

New Tools, New Roles



A Manager's Guide to Harnessing Generative AI for Marketing Insight

AUTHORS

Oded Netzer

Arthur J. Samberg Professor of Business
Columbia Business School

Simon J. Blanchard

Dean's Professor, McDonough School of
Business, Georgetown University

Nofar Duani

Assistant Professor of Marketing
Marshall School of Business
University of Southern California

Aaron M. Garvey

Associate Professor of Marketing
Haslam College of Business
University of Tennessee

Travis Tae Oh

Assistant Professor of Marketing
Sy Syms School of Business
Yeshiva University

KEYWORDS

Generative AI
Surveys
Experiments
Marketing Research
Text Analytics
Validation



GettyImages/Qi Yang



Picture this: It's Monday morning, and your team just came up with a new product idea. By Friday, you already have concept tests, customer reactions and a full insights deck – all powered by generative AI (GenAI). What once took months now happens in days. GenAI tools can scrape data, run adaptive interviews and summarize consumer sentiment in minutes. For organizations competing in fast-moving markets, this is game-changing. But with great power comes a simple truth: These tools can accelerate good research and flawed research equally fast. The difference lies in how managers use them.

How GenAI works in a nutshell

At the heart of many GenAI systems are large language models (LLMs). These models are trained on billions of words to predict the next word in a sequence. They excel at producing text that is coherent and contextually appropriate. This makes them powerful engines for generating survey items, summarizing reports or writing code. But LLMs are probabilistic. They select the most likely next word, not necessarily the factually correct one. This is why they sometimes “hallucinate” facts, misstate sources and findings, or drift ▶

away from strictly defined concepts. Modern GenAI systems extend beyond the core LLM by layering in additional capabilities. Retrieval-augmented generation (RAG) allows these systems to draw on outside documents, turning uploaded reports or datasets into usable context. Code execution environments allow GenAI systems to not only suggest but actually run code and statistical analyses. These system-level enhancements make GenAI more accessible to users, increasing its potential usefulness for researchers and businesses alike, but they do not change its underlying probabilistic nature, nor do they change the principles of rigorous research. For managers, understanding these basics is crucial. Just as one would not trust survey results without understanding the sampling

method, one should not rely on GenAI without recognizing its limits. In the following sections, we explore how GenAI capabilities can be used responsibly and diligently in different insight-generation settings.

Doing desk research with GenAI

Every research project begins with understanding what is already known. GenAI dramatically accelerates this stage of desk research.

Scanning the external knowledge landscape

A single query can produce a structured summary of dozens of academic articles, industry reports or trade publications, highlighting emerging themes in crowded fields. For early scoping, this speed is invaluable. Yet, the same efficiency creates risk.



Midjourney AI prompt by GCO

“
Treat GenAI summaries
as starting points,
not as definitive sources
of knowledge.
”

“
GenAI is a generator of
alternatives, but managers
must filter and refine
those alternatives.
”

Because GenAI reconstructs information rather than retrieving it directly, it may conflate findings or even cite the wrong sources. For managers, the takeaway is clear: Treat GenAI summaries as starting points, not as definitive sources of knowledge. They can help surface relevant themes and connections, but you must always verify them against the original documents before shaping strategic decisions. Used responsibly, GenAI can compress days of desk research into hours. Used uncritically, it can build a convincing case on shaky ground.

Unlocking insights from internal knowledge

Beyond scanning the outside world, many companies sit on a wealth of underused internal knowledge: past survey data – especially open-ended, qualitative data from large-scale surveys or focus groups –, customer studies, consulting reports or internal competitive analyses. GenAI can help extract patterns and lessons from this fragmented information base, bringing older insights back into active use. When applied to internal sources, however, confidentiality becomes paramount. Uploading proprietary data into public chat interfaces risks exposing sensitive company information. Managers should instead rely on enterprise APIs or localized LLM deployments that guarantee data security. With these safeguards, GenAI becomes a powerful tool for synthesizing years of accumulated knowledge and can turn siloed archives into actionable intelligence for current decisions. ▶

The art of prompting for precision: capturing the right construct

BOX 1

Generative AI can draft survey items, images and other stimuli in seconds – but only if it understands what you want to measure. When prompts are vague, models often drift toward related but different ideas. The result can sound polished yet miss the construct entirely.







Example

A retailer exploring consumer reactions to “misshapen produce” asked GenAI to create items measuring “visual attractiveness.” The system produced statements such as “This produce looks fresh and high quality.” While fluent, these responses shifted the construct from attractiveness to freshness and quality – two very different dimensions.

For visual clues, the challenge is similar: Providing clear definitions in prompts leads to dramatically different outputs. Figure 1 shows what GenAI produced when asked to create a “visually attractive” and “visually unattractive” apple. Providing clear definitions in prompts can lead to drastically different outputs and yield better results.

Drastically different outputs when prompting “visual attractiveness”

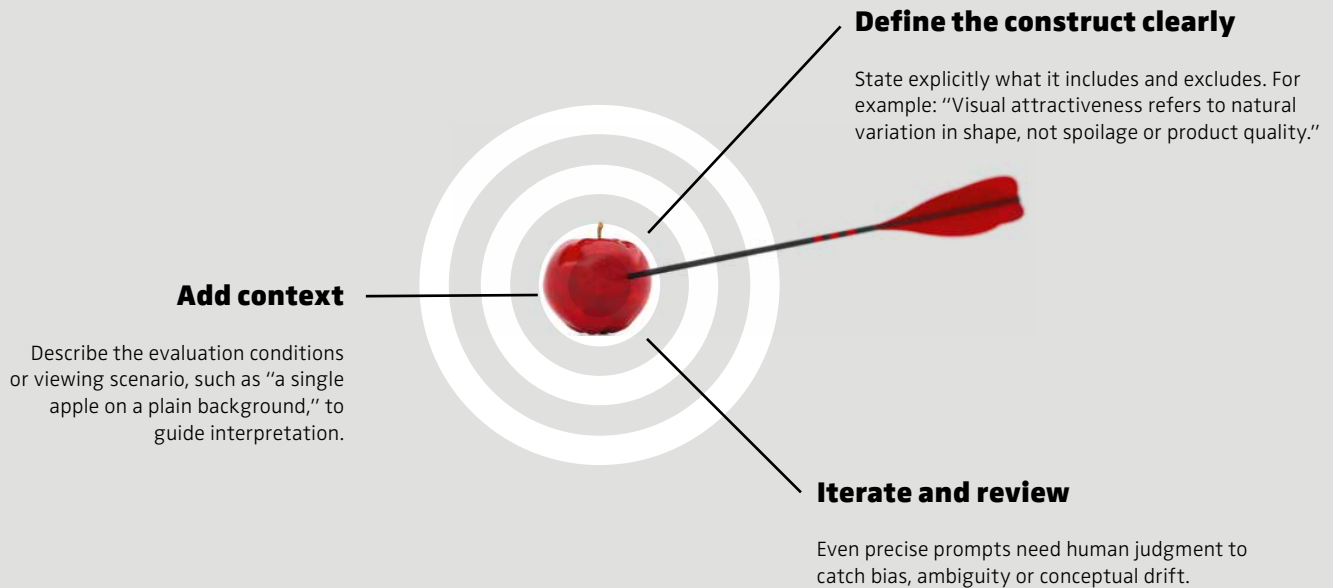
FIGURE 1

	Prompt approach	Attractive	Unattractive
VAGUE	“Create a visually attractive and a visually unattractive apple in two separate images.”		
WITH CONTEXT	Adding to row 1 the definition and context, for instance, specifying “visual attractiveness” as natural variation in shape rather than spoilage or quality.		
PRECISE*	Further specifying background, angle, color variations, surface texture, lighting, etc.		

*Precise Prompt: “A high-resolution, close-up photograph of a single red apple on a plain white background, taken at eye level. The apple has a very asymmetrical shape and is atypical. It features a mix of red and orange-yellow tones, with one large area near the middle where the red pigmentation fades into strong natural discoloration with less polish. Small, scattered discoloration dots are visible across the skin, giving it an organic, naturally imperfect look. The surface is smooth with very visible texture variations, but the apple remains fresh and undamaged. The lighting is soft and even, emphasizing its natural imperfections without making it appear bruised or spoiled.”

How to keep prompts on target

FIGURE 2



iStock/Evgeny Ostroushko

Designing surveys and measures

GenAI offers one of its greatest advantages in drafting survey items, scales and experimental stimuli. What once required time, creativity and often multiple iterations can now be generated in seconds. A single prompt can yield dozens of candidate items or ad copy variations, giving managers a breadth of options that would have been prohibitively costly to produce manually. The challenge is that these outputs are not always conceptually precise, as Figure 1 illustrates. One solution is to guide the model with clear definitions of what the focal construct includes and excludes (see Box 1). Even then, human review is critical to screen for any potential biases or ambiguities. For stimuli like ad copies or vignettes, the same principle applies: GenAI is a generator of alternatives, but managers must filter and refine those alternatives.

Collecting and analyzing data

GenAI transforms data collection and its conversations with respondents, coding and analysis.

Conversational research at scale

Customer surveys have long relied on fixed, scripted questions. GenAI changes that by enabling conversational formats in which an AI agent adapts its follow-ups based on what participants say. Companies can now embed AI interviewers into survey platforms to probe consumer motivations, simulate focus group-style discussions or deliver dynamic treatments in experiments. Such agents can also converse with the researcher during the process of designing or analyzing the research. The

benefit is qualitative depth at quantitative scale: Instead of hundreds of short, static answers, managers receive hundreds of adaptive conversations, each tailored to the respondent. This approach, though revolutionary and useful, can come with risks. Because GenAI generates responses probabilistically, two participants may receive slightly different prompts even when they give the same answer, creating noise in the data that undermines comparability. There is also the chance that the AI goes off-script entirely. To mitigate these risks, companies must pretest extensively, use reinforcement prompts to keep the AI focused and log full transcripts for review. With these safeguards, conversational research can expand what managers learn about customers without sacrificing rigor.

“
Automation does not
remove the need for
validation; it makes it even
more crucial.
”



“
 Without attention to
 accuracy, validity and
 reproducibility, GenAI can
 mislead as easily
 as it can enlighten.
 ”

Coding and analyzing data at scale

If AI interviewers expand the range of what can be asked, GenAI coding accelerates how answers are processed. For decades, managers struggled to classify customer reviews, social media posts, open-ended survey responses and interview transcripts at scale. GenAI now makes it possible to code text, images and videos in minutes, turning unstructured data into usable, actionable insights. But automation does not remove the need for validation; rather, it makes validation even more crucial. AI-coded variables must be tested just like any other: compared against self-reports when measuring internal states, benchmarked against human judges when capturing interpretations, or checked against outcomes when predicting behavior. GenAI can also draft code for analysis in R, Python or any other statistical software and run some of it, lowering the barrier for teams with limited programming expertise. Still, analysts should not stop at outputs produced entirely inside a chat window. Any generated code should be rerun and validated in a dedicated analytics package to ensure accuracy and reproducibility. Used carefully,

GenAI can dramatically expand the scope of both data collection and analysis, as long as managers pair speed with discipline with respect to replicating their workflows.

Accelerate insights – but verify every step

Insights that once took weeks to produce can now be obtained in days or even less. But speed without rigor is dangerous. Misleading desk research can steer strategy in the wrong direction, while poorly designed survey items can undermine segmentation studies; inconsistent AI interviews can compromise the reliability of insights, while unverified use of AI coding can erode reproducibility and generalizability. Managers must therefore adopt a new mindset: Treat GenAI not as an oracle but as an accelerator. It expands what is possible, but it does not replace the responsibility to check accuracy, ensure validity and protect privacy. Generative AI is not just another productivity tool; it is reshaping how surveys and experiments are conceived, executed and analyzed. For managers, the payoff is faster cycles of innovation, broader exploration of ideas and the ability to scale insights that were once bottlenecked by time and cost. But technology alone does not guarantee better decisions. Without attention to accuracy, validity and reproducibility, GenAI can mislead as easily as it can enlighten. The responsibility rests with managers to pair new tools with new rules. Those who do will find GenAI to be a genuine strategic asset, enabling them to move faster, without compromising rigor, in their decision-making. ◀

FURTHER READING

Blanchard, S. J., Duani, N., Garvey, A. M., Netzer, O., & Oh, T. T. (2025).

New tools, new rules: A practical guide to effective and responsible generative AI use for surveys and experiments in research. *Journal of Marketing*, 89(6), 119–139.

<https://doi.org/10.1177/00222429251349882>