

The Rise of Procurement: The Evolution from Chatbot Assistants to Strategic Copilots and Agents

By Pierre Mitchell

Sponsored by:



The Rise of Procurement: The Evolution from Chatbot Assistants to Strategic Copilots and Agents

The world of business operations, and procurement in particular, has become unbearably complex based on outside-in risks in supply markets and the broader volatile global geopolitical and environmental landscape. The situation also isn't helped by the internal complexity of organizations, processes and niche business systems littering the corporate landscape. The moment of truth and action has arrived: technology must become an enabler rather than a headwind for procurement organizations to tackle this complexity and free itself to evolve its value proposition to the business.

Recently, procurement organizations have been working in earnest to simplify the complexity of processes, policies and systems for the benefit of budget owners and casual requisitioners who find that procurement is too complex to easily engage with, which in turn leads to 'maverick' behaviors that destroy economic value and invite business risk.

CPOs are increasingly pursuing the goal of simplification by using evolving new orchestration capabilities for fine-grained automation of highly tailored operating models that overlay a complex systems landscape. They are also pushing a more simplified UX (user experience) and self-service capability to engage stakeholders, suppliers and other third parties. The goal is for these partners to help themselves while also freeing up procurement to pursue more aggressive automation within procurement (e.g., towards autonomous sourcing execution or Center of Excellence enablement) and thus create the five to tenfold 'ROI' that procurement organizations have historically delivered (which is an incredible return on capital, of course!) but are struggling to maintain as 'low-hanging fruit' savings opportunities have already been harvested.



What is 'orchestration' anyway?

The term 'orchestration' denotes a next-generation process management and automation approach that combines workflow/BPM (Business Process Management), system integration and even things like no-code/low-code (to democratize business rules management to end users) and embedded analytics (with AI included of course!) to orchestrate work in end-to-end cross-functional processes — including B2B processes with business partners too.

The Rise of Procurement: The Evolution from Chatbot Assistants to Strategic Copilots and Agents

In fact, this focus on UX is the top-rated non-process-specific driver for digital transformation by the enterprise, as recently confirmed by over 2,000 C-level executives by Economist Impact. UX is also second overall when process-specific capabilities are figured in — with spend visibility coming in at first.

Improved UX takes many forms, but the biggest recent trend is the explosion of Generative AI (GenAI) capabilities that have been served up in standalone chat bots like ChatGPT, but are also increasingly getting “baked into” the UX of business systems to make them more ‘intelligent’. Embedding intelligence directly into the UX is a trend that has manifested obviously in AI-powered intelligent assistants that leverage the power of large language models (LLMs) for end users to interact with directly and in context of the business process — SAP’s AI copilot Joule is a good example. Applications like chatbots and AI agents have captured user imaginations but have only whet the appetite for corporate users to use them more securely and more impactfully. As we’ve written previously:



“The real power comes not from building an LLM from a massive corpus of internet data that was assimilated to assemble a general-purpose chatbot without context. Rather, it comes from how rich highly-abstracted data-prediction models ... are increasingly training themselves to squeeze more value out of existing business data and systems.”

Moreover, chatbots have historically been limited in their capabilities and even viewed as annoyances — Microsoft’s Clippy assistant from 20 years ago comes to mind! They could only understand simple queries. Users had to be very careful in inputting the right prompt (which is improving but still an issue), and they were only intended for predefined and scripted interactions, making the supposedly conversational interaction a very unnatural and unusual one (e.g., the experience of ‘here is what you can ask me’ and/or a narrow set of clickable elements paired with a bail-out general purpose inquiry form).

Even so, the first generation of chatbots planted the seed for GenAI. Many use cases that people envisioned back then are now possible, and the technology sector learned a lot from these experiences in terms of what works and what doesn’t. With GenAI, we now have collaborative copilots and agents that

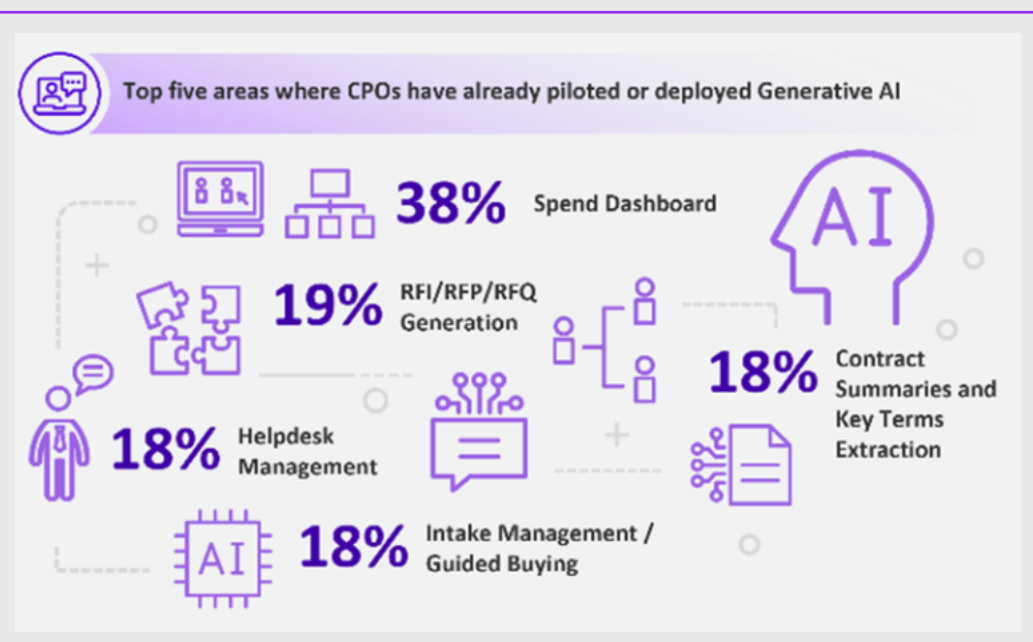
The Rise of Procurement: The Evolution from Chatbot Assistants to Strategic Copilots and Agents

know much more (from a public or private corpus) than past virtual assistants, understand language better (instant out-of-the-box multi-lingual presentation), are context-aware and have memory (in-context engagement/experience that leverages content from policy documents, FAQs, previous service requests, process/procedure information, training information, content from 'digital adoption platforms,' etc.).

And, more importantly, IT departments don't have to program a script for them since these systems are increasingly able to address most queries/requests from ad hoc employee stakeholders, professional procurement staff and suppliers alike (and support a mix of public and private datasets like past communications, internal policies, contracts and training documents).

GenAI can dramatically improve the quality of the process here from both a user and an outcome perspective (improved user requirements capture/intake) by using pre-trained intelligent assistants to ask tailored/optimized interview prompts and dynamically orchestrate workflows and integrations. Since GenAI is focused heavily on language models that are a natural fit for a chat-based UX, the use cases for improving the casual stakeholder user experience are obvious, but GenAI is much more than just enabling natural language in a UI (user interface) to improve the UX.

The result aims more at amplifying the value of existing procurement processes. According to a recent field study by Deloitte about GenAI adoption, CPOs have already deployed GenAI in spend dashboarding, Rfx generation, helpdesk management, intake management and contract management (e.g., where the language models help users understand and converse in legalese!).



Source: deloitte.wsj.com

These top 5 GenAI use case areas also translate nicely to supporting modern GenAI-based functionality:

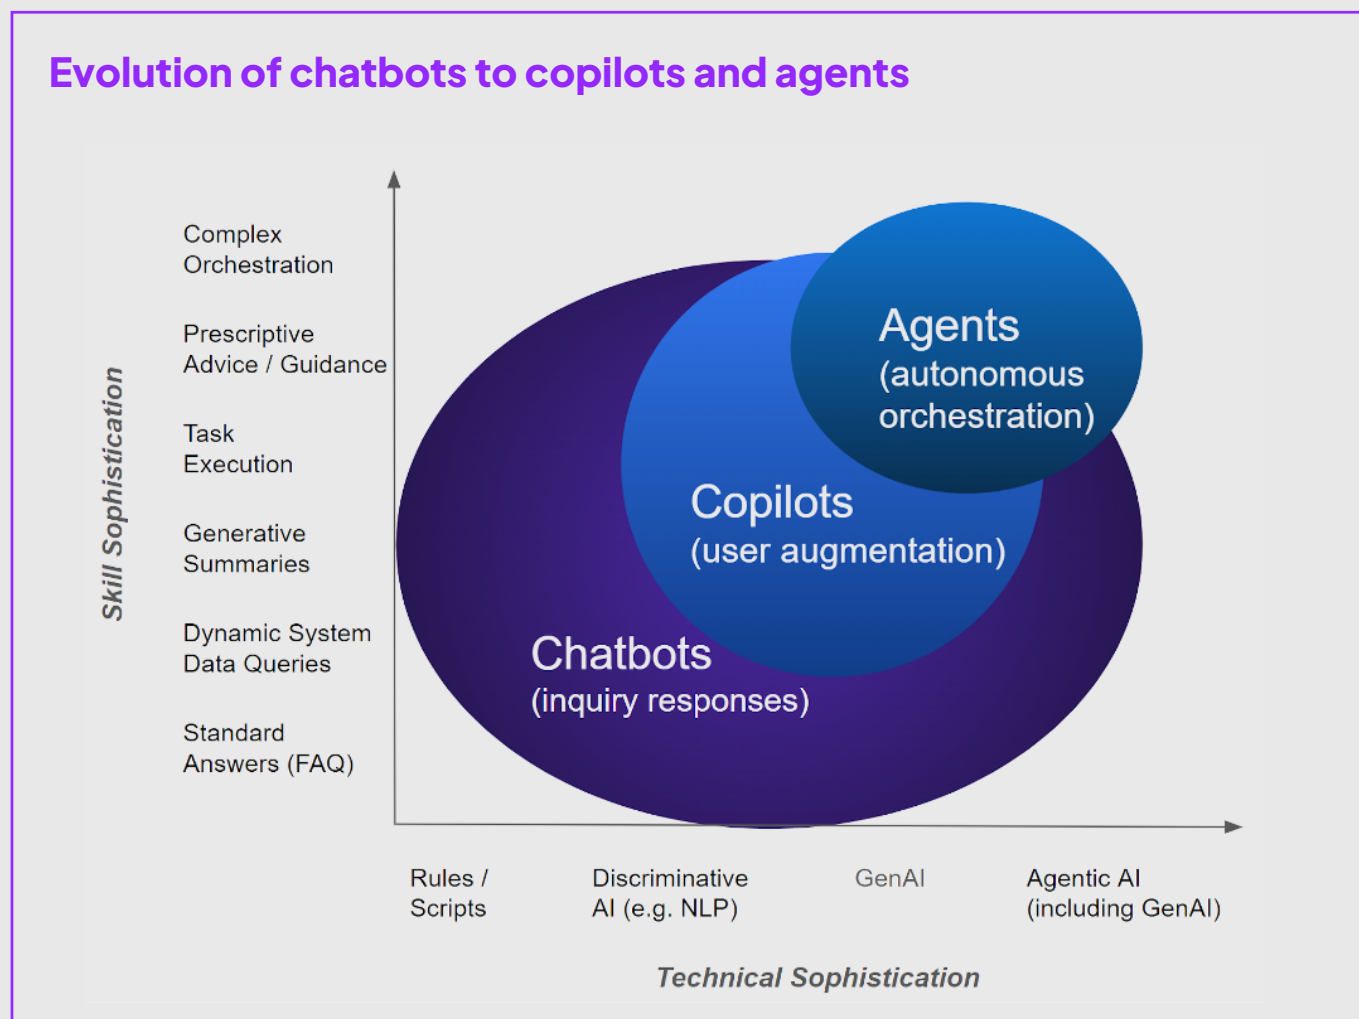
Generative AI adoption areas	How intelligent assistants can help
Analytics and dashboards	Chat with the app to query the latest data on spend, suppliers, S2P process status, job postings for external workers, invoices, etc. Storytelling to put 'raw data' into perspective, make it more digestible.
Text generation	Generate content like RFx documents (buyer side and supplier side), contract clauses, draft SOWs, prefill documents/ fields (e.g., supplier profiles), etc.
Summarize complex information	Summarize contracts (e.g., who is the supplier, what's the amount, what type of contract is it and what are the key dates, obligations and risks).
Answer user questions	Answer help desk questions for a procurement team based on the content of an entire document repository by finding the relevant documents and generating a helpful response (e.g., FAQs, contract terms or project schedules).
Requisitioner intake disposition and guidance	Guiding employees to preferred items/ suppliers/service providers and buying processes.

For example, chatbots can search for and retrieve information by focusing on better search experiences and subsequent user guidance/handoff. They can discover and access content and apps (permissions allowing) that relate to the user search prompt.

This 'talk to your apps' concept is made possible by GenAI working in concert with heavily API-embedded applications. In the same Deloitte study mentioned earlier, the number one GenAI capability that CPOs want to explore next is "Enabling Functions for Data and Analytics." GenAI can be embedded in standalone analytics applications, but the real value is in embedding them as a persistent UX across various applications. This is where chatbot functionality really shines — to serve up real-time analytics to end users (and maybe eventually automate the subsequent low-value activities).

The Rise of Procurement: The Evolution from Chatbot Assistants to Strategic Copilots and Agents

This evolution from traditional chatbots to more advanced augmentation-focused copilots and then to the most complex and discussed area in AI right now: Agentic AI that focuses on intelligent agents (with varying levels of intelligence that can range from narrowly scoped skills to agents that can adapt and reason to handle a broader set of user needs) that use AI on process execution/orchestration rather than just generating prescriptive content to help users help themselves:



We'll dive into how these increasing levels of technology- and skill-based sophistication in these intelligent assistants impact procurement, but for now, a simple way to view copilots is to compare them to a well-understood metaphor ... airplane cockpits!



What an airplane pilot may ask a human copilot	What a procurement practitioner may ask a virtual copilot
What's the status of our revised flight plan?	<p>Are we on track to hit our savings targets? Which stakeholder budgets are going to run dry?</p> <p>Are we on track with our staffing targets for this project? Which departments or projects might face a shortage of external workers?</p>
What's the procedure for fixing the failed sensor?	<p>What is our recovery playbook for our supplier that just filed for bankruptcy? What is the process for determining a new category strategy?</p> <p>What is our contingency plan for when a job posting for a critical role isn't attracting candidates or when a key contractor becomes unavailable? What is the process for sourcing a replacement quickly?</p>
Can you communicate with the crew regarding the weather delay?	<p>Please notify the eProcurement project team on MS Teams that we're going to need to accelerate our supplier outreach for getting them connected.</p> <p>Please update our hiring team on the status of job postings and alert them if we need to push new job ads or extend deadlines for receiving applications.</p>
What's the latest on the storm, and what are our options to avoid turbulence?	<p>Beyond cost savings, what are the market risk factors and sustainability factors for this spend category? And how will this affect our sourcing execution strategy?</p> <p>Which platforms are giving us the best application rates for specific job roles, and are there any market trends we should be aware of to improve our job postings?</p>
Can you fly the plane while I take a break?	<p>Please run a sourcing event with our preferred contractors for this marketing service and notify me when all bids are collected. Also, notify me if any suppliers don't respond by the response due date.</p> <p>Please initiate a new contractor/hiring process for this SOW/job posting, including sending out agreements and setting up necessary system access. Let me know if there are any issues with the onboarding process.</p>

The Rise of Procurement: The Evolution from Chatbot Assistants to Strategic Copilots and Agents

Helpful copilots that are implemented across the business application landscape are becoming more commonplace, most famously with Microsoft aggressively building its copilots in its well-known desktop applications. But it's also being implemented at a similar pace by large enterprise applications players in ERP like SAP, B2B in terms of the sell-side (e.g., Salesforce) and the buy-side (e.g., SAP Ariba, SAP Fieldglass, SAP Concur and SAP Business Network). Interestingly, Microsoft and SAP, along with several other players, are currently working together to integrate their digital assistants with each other to exchange information. This is a work in progress, but such integration is needed to help these players integrate their data and workflows beyond just low-level APIs and toward a more sophisticated coordinated orchestration of agents working together toward more autonomous processing.

This increasing level of sophistication of chatbots from traditional roles to copilot roles and toward agentic capabilities is an evolution:

Level of Capability*	Procurement Use Cases
<p>Level 1. Basic chatbots that follow predefined rules or scripts.</p> <p>They generate responses through keyword matching and other basic functions. This includes some connections to knowledge bases through support widgets or bots.</p> <p>Chatbots with semi-intelligent semantic parsing that can identify user intent.</p>	<p>A help desk for internal stakeholders and suppliers. The chatbot offers a predefined summarized response of a procurement policy, FAQs, processes/procedures, procurement contact details, system access information, etc. based on keywords or pre-defined questions.</p>

Level of Capability*	Procurement Use Cases
<p>Level 2: Copilots that ‘get you’ and offer conversational capabilities similar to Siri, Alexa or Google; they can interact via text or voice, provide context-specific help and guide users.</p> <p>Uses native access to data and knowledge repositories and GenAI techniques like Retrieval Augmented Generation (RAG*) to ground responses in a user’s actual business data, improving reliability and relevance of its assistance. This lets users ‘talk to the apps’ through inquiries and basic task execution.</p> <p>*Retrieval Augmented Generation is a technique that supplements user prompts for GenAI applications like chatbots with secure structured enterprise data to produce results that are more accurate, relevant and free of hallucinations.</p> <p>Copilots will increasingly leverage application vendor specific language models built from metadata of their application data models and associated (user-permissioned) data to create fit-for-purpose pre-trained foundation models.</p> <p>Future trend: Knowledge Graphs!</p> <p>Knowledge graphs (semantically modeled data built on graph databases) are beyond the scope of this paper but are a critical emerging technology foundation that helps build Knowledge Models that integrate Large Language Models to the mostly tabular data in relational Data Models of SaaS applications. They are becoming critical to next-generation application architectures. For more information, see SAPs excellent description of the topic and the work that it’s doing (see here and here) to bring the power of LLM-focused GenAI to help users interacting with traditional business applications. This capability is the essential capability of intelligent copilots that ‘speak your business language’!</p>	<p>The copilot understands user intent and guides users to provide guided procurement (beyond basic guided buying of requisitioners) to specific systems/forms, workflows, knowledge and agents within other tools (e.g., creating a job posting in a contingent workforce solution). Such guidance includes:</p> <ul style="list-style-type: none"> - Procurement applications, external workforce apps (e.g., timesheet tracking tools or onboarding tools), adjacent non-procurement applications (e.g., travel planning tools or expense reporting tools) and desktop tools from Microsoft, Google, Slack, etc. - Conversationally access data from systems (e.g., spending/savings, supplier performance, supplier risk status, status and rejection reasons and compliance status). - Asking a basic question and receiving a summarized response and a link to the source file to learn more. - Invoking simple tasks like approvals, workflow/team notifications, run reports, auto-invite suppliers to meetings, etc. - Guided sourcing to appropriate workflows and supporting tools. Examples include category strategy, RFx tools, supplier discovery and market intelligence sources. - Direct external workforce buyers to relevant workflows and tools to manage staffing needs, such as job posting platforms, skill-matching systems and talent discovery solutions. - Guided contracting to best-match playbooks, templates, authorized clause libraries (or lower-level agreement text) or SOW-generation tools. - A guided supplier management experience to help users find/use appropriate workflows and tools depending on context (e.g., review schedules and tracking compliance requirements). - General user guidance and ‘micro nudges’ to focus attention on outstanding work items in queues based on the highest value and risk.

Level of Capability*	Procurement Use Cases
<p>Level 3: <i>The next generation of copilots with built-in agents that can autonomously perform complex multi-step workflows without being explicitly guided by the user.</i></p> <p>AI agents are specialized experts that can achieve a wide range of user-defined goals through planning, self-reflection, reasoning and tool use. They allow users to request assistance with a high-level objective and leave it to the agent to determine and work through a set of steps to achieve it. A key attribute of such agents is ensuring users retain oversight and approval of the agent's actions.</p> <p>Some copilots, including SAP's AI copilot Joule, go a step further to enable their embedded agents to adaptively collaborate. These copilots allow agents from multiple business domains to self-assemble into ad-hoc teams and collaborate to achieve a mutual goal. This multi-agent collaboration allows copilots to take on even more complex workflows that span and combine business functions, such as finance, supply chain and procurement.</p> <p>We also expect that copilots will increasingly be able to learn user interactions with business systems akin to 'next-generation RPA' (Robotic Process Automation) that offers an alternative way to exchange data other than low-level APIs (that don't even exist in many applications). We also expect them to rationalize the best response over time based on myriad previous experiences. In other words, 'the UI becomes the API,' and it can also help with capturing and reusing user knowledge to augment users rather than replacing them — similar to how 'adoption management' tools (e.g., WalkMe) can help.</p>	<p>A new generation of copilots with multi-agent collaboration will help orchestrate S2P processes and tools which are not easily preconfigured (e.g., too many possible situations and resolutions) and normally require multiple steps to complete.</p> <p>Through intelligent process design, AI agents will work together as virtual teams to apply reasoning, planning and reflection to take on this next frontier.</p> <p>Individual agents themselves have limitations in the range of use cases they can address — much like how hard it is to train a human to be an expert in all aspects of a business. This is where multi-agent systems come in. They allow individual agents to combine their talents and work together to do more than any one agent can accomplish — like a team of humans that each bring certain domain expertise.</p> <p>For example, a sourcing system with a Level 3 AI agent would be able to determine a category strategy based on strategic reasoning (rather than picking from a static strategy repository). It would also recommend best-fit negotiation strategies during a sourcing event and maintain a comprehensive data model, such as supplier master data, in a knowledge graph that is aligned with various local data files across the enterprise and work with multiple agents to confirm key data points. This allows the GenAI-based intelligent copilots with built-in agents space to converse and manage data that is contextually relevant to the user (e.g., role or business unit).</p> <p>Another example is an external workforce solution that suggests optimal staffing strategies. Rather than relying on static role descriptions, it analyzes the project's scope, required skills and timeline to recommend the best types of contractors and provides strategic guidance on competitive hourly rates.</p>

*Criteria adapted from the 2024 Spend Matters Solution Map evaluation benchmark for Source-to-Pay technology

At Level 3, a bot could also generate the content of a supplier questionnaire based on internal and external corpus data (linked to multiple customer and regulatory requirements), create the questionnaire by generating an XML/CSV with code interpreter-like capabilities and upload it in the solution. Afterwards, it could trigger a workflow for the validation and activation of the new questionnaire.

The Rise of Procurement: The Evolution from Chatbot Assistants to Strategic Copilots and Agents

What you should look for in a copilot solution

Chatbots are already here and are getting baked into your browser and everyday digital experiences. And they are increasingly learning from you and getting trained by you, so knowing how they use your data and who is powering them is important. Beyond the big hyperscalers like Google, Amazon and Microsoft, players like OpenAI, Anthropic, Facebook and others are all serving their LLM-powered AI chatbots and copilots to the masses (and to the ecosystem of major application providers), and this doesn't even include the exploding open-source scene.

For everyday business application users/buyers who are seeing their favorite S2P application providers introduce chatbots or copilots (e.g., SAP's AI copilot Joule) on top of their current application stack, there are a few buying criteria to consider to help future-proof your successful adoption. The below list serves as a starting point before adopting one:

- **Require transparency** from your provider: how are the chatbots or copilots built in terms of language model creation (and sourcing of them more broadly); what fine-tuning and utilization via API or other techniques (e.g., [RAG](#)) is there?
- **Ask about integration** into the multiple AI/LLM/chatbot/copilot frameworks from the hyperscalers and mega application providers that you currently work with.
- **Understand how the copilots work in concert with other elements of the system** platform in terms of UX, low-code/no-code, workflow, integration, core data modeling and broader analytics and AI (e.g., discriminative AI). Your IT department will of course have some perspectives on this too!
- **Understand the portability of the capabilities** that you build into the chatbots/copilots. For example, ask your solutions vendor whether there is a focus on using knowledge graphs and graph databases to provide more explicit linkages between the LLMs (and associated action models) and the relational databases in the SaaS application data models/objects. Piloting with startups is fine, but scaling up means future proofing!
- **Get clear guidance on the roadmap** for chatbots/copilots/agents from the basic levels shown in this model to more sophisticated ones. Understand clearly how these can be

The Rise of Procurement: The Evolution from Chatbot Assistants to Strategic Copilots and Agents

permissioned incrementally at the cadence of your ability to pilot and scale them to bring value to your organization. And, of course, understand the commercial implications of the copilots, especially given the nature of token consumption within LLM-based environments.

- **Partner closely with your IT organization** to align procurement's business-centric requirements for chatbots, copilots and agents with IT's broader technology requirements.

For solution providers, we see also opportunities for copilots to continue evolving as more seamless and natural elements within an overall UX:

- **Structured onboarding:** A short onboarding sequence or tutorial can help users understand what the bot can do and how to interact with it.
- **Helping users understand the bots' capabilities** (and limitations) that are rapidly evolving towards more nuanced human-like communication. This includes proactively giving the users options that include sensing user struggles and redirecting them to humans (e.g., via a procurement/supplier/vendor support desk).
- **Personalization and context retention:** A conversational system should remember past interactions and personalize prefilled data and confirmations/responses based on a user's profile (e.g., location, role and organization). Copilots should also speak the 'language' of the user in terms of written/spoken language, company-specific terminology, role-specific semantics (e.g., 'Legalese' with Legal operations staff will be different than consumer-like engagement with casual requisitioners), etc. so that users don't have to struggle with the chatbot or become AI prompt engineers!
- **Inclusion of visual cues:** Copilots can offer suggested (clickable) commands or questions in the chat interface. This helps guide users by showing what kinds of queries the bot can handle and reduces the guessing game and speeds up the process.
- **Offering multimodal feedback and answers:** Even in a text-based conversation, the system could provide visual feedback, such as showing suggested buttons (quick-reply buttons after a user's input, such as 'Yes,' 'No' or 'Check status') within the chat window, progress bars or other visual cues. It could also display content beyond linear text-based answers and provide answers in various visual formats (charts, tables, links, etc.)
- **Error-resistant design:** Bots should be designed to handle mistakes by offering prompts or corrections. When users input something the chatbot doesn't understand, the bot should suggest other ways the user could phrase their question or provide clarification. The system should use deep learning to learn from these interactions and improve its future responses.
- **Seamlessly offer up [agentic] actions to take with the user human in the loop to trigger an action.** For example, including buttons like 'summarize with AI' or 'generate RFI questionnaire' in the solution's UI can then invoke further dialog (e.g., to gather any minimal additional needed data) and then invoke the relevant solutions' UX and/or triggering needed APIs — all in the context of the orchestrated workflow.

There's still plenty of work to be done here, but it's encouraging to really see the technology getting real and getting implemented

The Rise of Procurement: The Evolution from Chatbot Assistants to Strategic Copilots and Agents

Looking Forward

As procurement chatbots evolve from simple assistants to strategic copilots, their potential to enhance processes through intelligent data orchestration and system integration seems undeniable. Of course, there are issues to resolve: poor enterprise data quality; rapidly changing technology capabilities in GenAI; project prioritization (and business case development); seamlessly integrating chat-based experiences with other UX strategies; digital skill gaps with procurement talent; and others.

Even so, by leveraging Agentic AI, procurement systems can not only improve chatbot interactions but also enhance the entire workflow through sophisticated orchestration. With the right approach and evolving technology capabilities, chatbots will evolve from helpful conversational interfaces that are bolted on top of current traditional forms-and-fields-based interfaces into more seamlessly integrated tools that drive intelligent decision-making and autonomous operations within procurement.

Procurement professionals must carefully evaluate the integration and transparency of chatbots, copilots and agents within their existing systems. As these become increasingly embedded into applications, understanding how they use and process data, how they integrate with existing platforms and their long-term scalability will be key. The journey toward more advanced intelligent assistants, empowered by Agentic AI and capable of addressing data quality issues at the source, has only begun. Organizations that embrace this evolution and strategically adopt these tools will be well positioned to capitalize on the full potential of AI-driven procurement processes.

SAP helps bring simplicity and clarity to all types of spending from all across a company. With spend management and business network solutions from SAP, companies can guide employees to smart spending decisions, partner with the rest of the organization and work closely with suppliers to make sure every transaction is a strategic investment.

Interested in speaking with an
SAP representative?

[Contact us today](#)

