



# DON'T JUST ASK THE DATABASE DIRECTLY

Why Adding Generative and Agentic AI to a Fundraising  
CRM is a Dangerous Shortcut for the Nonprofit Sector



**Vered Siegel, CFRE**

Principal, Seagull Partners

## Executive Summary

Generative AI assistants embedded directly into fundraising CRMs are being promoted as transformative technology: a way for fundraisers to bypass technical barriers by asking natural language questions and receiving instant insights. In vendor messaging, these tools promise streamlined workflows, smarter donor engagement, and faster decisions.

We argue that generative AI-driven decision interfaces are structurally harmful when applied to the messy, incomplete, and highly contextual fundraising data most nonprofits carry. Rather than improving decision-making, this combination amplifies data errors, undermines human oversight, weakens professional judgment, and erodes institutional trust over time.



## The Appeal of “Just Ask the Database Directly”

Major nonprofit technology vendors today market AI capabilities that let users interact with donor data conversationally. Here are a few of the promises they make:

- Retrieve contact information and recent gift history via natural language prompts such as “Find John Doe’s contact information and latest gift.”
- Generate draft communications, suggest engagement strategies, and summarize donor histories through simple chat queries.
- Deliver next-best-action recommendations across a donor base.
- Automatically draft personalized emails, engagement plans, and reminders for development officers and major gift officers.
- Generate personalized content such as emails and call scripts tailored to an organization’s voice and audience.
- Assist in drafting donor outreach and engagement plans to save time on communication.

But fundraising databases are not clean systems of record like inventory counts or financial transactions; they are accretive, relational, and policy-dependent objects. They are manageable, but inherently messy.

And we know what happens when the data feeding AI tools is messy, incomplete, or contextually ambiguous: garbage in, garbage out...at scale.

## Advancement Services is a Layer of Interpretation AI Cannot Replicate

Real-world fundraising systems rely on interpretive human mediation. Advancement Services professionals are translators between a natural language question and the way a database tracks data and transforms it into information.

This human interpretive layer is invisible when it goes well but highly visible when it does not. Bypassing it entirely is irresponsible.



AI Agents, as currently designed, cannot replicate the human interpretive layer. Machine learning assumes input and output are stable and sufficient without interrogating underlying meanings. AI bypasses the safeguards provided by human expertise in favor of speed and volume.

In so doing, AI makes itself and consequently our sector vulnerable to common data quality issues:

- Records accumulate over decades and can include outdated info and legacy data from past systems.
- Contacts include individuals, households, and linked giving entities.
- Gift records may have soft credits, matches, pledges, or splits that complicate totals.
- Engagement histories reflect human behaviors, not machine-legible patterns.
- Databases exist in a chronic-but-manageable state of incompleteness because the data we collect is occurring in real time and the updates to the database are not.

Generative AI and AI Agents don't clean your data. AI won't detect your data hygiene issues and instead will operate as though the data "doesn't smell," *because it doesn't have a nose*.

Some AI readiness narratives suggest the answer is simply: "Clean up your data, then adopt AI."

But data hygiene is ongoing, not episodic. Advancement Services professionals manage this continuously, reconciling new gifts, updating contact info, merging duplicates, and contextualizing exceptions. AI systems assume stability in the data they query. When data quality is dynamic and imperfect, outputs can be misleading or even dangerous.<sup>1</sup>

---

<sup>1</sup> See Appendix A: Data Quality is a Process, not a State of Being



## Vendor-Promoted Queries vs. Real-World Data Behaviors

Below is a quick chart showing the kinds of queries vendors encourage and the risks those queries pose in the context of poor data quality.

What Fundraisers Are Encouraged to Ask	What the Fundraiser Thinks It Means	What Poor Data Quality Introduces	Real Risk When Action Is Taken
“Who has a giving anniversary coming up so I can call them?”	A clean list of donors at a meaningful moment	Multiple possible definitions of “anniversary,” outdated gifts, duplicates	Wasteful, embarrassing outreach; donor irritation
“What is John Q. Donor’s email?”	Correct, current contact info	Outdated or inaccurate email, mismatched household fields	Failed outreach, privacy risk, donor frustration
“Summarize a donor’s giving history and suggest next steps.”	Intelligent stewardship guidance	Missing soft credits, splits, legacy gifts, assumption that giving history is the only behavior that matters for strategizing next steps	Misaligned ask amounts; stewardship errors
“Identify our best major-gift prospects.”	Strategic prioritization	Multiple possible definitions of “best”; biased or incomplete history, inconsistent metrics	Misallocated portfolios; missed revenue
“Write a follow-up email to event attendees.”	Personalized, efficient outreach	Incorrect attendance, wrong names	Embarrassing errors; donor alienation

*(This chart builds on examples documented by vendors and real sector CRM behaviors.)*

## A Cascade of Compounding Consequences

AI doesn’t produce “bad results” and then stop and walk away leaving you to pick up the pieces. It keeps on going, amplifying systemic flaws in a self-perpetuating and increasingly harmful cycle:

1. Misleading outputs lead to decisions based on flawed assumptions.
2. Decisions get recorded back into the CRM as actions or notes.
3. Those records become inputs for future queries by a human or AI.

We call this “data entropy”: information decays cumulatively. The result isn’t just a bad list every once in a while (though it will absolutely result in those); it’s a bad



pattern of decision-making increasingly detached from best practices and increasingly embedded into institutional memory.

And anyone who works with data knows that end users tend to trust systems that produce beautifully formulated responses, even when the underlying logic or data is flawed. In cognitive science and human–AI interaction research, results that sound plausible are often mistaken for results that are correct.

Generative AI lacks a built-in sense of “epistemic humility;” in other words, it does not label its own uncertainty. Only humans do that.

AI is also very, very good at mimicking widely accepted frameworks for presenting data and making persuasive arguments and will apply those frameworks expertly atop its responses regardless of their quality.

Humans, on the other hand, *can* label their uncertainty when we need to pump the brakes to pause, question, and verify. It can be tempting to see this as “friction,” but it’s important friction. Removing human friction (the person who pauses, questions, verifies) creates a path toward blind trust rather than informed judgment.

## How Generative and Agentic AI Harms Fundraisers and Nonprofits

The promise of generative and agentic AI in fundraising is framed as empowerment: faster answers, fewer bottlenecks, and less reliance on intermediaries. These systems are marketed as conversational assistants that can interpret donor data, generate insights, recommend actions, and in some cases initiate or automate downstream tasks. In practice, unexamined use of generative and agentic AI often produces the opposite effect: it undermines fundraisers’ credibility, judgment, and professional dignity, while simultaneously degrading the systems they depend on to succeed.

At the individual level, fundraisers are exposed to new and largely invisible risks. Generative AI systems produce fluent, confident responses even when underlying



data is ambiguous or factually incorrect. Agentic AI systems go further by acting on those responses: prioritizing prospects, drafting communications, updating records, or shaping engagement strategies without human interpretation at each step.

When an AI system confidently provides poorly contextualized information (even simple errors like an outdated email address or a misinterpreted “giving anniversary”) the fundraiser bears the consequences. Donors do not experience these failures as “AI hallucinations.” They experience them as carelessness and lack of respect.

These mistakes also generate compounding technical debt. Actions taken based on AI-generated interpretations such as emails sent, calls logged, plans adjusted, and scores recalculated become embedded in the CRM itself. Agentic systems may then treat those outcomes as new “signals,” reinforcing errors through feedback loops that are difficult to detect or unwind. Corrections are rarely traced back to the original AI-mediated decision.

There is also a moral and vocational cost. Fundraisers are increasingly asked to defer their judgment to generative or agentic systems they do not control and cannot fully audit. When professionals are expected to act on machine-generated recommendations they know may be flawed, they are placed in an ethical bind: comply and risk donor harm, or resist and risk being labeled inefficient or obstructive. This erodes professional dignity and accountability. When these individual harms are repeated across organizations, the effects become structural.

As generative and agentic AI are positioned as replacements rather than supports, Advancement Services roles erode or are eliminated, removing the very people responsible for interpreting ambiguity, maintaining standards, and sustaining data quality over time. Without this human stewardship, CRM data degrades faster, not because staff are negligent, but because data hygiene is an ongoing, interpretive practice that cannot be automated.

Strategic decisions, from major-gift strategy to stewardship prioritization, increasingly rely on AI-mediated interpretations of flawed data. Agentic systems



accelerate this risk by enabling decisions to be made and acted upon at scale. The result is not better decision-making, but poor decisions made faster. Speed magnifies error; automation entrenches it.

Donors experience the consequences directly. Communications arrive at the wrong time, reference outdated or incorrect information, or miss obvious contextual cues. Outreach feels less informed, less personal, and less respectful because judgment has been outsourced to systems incapable of relational understanding.

At the sector level, this trajectory is dangerous. A fundraising ecosystem that substitutes generative and agentic AI for human stewardship, interpretive expertise, and accountability becomes increasingly fragile. Institutions lose trust in their own data, professionals lose confidence in their judgment, and donors lose confidence in the organizations they support.

## Conclusion: Resist the Shortcut

Generative AI embedded directly into fundraising CRMs is being marketed as a productivity improvement. In reality, when applied to typical real-world fundraising data, Generative AI:

- Accelerates data errors
- Encourages overconfidence without scrutiny
- Weakens human stewardship
- Harms donors, fundraisers, and institutions alike

Vendors push their technology to us because they've designed the technology to be desirable and profitable, and they set a tone such that we believe everyone is adopting and we're falling behind if we don't.

But we in Advancement need to stop and think who is calling the shots when it comes to new tool adoption: are we reacting to a real need, or to perceived (and externally implanted) needs?

Decisions made under pressure to innovate ("or die") should not result in a rush to adoption and must preserve the human systems that make responsible decisions possible.

**Let's stay connected:**  
**[visit seagullpartners.net](https://seagullpartners.net)**  
or **[follow us on LinkedIn](#)**.

## Appendix A

### Data Quality is a Process, not a State of Being

Some AI readiness narratives suggest the answer is simply: “Clean up your data, then adopt AI.”

There are three fundamental flaws with this narrative:

1. **Data cleanup is never “finished.”** Inherent complexity and an ongoing build mean there is always more ambiguity to resolve. Data quality is ongoing, not a preparatory task.
2. **Cleanup projects don’t build stewardship culture.** Data governance, consistent workflows, quality standards, and accountability are cultural prerequisites; the work done continuously by advancement services staff. A standalone clean-up project will not establish that culture.
3. **Cleaning without oversight still yields risky AI output.** Even relatively clean data can contain contextual nuances that require human interpretation. AI does not validate meaning; it only extrapolates patterns. Without human judgment, even “clean data” can support plausible sounding but incorrect recommendations. Is this a risk your organization is willing to take?

## References

- “AI Tool for Nonprofits - Engage by Kindsight.” Kindsight, November 6, 2025.  
<https://kindsight.io/engage/>.
- Blackbaud. “Blackbaud Showcases AI-Powered Innovation and Intelligent Workflows at Bi-Annual Product Update Briefings.” Blackbaud, November 24, 2025.  
<https://www.blackbaud.com/newsroom/article/blackbaud-showcases-ai-powered-innovation-and-intelligent-workflows-at-bi-annual-product-update-briefings>.
- Blackbaud. “What’s New in Blackbaud Raiser’s Edge NXT® - December 9, 2025.” Blackbaud Community, December 2025.  
<https://community.blackbaud.com/discussion/84284/whats-new-in-blackbaud-raisers-edge-nxt-december-9-2025>.
- Inel, Oana, Tim Draws, and Lora Aroyo. “Collect, Measure, Repeat: Reliability Factors for Responsible AI Data Collection.” arXiv.org, September 27, 2023.  
<https://arxiv.org/abs/2308.12885>.
- Virtuous. “Virtuous Unveils Major AI Advancements to Help Nonprofits Personalize Fundraising at Scale.” PRWeb, August 20, 2025.  
<https://www.prweb.com/releases/virtuous-unveils-major-ai-advancements-to-help-nonprofits-personalize-fundraising-at-scale-302534408.html>.