	Admin Operations
 aiR for Review Jobs 	Admin Repository - View

Assign these permissions under the Instance Details tab.

3.5.1.1 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.

Without these, you can see jobs from that workspace but cannot click on them 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.

3.5.2 Workspace-level permissions

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

Object Security	Tab Visibility
 aiR for Review Prompt Criteria View 	 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.

3.6 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
Document - View, Edit	 aiR for Review Jobs
View, Edit	
 aiR Issue Analysis - View, Edit 	
• aiR Key Analysis - View, Edit	

You can only clear or restore results for analysis types if you have Edit permissions for them.

For more information on clearing and restoring results, see <u>Managing jobs and document linking on</u> page 35.

3.7 Running the aiR for Review mass action

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

Object Security	Other Settings
 aiR for Review Prompt Criteria - View,	Mass Operations - aiR for
Edit, Add	Review

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

4 Best practices

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.

4.1 Tips for writing Prompt Criteria

The Prompt Criteria entered often aligns with a traditional review protocol or case brief in that they describe the matter, entities involved, and what is relevant to the legal issues at hand.

When writing Prompt Criteria, use natural language to describe why particular types of documents should be considered relevant. Write them as though you were describing them to a human reviewer.

- Write clearly—use active voice, use natural speaking phrases and terms, be explicit.
- Be concise—write as if "less is more," summarize lengthy text or only include key passages from a long review protocol. The Prompt Criteria have an overall length limit of 15,000 characters.
- Simply describe the case-do not give commands, such as "you will review XX."
- Use positive phrasing—phrase instructions in a positive way when possible. Avoid negatives ("not" statements) and double negatives.
- Use natural writing format styles—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.
- Is it important?—ask yourself will the criteria affect the results, it is essential.
- Avoid legal jargon or explanations—for example, don't use "including but not limited to" and "any and all" and don't include explanations of the law.
- Use ALL CAPS—helps identify essential information for the model to focus on, for example use "MUST" instead of "should."
- Identify internal jargon and phrases—the learning language model (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 internal to an organization.
- Identify aliases, nicknames, and uncommon acronyms—for example, a nickname for William may be Bill, or BT may be an abbreviation for the company name Big Thorium.
- Identify unfamiliar emails—normal company email addresses do not need identified, but unfamiliar ones should, for example Dave Smith may use Dave.Smith@AcmeCompany.com and skippy78@gmail.com.
- Iterate, iterate, iterate—test the prompt criteria and review the results, adjust it to obtain more accurate predictions and results.

Refer to the helper examples in the Prompt Criteria text boxes of the dialogs for additional guidance entering criteria in each field.

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
- <u>Al Advantage: Aiming for Prompt Perfection?</u>—on-demand webinar discussing Prompt Criteria creation

4.2 Prompt criteria iteration sample documents

Before setting up the aiR for Review project, create a saved search that contains a small sample of the documents you want reviewed.

For best results:

- Include roughly 50-100 test documents that are a mix of relevant, not relevant, and challenging documents.
- Make sure they highlight all the key features of your relevance criteria.
- Have human reviewers code the documents in advance.

For more information about choosing documents for the sample, see <u>Selecting a Prompt Criteria Iteration</u> <u>Sample for aiR for Review</u> on the Community site.

See Creating or editing a saved search in the Searching guide for details about saved searches.

4.3 Prompt criteria iteration workflow

We recommend the following workflow for crafting Prompt Criteria:

- 1. For your first analysis, run the Prompt Criteria on a saved search of 50-100 test documents that are a mix of relevant, not relevant, and challenging documents.
- 2. Compare the results to human coding. In particular, look for documents that the application 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 the application predicts coding decisions accurately for the test documents.
- 5. Test the Prompt Criteria on a sample of 50 more documents and compare results. Continue tweaking and adding documents until you are satisfied with the results for a diverse range of documents.
- 6. Finally, run the Prompt Criteria on a larger set of documents.

aiR for Review 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.

5 Creating an aiR for Review project

The instructions you give aiR for Review are called Prompt Criteria. Prompt Criteria often mimic a traditional review protocol or case brief in that they describe the matter, entities involved, and what is relevant to the legal issues at hand. 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. Starting with a small set lets you see immediately how aiR's coding compares to a human reviewer's coding and adjust the prompts accordingly.

See Best practices on page 14 for more information.

5.1 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. You'll need to determine which type of analysis is needed before starting a project.

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 of the issues you want to analyze should be represented by a choice on the field.

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

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

5.2 Choosing the data source

Before setting up the project, create a saved search that contains a small sample of the documents you want reviewed.

For best results:

- Include roughly 50 test documents that are a mix of relevant and not relevant.
- Have human reviewers code the documents in advance.

For more information about choosing documents for the sample, see <u>Selecting a Prompt Criteria Iteration</u> Sample for aiR for Review on the Community site.

For more information about creating a saved search, see Creating or editing a saved search in the Searching guide.

5.3 Overview of setting up an aiR for Review project

Setting up a project for the first time has three basic parts:

- 1. Create an aiR for Review project
- 2. Write the Prompt Criteria
- 3. Run the first analysis

At any point in this process, you can save your progress and come back later.

SeeBest practices on page 14 for tips and workflow suggestions.

5.4 Step 1: Creating an aiR for Review project

To create an aiR for Review project:

- 1. On the aiR for Review Project tab, select New aiR for Review Project.
- 2. Fill out the following fields:
 - Project Name—enter a name for the project.
 - Description—enter a project description.
 - Data source—select the saved search that holds your document sample.
 - Project Prompt Criteria—select one of the following:
 - Start blank—select this if you plan to write new Prompt Criteria from scratch.
 - **Copy existing**—select this to choose a previously created set of Prompt Criteria and copy it for this project.
 - **Prompt Criteria Name**—either leave as the default or click the Edit (^{CC}) icon to rename the Prompt Criteria. This name must be unique.
 - **Analysis Type**—select one of the following. For more information, see <u>Choosing an analysis</u> type on the previous page
 - **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.
- Project Use Case—choose the option that best describes the purpose of the project.
 - If none of the options describe the project, choose **Other** to type your own description in the field. It will only be used for this project. Keep this description generic and do not include any confidential or personal information.
 - This field is used for reporting and management purposes. It does not affect how the project runs.
- 3. Click Create Project.

After the project is created, the aiR for Review project dashboard appears.

5.5 Step 2: Writing the Prompt Criteria

The Prompt Criteria are a set of inputs that give aiR for Review 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. See <u>Best practices on page 14</u> for tips and workflow suggestions.

Note: The Prompt Criteria have an overall length limit of 15,000 characters.

Depending which type of analysis you chose, you will see a different set of tabs on the left-hand side of the aiR for Review dashboard. All Prompt Criteria include the Case Summary tab.

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.

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
- <u>AI Advantage: Aiming for Prompt Perfection?</u>—on-demand webinar discussing Prompt Criteria creation

5.5.1 Case Summary tab

The Case Summary gives the LLM 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 20 or fewer sentences overall, and 20 or fewer each of People and Aliases, Noteworthy Organizations, and Noteworthy Terms.

To fill out the Case Summary tab:

- 1. On the aiR for Review dashboard, click the Case Summary tab.
- 2. Fill out the following:
 - **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.

Prelativity one

- **People and Aliases**—list the names and aliases of key custodians who authored or received the documents. Include their role and any other affiliations.
- **Noteworthy Organizations**—list the organizations and other relevant entities involved in the case. Highlight any key relationships or other notable characteristics.
- **Noteworthy Terms**—list and define any relevant words, phrases, acronyms, jargon, or slang that might be important to the analysis.
- Additional Context—list any additional information that does not fit the other fields. This section is typically left blank.
- 3. Click Save.

Depending on which Analysis Type you chose in Step 1: Creating an aiR for Review project, the remaining tabs will be called Relevance, Key Documents, or Issues. Fill out those tabs according to the guide sections below.

5.5.2 Relevance tab

If you chose either *Relevance* or *Relevance and Key Documents* as the Analysis 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. On the aiR for Review dashboard, click the **Relevance** tab.
- 2. Fill out the following:
 - **Relevance Field**—select a single-choice field that represents whether a document is relevant or non-relevant. This selection cannot be changed after the first job run.
 - **Relevant Choice**—select the field choice you use to mark a document as relevant. This selection cannot be changed after the first job run.
 - **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.
 - **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 <u>Issues tab on the next page</u>.

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.

3. Click Save.

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), since those are often too long and complex to give good results. Instead, summarize it and include relevant 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."

5.5.3 Key Documents tab

If you chose *Relevance and Key Documents* as the Analysis 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. On the aiR for Review dashboard, click the Key Documents tab.
- 2. Fill out the following:
 - **Key Document Field**—select a single-choice field that represents whether a document is key to the case. This selection cannot be changed after the first job run.
 - **Key Document Choice**—select the field choice you use to mark a document as key. This selection cannot be changed after the first job run.
 - **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.
- 3. Click Save.

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

5.5.4 Issues tab

If you chose *Issues* as the Analysis 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. On the aiR for Review dashboard, click the Issues tab.
- 2. Fill out the following:
 - **Field**—select a multi-choice field that represents the issues in the case. This selection cannot be changed after the first job run.
 - **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 relating to that issue, such as falling outside a certain date range.

Note: The field choices cannot be changed after the first job run. However, you can still edit the summary in the text box.

3. Click Save.

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.

5.5.4.1 Removing issue choices

aiR for Review analyzes a maximum of 10 choices. If the issue field has more than 10 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.

These choices cannot be changed after the first job run.

5.6 Step 3: Running the first analysis

After filling out the Setup, Case Summary, and other tabs, review the job and run the analysis.

To run the analysis:

- 1. On the upper right of the dashboard, click **Analyze [X] documents**. A confirmation modal appears.
- 2. Review the confirmation summary:
 - Total Docs—number of documents to be analyzed.
 - Est. Run Time—estimated time it will take to analyze and return the results of the documents selected. This does not include time waiting in the job queue.
 - Est. Time to Start—estimated wait time from job submission to the start of the analysis process. Longer wait times occur when other jobs are queued across tenants.
- 3. Set the **Email Notification** toggle on to send email notifications to other recipients when the job completes, fails, or is canceled. Then, enter the email addresses in the text box, separating them with commas or semicolons. You will receive an email notification when the job completes regardless of this toggle setting.
- 4. Click Start Analysis.

After the analysis job starts, the results start to appear beside each document in the Analysis Results panel of the dashboard. The Analyze button is disabled, and a **Cancel** option appears.

When the job completes, the Analyze button is re-enabled.

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.

After the first analysis completes, use the results to fine-tune the Prompt Criteria.

- For information on the dashboard and the fine-tuning process, see <u>Iterating on the prompt criteria on</u> page 25.
- For information about the results fields, see aiR for Review results on page 39.

5.7 Running aiR for Review as a mass operation

If you want to run previously refined Prompt Criteria on a set of documents, you have the option of running the job as a mass operation from the Documents list page.

To run aiR for Review as a mass operation:

- 1. From the **Documents** tab, select the documents you want to analyze.
- 2. Under Mass Actions, select aiR for Review. A modal appears.

aiR for Review					₽ai₹
		Select a prompt cri	teria for the analysis. Reviev You can review progress in	v the following estimates before starting this job. the aiR for Review jobs page.	
Prom	n pt Criteria * (The perf Use of aiR	Select Existing formance of Relatin for Review with no	Create New vity alR for Review has b n-English language docu	Total Docs Est. Run Time Est. Time to Start 	3 2 minutes 0 minutes uments. ble results.
					Start Analysis Cancel

- 3. For Prompt Criteria, select one of the following:
 - Select Existing—select and load a set of previously created Prompt Criteria from your workspace. This only shows Prompt Criteria that have been run at least once, and it selects the most recent version of them.
 - **Create New**—this closes the modal and redirects you to the aiR for Review Projects tab. Keep in mind this option does not save your selected documents from the step above.

Prelativity one

- After you have loaded a set of Prompt Criteria, click Start Analysis. A banner appears at the top of the page, confirming that the analysis job has been queued. This banner also updates to show when the job is complete.
- 5. To view and manage jobs that are not part of an existing project, use the aiR for Review Jobs tab. For more information, see Managing aiR for Review jobs on page 35.

5.8 Editing and collaboration

If two users are editing the same Prompt Criteria version at the same time, the user who last saves their work will have that work override the other one's changes. Because of this, we recommend having only one user edit a project's Prompt Criteria at a time. You may find it helpful to define separate roles for users when iterating on prompt changes.

5.9 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.

5.9.1 Size limits

The documents and Prompt Criteria have the following size limits:

- The Prompt Criteria have an overall length limit of 15,000 characters.
- We recommend only including documents whose extracted text is between 0.05KB and 150KB. Although the LLM can handle some larger or smaller documents, most will receive an error.

Note: The size of a job encompasses both the document size and the Prompt Criteria size. If the combined total exceeds the per document limit, the job will be too large to process.

If a document receives an error, your organization will not be charged for it. For more information, see <u>How</u> document errors are handled on page 47.

5.9.2 Volume limits

The per-instance volume limits for aiR for Review jobs are as follows:

Volume Type	Limit	Notes
Max job size	150,000 doc- uments	A single job can include up to 150,000 documents.
Total documents run- ning per instance	600,000 doc- uments	Across all jobs queued or running within an instance, there is a max- imum of 600,000 documents.
Concurrent large jobs per instance	3 jobs	For jobs with over 200 documents, only 3 jobs can be queued or running at the same time within an instance.
Concurrent small jobs per instance	No limit	Jobs with 200 or fewer documents have no limit to how many can queue or run at the same time.

5.9.3 Speed

After a job is submitted, aiR analyzes roughly 25-50 documents per minute. Job speeds vary widely depending on the number of documents, the overall load on the LLM, and other factors such as the size of the documents.

5.10 Understanding documents and billing

For billing purposes, a document unit is a single document. The initial pre-run estimate may be higher than the actual units billed because of canceled jobs or document errors. To find the actual document units that are billed, see Cost Explorer on the Relativity documentation site.

A document will be billed each time it runs through aiR for Review, regardless of whether that document ran previously.

Caution: Customer may not consolidate documents or otherwise take steps to circumvent the aiR for Review Document Unit limits, including for the purpose of reducing the Customer's costs. If Customer takes such action, Customer may be subject to additional charges and other corrective measures as deemed appropriate by Relativity.

6 Iterating on the prompt criteria

After running an aiR for Review job for the first time, the initial results can be used as feedback for improving the prompt criteria. The cycle of examining the results, fine-tuning the prompt criteria, then running a new job on the sample documents is known as iterating on the prompt criteria. Refer to <u>Best practices on</u> page 14 for more information.

6.1 Navigating the aiR for Review dashboard

When you select a project from the aiR for Review Projects tab, a dashboard displays showing the project's Prompt Criteria, the list of documents, and controls for editing the project. If the project has been run, it also displays the results.

6.1.1 Project details strip

At the top of the dashboard, the project details strip displays:



- **Project name**—the name chosen for the project. The Analysis Type appears underneath the name.
- Version selector—if this is the first version of the Prompt Criteria, this will say Version 1 as a textonly field. For later versions, the version number becomes clickable, and you can choose older versions to view their statistics. For more information, see <u>How Prompt Criteria versioning works on</u> <u>page 32</u>.
- Data source—name of the saved search chosen at project creation and the document count.
 - If you add or remove documents from the saved search, those changes are not reflected in the aiR for Review project until you refresh the data source.
 - To refresh the data source, click the refresh symbol beside the name.
- **Run button**—press this to analyze the selected documents using the current version of the Prompt Criteria. If no documents are selected or filtered, it will analyze all documents in the data source.
 - If you are viewing the newest version of the Prompt Criteria and no job is currently running, this says **Analyze [X] documents**.
 - If an analysis job is currently running or queued, this button is unavailable and a **Cancel** option appears.
 - If you are viewing older versions of the Prompt Criteria, this button is unavailable.
- Feedback icon—send optional feedback to the aiR for Review development team.

6.1.2 Prompt Criteria panel

On the left side of the dashboard, the Prompt Criteria panel displays tabs that match the project type you chose when creating the project. These tabs contain fields for writing the criteria you want aiR for Review to use when analyzing the documents.

Case Summary ⊘	Relevance ⊘	Key Documents ⊘	«	1
Matter Overview (?) *				Î
This is an investigation is specializing in sustainable (RFP) by the City of Atland the city's energy efficient the contract. It is alleged win the contract.	into a potential bribery ole energy solutions, v ntis for a municipal er cy and reducing its ca d that employees of B	A. BigThorium, a company vas issued a Request for Pro hergy project aimed at enhan rbon footprint, and ultimatel igThorium bribed city official	posal cing y won s to	
People and Aliases (?)				
John Maxwell (Sales Re about offering gifts and of Sales, BigThorium): C including the purchase of	presentative, BigThori incentives to secure t coordinates the strate of a Rolex watch. Robe	um): Engages in discussions he contract. Sarah Thompso gy and approval of incentives ert Williams (City Treasurer,	s n (VP s,	

Possible tabs include:

- Case Summary—appears for all analysis types.
- Relevance—appears for Relevance and Relevance and Key Documents analysis types.
- Key Documents—appears for the Relevance and Key Documents analysis type.
- Issues—appears for the Issues analysis type.

For information on filling out the Prompt Criteria tabs, see Step 2: Writing the Prompt Criteria on page 18.

If you want to temporarily clear space on the dashboard, click the Collapse symbol (\leq) in the upper right of the Prompt Criteria Panel. To expand the panel, click the symbol again.

6.1.3 Aspect selector bar

The aspect selector bar appears in the upper middle section of the dashboard for projects that use Issues analysis or Relevance and Key Documents analysis. This lets you choose which metrics, citations, and other results to view in the Analysis Results grid.

• For a Relevance and Key Documents analysis: Two aspect tabs appear: one for the field you selected as the Relevant Choice, and one for the field you selected as the Key Document Choice.

Responsive	Кеу
Conflicts 1	Conflicts 0

• For Issues analysis:

An aspect tab appears for every Issues field choice that has choice criteria. They appear in order according to each choice's Order value. For information on changing the choice order, see Choices in the Admin guide.

Disparate Treatmen	Segregated Conditi	Coercion & Underpa	Retaliation	Theft and Criminal A
Relevant 52	Relevant 11	Relevant 2	Relevant 29	Relevant 49

When you select one of the aspect tabs in the bar, both the project metrics section and analysis results grid update to show results related to that aspect. For example:

• If you choose the Key Document tab:

The project metrics section shows how many documents have been coded as key. The Analysis Results grid updates to show predictions, rationales, citations, and all other fields related to whether the document is key.

• If you choose an issue from the aspect selector:

The project metrics section and analysis results grid both update to show results related to that specific issue. The total number of issue predictions in this section is calculated by multiplying the number of issues by the number of documents. For example, if there are five issues and 100 documents, there will be 500 issue predictions.

6.1.4 Project metrics section

In the middle section of the dashboard, the project metrics section shows the results of previous analysis jobs. There are two tabs: one for the current version's most recent results (<u>Version Metrics tab below</u>), and one for a list of historical results for all previous versions (<u>History tab on page 29</u>).

Version 12	Version 12 Metrics History								
Documents				aiR Analysis		Conflicts			
Reviewer Coded Issues	Analyzed	Not Analyzed	Errored	Not Relevant	Borderline	Relevant	Total	Relevant	Not Relevant
2	100	0	0	64	12	24	24	23	1

6.1.4.1 Version Metrics tab

The Version [X] Metrics tab shows metrics divided into sections:

Documents section:

• Reviewer Uncoded (for Relevance or Relevance and Key Documents analysis only)—documents that do not have a value assigned in the Relevance Field. When viewing the Key aspect, this shows

documents that do not have a value assigned in the Key Field.

- **Reviewer Coded Issues** (for Issues analysis only)—total number of documents that reviewers coded as having the selected issue.
- **Analyzed**—documents in the data source that have a prediction attached from this Prompt Criteria version.
- **Not Analyzed**—documents in the data source that do not have a prediction attached from this Prompt Criteria version.
- **Errored**—documents that received an error code during analysis. For more information, see <u>How</u> document errors are handled on page 47.

aiR Analysis section:

- Not Relevant—documents predicted as junk or not relevant to the current aspect.
- **Borderline**—documents predicted as bordering between relevant and not relevant to the current aspect.
- Relevant—documents predicted relevant or very relevant to the current aspect.

Conflicts section:

- **Total**—total number of documents that have a different coding decision from the predicted result. This is the sum of the Relevant Conflicts and Not Relevant Conflicts fields.
- **Relevant**—documents predicted as relevant or very relevant to the current aspect, but the coding decision in the related field says something else.
- Not Relevant—documents predicted as not relevant to the current aspect, but the coding decision in the related field says relevant.

Depending which type of results you view, the metrics base their counts on different fields:

- When viewing Relevance results, the relevance-related metrics base their counts on the Relevance Field.
- When viewing Key Document results, the relevance-related metrics base their counts on the Key Document Field.
- When viewing results for Issues analysis, the relevance-related metrics base their counts on whether documents were marked for the selected issue.

For example, if you view results for an issue called *Fraud*, the aiR Predicted Relevant field will show documents that aiR predicted as relating to Fraud. If you view Key Document results, the aiR Predicted Relevant field will show documents that aiR predicted as being key.

Filtering the Analysis Results using version metrics

To filter the Analysis Results table based on any of the version metrics, click the desired metric in the banner. This narrows the results shown in the table to only documents that are part of the metric. It also auto-selects those documents for the next analysis job. The number of selected documents is reflected in the Run button's text. This makes it easier to analyze a subset of the document set instead of selecting all documents every time. To remove filtering, click **Clear selection** underneath the Run button.

	Ran 3	on 13 × /28/2025 3:24 PM			اللہ ا	100 documents	>		Analyze 2 docume Clear selection	nts 🔎
Related to Specific Atla Conflicts 23	Rela	ted to a City Contra Conflicts 36	Gift Giving or Inc Conflicts	entiviz 2	. M	fentions Rolex Watche Conflicts 0	es Discussio	n of Food Pla onflicts 2	Highly Confider Conflicts 0	ntial Privileg
Version 13 Metrics	Version 13 Metrics History									
	D	ocuments				aiR Analysis			Conflicts	
Reviewer Coded Issues	Analyzed	Not Analyzed	Errored	Not Re	elevant	Borderline	Relevant	Total	Relevant	Not Relevant
2	99	0	1	6	0	16	23	23	22	1
Analysis Results										
								7 %	1 - 2 of 2 200	▼ per page K<>>
# Control Numb	er	Reviewer Coded	aiR Prediction	aiR	Score	aiR Rationale		aiR Consideration		aiR Citation 1
⊽ Filter		V Filter	V Filter	7	Filter	V Filter		🖓 Filter		▽ Filter
1 BTV5_000005099 Related to Specific Not Relevan Atlantis RFPs Specific Atla		Not Relevant to Relate Specific Atlantis RFPs	ed to 1		The document mentions both Atlantis and BigThorium, but there is no mention of a deal or RFP.		The context of the email is about planning a celebration, not discussing a deal or RFP.		Robert Williams <robert.williams@atlantis.go To: Linda Patel <lpatel@bigt< td=""></lpatel@bigt<></robert.williams@atlantis.go 	
2 D BTv5_000005	109	Related to Specific Atlantis RFPs	Relevant to Related to Specific Atlantis RFPs	3		The document mentions A BigThorium together, indic deal or communication be entities.	Atlantis and cating a potential stween the two	The document does deal, only compliand might not be directly	not explicitly mention a ce and communication. It y related to a deal.	To: Robert Williams <robert.williams@atlantis.go< td=""></robert.williams@atlantis.go<>

You can also filter documents in the Analysis Results grid by selecting them in the table. See <u>Filtering and</u> selecting documents for analysis on the next page.

6.1.4.2 History tab

The History tab shows results for all previous versions of the Prompt Criteria. This table includes all fields from the Version Metrics tab, sorted into rows by version. For a list of all Version Metrics fields and their definitions, see Version Metrics tab on page 27.

It also displays two additional columns:

- Version—the Prompt Criteria version that was used for this row's results.
- **Timestamp**—the time the analysis job ran.

Versi	ion 13 Metrics Hi	story									
Version	Timestamp	Docs w/ aiR Predictions	Docs w/o aiR Predictions	Uncoded	Errored	Predicted Relevant	Predicted Not Rele	Predicted Borderline	Total Conflict	Predicted Relevant	Predicted Not Rele
12	10/9/2024, 2:11 PM	100	0		0	24	64	12	24	23	1
11	10/2/2024, 11:13 AM	100	0		0	26	65	9	26	25	1
10	10/1/2024, 10:10 AM	100	0		0	23	66	11	24	23	1
9	9/26/2024, 1:40 PM	100	0		0	24	66	10	24	23	1
8	9/26/2024, 1:33 PM	100	0		0	22	66	12	23	22	1
7	9/26/2024, 1:37 AM	100	0		0	27	63	10	26	26	0
6	9/25/2024, 6:57 PM	100	0		0	26	62	12	25	25	0
5	9/24/2024, 2:58 PM	100	0		0	25	64	11	24	24	0
4	9/22/2024, 11:14 PM	100	0		0	7	88	5	9	7	2
3	9/17/2024, 6:11 PM	100	0		0	7	88	5	9	7	2
2	9/17/2024, 6:00 PM	100	0		0	8	88	4	10	8	2
1	9/17/2024, 5:54 PM	100	0		0	8	88	4	10	8	2

6.1.5 Analysis Results section

In the middle section of the dashboard, the Analysis Results section shows a list of all documents in the project. If the documents have aiR for Review analysis results, those results appear beside them in the grid.

A	nalysis R	lesults						
					∑ 🖗 1 - 100 of 100 200	▼ per page K <>>>		
#		Control Number	Reviewer Coded	aiR Prediction	aiR Score	aiR Rationale	aiR Consideration	aiR Citation 1
		√ Filter	√ Filter	√ Filter	▽ Filter	Filter	√ Filter	∀ Filter
1		BTv5_000005028		Relevant to Related to Specific Atlantis RFPs	3	The document mentions BigThorium and Atlantis discussing emergency preparedness and response plans, which could be related to a potential deal.	The document does not explicitly mention a deal or RFP, so it could be a preliminary discussion.	I wanted to follow up on the emergency preparedness that we discussed during Given the recent events in is essential for BigThoriur compliance requirements.
2		BTv5_000005032		Not Relevant to Related to Specific Atlantis RFPs	1	The document mentions BigThorium, but there is no mention of Atlantis or any potential deal.	There might be indirect references or implications not immediately apparent.	Subject : Enhancing Risk № Contingency Planning at E
3		BTv5_000005037		Not Relevant to Related to Specific Atlantis RFPs	1	The document does not mention Atlantis, BigThorium, and any potential deal together, nor does it relate to BigThorium responding to an RFP.	There is a mention of BigThorium and Atlantis, but no deal or RFP is discussed.	
4		BTv5_000005042		Not Relevant to Related to Specific Atlantis RFPs	1	The document does not mention Atlantis, BigThorium, and any potential deal together, nor does it discuss BigThorium responding to an RFP.	There might be indirect references or implications not immediately apparent.	
5		BTv5 000005044		Not Relevant to Related to	1	The document mentions BiaThorium but	There might be other parts of the document	CEO. BigThorium

The fields that appear in the grid vary depending on what type of analysis was chosen. For a list of all results fields and their definitions, see aiR for Review results on page 39.

Note: aiR for Review's predictions do not overwrite the Relevance, Key, or Issues fields chosen during Prompt Criteria setup. Instead, the predictions are held in other fields. This makes it easier to distinguish between human coding choices and aiR's predictions.

To view inline highlighting and citations for an individual document, click on the **Control Number** link. The Viewer opens and displays results for the selected Prompt Criteria version. For more information on using aiR for Review in the Viewer, see aiR for Review Analysis on page 49.

6.1.5.1 Filtering and selecting documents for analysis

To manually select documents to include in the next analysis run, check the box beside each individual document in the Analysis Results grid. The number of selected documents is reflected in the Run button's text. To remove filtering, click **Clear selection** underneath the Run button.

nvestigation	iion 13 ▼ 3/28/2025 3:24 PM	R 100 document	0	Analyze 3 documents Clear selection				
Related to Specific Atla Conflicts 23 Conflicts 36	ntra Gift Giving or Incentiviz Conflicts 2	Mentions Rolex Watches Conflicts 0	Discussion of Food Pla Conflicts 2	Highly Confidential Conflicts 0	Privileged Potentially Conflicts 9			
Documents		aiR Analysi	5	Cont	licts			
Reviewer Coded Issues Analyzed Not A 2 99	nalyzed Errored 0 1	Not Relevant Borderline 60 16	Relevant 23	Total Relevence 23 22	ant Not Relevant 2 1			
Analysis Results	nalysis Results							
# Control Number Reviewer Coded	aiR Prediction aiR :	iR S 🍞 aiR Rationale	aiR Consideration	aiR Citat	ion 1			
∀ Filter V Filter	V Filter	V 3 V Filter	V Filter	V Filte	,			
1 🕢 BTV5_000005028	Relevant to Related to 3 Specific Atlantis RFPs	The document mentions BigTho Atlantis discussing emergency p and response plans, which could a potential deal.	rium and The document doe reparedness deal or RFP; so it co l be related to discussion.	s not explicitly mention a Juld be a preliminary Given that we Given that is essent compliant	to follow up on the topic of cy preparedness and response plans liscussed during our last meeting, recent events in the city, I believe it ial for BigThorium to align with local cc requirements.			
2 2 BTv5_000005109 Related to Specifi Atlantis RFPs	c Relevant to Related to 3 Specific Atlantis RFPs	The document mentions Atlantis BigThorium together, indicating deal or communication between entities.	and The document doe a potential deal, only compliar the two might not be direct	s not explicitly mention a To: Robe ice and communication. It <robert.v ly related to a deal.</robert.v 	rt Williams villiams@atlantis.gov>			
3 🖌 BTv5_000005122	Relevant to Related to 3 Specific Atlantis RFPs	The document mentions a poter collaboration between Atlantis a BigThorium, which is relevant to a deal between the two entities.	tial The document doe nd details about the d the issue of collaboration, whic	s not provide specific eal or the nature of the h might limit its relevance. I would li- between	ke to coordinate our upcoming to discuss the potential collaboration Atlantis and BigThorium.			
4 BTv5_000005132	Relevant to Related to 3 Specific Atlantis RFPs	The document mentions both Bi and the City of Atlantis, indicatin deal or collaboration between th entities.	gThorium The document doe g a potential deal, but the conter e two collaboration which	s not explicitly mention a ct suggests ongoing h could imply a deal. CC:	njamin Reed in reed@bigthorium.com> rt Williams villiams@atlantis.gov>			
				Subject: Nuclear	Re: Insights and Strategies from Project Experiences			
					ł			

You can also filter the Analysis Results grid by clicking the metrics in the Version Metrics section. See Filtering the Analysis Results using version metrics on page 28.

6.1.6 Clearing selected documents

Click **Clear selection** underneath the Run button to deselect all documents in the Analysis Results grid. This resets the selections, and the next analysis will run on all documents in the data source.



6.1.7 Saving selected documents as list

To save a group of documents as a list, follow the steps below.

- 1. Select the box beside each individual document in the Analysis Results grid that you want to add to the list.
- 2. Click the Save as List button at the bottom of the grid.
- 3. Enter a unique Name for the document list.

Save as List 3 Documents		
Name*	1	
Notes		
		Save Cancel

- 4. Enter any **Notes** in the text box to help describe the list.
- 5. Click Save.

For more information on lists, see Lists in the Admin guide.

6.1.8 Exporting Prompt Criteria

Use the Export option to export the contents of the most recently saved Prompt Criteria tabs to an MS Word file. This can be helpful for reviewing the criteria, for saving it to use later, and for collaborating with others on it.

Note: Only the most recently saved version of the Prompt Criteria is exported. You cannot export earlier versions.

1. In the Prompt Criteria panel, click the **More** (

Case Summary ⊘	$\operatorname{Relevance} \oslash$	Key Documents ⊘	~	÷	
Matter Overview (?)*					
This is an investigation specializing in sustainal (RFP) by the City of Atla the city's energy efficien the contract. It is alleged win the contract.	into a potential bribery ble energy solutions, v ntis for a municipal er cy and reducing its ca d that employees of B	y. BigThorium, a company vas issued a Request for Pro hergy project aimed at enhan rbon footprint, and ultimatel igThorium bribed city official	posal icing y won s to		
People and Aliases (?)					
John Maxwell (Sales Re about offering gifts and of Sales, BigThorium): C including the purchase of	presentative, BigThori incentives to secure t coordinates the strates of a Rolex watch. Robe	um): Engages in discussions he contract. Sarah Thompso gy and approval of incentives ert Williams (City Treasurer,	3 n (VP 3,		

2. Select Export.

3. Click Save and Export to proceed with exporting the content.

The exported file is saved to your default browser download folder. It contains the information from all available criteria tabs for the most recently saved Prompt Criteria. Audit logs track all export actions from your project.

6.2 How Prompt Criteria versioning works

Each aiR for Review project comes with automatic versioning controls, so that you can compare results from running different versions of the Prompt Criteria. Each analysis job that uses a unique set of Prompt Criteria counts as a new version.

When you run aiR for Review analysis, the initial Prompt Criteria are saved as Version 1. Edits to the criteria create Version 2, which you can repeatedly modify until you finalize by running the analysis again to see the results. Subsequent edits follow the same pattern, creating new versions that finalize with each analysis run.

To see dashboard results from a earlier version, click the arrow next to the version name in the project details strip. From there, select the version you want to see.



6.2.1 How version controls affect the Viewer

When you select a Prompt Criteria version from the dashboard, this also changes the version results you see when you click on individual documents from the dashboard. For example, if you are viewing results from Version 2, clicking on the Control Number for a document brings you to the Viewer with the results and citations from Version 2. If you select Version 1 on the dashboard, clicking the Control Number for that document brings you to the Viewer with results and citations from Version 1.

When you access the Viewer from other parts of Relativity, it defaults to showing the aiR for Review results from the most recent version of the Prompt Criteria. However, you can change which results appear by using the linking controls on the aiR for Review Jobs tab. For more information, see <u>Managing aiR for</u> Review jobs on page 35.

6.3 Revising the Prompt Criteria

After you run the analysis for the first time on a sample set, use the dashboard to examine the results and refine the Prompt Criteria.

In particular, ask the following questions about each document:

- Did aiR for Review and the human reviewer agree on the relevance of the document?
- Read the aiR for Review rationale and considerations. Do they make sense?
- Do the citations make sense?

For all of these, if you see something incorrect, make notes on where aiR seems to be confused. Here are the most common sources of confusion:

- Insufficient context. For example, an internal acronym, key person, or code word may not have been defined. To fix this, add it to the proper section of the Case Summary tab.
- Ambiguous instructions or unclear language. To fix this, edit the instructions on the Relevance, Key Documents, or Issues tabs.

In general, consider how you would help a human reviewer making the same mistakes. For example, if aiR for Review is having trouble identifying a specific issue, try explaining the criteria for that issue with simpler language.

After you have revised the Prompt Criteria to address any weak points, run the analysis again. Continue refining the Prompt Criteria until results accurately predicts the human coding decisions for all test documents in the sample.

Note: aiR for Review only looks at the extracted text of each document. If a human reviewer marked a document as relevant because of an attachment or other criteria beyond the extracted text, aiR for Review will not be able to match that relevance decision.

6.4 Increasing the job size

When aiR for Review accurately matches human coding decisions on the initial sample documents, increase the sample size. Typically, we recommend starting with an initial sample of about 50 documents, then increasing it to include another 50. However, you may find a different number works better for your project.

To increase the job size:

- 1. Add the fresh documents to the saved search that acts as the project's data source. For more information about saved searches, see Creating or editing a saved search in the Searching guide.
- 2. Have a skilled human reviewer review the fresh documents. We recommend doing this before running aiR for Review, so that the reviewer is not biased by aiR's predictions.
- 3. On the aiR for Review Projects tab, select the project.
- 4. At the top of the project dashboard, click the refresh symbol next to the data source's name.
- 5. In the Project Metrics section, click **Not Analyzed**. This selects the new documents.
- After the document count has updated, click Analyze [X] documents. The analysis job runs on the new documents, while the previously run documents keep their old results.

After you have run the job on the larger sample, continue revising the Prompt Criteria until it returns satisfactory results. Continue to increase the job size incrementally until you feel satisfied with the Prompt Criteria. After that, use the refined Prompt Criteria on the larger set of documents. You can do this either from the dashboard, or as a mass operation.

For more information, see the following articles on the Community site:

- Selecting a Prompt Criteria Iteration Sample for aiR for Review
- Evaluating aiR for Review Prompt Criteria Performance

7 Managing aiR for Review jobs

Use the aiR for Review Jobs tab to monitor job progress, view prompt criteria details, or cancel the job. You can also view completed jobs and choose which analysis results are connected to the documents.

7.1 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 the most recent 100 jobs across all workspaces. It includes several extra columns to identify the workspace, matter, and client connected to each job.

The workspace-level tab shows all jobs for an individual workspace. 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 <u>Permissions on page 11</u>.

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.

7.2 How aiR for Review document linking works

When aiR for Review analysis has been run multiple times on the same document, each set of results is saved as part of a separate job. By default, when you look at a document's results, the results from the most recent analysis job are displayed. However, if you want to see the results from a previous job instead, you can use the aiR for Review Jobs tab to link an older job's results to the document. Each set of results can be linked or unlinked at any time without losing any data.

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

If you are viewing results from within the aiR for Review dashboard, the project version you select from the dashboard controls which job's results you see. If you are viewing results from other parts of Relativity, such as Review Center or the Document list, the job selected from the aiR for Review Jobs tab takes precedent. For more information on version selection in the dashboard, see <u>How version controls affect the Viewer on page 33</u>.

7.3 Managing jobs and document linking

You can use the aiR for Review Jobs tab to cancel jobs, hide job results from the document fields, and make hidden job results visible.

To manage jobs, use the following icons:

- **Cancel symbol** (⁽²⁾)—the job is currently queued or in progress. Clicking the symbol cancels the 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.
- Visible or Partially Hidden symbol ($^{\odot}$)—some or all of the job results are linked to documents.
 - If the Results column says Visible, this means that all documents from this job show this job's results in the Viewer. Clicking the symbol unlinks the job results and hides them from the Viewer.
 - If the Results column says Partially Hidden, this means only some of the documents from this
 job show this job's results in the Viewer. For example, if a few documents from this job were
 later included in a different job, they might have that more recent job's results showing instead.
 Clicking the symbol gives you a choice to either unlink and hide the results on all documents, or
 re-link and make the results visible on all documents.
- **Hidden symbol** (^{²⁰})—none of the job results are linked to documents. Clicking the symbol re-links the job results to the documents in the run and makes them visible in the Viewer. If the documents are currently linked to another job with the same result type, those results will be hidden.

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, hide the previous result set using the aiR for Review Jobs tab.

7.4 Viewing job details

To see the Prompt Criteria for an aiR for Review job, go to the aiR for Review Jobs tab and click within its row. A detail panel opens showing the setup details, case summary, fields, and criteria for analysis.

				aiR for Rev	iew aiR for Review	Jobs කි	Q Searc	h (CTRL	+ /)	\$ <u>₹</u>	Ask Al
aiR for Review Projects	aiR for Review Jobs										
₽aiR					▽ % 1 -1	00 of 477 100 🔻	per page KKN	В	T Investigation		
Results	Job ID	Project Name	Prompt Criteria Na	Version = • 🖓	Service Job ID	aiR for Review Vers	Job Status	AI R	etup nalysis Type ③ elevance		
No Results							CANCELLED	-			
🚿 Hidden	992	BT Investigation	BT Investigation	2	42ffe154-f8c5-4ed3- b2e6-fd584e53ae66	1.3.0	COMPLETED	C M	ase Summary latter Overview ③		
🚿 Hidden	005	ana tao tahun tahun da ta	ana kan bisadah cada ila	^	£1105000 6000	100	COMPLETED	m	atter.		
🚿 Hidden							COMPLETED	R	elevance		
💋 Hidden							COMPLETED	R	elevance Field esponsiveness		
💋 Hidden							COMPLETED	SI Ri	upportive Choice esponsive		
💋 Hidden							COMPLETED	Re	elevance Criteria ⑦ sponsive if it says anything funny		
n Hidden							COMPLETED				
🚿 Hidden					85ae-2h0h/108chhd		COMPLETED				

You can control a user's access to the detail panel using both item-level and workspace-level permissions. For more information, see <u>Viewing the aiR for Review Jobs tab on page 12</u> in <u>Permissions on page 11</u>.

7.5 Jobs tab fields

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

- **Results**—whether the job results are linked and visible on their corresponding documents. For more information, see Managing jobs and document linking on page 35.
- Job ID—the unique ID assigned to a job.
- **Project Name**—the name of the aiR for Review project associated with the job. To view the project, click on the project name.
- **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.
- Version—the Prompt Criteria version associated with the job. For more information, see <u>How Prompt</u> Criteria versioning works on page 32.
- **aiR for Review Version**—the version number of aiR for Review's internal model at the time the job ran. Please note that this is different from the large language model version.
- Job Status—the current state of the job. The possible statuses are:
 - Not Started
 - Queued
 - In Progress
 - · Completed
 - Cancelling
 - Errored
- Client Name (instance-level only)-the client associated with the job's workspace.
- Matter Name (instance-level only)-the matter name associated with the job's workspace.
- Matter Number (instance-level only)—the matter number associated with the job's workspace.
- Workspace ID (instance-level only)-the ID of the job's workspace.
- Workspace Name (instance-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 after the wait time.

8 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 as fields on document lists.

8.1 How aiR for Review analysis results work

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

- **aiR Prediction**—the relevance, key, or issue label that aiR predicts should apply to the document. See Predictions versus document coding below
- **aiR Score**—a numerical score that indicates how strongly relevant the document is or how well it matches the predicted issue. See Understanding document scores on the next page.
- aiR Rationale—an explanation of why aiR chose this score and prediction.
- **aiR Considerations**—a counterargument explaining why the prediction might possibly be wrong.
- **aiR Citation [1-5]**—excerpts from the document that support the prediction and rationale.

An	alysis R	esults						
							∑ % 1 - 100 of 100 200	▼ per page KC>>
#		Control Number	Reviewer Coded	aiR Prediction	aiR Score	aiR Rationale	aiR Consideration	aiR Citation 1
		∀ Filter	Filter Filter Second	√ Filter	🖓 Filter			V Filter
1		BTv5_000005028		Relevant to Related to Specific Atlantis RFPs	3	The document mentions BigThorium and Atlantis discussing emergency preparedness and response plans, which could be related to a potential deal.	The document does not explicitly mention a deal or RFP, so it could be a preliminary discussion.	I wanted to follow up on the emergency preparedness that we discussed during Given the recent events in is essential for BigThoriur compliance requirements.
2		BTv5_000005032		Not Relevant to Related to Specific Atlantis RFPs	1	The document mentions BigThorium, but there is no mention of Atlantis or any potential deal.	There might be indirect references or implications not immediately apparent.	Subject : Enhancing Risk N Contingency Planning at E
3		BTv5_000005037		Not Relevant to Related to Specific Atlantis RFPs	1	The document does not mention Atlantis, BigThorium, and any potential deal together, nor does it relate to BigThorium responding to an RFP.	There is a mention of BigThorium and Atlantis, but no deal or RFP is discussed.	
4		BTv5_000005042		Not Relevant to Related to Specific Atlantis RFPs	1	The document does not mention Atlantis, BigThorium, and any potential deal together, nor does it discuss BigThorium responding to an RFP.	There might be indirect references or implications not immediately apparent.	
5		BTv5 000005044		Not Relevant to Related to	1	The document mentions BioThorium but	There might be other parts of the document	CEO. BigThorium

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-scoring 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 does not consider this a true change of venue request.

8.1.1 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. This makes it easier to compare aiR's predictions to human coding while refining the Prompt Criteria.

If you have refined a set of Prompt Criteria to the point that you are comfortable adopting those predictions, you can copy those predictions to the coding fields using mass-tagging or other methods.

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 46.

8.1.2 Variability of results

Due to the nature of large language models, output results may vary slightly from one run to another, even using the same inputs. aiR's scores may shift slightly, typically between adjacent levels, such as from 1-not relevant to 2-borderline. Significant changes, like moving from 4-very relevant to 1-not relevant, are rare.

8.2 Understanding document scores

aiR scores 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. A score of -1 is assigned to any errored documents. Because these documents were not properly analyzed, they cannot receive a normal score.

The aiR for Review scores are:

Score	Description
-1	The document either encountered an error or could not be analyzed. For more information, see How document errors are handled on page 47.
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.

8.3 Viewing results from the dashboard

Within a project, you can view results using the dashboard. This dashboard includes not just results fields, but calculated metrics such as the number of documents with predictions that conflict with human coding.

To view the dashboard, select a project from the aiR for Review Projects tab. For detailed information on the dashboard layout, see <u>Navigating the aiR for Review dashboard on page 25</u>.

8.4 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 in the Viewer, click on the **aiR for Review Analysis** icon (⁽⁾) to expand the pane. The aiR for Review Analysis pane displays the following:

- 1. Prompt Criteria version
- 2. Analysis Name
- 3. Prediction
- 4. Rationale and Considerations
- 5. Citation

For more information, see Viewer documentation.

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, hide the previous result set using the aiR for Review Jobs tab.

8.4.1 Citations and highlighting

A maximum of five citations will be displayed with the document.

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 Color Map on the Relativity documentation site.



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

8.4.2 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.

8.5 Creating document views and saved searches

In addition to using the dashboard, you can view and compare the 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, hide the previous result set using the aiR for Review Jobs tab.

8.5.1 Creating an aiR for Review results view

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

Prelativity one

- Edit
- Control Number
- <Review Field>
- aiR <Review Type> Analysis::Score
- 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 below.

8.5.2 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 several 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.

8.6 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 Issue Analysis
- aiR Key Analysis

Additionally, the results are linked to each of the analyzed documents. These linked fields, called reflected fields, update to link to the newest results every time the document is analyzed. However, the application 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 35.

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.

8.6.1 aiR Relevance Analysis fields

The fields for aiR Relevance Analysis are:

Field name	Field type	Description
Name	Fixed-	The name of this specific result. This formatted as <document artifact<="" td=""></document>

Field name	Field type	Description
	length Text	ID>_ <job id="">.</job>
Job ID	Fixed- length Text	The unique ID of the job this result came from.
Score	Whole Number	Numerical score indicating how strongly relevant the document is. For more information, see <u>Understanding document scores on page 40</u> .
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 score and prediction.
Considerations	Fixed- length Text	A counterargument explaining why the prediction 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- 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 <u>How document errors are handled on page 47</u> .

8.6.2 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></document>
Job ID	Fixed- length Text	The unique ID of the job this result came from.
Choice Analyzed	Fixed- length 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.
Score	Whole Number	Numerical score indicating how well the document matches an issue. For more information, see <u>Understanding document scores on page 40</u> .
Prediction	Fixed- length Text	aiR's predicted issue choice for this document.
Rationale	Fixed- length Text	An explanation of why aiR chose this score and prediction.
Considerations	Fixed- length Text	A counterargument explaining why the prediction 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 <u>How document errors are handled on page 47</u> .

8.6.3 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></document>
Job ID	Fixed- length Text	The unique ID of the job this result came from.

Field name	Field type	Description
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.
Score	Whole Number	Numerical score indicating how strongly relevant the document is. For more information, see <u>Understanding document scores on page 40</u> .
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 score and prediction.
Considerations	Fixed- length Text	A counterargument explaining why the prediction 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- 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 <u>How document errors are handled on the next page</u> .

8.7 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 scored as borderline or above for relevance. To do that:

- 1. Set up a saved search for documents where **aiR Relevance Analysis::Score** is greater than **1**. This returns all documents scored 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.

8.8 How document errors are handled

If aiR encounters a problem when analyzing a document, it will not return results for that document. Instead, it scores the document as -1 and returns an error message in the Error Details column. Your organization is not charged for any errored documents, and they do not count towards your organization's aiR for Review total document count.

The possible error messages are:

Error message	Description	Retry?
Failed to parse com- pletion	The large language model (LLM) encountered an error.	
Completion is not valid The LLM encountered an error. JSON		Yes
Ungrounded citations detected in completion	The results for this document have a chance of including an ungroun- ded citation. For more information, see <u>Ungrounded citations below</u> .	Yes
Document text is empty	The extracted text of the document was empty.	No
Document text is too short	There was not enough extracted text to analyze in the document.	No
Document text is too long	The document's extracted text was too long to analyze.	No
Model API error occurred	A communication error occurred between the large language model (LLM) and Relativity. This is usually a temporary problem.	Yes
Uncategorized error occurred	An unknown error occurred.	Yes

If the Retry? column says **Yes**, you may get better results trying to run that same document a second time. For errors that say **No** in that column, you will always receive an error running that specific document.

If you retry a document and keep receiving the same error, the document may have permanent problems that aiR for Review cannot process.

8.9 Ungrounded citations

An ungrounded citation may occur for two reasons:

- When the aiR results citation cannot be found anywhere in the document text. This is usually caused by formatting issues. However, just in case the LLM is citing sentences without a source, we mark it as a possible ungrounded citation.
- 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 ungrounded citations are marked as errors, and they receive a score of -1 instead of whatever score they were originally assigned. If retrying documents with these errors does not succeed, we recommend manually reviewing them instead.

Actual ungrounded citations are extremely rare. However, highly structured documents, such as Excel spreadsheets and PDF forms, are more likely to confuse the detector and trigger these errors.

9 aiR for Review Analysis

Once you have run aiR for Review on a set of documents, you can open a document in the Viewer and see the aiR results highlighted as well as analysis for each one that explains why it might and might not be relevant to your case. This information helps you make decisions on what documents are useful as part of the review process.

9.1 Permissions

The following permissions are required to use aiR for Review Analysis:

Object Security	Tab Visibility	Other Settings
 aiR Relevance Analysis View 		
 aiR Issue Analysis - View 		
• aiR Key Analysis - View		

Note: If a user is not granted any of these permissions, they will not see the aiR for Review Analysis icon. If a user is granted some but not all of the permissions above, the aiR for Review Analysis icon is visible and the user will be able to see highlights and analysis for the granted permissions.

9.2 Reviewing aiR for Review Analysis results

To view the analysis of aiR results, click on the **aiR for Review Analysis** icon to expand the pane. The aiR for Review Analysis pane displays the following:

- 1. Prompt Criteria version
- 2. Analysis Name
- 3. Prediction
- 4. Rationale and Considerations
- 5. Citation



To adjust or interact with the highlighting:

• To navigate to a specific citation in the document, click on the citation card.

Note: Citation highlighting works best in the Extracted Text Viewer. If citation highlights do not appear in the Native Viewer, we recommend trying the Extracted Text Viewer.

• To turn highlighting on or off, click on the aiR for Review Analysis toggle.

9.3 aiR for Review Analysis highlighting

The highlight colors vary depending on the type of citation. Relevance citations are orange, key document citations are purple, and issues are assigned a range of colors from the Color Map application. If two issues use the same citation, the highlight blends their colors.

Prelativity one

If two or more highlights overlap in the Viewer, the overlapping parts will be a blend of the highlight colors. For example, if you have a yellow and green highlight that overlap, the overlapping text will be highlighted in blue.

Bill

Thank you for the update. You have my permission to increase the housing cost for the Indian workers. Inform them that due to not meeting construction deadlines on time, they will have to make up for the additional housing cost increase by having their pay docked. Inform the workers that if they don't step up production and stop complaining they will have to return home. Of course we know that won't happen.

Let me know if this resolves the issue.

Frank

9.4 aiR for Review Analysis results order

The results in the aiR for Review Analysis pane are first ordered by:

- Relevance citation
- Key Document citation
- Issue citation

The Issue results are ordered according to each issue choice's Order value. For information on changing the choice order, see Choices in the Admin guide.

Finally, duplicate results are ordered from most recent to oldest.

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.

©2025. Relativity ODA LLC. All rights reserved. Relativity® is a registered trademark of Relativity ODA LLC.