



IPSOS VIEWS

CONVERSATIONS WITH AI PART IV

AI-assisted knowledge libraries and curation,
the search for a trusted output

Ajay Bangia
Jim Legg



At Ipsos, we champion the unique blend of Human Intelligence (HI) and Artificial Intelligence (AI) to propel innovation and deliver impactful, human-centric insights for our clients.

Our Human Intelligence stems from our expertise in prompt engineering, data science, and our unique, high quality data sets – which embeds creativity, curiosity, ethics, and rigor into our AI solutions, powered by our Ipsos Facto Gen AI platform. Our clients benefit from insights that are safer, faster and grounded in the human context.

Let's unlock the potential of HI+AI!

#IpsosHiAi



When our clients are asked by their internal business partners what they know about a specific topic, they often have the information... somewhere. The challenge is locating the right sources and writing a thorough and succinct report that's ready to socialize.



AI-assisted knowledge libraries and curation, the search for a trusted output

Many of us remember going to the library as students to do research. It was as if we were on an information-gathering expedition. We mentally prepared ourselves for a long and arduous journey that included locating resources and scouring through them for relevant nuggets of knowledge.

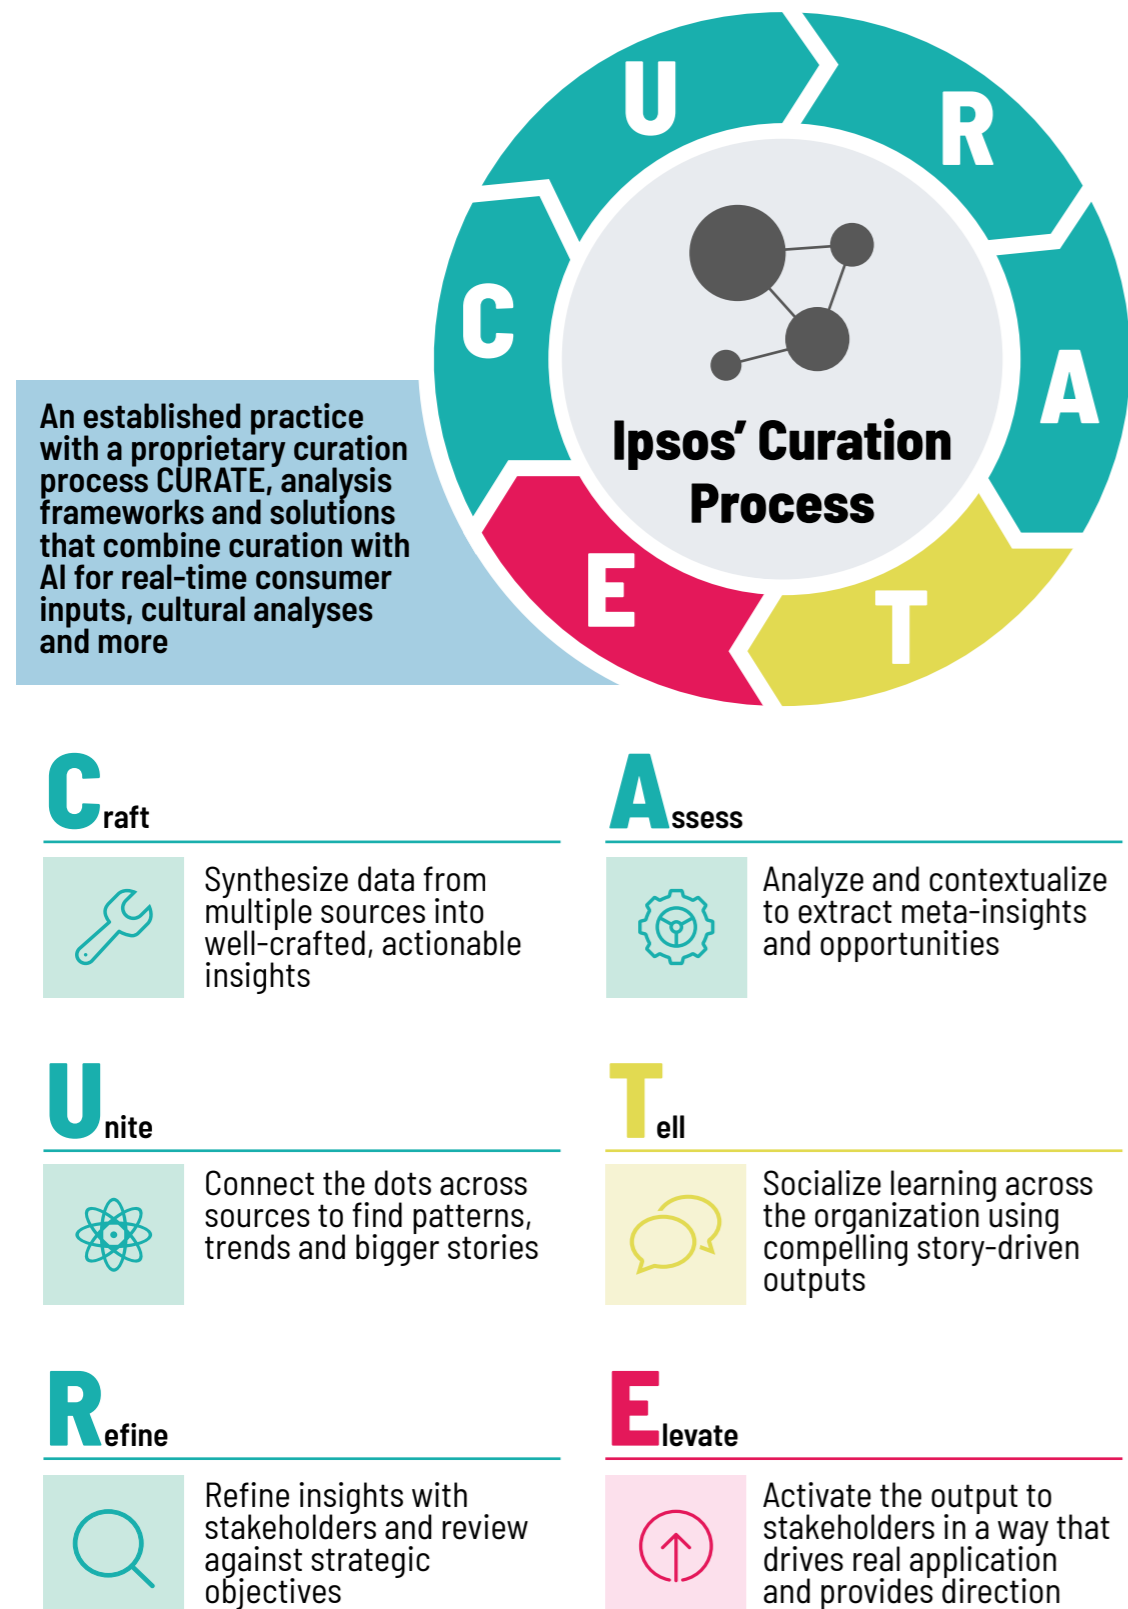
We weren't always successful, and the pile of books to be reshelfed was often taller than the pile of books bristling with pieces of paper marking the pages with relevant information.

After all the time searching, we'd switch gears and start writing. Our teachers wanted us to turn our learnings into a story that was easy to read and understand. The story, in the form of a

report, had to be sourced so the teacher felt confident in its accuracy. We have all been trained, to some extent, to be curators, but some have taken this foundation and made a career of it.

Ipsos executes tens of thousands of studies throughout 90 countries for its clients each year, some with multiple objectives per study. Our clients house not only Ipsos reports and data, but also research done in-house or by other agencies, as well as syndicated data or reports. When our clients are asked by their internal business partners what they know about a specific topic, they often have the information... somewhere. The challenge is locating the right sources and writing a thorough and succinct report that's ready to socialize. The experts who

Figure 1: Ipsos Curation Framework



Source: Ipsos




do this are curators. Ipsos has a global network of curators, trained in a proven curation framework (see left) as well as best practices in AI curation.

Ipsos clients rely on our expert curators not only to curate information but also to advise on how information is organized. Our clients want to stay focused on their core business. They don't have time to become librarians or expert curators of their knowledge base. They need to be able to ask a question and depend on an outside expert who will synthesize, interpret, and help them activate insights.




All of this being said, existing curation platforms are not moving fast enough with AI. The start-up solutions entering the marketplace are missing critical features, don't empower expert curators, or are not flexible enough for unique client needs. We also see the need to evolve our curators into certified AI curators who are experts at applying the latest advances

in Generative AI to the curation process and who know how to apply analytical frameworks to a large corpus of data. The headline: we see speed as the primary benefit that AI could deliver.

Ipsos is looking to disrupt the insights curation space. We have decided to develop our own custom knowledge library and curation solution. We are focused on aligning the best capabilities of Generative AI with our clients' needs. To that end, we asked ourselves four important questions about the future of custom knowledge libraries and curation.

- 
Finding relevant data in an AI-boosted library: With properly ingested unstructured data¹ and proprietary analytical frameworks, can both classic and Generative AI effectively and accurately identify relevant information to accelerate the time to insight?

¹ a variety of reports, transcripts, data tables, videos, etc. that don't follow a consistent organizational structure

- 
Leveraging AI to summarize relevant information: Once relevant sources are identified, can Generative AI effectively and accurately summarize relevant information? Is human validation still needed?
- 
From summaries to a story: How does Ipsos connect AI-generated summaries into a coherent report?
- 
AI-assisted curation benefit: What benefit should Ipsos clients expect from AI-assisted knowledge libraries and curation solutions?

data (not images like tables, charts, and graphs, which Ipsos is currently testing now). Each pilot included two workstreams: a human-only curation approach and an AI-assisted approach. The human-only approach relied on the traditional curation model, in which the human curator reviewed a subset of the documents and then curated information into a coherent report. The AI-assisted approach consisted of leveraging the AI-assisted solution to find relevant information, summarize relevant documents, and combine the summaries into a coherent report. During the AI-assisted workstream, the human-only output and curator were leveraged to validate or identify hallucinations/gaps (in either output), as well as to elevate the findings and add activation ideas within Ipsos's storytelling format.

To answer these questions, Ipsos ran two pilots which are showcased in this paper. Our focus was primarily on textual



Curation is often used for low-error-tolerance tasks, which is why we went on a path to see what we could do to reduce hallucinations further and thus build our learnings.

Finding relevant data in an AI-boostered library

As the foundation, Ipsos is using multiple retrieval augmented generation (RAG) techniques. RAG² optimizes the output of a large language model (LLM) by allowing it to reference an authoritative knowledge base outside its training data sources before generating a response.

While this reduces hallucinations, it doesn't eliminate them. Curation is often used for low-error-tolerance tasks, which is why we went on a path to see what we could do to reduce hallucinations further and thus build our learnings.

Part one: Finding the common denominators

When curating a set of data sources, it's rare to find similarly structured data. The first step in building a library of documents is structuring the diverse data sources. This might include a large set of documents, some created in PowerPoint, MS Word, Google Slides, Google Docs, or Adobe PDF. Each document contains different information and is structured in a specific way. Documents range from audio transcripts of online interviews to presentations that are image-heavy and light on copy. There might be national, regional, or global reports. Some reports might be very detailed. Most will include both text and images. Images might be pictures, or they could be charts, tables, or graphs. And you can be assured that each document organizes information in a unique way.

Early in our testing of the AI-assisted solution's ability to find relevant content, we determined that simply searching within the body of a document's slide or page was not enough. Critical information was often in other places. So it was decided to enhance our algorithms to ingest documents based on their common denominators, including title, body, footnotes, notes, comments, hidden slides, etc.

Part two: Building indexes

Imagine a reference book without an index to help you find information within the book. Ipsos needed to find a way to effectively catalog what was found in each document at the page/slide level, as well as throughout the entire document. We explored two distinct summarization methods: a refined summary approach, in which slides within a document are progressively summarized (i.e. generating a summary of two slides as step one, generating a summary of the first step with the third slide, generating a summary of the third step with the fourth slide, and so on), and a

² RAG stands for Retrieval-Augmented Generation, an AI development technique where a LLM is connected to an external knowledge base to improve the accuracy and quality of its responses. Techopedia February 2024

summary of summaries approach (generating one summary of all slides at one time). We found that when summarizing an entire document, refined summaries performed better.

But even with refined summaries for each page of a document added to aid searchability, LLMs were still not reliable enough in finding and accessing a comprehensive list of relevant documents. Through a manual review, we found that relevant documents were still missing. To enhance the outcome of the curation process and how the solution ranks relevant content, we added custom meta-data based on each client’s needs. Most will include the author, country, or date on which the document originated. Ipsos also added, as meta-data, the summaries that were created from each page/ slide of the documents. Ipsos found that a combination of meta-date (author name, date, etc.) and micro summaries maximized the output’s quality in response to a researcher’s query.

Leveraging AI to summarize relevant information

Part one: What to summarize

Once a set of relevant documents was identified, Ipsos tested whether Generative AI could effectively summarize those sources. The fundamental principle was to pinpoint the documents or slides that best answered a query and then summarize them without significant loss of information. But could we trust that all relevant documents were identified?

A semantic similarity algorithm was developed and refined to identify relevant documents or slides based on the similarity to the prompt. Ipsos learned that, even with a finetuned similarity score, some relevant documents were missed. Because of this, document validation protocols were deployed to ensure a full set of relevant documents was displayed.

Part two: A better question, a better answer

Ipsos found that the more comprehensive the prompt to find relevant information, the higher the success rate. Ipsos leveraged a version of its C.R.E.A.T.E prompt framework (see below). In the request portion of the prompt, we included as many of the meta-data variables as possible, i.e.,

author, date of origination, geography, etc. A clear and specific prompt yields a better response than a vague one.

Without proper grounding techniques, we observed multiple issues. One key example involves hallucinations³ for uncommon industry terminology or acronyms found within the library’s documents. The large language models we tested attempted to define these terms or acronyms with varying degrees of success. We found that less critical statements within documents were often given more weight in the summaries than expected, and, if the content was not well written in curated documents, there was a direct correlation to lower quality summaries or more hallucinations.

Based on these learnings, Ipsos has developed a framework for content curation that includes summary validation protocols to ensure higher-quality summaries. We are also working with curation clients on a key term and acronym memory feature to instruct the LLMs to consider when summarizing. A ‘memory’ is a constantly evolving set of words, terms, and acronyms that have clear definitions to avoid hallucinations in future queries.

³ where a LLM produces incorrect or misleading results that seem convincing


Figure 2: Ipsos C.R.E.A.T.E. prompt framework

A basic example of an Ipsos AI curation prompt



>> From the identified documents, act as a researcher analyzing the main themes around food and drinks discussed. List the top 5 reasons in a table format. Use the Ipsos Censydiam framework to group the themes. In the table, also list why they pick these items on a specific holiday.

Source: Ipsos



When challenged to address broader and more ambiguous questions in a curation report, the human curator was irreplaceable.

From summaries to stories

Part two: AI wins in some areas

Ipsos found that AI consistently delivers when asked straightforward questions. It summarizes basic information and generates solid summaries of what is available, and it provides mostly reliable direction on where to go when more information is needed. Ipsos also has enjoyed efficiencies in leveraging Generative AI to improve human-written content via grammar and sentence structure checks.

Part two: Human storytelling remains critical

When challenged to address broader and more ambiguous questions in a curation report, the human curator was irreplaceable. In multiple pilots, the curator clearly demonstrated the ability to identify and report on nuanced findings. An example is having background client knowledge and applying that knowledge to connect the dots from information uncovered during curation. Lastly, the human was critical in resolving conflicting findings when summaries didn't align as well as editing and connecting the summaries in a coherent and digestible story.

AI-assisted curation benefit

Speed. It's a simple but powerful benefit. In time, our goal is to enable researchers and clients to get the answers they need within minutes, but not at the expense of accuracy. Like with learning to dance, however, we have to learn to trust our


partner. 2024 will be a year of piloting and validation. Yes, speed will be the primary benefit that clients enjoy, but we don't want to deliver summaries (or reports on summaries) with hallucinations.

What's coming next?

Newer Generative AI models, agentic workflows, grounding techniques, specialized tools for LLMs, etc. are rapidly evolving in the market, especially multimodal solution designs that enable text and images (charts, tables, graphs, and pictures) to be summarized alone and combined with the text on the same slide/page. Testing completed to date has shown limited success in the combined

summarization of text survey data tables, and deeply hierarchical datasets. Ipsos plans to continue to test and iterate its custom library and curation solution.

Ipsos is watching new Generative AI releases closely. We expect that enhancements will lead to better multimodal (text and images) results in the coming months, not years.



Key learnings

Part one – Solutions: Human validation features are critical



How knowledge management platforms are built to find information is critical.

While it may be tempting to embrace a start-up promising AI-enabled speed, a tech-driven organization may not have the data science expertise to deliver the best results. Instead, we suggest partnering with an organization with deep expertise in content curation and knowledge sharing with a strong data science organization behind it. The ideal solution maximizes output quality through custom meta-data that includes

slide/page summaries as well as the author, date, relevant geography, etc., and that is customizable to meet your unique needs. Enabling the human curator to easily evaluate whether the documents that been found are relevant is key. The curator should be able to easily validate summaries through a quick review of the original content. You should not trust that the right documents have been identified or that the summaries generated are without hallucinations. Platforms that don't have native validation features designed for the human curator to check the results should be avoided.



Part two – Curation: Keeping expert humans in the loop



Ipsos also recommends a 'human in the loop' approach by partnering with an organization that has expert curators whose role is to help build

strong protocols around the organization of your library. This partner should also be an expert in developing AI-assisted solutions with proven analytical frameworks. The ideal AI curator is one who inherently distrusts what is generated by AI, not because they don't see the value that Generative AI offers, but because the AI curator is focused on delivering truth (not hallucinations) and uncovering insights. The curator should be a competent storyteller who

can connect multiple summaries into a coherent and digestible story. Finding a partner who employs both AI + HI (Human Intelligence) will lead to success.

Ipsos is investing in developing an AI-infused custom library and curation solution designed for enterprise search, document interrogation/summarization, and expert curation. Many of Ipsos' clients have already joined us in a shared learning approach designed to meet their current curation needs AND inform the AI-assisted curation solutions that Ipsos is building. If you'd like to learn more, contact Jim.Legg@ipsos.com and Valentina.Rodriguez@ipsos.com.


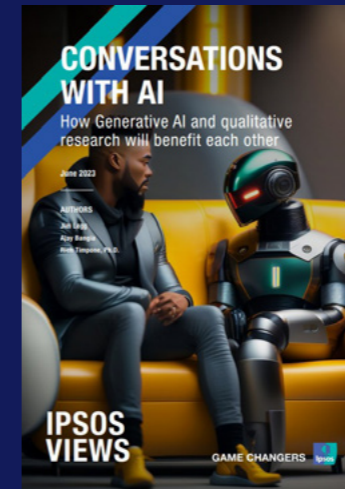


The ideal AI curator is one who inherently distrusts what is generated by AI, not because they don't see the value that Generative AI offers, but because the AI curator is focused on delivering truth, not hallucinations.





Further Reading



AUGUST 2024

CONVERSATIONS WITH AI PART IV

AI-assisted knowledge libraries and curation,
the search for a trusted output.

AUTHORS

Ajay Bangia

Global Scale Leader,
UU, Ipsos

Jim Legg

Global Head of Operations,
UU, Ipsos

CONTRIBUTORS

**Manali Vivek, Robyn Clayton, Valentina Rodriguez,
Karin O'Neill, Luciana Obniski**

The **IPSOS VIEWS** white
papers are produced by the
Ipsos Knowledge Centre.

www.ipsos.com

@Ipsos

