Figma

AI PRODUCT SYSTEM CARD

FigJam

LAST UPDATED DATE, 10/13/26

PRODUCT STATUS	AVAILABLE ON THE FOLLOWING CUSTOMER TIERS	WHAT IT DOES
Generally available. FigJam Agent is in closed beta.	Paid plans. Full, Dev, and Collab seats. FigJam Agent is enabled plan by plan.	FigJam AI helps teams summarize and organize board content, create templates and visuals, work with Jambot, and use the closed beta FigJam Agent to build and rearrange board content.

About this document

An AI feature is more than the model behind it. How well it works for you depends on the people using it and the tasks they use it for. This system card explains how FigJam AI works, best practices for performance, and how your data is used. This document is for information only. It is not a contract, a warranty, or legal advice.

Contents

1 Overview at a glance

2 What it does and how it works

3 Models and providers

4 Intended use

5 Limitations

6 Your data and privacy

7 Content moderation

8 Safeguards and security

9 Performance, evaluation, and fairness

10 Human oversight, feedback, and control

11 Transparency

1 Overview at a glance

This section summarizes the main facts described in the rest of this system card.

What it does	Generates and organizes board content from your instructions and the board context available to the selected capability.
Who it is for	Full, Dev, and Collab seats on paid plans.
Where it appears	On a FigJam board.
Availability	FigJam AI capabilities are available on paid plans. FigJam Agent is a closed beta that Figma enables plan by plan.
Model selection	Configured by Figma. There is no model picker.
Data processing and storage	Figma stores your content in the United States. Model providers do not store it at rest.
Training on your content	Model providers never train on your content. Figma content training follows the plan level administrator setting described in Section 6.6.2.
Content moderation	Screening varies by capability. Supported routes screen selected input before generation, and providers apply safety classifiers to generated content.
Human oversight	You start every action. Legacy actions and Jambot return results for review. FigJam Agent can apply board changes while it works, and you can stop the run and review its effects.
FigGov Experience	Not available.
Gov+ Experience	Available. FigJam Agent is unavailable when a Governance+ customer enables Bedrock pinning.
EDU Experience	Available to higher education and bootcamp users, but not to K-12 users. FigJam Agent is excluded from the closed beta for education accounts.

2 What it does and how it works

2.1 Description

FigJam AI works from your instructions and the board context available to the capability you start. It can summarize a board, cluster or classify sticky notes, match content to a template, create or update a visual, rewrite text, generate source code as text, and reply through Jambot. FigJam Agent can build and rearrange diagrams and templates and can group or summarize sticky notes while carrying out your requested task.

Legacy actions and Jambot place editable results on the board for you to review, keep, change, undo, or replace. FigJam Agent can apply board changes during a run, and you can stop the run and review its effects.

2.2 What the AI generates or predicts

FigJam AI generates structured board content, conversational replies, rewritten text, and source code shown as text. Structured content includes summaries, related groups of sticky notes, classifications, templates, and diagrams. A supporting embedding step converts board text into numeric representations so the feature can match the board to a template. It does not generate text.

FigJam AI does not generate photographs, video, audio, or realistic images of people.

2.3 How it works with other Figma products

FigJam AI works inside the FigJam board you have open. It uses your instruction and the board content available to the capability, then returns or applies editable content on that board. FigJam Agent can read the board and apply changes while it works. Access and administrator controls are described in Sections 6.2 and 8.2.

3 Models and providers

3.1 Models and where they run

FigJam AI uses language, embedding, and moderation models built by Anthropic, OpenAI, Meta, and Amazon. Language models predict likely text from their training data and the context sent with your request. The same instruction can produce different results, and output can be inaccurate, incomplete, or confidently wrong.

MODEL AND DEVELOPER	ROLE IN THIS FEATURE	WHERE IT RUNS
Claude models, Anthropic	Generate summaries, clusters, templates, and visuals, and power the FigJam Agent.	Anthropic, AWS Bedrock, and Google enterprise platform
GPT models, OpenAI	Classify board content and power Jambot.	OpenAI and Microsoft Azure OpenAI
Text embedding model, OpenAI	Converts board text into numeric representations so a board can be matched to a template.	OpenAI and Microsoft Azure OpenAI
Meta Llama, Meta	Provides a fallback path for Jambot responses.	AWS Bedrock
Moderation models, OpenAI and Amazon	Apply safety checks to supported FigJam routes.	OpenAI, Microsoft Azure OpenAI, and AWS Bedrock

3.2 Can I select the models?

No. Figma configures the model and hosting path for each capability, and there is no model picker. Moderation also uses models configured by Figma.

3.3 What the providers do with your content

Providers do not train on your content. They receive your instruction and relevant board content only to process the request and apply safety checks, and their contracts prohibit training on Figma customer content. Every provider operates under zero data retention or an equivalent no-logging configuration, so nothing is retained once your request has been served.

3.4 How a model change reaches you

Figma pins FigJam AI to specific model versions, so a provider update does not reach you automatically. Before adopting a replacement, Figma compares its quality and safety with the version already in use, uses sampled prompts, and stages the rollout. Figma tracks provider release and deprecation notices and keeps a tested rollback available.

4 Intended use

4.1 Primary intended users

FigJam AI is for designers, product managers, researchers, facilitators, and other people who plan or run collaborative work on a FigJam board. Access requires a Full, Dev, or Collab seat on a paid plan. FigJam Agent is available only while Figma enables its closed beta plan by plan.

4.2 Primary intended uses

- Summarize a board into takeaways after a session.
- Group, classify, and label sticky notes and other board content.
- Create or update templates and diagrams, ask Jambot for help, or use FigJam Agent to build and rearrange board content.

4.3 Where it works best

- Your instruction is specific and limited to one task.
- The content needed for the result is present on the board and available to the capability.
- You select the relevant objects and plan to review and adjust the result.

5 Limitations

5.1 Known limitations

- Output can be misleading or wrong, so do not treat a summary, classification, reply, diagram, or code result as fact.
- A summary or classification can miss nuance or group items in a way a person would not.
- Most capabilities work only from the board context they receive, so they can miss information in other files or information your team has not written down.

5.2 Factors that affect output quality

- The specificity and scope of your instruction.
- The clarity and completeness of the board content available to the capability.
- The objects you select and the capability you choose.

5.3 Best practices for performance

- Check a result against the visible board content before you rely on it.
- Keep, edit, undo, or replace a result when it does not meet your needs.
- Watch changes made by FigJam Agent, stop a run when needed, and review the board before using the result.

6 Your data and privacy

6.1 Data that is Processed

DATA	WHAT IT INCLUDES
Your instruction	The prompt you provide to a FigJam AI capability.
Relevant board content	Text on sticky notes and shapes, section and connector labels, connected Jambot content, and the objects selected for the action.
Operational metadata	File, session, and request identifiers used to route the response back to your board.

6.2 Who at Figma can access your data

Access follows the permissions set in your Figma organization, and Figma confirms your identity and your permission to act on the target file before a build runs.

Figma tightly restricts internal access to customer data on a least-privilege basis, granting it only to employees whose roles require it and only for the specific data needed.

6.3 Who at our model providers can access your content

Your prompt and data is shared with model providers solely so that they can process your request and return the requested output. Provider staff cannot access this content. Every provider operates under zero data retention or an equivalent no-logging configuration, and nothing is retained once your request has been served. No third party is permitted to train on your content.

6.4 Where your data is stored

Figma and the model providers handle processing and storage separately, so each has its own subsection below.

6.4.1 Where Figma stores your content

Figma stores your content in the United States.

6.4.2 Where the model providers store your content

Not applicable. No model provider that powers this feature stores your data at rest.

6.5 Retention and deletion

6.5.1 What is retained

WHAT IS RETAINED	PURPOSE	RETENTION PERIOD
Your files, including prompts and generated content	Your work product, stored as customer content	For the life of the account or until deletion
Usage and telemetry events, and metering aggregates	AI billing	For the life of the account
Operational telemetry and error context	Debugging and incident response	For the life of the account

6.5.2 Deleting your data

You can submit a deletion request. Figma propagates the request across the AI datasets and usage data it stores. FigJam files, including prompts and generated board content, are customer content and are kept for the life of the account or until deletion. Providers ordinarily have nothing to delete because they do not retain prompts or outputs.

6.6 Training on your content

6.6.1 Whether the model providers train on your content

Providers never train on your data. There is no separate opt-out for provider training, because your content is already excluded by contract and by the no-retention configuration Figma runs on every provider.

6.6.2 Whether Figma trains on your content

Whether Figma uses your content to improve its own AI is governed by an account-level content training setting that an administrator manages. This setting is independent of the control that turns AI on or off. On Starter and Professional, the team-level setting is on by default and an administrator can turn it off. On Organization and Enterprise, content training is disabled and cannot be turned on. Content from Figma for Government and education accounts is never used for training, regardless of the setting.

6.7 Sensitive and confidential data

FigJam AI does not require special category personal data, and it does not ask for it. The Help Center carries the same reminder. Apply your own policies for confidential, regulated, and personal data, and confirm that you have the right to use the content you provide.

7 Content moderation

7.1 What is screened

Screening varies by capability. Summarize, cluster, classification, and Jambot screen selected text before generation. Template and visual generation do not run this pre-generation step. Providers apply their own safety classifiers to generated content. Figma validates the structure and length of generated board objects before returning them where the capability uses structured output.

7.2 What is blocked

Prompts and generated results are blocked when they violate our acceptable use policy. The following categories are examples of content the policy prohibits.

- Child sexual abuse material, including any depiction, simulation, or facilitation of sexual content involving minors.
- Pornographic or graphic sexual content involving adults.
- Harassment, hate, dehumanization, and credible threats of violence toward individuals or protected classes.
- Step-by-step instructions for committing crimes.
- Instructions for building weapons or explosives, or for synthesizing controlled substances.
- Malicious code, exploits, backdoors, and instructions enabling unauthorized access or system abuse.

– Encouragement or instruction of self-harm or suicide.

7.3 What happens when content is blocked

Blocked content is not returned. Figma shows a message stating that potentially harmful results were blocked. Moderation is automated, so a request can be blocked even when it was not intended to violate the policy. You can report a block through the channels in Section 10.

7.4 Whether moderation can be edited or configured

You cannot turn content moderation off or adjust its thresholds in the product. Figma controls the safety settings on the server.

8 Safeguards and security

8.1 Security controls

– Traffic to model providers is encrypted in transit with HTTPS and Transport Layer Security. Requests to managed hosting paths are signed and encrypted end to end.

– Production data at rest is encrypted. Keys are managed in a dedicated key management service, and credentials are held separately.

– Every AI request passes through Figma's AI gateway, which applies routing, moderation entry, rate limiting, structured logging, and alerts.

– Server-side instructions and provider credentials remain on Figma servers and are not sent to your client.

– Per person and per organization rate limits protect against abuse and overload, and Figma can disable the feature to contain a problem.

– Internal data stores are isolated and reachable only through approved roles and network policies.

8.2 Administrator controls

– Enable or disable AI features, including FigJam AI, at the organization or team level.

– Manage the content training setting described in Section 6.6.2. It applies at the organization level on Organization and Enterprise and at the team level on Professional.

– Use Governance+ organization controls, including AI feature enablement. FigJam Agent is unavailable when a Governance+ customer enables Bedrock pinning.

– FigJam AI is unavailable in Figma for Government. It is available to higher education and bootcamp users, but not to K-12 users. FigJam Agent is excluded from the education closed beta.

9 Performance, evaluation, and fairness

9.1 How we evaluate this feature

Figma validates model changes before they reach you and stages their rollout as described in Section 3.4. A release is blocked by systemic quality regressions, elevated policy violations, repeated timeouts, defects that stop results from rendering, and unsafe content.

In production, Figma monitors request success, latency, crashes, severe issues, and infrastructure load. Periodic human and model-graded review compares sampled outputs with task specific criteria for accuracy, instruction adherence, and usefulness. Critical regressions can trigger rollback or a provider switch.

9.2 What we measure

MEASURE	WHAT IT MEASURES	WHY IT MATTERS TO YOU
Request success and latency	How often a request produces a result and how long the result takes.	Changes can identify a model or hosting problem before it disrupts a session.
Client crash rate	How often a FigJam session crashes during an AI action.	A rise can reveal a defect that prevents you from completing work.
Active severe issues	Whether a rollout cohort has an unresolved severe defect.	A severe issue can block a rollout or trigger containment.
Infrastructure load	Capacity used against the capacity available for the rollout.	Capacity pressure can reduce reliability or increase latency.
Sampled quality review	Human and model-graded review of sampled output for fabrication, instruction adherence, and functional validity.	Review can catch a well formed result that is not grounded in your board or does not complete the task.

Figma sets internal baselines for these measures and reviews the feature when a baseline is crossed.

9.3 Fairness and bias

9.3.1 What Figma reviewed

Figma uses red teaming with demographic and sensitive prompts and sampled human review of generated content for stereotyping. These qualitative and adversarial methods fit a feature that organizes board content rather than people and does not use or infer sensitive demographic attributes. Findings are logged and routed to the feature team. Neutral prompts that produce placeholder text, names, or examples can still default to a dominant demographic.

9.3.2 What the model providers reviewed

Figma also relies on bias evaluations published by the model providers. Anthropic describes political bias evaluations using paired opposing prompts and a question answering bias benchmark, along with child safety evaluations and harmful content filtering during training. OpenAI describes bias testing and harmful content filtering at a high level.

9.3.3 Known differences in performance

Result quality varies by language and cultural context. The underlying models are multilingual, and sampled reviews include boards that are not in English, but generation can be strongest in English and Western contexts. Figma is not aware of systematic performance differences by demographic group or geographic region.

9.4 Evaluate it for your own use

Before you depend on FigJam AI, test it with your own boards, tasks, and reviewers. Compare a result with the visible board content that shaped it, and decide which outputs require review before use. Include FigJam Agent changes in this review because the agent can apply changes during a run. The same instruction can produce different results, so Figma does not guarantee exact reproduction.

10 Human oversight, feedback, and control

10.1 Human oversight

You start every FigJam AI capability. Legacy actions and Jambot return a result that you can keep, edit, undo, or replace. FigJam Agent can apply board changes while carrying out your requested task, and you can stop the run and review its effects.

FigJam AI does not make or enforce decisions about people, restrict accounts, or report people to outside parties. Only signed in Figma customers with access to a board can use it. A board you share can be viewed by others and remains subject to Figma's usage policies.

10.2 Reporting problems and giving feedback

- Report a failure, an unsafe result, or another concern through the Figma Help Center or your usual support channel.
- Enterprise customers can raise concerns through Customer Success.
- Support reviews reports and escalates them to product, security, and legal teams when needed.

10.3 Turning the feature off

Organization and team administrators can turn off AI features at the organization or team level. The feature stays off until an administrator enables it again.

11 Transparency

11.1 How AI involvement is disclosed

FigJam AI is presented as named and labeled AI capabilities on the board. You start each capability deliberately, and Jambot is labeled as AI rather than styled as a person. Generated content appears as editable board content, and FigJam Agent changes remain visible on the board for review.

11.2 How we communicate changes

Figma communicates material changes to functionality, intended use, and reliability through release notes, Help Center documentation, in-product messaging, and standard release communications.