Description of Methodology

Corvane Context Platform

Document owner	Measurement Governance
Document ID	CMI-GOV-DM-001
Version	3.2
Effective date	March 2, 2026
Review cycle	Annually, and before any material methodology change takes effect
Related documents	Transparency and Disclosure Statement (CMI-GOV-DS-004); Video Analysis API Reference; Taxonomy Reference

> **Note on figures**
> All figures in this sample, including accuracy results and processing times, are invented to illustrate how such a document presents them.

1 Introduction

This document describes how the Corvane Context Platform (the “Platform”) analyzes web pages and video, how it assigns topics and brand safety risk levels, how its models are built and checked, and where its results are limited. It is written for clients who use Platform data to make advertising decisions, for partners who integrate with the Platform, and for auditors who assess it.

This document explains how the Platform works. The Transparency and Disclosure Statement sets out Corvane’s commitments about accuracy, error correction, and notice of changes. The two documents should be read together.

2 What the Platform Measures

The Platform analyzes content, not audiences. For each web page or video segment it analyzes, it produces four kinds of result:

Result	What it describes
Topics	What the content is about, using the IAB Tech Lab Content Taxonomy. Content can have several topics.
Brand safety risk	Whether the content includes material in any of the Platform’s risk categories, and how severe that material is (Section 5).
Sentiment	Whether the overall tone of the content is positive, neutral, or negative.
Keywords	Notable terms that appear in the content, provided to support custom targeting and review.

Results are delivered through the API, the client dashboard, and scheduled reports. All three draw on the same stored results.

3 How Content Is Analyzed

3.1 Web pages

A web page is analyzed when a client requests it, either directly through the API or through a page tag that sends the page URL when the page loads. The process is:

1. The Platform checks whether the page has a current result. Results are considered current for 24 hours for pages on news sites and 14 days for all other pages. If a current result exists, it is returned and the process ends.
2. The Platform requests the page using a crawler that identifies itself as Corvane in its user agent. It follows the site's robots.txt rules and does not sign in, accept paywall terms, or bypass access controls.
3. The Platform renders the page, waiting up to 10 seconds for content that loads with scripts.
4. It extracts the page title, main body text, image files, and descriptive metadata, and removes navigation, advertisements, comments, and other content that is not part of the article or main content.
5. Text and images are sent to the classification models described in Section 4.
6. The results are combined into a page-level result, stored, and returned to the client.

3.2 Video

Video is analyzed when a client submits a media URL through the API. The process is:

1. The Platform downloads the media file or, for HLS playlists, the full set of media segments.
2. Audio is transcribed into text using a third-party speech-to-text service (see Section 7).
3. One frame is sampled every second. Text that appears on screen is read from the sampled frames, and objects and scenes are identified.
4. The video is divided into segments of equal length (15 seconds by default). The transcript, on-screen text, and visual results for each segment are classified together.
5. Segment results are combined into a video-level result as described in Section 4.4.

4 Classification Models

4.1 Topic classification

Separate models classify text and images. Each model returns a confidence score from 0 to 1 for each topic. A topic is reported only if its combined confidence is 0.60 or higher. Topics are mapped to IAB Content Taxonomy 3.0, and any place where Corvane's categories differ from the IAB definitions is listed in the Taxonomy Reference.

4.2 Brand safety classification

Brand safety models assess content against 11 risk categories, which are based on categories widely used across the advertising industry. For each category, the model decides whether relevant material is present and, if so, assigns a risk level using the definitions in Section 5.

4.3 Sentiment

A text model scores the overall tone of the content from -1.0 (strongly negative) to 1.0 (strongly positive). Scores between -0.2 and 0.2 are labeled neutral. Sentiment is not assessed for content without text or speech.

4.4 Combining results

- A page or video takes the highest risk level found in any of its parts. One high-risk segment makes the whole video high risk, even if the rest of the video has no risk.
- Video-level topics are weighted by how much of the video each topic covers.
- Video-level sentiment is the average of segment scores, weighted by segment length.

5 Risk Categories and Risk Levels

Risk categories (1 of 2)	Risk categories (2 of 2)
Adult and explicit sexual content	Illegal drugs, tobacco, and vaping
Arms and ammunition	Military conflict
Crime and harmful acts	Online piracy
Death, injury, and graphic violence	Spam and harmful software
Hate speech and harassment	Terrorism
Obscenity and profanity	

Risk levels

Level	Definition
Floor	Content that no advertiser would knowingly support, such as content that promotes violence or exploitation. Floor content is reported regardless of client thresholds.
High	Explicit or graphic material in a risk category, or content whose main purpose is related to a risk category.
Medium	Material in a risk category that is discussed in a news, educational, or informational context.
Low	Brief or incidental references to a risk category.
None	No material in any risk category was found.

Clients may choose which risk levels to avoid for each category. Default thresholds avoid Floor and High content in every category.

6 Model Development and Quality Control

6.1 Training data and labeling

Models are trained on content labeled by trained human reviewers following written labeling guidelines. Guidelines are kept under version control and updated whenever the taxonomy or risk definitions change. Each item in an evaluation set is labeled by at least two reviewers; disagreements are resolved by a senior reviewer, and agreement rates are tracked for each category.

6.2 Evaluation before release

Each new model version is tested against a held-out set of labeled content that was not used in training. A new version is released only if it meets both of these criteria:

- Precision and recall for every Floor and High category are at or above the targets set by Measurement Governance.
- Recall for any risk category does not fall by more than one percentage point compared with the current production model.

Figure 1. Example evaluation results for selected categories (illustrative)

Category	Precision	Recall	Evaluation items
Death, injury, and graphic violence	0.93	0.89	4,200
Hate speech and harassment	0.90	0.86	3,850
Illegal drugs, tobacco, and vaping	0.94	0.91	3,100
Topic: Personal Finance	0.88	0.84	2,750

6.3 Monitoring in production

Each week, reviewers label a random sample of 2,000 items that the Platform has classified, and the results are compared with the Platform's output. If precision or recall for any category falls below its target for two consecutive weeks, Measurement Governance opens an investigation and decides whether retraining is needed.

6.4 Releasing model changes

New models are released to 10% of traffic for at least seven days, and results are compared with the current model before full release. A change that could reasonably alter reported results is treated as a material change and announced to clients at least 30 days in advance, as described in the Transparency and Disclosure Statement.

7 Third-Party Technology

Function	Provider type	Data shared
Speech-to-text transcription	Commercial cloud transcription service	Audio tracks from submitted video. The provider is contractually barred from retaining or reusing the audio.
Page rendering	Corvane-operated rendering service on a commercial cloud platform	Page URLs and page content
Model hosting	Commercial cloud platform	Extracted text and sampled frames, processed in Corvane-controlled accounts

8 Processing Times

Request	Typical time to result
Web page with a current result	Under 200 milliseconds
Web page without a current result	Within 5 minutes; results for the first request may be returned as Not Analyzed while processing completes
Video	About half the length of the video, plus up to 5 minutes in the queue

9 Limitations

The conditions below affect whether, or how fully, the Platform can analyze content. When content cannot be analyzed, it is reported as Not Analyzed with a reason code, and it is never reported as having no risk.

Condition	Effect on results	How it is reported
Page requires sign-in or is behind a paywall	Only content visible without signing in is analyzed, which may be just a headline or summary.	Not Analyzed (`access_restricted`) if under 100 words are visible
Site's robots.txt excludes the Corvane crawler	The page is not retrieved.	Not Analyzed (`robots_excluded`)
Page content takes more than 10 seconds to load	Content that loads late may be missed.	Analyzed; flagged `partial_render`
Text is in an unsupported language	Text is not classified; images are classified.	Analyzed; flagged `language_unsupported`
Video has no speech	Segments are classified from visual content and on-screen text only.	Analyzed; `transcript_available` is false
Content changes after analysis	Changes are not reflected until the result expires and the content is analyzed again.	Result date is included with every result
Satire, irony, or coded language	Models may misjudge meaning or risk.	Not flagged; included in measured error rates

10 Privacy

The Platform analyzes content only. It does not set or read cookies, collect device identifiers, or receive information about the people who view the content it analyzes. Page tags send only the page URL. Request logs, which contain URLs and timestamps, are retained for 13 months to support audits and dispute resolution.

11 Revision History

Version	Date	Summary of changes
3.0	April 7, 2025	Restructured the document. Added video analysis (Section 3.2) and segment-level results.
3.1	October 6, 2025	Added sentiment (Section 4.3). Updated topic mapping to IAB Content Taxonomy 3.0.
3.2	March 2, 2026	Added staged release of model changes (Section 6.4). Added Not Analyzed reason codes to Limitations (Section 9).