When Your Virtual Agent has got to be right

kama.ai: Trust. Empathy. Accuracy.

kama.ai

Knowledge Management in the Virtual Agent Era:

How does Conversational AI ensure your knowledge base is accurate, scalable, and accessible?

https://kama.ai

Virtual Agents

Conversational AI is increasing in popularity. Commonly called Virtual Agents, this technology is evolving with the changing needs of enterprises. Large organizations need to efficiently disseminate timely, reliable, procedural information, policy guidelines, and other material – accurately. In fact, the global knowledge management (KM) market is expected to grow from $773.6Billion USD in 2024 to $3.5Trillion USD by 2034.[1] That's an astounding compounded annual growth rate (CAGR) of 16.5%.[1] With such growth in the KM market, there is a burgeoning opportunity to combine enterprise KM with Virtual Agents. This melding of technologies will help employees, prospective customers, and existing clients access the organizational information they need. Not only will this boost productivity, it will transform customer support by providing accurate, and better informed personal assistance.

KM and AI Intersect

Neither knowledge management (KM) nor knowledge sharing are new. Technical approaches to knowledge management have been enterprise goals for at least 3 decades. Further, knowledge keeping and sharing is as old as humankind. What is new to KM is using a digital conversational medium (chat or voice) to accelerate onboarding and activating knowledge. Merging conversational Virtual Agents and the KM function can significantly improve employee work productivity. This done accurately and efficiently, can not only improve efficiency but cut costs, reduce risk of errors, and curtail flawed decision making.

Supporting this goal, Artificial Intelligence (AI) is increasing the ability to correctly understand natural language user requests. Meanwhile Large Language Models (LLMs) and Retrieval Augmented Generation (RAG) improve the ability to efficiently research, summarize and draft different documents into human readable answers. Having stated this, there remain weaknesses in using LLMs. These are probabilistic technologies. Answers derived from LLM's are phrased convincingly, but as we now know they can create false narratives or 'hallucinations'. Worse, they may even contain biased or offensive language to certain populations. While AI, especially LLMs, significantly enhances knowledge curation, kama's concept of *GenAI's Sober Second Mind*® ensures the necessary human oversight and validation, protecting against hallucinations and inaccuracies.

"We were extremely pleased with how fast we were able to launch our virtual agent and how easy it was to train for new inquiries." **Robert Killin, CEO SmartDesk CRM**

KM Goals, Curation and Delivery

Each organization has its own goals for setting up a knowledge management framework. This paper does not develop a business case nor a return on investment for implementing a new knowledge curation and dissemination system. Instead, it investigates the practical approaches enterprises can take to curate and deliver knowledge effectively through modern AI-driven systems. It explores how leveraging tools like knowledge graphs, and Virtual Agents can enhance knowledge access while mitigating risks. These risks associated with large language models include flawed, biased, or inaccurate responses. By focusing on a two-stage process—using LLMs and Retrieval-Augmented Generation (RAG) for knowledge curation and Knowledge Graph databases for secure, scalable knowledge storage—organizations can ensure accurate and timely knowledge delivery. This blend of technologies offers enterprises a reliable and efficient way to share their core knowledge such that it is accessible, and aligned with business objectives, governance, and compliance needs.

Knowledge Management Goals

Sharing a single source of truth across an organization and ensuring everyone has access to it brings significant benefits with it. These benefits include economic efficiency, tighter compliance, and improved governance. Alternatively, they might even provide a significant societal impact, depending on the organization's goals. In setting goals, it is important to distinguish between gathering knowledge and effectively sharing it to achieve the organization's objectives.

Efficient Curation

58%

USING OPENAI'S CHATGPT 4 AND OTHER PUBLIC MODELS, WE SHOW THAT LLMS HALLUCINATE AT LEAST 58% OF THE TIME, STRUGGLE TO PREDICT THEIR OWN HALLUCINATIONS, AND OFTEN UNCRITICALLY ACCEPT USERS INCORRECT ASSUMPTIONS. [2]

OXFORD UNIVERSITY

Sound knowledge curation starts with an organization identifying the knowledge needed to achieve its goals. Next, it must organize that data into a searchable format. This collection is a repository holding operational knowledge. But such a repository is not always actionable nor is it always easy to access.

To work efficiently, we need specific information related to our current tasks. This is where LLMs and RAG are powerful tools. They should be part of the knowledge curation process. But, generative AI carries risk in the accuracy of its responses, as highlighted in the cited Oxford study.[2] As mentioned earlier LLM solutions can introduce false narratives into the mix. It is risky if AI-created answers are given directly to users who need accurate and actionable data. Applying guardrails reduces this risk. Copyright detection is an example. Despite a lower risk, it isn't zero. Zero risk and 100% alignment with company values, vocabulary, and tone is virtually impossible with probabilistic response generation.

Given the above, the best way to share knowledge with zero risk is to take the previously stated two-stage approach. Use LLM and RAG technologies for curation. Next, store verified information in a scalable, natural language knowledge base. Then use graph databases with organized and linked knowledge - which resembles how human minds work. This is where *GenAI's Sober Second Mind*® can play a crucial role. It ensures the shared knowledge is fact-checked and governed in advance. It gives you with a higher level of trust in the AI's output.

Graph databases, like regular relational databases, are stable and reliable. They store accurate information. Although they have been around for decades, they are not designed for conversational AI, and that's the challenge. We need the accuracy, scalability, and performance of a graph database, combined with the ease of conversational access for reliable AI.

Knowledge Delivery

50%

THE AMOUNT OF TIME THAT KNOWLEDGE WORKERS WASTE IN HIDDEN DATA FACTORIES, HUNTING FOR DATA, FINDING AND CORRECTING ERRORS, AND SEARCHING FOR CONFIRMATORY SOURCES FOR DATA THEY DON'T TRUST. [4]

HARVARD BUSINESS REVIEW

Once we have enough validated and useful information for users, we can focus on delivering it efficiently. As texting has become one of the most popular communication methods, live chat has surpassed voice for enterprise-client interactions. With chat-based generative AI tools, it's clear that natural language 'chat' is the main medium to target. By adding text-to-speech technology, voice can also be used for phone and smart speaker applications. This creates a multi-channel communication system for your organization.

What we need is a natural language, text-based interface combined with a flexible, reliable knowledge base. Then we need to ensure easy access to this information for either employee or customer facing access. Herein lies the challenge. Chat via LLMs as the main tool for natural language access brings the risk of inaccurate or off-brand responses. Simply put, using an LLM to query a knowledge base can lead to unreliable answers. It forces users to fact-check, or continually question the answers provided. This defeats the purpose and ideal of improving efficiency, as 50% of the knowledge worker's time is wasted searching and verifying for data accuracy, according to an HBR analysis.[4]

"Responsible AI isn't just about technology—it's about trust. Every interaction with a Virtual Agent needs to be designed to provide accurate, human-like responses that align with your organization's values. That builds credibility with customers, improves employees productivity, and safeguards your brand in market." **Brian Ritchie, CEO kama.ai**

Challenge Accepted

61%

OF RESPONDENTS EXPECT CHATBOTS TO ANSWER QUESTIONS AND BE BOTH CAPABLE OF PROVIDING DETAILS AS WELL AS EXPLANATIONS THAT ARE ACCURATE AND DEPENDABLE. [5]

SALESFORCE

The challenge is delivering accurate knowledge to the right people and systems. At kama.ai, we've embraced the goal of "***Getting the right information to the right people at the right time, for the right reasons.***" It is about creating a platform that's easy to use for everyday users, business teams, and communities. It is also about providing a flexible enough tool to work with existing enterprise systems like SharePoint or OneDrive. The goal is to connect curated information and systems with those who need them, to work accurately and productively. Finally, the ideal system would offer a human-like understanding of inquiries and providing personalized responses beyond what simple search or modern chat solutions offer, today.

Kama.ai's solution to this challenge is called the *Designed Experiential Intelligence*® platform, or kama DEI. The kama DEI system provides 5 key components:

1. Proprietary patented knowledge graph technology (graph database) organizes the enterprise information. It supports contextual machine understanding based on human values or other relevant prioritizations and human-like understanding criteria.

2. An administrative system that is operated by Knowledge Managers who enter, link, and organize enterprise information into Knowledge Records. This is the basis of the knowledge graph, along with the relevant Extended Data payloads. Extended Data can be short introductory answers, links to larger related manuals, practices or procedures, instructional videos, or other content that is helpful in knowledge transfer to keep people informed and productive.

3. A Natural Understanding Layer (NLU) that interprets the user's input and transforms it into natural language machine understandable graph-like components called 'triples' (eg. I,am,hungry;and;I,am at,home)

4. An Emotional Intelligence (EI) Controller that understands user inquiries through the above triples in human-like ways. It ties to human values like 'comfort', 'health', 'shelter', 'convenience' and other such options that contextualize almost any human or organizational situation.

5. A convenient chat interface (chatbot framework) that can be rendered on any digital device. This includes laptops, tablets, and smart phones. It can also extend through a common API to customized mobile applications, intelligent

Delivering Virtual Agents

Knowledge management isn't new, yet making it easy to access and search has always been a challenge. The internet and mobile devices have made information widely available. Yet, web search still hasn't fully solved the problem. Now, with AI-powered chat interfaces connecting to enterprise knowledge systems, the idea of a conversational virtual coach or 24/7 automated help desk is becoming a reality. In the fields of knowledge management and AI, this coach is known as a 'Virtual Agent'. In effect, it is like a virtual employee, or mentor that helps users find information through text or speech - anywhere and in almost any language.

With integrated Robotic Process Automation (RPA), the same Virtual Agent can handle complex tasks like providing a life insurance quote based on user input. Another example is helping update a home address in the HR system—simply by answering a few questions in a chat.

The kama.ai Virtual Agent platform uses human knowledge to ensure every piece of information delivered, is accurate. In this respect, every process it handles can be trusted by the consumer or employee within your enterprise. No more wasted time and effort.

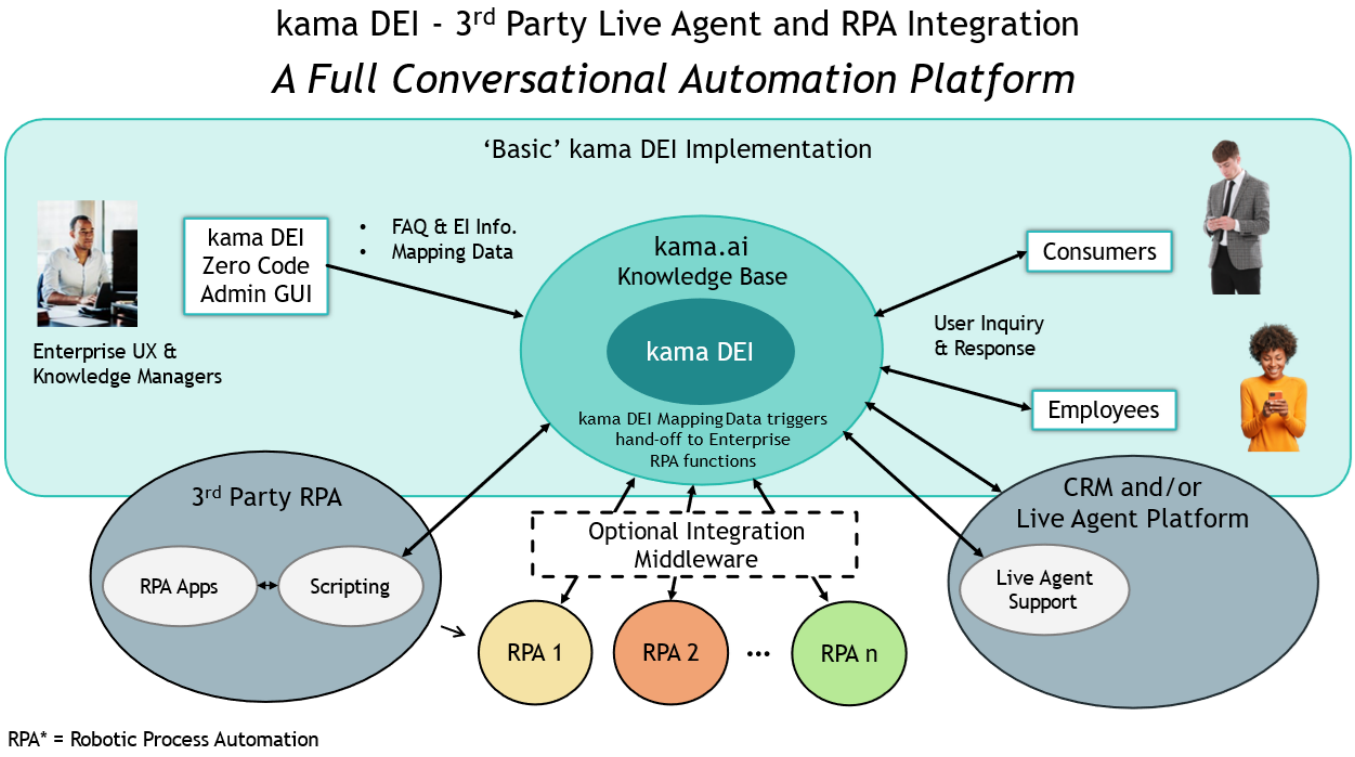

Figure 1

22%

THE GLOBAL ROBOTIC PROCESS AUTOMATION (RPA) MARKET SIZE WAS VALUED AT $4.14 BILLION USD IN 2022 AND IS EXPECTED TO GROW AT A CAGR OF 21.85% TO $13.6 BILLION USD BY 2028. [3]

360 INDUSTRY INSIGHTS

Figure 1 shows how kama DEI distributes knowledge both within and outside an enterprise. The key functions are:

1. Knowledge Managers and/or Subject Matter Experts use the kama DEI admin system to enter knowledge into it. This done as questions and answers that are represented as different Knowledge Records.
2. Knowledge Managers then configure the question and answer records by adding Emotional Intelligence information (i.e. human value ratings), and Extended Data content. Extended data can be short text answers, links to related documents, or a link to a related explainer video. Each link can have further descriptions such as '...see page 10 in the operations manual...'.
3. The Knowledge Manager adds synonym records to the graph to improve the Virtual Agent's ability to understand and find information. Other records can also be added to link this knowledge to related topics, helping users find it under different categories. Graph databases are perfect for organizing and connecting this type of information.
4. The Knowledge Manager adds the new Knowledge Records to one or more "Personas" such as 'Marketing', 'Sales Team' or 'Client Services'.
5. Staff/Users conveniently use a standard web chatbot, or a custom mobile application to ask questions in natural language text. This can be done on any smart phone, tablet, computer, or even through speech interfaces. The system then retrieves the relevant needed information.
6. If the users are not able to find the information they seek, they can provide 'thumbs down' feedback and add specific details of what other information they wanted.

7. Whether or not the user provides specific feedback, the Knowledge Manager can review both the list of unanswered questions and specific user feedback to:

    a. Insert additional synonyms. This addresses the different ways that the user inquired about information already existing on the knowledge graph.

    b. Work with Subject Matter Experts to find new information for the Virtual Agent, if the information was not in the knowledge graph. In addition, they can use 'Retrieval Augmented Generation' (RAG), or kama DEI "Draft Assist", to create a draft answer. This from existing enterprise documents, information on the internet, or a specific website. The draft answer is then reviewed and approved by the right enterprise personnel, and is stored in the knowledge graph. The new information is added to one or more of kama DEI personas to expand the capability of the Virtual Agent(s). These steps are performed in a zero code natural language model, with virtually no technical skills required. The work for a single Frequently Asked Question (FAQ) can be accomplished in about 5 to 10 minutes.

"...despite their best efforts, building AI responsibly isn't enough to guarantee trust. That's up to stakeholders. It's their perceptions and experiences that determine whether trust is earned, and whether AI initiatives are ultimately successful." [6] **PwC 2024 Responsible AI Survey**

Governed in Advance

93%

OF BUSINESS LEADERS BELIEVE HUMANS SHOULD BE INVOLVED IN ARTIFICIAL INTELLIGENCE DECISION-MAKING. [8]

WORKDAY

As noted, LLMs and RAG have improved efficiency in researching, summarizing, and organizing information for enterprises. But, their ability to generate false, biased, or even offensive comments means they alone cannot be fully trusted as knowledge management systems. True knowledge must be factual and reliable. Technologies prone to even small errors simply cannot serve as dependable sources.

Despite this, LLMs and RAG are valuable tools for research and drafting responses as long as the content is validated by experts. This is where kama's *GenAI Sober Second Mind*® can ensure that all AI-generated content is checked before being shared. This maintains the integrity and trustworthiness of the enterprise knowledge. The Knowledge Manager (KM) does not have to second guess every interaction or response. They only need to ensure the right data, information, and documents are approved for a particular knowledge database. This way, only fact-checked, truthful information, and references to approved documents are included as the knowledge set from which to develop user answers.

Kama DEI;s information is completely trustworthy because it uses its own full-stack natural language conversational AI which does not rely on LLMs or RAG for real-time inquiries. Requests without pre-validated knowledge triggers a human-in-the-loop process. Here, critical responses are provided within minutes to a few hours. Enterprise knowledge managers can also provide direct answers to close the loop with users that did not get a real-time response, earlier.

Typical Knowledge Delivery shown in figures 2 (next page) shows the general operation for question asking (prompting) and answering/knowledge delivery, is done in steps 1-3:

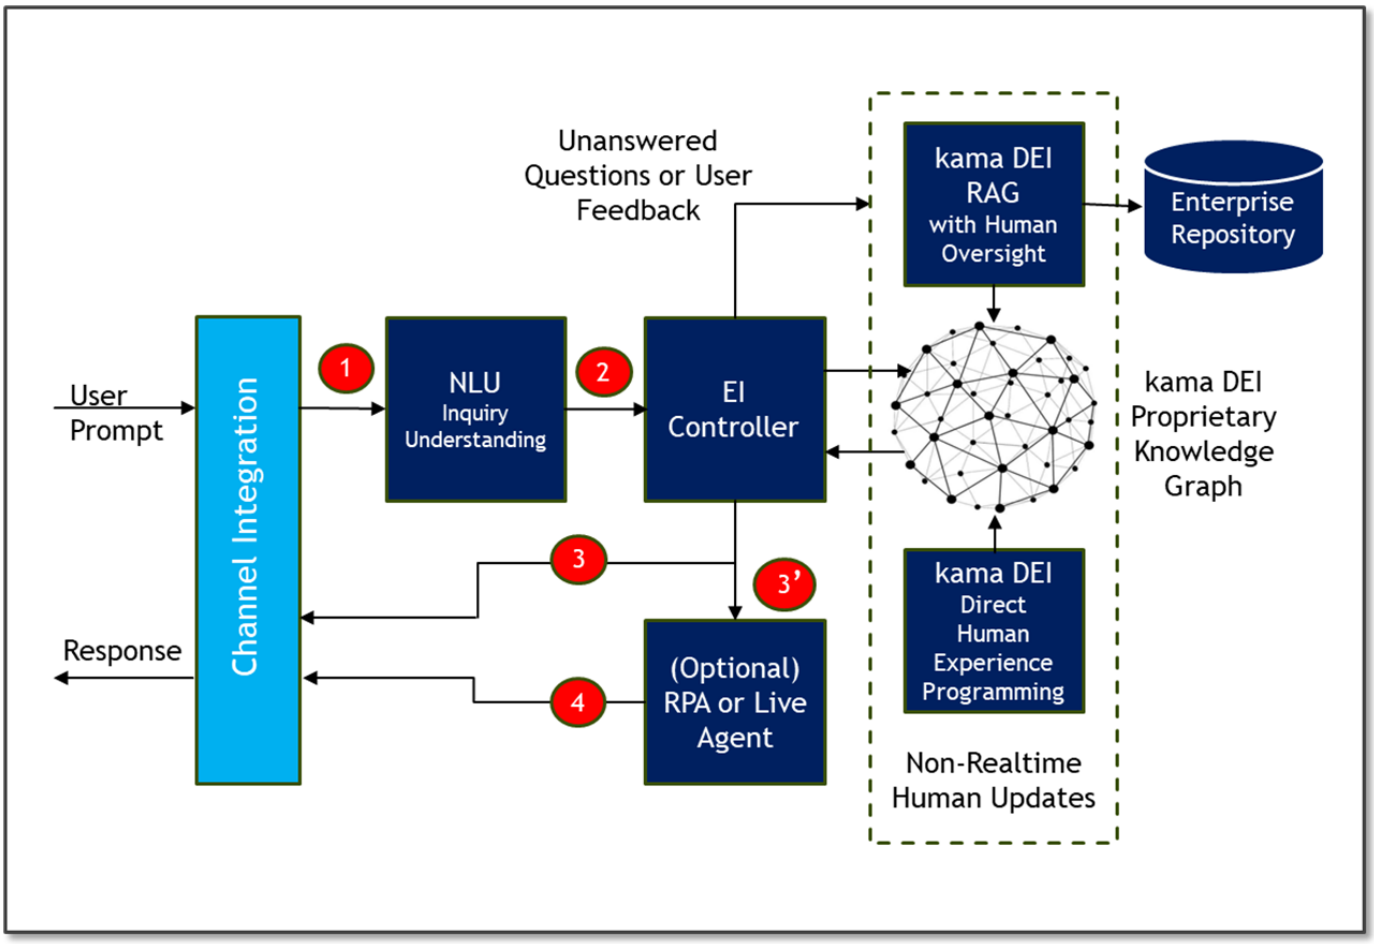

Figure 2: Process in practice with user inquiries to the kama DEI platform.

1. A question is typed into the Virtual Agent interface (i.e. chatbot) in a natural language inquiry. The question is 'understood' grammatically using the Natural Language Understanding (NLU) layer.

2. The 'problem' is located in the knowledge graph.

3. The pre-validated and approved response or responses are delivered in a human value prioritized order to the user in convenient natural language responses or descriptive natural language option 'buttons' that are selected by the user for the most relevant information they were seeking.

Live Agent Handoff - In some cases, the system detects that an inquiry is serious enough to need direct human support. In these cases the request is handed off to a live agent in real-time. It is transferred within the same chat. Once the live agent confirms the user is satisfied, that user flows back to the Virtual Agent for any further questions, as needed.

Handoff to RPA - In addition, kama DEI supports hand off to 3rd party Robotic Process Automation (RPA) solutions like Blue Prism, AWS Lex, or Automation Anywhere. In this case, the user continues to engage with the chat but the conversation is temporarily handed off to a third party bot. For this example, the handoff gathers and validates information needed for the RPA function. It then performs the desired outcome. It might be gathering information for a process like an address change request or a service quotation. After the process is complete the conversation goes back to kama DEI which asks "Is there anything else I can help you with?" The above two handoff processes are represented by 1 to 3' and 4 in figure 2.

Expanding Knowledge and Process Capability

The dotted section in figure 2 shows cases where the system responds with "I'm sorry, I don't have an answer for that" or the user states that their question was not addressed. In kama DEI, these options trigger the creation of new content or a new process. This new information is not generated in real-time; it is developed by Knowledge Managers and Subject Matter Experts who may choose to use LLMs and RAG. This to ensure accuracy and trustworthiness. Once fact-checked and approved, the information is added to the knowledge graph. Human oversight like this ensures higher compliance, lower risk, and reduced overall costs for the enterprise.

Be a Knowledge Champion

Share this Whitepaper with your peers and let's build smarter, more efficient organizations together.

Takeaways

58%

ONLY 58% OF RESPONDENTS HAVE COMPLETED A PRELIMINARY ASSESSMENT OF AI RISKS IN THEIR ORGANIZATION. [7]

PRICEWATERHOUSE COOPERS

Generative AI has brought significant progress to enterprise knowledge management. But challenges remain, particularly with the accuracy of LLM based output. An April 2024 Gartner article, "When Not to Use Generative AI," highlights these limitations, noting the potential for errors, bias, and unreliable information.[10] To address these issues, organizations should complement GenAI with dependable tools like those based on Knowledge Graphs, and rules-based AI to ensure accurate, fact-based knowledge sharing.

To achieving effective and timely dissemination of enterprise knowledge organizations must:

- **Leverage Virtual Agents for Knowledge Mgmt:** Virtual Agents powered by conversational AI are a new and effective approach to making knowledge accessible. This technology is useful for both internal use (employee inquiries) and external use (customer service). A scalable option for knowledge delivery.

- **Use Generative AI for Curation, Not Dissemination:** Generative AI is excellent at knowledge curation tasks like search, summarization, and drafting. However, it should not be directly used to distribute actionable enterprise information. This is where *GenAI's Sober Second Mind*® comes into play. It ensures human oversight validates and governs the content - curtailing hallucination risks, for example.

- **Deploy Knowledge Graphs for Scalability and Accuracy:** Knowledge graphs provide a scalable, flexible, and highly reliable system for organizing and link-

ing enterprise knowledge. They provide accessible and accurate information with a stable foundation for enterprise knowledge management.

- **Adopt No-Code for Easy Virtual Agent Creation:** No-code platforms let business users create and manage Virtual Agents without technical expertise. It empowers administrative staff to make enterprise knowledge accessible through conversational AI. This without needing specialized developers or contractors.

- **Expand with RPA and Live Agent Handoff:** Integrating live agent handoff and Robotic Process Automation (RPA) enhances the Virtual Agent's capabilities, enabling it to handle more complex tasks. These tools let the knowledge system not only answer questions but also complete tasks like processing HR updates or delivering quotes, turning the platform into a full enterprise automation solution.

Additional Insights

80%

OF RESPONDENTS AGREED THAT AI AND ML HELPS EMPLOYEES WORK MORE EFFICIENTLY AND MAKE BETTER DECISIONS. THIS IS ESPECIALLY TRUE WHEN EQUIPPED WITH ACCURATE AND RELIABLE KNOWLEDGE MANAGEMENT TECHNOLOGY THEY CAN ACCESS EASILY. [9]

WORKDAY

WITH RESPONSIBLE AI VIRTUAL AGENTS, YOU CAN DRAMATICALLY IMPROVE YOUR EMPLOYEE PRODUCTIVITY, AND CUSTOMER ENGAGEMENT. PROVIDE ACCURATE ANSWERS, MINIMIZE RISK, AND RELIABLY PROTECT YOUR BRAND.

- **Graph Databases for Stability and Flexibility:** As emphasized throughout the report, graph databases offer a stable, reliable, and scalable way to organize enterprise knowledge. These databases function similar to how the human brain links information. It makes them ideal solutions for managing large, complex bodies of knowledge.

- **Human Oversight for Accuracy:** To ensure trust and reduce risks, kama.ai integrates a human-in-the-loop process. This approach involves human experts to validate AI-generated content that is critical, before it is shared. Doing so ensures the information delivered is accurate, on brand, and aligned with enterprise standards.

- **Personalization through Human-Like Understanding:** The kama.ai platform provides a high level of personalization. It aligns responses with human values and specific user needs. This makes the system more effective at delivering relevant, accurate, and personalized information to users, whether they are employees or customers.

In summary, organizations can safely and reliably harness the power of Generative AI by combining it with deterministic knowledge graph technology, and human oversight. This approach provides an effective, governed-in-advance, ultra-low-risk knowledge management solution that is easily accessible by customers and employees through trustworthy Virtual Agents. This hybrid approach enables scalable, efficient, and accurate delivery of core enterprise knowledge for maximum productivity, reduced risk, and optimal return on investment.

Let's Talk

Think you need a Virtual Agent, AI, or knowledge management solution?

Let's connect to review your options.

Ready to elevate your customer interactions with a powerful, human-centered AI? At kama.ai, we're revolutionizing the way businesses connect with their customers through our *Designed Experiential Intelligence*® (kama DEI) platform. Our technology goes beyond typical AI—it's emotionally intelligent, personalized, deterministic (does not hallucinate) and is continually evolving with each interaction. Whether it's automating routine processes, improving customer engagement, or delivering 24/7 support, kama DEI brings the human touch to your AI and Virtual Agent automation.

How can better knowledge management transform your business? Let's have a pressure free conversation. Kama.ai helps organizations like yours drive sales, improve customer loyalty, and improve operational efficiency with the AI that's as human as it gets.

Think kama.ai for trust, empathy, and accuracy.

E: INQUIRIES@KAMA.AI
M: +1 (416) 275 - 1780
W: KAMA.AI

Responsible AI - PDF

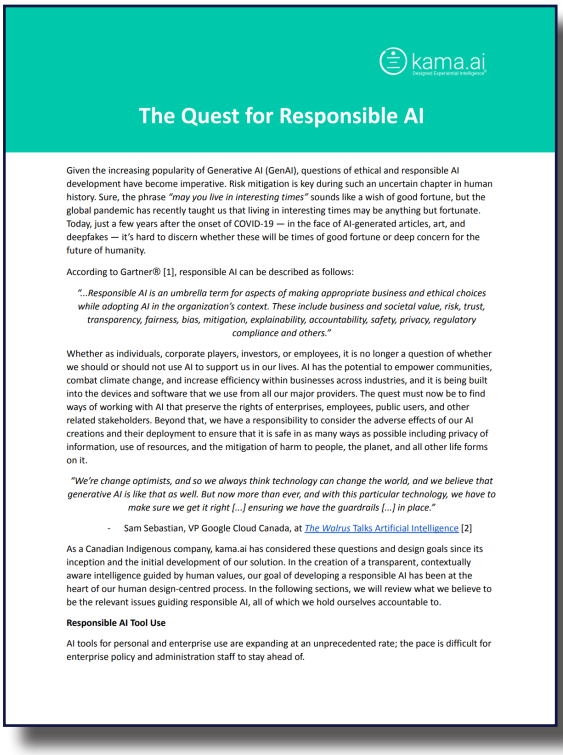

kama.ai

The Quest for Responsible AI

Given the increasing popularity of Generative AI (GenAI), questions of ethical and responsible AI development have become imperative. Risk mitigation is key during such an uncertain chapter in human history. Sure, the phrase *"may you live in interesting times"* sounds like a wish of good fortune, but the global pandemic has recently taught us that living in interesting times may be anything but fortunate. Today, just a few years after the onset of COVID-19 — in the face of AI-generated articles, art, and deepfakes — it's hard to discern whether these will be times of good fortune or deep concern for the future of humanity.

According to Gartner® [1], responsible AI can be described as follows:

"...Responsible AI is an umbrella term for aspects of making appropriate business and ethical choices while adopting AI in the organization's context. These include business and societal value, risk, trust, transparency, fairness, bias, mitigation, explainability, accountability, safety, privacy, regulatory compliance and others."

Whether as individuals, corporate players, investors, or employees, it is no longer a question of whether we should or should not use AI to support us in our lives. AI has the potential to empower communities, combat climate change, and increase efficiency within businesses across industries, and it is being built into the devices and software that we use from all our major providers. The quest must now be to find ways of working with AI that preserve the rights of enterprises, employees, public users, and other related stakeholders. Beyond that, we have a responsibility to consider the adverse effects of our AI creations and their deployment to ensure that it is safe in as many ways as possible including privacy of information, use of resources, and the mitigation of harm to people, the planet, and all other life forms on it.

"We're change optimists, and so we always think technology can change the world, and we believe that generative AI is like that as well. But now more than ever, and with this particular technology, we have to make sure we get it right [...] ensuring we have the guardrails [...] in place."

- Sam Sebastian, VP Google Cloud Canada, at *The Walrus* Talks Artificial Intelligence [2]

As a Canadian Indigenous company, kama.ai has considered these questions and design goals since its inception and the initial development of our solution. In the creation of a transparent, contextually aware intelligence guided by human values, our goal of developing a responsible AI has been at the heart of our human design-centred process. In the following sections, we will review what we believe to be the relevant issues guiding responsible AI, all of which we hold ourselves accountable to.

Responsible AI Tool Use

AI tools for personal and enterprise use are expanding at an unprecedented rate; the pace is difficult for enterprise policy and administration staff to stay ahead of.

A serious look at Responsible AI and why it is critically important to your Enterprise business and brand.

Dive into the world of Responsible AI with this comprehensive white paper, where ethics meets innovation. In today's fast-evolving AI landscape, responsible AI is no longer optional—it's a must. This paper highlights key areas like transparency, data privacy, and explainability, providing a roadmap for companies to implement AI ethically. Learn how generative AI, when combined with human oversight, can empower enterprises while also minimizing the risks of misinformation, hallucinations, and bias.

Pave the way for a future where AI serves both people and the planet responsibly.

Find it here:
https://bit.ly/3Bj9lyo

FREE to download
NO FORMS.
NO STRINGS ATTACHED !

Endnotes

1. Fact.MR, “Knowledge Management Market Study by Systems, Mechanisms & Technologies, and Processes, Infrastructure for Small & Medium Enterprises and Large Enterprises from 2024 to 2034”, 2024, https://bit.ly/47WbIDD

2. Oxford Academic, “Large Legal Fictions: Profiling Legal Hallucinations in Large Language Models”, Journal of Legal Analysis, Vol 16. Issue 1. Matthew Dahl, Varun Magesh, Mirac Suzgun, Daniel E. Ho, Ferris Ayar, Jun 2024, https://bit.ly/4gBcd9T

3. 360 Industry Insights, “Global Robotic Process Automation (RPA) Market Size, Share and Growth Report 2030”, Feb 2024, https://bit.ly/4gGFbFh

4. Harvard Business Review, “Bad Data Costs the U.S. $3 Trillion Per Year” Thomas C. Redman, Sep 2016, https://bit.ly/3TJMKkN

5. Salesforce & Drift, “Key Chatbot Statistics to Know in 2019” Mathew Sweezey, Aug 2019, https://sforce.co/3XCvD5K

6. PwC, “PwC’s 2024 US Responsible AI Survey”, 2024, https://pwc.to/4dqrXcR

7. PwC, “PwC’s 2024 US Responsible AI Survey”, 2024, https://pwc.to/4dqrXcR

8. Workday, “Workday Global Survey: Majority of Business Leaders Believe Humans Should be Involved in AI Decision-Making; Cite Ethical and Data Concerns”, Jun 2023, https://bit.ly/3N0b4uS

9. Workday, “Workday Global Survey: Majority of Business Leaders Believe Humans Should be Involved in AI Decision-Making; Cite Ethical and Data Concerns”, Jun 2023, https://bit.ly/3N0b4uS

10. Gartner, “When Not to Use Generative AI”, Leinar Ramos, Ben Yan, Haritha Khandabattu, and Gabriele Rigon, Mar 2024, ID G00806238

Photography Thanks To:

1. Unsplash: Stefan Stefancik
2. Unsplash: Evangeline Shaw
3. Unsplash: LinkedIn Sales Solutions
4. Adobe Firefly
5. Unsplash: Christine Woc
6. Unsplash: Mimi Thian

kama.ai

E: INQUIRIES@KAMA.AI
M: +1 (416) 275 - 1780
W: HTTPS://KAMA.AI

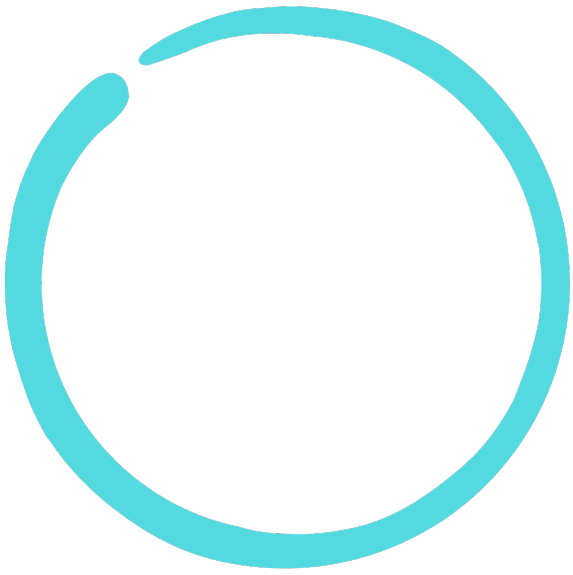

Visit us on **LinkedIn**