

present

Contextual Intelligence: Using Proprietary Language Models for Presentation Generation



Table of Contents

Challenges with AI Presentation Creation		3
Using Proprietary Language Models for a Better Presentation Generation Experience		5
Levels of Contextual Intelligence		6
Example of How Contextual Intelligence Works		7
Evaluating Contextual Intelligence in the Market for Business		9
FAQs		13

Challenges with AI Presentation Creation

Presentations made from AI models can be underwhelming, full of hallucinations, include irrelevant information, and lack critical context. This creates major challenges for business professionals, who are expected to communicate with clarity, consistency, and speed.

Most AI presentations:

Aren't personalized to your specific audience

1

Don't incorporate your company and team content or terminology

2

Don't capture your company and team branding

3

Don't follow best practice layouts by industry function

4

Can be too generic and error prone for industry and subject matter experts

5

Lacking audience personalization

1

We've all experienced sitting through a presentation that felt irrelevant or failed to connect. Generic presentations lose attention fast. To truly engage, presentations must be tailored to the audience making them relevant, impactful, and effective. For example, different roles require different approaches. A data scientist might need slides filled with data and analytics, while a VP of Sales may prefer a concise summary that gets straight to the bottom line.

Missing critical team and company content

2

An AI tool is only as effective as the inputs it receives. Every company has its own unique style, language, and acronyms, but today's AI tools fall short. Without the ability to incorporate proprietary assets, documents, and organization-specific data, they can only deliver generic results, limiting the impact and relevance of the presentations they generate. Imagine

a biopharma company preparing a presentation for an FDA advisory committee. The company has proprietary frameworks, specific acronyms like “AUC-PK” (Area Under the Curve for Pharmacokinetics) and “iPSC” (induced Pluripotent Stem Cells) and they need to incorporate confidential clinical trial data, and patient stratification methodologies to support their drug approval application. The result with most AI tools? A generic presentation that lacks the company's identity and fails to convey expertise—leaving the committee unimpressed and the message diluted.

Not brand-aligned to your team and company guidelines

3

Consistent branding is critical for professional credibility, yet both AI tools and humans often struggle to maintain it. AI tools, in particular, can exacerbate the issue through errors like inconsistent branding or hallucinations—where factually incorrect or fabricated information is generated due to AI's reliance on patterns rather than real-time verification. This is especially significant in heavily regulated industries like biopharma or finance, where strict guidelines for fonts, layouts, logos, and colors must be followed. When these tools fail to uphold branding standards or are inaccurate, the results are bad presentations that weaken a brand identity and credibility.

Take a biopharma company for example that is preparing a presentation for potential investors to secure funding for their next-stage of clinical trials. This is a high-stakes presentation that needs to align with the company's branding, including a proprietary template with defined fonts and colors. It can take an internal designer weeks to do this due to the necessary attention to detail or it will cost an arm and a leg and time to ask traditional agencies. Many AI tools will misapply branding, misplace logos, and generate fabricated data. These errors and attention to detail risk confusing investors and delaying funding.

Mismatched layouts by industry-function

4

Pre-designed templates often miss the mark on industry-specific or functional best practices, reducing the impact of presentations. Have you ever searched for a slide template online and found that none of them quite fit? That's because every industry and company have their own unique approach to presenting information—an approach that generic AI tools can't replicate. A tech company preparing a pitch relies on specific layouts, like a “Feature Comparison” slide

and a “Technical Architecture” diagram, to clearly present product details. A generic AI tool misaligns these layouts, cluttering information and simplifying complex diagrams. These errors dilute the presentation’s impact and risk losing interest.

Too generic and error-prone for industry and functional experts

5

Other AI presentation businesses fail to provide industry-specific accuracy, often generating content that is overly generic and failing to capture the specialized knowledge and terminology required in fields like biopharma, technology, or finance. For example a tech company might require advanced product roadmaps or architecture diagrams that need to adhere to specific standards of technical detail—elements that basic AI templates are often too simplistic to handle.

Alternatively, agencies offer expertise but they often charge premium rates that are not sustainable for many businesses. This creates a gap for organizations seeking effective communication and presentation solutions at scale.

Contextual Intelligence: Using Proprietary Language Models for a Better Presentation Generation Experience

Context bridges the gap between AI’s capabilities and human centered communication and design, creating a powerful intersection between technology and humanity. Turning important and relevant context into referenceable intelligence enhances communication clarity and effectiveness. This is called Contextual Intelligence (CI).

Contextual Intelligence is Prezent’s proprietary approach to creating precise, relevant, and impactful communication. Allowing users to generate presentations that are fully customized at company, team, and individual levels, contextual intelligence is the missing piece between a generic, vague AI response and a hyper personalized presentation, so tailored to your prompt, you’re 80% complete in minutes rather than starting from zero.

Contextual Intelligence enables teams to craft highly effective presentations, ensuring they are not only relevant and impactful for their specific audience but also maintain consistent brand messaging. While most AI presentation tools generate content based on patterns in large datasets, they often fail to incorporate the nuanced needs of a particular industry or company. By feeding Prezent's AI platform with relevant context, such as industry-specific terminology, brand guidelines, individuals preferences and much more, generated presentations are precise and compelling—enabling more effective business communication.

While this context includes audience preferences, specific company content such as acronyms and tone and of course company branding, this company or member data is not stored or used by Prezent's AI. They are specific to an individual company's experience within the software.

The Power of Contextual Intelligence



Tailored for Precision

Prezent's AI models are fine-tuned specifically for your industry, understanding sector-specific terminology and nuances.



Consistency and Branding

Presentations adhere strictly to your company's branding and voice, ensuring consistent messaging across all platforms



Transparent and Trustworthy

Using Retrieval-Augmented Generation (RAG), Prezent maintains accuracy by grounding content in verified sources, clearly cited for transparency.

Levels of Contextual Intelligence

Prezent has created Contextual Intelligence across 5 levels to transform how teams create and deliver impactful presentations and narratives. These combined levels result in tailored presentations that resonate with audiences, drive engagement, and elevate business outcomes.

	Levels of Context	Benefit
L5	Individual	Personalize to audience preferences
L4	Team	Capture team voice
L3	Company	Capture company branding
L2	Function	Best-practice layouts
L1	Industry	Reduce hallucinations + industry specific

Example of How Contextual Intelligence Works

Step 1: Understanding your needs

When you request a presentation, our AI-powered Input Preprocessor Agent quickly identifies key topics from your prompt. For example, a prompt such as "Create a slide on the Roche-Gilead collaboration on Tamiflu" is distilled to essential phrases like "Roche-Gilead collaboration" and "Tamiflu".

Step 2: Finding Relevant Information

Those extracted keyphrases are passed to our Retriever Agent which then searches both your internal company resources and the wider web to gather detailed and relevant information. Using proprietary algorithms, it provides a rich source of contextual knowledge. For the Tamiflu example, it would identify and retrieve critical documents outlining specifics like collaboration details specified in Gilead and Roche News or about the accelerated timeline for bringing the drug to market and revenue sharing.

Step 3: Ensuring Accuracy and Relevance

Complex prompts can generate numerous search results, not all equally valuable. The Reranker Agent ensures only the most relevant documents make it through, prioritizing information most closely aligned with your intended message. It addresses scenarios where the retriever agent provides many sources, including irrelevant or partially relevant ones and where important information may not be among the top documents retrieved.

How the Reranker Agent works:



Acts as a quality
assessor for
retrieved content



Assigns high scores
to documents closely
aligned with the
user's intent.



Filters and prioritizes
documents for
accuracy and relevance.

The highest-scoring documents are then passed to the Context Builder Agent.

Step 4: Concise and Clear Summaries

Once relevant documents are identified, our Context Builder Agent uses a fine-tuned small language model (SLM) to synthesize this information into clear, accurate, and concise summaries. These summaries precisely reflect your intent, using language tailored specifically for your business sector. For example, when adapting the model to specific domains it understands terminology and context meaning, "BP" in a medical context means "blood pressure," not "business process".

These small language models ensure that summaries accurately reflect the intent and critical information of the original documents.

Step 5: Generating Your Presentation

Finally, your prompt and summarized insights are combined to create a professional and impactful presentation. This process ensures every slide is accurate, relevant, and visually aligned with industry best practices.

Contextual Intelligence simplifies and accelerates the presentation creation process, turning complex requests into precise, powerful, and professional presentations making Present the enterprise-grade AI solution for business presentations.

Evaluating Contextual Intelligence in the Market for Business

Evaluating the effectiveness and efficiency of Contextual Intelligence (CI) for business communication (specifically focusing on presentations generated by an AI system) is best done by categorizing elements of CI into the following groups: Industry Expertise, Hallucinations, Functional Accuracy & Data Visualizations, Company Compliance, Structured Storytelling and Hyper-personalization to Audience.

Category Definitions:

1

Industry Expertise

Measures how well the content addresses the user's prompt, terminology, and implicit requirements. Checks for fabricated or unsupported details in the generated presentations.

2

Functional Accuracy & Data Visualizations

Assesses how well the presentation aligns with the provided context and evaluates the accuracy and clarity of accompanying scientific or data visualizations.

Category Definitions Con't:

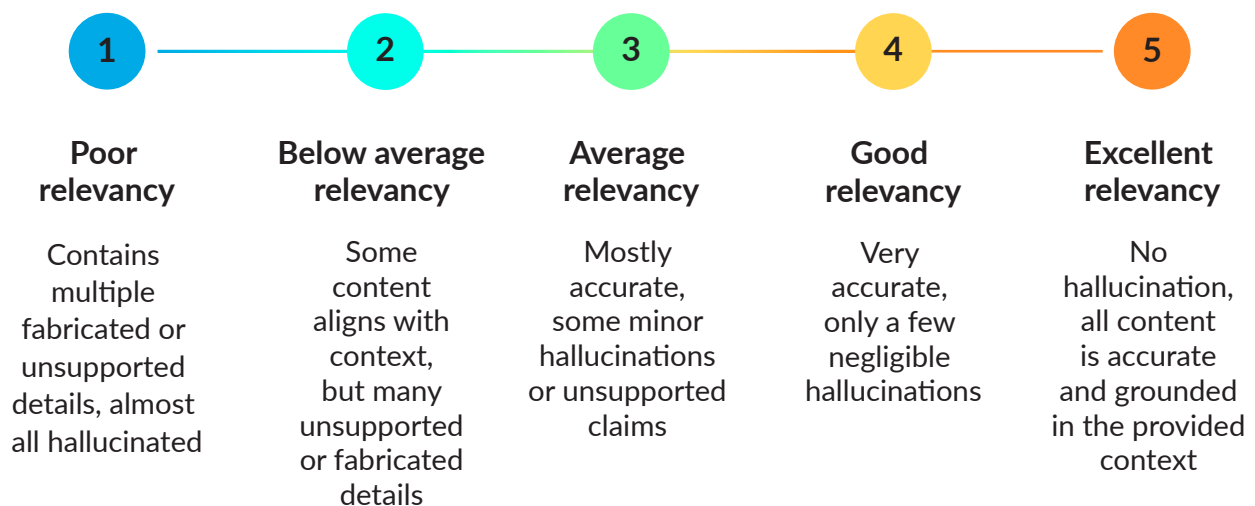
3 **Company Compliance**
Ensure adherence to brand and team-specific guidelines ie. template compliance, color scheme, typography, logo usage, imagery selection relevancy and brand voice.

4 **Structured Storytelling**
Scores the logical progression of information and connections between concepts ie. titles that are relevant and explanatory of the storyline while also capturing team specific terminology and voice.

5 **Hyper-personalization to audience**
Presentation visuals and content aligned to audience preferences or archetypes such as: Architect, Director, Navigator, Performer, Producer, Scholar, Scientist, and Surgeon.

Scoring Reasoning & Considerations

Each category should be scored on a scale of 1 to 5, with specific guidelines for each score level. Reasoning for each score is required, often including lists of deviations, inconsistencies, or hallucinations graded on a sliding scale from poor (1) to excellent (5), as defined below.



Reasoning Guidelines:

- + Checking how well it addresses all explicit requirements in the prompt
- + Assessing the consistency and accuracy of scientific terminology used
- + Determining how well it captures implicit/inferred requirements
- + Providing a detailed list of any deviations, inconsistencies, or missed requirements (both explicit and implicit)

This unified metric combines the three original dimensions (content relevance, textual inference capture, and scientific terminology consistency) into a single comprehensive evaluation of how well the content meets both the letter and spirit of the user's requirements.

Hallucinations

Score the generated slide text based on the presence of content hallucination on a scale from 1 to 5. Use the following reasoning: Determine whether all the information in the generated text aligns with the provided context, or whether there are any unsupported claims or fabricated details. Provide a list of all hallucinations in your reasoning.

Functional Accuracy & Data Visualizations

Score the generated slide text and visualizations based on its context alignment on a scale from 1 to 5. Use the following reasoning: Evaluate how well the generated text incorporates information from the context. Note if any key context elements are ignored or contradicted. Provide a list of all misalignments in your reasoning.

- + Assess the accuracy and clarity of scientific concept visualizations (mechanisms of action, molecular structures, etc.)
- + Evaluate the effectiveness of data presentations through graphs, charts, and tables

- + Identify any scientific inaccuracies, clarity issues, or misleading elements
- + Determine how well the visualizations enhance understanding rather than requiring additional explanation
- + Provide specific examples of visualization strengths or areas for improvement

This unified metric combines the evaluation of both scientific concept visualization and data presentation effectiveness into a single comprehensive assessment of visual information clarity and accuracy.

Company Compliance

Score the slide design based on adherence to team-specific slide templates on a continuous scale from 1 to 5. Use the following reasoning: Explain how well the slides follow the provided team-specific templates. Identify any slides that deviate from required templates and describe the nature of these deviations.

Structured Storytelling

Score the presentation based on how information progresses in a logical sequence with clear connections between concepts on a continuous scale from 1 to 5. Use the following reasoning: Evaluate how logically information is structured and how well concepts connect to each other. Note any instances where the logical flow is disrupted or concepts appear disconnected. Consider whether any specific team or company terminology was included or left out.

FAQ

How many existing LLMs are we using?

When a user submits a prompt, we follow a structured process to ensure accurate and relevant responses:

- **Search & Retrieval:** The prompt is processed to generate search queries, retrieving relevant information from internal and external source
 - **Reranking:** Retrieved results are ranked for relevance using fine-tuned models, ensuring the most accurate content is prioritized.
 - **Summarization:** The top results are summarized using Claude family models, with references included for transparency.
 - **RAG Approach:** This Retrieval-Augmented Generation (RAG) method reduces inaccuracies by grounding outputs in verified sources. This process is central to delivering precise and context-rich outputs.
-

What is the finetuning process?

Fine-tuning is the process of customizing models to perform specific tasks with higher accuracy. Here's how we approach it:

1. **Context Retrieval:** If a user doesn't provide additional information (e.g., text, URLs, or documents), we retrieve and prioritize relevant content using a fine-tuned reranker.
 2. **Fine-Tuning Process:**
 - **Data Collection:** We create datasets with queries and corresponding positive and negative search results, often using publicly available pharma data.
 - **Optimization:** Models are fine-tuned by adjusting parameters like batch size, learning rate, and epochs to improve performance.
 3. **Reranker & Summarizer Tuning:**
 - The reranker has been fine-tuned using this process to rank retrieved content effectively.
 - Smaller LLMs (<7 billion parameters) have been fine-tuned for summarization tasks using synthetic datasets enriched with reasoning and problem-solving examples (e.g., PubMed articles for pharma, CNN Daily for news).
 4. **Training Methods:**
 - Full precision training was applied to Phi and Qwen models.
 - PEFT (Parameter-Efficient Fine-Tuning) was used for Qwen to enhance efficiency.
 5. **Ongoing Benchmarking:** We benchmark our in-house rerankers and summarizers against external models like Cohere Reranker and Claude family models to measure performance..
-

What is the specialized part?

The specialized part of our solution lies in the targeted data and models designed for the pharmaceutical domain:

1. **Domain-Specific Data:** The training data for the reranker and summarizer is curated by our ML team to focus specifically on pharmaceutical use cases.

2. **Specialized Database:** The internal database used during retrieval is tailored to the biopharma domain, ensuring relevance and accuracy.
 3. **Custom Tokenizer Model:** A tokenizer model, trained on user prompts and inspired by trends from historic AutoGenerate data, has been developed. This model identifies key phrases in user prompts and serves as the foundation of our entire contextual intelligence pipeline.
-

What is the presentation layout?

The presentation layout ensures slides are visually and textually aligned to meet user needs:

1. **Best Practice Layouts:**
 - We have many internal experts from the industries we serve. Layouts are created by our experts based on their knowledge and public data available—sourced by researching content relevant to each sector — including whitepapers, case studies, and consulting decks — aligned to the storyline or industry being used.
 - The system selects layouts that best match the user's prompt and storyline.
 2. **Enhanced Content Retrieval:**
 - For suggested layouts, additional searches use layout details to find the most relevant content.
 3. **Structured Slides:**
 - The retrieved content is formatted to fit the layout using a generator LLM, ensuring the slides are clear, impactful, and visually appealing.
-

How does Prezent deal with company information?

We handle company information by using our fine-tuned models and internal workspace to create personalized and branded presentations:

1. **Retrieval:**
 - We access the internal workspace, which includes company-specific data, templates, and branding guidelines.
 - Fine-tuned models ensure the most relevant company information is prioritized.
2. **Personalization:**
 - Retrieved content is aligned with the company's brand, voice, and tone for consistent messaging.
3. **Integration:**
 - Summarized and refined content is used to create presentations that reflect the company's style, using branded templates and team-specific designs.

This approach ensures every presentation is accurate, relevant, and tailored to the company's identity.

How does Prezent use RAG models?

We use RAG models to enhance presentations by retrieving relevant content as needed in L1 Contextualization (Content Retrieval):

- The user prompt triggers a keyword search.
 - Retrieved results are reranked for relevance and sent to a generator (like Claude or a fine-tuned LLM).
 - This fills knowledge gaps in the LLMs used for content generation.
-

Which levels of Prezent's 5 levels of context are in Specialized Presentation Models (SPM)?

Specialized Presentation Models (SPMs) are AI powered and are the backbone of Astrid, our 3-in-1 presentation AI agent. They help in making presentations industry-specific to ensure accuracy, specificity and reduce hallucinations. L1 and L2 Contextualization are in our SPM.

Can SPMs have inference?

Yes, SPMs can definitely perform inference, our proprietary models are served using optimized inference engines.

How are SPMs trained?

Our SPMs training process adapts language models for specific domains like tech, or biopharma by fine-tuning both general models and specialized components using domain-specific data and curated knowledge graphs. A fine-tuned model retrieves and prioritizes relevant content when additional context isn't provided, while training data is gathered from public pharma sources to enhance search accuracy. Optimization techniques, including adjustments to batch size, learning rate, and the use of PEFT, improve model efficiency and performance. Additionally, a "growing database of industry-specific layouts." trains the system to select and format content into visually effective slide templates. The SPM is trained to suggest the most suitable layout for a given set of content. The system performs "additional searches using layout details to find the most relevant content." For instance if the storyline suggests discussion on clinical trial results, then a layout with specific complex graphs like Kaplan-Meier are chosen and content corresponding to it is used for slide generation.

How are our sub-agents trained on empathy, storylines, layouts and brand alignments?

What information are they given?

The sub-agents are fine-tuned to incorporate empathy by leveraging user fingerprint prompts that guide emotional tone and audience targeting. They utilize domain experts' examples to craft coherent storylines that drive effective business communication. For example, a tech company preparing a pitch relies on specific layouts, like a "Feature Comparison" slide and a "Technical Architecture" diagram, to clearly present product details. A generic AI tool misaligns these layouts, cluttering information and simplifying complex diagrams. These errors dilute the presentation's impact and risk losing interest. Best practices for visual communication inform the layout selection process, ensuring that content is both engaging and well-structured. Training of the sub-agents is done to predict the most suitable layout for a given set of content. Proprietary models are invoked for generating layouts.

Additionally, the model is refined to align with specific brand voices. This includes using company-specific branding guidelines, including logos, color palettes, font choices, and messaging styles.