



aiR Review Guide

October 7, 2026

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

Table of Contents

1 aiR Review (air for Review)	6
1.1 aiR Review process overview	7
1.2 aiR Review phases	7
1.3 Analysis review types	8
1.4 How it works	9
1.5 Security and privacy	9
1.6 Generating document summaries	9
1.7 aiR units and consumption	9
1.8 Language support	10
1.9 Analyzing emojis	10
2 Installing aiR Review	10
2.1 Object types available	11
2.2 Tabs available	11
2.3 Configuring aiR Review job size limits	12
3 Permissions	12
3.1 Instance-level permissions	12
3.2 Workspace-level permissions	12
4 Best practices	15
4.1 Tips for writing prompt criteria	16
4.2 General prompt criteria iteration workflow	17
4.3 Tips for creating sample documents set	17
4.4 Best practices for writing Custom insight prompts	18
4.4.1 General tips	18
4.4.2 Extraction tips	21
4.4.3 Classification tips	23
4.4.4 Narrative/Insight tips	24
5 Common analysis workflows	26
5.1 Custom analyses	27
5.1.1 Considerations	27
5.1.2 How Custom analyses insights work	28
5.1.3 Custom analyses workflow	29
5.1.4 Setting up a project	29

5.1.5 Adding Custom insights and prompt criteria	30
5.1.6 Running a Custom analyses job	33
5.1.7 Viewing & iterating on the analysis results	34
5.1.8 Applying the results and publishing to the Document Object	35
5.1.9 Exporting Custom analyses results	37
5.1.10 Handling document errors in Custom analyses	37
5.1.11 Common use cases	39
5.1.12 Best practices	39
5.1.13 Known limitations	39
6 Navigating the aiR Review dashboard	40
6.1 Project details strip	40
6.2 Prompt criteria panel	40
6.3 Aspect selector bar	41
6.4 Project metrics section	42
6.4.1 Version Metrics tab	42
6.4.2 History tab	45
6.5 Results and summaries	45
6.5.1 Analysis Results tab	45
6.5.2 Summaries tab	46
6.5.3 Viewing inline highlighting and citations in the Viewer	46
6.5.4 Filtering and selecting documents for analysis	46
6.5.5 Clearing selected documents	47
7 Creating an aiR Review project	47
7.1 Project workflow	47
7.2 Job capacity, size limitations, and speed	48
7.2.1 Size limits	48
7.2.2 Volume limits per instance	49
7.2.3 Speed	49
7.3 Setting up the project	50
7.3.1 Choosing the data source	50
7.3.2 Choosing an analysis type	50
7.3.3 Setting up an aiR Review project	51
7.4 Deleting an aiR Review project	53
7.4.1 Hiding document analysis results in the Viewer after deleting a project	54

7.5 Using project sets	55
7.5.1 Navigating the project set details strip	55
7.5.2 Creating a new project set	56
7.5.3 Validating the prompt criteria	56
7.5.4 Applying the prompt criteria	57
7.5.5 Refreshing a data source	57
8 Developing prompt criteria	58
8.1 General prompt criteria workflow	58
8.2 Entering prompt criteria	58
8.3 Prompt criteria tabs and fields	59
8.3.1 Case Summary tab	59
8.3.2 Relevance tab	60
8.3.3 Key Documents tab	61
8.3.4 Issues tab	62
8.3.5 Confidential Business Information tab	62
8.3.6 Custom tab	64
8.4 Using prompt kickstarter (Draft with AI)	64
8.5 Editing and collaboration	67
8.6 How prompt criteria versioning works	68
8.7 How prompt criteria versions affect the Viewer	68
8.8 Revising the prompt criteria	68
8.9 Exporting prompt criteria	69
9 Running the analysis	70
10 Analyzing aiR Review results	72
10.1 How aiR Review analysis results work	72
10.1.1 Predictions versus document coding	73
10.2 Issues Field versus Issues analysis type for aiR Relevance Analysis::Issues	73
10.3 Document summaries	73
10.3.1 Documents with an empty summary field	74
10.4 Custom insight analyses results	74
10.5 Understanding document scores	75
10.6 Handling document errors	75
10.7 Variability of results	76
10.8 Ungrounded citations	76

10.9 Viewing and filtering analysis results	77
10.9.1 Viewing results from the aiR Review dashboard	77
10.9.2 Viewing results for individual documents from the Viewer	77
10.10 Exporting analysis results	80
10.10.1 Export capacity	81
10.10.2 Exported file structure	81
10.10.3 Exporting document summaries	81
10.10.4 Exporting Custom analyses results	82
10.11 Saving selected documents as a list	82
10.12 Creating a saved search	83
11 aiR Review prompt criteria validation	83
11.1 Prerequisite	84
11.2 Benefits of validation	84
11.3 How validation fits into aiR Review	84
11.4 High-level prompt criteria validation workflow	85
11.5 Setting up aiR Review prompt criteria validation	88
11.5.1 Applying the validated results or developing different prompt criteria	90
11.5.2 Validation Settings fields	91
11.6 Prompt Criteria validation in Review Center	92
11.6.1 How Review Center fits into the validation process	92
11.6.2 Managing the coding process	92
11.6.3 Reviewing validation statistics	93
11.6.4 Accepting or rejecting results	94
11.6.5 How changes affect the validation results	95
11.7 Prompt Criteria validation statistics	96
11.7.1 Defining the validation statistics	96
11.7.2 How documents are categorized for calculations	97
11.7.3 Prompt Criteria validation metric calculations	97
11.7.4 How the confidence interval works	99
11.7.5 How Prompt Criteria validation differs from other validation types	99
12 Managing aiR Review jobs	99
12.1 Setting aiR Review job size limits	100
12.2 Monitoring Review Jobs	101
12.2.1 Jobs list columns	101

12.2.2 Project details panel fields	101
12.3 Linking/Unlinking analysis types	103
13 Creating document views	103
13.1 Creating an aiR Review results view	104
13.2 aiR Review results fields	104
13.2.1 aiR Summaries fields	105
13.2.2 aiR Relevance Analysis fields	105
13.2.3 aiR Issues Analysis fields	106
13.2.4 aiR Key Analysis fields	107
13.2.5 aiR Confidential Business Information Analysis fields	108
13.2.6 aiR Custom Analyses fields	109
14 Using aiR Review with Review Center	110
14.1 Using prompt criteria validation	110
14.2 Using aiR to create a Review Center queue	110
14.2.1 Selecting documents to add to a queue	110
14.2.2 Creating a saved search queue through the aiR Review dashboard	110
15 Archiving and restoring workspaces with aiR Review	111
15.1 Archiving workspaces	111
15.2 Restoring workspaces	111

1 aiR Review (air for Review)

Note: aiR for Review is now aiR Review. References to aiR for Review may still appear in this documentation, in the product, and in other sources such as marketing materials, published content, and commercial or legal agreements. Both names refer to the same product/feature.

aiR Review is powered by generative AI, large language models (LLM), and natural language processing to examine extracted document text based on your prompt instructions. It helps pinpoint significant documents, explains their significance, and provides citations from those documents where applicable.

A few benefits of the application include:


- Highly efficient, low-cost document analysis
- Quick discovery of important documents to the case based on prompt criteria
- Early detection of documents containing confidential business information based on prompt criteria
- Trustworthy predictions, explanations, and citations
- Consistent, cohesive analysis across all documents

Below are some common use cases for it:

- **Beginning 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**—help reduce the effort to find important material and get it into the hands of decision makers.
- **Quality control for traditional review**—compare aiR Review's coding predictions to decisions made by reviewers to accelerate QC and improve results.

1.1 aiR Review process overview


The aiR Review process uses Relativity instructions, SME-developed prompt criteria, and extracted document text as input. It processes this data to produce document citations with associated rationale, considerations on possible prediction errors, and its recommendation (such as relevant, not relevant, or borderline). aiR then verifies citations per document to help assess accuracy and reduce the likelihood of hallucinated content (up to five citations per document are allowed for Relevance and Relevance and Key Documents analyses, and one citation per issue for Issues analysis).



1.2 aiR Review phases

The workflow is similar to training a human reviewer: explain the case and its criteria, hand over the documents, and check the results. If the application appears to misunderstand aspects of the prompt criteria, refine the instructions and iterate to improve results.

The workflow has three phases:



1. **Develop**—write and iterate on the prompt criteria (review instructions) and test it on a small document set until aiR's recommendations are reasonably aligned with expected results based on validation criteria.
2. **Validate**—run the prompt criteria on a slightly larger set of documents and compare to results from reviewers to

evaluate alignment and identify differences.

3. **Apply**—use the validated prompt criteria on larger sets of documents, or the full set.

Within Relativity aiR, the main steps are:

1. **Select the documents to review.** See [Choosing the data source on page 50](#) for more information.
2. **Create an aiR Review project.** See [Creating an aiR Review project on page 47](#) for more information.
3. **Write and submit the prompt criteria (review instructions).** See [Developing prompt criteria on page 58](#) for more information.
4. **Review the results.** See [Analyzing aiR Review results on page 72](#) for more information.

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.

For additional workflow help and examples, see [Workflows for Applying aiR Review](#) on the Community site (login credentials required).

For information on document and prompt size limits, as well as overall job volume limits, see [Job capacity, size limitations, and speed on page 48](#).

1.3 Analysis review types

aiR Review offers the following analysis types, each suited to a specific review or investigation phase.

Analysis Type	Description
Relevance	Use for relevance reviews to help identify documents that are relevant to a case matter or situation that you describe, such as documents responsive to a production request.
Key Documents	Use to help identify documents that may be considered <i>hot</i> or important to a case or investigation, such as those that might be critical or embarrassing to one party or another.
Issues	Use for issues reviews to find documents that include content that falls under specific categories or legal issues that are relevant to the case. For example, you might use this to check whether documents involve coercion, retaliation, or a combination of both.
Confidential Business Information (CBI)	Use for a confidentiality review to find and classify documents that include confidential/sensitive information on the selected CBI types to ensure proper handling and protection. For example, use to locate documents containing business strategy, trade secrets, or non-public financial data. Note: This functionality is being released through a phased rollout and is only available to clients in the US and EU Data Zone at this time. For more information on phased rollouts, see the Phased Rollout FAQ article available on the Community.
Custom	Use to create your own custom prompt-driven insights per project to analyze document extracted text and native images based on your current needs. Each custom insight acts as an independent analysis for questions, extractions,

Analysis Type	Description
	classifications, or evaluative instructions applied to every eligible document. For example, extracting themes from documents and object detection in images.

1.4 How it works

aiR Review uses an Azure OpenAI LLM to analyze documents based on prompt criteria. Each document is reviewed separately, with results returned before the next document is sent. The LLM bases its predictions solely on prompt criteria and internal training, unlike Review Center, which learns from the document set.

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

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

1.5 Security and privacy

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.

Also see AI products and features in Relativity aiR on the help site for additional details.

1.6 Generating document summaries

aiR Review automatically generates descriptive topics and summaries for each document in a project, except for Custom Analyses projects. Each document is sent individually to Azure OpenAI and analyzed independently, without context from previously reviewed documents. Because these system-generated summaries are based solely on the document's content and not prompt criteria or case context, they remain objective and consistent across projects.

Notes:

- Document topics and summaries are generated automatically when you run an analysis type, except Custom Analysis, which does not produce them. They cannot be generated independently of running an analysis project.
- Document summaries do not affect the number of aiR Units used. For more information on aiR Units and billing, see [aiR Units in Relativity aiR standard workspaces](#) on the Relativity documentation site and [aiR Units, Thresholds, and Included Products](#) on the Community (valid login credentials required).
- aiR Review document summaries differ from witness summaries in Case Strategy, which support fact extraction and case-building workflows. When generating facts, Case Strategy summarizes a document based on its relevance to the case prompt. See Case Strategy documentation for more information on Case Strategy summaries.

1.7 aiR units and consumption

For billing purposes, what an aiR Unit represents depends on the analysis type:

- **Relevance, Relevance & Key Documents, Issues, and Confidential Business Information**— an aiR Unit is the extracted text of a single document. Each document consumes one aiR Unit every time it runs through analysis, regardless of whether that document ran before.
- **Custom Analyses**— aiR Units are consumed per insight, where one insight is a single field result on a single document. Every 5 successful insights consume 1 aiR Unit, rounded down to the nearest whole aiR Unit at the job level, with a minimum of 1 aiR Unit per job. Because Custom Analyses runs a separate analysis for each insight field, a document consumes as many insights as there are fields configured. If a single field errors on a document while others succeed, only the failed insight is excluded — not the whole document.

The initial pre-run estimate may be higher than the actual aiR Units consumed because of canceled jobs, document errors, or errored insights (for Custom Analyses). To find the actual aiR Units consumed, use Cost Explorer. For more information, see Cost Explorer on the Relativity documentation site..

Document summaries do not affect the number of aiR Units consumed. For more information on aiR Units, see [aiR Units in Relativity aiR standard workspaces](#) on the Relativity documentation site and Community article [aiR Units, Thresholds, and Included Products](#) (valid login credentials required).

Caution: Customer may not consolidate documents or otherwise take steps to circumvent the aiR 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.

1.8 Language support

The underlying LLM used by aiR Review has been evaluated for use with 83 languages. While aiR Review itself has been primarily tested on English-language documents, unofficial testing with non-English datasets shows encouraging results. For more information, see see Language support in aiR products on the Relativity help site.

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

- Rigorously follow best practices for writing. For more information, see [Best practices on page 15](#).
- Iterate on the prompt criteria. For more information, see [Revising the prompt criteria on page 68](#).
- 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 reviewing analysis results, all citations remain in the language of their respective source documents. By default, rationales and considerations are provided in English. To request rationales and considerations in another language, include this sentence within the Additional Context field of your prompt criteria: *Write rationales and considerations in <Language>*. While this approach has been effective in many instances, the LLM may not consistently generate output in the specified language.

1.9 Analyzing emojis

aiR 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.

2 Installing aiR Review

aiR Review is available as a secured application from the Application Library. You must have an active aiR Review contract to use it. It 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 the ellipsis (...) button in the **Choose from Application Library** field.
6. Select the **aiR 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 [Permissions on the next page](#) for the list of permissions required for using aiR Review.

2.1 Object types available

The following object types will appear in your workspace:

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

2.2 Tabs available

The following tabs will appear:

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

- **aiR Review Projects** (workspace level)—create and manage aiR Review projects and view the project dashboard.

2.3 Configuring aiR Review job size limits

Authorized administrators can configure job size limits per user group at the instance level to help control and manage the size of jobs that users execute. For more information on setting job sizes, see [Setting aiR Review job size limits on page 100](#).

3 Permissions

aiR Review permissions are set at the instance level and workspace level. The per group user roles available at the workspace-level include: Administrators, Power Users, Basic Users, and Reviewers.

3.1 Instance-level permissions

You can only grant or revoke access to aiR Review by enabling or disabling the aiR Review feature permission directly. The aiR Review Jobs tab at the instance level displays all jobs across all workspaces, with additional columns for workspace, matter, and client per job. For details on accessing the Admin Security, see Instance Details.

The screenshot shows the 'Permissions' configuration window for 'aiR Review'. On the left is a sidebar with navigation links: Admin Security, Groups, Features (highlighted), Object Security, Tab Visibility, and Admin Operations. The main panel is divided into 'Selections' and 'Features'. Under 'Selections', 'Selected Group' is set to 'Tst Level Review Tea'. The 'Features' section has a search bar and a list of feature permissions. 'aiR Review' is checked, while 'Application Management', 'Archive and Restore (ARM)', 'Audit', 'Authentication', and 'Case Metrics' are unchecked. Each feature has a count of actions in parentheses. At the bottom right, there is a link for 'Granular permissions mapping' and three buttons: 'Review and Save', 'Revert All', and 'Close'.

3.2 Workspace-level permissions

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

Permissions for aiR Review include:

- **View aiR Review Analysis Results in Document Viewer**—select to allow group members to view aiR analysis results highlighted in the Document Viewer. They cannot access the aiR Review Projects or Jobs tabs, and cannot run analysis jobs.
- **Manage aiR Review Analysis Results in Document Viewer**—select to allow group members to view aiR analysis results highlighted in the Document Viewer, and edit document metadata. They can clear or restore aiR analysis results using the aiR Review Jobs tab. They cannot access the aiR Review Projects tabs, and cannot run analysis jobs.
- **View aiR Review Projects & Prompt Criteria**—select to allow group members to view existing aiR Review projects and prompt criteria. They can access to the aiR Review, Projects, and Jobs tabs. They cannot create, edit, or delete projects or prompt criteria.
- **Manage aiR Review Projects & Prompt Criteria**—select to allow group members to create, edit, and delete aiR Review projects and prompt criteria. They can access to the aiR Review, Projects, and Jobs tabs.
- **Run aiR Review Analysis Jobs**—select to allow group members to run aiR Review analysis jobs against existing projects and prompt criteria. They can access to the aiR Review, Projects, and Jobs tabs.

For more detailed information on what is included in each permission, see Workspace-level permissions.

Granular permissions settings are provided below, as needed.

Note: The recommended method of assigning permissions is to use the Features tab.

Viewing the aiR Review Jobs tab

To view the aiR Review Jobs tab at the workspace-level, you need the following permissions. Assign these permissions under the **Workspace Details** tab within the chosen workspace.

Object Security	Tab Visibility
• aiR Review Prompt Criteria - View	• aiR Review Jobs

You can also use item-level permissions to restrict access to a specific job's aiR Review prompt criteria. For more information, see Security and permissions in the Admin guide.

Viewing the aiR Review dashboard

To view the aiR Review dashboard, you need the following permissions. You can view results only for the analysis types you have permission to view.

Object Security	Tab Visibility
• aiR Review Project - View • aiR Review Prompt Criteria - View • aiR Relevance Analysis - View • aiR Issue Analysis - View • aiR Key Analysis - View • aiR Summaries - View • aiR Confidential Business Information Analysis - View • aiR Custom Analysis - View	• aiR Review Projects

Creating and running an aiR Review project
To create an aiR Review project and run the analysis, you need the following permissions.

Notes:

- To prevent users from running jobs, withhold the **Run aiR Review Job** permission under Other Settings.
- You can run the job without permissions for the analysis types, but you will not be able to see the results.

Object Security	Tab Visibility	Other Settings
• aiR Review Project - View, Edit, Add • aiR Review Prompt Criteria - View, Edit, Add • aiR Relevance Analysis - View • aiR Issue Analysis - View • aiR Key Analysis - View • aiR Summaries - View • aiR Confidential Business Information Analysis - View • aiR Custom Analysis - View	• aiR Review Projects	• Run aiR Review Job

Editing and running an existing aiR Review project
To edit an existing aiR Review project and run the analysis, you need the following permissions.

Notes:

- To prevent users from running jobs, withhold the **Run aiR Review Job** permission under Other Settings.
- You can run the job without permissions for the analysis types, but you will not be able to see the results.

Object Security	Tab Visibility	Other Settings
• aiR Review Project - View, Edit • aiR Review Prompt Criteria - View, Edit	• aiR Review Projects	• Run aiR Review Job

Object Security	Tab Visibility	Other Settings
• aiR Relevance Analysis - View • aiR Issue Analysis - View • aiR Key Analysis - View • aiR Summaries - View • aiR Confidential Business Information Analysis - View • aiR Custom Analysis - View		

Viewing highlights in the Viewer

To see aiR Review results highlighted in the Viewer, you need the following permissions. You will see highlighting only for the analysis types you have permission to view. If you are not granted any of these permissions, you will not see the aiR Review Analysis icon in the Viewer.

Object Security
• aiR Relevance Analysis - View • aiR Issue Analysis - View • aiR Key Analysis - View • aiR Summaries - View • aiR Confidential Business Information Analysis - View • aiR Custom Analysis - View

Clearing and restoring job results

To clear or restore job results using the aiR Review Jobs tab, you need the following permissions. 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 [Managing aiR Review jobs on page 99](#).

Object Security	Tab Visibility
• Document - View, Edit • aiR Relevance Analysis - View, Edit • aiR Issue Analysis - View, Edit • aiR Key Analysis - View, Edit • aiR Confidential Business Information Analysis - View • aiR Custom Analysis - View	• aiR Review Jobs

4 Best practices

Refer to the tips and recommendations below for prompt criteria and sample document sets to effectively use aiR Review for all analysis types.

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 pertinent to the legal issues at hand.

When writing prompt criteria, use natural language to describe why particular types of documents should be considered. 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 some sections, such as a list of names or email addresses, but paragraphs might make sense in others. Clear formatting, layout, and spelling of information aids aiR's logical processing.
- **Is it essential?**—ask yourself is this essential to understanding and locating pertinent document or matters; will the criteria affect the results.
- **Avoid legal jargon or explanations**—for example, do not use legal jargon, such as "including but not limited to" and "any and all." Do not include explanations of the law or other common descriptions that could be found through a web search. aiR Review already understands them.
- **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 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 key person in the case named Dave may commonly go by the nickname Skip, 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 for additional guidance entering criteria in each field.

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

- [aiR Review Prompt Writing Best Practices](#)—downloadable .pdf of writing guidelines
- [aiR Review example project](#)—detailed example of adapting a review protocol into prompt criteria
- [AI Advantage: Aiming for Prompt Perfection?](#)—on-demand webinar discussing prompt criteria creation

4.2 General prompt criteria iteration workflow

We recommend the following workflow for developing and iterating on prompt criteria for all analysis types:

1. For your first analysis, run the prompt criteria on a saved search of 50-100 test documents that contains a diverse range of 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 produces results that reasonably align with coding decisions 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.
6. Finally, run the prompt criteria on a larger set of documents.

aiR 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 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.

4.3 Tips for creating sample documents set

Before setting up the aiR Review project, create a saved search that contains a small representative set of documents for testing the aiR Review prompt criteria during the initial development phase. This sample set is intended to be similar to those used when training a review team. See [Creating or editing a saved search in the Searching guide](#) for details about saved searches.

Here are some guidelines:

- **Deduplication**—Reduce the size of the population. For example, remove duplicate information, documents with no text, too much text, low extracted text, and password protected files. Possible tools to use: Email threading to focus on only inclusive emails, Textual near duplicate (TND), Hash duplicates of attachments and loose records, run extracted text size script, and remove documents with no text or too much text.
- **Diversity**—Ensure diversity in the sample so that it covers all categories of interest to provide representation across what you're likely to see. Possible tools to use: Clustering to pull documents from a variety of clusters, Stratified sampling script, Targeted searching, and Random sampling of rich populations.
- **Richness**—aim for a rich sample containing a good representation of matters. Recommendation is 50-75% pertinent documents and 25-50% not. Possible tools to use: Date filters, Search terms report (STR), and Key custodians to focus on, and already identified documents.
- **Size**—include 50-100 documents to start that are a mix of applicable, not applicable, and challenging documents, then expand the number gradually as needed with the prompt iterations.
- **Key features**—make sure the documents highlight all the key features of your prompt criteria.
- **Human reviews**—have human reviewers code the documents in advance.

For more information about choosing documents for the sample, see [Selecting a Prompt Criteria Iteration Sample for aiR Review](#) on the Community site.

4.4 Best practices for writing Custom insight prompts

Use this guide as you create Custom insights. Small changes in how you write your insight prompts can dramatically improve quality and consistency. See [Custom analyses on page 27](#) to add Custom insights to a project.

How to use this guide:

- **Workflow:** Start with the [General tips below](#), then explore [Extraction tips on page 21](#), [Classification tips on page 23](#), and [Narrative/Insight tips on page 24](#).
- **Example User Goal:** The outcome you might try to achieve with a custom analysis field.
- **Don't / Do example table:** Shows how vague versus explicit field descriptions affect the system's results.

4.4.1 General tips

General tips apply to all analysis types: extraction, classification, and narrative alike. Start here before diving into type-specific guidance.

4.4.1.1 Ask for one thing at a time

Focus on one detail at a time, such as one question, one field, or one answer format. Bundling multiple tasks into a single open-ended question leads to inconsistent structure. Separate each field or use a clearly defined schema.

EXAMPLE USER GOAL: Determine whether a document contains handwritten text, with a concise Yes or No answer you can use for filtering or routing.

Click to view Do and Don't examples.

Don't	Do
EXAMPLE PROMPT: Does this document have any handwriting, annotations, stamps, sticky notes, margin notes, correction marks, or other manual markings? Describe what you find. EXAMPLE RESPONSE: The document contains a small handwritten note in the left margin that reads "See clause 4.2 — RU." It also has what appears to be a stamped "RECEIVED" date in the upper right corner. I did not find any sticky notes, correction tape, or other manual markings. <i>By asking about six different things at once and requesting description, you get a multi-paragraph response that's hard to parse. The system conflates handwriting with stamps and annotations because the prompt bundles them together.</i>	EXAMPLE PROMPT: Does this document contain any handwritten text? Answer only Yes or No. EXAMPLE RESPONSE: Yes <i>One question, one answer format. The response is a clean Yes or No you can act on. If you also need stamp detection, create a separate analysis for it.</i>

4.4.1.2 Describe how to handle edge cases

Documents are messy. Tell the system what to do when data is missing, partial, or ambiguous.

EXAMPLE USER GOAL: Extract the primary document date in YYYY-MM-DD format, correctly handling documents where dates are missing, partial, or appear multiple times.

Click to view Do and Don't examples.

Don't	Do

Don't	Do
EXAMPLE PROMPT: Extract the date from this document. EXAMPLE RESPONSE (NO DATE PRESENT): I'm sorry, I couldn't find a specific date in the document. However, based on the context, it may have been written around mid-2024. <i>Without edge-case rules, the system guesses when data is missing. It might hallucinate a date from context or return inconsistent formats across documents. You can't rely on the output being consistent.</i>	EXAMPLE PROMPT: Extract the primary document date, formatted as YYYY-MM-DD. Follow these rules: - If multiple dates appear, return the one that represents the document's execution, issue, or effective date. - If no date is found, respond with exactly: N/A - If the date is partial (e.g. only a month and year), return what is available using XX for unknown parts, e.g. 2024-10-XX. EXAMPLE RESPONSE (NO DATE PRESENT): N/A Every edge case has a defined behavior. The system knows to return N/A (not apologize), pick the execution date (not a body reference), and use XX for unknowns. You get predictable output regardless of document quality.

4.4.1.3 Give brief positive and negative examples

When asking the system to make a judgment call, like classifying a field, a couple of short examples go a long way toward defining the expected behavior.

EXAMPLE USER GOAL: Detect whether a document contains a formal signature (handwritten, /s/ conformed, or digital), distinguishing it from email sign-offs and typed names.

Click to view Do and Don't examples.

Don't	Do
EXAMPLE PROMPT: Does this document contain a signature? Answer Yes or No. EXAMPLE RESPONSE: Yes <i>Without boundary examples, "signature" is ambiguous. The system treats email sign-offs and typed names as signatures, producing false positives that defeat the purpose of the detection.</i>	EXAMPLE PROMPT: Does this document contain a formal signature? Answer only "Yes" or "No". A formal signature is a handwritten signature, a /s/ (conformed signature) line, a digital signature block, or a witnessed mark. Examples: - /s/ John Smith → Yes - Handwritten cursive name above a signature line → Yes These are NOT signatures: - "Best regards, John Smith" at the end of an email → No - A typed name at the bottom of a report (e.g. "Prepared by: Jane Doe") → No EXAMPLE RESPONSE: No

4.4.1.4 Specify the output format

Tell the system exactly how you want the answer formatted, such as a Yes/No answer, a date, or a category count. Without this, responses will vary in style and be harder to use consistently.

Note: Date output format will always display as YYYY-MM-DD for Custom Analyses.

EXAMPLE USER GOAL: Extract vendor, document number, date, total, and currency from financial documents into a consistent, structured format.
Click to view Do and Don't examples.

Don't	Do
EXAMPLE PROMPT: What are the financial details of this document? EXAMPLE RESPONSE: This invoice was issued by Bright Horizon Consulting LLC on January 15, 2025. The total amount due is \$17,320.00 for consulting services including strategic planning and executive coaching. <i>This prompt doesn't name the fields you need or specify a format. The system returns a narrative. It mentions the vendor and total, but not in a consistent structure. Run to run, the format will vary.</i>	EXAMPLE PROMPT: Extract the following from this financial document and respond in this format: Vendor: [company that issued the document] Document Number: [invoice number, PO number, or other identifier] Total: [total amount as a number, no currency symbol] Currency: [three-letter code like USD] If a field is not present, write "N/A". EXAMPLE RESPONSE: No Vendor: Bright Horizon Consulting LLC Document Number: BHC-2025-0187 Total: 17320 Currency: USD <i>Every field is named, the format is defined, and missing-value handling is explicit. The system knows exactly what you need and returns it in a consistent structure every time.</i>

Format spectrum: You can make your format as structured as you need. The example above uses labeled fields, simple and readable. For more technical workflows you can go further:

Labeled fields → Vendor: Acme Corp

Delimited → Acme Corp; BHC-2025-0187; 2025-01-15; 17320; USD

JSON → {"vendor": "Acme Corp", "total": 17320, ...}

Pick whatever fits your workflow. The key principle is the same—tell the system exactly what format you want.

4.4.1.5 Assign a role for context

Assigning a domain role when the question needs subject matter judgment. This can improve how the system interprets domain-specific language and what level of precision it applies.

EXAMPLE USER GOAL: Classify internal communications into HR conduct categories (Hostile, Policy Concern, or Clean) from a compliance analyst's perspective, catching subtle policy violations that a general reader might miss. Click to view Do and Don't examples.

Don't	Do
EXAMPLE PROMPT: Is there anything problematic in this document? EXAMPLE RESPONSE (EMAIL WITH VEILED BIAS): There's nothing overtly problematic, though the tone could be interpreted as slightly cynical. The email discusses hiring preferences but doesn't contain anything explicitly inappropriate. <i>Without a role or taxonomy, the system applies a general-audience lens. Subtle policy concerns like favoritism or hiring bias get dismissed as "slightly cynical." An HR compliance analyst would flag this immediately.</i>	EXAMPLE PROMPT: You are an HR compliance analyst reviewing internal communications for workplace conduct issues. Classify this document into exactly one of these categories: - Hostile: aggressive, demeaning, or threatening language directed at an individual - Policy Concern: potential violation of company policies (hiring, discrimination, retaliation) without overtly hostile language - Clean: no workplace conduct issues detected Respond with only the category name. EXAMPLE RESPONSE (EMAIL WITH VEILED BIAS):: Policy Concern <i>The role ("HR compliance analyst") activates domain expertise. The taxonomy gives the system three precise labels instead of a vague yes/no. Veiled references to favoritism now correctly map to "Policy Concern."</i>

4.4.2 Extraction tips

Extraction tips help when your analysis pulls specific values out of documents: dates, amounts, names, identifiers. The challenge with extraction is getting consistent, structured output across diverse document formats.

4.4.2.1 Normalize extracted values

Documents express the same value in different formats: "\$1,234.56" vs "1.234,56 EUR" vs "one thousand two hundred thirty-four dollars." Tell the system exactly how to normalize.

EXAMPLE USER GOAL: Extract the final total amount as a plain number with two decimal places and the three-letter currency code, with no symbols or locale-specific formatting, ready for direct ingestion into a database or spreadsheet. Click to view Do and Don't examples.

Don't	Do
EXAMPLE PROMPT: What is the total amount on this document? EXAMPLE RESPONSE (US INVOICE): The total amount is \$35,984.90. EXAMPLE RESPONSE (EU INVOICE): The total is €51.480,00. <i>The system mirrors the document's formatting: dollars with commas, euros with periods for thousands. You'd</i>	EXAMPLE PROMPT: Extract the final total amount from this document. Follow these formatting rules exactly: - Return only a plain number with exactly two decimal places (e.g. 1234.56) - Do not include currency symbols (\$, €, £), commas, or text - Use a period as the decimal separator regardless of the document's locale - On a new line, return the three-letter currency code (e.g. USD, EUR, GBP) EXAMPLE RESPONSE (US INVOICE):

Don't	Do
<i>need to account for every locale variation, and currency is sometimes missing entirely.</i>	35984.90 USD EXAMPLE RESPONSE (EU INVOICE): 51480.00 EUR <i>Explicit formatting rules produce identical output structure regardless of the source document's locale. Every value uses period decimals, no symbols, and always includes the currency code on a separate line.</i>

4.4.2.2 Handle multi-value fields

A document may contain multiple phone numbers, addresses, or dates. Tell the system which one to return, or how to return all of them in a consistent structure.

EXAMPLE USER GOAL: Extract all instances of a repeated field (such as, phone numbers) in a uniform format, or select the right one when only one is needed.

Click to view Do and Don't examples.

Don't	Do
EXAMPLE PROMPT: Extract the phone number from this document. EXAMPLE RESPONSE (3 NUMBERS IN DOC): The document contains several phone numbers: (212) 555-0147 belongs to the sender, (312) 555-0298 appears to be the office line, and 1-800-555-0199 is the customer service number.	EXAMPLE PROMPT: Extract all phone numbers from this document. Return each number on a separate line in the format +1-XXX-XXX-XXXX. If no phone numbers are found, respond with: N/A EXAMPLE RESPONSE (3 NUMBERS IN DOC): +1-212-555-0147 +1-312-555-0298 +1-800-555-0199

4.4.2.3 Anchor to document structure

When you know your documents share a consistent layout (such as, you are processing a batch of purchase orders or a standard form), it can be worth trying to tell the system where in the document to look. This is most useful when analyzing many similar documents whose structure is predictable. Positional anchoring can sometimes reduce ambiguity when the same value appears in headers, footers, and body text.

EXAMPLE USER GOAL: Extract a specific value from a known location in the document (such as, the approval block at the bottom), not from a mention elsewhere in the text.

Click to view Do and Don't examples.

Don't	Do
EXAMPLE PROMPT: Who approved this purchase order?	EXAMPLE PROMPT: Extract the names of the people who approved this pur-

Don't	Do
EXAMPLE RESPONSE: The purchase order appears to have been approved by several people. John Miller is mentioned in the header as the requesting manager, and there are also references to Sarah Kim in the vendor notes section.	chase order. Look specifically in the approval block at the bottom of the document (after "Authorized By" or similar closing section). Return each approver's name and title on a separate line. EXAMPLE RESPONSE (3 NUMBERS IN DOC): John Miller, Procurement Manager, Acme Corp Sarah Kim, VP of Finance, Acme Corp

4.4.3 Classification tips

Classification tips help when your analysis assigns a label or category to documents. The challenge is getting consistent labels from a controlled vocabulary that can be used for filtering, routing, and reporting at scale.

4.4.3.1 Define all valid categories up front

Provide the complete list of valid labels. Open-ended classification ("what type is this?") produces inconsistent taxonomies like "business agreement" vs "service arrangement" vs "vendor document" for the same document.

EXAMPLE USER GOAL: Classify documents into exactly one label from a predefined 6-category taxonomy (Agreement, Correspondence, Legal Filing, Financial Record, Report, Other) for document routing and reporting. Click to view Do and Don't examples.

Don't	Do
EXAMPLE PROMPT: What type of document is this? EXAMPLE RESPONSE (SAME DOC TYPE): Run 1: "This is a vendor agreement" Run 2: "Service arrangement" Run 3: "This appears to be a binding agreement between two companies" <i>Without a taxonomy, the system invents labels on each run. You can't filter or route on these because the same document gets different labels every time. You need exact, repeatable categories.</i>	EXAMPLE PROMPT: Classify this document into exactly one of these categories. Respond with only the category name. - Agreement: any binding arrangement, including service agreements, leases, licenses - Correspondence: emails, letters, memos, or other person-to-person communications - Legal Filing: court documents, motions, briefs, complaints, subpoenas - Financial Record: invoices, receipts, bank statements, tax filings - Report: incident reports, audit reports, inspection reports - Other: anything that does not fit the above categories EXAMPLE RESPONSE: Agreement <i>A closed taxonomy with definitions eliminates label drift. The system picks from exactly 6 options, and each option has boundary examples so ambiguous documents (like an email about a project) get classified consistently.</i>

4.4.3.2 Include a catch-all category

Without "Other" or "None of the above," the system force-fits every document into the closest label, even when none apply. A catch-all can help reduce misclassifications.

EXAMPLE USER GOAL: Correctly classify documents that do not fit your taxonomy into a catch-all category, rather than forcing them into the wrong bucket.
Click to view Do and Don't examples.

Don't	Do
EXAMPLE PROMPT: Classify this document: Invoice, Agreement, or Legal Filing. EXAMPLE RESPONSE (HR MANUAL POLICY): Agreement	EXAMPLE PROMPT: Classify this document into one of these categories: Invoice, Agreement, Legal Filing, or Other. If the document does not clearly fit Invoice, Agreement, or Legal Filing, use "Other." EXAMPLE RESPONSE (HR MANUAL POLICY): Other

4.4.3.3 Define decision boundaries for ambiguous cases

Do not just list labels. Describe what distinguishes borderline cases. An email about a project is Correspondence, not a Report. A memo containing financial figures is not necessarily a Financial Record.

EXAMPLE USER GOAL: Correctly classify borderline documents where content from multiple categories overlaps (such as, an email discussing a budget, a memo with financial figures).
Click to view Do and Don't examples.

Don't	Do
EXAMPLE PROMPT: Is this a financial document? Answer Yes or No. EXAMPLE RESPONSE (EMAIL WITH BUDGET NUMBERS): Yes	EXAMPLE PROMPT: Is this document a Financial Record? Answer only "Yes" or "No". A Financial Record is a document whose primary purpose is recording a financial transaction or reporting financial position. Examples: - Invoice, receipt, bank statement, tax return → Yes - Email discussing a budget or mentioning dollar amounts → No - Inspection report with cost estimates → No (the primary purpose is the inspection findings, not the financial transaction) EXAMPLE RESPONSE (EMAIL WITH BUDGET NUMBERS): No

4.4.4 Narrative/Insight tips

Narrative / Insight tips help when your analysis generates summaries, descriptions, or open-ended assessments. The challenge is getting consistent length, appropriate depth, and grounded claims across a large document population.

4.4.4.1 Constrain output length and focus

Without bounds, summaries range from one sentence to five paragraphs unpredictably. Set explicit length and tell the system what to prioritize.

EXAMPLE USER GOAL: Produce a 2-3 sentence summary that includes specific names, dates, and numbers, focused on the key finding and next steps, not background or methodology.
Click to view Do and Don't examples.

Don't	Do
EXAMPLE PROMPT: Summarize this document. EXAMPLE RESPONSE: This compliance audit report was prepared by Stonebridge Compliance Partners for Pacific Northwest Credit Union covering the period from January 1 to December 31, 2024. The audit reviewed 450 consumer loan files and 120 mortgage applications across six branches, focusing on compliance with ECOA, TILA, and the Fair Housing Act. Four key findings were identified. Finding #1 (High Risk) involved ECOA adverse action notice failures in 4% of files, disproportionately affecting younger applicants... [continues for 3 more paragraphs] Without length or focus constraints, the system dumps everything: methodology, background, and every finding. A reviewer scanning 500 documents needs the essential takeaway in seconds, not a five-paragraph essay.	EXAMPLE PROMPT: Summarize this document in exactly 2-3 sentences. Focus on: 1. What the document is and who created it 2. The most important finding, decision, or event 3. Any required next steps or actions Do not include background or methodology. Be specific: include names, dates, and numbers. EXAMPLE RESPONSE: Stonebridge Compliance Partners' 2024 audit of Pacific Northwest Credit Union found two high-risk issues: ECOA adverse action notice failures in 4% of loan files and mortgage pricing disparities of 38 basis points in majority-minority census tracts. The overall risk rating is Elevated, and the auditor recommends engaging an independent fair lending consultant and implementing automated notice tracking. The sentence count, focus areas (who/what/next steps), and specificity requirement ("include names, dates, and numbers") produce a dense, actionable summary every time.

4.4.4.2 Require evidence and citations

Tell the system to reference specific passages or facts from the document. This prevents hallucination, aids reviewer verification, and surfaces which parts of the document support the conclusion.

EXAMPLE USER GOAL: Get claims backed by verbatim quotes from the document, so reviewers can verify conclusions without re-reading the full text.
Click to view Do and Don't examples.

Don't	Do
EXAMPLE PROMPT: Does this document contain evidence of bid rigging? EXAMPLE RESPONSE:	EXAMPLE PROMPT: Does this document contain evidence of bid rigging or market allocation? Answer "Yes" or "No", then on a new line quote the most relevant passage (verbatim, in quotation marks) from the document that supports your answer. If no relevant passage exists, write "No supporting passage found."

Don't	Do
Yes, the document contains some indicators that could suggest coordinated bidding behavior between the parties.	EXAMPLE RESPONSE: "I will focus our bids on the Illinois and Wisconsin RFPs, and you take the lead on Indiana and Ohio. That way we're not undercutting each other on price."

4.4.4.3 Separate observations from inferences

In investigation and compliance contexts, it's critical to distinguish what the document explicitly states from what can be inferred or concluded. Ask the system to separate the two.

EXAMPLE USER GOAL: Get a structured analysis that clearly separates stated facts (with quotes) from potential inferences, so reviewers know what is documented versus interpreted.

Click to view Do and Don't examples.

Don't	Do
EXAMPLE PROMPT: What are the key takeaways from this incident investigation report? EXAMPLE RESPONSE: The maintenance team was clearly trying to cover up the equipment failure and delayed reporting it to management, likely because they knew the preventive maintenance schedule had been skipped.	EXAMPLE PROMPT: Analyze this incident investigation report. Provide two clearly labeled sections: STATED FACTS: What the report explicitly documents (with direct quotes where possible). POTENTIAL INFERENCES: What a reviewer might reasonably conclude from the findings, clearly marked as inference rather than fact. EXAMPLE RESPONSE: STATED FACTS: The report documents that elevated vibration readings were first recorded on August 12, 2023 but were not escalated to plant management until October 3. The shift supervisor described the initial readings as "within acceptable tolerance." A maintenance log entry from Sept 18 noted "Martinez aware of continued vibration anomaly." POTENTIAL INFERENCES: The seven-week gap between initial detection and management escalation, combined with the Sept 18 log entry, could suggest the shift team had more awareness of the severity than the timeline of formal reports implies.

5 Common analysis workflows

Here a few common analysis workflows for aiR Review:

- First pass review for responsive docs—this workflow uses Relevance analysis.
- Hot documents identification—this workflow uses Relevance & Key analysis.
- Legal issues tagging—this workflow uses Issues analysis.

- Confidentiality review—this workflow uses Confidential Business Information (CBI analysis)
- Custom insights—this workflow uses Custom analyses. See [Custom analyses below](#) for general steps.

5.1 Custom analyses

Custom analyses provide a flexible framework for creating and executing custom AI-powered insights per project within Relativity. Each insight acts as an independent instruction, enabling customizable results across all your legal data intelligence use cases.

Instead of relying solely on the predefined analysis types in aiR Review (such as Relevance and Key Documents), you can create up to five Custom analysis insights per project, each acting as independent analyses for questions, extractions, classifications, or evaluative instructions applied to every eligible document. Reviewers receive per-document outputs, dependent on their unique Custom analyses prompt.

Custom analyses support both Text and Vision analysis models:

Custom Analyses Model	Description
Text	• Analyzes the extracted text of each document. • Best for emails, contracts, reports, and most standard documents. • Usually the starting point for most review matters.
Vision	• Analyzes the document as a native image file; not extracted text. • Best for scanned materials, photos, screen shots, and handwritten documents, for example. • Can read text, identify objects, locate handwriting, and detect logos on this image, for example. • Surfaces information that does not exist in extracted text.

Each Custom insight analyses produces one text-based output, which may include:

- Summaries
- Yes/No responses
- Classification labels
- Short explanations
- Identification

Each Custom insight consists of a unique title and prompt instructions. The prompt instructions tell aiR Review how to analyze the documents and how to present the results. Examples include summarizing information, extracting items, identifying attributes, or applying classifications. These instructions form the core logic for what the aiR system model returns for each document in the Analysis Results. Each Custom insight will be listed in its own column in the Analysis Results grid and document viewer. For guidance on writing effective custom insight prompts, see [Best practices for writing Custom insight prompts on page 18](#).

5.1.1 Considerations

Review the following considerations for using Custom analyses in aiR Review:

- Unlike other analysis types in aiR Review, Custom analysis results are not available on the Relativity Document Object until you apply them. They are initially only available in your aiR Review project and the Viewer. To publish them to the Relativity Document Object fields, run the Apply step (see [Applying the results and publishing to the Document Object on page 35](#)).
- After the Apply process is complete, the Custom analysis results are available for export. See [Exporting Custom analyses results on page 82](#) for more information.
- Validating prompt criteria in Review Center is not supported for Custom analyses. See [Prompt Criteria validation in Review Center on page 92](#) for more information on validation.
- Insight Suggestions have no effect on how your analysis runs or its cost. They are a starting point you can use, edit, or ignore. See [Using Insight Suggestions on page 31](#) for more information.

5.1.2 How Custom analyses insights work

A Custom analysis lets you define one or more insights, each with its own prompt criteria. Each insight runs independently against a set of documents. Understanding how insights are processed will help you write effective prompts and plan your jobs.

5.1.2.1 How insights run

Each Custom insight runs as its own independent job. This has two practical consequences:

1. **Insights do not share context.** One insight cannot see another insight's prompt, criteria, or results. Each insight evaluates a document using only its own prompt criteria and the document text.
2. **Each insight counts toward your concurrent job limit.** A Custom analysis is not a single job. See [Job limits on the next page](#) below.

5.1.2.2 Writing prompts for independent insights

Because insights run separately, each prompt must contain everything needed to evaluate a document on its own. Do not write a prompt that refers to another insight, its criteria, or the documents it returned. If you need to apply the same definition or criteria across insights, repeat it in each insight's prompt rather than referring to another insight.

Example: Finding documents that meet multiple criteria

Avoid: Referencing another insight

Do not write a prompt that depends on the results of another insight. For example, the following approach does not work:

Insight	Prompt
Insight A	Identify documents discussing the merger.
Insight B	Of the documents identified in Insight A, flag those that mention pricing.

Insight B cannot access the results of Insight A. Instead, it evaluates every document using a reference it cannot resolve, which can produce unreliable results.

Use one of these approaches

To find documents that discuss both the merger and pricing, make each insight self-contained. Choose the approach that best supports how you want to review the results.

- Option 1: Create independent insights

Insight	Prompt
Insight A	Identify documents discussing the merger.
Insight B	Identify documents that mention pricing.

After the analysis completes, review the results together (for example, by filtering on both insights) to find documents that meet both criteria.

- Option 2: Combine the criteria in one insight

Insight	Prompt
Insight A	Identify documents that discuss the merger and also mention pricing.

5.1.2.3 Job limits

Each insight in a Custom analyses counts as a separate job toward your concurrent job limit. For example, a Custom analysis with five insights submits five jobs, not one.

Plan the number of insights accordingly to avoid exceeding the limit of 20 concurrent jobs. See [Job capacity, size limitations, and speed](#) for information on limits.

5.1.3 Custom analyses workflow

The workflow for Custom analyses is similar to the other analysis types:

1. [Setting up a project below](#)
2. [Adding Custom insights and prompt criteria](#)
3. [Running a Custom analyses job on page 33](#)
4. [Viewing & iterating on the analysis results on page 34](#)
5. [Applying the results and publishing to the Document Object on page 35](#)

5.1.4 Setting up a project

Begin by setting up the aiR Review project as described below. For additional information on project set up, see [Setting up the project on page 50](#).

1. Navigate to the **aiR Review Project** tab.
2. Click **New aiR Review Project**.

3. Fill out the following fields on the Setup Project modal:

The following settings cannot be changed after the Project has been created.

Project Name *

Description

Data Source *

Analysis Type *

Project Prompt Criteria * ☒ Start Blank ☐ Copy Existing

Prompt Criteria Name *

Project Use Case *

- **Project Name**—enter a name for the project.
- **Description**—enter a useful short description of the project.
- **Data source**—select the saved search that holds your document sample set. Refer to [Choosing the data source](#) for more information.
- **Analysis Type**—select **Custom**.
- **Project Prompt Criteria**—defaults to **Start blank** to write new prompt criteria from scratch.
- **Prompt Criteria Name**—skip this field for Custom Analysis Type.
- **Custom Analysis Model**—choose the analysis model desired:
 - **Text**—run analysis on the extracted text of documents.
 - **Vision**—run analysis on native image file types of JPEG, GIF, or PNG.
- **Project Use Case**—choose the option that best describes the purpose of the project. If none of the options describe the project, choose Other and 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. The Project Use Case field is used for reporting and management purposes and does not affect how the project runs.

4. Click **Create Project**.

5.1.5 Adding Custom insights and prompt criteria

On the Custom panel, you can create and edit Custom insights and prompts. You can also delete ones that are no longer needed.

1. Add up to five Custom insights using the methods below:

- **Manually enter Custom insights and prompts:** Enter a unique **Title** (up to 30 characters) for the insight. The title displays in the Analysis Results table, Custom Projects card in the Viewer, and the published output. Next, manually enter clear and descriptive instructions in the **prompt text box** (up to 1000 characters) that will guide aiR in analyzing documents. See [Best practices for writing Custom insight prompts on page 18](#) for more information.
- **Use preset Insight Suggestions:** Click the **+Add** button next to a Suggested Insight to automatically add it. Afterward, if needed, edit the suggestion's prompt to fit your matter. Editing an Insight Suggestion is the same as editing an insight you manually created. See [Using Insight Suggestions on the next page](#)

for more details.

2. Select the **Field Type** output format for this insight's result: Fixed Length Text (default), Single Choice, Date, or Whole Number. This selection determines how the result is validated, how it displays in the Analysis Results table and Viewer, and which native Relativity field type is published on the Document object. Field Type cannot be changed after the first Apply & Publish project set is run for the insight.

Note: Date output format will always display as YYYY-MM-DD for Custom Analyses.

3. When Field Type is set to *Single Choice*, define 2–5 **Choices** for aiR to select from. Click **+ Choice** to add another choice to the list. Add choices in the desired order as they cannot be reordered later. For each choice:
 - a. Enter a unique, case-insensitive name in the **Name Choice [#]** field (up to 30 characters).
 - b. Enter criteria in the **Describe Choice [#]** text box (up to 300 characters) explaining when a document matches that choice.

If a result does not conform to the selected Field Type, the cell is marked as an error in the Analysis Results table. For example, if the result is a non-numeric response for a Whole Number field or is a value outside the defined choices for a Single Choice field. See [Handling document errors in Custom analyses on page 37](#) for more information.

4. To add more Custom Insights, click **+ Insight** to add one manually, or click **+ Add** next to a suggested insight. Repeat the previous step for each additional insight.
5. To delete a Custom insight from the list, click the **vertical ellipses** next to the desired title and select **Delete**.
6. Click **Save** when finished adding Custom insights.

Custom insight information can be edited or deleted prior to running the analysis.

5.1.5.1 Using Insight Suggestions

Insight Suggestions are preset prompt templates that appear at the top of the prompt configuration panel. They help you get started quickly by providing automatic insights and related instructions that you can add to the panel.

Note: The suggestions have no effect on how your analysis runs or its cost. They are a starting point you can use, edit, or ignore.

The suggestions displayed correspond the analysis model you selected when setting up the project (Text or Vision):

- Text projects show text-based suggestions. For example, summarizing a document or extracting key facts, people, dates, and amounts.
- Vision projects show vision-based suggestions. For example, summarizing an image, identifying key objects, or extracting and locating visible text.

These Insight Suggestions for Text analysis and Vision (images) analysis cover some common requested needs.

Title	Prompt description
Summarize	Summarize this document in exactly 2-3 sentences. Focus on: 1. What the document is and who created it 2. The most important finding, decision, or event 3. Any required next steps or actions Do not include background or methodology. Be specific: include names, dates, and numbers.
Categorize document type	Classify this document into exactly one of these categories. Respond with only the category name. • Agreement: any binding arrangement, including service agreements, leases, licenses • Correspondence: emails, letters, memos, or other person-to-person communications • Legal Filing: court documents, motions, briefs, complaints, subpoenas • Financial Record: invoices, receipts, bank statements, tax filings • Report: incident reports, audit reports, inspection reports • Other: anything that does not fit the above categories
Identify themes	Identify the top 3-5 high-level themes in this document. Focus on themes that reflect the document's core ideas, not minor details. Output requirements: • 1-2 words per theme • One theme per line • No explanations, bullets, or numbering If no clear themes are present, return exactly: N/A
Identify individuals	Identify up to 5 real individuals (people) named in this document. Rank individuals (people) by: 1. Number of mentions

Title	Prompt description
	2. Earliest appearance in the document Use the earliest appearance as a tiebreaker. Include only real people, not organizations, departments, email domains, placeholder/test values, or redacted text. Return one full name per line. If none appear, return exactly: N/A

Title	Prompt description
Total amount	If explicitly marked in the document, what is the total financial amount? Otherwise, respond with exactly: N/A
Dates	What are the dates mentioned in this document? List them in order of appearance, using YYYY-MM-DD format. Do not include partial dates such as "January 1985" or "06/12"; only include dates where day, month and year are all shown together. If no dates are mentioned, respond with exactly: N/A
Logos	Does this document contain any logos? Answer Yes or No. Text-only company names are not considered logos.
Signature	Does this document contain a signature? Answer Yes or No.

5.1.6 Running a Custom analyses job

After setting up all the custom insights, use the steps below to run the analysis job. For additional information, see [Running the analysis on page 70](#).

Note: This step only generates results inside your aiR Review project and the Viewer. Your Custom analysis fields will not appear on the Document Object, and cannot be exported, until you separately create and run an Apply & Publish project set. See [Applying the results and publishing to the Document Object on page 35](#).

1. Click **Analyze [X] documents**.
2. Review the confirmation summary modal.

aiR Review

Review and confirm the following estimates before starting this analysis.

	Total Docs	180
	Total Insights	360
	Estimated max aiR units ?	72*
	aiR Version	2026-08-11-a4a88b

*The performance of Relativity aiR Review has been tested on primarily English language documents.
 Use of aiR Review with non-English language documents may result in unexpected or unverifiable results.*

Start Analysis
Cancel

- **Total Docs**—number of documents to be analyzed.
- **Total Insights**—number of insights to be analyzed.
- **Estimated max aiR units**—the estimated maximum number of aiR units calculated for this project set analysis. The final count is determined only when the job completes or is canceled. This estimate is based on the total number of Insights in this project set. Any insights with existing results are included in the estimate but do not count towards the final aiR unit cost. For detailed information on aiR units for Custom analyses jobs, see aiR Units in Relativity aiR standard workspaces documentation on the Relativity documentation site.
- **aiR Version**—a unique ID that identifies the underlying aiR system used to process each job. This ID does not map directly to a single large language model version used for the job. You can also find the aiR Version on the aiR Review Jobs list. See [Managing aiR Review jobs on page 99](#).

3. Click **Start Analysis**. Custom insight information cannot be edited or deleted after running the analysis.

The analysis results display in the Analysis Results grid. See [Viewing & iterating on the analysis results below](#).

Documents that cannot be analyzed will display with errors. Refer to [Handling document errors in Custom analyses on page 37](#) for more information.

Results are written to Fixed Length Text fields and will have a maximum of 3,000 characters per output.

Each Custom insight runs independently across eligible documents that meet the following criteria. Refer to [How Custom analyses insights work on page 28](#) for more details:

- Match the document requirements as stated in [Job capacity, size limitations, and speed on page 48](#).
- Do not have prior results with the same prompt.

5.1.7 Viewing & iterating on the analysis results


Once the analyses complete, each custom insight displays as individual columns in the Analysis Results table and in a dedicated Custom Projects card within the Viewer. For additional information, see [Viewing and filtering analysis results on page 77](#).

Iterate on the results as needed by refining instructions, rerunning samples, and reviewing output until you are confident to run it on larger document sets.

Unlike other aiR Review analysis types, Custom insight column results populate as they become available rather than the order in which they were added to the Custom panel. This allows results to appear live as each field completes, without waiting for all fields to finish.

Notes:

- You cannot filter Custom analyses results at this time. This applies to the Analysis Results table only. Once results are published to the Relativity Document Object (RDO) by running the Apply & Publish step (see [Applying the results and publishing to the Document Object below](#)), they can be filtered and sorted according to each insight's selected Field Type.
- Custom analyses results are initially only available in your aiR Review project and the Viewer. To publish them to the RDO, run the Apply & Publish step (see [Applying the results and publishing to the Document Object below](#)). After the apply process is complete, the results are available for export (see [Exporting Custom analyses results on page 82](#)) and where they can be used across the review workflow for filtering, sorting, and display alongside other document fields in the Document List.
- Custom analyses results do not enforce RDO document-level permissions in the same way as individual document viewing. This means users may see aggregated dashboard results from documents they cannot access directly. To manage or restrict selective document access, you can configure user permissions to control who can view or create Custom analyses projects. See [Permissions on page 12](#) for details.



To open the document in the Viewer, click on the **Control Number** link. See Viewer documentation for more information on using the Viewer.

If an error displays for an insight, see [Handling document errors in Custom analyses on page 37](#) for guidance.

If you modify Custom insights, change the data source, or create a new project set, only the updated insights will re-run on the next analyses.

For more information on analyzing and filtering results, see [Analyzing aiR Review results on page 72](#).

To export Custom analyses results, you must first apply the results, which saves the data to the Relativity Document Object. See [Applying the results and publishing to the Document Object below](#) and [Exporting Custom analyses results on page 82](#) for details on each process.

5.1.8 Applying the results and publishing to the Document Object

After confirming that the results meet your expectations, apply the prompt criteria to the document population and publish the results to the Relativity Document Object, where they can be:

- used across the review workflow for filtering, sorting, and display alongside other document fields in the Document List.
- exported (see [Exporting Custom analyses results on page 82](#)).

Since the iteration process occurs inside the project, you may rename, add, delete, or change insight fields numerous times as you are iterating on the output. The apply and publish process is separate because applying the final results on the sample or full data set ensures the final insight fields are created correctly in the workspace Document Object.


Notes:

- Unlike other analysis types in aiR Review, Custom analysis results are not available on the Document Object until you apply and publish them in this step. Prior to that step, they are initially only available in your aiR Review project and the Viewer.
- Validating prompt criteria in Review Center is not supported for Custom analyses. See [Prompt Criteria validation in Review Center on page 92](#) for more information on validation.
- See [How are aiR Units handled when reapplying and republishing the same prompt criteria on the next page](#) for information surrounding aiR units.


5.1.8.1 Running Apply & Publish

To apply the prompt criteria and publish results to the Document Object:

1. Click the **Create a new Apply & Publish project set** link.




2. Click **I want to apply my prompt criteria and publish results to the Document list, where they'll appear on each document.**



3. Select the **Data Source** from the list.

4. Click **Create Apply & Publish Set.**



The project set type changes to **Apply & Publish**. The results are published to the Document Object.

5. Click **Analyze [#] documents.**
6. Review the confirmation summary, including the estimated number of documents and insights to be processed, and add notification email addresses, if needed.
7. Click **Start Analysis.**

To export Custom analyses results after applying the prompt criteria, see [Exporting Custom analyses results below](#) for details.

5.1.8.2 How the Apply & Publish process works

When you run Apply & Publish on a project set, each insight result is written to the Document Object automatically. For every insight in your project, a field is created and populated using the naming convention **aiRCustom::[Insight Name] ([Project Name])**.

The Apply Set will generate new results for documents that meet the following criteria:

- Match the document requirements as stated in [Job capacity, size limitations, and speed on page 48](#).
- Do not have prior results with the same prompt.

Existing results on the latest prompt version will not be re-written or re-charged.

After publishing, the results are immediately filterable like any other document field.

- **One record per document:** All of a document's Custom analyses results are written to a single shared object on that document.
- **Free-form results:** Each insight publishes its own text output under its own field, reflecting the prompt you wrote.
- **Errored documents are skipped:** A document that errored on a given insight is not written for that field.

5.1.8.3 Re-applying a project

You can re-run and re-apply a project as your prompts evolve.

- **Updates in place:** Re-applying a project updates that project's results on each document rather than creating duplicates.
- **Other projects are preserved:** Re-applying one Custom analyses project does not affect the published results of any other project on the same document.

See [How are aiR Units handled when reapplying and republishing the same prompt criteria below](#) for additional information surrounding aiR units.

5.1.8.4 How are aiR Units handled when reapplying and republishing the same prompt criteria

When reapplying and republishing results for the same prompt criteria, no aiR Units are consumed as long as: (1) analysis results *already exist* for the documents, *and* (2) the prompt criteria have *not changed at all* since the initial analysis run. If the *exact same* prompt is reused across versions or even across different projects in the same workspace, the system recognizes it and displays the prior analysis result without consuming additional aiR Units. However, be aware that the prompt criteria must be an *exact match*, even an accidental extra space would count as a different prompt and, therefore, would consume aiR Units. For more information on aiR Units, see aiR Units in Relativity aiR standard workspaces documentation on the Relativity documentation site and Community article [aiR Units, Thresholds, and Included Products](#) (valid login credentials required).

5.1.9 Exporting Custom analyses results

Exporting Custom Analyses results directly from the Analysis Results tab is not currently supported. To export Custom analyses results, you must publish your results to Relativity Document Object using the Apply & Publish workflow (see [Applying the results and publishing to the Document Object on page 35](#)). After publishing, you may export your results from the Document List. See [Exporting analysis results on page 80](#) for exporting procedures.

5.1.10 Handling document errors in Custom analyses

Custom analysis errors are reported at the insight level for each document, rather than for the document as a whole. Because each insight runs independently, one insight can succeed on a document while another fails. Each insight column shows its own success or error status.

5.1.10.1 Non-responses

Sometimes, the aiR System model analyzes a document but intentionally does not return an answer. This is expected behavior, not an error. It helps ensure high-quality results when the model cannot provide a conclusive response.

A non-response can occur when:

- a document is blank
- the text or image is too faint or illegible
- the prompt instructions do not apply
- the model cannot find relevant content

In these cases, the non-response message may appear in one of the following ways:

- No relevant content found.
- This document appears to be empty and could not be analyzed.
- This document could not be read or interpreted.
- A blank (null) cell appears.

5.1.10.2 Error reference

The following table provides some of the errors you may encounter, along with information on how to resolve them:

Error	What it means	Likely resolution
Failed to parse completion	The aiR system model produced output that couldn't be read in the expected format.	Retry the document. If it persists, simplify or clarify the insight's prompt, or contact Relativity Support.
No content to analyze	The document had no extractable text for a text analysis.	Confirm the document has extracted text. For image-based documents, use a Vision project instead.
Content exceeds limits	The document's content is too large to analyze in a single pass.	Reduce the size of the document set or split large documents. See Job capacity, size limitations, and speed on page 48 for additional information.
File provides insufficient context	The document has too little content for the aiR system model to act on.	Verify it is the intended document. Near-empty files may return a limited result or none.
File is not supported	The file type provided as the data source is not supported.	Convert the file to an acceptable file format. Make sure image files are native images (JPEG, GIF, PNG) for Vision analysis. See Job capacity, size limitations, and speed on page 48 for additional information.
Uncategorized error occurred	An unexpected error occurred.	Retry the document. Also, make sure image files are native images for Vision analysis. If it persists, contact Relativity Support.

5.1.11 Common use cases

Custom analyses is most effective for per-document tasks that you can express as a clear instruction. Customers have explored use cases such as:

- Summarizing documents and surfacing key facts, people, and dates
- Extracting and structuring entities — names, dates, and monetary amounts
- Classifying or categorizing documents
- Describing and summarizing non-English text
- Vision model tasks such as, describing images, identifying objects, and extracting visible or handwritten text

5.1.12 Best practices

To get the best results:

- **Be clear and specific**—clear, well-structured prompts consistently outperform vague ones. State exactly what to return and in what form.
- **Define the output format**—if you want results returned in a specific structure (for example, a label, a yes/no, or a short list), define it. Vague asks tend to return more than you want.
- **Don't over-constrain**—piling on many conditions and caveats can sometimes reduce quality. If results degrade, simplify and iterate.
- **Use one insight per question**—each insight runs independently; keeping each to a single, focused task improves consistency.
- **Provide good inputs for Vision analysis**—higher-resolution, well-oriented images produce markedly better and more reliable results.

For additional guidance, see [Best practices for writing Custom insight prompts on page 18](#).

5.1.13 Known limitations

Custom Analysis is powerful but, like any AI feature, has limits. Knowing these helps you set expectations and review results appropriately.

- **Output format can vary**—the aiR system model may not follow the requested structure or conditional logic perfectly across every document (for example, occasionally varying formatting or dropping a requested element). Defining the output format clearly reduces this.
- **Image quality drives reliability**—low-resolution, very small, or partially obscured images are the most common source of mis-identification in Vision analysis. Text in unusual orientations or scripts may also be missed.
- **Handwriting and dense extraction**—handwriting transcription may occasionally complete or guess characters. Output is limited to about approximately one page of text. Review handwritten and hard-to-read extractions.
- **Run-to-run variation**—results can vary slightly between job runs on the same document, particularly for subjective or estimative tasks. Perform quality assurance accordingly.
- **Multiple values in one document**—when a document contains several candidates (for example, multiple dates), the aiR system model may pick one without flagging the ambiguity.
- **Mental math**—the aiR system model is not a calculator and may produce unreliable results for math-related tasks, such as adding figures or converting units.


- **Configuration limits**—maximum of five insights per project, prompt and output lengths are capped, and very large documents may exceed processing limits.

6 Navigating the aiR Review dashboard

When you select a project from the aiR 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 Project details strip

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



Features		
1	Project name	Name given to the project during set up. The Analysis Type appears underneath the name.
2	Project set type	Indicates whether the project set and criteria are Develop, Validate, or Apply.
3	Data Source Name	Name of the data source (saved search) chosen at project set up and the document count. Click the link to view the documents in the data source.
4	Version number	Version number of the prompt criteria for the set and last run or saved date. For more information, see How prompt criteria versioning works on page 68 .
5	Down-arrow	Click to move between project sets to view their statistics.
6	+ plus icon (Project set)	Click to create a new project set using a different data source for the prompt criteria, to validate the prompt criteria on a target document population, or to apply it to the larger document set. See Using project sets on page 55 for more information.
7	Analyze documents button	Click 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.
8	Feedback icon	Click to send optional feedback to the aiR Review development team.

6.2 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 Review to apply when analyzing the documents.

Matter Overview ? *

This is an investigation into a potential bribery. BigThorium, a company specializing in sustainable energy solutions, was issued a Request for Proposal (RFP) by the City of Atlantis for a municipal energy project aimed at enhancing the city's energy efficiency and reducing its carbon footprint, and ultimately won the contract. It is alleged that employees of BigThorium bribed city officials to win the contract.

People and Aliases ?

John Maxwell (Sales Representative, BigThorium): Engages in discussions about offering gifts and incentives to secure the contract. Sarah Thompson (VP of Sales, BigThorium): Coordinates the strategy and approval of incentives, including the purchase of a Rolex watch. Robert Williams (City Treasurer,

Possible tabs include:

- **Case Summary**—appears for *Relevance*, *Relevance and Key Documents*, and *Issues* 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.
- **Confidential Business Information (CBI)**—appears for the *Confidential Business Information* analysis type.
- **Custom**—appears for the *Custom* analysis type.

For information on filling out the prompt criteria tabs, see [Developing prompt criteria on page 58](#).

For information on building prompt criteria from existing case documents, like requests for production or review protocols, see [Using prompt kickstarter \(Draft with AI\) on page 64](#).

To export the prompt criteria displayed, see [Exporting prompt criteria on page 69](#).

If you want to temporarily clear space on the dashboard, click the Collapse symbol (<<) in the upper right of the prompt criteria panel. To expand the panel, click the symbol again.

6.3 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 document grid.

- **For a Relevance and Key Documents analysis:**

Responsive Conflicts 1	Key Conflicts 0
---------------------------	--------------------

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*.

- **For Issues analysis:**

All Issues Relevant 263	Disparate Treatment... Relevant 52	Segregated Condi... Relevant 11	Coercion & Underpa... Relevant 2	Retaliation Relevant 29	Theft and Criminal A Relevant 49
----------------------------	---------------------------------------	------------------------------------	-------------------------------------	----------------------------	-------------------------------------

- The **All Issues** tab displays first. It offers a comprehensive overview of predictions and scores for all issues and documents within the project. The total number of issue predictions displayed on the *All Issues* tab is calculated by multiplying the number of issues by the number of documents. For example, if you had 10 issues and 100 documents, issue predictions would equal 1000.
- A tab appears for every Issues field choice that has choice criteria. Click each to view the corresponding data. The tabs appear in order according to each choice's Order value. For information on changing the choice order, see Choices in the Admin guide.

When you select one of the aspect tabs in the bar, both the project metrics section and document 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 document 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 document 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.

6.4 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 ([Version Metrics tab below](#)), and one for a list of historical results for all previous versions ([History tab on page 45](#)).

Metrics		History							
Documents				aiR Analysis			Conflicts		
Reviewer Uncoded	Analyzed	Not Analyzed	Errored	Not Relevant	Borderline	Relevant	Total	Relevant	Not Relevant
88	112	0	0	55	39	18	1	0	1

6.4.1 Version Metrics tab

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

Section	Field Description
Documents	• Reviewer Uncoded(<i>for Relevance or Relevance and Key Documents analysis only</i>)—total number of documents that have not been coded by reviewers. These documents need to be coded to identify conflicts between human reviewers and aiR Review results. • Reviewer Coded Issues(<i>for Issues analysis only</i>)—total number of documents that reviewers coded as having the selected issue. Be aware that when the <i>All Issues</i> tab is selected, the counts for each issue are combined, which may result in a number that is more than the total document count. • 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 (-1 score) during analysis. For more information, see Handling document errors on page 75.
Issue Predictions	Note: The Issue Predictions section displays for Issues analysis when the <i>All Issues</i> tab is selected. • Not Relevant—issues predicted as <i>junk</i> or <i>not relevant</i> to the current aspect. • Borderline—issues predicted as <i>bordering</i> between relevant and not relevant to the current aspect. • Relevant—issues predicted <i>relevant</i> or <i>very relevant</i> to the current aspect. • Errored—issues that received an <i>error code</i> during analysis.
aiR Analysis	• Not Relevant—documents predicted as <i>junk</i> or <i>not relevant</i> to the current aspect. • Borderline—documents predicted as <i>bordering</i> between relevant and not relevant to the current aspect. • Relevant—documents predicted as <i>relevant</i> or <i>very relevant</i> to the current aspect. • No CBI—documents predicted as not containing confidential business information for the current aspect. This applies to the Confidential Business Information analysis type. • CBI—documents predicted as containing confidential business information for the current aspect. This applies to the Confidential Business Information analysis type.
Conflicts	Note: Conflicts are not displayed after running Apply on a project set. See Applying the prompt criteria on page 57 for more information on running an Apply function. Conflicts are when the human reviewers' coding and aiR's predictions are not in alignment. • Total—total number of documents that have a different coding decision from the predicted result by the human reviewer. This is the sum of the <i>Relevant Conflicts</i> and <i>Not Relevant Conflicts</i> fields. • Relevant—documents predicted as <i>relevant</i> or <i>very relevant</i> to the current aspect, but the coding decision in the related field says something else. • Not Relevant—documents predicted as <i>not relevant</i> to the current aspect, but the coding

Section	Field Description
	decision in the related field says <i>relevant</i>. • CBI—documents predicted as containing confidential business information to the current aspect. • No CBI—documents predicted as not containing confidential business information to the current aspect.

The metrics adjust their counts based on the type of results displayed:

- For *Relevance* analysis results, relevance-related metrics use the Relevance Field for their counts.
- For *Key Document* analysis results, relevance-related metrics use the Key Document Field for their counts.
- For *Issues* analysis results, relevance-related metrics count documents marked for the selected issue.

For instance, when viewing results for an issue such as Fraud, the aiR Predicted Relevant field displays documents identified as associated with Fraud. When viewing Key Document results, the aiR Predicted Relevant field displays documents identified as key documents.

6.4.1.1 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 Version Metrics 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.

Develop
BigThorium Docs
112 analyzed | 112 total
Version 4
Ran 1/14/2026
Analyze 39 documents
Clear selection

ResponsiveConflicts 1

KeyConflicts 0

Metrics
History

Documents

aiR Analysis

Conflicts

Reviewer Uncoded

Analyzed

Not Analyzed

Errored

Not Relevant

Borderline

Relevant

Total

Relevant

Not Relevant

88

112

0

0

55

39

18

1

0

1

Analysis Results

Summaries

#

Control Number

Reviewer Coded

aiR Prediction

aiR Score

aiR Rationale

aiR Consideration

aiR Citation 1

1

BTV5_000005028

Responsive

Borderline Responsive

2

This document contains communication between a City of Atlantis official and BigThorium employees, discussing compliance with local regulations and financial implications, which could be relevant to the allegations of unethical practices and bribery in securing contracts. The mention of aligning with local compliance requirements and briefing the sales team on protocols may hint at processes related to securing contracts.

The document does not explicitly mention the RFP, bribery, or unethical practices, and the discussion appears to focus on emergency preparedness rather than the municipal energy project contract. It is possible that this communication is unrelated to the allegations in the case.

Given the recent events in the is essential for BigThorium to compliance requirements.

2

BTV5_000005037

Non-Responsive

Borderline Responsive

2

This document includes communication between a BigThorium employee and an Atlantis City official, which falls under the responsive criteria of 'Any conversations that include ANYONE from Atlantis and ANYONE from BigThorium.' While the content of the email is unrelated to the allegations or the RFP, the content of communication.

The content of the email does not directly address the allegations, the RFP, or any relevant business dealings. It is purely social and unrelated to the case, which could argue against its responsiveness.

From: Samuel Jenkins <sjenkins@bigthorium.com>
To: Robert Williams <robert.williams@atlantis.gov>

You can also filter documents in the document grid by selecting them in the table. See [Filtering and selecting documents for analysis on page 46](#).

6.4.2 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 42](#).

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.

Metrics	History						
Version	Timestamp	Docs w/ aiR Predict...	Docs w/o aiR Predi...	Uncoded	Errored	Predicted Relevant	Pre
3	11/26/2025, 3:40 AM	1	0	0	0	0	1
2	11/19/2025, 7:04 AM	107	1993	1904	0	1	70
1	11/19/2025, 6:07 AM	25	2075	1904	0	0	24

6.5 Results and summaries

In the middle section of the dashboard, a document grid appears with two tabs: Analysis Results and Summaries. The Analysis Results tab shows documents alongside their scores and results, while the Summaries tab shows the descriptive topics and summaries that aiR generated for each document.

6.5.1 Analysis Results tab

The Analysis Results tab shows a list of all documents in the project and their results. If the documents have aiR Review analysis results, they appear in the grid. 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 [Analyzing aiR Review results on page 72](#).

Each of the objects has a group of associated fields: a prediction, a corresponding numerical score, the rationale supporting the prediction, any considerations that might prove the prediction wrong, any citations found within the document (maximum 5) to support the prediction, and any errors encountered. Use these fields to understand why aiR scored a document a certain way, to provide insights into your documents, and to provide insights into ways you may want to iterate on the prompt criteria.

Analysis Results								Summaries ?
								1 - 112 of 112 200 per page
#	Control Number	Reviewer Coded	aiR Prediction	aiR Score	aiR Rationale	aiR Consideration	aiR Citation	
1	BTv5_000005028	Responsive	Borderline Responsive	2	This document contains communication between a City of Atlantis official and BigThorium employees, discussing compliance with local regulations and financial implications, which could be relevant to the allegations of unethical practices and bribery in securing contracts. The mention of aligning with local compliance requirements and briefing the sales team on protocols may hint at processes related to securing contracts.	The document does not explicitly mention the RFP, bribery, or unethical practices, and the discussion appears to focus on emergency preparedness rather than the municipal energy project contract. It is possible that this communication is unrelated to the allegations in the case.	Given the re is essential compliance	
2	BTv5_000005032	Non-Responsive	Not Responsive	1	Das Dokument enthält keine relevanten Informationen zu den Kriterien der Relevanz, da es sich ausschließlich auf allgemeine Risikomanagement- und Notfallplanungsstrategien konzentriert und keine spezifischen Hinweise auf die RFP, Bestechungsvorwürfe oder Interaktionen mit der Stadt Atlantis enthält.	Es könnte argumentiert werden, dass ein verbessertes Risikomanagement indirekt mit der Fähigkeit von BigThorium zusammenhängt, sich auf Ausschreibungen wie die RFP der Stadt Atlantis vorzubereiten. Allerdings fehlen direkte Bezüge zu den relevanten Themen.		
3	BTv5_000005037	Non-Responsive	Borderline Responsive	2	This document includes communication between a BigThorium employee and an Atlantis City official, which falls under the responsive criteria of 'Any conversations that	The content of the email does not directly address the allegations, the RFP or any relevant business dealings. It is purely social and unrelated to the case, which could argue	From: Samu <sjenkins@ To: Robert V <robert.willi;	

Note: aiR 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.

6.5.2 Summaries tab

The Summaries tab shows a list of all documents in the project and their aiR-generated topics and summaries. These topics and summaries are independent of the prompt criteria, and they are the same for each document across all projects.

Analysis Results <u>Summaries ?</u>			
1 - 112 of 112 200 per page			
#	Control Number	aiR Topic	aiR Content Summary
Filter Filter Filter			
1	BTv5_000005028	Emergency Preparedness and Response Plans Email	This is an email exchange between Robert Williams from Atlantis government and John Maxwell from BigThorium discussing compliance with emergency preparedness and response plans. The correspondence covers briefing the sales team on revised protocols and addressing financial implications and legal compliance.
2	BTv5_000005032	Email Exchange on Risk Management Planning	This is an email correspondence between Daniel Carter, a Senior Partner at LawFirmX, and Laura Smith, a PR Specialist at BigThorium, discussing enhancements in risk management and contingency planning strategies at BigThorium. They consider implementing a centralized digital platform and plan to meet to discuss further steps.
3	BTv5_000005037	Inspirational Quotes Email Chain	This is an email between Samuel Jenkins and Robert Williams discussing sharing inspirational quotes to spread positivity within their teams. Robert appreciates a Gandhi quote and plans to share it at an upcoming town hall meeting.
4	BTv5_000005042	Email Exchange on Site Selection Criteria	This is an email conversation among BigThorium employees Kayla Fry, Mark Simmons, and Linda Clark discussing criteria for site selection and evaluation of potential nuclear power plant locations. The emails cover proposed evaluation factors, input requests, and planning for stakeholder outreach.
5	BTv5_000005044	Urgent Compliance Review of Nuclear Standards	This is an email exchange between Emily Roberts, Legal Counsel, and Francis Ham, CEO of BigThorium, discussing urgent compliance issues related to national and international nuclear standards. They address discrepancies in assessments and the need for a comprehensive report to prepare for upcoming inspections and stakeholder

For more information on aiR topics and summaries, see [Document summaries on page 73](#).

6.5.3 Viewing inline highlighting and citations in the Viewer

To view inline highlighting and citations for an individual document, click on the document **Control Number** link. The Viewer opens and displays results for the selected prompt criteria version. For more information on viewing aiR Review results in the Viewer, see aiR Review Analysis information in the Viewer documentation.

6.5.4 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 grid. The number of selected documents is reflected in the Run button's text. To remove filtering, click **Clear selection** underneath the Run button.

Develop

BigThorium Docs

112 analyzed | 112 total

Version 4

Ran 1/14/2026

Analyze 3 documents

Clear selection

Responsive

Conflicts 1

Key

Conflicts 0

Metrics

History

Documents				aiR Analysis			Conflicts		
Reviewer Uncoded	Analyzed	Not Analyzed	Errored	Not Relevant	Borderline	Relevant	Total	Relevant	Not Relevant
88	112	0	0	55	39	18	1	0	1

Analysis Results

Summaries

1

- 112 of 112

200

per page

#

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

Responsive

Borderline Responsive

2

This document contains communication between a City of Atlantis official and BigThorium employees, discussing compliance with local regulations and financial implications, which could be relevant to the allegations of unethical practices and bribery in securing contracts. The mention of aligning with local compliance requirements and briefing the sales team on protocols may hint at processes related to securing contracts.

The document does not explicitly mention the RFP, bribery, or unethical practices, and the discussion appears to focus on emergency preparedness rather than the municipal energy project contract. It is possible that this communication is unrelated to the allegations in the case.

Given the recent events in the is essential for BigThorium to compliance requirements.

2

☒

BTV5_000005032

Non-Responsive

Not Responsive

1

Das Dokument enthält keine relevanten Informationen zu den Kriterien der Relevanz, da es sich ausschließlich auf allgemeine Risikomanagement- und Notfallplanungsstrategien konzentriert und keine spezifischen Hinweise auf die RFP, Bestechungsvorwürfe oder Interaktionen mit der Stadt Atlantis enthält.

Es könnte argumentiert werden, dass ein verbessertes Risikomanagement indirekt mit der Fähigkeit von BigThorium zusammenhängt, sich auf Ausschreibungen wie die RFP der Stadt Atlantis vorzubereiten. Allerdings fehlen direkte Bezüge zu den relevanten Themen.

3

☒

BTV5_000005037

Non-Responsive

Borderline Responsive

2

This document includes communication between a BigThorium employee and an Atlantis City official, which falls under the responsive criteria of 'Any conversations that include ANYONE from Atlantis and ANYONE from BigThorium'. While the content of the email is unrelated to the allegations or the RFP, the mere fact of communication between these parties makes it relevant.

The content of the email does not directly address the allegations, the RFP, or any relevant business dealings. It is purely social and unrelated to the case, which could argue against its responsiveness.

From: Samuel Jenkins <sjenkins@bigthorium.com>
To: Robert Williams <robert.williams@atlantis.gov>

4

☐

BTV5_000005042

Non-Responsive

Borderline Responsive

2

The document discusses site selection and evaluation criteria for nuclear power plant locations, including community outreach and regulatory compliance. While it does not

The document predates the RFP-ATL-22-437 and does not mention Atlantis or the specific municipal energy project. It focuses on general site selection criteria rather than

I have drafted a set of criteria selection and evaluation of ou nuclear power plant locations

You can also filter the grid by clicking the metrics in the Version Metrics section. See [Filtering the Analysis Results using version metrics on page 44](#).

6.5.5 Clearing selected documents

Click the **Clear selection** link below the Analyze button to deselect all documents in the grid. This resets your selections so the next analysis includes all documents in the data source.

Analyze 2 documents

Clear selection

7 Creating an aiR Review project

The instructions you give aiR Review are called prompt criteria. They 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. To help improve 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 allows you to evaluate how aiR Review's coding compares to a human reviewer's coding and adjust the prompts accordingly.

Refer to [Best practices on page 15](#) for more information on prompt criteria tips, sample workflow guidance, and iteration workflow.

7.1 Project workflow

The aiR Review project workflow has three basic parts:

1. [Setting up the project](#)—defining the project name, description, data source (saved search), and analysis type.
2. [Developing the prompt criteria](#)—specifying the criteria that the AI should apply when reviewing documents according to the chosen analysis type.
3. [Running the first analysis](#)—running AI analysis on documents according to prompt criteria.

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

7.2 Job capacity, size limitations, and speed

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

Note: Authorized administrators can configure job size limits per user group at the instance level to help control and manage the size of jobs that users execute. For more information on setting job sizes, see [Setting aiR Review job size limits on page 100](#).

7.2.1 Size limits

The documents and prompt criteria have the following size limits:

Category	Applies to	Maximum	Description
Prompt kickstarter	Documents	10 documents	Maximum of 10 documents with a maximum combined total of 300,000 characters. See Using prompt kickstarter (Draft with AI) on page 64 for more information on using the prompt kickstarter feature for some analysis types.
	Combined document content	300,000 combined total	
Prompt criteria	Combined content across all tabs	15,000 characters	The combined content of all prompt criteria tabs cannot exceed 15,000 characters to avoid errors. See Best Practices for tips on writing prompt criteria and removing excess information.
Document extracted text	Text per document	0.05 - 300KB	Documents must fall within this range to be eligible for aiR Review analysis. Although aiR Review can handle extracted text up to 300KB, use caution above 150KB as processing errors may prevent results. This range also serves as the limits of an aiR Review "aiR Unit" that is available for aiR Unit Analysis.
Custom Analyses limits	Prompt criteria	1,000 characters	Custom insight/prompt

Category	Applies to	Maximum	Description
			cannot exceed 1,000 characters.
	Insight title	30 characters	Custom insight title cannot exceed 30 characters.
	Vision analysis native image	20MB	- Vision analysis uses native JPEG, GIF, and PNG image files. - Text analysis uses extracted text.

Notes:

- 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 [Handling document errors on page 75](#).

7.2.2 Volume limits per instance

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

Limit per instance	Applies to	Maximum
Job size	Relevance, Key Document, Confidential Business Information (CBI), and Custom analyses	300,000 documents
	Issues analysis	250,000 documents
Total documents	All queued and running jobs in an instance	600,000 documents
Concurrent large jobs	Jobs with more than 200 documents	- Relevance, Keys, Issues, and CBI analyses projects = 10 queued or running jobs per instance - Custom analyses projects = 20 queued or running jobs per instance
Concurrent small jobs	Jobs with 200 or fewer documents	No limit

7.2.3 Speed

After a job is submitted, aiR Review can analyze up to 3 million documents per day.

For individual jobs, aiR Review averages between 10,000 and 13,000 documents per hour.

Job speeds vary widely depending on the region, the number and size of documents, the overall load on the LLM, and other similar factors.

7.3 Setting up the project

Begin by configuring the aiR Review project, selecting the appropriate type of analysis, and specifying the desired data source (saved search). The aiR Review project serves as a repository for all iterative project sets—develop, validate, and apply—that are created during the refinement of prompt criteria.

After the project is set up, proceed with [developing prompt criteria](#) and [running an analysis](#).

Note: Authorized administrators can configure job size limits per user group at the instance level to help control and manage the size of jobs that users execute. For more information on setting job sizes, see [Setting aiR Review job size limits on page 100](#).

7.3.1 Choosing the data source

Before setting up the project, create a saved search that contains a small sample of the documents you want reviewed. You will choose this data source (saved search) when setting up the project.

To help improve results:

- Include roughly 50 test documents, ensuring a mix of both documents that align with the intended analysis type and those that do not.
- Have human reviewers code the documents in advance.
- See [Tips for creating sample documents set on page 17](#) for more information.
- See [Selecting a Prompt Criteria Iteration Sample for aiR Review](#) on the Community for suggestions on choosing documents for the sample.

For details on creating a saved search, see [Creating or editing a saved search](#) in the Searching guide.

7.3.2 Choosing an analysis type

aiR Review offers the following analysis types, each suited to a specific review or investigation phase. Choose the appropriate type when setting up your project. Then, based on the analysis type chosen, you will need the fields indicated in the table. aiR Review does not actually write to these fields. Instead, it uses them for reference when reporting on its predictions.

Analysis Type	Description	Fields required
Relevance	Analyzes whether documents are relevant to a case or situation that you describe, such as documents responsive to a production request.	One single-choice results field. The field must have at least one choice.
Relevance and Key Documents	Analyzes documents for both relevance and whether they are "hot" or key to a case or investigation.	Two single-choice results fields. These should have distinct names, such as "Relevant" and "Key," and each field must have at least one choice.
Issues	Analyzes documents for whether they include content that falls under specific issues or categories. For example, you might use this to check whether documents involve coercion, retaliation, or a combination of both.	One multi-choice results field. Each of the issues you want to analyze should be represented by a choice on the field.

Analysis Type	Description	Fields required
		Note: Currently, aiR Review analyzes a maximum of 10 issues per run. You can have as many choices for the field as you want, but you can only analyze 10 at a time. To analyze more, run multiple jobs. If multiple jobs are run, the previous results are not overwritten.
Confidential Business Information	Analyzes and detects whether documents contain Confidential Business Information (CBI) based on your selected Confidential Business Information Types, then classifies documents by confidentiality level. Note: This functionality is being released through a phased rollout and is only available to clients in the US and EU Data Zone at this time. For more information on phased rollouts, see the Phased Rollout FAQ article available on the Community.	One single-choice results field. The field must have at least one choice. All the choices signaling the presence of confidential information in the document must be selected.
Custom	Analyzes text and image documents based on your custom insight prompts. Each custom insight acts as an independent analysis for questions, extractions, classifications, or evaluative instructions applied to every eligible document.	One single-choice results field. The field must have at least one choice. Each of the custom insights you want to analyze will contain prompt instructions and a selected Field Type (Fixed Length Text, Whole Number, Single Choice, or Date), which determines how the result is validated and published.

7.3.3 Setting up an aiR Review project

The aiR Review project houses all of the project iterations (project sets) as you go through the process of refining the prompt iterations. See [Using project sets on page 55](#) for more information on using project sets for iteration. After setting up the foundation of the project below, you can begin developing prompt criteria. See [Developing prompt criteria on page 58](#) for the steps on writing the set of prompt criteria for the project analysis.

To set up an aiR Review project:

1. Navigate to the **aiR Review>aiR Review Project** tab.
2. Click **New aiR Review Project**.

3. Fill out the fields on the Setup Project modal to configure the project. See [Setup project fields below](#) for descriptions of each one.

aiR Setup Project

The following settings cannot be changed after the Project has been created.

Project Name *

Description

Data Source * ? Select Clear

Analysis Type * ? Choose...

Project Prompt Criteria * ☒ Start Blank ☐ Copy Existing

Prompt Criteria Name * ?

Project Use Case * ? Choose...

Create Project Cancel

4. Click **Create Project**.

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

Next, refer to [Developing prompt criteria on page 58](#) for information on the next phase on developing prompt criteria for project analysis.

7.3.3.1 Setup project fields

- **Project Name**—enter a name for the project.
- **Description**—enter a useful short description of the project.
- **Data source**—select the saved search that holds your document sample set. Refer to [Choosing the data source on page 50](#) for more information.
- **Analysis Type**—select one of the following analysis types. For more information, see [Choosing an analysis type on page 50](#).
 - **Relevance**—analyzes whether documents are relevant to a case or situation that you describe, such as documents responsive to a production request. See [Relevance tab on page 60](#) for more information on the fields involved with developing the prompt criteria.
 - **Relevance and Key Documents**—analyzes documents for both relevance and whether they are “hot” or important (key) to a case. See [Relevance tab on page 60](#) and [Key Documents tab on page 61](#) for more information on the fields involved with developing the prompt criteria.
 - **Issues**—analyzes documents for whether they include content that falls under specific issues or categories. See [Issues tab on page 62](#) for more information on the fields involved with developing the prompt criteria.

- **Confidential Business Information**—analyzes documents that contain confidential business information based on your selected CBI types. For more information, see [Confidential Business Information tab on page 62](#) and [Confidential Business Information types on page 63](#).
- **Custom**—enables you to create up to five custom insights to analyze documents, each with its own Field Type defined on the Custom tab. See [Custom analyses on page 27](#) and [Custom tab on page 64](#) for more information.
- **Project Prompt Criteria**—select one of the following:
 - **Start blank**—select this option if you plan to write new prompt criteria from scratch or to use the **Draft with AI** (prompt kickstarter) feature. See [Using prompt kickstarter \(Draft with AI\) on page 64](#) for more details. This option is the only available option if *Custom* is selected for Analysis Type.
 - **Copy existing**—select this option to use a previously created set of prompt criteria in this workspace and copy it into this project for modifying as needed. This option does not display if *Custom* is selected for Analysis Type.
- **Prompt Criteria Name**—either keep the default name that automatically populated or click the pencil icon to rename it to identify the prompt criteria used by the job. The name must be unique and will display on the aiR Review Jobs list. Skip this field if *Custom* is selected for Analysis Type.
- **Custom Insight Model**—this field displays if *Custom* is selected for Analysis Type. Choose **Text** to run the analysis on the extracted text of documents or **Vision** to run it on documents where their native file is either JPEG, GIF, or PNG. For more information, see [Custom tab](#).
- **Project Use Case**—choose the option that best describes the purpose of the project. If none of the options describe the project, choose **Other** and 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. The Project Use Case field is used for reporting and management purposes and does not affect how the project runs.

7.4 Deleting an aiR Review project

Use Mass Operations to delete one or more projects from the project list. For more information on Mass delete, see Mass delete documentation.

Note: Mass delete is permanent action, so be sure to double-check your selections before proceeding. This operation cannot be undone.

1. Navigate to the **aiR Review Projects** tab.
2. Select the check box next to one or more projects to be deleted from the projects list.
3. Click **Delete** from the Mass Operations option at the bottom of the grid.
4. Click **View Dependencies** to review dependencies as needed on the confirmation modal.

Delete 1 aiR for Review Project



Warning: Clicking delete will permanently delete this object and any children and unlink any associated objects. Click View Dependencies below for details.

> [View Dependencies](#)


Delete

Cancel


5. Click **Delete** to confirm.

7.4.1 Hiding document analysis results in the Viewer after deleting a project

Even though an aiR Review project is deleted, the associated analysis results remain linked to the document and are not deleted. These results will still be displayed in the Document Viewer. For more information on document linking, see [Managing aiR Review jobs on page 99](#) and see Viewer documentation for more information on aiR Review analysis results in the Viewer.




To hide the results from the Viewer, go to the **aiR Review Jobs** tab and click the **Visible eye** icon beside the Job ID to change the setting to **Hide results**. The Results column will update from Visible to **Hidden**, and the eye icon will have a slash through it. This removes the link to the analysis results so they do not display in the Document Viewer's panel. Click the icon again to relink them and show results in the Viewer. For more information on document linking, see [Managing aiR Review jobs on page 99](#).



7.5 Using project sets

Project sets make it easier to manage and refine prompt criteria for aiR Review projects. They let you create new sets within a single aiR project so you can iterate without having to start over with a new aiR project. You can use these sets to change saved searches (data sources) during prompt criteria development, to switch between sets to review results, or to start prompt criteria validation in Review Center.




Develop	BigThorium Bribery 0 analyzed 115 total	Version 4 Saved 8/15/2025	
Develop	BigThorium Bribery 13 analyzed	Version 3 Ran 5/8/2025	
Develop	BigThorium Bribery 14 analyzed	Version 2 Ran 3/18/2025	
Develop	BigThorium Bribery 20 analyzed	Version 1 Ran 3/18/2025	

The project set workflow outlined below aligns with aiR Review’s phases (see [aiR Review phases on page 7](#)).

1. **Develop:** This set type is used to develop your prompt criteria on a small document set. With develop sets, the data source (saved search) being used in the project can be changed to allow greater flexibility to run multiple jobs within a single aiR project. See [Creating a new project set on the next page](#) for more information.
2. **Validate:** This set type is used to validate your prompt criteria leveraging Review Center. See [Validating the prompt criteria on the next page](#) for more information.
3. **Apply:** This set type is commonly used after validation to apply your prompt criteria to a larger document population. See [Applying the prompt criteria on page 57](#) for more information.

7.5.1 Navigating the project set details strip

The project set details strip contains the following:



Develop	ALL DOCUMENTS 13 analyzed 4,989 total	Version 2 Ran 6/15/2026	
---------	--	----------------------------	--

- **Project set type**—identifies whether the project set is in the Develop, Validate, or Apply phase.
- **Name**—identifies the name of the saved search used for the set, number of documents analyzed, and total number of documents in the saved search.
- **Refresh button**—click the circular arrows to update the selected data source (saved search) so it reflects any changes made to the saved search document list after the project was created. For example, to ensure it reflects newly added or deleted documents or updated saved search criteria. This option is only available for Develop project sets. See [Refreshing a data source on page 57](#).
- **Version number**—identifies the version number of the prompt criteria for the project set and the date it ran or was saved.
- **Down arrow button**—click the down arrow to switch between project sets to review analysis results and to track project activity.
- **+ sign**—click the plus sign to create a new project set using a different data source, to validate the currently selected prompt criteria version, or to apply it to a larger document population.

7.5.2 Creating a new project set

Use the steps below to create (develop) a new project set within the current aiR Review project. This enables you to change the data source (saved search) used to iterate on the prompt criteria and view results.

1. Click the project sets **+** **sign**.
2. Click **I want to develop my prompt criteria using a different data source**.

Create a New Set

I want to **develop** my prompt criteria using a different data source.

I want to **validate** the prompt criteria on a target document population.

I want to **apply** the prompt criteria to a document population.

Cancel

3. Click **Select** and choose the data source (saved search) to use for this project set.
4. Click **Create Develop Set**.
5. Proceed with modifying the prompt criteria as needed.

7.5.3 Validating the prompt criteria

Use the steps below to validate the project set's prompt criteria within the current aiR Review project.

Note: Validating prompt criteria is not supported for Custom analyses.

1. Within the desired project set, click the project sets **+** **sign**.
2. Click **I want to validate the prompt criteria on a target document population**.

Create a New Set

I want to **develop** my prompt criteria using a different data source.

I want to **validate** the prompt criteria on a target document population.

I want to **apply** the prompt criteria to a document population.

Cancel

3. Click **Select** and choose the data source (saved search) to use for this project set.

Note: We strongly recommended that you avoid modifying the data source (saved search) after you start the validation process as this will interfere with the results.

4. Fill out the **Validation Settings fields** to determine the validation sample. For more information on each field, see [Validation Settings fields on page 91](#).
5. Click **Analyze [#] documents**.
6. Review the confirmation summary modal and provide email addresses, if needed.
7. Click **Start Analysis**.

- When the analysis is finished, click **Go to Review Center** to begin the human review and coding of the validation sample in Review Center. Refer to [Prompt Criteria validation in Review Center on page 92](#) for the steps to follow in Review Center.

For more information on validating prompt criteria, see [Setting up aiR Review prompt criteria validation on page 88](#).


7.5.4 Applying the prompt criteria

Use the steps below to apply the project set's prompt criteria within the current aiR Review project.

Note: For Custom analyses projects, see [Running Apply & Publish on page 36](#) and [Applying the results and publishing to the Document Object on page 35](#) for more information on the apply process.

- Within the desired project set, click the project sets **+ sign**.
- Click **I want to apply my prompt criteria to a document population**.

Create a New Set



I want to **develop** my prompt criteria using a different data source. >

I want to **validate** the prompt criteria on a target document population. >

I want to **apply** the prompt criteria to a document population. >

Cancel

- Click **Create Apply Set**.
- Click **Analyze [#] documents**.
- Review the confirmation summary modal and provide email addresses, if needed.
- Click **Start Analysis**. For more information, see [Running the analysis on page 70](#).

Note: Conflicts are not displayed in the Project Metrics section of the dashboard after running Apply on a project set. See [Version Metrics tab on page 42](#) for more information on the Conflicts tab.

After the prompt criteria are applied to the document population, they cannot be changed. The only way to revise the criteria is by creating a new develop project set and updating it. See [Creating a new project set on the previous page](#) for the details.

Note: Please note that if the project set has already completed the prompt criteria validation process, you cannot change it by creating a new project set. Instead, you must start a new aiR Review project. For additional guidance, see [Validating prompt criteria](#) and [Setting up the project](#).

7.5.5 Refreshing a data source

If the document list in the selected data source saved search changes after you create a project, click the **Refresh** button (circular arrows) next to the data source in the project set details strip. This updates the project analysis to use the most current document list. For example, documents may have been added or removed, or the saved search criteria may have changed after the project started. This option is available only for Develop project sets.

After a project set begins validation, the Refresh option is no longer available. In this situation, you need to create a new project to use the updated data source.

8 Developing prompt criteria

The prompt criteria are a set of inputs that provide aiR Review with the necessary context it needs to understand the matter and evaluate each document based on the provided criteria. Developing these criteria is essential for training aiR Review as a "reviewer," paralleling the process of preparing a human reviewer. See [Best practices on page 15](#) for tips and workflow suggestions.

Depending on the selected analysis type during [set up](#), the prompt criteria panel in the [dashboard](#) will display different tabs. To assist in writing your initial prompt criteria, fields include italicized, grayed-out helper text offering example entries. Use it for reference when drafting your own criteria.

For certain analysis types, you may automatically generate initial prompt criteria from existing case documents, such as review protocols, requests for production or case memos, by using the prompt kickstarter feature. See [Using prompt kickstarter \(Draft with AI\) on page 64](#) for more information.

If needed, you can create project sets within a single aiR Review project. This enables you to develop, validate, and implement prompt criteria without needing to create new projects for each iteration. See [Using project sets on page 55](#) for more information.

To learn more about how prompt criteria versioning works and how versions affect the Viewer, see [How prompt criteria versions affect the Viewer on page 68](#).

8.1 General prompt criteria workflow

We recommend the following workflow for developing and iterating on prompt criteria for all analysis types:

1. For your first analysis, run the prompt criteria on a saved search of 50 test documents that are a mix of relevant and not relevant.
2. Compare the results to human coding. In particular, look for documents that aiR Review coded differently than the humans did and investigate possible reasons. This could include unclear instructions, needing to define an acronym or code word, or other blind spots in the prompt criteria.
3. Tweak the prompt criteria to adjust for blind spots.
4. Repeat steps 1 through 3 until aiR Review produces results that reasonably align with coding decisions 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 Review only sees the extracted text of a document. It does not see any non-text elements, such as advanced formatting, embedded images, or videos. We do not recommend using aiR 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.

8.2 Entering prompt criteria

The tabs that appear on the prompt criteria panel depend on the analysis type you selected during set up. Refer to [Setting up the project](#) for more information.

1. On the aiR Review dashboard, you can enter data in the prompt criteria tabs using any of the methods below. Grayed-out, italicized helper text displays in the boxes as a guide for your entries.
 - Manually enter the required information as outlined in the [Prompt criteria tabs and fields on the next page](#) section.

- Click the **Draft with AI** button (prompt kickstarter) to upload and use existing documentation, like requests for production or review protocols, to automatically fill in the tabs with draft criteria. This option only applies to Relevance, Key Documents, and Issues analysis types. See [Using prompt kickstarter \(Draft with AI\) on page 64](#) for more information.
 - Use project sets to develop, validate, and apply prompt criteria using a new data source (saved search) without having to create a new aiR project. See [Using project sets on page 55](#) for more information.
2. Click **Save** after entering data on each tab or after all tabs are completed.
 3. After you have the desired prompt criteria set, click **Analyze [#] documents** to analyze documents. For more information on analyzing documents, see [Running the analysis on page 70](#).

Note: After job analysis starts, tab fields cannot be edited, unless noted otherwise.

4. To validate the prompt criteria, see [aiR Review prompt criteria validation on page 83](#) and [Setting up aiR Review prompt criteria validation on page 88](#) for more information.

Note: Performing prompt validation through aiR Review and Review Center does not apply to Confidential Business Information and Custom analysis projects.

5. If you want to use different saved searches within the project, see [Creating a new project set on page 56](#).

8.3 Prompt criteria tabs and fields

Use the sections below to enter information in the necessary fields.

Note: The combined total across all prompt criteria tabs must not exceed 15,000 characters.

8.3.1 Case Summary tab

The Case Summary gives the Large Language Model (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. This tab only appears for Relevance, Key Documents, and Issues analysis types.

Considerations:

- To ensure optimal results, limit the Case Summary content to 20 or fewer sentences overall, and 20 or fewer entries each of People and Aliases, Noteworthy Organizations, and Noteworthy Terms. Although longer entries can be processed by the system, exceeding these limits may reduce result quality.
- If using prompt kickstarter to automatically generate initial prompt criteria, these fields will automatically populate based on the documents uploaded or text manually added using the **Draft with AI** feature. See [Using prompt kickstarter \(Draft with AI\) on page 64](#) for more information. Refer to this section as needed for field descriptions.

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.
- **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, entities, or other notable characteristics.

- **Noteworthy Terms**—list and define any relevant words, phrases, acronyms, jargon, project code names, or slang that might be important to the analysis.
- **Additional Context**—list any additional information that does not fit the other fields. This section can be left blank.

Depending on which Analysis Type you chose when [setting up the project](#), the remaining tabs will be *Relevance*, *Key Documents*, or *Issues*. Refer to the appropriate tab section below for more information on each one.

For best results when entering information in the Case Summary tab:

- Keep Case Summaries to 20 sentences or fewer. Longer summaries often include excess details, weakening prompt quality and increasing borderline documents. Concise, direct summaries aligned with best practices lead to better results.
- For People and Aliases, avoid listing every person in the case. The system picks up common information, such as work emails. Instead, focus on including people and aliases a typical person could not easily infer, such as unusual or personal email addresses, and only when those details matter for determining document relevance.

8.3.2 Relevance tab

This tab defines the fields and criteria used for determining if a document is relevant to the case. It appears if you selected *Relevance* or *Relevance and Key Documents* as the Analysis Type during project setup.

Fill out the following:

Note: If using prompt kickstarter to automatically generate initial prompt criteria, these fields will automatically populate based on the documents uploaded or text manually added using the **Draft with AI** feature. See [Using prompt kickstarter \(Draft with AI\) on page 64](#) for more information. Refer to this section as needed for field descriptions.

- **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 you want to use as criteria for the Relevance analysis. The Relevance analysis uses these issues to assess document relevance, but it does not assign issue choices to documents. This selection cannot be changed after the first job run.

Note: When using the **Issues Field** on the Relevance tab for Relevance or Relevance and Key Document projects, aiR does not make Issue predictions or coding and does not write to *aiR Relevance Analysis::Issues* on the Documents list. If you want aiR to provide this information, you must run another project using **Issues** as the analysis type on the project set up dialog. See [Setting up an aiR Review project on page 51](#) for more information.

The Issues Field selections on the Relevance tab are used to assess document relevance, but aiR does not assign issue choices to documents. For example, you could tell aiR that documents containing any of the defined issues are relevant.

The *aiR Relevance Analysis::Issues* column on the Documents list may show results for jobs run before the early August 2026 system fix that removed Issue Field results from this field. For jobs run after the fix, the *aiR Relevance Analysis::Issues* column is empty.

-
- **Choice Criteria**—this option displays when an Issues Field is selected. Individually select each field Choice to be analyzed. A maximum of 10 choices can be analyzed at a time. To remove a selected choice, click the X in its row. For each choice, enter criteria description in the text box that determine whether that issue applies to a document.

Notes:

- You can edit the Choice Criteria description in the text box after the first job is run, but you cannot change the Issue Field selection.
- If you delete or modify the name of a field Choice used in the project before running an analysis, an error message appears indicating the choices configured for the Issues Field are not valid. To resolve this, return to the Choice Criteria field within the prompt criteria tab and update it as necessary. For further details about deleting or modifying Choices, refer to the Review Guide for documentation on Choices.

For best results when writing the Relevance Criteria:

- Limit it to 5-10 sentences.
- Do not use a plain number for the Choice Criteria name, such as 12.
- 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.”

8.3.3 Key Documents tab

This tab defines the fields and criteria used for determining if a document is "hot" or key to the case. It appears if you selected *Relevance and Key Documents* as the Analysis Type during project setup.

Fill out the following:

Note: If using prompt kickstarter to automatically generate initial prompt criteria, these fields will automatically populate based on the documents uploaded or text manually added using the **Draft with AI** feature. See [Using prompt kickstarter \(Draft with AI\) on page 64](#) for more information. Refer to this section as needed for field descriptions.

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

8.3.4 Issues tab

This tab defines the issue fields and criteria used for determining whether a document relates to the set of topics or issues specified. It appears if you selected *Issues* as the Analysis Type during project setup.

Fill out the following:

Note: If using prompt kickstarter to automatically generate initial prompt criteria, these fields will automatically populate based on the documents uploaded or text manually added using the **Draft with AI** feature. See [Using prompt kickstarter \(Draft with AI\) on page 64](#) for more information. Refer to this section as needed for field descriptions.

- **Field**—select a multi-choice field that represents the issues in the case.
- **Choice Criteria**—select each of the field Choices one by one to be analyzed. A maximum of 10 choices can be analyzed at a time. To remove a selected choice, click the X in its row. For each choice, write criteria description in the text box listing the criteria that determine whether that issue applies to a document, or this information can be populated using prompt kickstarter.

Notes:

- The field choices cannot be changed after the first job run. However, you can still edit the criteria description in the text box.
 - If you delete or modify the name of a field Choice used in the project before running an analysis, an error message appears indicating the choices configured for the Issues Field are not valid. To resolve this, return to the Choice Criteria field within the prompt criteria tab and update it as necessary. For further details about deleting or modifying Choices, refer to the Review Guide for documentation on Choices.
-

For best results when writing the Choice Criteria:

- Include keywords, phrases, legal concepts, parties, entities, and legal claims.
- Include any criteria that would exclude a document from relating to that issue, such as falling outside a certain date range.
- Be clear and specific; avoid ambiguity.
- Limit the criteria description for each choice to 5-10 sentences.
- Each choice must have its own criteria. If a choice has no criteria, either fill it in or delete the choice.
- Do not use a plain number for the Choice Criteria name, such as 12.

8.3.5 Confidential Business Information tab

This tab defines the fields and criteria used for identifying Confidential Business Information (CBI) types and classifying documents by confidentiality level, such as Confidential, Highly Confidential, Attorneys' Eyes Only. It appears if you selected *Confidential Business Information* as the Analysis Type during project setup.

Note: This functionality is being released through a phased rollout and is only available to clients in the US and EU Data Zone at this time. For more information on phased rollouts, see the [Phased Rollout FAQ](#) article available on the Community.

Considerations:

- The selected types, their associated positive choice, and any additional details can be changed between job runs. However, the CBI Field and Positive Choice Rank cannot be changed.
- Confidentiality ranking is controlled by the order of choices on the selected CBI field in the workspace, not within the prompt criteria panel in the aiR Review application.

Fill out the following:

- **Confidential Business Information Field**—select the desired CBI field from the list. This field represents whether a document is sensitive/confidential or not. This selection cannot be changed after the first job run.
- **Positive Choice Rank**—depending on which CBI Field was selected, this list displays the relevant options in the same order as they appear in the workspace, with top choice being most confidential/highly sensitive.
 - The ranking is crucial as it will impact document results. You can adjust these rankings by navigating to **Fields** in the workspace, opening the **CBI field**, and reordering the **Choices**.
 - Use the X icon to delete non-confidential choices.
 - The CBI field and Positive Choice Rank cannot be modified after the first job run.
- **Information Types**—select one or more Information Types that you are interested in identifying in your data-set. Only the selected types will be analyzed in your documents. Be selective in choosing these types to avoid overlap. If there are multiple hits, only one citation will be displayed for each Information Type in the Viewer. See [Confidential Business Information types below](#) for descriptions of each one.
- **Positive Choice**—select the confidential designation corresponding to the Information Type chosen. The choices in this field reflect what you selected for Positive Choice Rank.
- **Additional Details**—we strongly recommend providing additional context in this box to further clarify and describe each Information Type. Each Information Type selected previously has a broad, built-in system description, so the Additional Details entered here add to that description rather than replace it. This information assists the model's ability to make more accurate classification predictions. Suggested guidelines:
 - Include matter-specific context on how to identify the documents you are looking for, especially key words and phrases.
 - Be clear and specific; avoid ambiguity.
 - Limit the context added for each type to 5-10 sentences.

8.3.5.1 Confidential Business Information types

The Information types eligible for reporting are described below. When a document contains multiple hits, only one citation per type is shown.

- **Business Plans & Strategy**—relates to internal materials discussing business direction, priorities, or planning activities.
- **Customer Testimonials**—includes customer-provided statements or feedback regarding products or services.
- **Financial Data**—refers to non-public financial-related materials used internally by a business.

- **General Confidential Business Data**—a broad category intended to capture internal business information that does not fit within a more specific category.
- **Intellectual Property (General)**—relates to materials associated with a company's intellectual property or innovation efforts.
- **Marketing Plans & Campaigns**—includes information related to marketing initiatives, planning, or promotional efforts.
- **Sales Data**—refers to internal information related to sales activities or performance.
- **Technical Information**—includes information describing technical aspects of products, services, or operations.
- **Trade Secrets**—refers to sensitive business information that a company treats as secret.

8.3.6 Custom tab

This tab defines the custom insights and criteria used to analyze documents. It appears if you selected *Custom* as the Analysis Type during project setup.

Fill out the following:

- **Title**—enter a unique title (name) for the custom insight (up to 30 characters). The title will act as the label for displaying results in Analysis Results tables, Custom Projects card in the Viewers, and the published output. Or use one of the Suggested Insights listed in the panel. See [Adding Custom insights and prompt criteria on page 30](#) for more information.
- **Prompt instructions**—enter clear and descriptive instructions in the prompt text box (up to 1000 characters) that will guide aiR Review in analyzing documents. See [Best practices for writing Custom insight prompts on page 18](#) for more information.
- **Field Type**—select the output format for this insight's result: Fixed Length Text (default), Single Choice, Date, or Whole Number. This selection determines how the result is validated, how it displays in the Analysis Results table and Viewer, and which native Relativity field type is published on the Document object. Field Type cannot be changed after the first job run for this insight.
 - **Choices** (for Single Choice only)—when Field Type is set to Single Choice, define 2–5 choices for aiR to select from. Click **+ Choice** to add another choice to the list. Add choices in the desired order as they cannot be reordered later.
For each choice, enter a unique, case-insensitive name in the **Name Choice [#]** field (up to 30 characters), and criteria in the **Describe Choice [#]** text box (up to 300 characters) explaining when a document matches that choice.
 - If a result does not conform to the selected Field Type, the cell is marked as an error. For example, if the result is a non-numeric response for a Whole Number field or is a value outside the defined choices for a Single Choice field. See [Handling document errors in Custom analyses on page 37](#) for more information.

To add more custom insights (up to five): Click the **+ Insight** button and repeat the above process for each one.

To delete a custom insight from panel: Click the vertical ellipses next to the desired title and select **Delete**.

8.4 Using prompt kickstarter (Draft with AI)

aiR Review's prompt kickstarter (Draft with AI option) lets you automatically generate initial prompt criteria for Relevance, Key Documents, and Issues analysis projects using existing case documents, such as review protocols, requests for production, complaints, or other foundational case documents, instead of doing so by scratch. By uploading up to 10 documents (with a total character count of up to 300,000), aiR Review analyzes them to fill out the


needed initial prompt criteria. This makes it easy to launch a new project with minimal effort. See [Job capacity, size limitations, and speed on page 48](#) for more information on document and prompt limits.

You can repeat this process as needed to refine the prompt criteria before starting the first job analysis. After the analysis begins, the Draft with AI option is disabled.

- Notes:**
- There are no additional charges to use prompt kickstarter.
 - Prompt kickstarter does not apply to Confidential Business Information or Custom analyses projects.
 - To use prompt kickstarter for *Issues* projects, be sure to create an Issues field in the workspace beforehand starting the project.


To use prompt kickstarter:

1. On the prompt criteria panel, click the **Draft with AI** button. You can also click the three vertical ellipses **More** () icon and select **Draft with AI**.



2. Upload content using any of the methods below for AI to draft prompt criteria. The total number of selected files and uploaded characters counters are displayed at the top and update automatically as changes occur.

Note: The maximum number of documents is 10, with a combined total character limit of 300,000.



- **Drop file here or browse files**—drag or upload document files. Supported file formats are .txt, .docx, and .pdf.
 - **Add text manually**—click to manually add text or to copy and paste from another source. You can also use this option to include additional case context for the AI assistant to use.
3. Select the **Document Type** for each uploaded file. Options include: Review Protocol, Request for Production, General Case Memo, Complaint, Key Document, and Other. If you select Other, enter a document type description in the Custom Type box. The number of characters in each file appears below the filename to help keep

track of the 300,000 character limit.

Big Thorium - Bribery Case (v5) - Complaint.pdf 4,614 characters	Document Type *	Complaint	ⓧ
Big Thorium - Bribery Case (v5) - Review Protocol ... 12,433 characters	Document Type *	Review Protocol	ⓧ
Big Thorium - Bribery Case (v5) - RFP.docx 7,025 characters	Document Type *	Request for Produ...	ⓧ
New insight to the case copied from an email ✎ 719 characters	Document Type *	Other	ⓧ
	Custom Type *	Additional case contex	

4. Repeat steps 2-3 to upload more files.
5. To delete a file from the list, click its corresponding **circle X** icon.
6. To use prompt kickstarter to automatically draft **Issue descriptions** for your prompt criteria, do the following:
 - a. For *Issues* projects, use the **Select Issues Field** drop-down list below the documents list. For *Relevance* or *Relevance and Key* projects, enable the **Include Issues in Relevance Criteria** toggle below the documents list.

Note: Case issues will be included as part of the criteria for a Relevance analysis. However, documents will not be labeled with issue choices as part of a Relevance analysis.

- b. Set up the **Issues** field by selecting a single-choice or multi-choice field to indicate the issues involved in the case. Each option should identify a separate issue for which prompt kickstarter will generate a description. See [Issues tab on page 62](#) for more details on the fields. For best results drafting Issue descriptions:
 - Name each issue choice concisely and descriptively to help the AI model understand and generate accurate outputs. For example, using the Choice Criteria name, *Anti-bribery policy in Acme Company internal communication* will help the AI generate a description such as *Documents, policies, or internal communications describing or outlining Acme Company's corporate policies related to ant-bribery or anti-corruption with public officials*.
 - Correct any misspellings.
 - Do not use a plain number for the Choice Criteria name, such as 12.
 - Upload only relevant documents, and if necessary, include additional case context using the **Add text manually** option. The AI assistant uses contextual cues from the uploaded files. Providing unrelated or contradictory materials may degrade the quality of the generated output.

Note: If the AI assistant is unable to find relevant information in the uploaded documents for at least one of the selected issues, a message appears. You can try to manually enter the choice criteria description. Alternatively, you can run Draft with AI again after updating the choice name or uploading more content-rich documents.

7. Click **Draft** to begin automatically drafting prompt criteria from the uploaded documents or the manually entered content. Typically results display within 1-2 minutes.

Note: You cannot run document analysis during the drafting process.

8. Review and edit the draft prompt criteria entered in the available tabs. Click **Save** to keep the changes or **Discard** to delete them.

The screenshot shows the Relativity aiR interface with three tabs: 'Case Summary' (checked), 'Relevance' (star), and 'Key Documents' (star). The 'Matter Overview' tab is active, showing a text area with a draft prompt criteria. Below the text area is a button labeled '* My own edits to the prompt'. The 'People and Aliases' tab is also visible, showing a list of individuals. At the bottom, there is a status bar indicating 'You have unsaved changes' and '11,074 characters left', along with 'Save' and 'Discard' buttons.

9. Repeat these steps with other documents and information as needed until the desired set of prompt criteria is achieved.
10. Review the available tabs and enter any required information on each.
11. Once the desired prompt criteria is set, click **Start Analysis** to analyze documents. For more information on analyzing documents, see [Running the analysis on page 70](#)

Note: The Draft with AI option is no longer available after the analysis process begins.

8.5 Editing and collaboration

If two users edit the same prompt criteria version simultaneously, the most recent save will overwrite previous changes. For this reason, it can be beneficial to have only one user edit a project's prompt criteria at any given time. Defining distinct roles for users when updating prompt criteria may help streamline the process.

To facilitate team collaboration outside of Relativity aiR, use the Export option. It exports the contents of the currently displayed prompt criteria to an MS Word file. For more information, see [Exporting prompt criteria on page 69](#).


Note: Exporting Custom Analyses results directly from the Analysis Results tab is not currently supported. Refer to [Exporting Custom analyses results on page 37](#) for the necessary steps.

8.6 How prompt criteria versioning works

Each aiR 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 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 down arrow next to the version name in the project details strip. From there, select the version you want to see.



The screenshot shows a dashboard header with a 'Validate' button, a project name 'BigThorium Bribery' with '10 analyzed | 10 total', and a version selector 'Version 4' with 'Ran 9/16/2025'. A red box highlights a dropdown arrow next to the version name. Below this is a table with five rows, each representing a different version of the project.

Validate	BigThorium Bribery 10 analyzed 10 total	Version 4 Ran 9/16/2025
Develop	BigThorium Bribery 21 analyzed	Version 4 Ran 9/11/2025
Develop	BigThorium - Not Relevant 21 analyzed	Version 3 Ran 9/11/2025
Develop	BigThorium Bribery 41 analyzed	Version 2 Ran 9/11/2025
Develop	BigThorium Bribery 50 analyzed	Version 1 Ran 9/11/2025

8.7 How prompt criteria versions 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 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 Review Jobs tab. For more information, see [Managing aiR Review jobs on page 99](#).

8.8 Revising the prompt criteria

After running an aiR Review job for the first time on the sample set, the initial results on the dashboard can be used as feedback for refining the prompt criteria. The cycle of examining the results, revising the prompt criteria, then running a new job on the sample documents is known as iterating on the prompt criteria. Refer to [Best practices on page 15](#) for more information. Also see [Job capacity, size limitations, and speed](#) for details on document and prompt limits.

In particular, ask the following questions about each document:

- Did aiR Review and the human reviewer agree on the relevance of the document?
- Read the aiR 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 and rephrase the prompts. 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 to clarify meaning and intent.

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

General steps for revising prompt criteria:

1. Revised the prompt criteria to address any weak points.
2. Run the analysis again.
3. Continue refining the prompt criteria until results show reasonable alignment with human coding decisions for the test documents in the sample.
4. To run validation on the prompt criteria in Review Center to compare it to aiR Review results before using it on a larger target population, refer to [Setting up aiR Review prompt criteria validation on page 88](#) for more information.

Note: aiR 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 Review will not be able to match that relevance decision.

8.9 Exporting prompt criteria

You can export the contents of all available criteria tabs for the currently displayed prompt criteria to an MS Word (.docx) file. This can be helpful for reviewing criteria, saving it to use later, and collaborating with others on it.

1. In the prompt criteria panel, click the three vertical ellipses **More** () icon next to the double left arrows Collapse (<<) icon.

Matter Overview ? *

This is an investigation into a potential bribery. BigThorium, a company specializing in sustainable energy solutions, was issued a Request for Proposal (RFP) by the City of Atlantis for a municipal energy project aimed at enhancing the city's energy efficiency and reducing its carbon footprint, and ultimately won the contract. It is alleged that employees of BigThorium bribed city officials to win the contract.

People and Aliases ?

John Maxwell (Sales Representative, BigThorium): Engages in discussions about offering gifts and incentives to secure the contract. Sarah Thompson (VP of Sales, BigThorium): Coordinates the strategy and approval of incentives, including the purchase of a Rolex watch. Robert Williams (City Treasurer,

2. Select **Export Prompt**. The exported file is saved to your default browser download folder.

Audit logs track all export actions from your project.

9 Running the analysis

After [setting up the project](#) and [developing the prompt criteria](#), you are ready to run the analysis against the prompt criteria. For Custom analyses projects, refer to [Running a Custom analyses job on page 33](#).

Note: Authorized administrators can configure job size limits per user group at the instance level to help control and manage the size of jobs that users execute. For more information on setting job sizes, see [Setting aiR Review job size limits on page 100](#).

To run the analysis:

1. On the upper right of the dashboard, click **Analyze [#] documents**.

2. Review the confirmation summary:

aiR Review



Review and confirm the following estimates before starting this analysis. You can review progress in the aiR Review Jobs page. The prompt criteria and metrics will be saved as **Kickstart Case Version 1**.

Total Docs	77	Email Notification	☐
Est. Time to Start	0 minutes		
Est. Run Time	2 minutes		
aiR Version	2026-04-14-DE00FC		

The performance of Relativity aiR Review has been tested on primarily English language documents. Use of aiR for Review with non-English language documents may result in unexpected or unverifiable results.

Start Analysis

Cancel

- **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.
- **Email Notification**—enable the toggle if you want to send email notifications when a job completes, fails, or is canceled. Your email address is automatically entered in the text box. Proceed with entering more recipient email addresses, separated by commas or semicolons. Email notifications are only sent when the toggle is enable.
- **aiR Version**—a unique ID that identifies the underlying aiR system used to process each job. This ID does not map directly to a single large language model version used for the job. You can also find the aiR Version on the aiR Review Jobs list. See [Managing aiR Review jobs on page 99](#).

3. Click **Start Analysis**.

Once the analysis job starts, the results begin to appear beside each document in the Analysis Results panel of the dashboard. For more information on analyzing the results or seeing results in the Viewer, see [Analyzing aiR Review results on the next page](#).

The Analyze button is disabled during the analysis process, and a **Cancel** option appears. When the job completes, the Analyze button is active again.

Note: If you try to run a job that is too large or when too many jobs are already running, an error appears. 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, size limitations, and speed on page 48](#) and the Community article [aiR Review - Error Message of Text is too long, but Extracted Text Size is within limits](#).

After the first analysis completes, use the results to fine-tune the prompt criteria. For more information on:

- the dashboard, see [Navigating the aiR Review dashboard on page 40](#).
- the fine-tuning process, see [Revising the prompt criteria on page 68](#)
- the results fields, see [Analyzing aiR Review results on the next page](#).

10 Analyzing aiR Review results

When aiR 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.

Additionally, you can export the results to view data offline, share it with others, and compare information. See [Exporting analysis results on page 80](#) for details.

10.1 How aiR Review analysis results work

After analyzing a document, aiR Review provides a Prediction, a Score with recommendations on how to categorize the document, and the supporting Rationale for prediction. You can use this information to help update and improve your prompt criteria.

This analysis has several parts:

Analysis Results

Summaries

1

- 112 of 112

200

per page

#	Control Number	Reviewer Coded	aiR Prediction	aiR Score	aiR Rationale	aiR Consideration	aiR Citation
	Filter	Filter	Filter	Filter	Filter	Filter	Filter
1	BTv5_000005028	Responsive	Borderline Responsive	2	This document contains communication between a City of Atlantis official and BigThorium employees, discussing compliance with local regulations and financial implications, which could be relevant to the allegations of unethical practices and bribery in securing contracts. The mention of aligning with local compliance requirements and briefing the sales team on protocols may hint at processes related to securing contracts.	The document does not explicitly mention the RFP, bribery, or unethical practices, and the discussion appears to focus on emergency preparedness rather than the municipal energy project contract. It is possible that this communication is unrelated to the allegations in the case.	Given the re is essential compliance
2	BTv5_000005032	Non-Responsive	Not Responsive	1	Das Dokument enthält keine relevanten Informationen zu den Kriterien der Relevanz, da es sich ausschließlich auf allgemeine Risikomanagement- und Notfallplanungsstrategien konzentriert und keine spezifischen Hinweise auf die RFP, Bestechungsvorwürfe oder Interaktionen mit der Stadt Atlantis enthält.	Es könnte argumentiert werden, dass ein verbessertes Risikomanagement indirekt mit der Fähigkeit von BigThorium zusammenhängt, sich auf Ausschreibungen wie die RFP der Stadt Atlantis vorzubereiten. Allerdings fehlen direkte Bezüge zu den relevanten Themen.	
3	BTv5_000005037	Non-Responsive	Borderline Responsive	2	This document includes communication between a BigThorium employee and an Atlantis City official, which falls under the responsive criteria of 'Any conversations that	The content of the email does not directly address the allegations, the RFP, or any relevant business dealings. It is purely social and unrelated to the case, which could argue	From: Samu <sjenkins@ To: Robert V <robert.will

- **aiR Prediction**—the relevance, key, or issue label that aiR predicts should apply to the document. See [Predictions versus document coding on the next page](#)
- **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 page 75](#).
- **aiR Rationale**—an explanation of why aiR chose this score and prediction.
- **aiR Consideration**—aiR is asked to argue against itself, identifying what assumptions it is making or additional facts that could reverse the prediction.
- **aiR Citation**—excerpts from the document that support the prediction and rationale. For the Relevance and Relevance and Key Documents analysis types, a maximum of five citations will be presented per document. For the Issues analysis type, only one citation will be shown for each issue.

Note: In general, citations are expected for borderline and relevant documents, and are left empty for non-relevant documents and documents that do not 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.

10.1.1 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. It is not coding the documents or writing to the coding 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 Review results into a larger review workflow, see [Using aiR Review with Review Center on page 110](#).

10.2 Issues Field versus Issues analysis type for aiR Relevance Analysis::Issues

When using the Issues Field on the Relevance tab for Relevance or Relevance and Key Document projects, aiR does not make Issue predictions or coding and does not write to *aiR Relevance Analysis::Issues* on the Documents list. If you want aiR to provide this information, you must run another project using **Issues** as the analysis type on the project set up dialog. See [Setting up an aiR Review project on page 51](#) for more information.

The Issues Field selections on the Relevance tab are used to assess document relevance, but aiR does not assign issue choices to documents. For example, you could tell aiR that documents containing any of the defined issues are relevant, but the results will not display for *aiR Relevance Analysis::Issues*. For more information on the Issues Field, see [Relevance tab on page 60](#).

Note: The *aiR Relevance Analysis::Issues* column on the Documents list may show results for jobs run before the early August 2026 system fix that removed Issue Field results from this field. For jobs run after the fix, the **aiR Relevance Analysis::Issues** column is empty.

10.3 Document summaries

The Summaries tab in the document grid displays a descriptive topic and summary of each document. These remain the same for each document across all projects, except Custom Analysis, which does not generate summaries and topics. They are not affected by the project settings, analysis type, prompt criteria, or other documents.

The generated columns are:

- **aiR Topic**—the descriptive topic aiR generated for this document.
- **aiR Content Summary**—the summary aiR generated for this document's content.

These columns may populate either faster or slower than the analysis results.

For more information on document summaries, see [Generating document summaries on page 9](#).

Notes:

- Document topics and summaries are generated automatically when you run an analysis type, except Custom Analysis, which does not produce them. They cannot be generated independently of running an analysis project.
- Document summaries do not affect the number of aiR Units used. For more information on aiR Units and billing, see [aiR Units in Relativity aiR standard workspaces](#) on the Relativity documentation site and [aiR Units, Thresholds, and Included Products](#) on the Community (valid login credentials required).
- aiR Review document summaries differ from witness summaries in Case Strategy, which support fact extraction and case-building workflows. When generating facts, Case Strategy summarizes a document based on its relevance to the case prompt. See Case Strategy documentation for more information on Case Strategy summaries.

10.3.1 Documents with an empty summary field

If the document has minimal content or unintelligible text, the aiR Topic will be “Poor Quality Content,” and it will not have an aiR Content Summary. Re-analyzing these documents will return a similar result.

If there is no content for either the aiR Topic or the aiR Content Summary, this means an error occurred when generating them. These errors are rare, and they do not affect whether the main document analysis succeeded. For more information, see [Handling document errors on the next page](#).

10.4 Custom insight analyses results

Once the analyses complete, each custom insight displays as individual columns in the Analysis Results table and in a dedicated Custom Projects card within the Viewer.

Notes:

- You cannot filter Custom analyses results at this time. This applies to the Analysis Results table only. Once results are published to the Relativity Document Object (RDO) by running the Apply step (see [Analyzing aiR Review results on page 72](#)), they can be filtered and sorted according to each insight’s selected Field Type.
- Custom analysis results are initially only available in your aiR Review project and the Viewer. To publish them to the RDO fields, run the Apply step (see [Analyzing aiR Review results on page 72](#)). After the apply process is complete, the results are available for export (see [Exporting Custom analyses results on page 82](#)).
- Custom analyses results do not enforce RDO document-level permissions in the same way as individual document viewing. This means users may see aggregated dashboard results from documents they cannot access directly. To manage or restrict selective document access, you can configure user permissions to control who can view or create Custom Analysis projects. See [Permissions on page 12](#) for details.

The screenshot displays the Relativity aiR interface. On the left, the 'Custom' sidebar shows 'Insight Suggestions' and 'Custom Insights'. Under 'Custom Insights', there is a 'Summarize' insight selected, with a description: 'Summarize this document in exactly 2-3 sentences. Focus on: 1. What the document is and who created it. 2. The most important finding, decision, or event. 3. Any required next steps or actions. Do not include background or methodology. Be specific: include names, dates, and locations.' On the right, the 'Analysis Results' table is shown. It has columns for '#', 'Control Number', and 'Summary'. The table contains 7 rows of results, each with a control number and a brief summary of the document content.

#	Control Number	Summary
1	BT-V-000000018	This document is an email from Francis Han, CEO of BigThrum, sent on 10/10/2017 to Linda Patel regarding the upcoming board meeting. Summary: Francis Han is requesting feedback on the upcoming board meeting. The next step is for Linda Patel to provide feedback if necessary.
2	BT-V-000000019	This is an email from Linda Patel to Francis Han dated June 18, 2017. It is a plan to visit a new site and eat at a local place, which Francis Han is supposed to coordinate and finalize the restaurant plan.
3	BT-V-000000020	This is an email thread from August 18, 2017, initiated by Francis Han, CEO of BigThrum, regarding the upcoming board meeting. Summary: Francis Han is requesting feedback on the upcoming board meeting. The next step is for Linda Patel to provide feedback if necessary.
4	BT-V-000000021	This document is an email exchange from June 23, 2017, between Kayla Piro, BigThrum, discussing new company policy. Summary: Kayla Piro is requesting feedback on the upcoming board meeting. The next step is for Linda Patel to provide feedback if necessary.
5	BT-V-000000022	This document is an email from Francis Han, CEO of BigThrum, regarding the upcoming board meeting. Summary: Francis Han is requesting feedback on the upcoming board meeting. The next step is for Linda Patel to provide feedback if necessary.
6	BT-V-000000023	This document is an email exchange between Kayla Piro, BigThrum, regarding scheduling a meeting on April 13, 2017, at 10:00 AM. Summary: Kayla Piro is requesting feedback on the upcoming board meeting. The next step is for Linda Patel to provide feedback if necessary.
7	BT-V-000000024	This email thread from April 17, 2017, involves Laura Smith, Senior Analyst.

Click on the **Control Number** link to open the document in the Viewer. See Viewer documentation for more information on using the Viewer.

If an error displays for an insight, see [Handling document errors in Custom analyses on page 37](#) for guidance. A result that does not conform to its Field Type (for example, a non-numeric response returned for a Whole Number field) is treated as a cell-level error using this same handling. Errored cells are automatically included in the next re-analysis of that insight.

If you modify custom insights, change the data source, or create a new project set, only the updated insights will re-run on the next analyses.

10.5 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. They cannot receive a normal score because they were not properly analyzed.

Below are the aiR Review scores:

Score	Description
4	Highly Relevant: 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.
3	Relevant: The document is predicted relevant to the issue. Citations show the relevant text.
2	Borderline Relevant: The document is predicted borderline relevant. aiR found some content that might relate to the case or issue. It is expected to have citations.
1	Not Relevant: The document is predicted not relevant. aiR did not find any evidence that it relates to the case or issue.
0	Junk: <i>(For aiR versions below 2026-04-14)</i> The document contains no useful information or is considered “junk” data, such as system files, an empty document, or sets of random characters. Refer to the aiR Review Jobs tab for the aiR version used for the job. Highly Not Relevant: <i>(For aiR versions 2026-04-14 and above)</i> The document either contains no useful information or is predicted highly not relevant to the case or issue. Refer to the aiR Review Jobs tab for the aiR version used for the job.
-1	Error: The document either encountered an error or could not be analyzed. For more information, see Handling document errors below .

10.6 Handling document errors

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 Review total document count.

Note: For information on errors when using Custom analyses, see [Handling document errors in Custom analyses on page 37](#).

The possible error messages are:

Error message	Description	Retry?
Completion is not valid JSON	The aiR system encountered an error.	Yes
Failed to parse com-	The aiR system encountered an error.	Yes

Error message	Description	Retry?
pletion		
Document text is empty	The extracted text of the document was empty.	No
Document text is too long	The document's extracted text was too long to analyze.	No
Document text is too short	There was not enough extracted text to analyze in the document.	No
Model API error occurred	A communication error occurred within the aiR system. This is usually a temporary problem.	Yes
Uncategorized error occurred	An unknown error occurred.	Yes
Ungrounded citations detected in completion	The results for this document have a chance of including an ungrounded citation. For more information, see Ungrounded citations below .	Yes
Missing citations	The results for this document do not contain any citation, while citations are expected for borderline and relevant documents.	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 Review cannot process.

Note:

In rare cases, there may be an error generating the document's topic or summary. As long as the main document analysis succeeded, these are not considered errored documents. An error in the aiR Topic or aiR Content Summary field does not affect any of the other results. You can re-run the document to try for better results, but the document aiR Units will be consumed normally.

10.7 Variability of results

Due to the nature of large language models and how they analyze each document individually, 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. Similarly, when running prompts on Custom Analyses, tweaks to prompts between different iterations will result in different responses for those insights. For more information on variability of results, see the Community article, [Why might aiR Review \(a4R\) results change even when using the same prompt?](#)

10.8 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.

10.9 Viewing and filtering analysis results

aiR Review analysis results can be viewed from the dashboard or the Viewer.

10.9.1 Viewing results from the aiR Review dashboard

Within a project, you can view results using the aiR Review dashboard. This dashboard includes not just results fields, but may also include 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 Review Projects** tab. For detailed information on the dashboard layout, see [Navigating the aiR Review dashboard on page 40](#).

10.9.2 Viewing results for individual documents from the Viewer

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

Notes:

- You will only see analysis highlights if you have the necessary permissions. Without these, the aiR Review Analysis icon does not display. For more information, refer to [Permissions](#).
- 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. See View documentation on Extracted Text Viewer for more information.

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

1. Document topic and summary
2. Project name and version of the analysis
3. Field choice for which the document was being analyzed
4. aiR's Prediction
5. aiR's Rationale and Considerations
6. Citations found in the document (click each to view it in the document)

H

aiR for Review Analysis

1

Summaries

▼

aiR Topic

Email Exchange on Nuclear Plant Regulations

aiR Content Summary

This is an email exchange between Robert Williams of Atlantis City Council and Samuel Jenkins of BigThorium discussing regulatory and licensing requirements for constructing a nuclear power plant. The correspondence clarifies necessary guidelines including environmental assessments, public consultations, and compliance with local, state, and federal regulations.

2

Viewing BigThorium Version 1

3

Responsive

4

Borderline

▼

5

✓

Rationale

The document involves communication between Robert Williams, a city official, and Samuel Jenkins from BigThorium regarding regulatory and licensing requirements for a nuclear power plant construction. While it does not directly address the RFP for the municipal energy project or bribery allegations, it pertains to BigThorium's process for engaging with city officials and adhering to regulatory guidelines, which could be tangentially relevant to the investigation.

?

Considerations

The document does not directly address the RFP, bribery, or the purchase of luxury goods. It focuses on regulatory compliance for a nuclear power plant, which may not be directly related to the municipal energy project in question. The lack of direct connection to the bribery allegations weakens its relevance.

6

1.

In response to your inquiry regarding the regulatory and licensing requirements for the nuclear power plant construction, I would like to confirm that the city has established clear guidelines that must be followed.

Page 2

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 Review Jobs tab.

Relativity aiR

aiR Review Guide

78

10.9.2.1 Citations and highlighting


For the Relevance and Relevance and Key Documents analysis types, a maximum of five citations will be presented per document. For the Issues analysis type, one citation will be shown for each issue.

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 Review Analysis pane.

Citation colors

The highlight colors depend on the type of citation:

- **Relevance citation**—orange.
- **Key Document citation**—purple.
- **Issue citation**—color is 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.

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

Citation order

The results in the aiR 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.

10.9.2.2 Adding aiR Review fields to layouts

Because of how aiR 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.

10.10 Exporting analysis results


The Export Results option in Mass Ops lets you easily export your analysis results to an Excel (.xlsx) file so you can view data offline, share it with others, and compare information.

Notes:

- To export results for Document Summaries and Custom analyses, see [Exporting document summaries on the next page](#) and [Exporting Custom analyses results on page 82](#).
- If you want to send analysis results to Review Center, see [Using aiR Review with Review Center on page 110](#).

Use the steps below to export analysis results data to an Excel (.xlsx) file.

1. On the **Analysis Results** tab, select up to 1,000 documents for export by checking individual boxes or applying filters. If you do not select or filter, all documents currently visible in the results view are exported, up to the 1,000 limit.
2. Select **Export Results** from the mass operations list at the bottom of the grid.



- On the **Exporting analysis results** modal, choose the desired **Data Source** from the list. The currently selected data source defaults.

Exporting analysis results

Data Source ★

Significant Documents ▼

Prompt Criteria Version ? ★

Version 4 (x) Version 3 (x) Version 2 (x) (x) ▲

3 / 4 selected

Version 1

✓ Version 2

✓ Version 3

✓ Version 4

- Choose up to three versions from the **Prompt Criteria Version** list.
- Click **Export**.
Once confirmed, a progress modal appears in the bottom-right corner of the application window showing the export status.

The Excel file automatically saves to the default download folder configured in the user's web browser.

10.10.1 Export capacity

- A maximum of 1,000 documents can be selected for exporting, regardless of the number of prompt criteria versions selected, to ensure performance and stability.
- A maximum of three prompt criteria versions can be selected for exporting to avoid overloading Excel with excessive sheets.

10.10.2 Exported file structure

- Each analysis type is exported to a separate sheet in the Excel file. Please note that if Relevance and Key Document fields in a project have the same name, their results export to the same sheet.
- Up to 10 Excel sheets can be generated for Issue Analysis results, since 10 is the limit for number of issues analyzed in aiR Review project.
- Field names reflect the criteria selected during analysis setup.
- Data on the Insights tab is currently not included in the exported file.

10.10.3 Exporting document summaries

When aiR Review generates document topics and summaries, these are the same for each document across all projects. They are not attached to a specific project or prompt criteria version, and they are not included in the main analysis results export.

You can export document topics and summaries by creating a document view that includes those fields, then exporting data from the Documents tab.

To add topics and summaries to a view:

1. Create a new document view or enter edit mode on an existing one. For instructions, see [Views in the Admin guide](#).
2. Add the following fields:
 - aiR Summaries::Topic
 - aiR Summaries::Summary
3. Add any other fields you want in the export. For a list of aiR Review-related fields, see [Creating document views on page 103](#).
4. Click **Save**.

To export topics and summaries using the view:

1. At the top of the Documents tab, select the view that contains the Topic and Summary fields.
2. Select the documents whose summaries you want to export.
3. Under the mass actions drop-down, select **Export to File**.
4. Click **Export**.
The file will begin downloading.

For more information on document summaries, see [Document summaries on page 73](#).

10.10.4 Exporting Custom analyses results

Exporting Custom analyses results directly from the Analysis Results tab is not currently supported.


To export these results, you must publish the results to the Relativity Document Object using the Apply workflow (see [Applying the results and publishing to the Document Object on page 35](#)). Following publishing, you may export your results from the Document List.

For more information on Custom analyses, see [Custom analyses on page 27](#).

10.11 Saving selected documents as a list

To save a group of documents in the Analysis Results grid as a list, follow the steps below.

1. Select the box beside each individual document in the grid that you want to add to the list.
2. Click **Save as List** from the mass operations list at the bottom of the grid.



3. On the Save as List modal, enter a unique **Name** for the document list.

Save as List 3 Documents

Name*

Notes


4. Enter any **Notes** in the text box to help describe the list.
5. Click **Save**.
A green confirmation banner appears.

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

10.12 Creating a saved search

To save a group of documents in the Analysis Results grid as a saved search, follow the steps below.

1. Select the box beside each individual document in the grid that you want to add to the saved search.
2. Click the **Create Saved Search** option from the mass operations list at the bottom of the grid.



3. Enter a unique **Name** for the saved search.

Create Saved Search

Name*

4. Click **Save**.

A green banner displays confirming the save and providing a link to the saved search in the Document list.

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

11 aiR Review prompt criteria validation

Using aiR Review and Review Center in tandem, the prompt criteria validation process gathers statistical metrics to check whether the prompt criteria are effective and defensible before using them on a larger data set. You can set up a smaller document sample, oversee reviewers, compare aiR's relevance predictions to actual coding results, and iterate on prompt criteria as needed.

Note: This functionality is currently only enabled for the Relevance and Relevance & Key analysis types. See [Choosing an analysis type on page 50](#) for descriptions of these analysis types.

11.1 Prerequisite

To run validation with aiR Review, the Review Center application must be installed, and users must have the appropriate permissions to each application. See [aiR Review permissions](#) and the Review Center Guide for more information.

11.2 Benefits of validation


Using aiR Review validation helps streamline document review by improving accuracy, reducing rework, boosting collaboration, and ensuring consistent results through AI validation and performance metrics.

The following points highlight key benefits of implementing robust validation and collaboration strategies with aiR Review.

- **Demonstrate defensibility with measurable evidence:** Provide concrete evidence of aiR Review's accuracy by calculating precision, recall, and elusion directly in Relativity aiR.
- **Prevent costly rework:** Avoid rework and identify early on if your prompt criteria should be adjusted before it is used at scale.
- **Enhance collaboration between humans and AI:** Work in sync with aiR Review by pinpointing where AI predictions diverge from those of SMEs and adjusting prompt criteria as needed.
- **Ensure consistency and foster trust:** Build confidence with clients and stakeholders. Validation is intended to help improve consistency of outcomes, reducing worries about unpredictability.

11.3 How validation fits into aiR Review

aiR Review leverages the prompt criteria across a three-phased workflow:




1. **Develop**—user writes and iterates on the prompt criteria and tests on a small document set until aiR's recommendations align sufficiently with expected relevance and issue classifications. See [Developing prompt criteria on page 58](#) for more information.
2. **Validate**—user leverages the integration between aiR Review and Review Center to compare results and validate the prompt criteria. See [Setting up aiR Review prompt criteria validation on page 88](#) for more information.
3. **Apply**—user applies the validated prompt criteria on much larger sets of documents. See [Applying the validated results or developing different prompt criteria on page 90](#) for more information.

The prompt criteria validation process covers phase 2, Validate.

11.4 High-level prompt criteria validation workflow

This diagram details the steps (in orange) for prompt criteria validation, which you can perform once the develop phase is complete.



The Validate and Apply phases involve several steps that cross between aiR Review and Review Center:

Process Flow	Application Used
1. Identify target review set & choose validation sample settings. Set up the validation sample by choosing the sample size, desired margin of error for the validation statistics, and other settings. See Setting up aiR Review prompt criteria validation on page 88 for details. 2. Run aiR Review on the sample to receive predictions. When a validation sample is created, a	aiR Review

Process Flow	Application Used
corresponding Review Center queue is automatically created for their review. Run the sample documents through aiR Review to obtain the relevance predictions that will be used for comparison with the manual human reviews in Review Center. See Applying the prompt criteria on page 57 for details.	
3. Manually review and code the documents in the validation sample. Human reviewers code the documents in the sample for comparison to the aiR predictions. For the sake of validation, the human coding decisions are considered "correct." The Review Center dashboard tracks reviewers' progress and compares their choices to aiR. See Prompt Criteria validation in Review Center on page 92 for details. 4. Evaluate built-in	Review Center

Process Flow	Application Used
statistical results of the validation. After human reviewers finish reviewing the validation sample, final validation statistics display comparing their results with aiR Review's AI predictions. These results are then evaluated. See Reviewing validation statistics on page 93 and Prompt Criteria validation statistics on page 96 for details. 5. Accept or reject the validation results. After reviewing the results, decide whether to accept or reject them. If the results are accepted, the validated prompt criteria can be used for all documents in the target data source, and the process moves to aiR Review to apply the criteria to the larger data set. If they are rejected, go back to the iteration phase and adjust the prompt criteria	

Process Flow	Application Used
on additional document samples. See Accepting or rejecting results on page 94 for details.	
6. Optionally apply or improve the prompt criteria. Return to aiR Review to either run the validated prompt criteria on larger sets of documents, or to improve the prompt criteria and validate again. See Applying the validated results or developing different prompt criteria on page 90 for details.	aiR Review

11.5 Setting up aiR Review prompt criteria validation

After developing the prompt criteria, you can run it through the validation process to gather metrics and evaluate their effectiveness and support defensible workflows, before using them on a larger data set. Begin by configuring the validation sample settings in aiR Review. Based on these parameters, aiR will pull the random sample and create a Review Center queue for their evaluation. After humans review and code the documents in Review Center, you can return to aiR Review to run the accepted prompt criteria on the larger document set.

Note: Prompt criteria validation is currently only available for the Relevance and Relevance & Key analysis types. See [Choosing an analysis type on page 50](#) for descriptions of these analysis types.


The general workflow involves:

- 1) Setting up the validation sample by choosing the sample size, desired margin of error for the validation statistics, and other settings.
- 2) Running aiR Review on the sample to receive relevance predictions.
- 3) Going to Review Center to begin the human review and coding of the validation sample.
- 4) Returning to aiR Review to run the accepted prompt criteria on the larger document set if desired.


For information on coding documents and viewing statistics in Review Center, refer to the Review Center Guide.

Follow the steps below to use aiR Review for prompt criteria validation:

1. Choose the desired project set version number to validate.
2. Click the project set **+ sign**.



3. Click **I want to validate the prompt criteria to a document population**.



4. Fill out the **Validation Settings fields** to determine the validation sample. For more information on the fields, see [Validation Settings fields on page 91](#).

Note: It is strongly recommended that the data source selected does not change after validation begins.

Setup Validation

This process will estimate prompt criteria accuracy and effectiveness against a large data set. A Review Queue will be created to review the sample. Control queue access and monitor review progress in Review Center.

Target Data Source for Review
The target population of documents for the full review. The validation sample is pulled from this data source. The sample will default to 1,000.

Select

Validation Reviewer Groups
The user group(s) who will review the validation sample. Reviewers access the sample and code documents via the Review Queues tab.

Select

Sample Settings
The sample is taken from the target data source for review and validation reviewers must code all sample documents. Select a richness estimate for the relevant documents in the data set. The Recall margin of error estimates how well the sample will estimate recall of the review. A lower margin will increase the sample size and elusion margin of error.

Size	Richness Estimate	Margin of Error Estimate (recall)
1000	25%	±5%

> Show Advanced Settings

All statistics are calculated for a 95% confidence interval.
[Learn more about validation settings.](#)

Create Validation Set & Queue Cancel

5. Click **Create Validation Set & Queue**. A validation sample set is generated in aiR Review, and a corresponding queue is established in Review Center for the prompt validation. The project set state transitions from *Develop* to *Validate*. For more information on Review Center queues, see the Review Center Guide.

6. Review the confirmation summary modal and provide email addresses, if needed, then click **Start Analysis**.

aiR Review



Review and confirm the following estimates before starting this analysis. You can review progress in the aiR Review Jobs page. The prompt criteria and metrics will be saved as **Kickstart Case Version 1**.

 **Total Docs**

77

 **Est. Time to Start**

0 minutes

 **Est. Run Time**

2 minutes

 **aiR Version**

2026-04-14-DE00FC

 **Email Notification** 

☐

Start Analysis

Cancel

The performance of Relativity aiR Review has been tested on primarily English language documents. Use of aiR for Review with non-English language documents may result in unexpected or unverifiable results.

The aiR Review predictions appear in the Validate Metrics tab. The banner lets you know that a validation queue has been created in Review Center to facilitate review of the validation set there. The Analyze button at the top changes to Active Validation Review to let you know this project set is now in an active validation review process.

7. Click the **Go to Review Center** button or the **View in Review Center** link to begin the human review and coding of the validation sample in Review Center.
8. The Review Center application opens to the active queue. Refer to [Prompt Criteria validation in Review Center on page 92](#) for the steps to follow in Review Center.

Note: The Review Center application must be installed to run the validation workflow with aiR Review. Otherwise, an error message displays.

In Review Center, human reviewers manually code the sample set of documents. The Review Center dashboard generates statistics comparing the human coding to aiR's predictions.

After reviewers have finished coding all the documents in the sample, you can evaluate the Review Center results against the ones in aiR Review and decide whether to accept or reject them. Refer to [Applying the validated results or developing different prompt criteria below](#) for more information.

11.5.1 Applying the validated results or developing different prompt criteria

The following outlines the next steps in aiR Review, depending on whether the validation results were accepted or rejected in Review Center.

If accepted:

- A green check mark displays next to the project set version number in aiR Review, and it's labeled "Validation Accepted."
- To apply the validated prompt criteria on the larger document set, either click the **Run Target Population** button or click the project set **+ sign** and select **I want to apply the prompt criteria to the target document population**. Then go through the process of reviewing the confirmation summary and clicking **Start Analysis**. For details on running an analysis and project sets, refer to [Running the analysis](#) and [Using project sets on page 55](#).

Note: If the data source (saved search) exceeds the maximum document limit for a single job, it is necessary to split the data into multiple saved searches and apply each one separately to prevent errors. For information on document and prompt size limits, as well as overall job volume limits, see [Job capacity, size limitations, and speed on page 48](#).

- After running the accepted prompt criteria version on the target population in aiR Review, a green check mark and "Review Complete" notification appear next to the project set. No additional validations can be run on the project after validation is accepted.

If rejected:

- A red X displays next to the version number.
- Click the **+ sign** to create a new project set and repeat the prompt criteria development and validation process again.

11.5.2 Validation Settings fields

Enter information into the following fields:

- **Target Data Source for Review**—select the source from which documents are selected for the validation sample. Be sure the source is the entire set of documents that aiR Review will eventually analyze to ensure that the sample is representative of the whole review project, regardless of [maximum job size](#). Documents are randomly chosen from this group and may include those used during the prompt criteria development process.
- **Validation Reviewer Groups**—select the user groups who will review the validation sample. They will access the sample set and code the documents using the Review Center queue tab. We recommend using experienced reviewers for the validation. It can also be helpful if the reviewers were not involved in the prompt criteria creation and iteration process.
- **Sample Settings**—these settings determine the characteristics of the sample needed to achieve desired validation results.
 - **Size**—enter the number of documents for the validation sample. The default is set to 1000.
 - **Richness Estimate**—enter the estimated percentage of relevant documents in the target data source used to determine sample size and confidence levels. A default value of 25% is commonly used after data culling for production reviews. A lower value will increase the recommended sample size for a given margin of error.
 - **Margin of Error Estimate (Recall)**—click the plus and minus symbols to adjust the desired margin of error for the Recall statistic. The default of 5% MoE is a widely accepted threshold for defensible validation and recall calculations.

Advanced Settings

- **Display aiR Results in Viewer**—the toggle is enabled by default. When enabled, reviewers will see aiR Review's predictions in the Viewer while reviewing the validation sample. While seeing these decisions can help spot true positives, it may also bias reviewers towards agreeing with the AI-generated predictions. You can disable the toggle to hide predictions.
- **Email Notification Recipients**—enter email addresses, separated by semicolons, for individuals to be notified when manual queue preparation finishes, a queue becomes empty, or an error occurs during queue population.

11.6 Prompt Criteria validation in Review Center

When you validate aiR Review Prompt Criteria, the validation process compares aiR Review's AI-based relevance predictions to human coding decisions. Review Center calculates the validation statistics and helps you organize, track, and manage the human side of the coding process.

11.6.1 How Review Center fits into the validation process

When you create a validation set, there are several steps that cross between aiR Review and Review Center:

1. **Set up the validation sample**—choose the sample size, settings, and desired margin of error for the validation statistics.
2. **Run aiR Review on the sample**—this creates the relevance predictions that will be used for comparison.
3. **Code the sample using skilled human reviewers**—this records human coding decisions for comparison to the AI predictions. For the sake of validation, these are considered the "correct" decisions.
4. **Review the validation statistics**—these statistics measure any differences between the AI predictions and the human decisions. It also measures what percentage of the overall document set is likely relevant.
5. **Accept or reject the results**—this either confirms the Prompt Criteria as effective for use with a larger document set, or it returns you to the aiR Review dashboard for editing and improvement.
6. **Apply or improve the Prompt Criteria**—return to aiR Review to either run the Prompt Criteria on larger sets of documents, or to improve the Prompt Criteria and try again.

During validation, steps 3 through 5 take place in Review Center. For information on the other steps, see [Setting up aiR Review prompt criteria validation on page 88](#).

For a general overview, see [aiR Review prompt criteria validation on page 83](#).

11.6.2 Managing the coding process

After you create the Prompt Criteria validation queue, the coding process is similar to any other validation queue in Review Center. Reviewers code documents using the Review Queues tab, and administrators track and manage the queue through the main Review Center tab.

11.6.2.1 Administering the queue

As coding progresses, the Review Center dashboard displays metrics and controls related to queue progress. The main validation statistics will not appear until all documents have been coded and the validation process is complete. From the dashboard, the queue administrator can pause or cancel the queue, view coding progress, and edit some settings.

- For information on the Review Center dashboard, see [Monitoring a Review Center queue in the Review Center guide](#).
- For information on validation queue settings, see [Monitoring a validation queue in the Review Center guide](#).

Note: The statistics produced during Prompt Criteria validation are similar to the ones produced for a regular Review Center queue, but not identical. For more information, see [Prompt Criteria validation statistics on page 96](#).

11.6.2.2 Coding in the queue

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

During review:

- Have reviewers code documents from the sample until all documents have been served up.
- Sample documents that are coded outside the Review Queues tab will still count towards the validation statistics, but they might not be reflected in dashboard charts and tables.
- We strongly recommend coding every document in the validation queue. Skipping documents lowers the accuracy of the validation statistics.

For full reviewer instructions, see [Reviewing documents using Review Center](#) in the Review Center guide.

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

11.6.3 Reviewing validation statistics

When reviewers have finished coding all the documents in the queue, review the validation statistics. You can use these to determine whether to accept the validation results, or reject them and try again with a different set of Prompt Criteria. You can also see how these statistics change based on the cutoff score you select.

The statistics for Prompt Criteria validation include:

- **Elusion rate**—the percentage of documents that aiR predicted as non-relevant, but that were actually relevant.
- **Precision**—the percentage of documents that aiR predicted as relevant that were truly relevant.
- **Recall**—the percentage of truly relevant documents that were found using the current Prompt Criteria.
- **Richness**—the percentage of relevant documents across the entire document set.
- **Error rate**—the percentage of documents that received errors in aiR Review.


You can click on each statistic in the dashboard to see how it was calculated. The ranges listed below each statistic apply the margin of error.

Note: The exact criteria for whether to accept or reject may vary depending on your situation, but the goal is to have the AI predictions match the decisions of the human reviewers as closely as possible. In general, look for a low elusion rate and high recall.

For more information on how the statistics are calculated, see [Prompt Criteria validation statistics on page 96](#).

11.6.3.1 Choosing a cutoff score

At the top of the Validation Stats panel, a slider appears for adjusting the cutoff score. This is the aiR score that Review Center uses as the dividing line between positive and negative document predictions.




The screenshot shows a user interface for 'Validation Stats'. Below the title, there is a section titled 'Adjust Cutoff Score' with a descriptive text: 'Customize the cutoff score to determine what score and above is considered Relevant in the validation results.' To the right of the text is a horizontal slider with a blue handle. The slider is positioned at the value '3', which is also displayed as a number at the far right end of the slider track.

By default, the cutoff score is set at 2 (Borderline Relevant). This means that aiR document scores of 2 or higher are counted as positive predictions, and lower scores are counted as negative. If you change the slider to 3 (Relevant), the statistics will update to count only scores of 3 or higher as positive. You can set the slider to any value between 1 and 4. You can also change the cutoff score as many times as you want before accepting or rejecting results. Each time you move the slider, the validation statistics update to reflect the new cutoff.

Notes:

- The most commonly used cutoffs are 2 (Borderline Relevant) and 3 (Relevant).
- When you choose a cutoff, the documents' underlying aiR scores do not change. For example, if you set the cutoff to treat Borderline Relevant documents as negative, their score remains a 2. It does not change their underlying score to a 1.

When you accept the validation results, the chosen cutoff score becomes permanent. The final cutoff score appears in the upper right corner of the Validation Stats panel.




For more information on how the cutoff score affects the statistics, see [How documents are categorized for calculations on page 97](#).

11.6.4 Accepting or rejecting results

When the human coding decisions are complete, you can review how effectively the AI matched human decisions, then decide whether to accept the results and use the Prompt Criteria as-is, or whether to reject the results and improve the Prompt Criteria.

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



If you click Accept:

- The cutoff score you selected becomes permanent for the queue. For more information, see [Choosing a cutoff score on the previous page](#).
- In aiR Review, you can no longer create new Develop sets.
- The queue status changes to Validation Complete.

If you click Reject:

- In aiR Review, you can create a new Develop set.
- The queue status changes to Rejected.

After you make the choice, the Validation Progress strip on the dashboard displays the final validation statistics and a link back to the aiR Review project. From there, you can either use the finalized Prompt Criteria on a larger document set, or create a new Develop set, edit the Prompt Criteria, and continue improving it.

For information on continuing work in the aiR Review tab, see [Setting up aiR Review prompt criteria validation on page 88](#).

Note: If you reject this validation, you can run validation again later. Even if you reject the results, Review Center keeps a record of them. For more information, see [Viewing results for previous validation queues below](#).

11.6.4.1 Manually rejecting validation results

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

To reject the results after accepting them:

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


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

11.6.4.2 Viewing results for previous validation queues

After you have run validation, you can switch back and forth between viewing the statistics for the current validation attempt and any previous validation queues that were completed or rejected. These queues are considered linked. Viewing the statistics for linked queues does not affect which queue is active or interrupt reviewers.

To view linked queues:

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



A drop-down menu listing all linked queues appears.

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

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

11.6.5 How changes affect the validation results

The validation process assumes that the Prompt Criteria, document set, and coding decisions will all remain the same. If any of these things change, the validation results will also change. Sometimes this can be solved by recalculating the validation statistics, but often it means creating a new validation queue.

11.6.5.1 Scenarios that require recalculation

The following scenarios can be fixed by recalculating statistics:

- Changing coding decisions on documents within the validation sample
- Re-running aiR Review to fix errored documents

In these cases, the sample itself is still valid, but the numbers have changed. For these situations, recalculate the validation results to see accurate statistics. For instructions on how to recalculate results, see [Recalculating validation results on the next page](#).

11.6.5.2 Scenarios that require a new validation queue

The following scenarios require a new validation queue:

- Changing the Prompt Criteria
- Adding or removing documents from the document set after validation starts

In these cases, the sample or the criteria themselves have changed, so recalculating does not help. For these situations, create a new validation queue.

11.6.5.3 Recalculating validation results

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

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

To recalculate validation results:

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

11.7 Prompt Criteria validation statistics

aiR Review Prompt Criteria validation provides several metrics for evaluating your Prompt Criteria. Together, these metrics can help you determine whether the Prompt Criteria will work as expected across the full data set.

Because Prompt Criteria validation is checking the effectiveness of the Prompt Criteria, rather than checking completeness of a late-stage review, these statistics are calculated slightly differently than standard Review Center validation statistics. For the standard Review Center metrics, see Review validation statistics in the Review Center guide.

For more details on the differences between Prompt Criteria validation statistics and standard Review Center statistics, see [How Prompt Criteria validation differs from other validation types on page 99](#).

11.7.1 Defining the validation statistics

Prompt Criteria validation centers on the following statistics. For all of these, it uses a 95% confidence interval:

- **Elusion rate**—the percentage of documents that aiR predicted as non-relevant, but that were actually relevant.
- **Precision**—the percentage of documents that aiR predicted as relevant that were truly relevant.
- **Recall**—the percentage of truly relevant documents that were found using the current Prompt Criteria.
- **Richness**—the percentage of relevant documents across the entire document set.
- **Error rate**—the percentage of documents that received errors in aiR Review.

In everyday terms, you can think of these as:

- Elusion rate: "How much of what we're leaving behind is relevant?"
- Precision: "How much junk is mixed in with what we think is relevant?"
- Recall: "How much of the relevant stuff can we find?"
- Richness: "How much of the overall document set is relevant?"
- Error rate: "How many documents aren't being read at all?"

For each of these metrics, the validation queue assumes that you trust the human coding decisions over aiR's predictions. It does not second-guess human decisions.

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

11.7.2 How documents are categorized for calculations

aiR Review tracks one field that represents whether a document is relevant or non-relevant, with only one choice for Relevant. Any other choices available on that field are considered non-relevant.

When you validate Prompt Criteria, Review Center uses that field to categorize coding decisions:

- **Relevant**—the reviewer selected the relevant choice.
- **Non-relevant**—the reviewer selected a different choice.
- **Skipped**—the reviewer skipped the document.

aiR Review's relevance predictions have a similar set of possible values:

- **Positive**—The document's score is greater than or equal to the cutoff.
- **Negative**—The document's score is at least zero, but less than the cutoff.
- **Error**—The document's score is -1, meaning that aiR Review could not process the document.

At the end of validation, a slider appears in the Validation Stats area to adjust the cutoff score. The cutoff you choose determines which scores are considered positive. For example, if you set the cutoff score at 2 (Borderline Relevant), all document scores of 2, 3, or 4 are considered positive predictions. If you set the cutoff score to 3 (Relevant), only scores of 3 or 4 are considered positive. For more information, see [Choosing a cutoff score on page 93](#).

11.7.2.1 Variables used in the metric calculations

Based on the possible combinations of coding decisions and relevance predictions, Review Center uses the following categories when calculating statistics. Each of these categories, except for Skipped, is represented by a variable in the statistical equations.

Document Category	Variable for Equations	Coding Decision	aiR Prediction
Error	E	(any)	error
Skipped	(N/A)	skipped	(any)
True Positive	TP	relevant	positive
False Positive	FP	non-relevant	positive
False Negative	FN	relevant	negative
True Negative	TN	non-relevant	negative

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

11.7.2.2 How validation handles skipped documents

We strongly recommend coding every document in the validation queue. Skipping documents reduces the accuracy of the validation statistics. The Prompt Criteria validation statistics ignore skipped documents and do not count them in the final metrics.

11.7.3 Prompt Criteria validation metric calculations

When you validate a set of Prompt Criteria, each metric is calculated as follows.

11.7.3.1 Elusion rate

This is the percentage of documents that aiR predicted as non-relevant, but that are actually relevant.

Elusion Rate =
FORMULA

$$\frac{FN}{FN + TN}$$

Elusion = (False negatives) / (False negatives + true negatives)

The elusion rate gives an estimate of how many relevant documents would be missed if the current Prompt Criteria were used across the whole document set.

11.7.3.2 Precision

This is the percentage of documents that aiR predicted as relevant that were truly relevant.

Precision =
FORMULA

$$\frac{TP}{TP + FP}$$

Precision = (True positives) / (True positives + false positives)

Precision shares a numerator with the recall metric, but the denominators are different. In precision, the denominator is "what the Prompt Criteria predicts is relevant;" in recall, the denominator is "what is truly relevant."

11.7.3.3 Recall

This is the percentage of truly relevant documents that were found using the current Prompt Criteria.

Recall =
FORMULA

$$\frac{TP}{TP + FN}$$

Recall = (True positives) / (True positives + false negatives)

Recall shares a numerator with the precision metric, but the denominators are different. In recall, the denominator is "what is truly relevant;" in precision, the denominator is "what the Prompt Criteria predicts is relevant."

11.7.3.4 Richness

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

Richness =
FORMULA

$$\frac{TP + FN}{TP + FP + FN + TN}$$

Richness = (True positives + false negatives) / (True positives + false positives + false negatives + true negatives)

The richness calculation counts the final number of positive documents, then divides it by all groups except errored or skipped documents.

11.7.3.5 Error rate

This is the percentage of documents that could not be processed by aiR Review and received a score of -1.

Error Rate =
FORMULA

$$\frac{E}{E + TP + FP + FN + TN} = \frac{E}{\text{Sample Size}}$$

Error rate = (Errored documents) / (Errors + true positives + false positives + false negatives + true negatives)


This could also be written as:

Error rate = (Errored documents) / (Sample size)

The error rate counts all errors, then divides them by the total number of documents in the sample.

11.7.4 How the confidence interval works

When Review Center reports validation statistics, it includes a range underneath the main estimate. The main estimate is the "point estimate," meaning the value we estimated from the sample, and the range beneath it is the confidence interval.



Prompt Criteria validation uses a 95% confidence interval when reporting statistics. This means that 95% of the time, the range contains the true statistic for the entire document set.

This interval uses the Clopper-Pearson calculation method, which is statistically conservative. It is also asymmetrical, so the upper and lower limits of the range may be different distances from the point estimate. This is especially true when the point estimate is close to 100% or 0%. For example, if a validation sample shows 99% recall, there's lots of room for that to be an overestimate, but it cannot be an underestimate by more than one percentage point.

11.7.5 How Prompt Criteria validation differs from other validation types

aiR Review Prompt Criteria validation queues have a few key differences from prioritized review validation queues in Review Center. The main metrics are the same, but some of the details vary.

In Prompt Criteria validation:

- The validation typically takes place early in the review process, instead of towards the end.
- Skipped documents are excluded from calculations entirely, instead of counting as an unwanted result.
- The validation sample is taken from all documents in the data set, regardless of whether they were previously coded. This leads to some slight differences in the metric calculations.

12 Managing aiR Review jobs

Use the aiR Review Jobs tab to view and monitor job details at the instance and workspace levels.

- **Monitoring in-flight jobs**—see which jobs are queued, running, completed, or errored across every analysis type.
- **Troubleshooting failures**—quickly identify errored jobs and open the side panel to see failure reasons, document counts, and model or prompt versions.
- **Administering aiR jobs at the instance level**—review the most recent jobs across every workspace, including matter, workspace, and user context.

Additionally, authorized administrators can configure job size limits per user group at the instance level to help control and manage the size of jobs that users execute.

12.1 Setting aiR Review job size limits

At the instance level, administrators can limit how many documents are included in an aiR Review job per user group. This helps manage costs and prevents users from accidentally running large jobs.

Note: Job size limits are set at the instance level from the **Review Jobs** page. Only users with the required administrative permissions can configure these limits.

The following points outline how job size limits are managed for user groups in the system:

- Limits are managed at the user group level, so every member of a group shares the same maximum job size.
- Different limits can be assigned to different groups.
- Multiple groups can be assigned the same job size limit
- Each group can only have one job size limit assigned at a time.
- If a user belongs to multiple groups with different limits, the highest limit among those groups applies when running a job.

To configure job size limits for user groups at the instance level:

1. At the instance level, navigate to the **Review Jobs** page.
2. Click **Manage Job Limits** above the table.
3. On the **Manage User Group Job Limits** model, do the following:

Manage User Group Job Limits

Set maximum aiR Units allowed per job for specific user groups.
Users assigned to these groups won't be able to start a job that exceeds the defined limit.

[+ Add Job Limit](#)

User Group(s)	Job Limit
aiR for Revi...	1000

[Save](#) [Cancel](#)

- **User Group(s)**—select one or more user groups from the list. To remove a group from the field, click its corresponding circled X icon.
 - **Job Limit**—enter the maximum number of documents (aiR Units) that users in that group can submit in a single job.
4. To add more user groups and limits, click **+Add Job Limit** and repeat the previous step.
 5. To delete a job limit entry and its associated groups, click its corresponding **trashcan** icon.
 6. Click **Save** when done.

If users attempt to run a job that exceeds their document limit, an error appears for them showing the maximum size allowed, and the job does not execute.

12.2 Monitoring Review Jobs

There are two versions of the Jobs tab: one at the instance level, and one at the workspace level.

- **Instance-level tab**—System administrators can see the most recent 100 jobs across all workspaces in the instance on the **Review Jobs** page. It includes several extra columns to identify the workspace, matter, and client name connected to each job.
- **Workspace-level tab**—users at the workspace level can see all jobs within a single workspace on the **aiR Review Jobs** tab.

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

For information on managing tab permissions, see [Permissions on page 12](#).

To open a project from the jobs list, click the **Project Name** link. The aiR Review project opens for viewing and iteration editing.

To view details for a project, click its row to open the details panel. It displays attributes specific to the analysis type, so the fields displayed may vary.

12.2.1 Jobs list columns

- **Analysis Type**—analysis type of the job: Relevance, Key Documents, Issues, Confidential Business Information, or Custom.
- **Project Name**—name of the aiR Review project associated with the job. Click the project name link to view project details.
- **Job Status**—current state of the job. The possible statuses are: Not Started, Queued, In Progress, Completed, Cancelling, and Errored.
- **Doc Count**—number of documents submitted for analysis
- **User Name**—name of the user who submitted the job.
- **Matter Name** (*instance-level view only*)—name of the matter associated with the job's workspace.
- **Matter Number** (*instance-level view only*)—number of the matter associated with the job's workspace.
- **Workspace ID** (*instance-level view only*)—ID of the job's workspace.
- **Workspace Name** (*instance-level view only*)—name of the job's workspace.
- **Submitted Time**—time the job was submitted for analysis.
- **Completed Time**—time the job successfully completed. If the job failed or was canceled early, this field is blank.
- **Terminated Time**—time the job stopped running, regardless of whether it was canceled, failed, or completed successfully.

12.2.2 Project details panel fields

Click a row to open the details panel for the project. It displays attributes specific to the analysis type, so the fields displayed may vary. The list below contains all possible fields. Only the fields specific to the project's analysis type will

display.

- **[Project Name]**—name of the project displays in the panel heading.
- **Result Visibility**—use this field to link and unlink analysis types to a job. Options include Hidden, Partially Hidden, and Visible. See [Linking/Unlinking analysis types on the next page](#) for more information. This option is not available for Custom analyses projects.
- **Insights Successful**—appears only for Custom analyses projects. Number of insights that generated successful outputs.
- **Insights Errors**—appears only for Custom analyses projects. Number of insights failed to generated outputs.
- **Job Failure Reason**—reason the job failed.
- **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.
- **aiR Version**—the unique ID that identifies the underlying aiR system used to process each job. This ID does not map directly to a single large language model version used for the job.
- **Prompt Version**—the prompt criteria version associated with the job. For more information, see How prompt criteria versioning works.
- **Docs Successful**—number of documents that were successfully analyzed.
- **Docs Errored**—number of documents that encountered an error during analysis.
- **Docs Pending**—number of documents that are waiting to be analyzed.
- **Docs Skipped**—the number of documents that aiR did not return results for. This can happen for reasons such as canceling a job, network errors, and partial or complete job failures.

- **Estimated Wait Time**—estimated wait time for job to begin.
- **Estimated Run Time**—estimated wait time for job to complete.

The analysis type project information repeats from the Jobs list page.

12.3 Linking/Unlinking analysis types

When aiR 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 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 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 Review Jobs tab takes precedent. For more information on version selection in the dashboard, see [How prompt criteria versions affect the Viewer on page 68](#).

Note: This option is not available for Custom analyses projects.

To link (show) and unlink (hide) analysis types, do the following:

1. Navigate to the **aiR Review Jobs** tab.
2. Click **See details** for the job.
3. From the **Result Visibility** list, select one of the following:
 - **Visible**—links all documents from this job so their results are visible in the Viewer. See Viewer documentation for more information on aiR Review analysis results in the Viewer.
 - **Partially Hidden**—only some of the documents from this job show 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.
 - **Hidden**—unlinks all documents from this job so their results are hidden in the Viewer. See Viewer documentation for more information on aiR Review analysis results in the Viewer.

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 Review Jobs tab.
-

13 Creating document views

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.

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

For a full field list, see [aiR Review results fields on the next page](#). For more information on views, see Views 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 Review Jobs tab.
-

13.1 Creating an aiR Review results view

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

- Edit
- Control Number
- aiR Summaries::Topic
- aiR Summaries::Summary
- <Review Field>
- aiR <result type>::Score
- aiR <result type>::Prediction
- aiR Confidential Business Information::Types (*applicable to CBI projects*)
- aiR Custom::[Insight Name] ([Project Name]) (*applicable to Custom analyses projects*)

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 Review results fields below](#).

13.2 aiR Review results fields

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

- [aiR Summaries](#)
- [aiR Relevance Analysis](#)
- [aiR Issue Analysis](#)
- [aiR Key Analysis](#)
- [aiR Confidential Business Information Analysis](#)
- [aiR Custom analyses](#)

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 aiR Review jobs on page 99](#).

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

13.2.1 aiR Summaries fields

The Summaries fields include the topic and document summary that aiR automatically generates for each document, regardless of analysis type chosen for the project. These remain the same for each document across aiR projects.

When adding these fields to a view or saved search, use the reflected fields. The reflected fields are shown as `aiR Summaries::<fieldname>`.

The fields for aiR Summaries are:

Field name	Field type	Description
Topic	Fixed-length Text	The descriptive topic aiR generated for this document.
Summary	Fixed-length Text	The overall summary aiR generated for this document.

13.2.2 aiR Relevance Analysis fields

When adding these fields to a view or saved search, use the reflected fields. The reflected fields are shown as `aiR Relevance Analysis::<fieldname>`.

The fields for aiR Relevance Analysis are:

Field name	Field type	Description
Name	Fixed-length Text	The name of this specific result. This formatted as <code><Document Artifact ID>_<Job ID></code> .
Job ID	Fixed-length Text	The unique ID of the job this result came from.
Score	Whole Number	Numerical score indicating how strongly relevant the document is. For more information, see Understanding document scores on page 75 .
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	First 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-	Third excerpt from the document that supports the prediction and rationale. This may

Field name	Field type	Description
	length Text	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 Handling document errors on page 75 .

13.2.3 aiR Issues Analysis fields

When adding these fields to a view or saved search, use the reflected fields. The reflected fields are shown as `aiR Issues Analysis::<fieldname>`.

The fields for aiR Issues Analysis are:

Field name	Field type	Description
Name	Fixed-length Text	The name of this specific result. This formatted as <Document ID>_<Job ID>.
Job ID	Fixed-length Text	The unique ID of the job this result came from.
Choice Analyzed	Fixed-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 Understanding document scores on page 75 .
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-	Excerpt from the document that supports the prediction and rationale. This may be

Field name	Field type	Description
	length Text	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 Handling document errors on page 75 .

13.2.4 aiR Key Analysis fields

When adding these fields to a view or saved search, use the reflected fields. The reflected fields are shown as `aiR Key Analysis::<fieldname>`.

The fields for aiR Key Analysis are:

Field name	Field type	Description
Name	Fixed-length Text	The name of this specific result. This formatted as <Document ID>_<Job ID>.
Job ID	Fixed-length Text	The unique ID of the job this result came from.
Document	Multiple Object	The Control Number of the document this result is linked to. If the result is not currently linked to any documents, this field is blank.
Score	Whole Number	Numerical score indicating how strongly relevant the document is. For more information, see Understanding document scores on page 75 .
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	First 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.

Field name	Field type	Description
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 Handling document errors on page 75 .

13.2.5 aiR Confidential Business Information Analysis fields

When adding these fields to a view or saved search, use the reflected fields. The reflected fields are shown as `aiR Confidential Business Information Analysis::<fieldname>`.

The fields for aiR Confidential Business Information (CBI) Analysis are:

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.
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 prediction.
Considerations	Fixed-length Text	A counterargument explaining why the prediction might possibly be wrong.
Types	Fixed-length Text	The name of the CBI Type.
Citation 1	Fixed-length Text	First 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.

Field name	Field type	Description
Citation 6	Fixed-length Text	Sixth excerpt from the document that supports the prediction and rationale. This may be blank for some documents.
Citation 7	Fixed-length Text	Seventh excerpt from the document that supports the prediction and rationale. This may be blank for some documents.
Citation 8	Fixed-length Text	Eighth excerpt from the document that supports the prediction and rationale. This may be blank for some documents.
Citation 9	Fixed-length Text	Ninth excerpt from the document that supports the prediction and rationale. This may be blank for some documents.
Citation 10	Fixed-length Text	Tenth 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 Handling document errors on page 75 .

13.2.6 aiR Custom Analyses fields

When adding these fields to a view or saved search, use the reflected fields. The reflected fields are shown as `aiR Custom::<fieldname>(Project Name)`.

Field results will only be populated when a user creates and runs an “Apply Set.” See [Creating document views on page 103](#) for more information.

The field type of each <Custom Field Name> column below matches the Field Type selected for that insight on the Custom tab. See [Custom tab on page 64](#) for more information.

The fields for aiR Custom Analyses are:

Field name	Field type	Description
Name	Fixed-length Text	The name of this specific result. This formatted as <Document Artifact ID>_<Job ID>.
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.
<Custom Field Name> (Project Name)	Varies - matches the insight's selected Field Type (Fixed Length, Single Choice, Date, or Whole Number)	For Single Choice fields, choices are auto-created on this field at publish time; renamed choices update in place (tracked by Artifact ID) and new choices appended on later versions are not retroactively applied to previously coded documents.

14 Using aiR Review with Review Center

You can use aiR Review and Review Center in tandem as part of a larger review workflow. When you use them together, aiR Review provides the set of documents to Review Center, and Review Center helps you serve the selected documents to human reviewers.

14.1 Using prompt criteria validation

Validating aiR Review prompt criteria involves comparing AI-generated relevance predictions generated in aiR with human coding results in Review Center on a sample of the document set. After it's validated, you can choose to run the prompt criteria on the larger document population. For more information, see [aiR Review prompt criteria validation on page 83](#).

14.2 Using aiR to create a Review Center queue

At any point after the first analysis, you can select documents from the aiR Review dashboard and send them to Review Center. These documents get grouped together as a saved search queue.

14.2.1 Selecting documents to add to a queue

After you have run the aiR Review analysis, you can use the filters on the dashboard to select groups of documents for review. You can do this by clicking the document categories in the Metrics section, filtering the Analysis Results column headers, or a combination of both. You can also check or uncheck the boxes beside individual documents.

For example:

- If you want to review all documents that have conflicts, click the **Total** metric under Conflicts in the Metrics section.
- If you want to review all documents that have borderline or above relevance predictions, filter the **aiR Score** column for values of **2**, **3**, and **4**.


14.2.2 Creating a saved search queue through the aiR Review dashboard

After you have selected the documents you want to review, use a mass action to send them to Review Center and create a saved search queue. Review Center will also automatically create a saved search based on these documents.

Note: To use this mass action, you need permissions to create a Review Center queue. For more information, see Review Center security permissions in the Review Center guide.

To create the saved search queue:

1. On the aiR Review dashboard, select the documents you want in the queue.
2. Select **Create Queue** from the mass operations list at the bottom of the grid.



3. On the Create Queue in Review Center modal, fill out the setup fields. For descriptions of each field, see Queue and template settings list in the Review Center guide.

Note: If you set a relational field when creating the queue, set the same field in the Related Items drop-down of the saved search Conditions tab. Only relational group members returned by the saved search will be included in the queue.

4. Click **Create Queue**.

A green confirmation banner appears.

After you create the queue, you can manage it from the Review Center dashboard. The saved search will also be available after the first queue refresh.

For more information, see the Review Center guide.

14.2.2.1 Managing the saved search

After the first queue refresh, the saved search will be available in the saved search browser. It will have the same name as the queue.

The saved search has a condition that points to documents tagged as being in the queue. If you want to add or remove documents from the queue, you can modify the conditions of the saved search.

Note: If you set a relational field when creating the queue, set the same field in the Related Items drop-down of the saved search Conditions tab. Only relational group members returned by the saved search will be included in the queue.

For more information on editing saved searches, see the Searching guide.

15 Archiving and restoring workspaces with aiR Review

Workspaces with aiR Review installed can be archived and restored using the ARM application. For more information on using ARM, see ARM Overview on the Relativity documentation site.

15.1 Archiving workspaces

When you use ARM to archive workspaces with aiR Review installed, select the **Include Extended Workspace Data** option under Extended Workspace Data Options. If this option is not selected, the aiR Review features in the restored workspace will not be fully functional. If this happens, you must reinstall aiR Review in the restored workspace. See [Installing aiR Review on page 10](#) for installation steps.

15.2 Restoring workspaces

When you use ARM to restore workspaces that have aiR Review installed and include prior aiR Review jobs, those earlier jobs will not show up on the instance-level aiR Review Jobs tab. However, both the jobs and their results will remain accessible from the aiR Review Jobs tab at the workspace-level.

Proprietary Rights

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

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