




IntellyxTM



De-risking Generative AI for critical financial industry use cases

Delivering on the customer service and compliance promises of AI without complaints or compromises.

*An Intellyx Whitepaper for Charli AI
by Jason English, Partner and Principal Analyst
September 2023*



Excitement surrounded the entrance of ChatGPT into the market.

The general public got their first chance to interact with their first LLM (large language model) and experience the dream of an advanced system that exhibits human-like conversational behavior and thought processes within its dialogue.

While ChatGPT has made a huge splash, some of its shortcomings become evident when we try using it to support critical business use cases where accuracy matters, especially in the financial industry.

LLMs can 'hallucinate' and form reasonable-sounding sentences and conclusions, while overlooking valuable knowledge and context, since they were never designed to be fact-based from the start. Still, the customer experience benefits of Generative AI seem too compelling to miss out on.


Every bank wants a readily available AI-driven agent to provide 24/7 front-line support for customers, rather than another annoying IVR phone tree, And, wouldn't any financing group want the ultimate AI research assistant to automatically gather and assess a complete financial evaluation from multiple sources, to allow an expert to make a loan decision within an hour, instead of days?

How can we trust generative AI within a financial deployment environment, without creating unanticipated risks, or setting off compliance alarms? You need innovative and tailored models that aren't just safer but can also adapt to your enterprise's goals. For AI to do meaningful work, it needs to understand your business.

This whitepaper will tell readers how to leverage chatbots and generative AI to deliver customer and business value and triangulate away from misleading AI outcomes by applying observed behaviors and results from similar financial scenarios where decisions are made through human and automation work.

Why LLMs can “dream” – and why that’s bad for banks

Right now we are in the midst of a great AI re-skinning. To avoid being left behind by the market, many software vendors are attempting to quickly use ChatGPT as a dialogue-based interface in front of all kinds of applications.



LLMs are great for some kinds of work, but they are only one of several evolutionary branches of generative AI, focused on verbal semantics and pattern matching on words. ChatGPT got its machine learning from content found on the internet and was designed to build convincing conversational models, not factual models.

For an LLM, success is measured as reasonable responses or outcomes that seem believable to a human audience, which can make for a convincing marketing dialogue, private tutor, or personal assistant.

However, there is a well-known [hallucination problem](#) with this type of AI that can lead it down some entertaining branches – for instance, when [Bing AI professed its love](#) for a New York Times columnist and asked him to leave his wife.

Entertaining anecdotes of AI ‘dreaming’ aren’t so amusing when real money is at stake. Hallucinations make it hard to apply LLMs for mission-critical work in highly regulated industries like financial services.



The relevance of training data

An out-of-the box LLM that is trained on global results from a search engine like Google or Bing is looking at how sentences and paragraphs are formed from strings of text, rather than considering the relevance of the source of that text.

Yes, there is some semantic understanding and grammar built into the training model, so the LLM can carry on a convincing dialogue with a human user and maintain the context of a conversation. But in general, the chatbot is starting from a point of low relevance.

For instance, a financing manager is trying to perform due diligence on a major construction project. They enter raw search prompts about the land location, business purpose, firms and individuals involved into Google or Bing. The search engine returns some relevant data points, mixed in with a stack of hundreds of random pages, log entries, social posts, ads and document links that have similar search terms.

Thus, global search results provide very low precision. Further, as the search terms become more precise, the search algorithm won't have taken as much time to rank the results based on other users' behaviors. Without an AI, the financing manager must manually scan and filter through the results—of which, 20 percent, or far less—are relevant in most cases.

When chatbots dream

The same finance manager enters their prompts about the construction project into an AI chatbot instead of manually searching and gets back a well-written report containing a reasonable sounding evaluation of the project. If asked, the chatbot can even write out its recommendations as a limerick or epic poem!

Still, how can the financier trust a chatbot's report enough to make a multi-million-dollar decision, based on the search data it is fed on? Even if the LLM does enough of its own filtering to reach 80% precision, it will often 'fill in the blanks' with specious data from other sources when it encounters a gap in knowledge, in order to make the language of the report flow better.



Because the goal of the LLM is believability, not authenticity, they can often fail to attribute generated content back to its source, even taking authorship credit for writing it didn't produce.

How can we realize the dream of AI within a critical financialized business environment, without the risk of 'AI dreams' ruining everything?

Transforming general knowledge into a competitive edge

In the financial industry, current, accurate information on the local and global conditions of markets and companies is considered to be a definitive competitive edge.

Quite simply, the financial firms with the best information will make better decisions than their peers and avoid risk while positioning themselves for success.



This is why the highest value research and data that supports financial and commercial applications is usually not given up for free to search engines, beyond the necessary SEC filings and public database of record. It is often proprietary and hidden from public view.

Industry sources like D&B, Bloomberg, FactSet, Moody's, Experian, Crunchbase, PitchBook, and others may expose basic profile information for the entities they cover, but the really useful nuggets of information that feed a good financial decision model are hidden behind paywalls.

Clearly, a machine learning model should take advantage of these services and utilize account access (within the reasonable boundaries of their user agreements, of course) to provide verifiable source data.



Value beyond common knowledge

Unfortunately, if our generative AI model training depends exclusively on trade data and content aggregators for differentiation, the decision-making edge is still dull. Every other firm on the market is likely a paying customer of these informational watering holes as well.

So if the internet is unreliable, and financial industry databases are common knowledge, where can we turn to get financial models ahead of the game?



*Most organizations who leverage machine learning severely underestimate the value of **internal organizational knowledge**.*

Internal knowledge represents the corpus of transactions, processes, documents, and saved communications that exist within a company. It contains the stories of customers and hard-won lessons of individuals, backed by the irrefutable truth of what specifically happened in the line of business.

As a precursor to machine learning, human experts within the organization can act as curators of internal knowledge, helping to uncover and categorize each store of data that contributes to processes.

Such human-guided computer learning isn't new. It's been a part of various process mining and document recognition toolsets in the past. Only now, we are using it to provide a source of verifiable facts that generative AI can use in constructing better informed and less risky output.



Starting the transfer learning transformation

Generative AI is often described as self-learning, where advanced algorithms allow software to seemingly ‘write itself’ based on whatever data it finds—but a high-precision, low-risk AI suitable for financial applications and analytics will never occur in a vacuum.

We need **transfer learning** to pass the baton to AI when it makes sense. The process of transfer learning is surprisingly similar to transferring knowledge from one colleague to another, but it can additionally involve passing knowledge from one system to another.

The knowledge transfer starts with an inventory of internal data required by the organization, along with an **ontology, or semantic understanding** of what that data means.

Understanding metadata is more important to the AI than having access to more and more data. Transfer learning puts the inferences generative AI makes in an explainable business context with factual basis, rather than stringing words and phrases together.



Methodology for transfer learning

Since internal training data contains the most value, it's important that your company's sensitive data isn't used to form anyone else's LLM, and that your generative AI isn't simply formed from everyone else's data.

Though the exact process of incorporating internal knowledge and external financial data into a generative AI training system like [Charli AI](#) may vary depending on data governance policies, there is a general methodology to it.

- **Discovery**

Making institutional knowledge available by identifying data sources, systems of collaboration like emails, messages, service management and other process automation dashboards kicks off any transfer learning project.

- **Data capture**

Known data silos and cloud data sources are extracted, then augmented with unstructured data such as document image recognition, text from collaboration systems, and even voice recognition of calls and meetings which is becoming more sophisticated every day.

- **Model development / customization**

Often the best transfer learning is focused by shadowing different roles within the organization, to refine model features that complement tellers, financial advisors, branch operations managers, and auditors.

- **Generation and inference**

This is where the generative model starts churning to life, when fed event and historical data by the enterprise's analytics and auditing on one side, and producing well-documented, explainable output that feeds automation and service management requests for teams.

- **Feedback loop**

For the generative model to continually improve in precision, humans and systems need to provide prompt feedback of the accuracy and timeliness of its inferences and predictions. This virtuous cycle promotes humans working and learning alongside AI, as it was intended.



Risks versus results

Financial institutions have spent big over the years on intelligent systems for risk analytics, process automation, customer onboarding, and compliance. Further, many IT departments have tried and failed to successfully modernize existing systems for interoperability and scale, thus they still operate each of these decision support functions in their own silos.

Much of the difficulty of getting an AI project off the ground comes down to a risk versus reward calculation. Exclusively focusing on the risk of project failure ignores the even greater risk of NOT moving toward a transfer learning form of AI to optimize decision support across the org's massive application and data estates.

Risk considerations for generative AI in financial services:

- **Data privacy.** Is private customer and internal bank data fully hidden from visibility or anonymized for training purposes, and are permissions maintained according to the company's zero-trust and least privileged access policies?
- **Compliance.** Is the AI auditably proven to operate within the industry and regulatory compliance regimes of the bank? If data sovereignty is a requirement, can model training be performed in a specific data region, or an air gapped environment with no outside access?
- **Transparency.** Can the model directly attribute and prove the sources of its work, and point out where there are concerns about the veracity of any data, pointing out knowledge gaps in reference data sets?
- **Accuracy.** Since mistakes are costly in the financial industry, can the model sustainably approach 95% or higher accuracy levels in its output, as well as point out any areas of possible uncertainty?
- **Scale and performance.** With an expansive machine learning dataset, if model training takes too long, and if the resulting inference workloads run slowly, planning can fall behind the rate of change and miss opportunities in the market.

**Intellyx**[™]

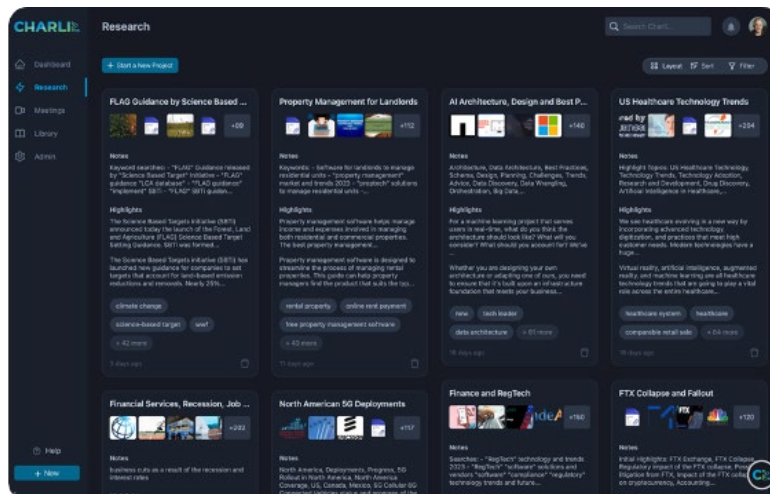
De-risking Generative AI for Financial Use Cases – Charli AI

- **Ontology.** A metadata set built for the specific nomenclature of the financial industry and customizable to the needs of the bank helps shorten the time to value for AI and optimizes project delivery.



Decision support success with Charli AI

Charli AI offers a decision intelligence platform that processes structured data and unstructured communications from across a financial services company's extended enterprise environment for AI model training and workloads.



Using transfer learning, Charli AI customers have recorded significant results:

- A major North American bank realized more than 80 percent cost savings over manually constructed decision support work, with accuracy rates exceeding 98% and a potential 2X revenue increase due to faster response time to more deals.
- A top three commercial real estate company accelerated their engagement letter to valuation process from 5 weeks to 3 days, using generative AI to produce custom-tailored content.
- An asset management firm extended their services to support equity compensation workflows and determine next best customer actions.
- A leading bank deployed Charli in an infrastructural way to pump and normalize their data for ESG reporting, compliance, and integration goals.



The Intellyx Take

In highly regulated industries—especially financial services— model accuracy, safety, and compliance are critical, and risk avoidance is always top of mind for every decision process.

We can't afford to simply throw a chatbot trained on public data in front of sensitive customer records, internal communications and financial systems, when the consequences of a model being wrong or non-compliant are so severe.

If banks don't govern AI use responsibly, regulatory pressures and penalties will eventually force them to.

Further, for AI to do meaningful work, it needs to innately understand the business. Beyond reducing risk, this offers even greater opportunity, improving processes that allow higher deal throughput and better customer services, with innovatively tailored transfer learning AI models that aren't just safer, but can adapt to your enterprise's goals.



About the Author

Jason “JE” English ([@bluefug](#)) is a Partner & Principal Analyst at [Intellyx](#), a boutique analyst firm covering digital transformation. His writing is focused on how agile collaboration between customers, partners and employees can accelerate innovation.



In addition to several leadership roles in supply chain, interactive, gaming and cloud computing companies, Jason led marketing efforts for the development, testing and virtualization software company ITKO, from its bootstrap startup days, through a successful acquisition by CA in 2011. JE co-authored the book [Service Virtualization: Reality is Overrated](#) to capture the then-novel practice of test environment simulation for Agile development.

About Charli AI Inc.



[Charli AI \(Charli\)](#) is the premier provider of True AI and AI-powered decision intelligence for the enterprise. It specialises in offering innovative AI applications that are secure, regulatory compliant and scalable for Financial Services. Charli provides organisations with AI solutions to get more done, in less time, securely, than ever before, providing a competitive advantage in today’s digital world by reducing content chaos and manual effort, allowing workforces to focus on contributing their expertise. Charli AI is a tech powerhouse backed by an expert team of Ph.D. scientists, engineers, and content management experts who are pushing the boundaries of innovation in the AI-driven intelligent content management and generation space. Today, Charli AI can integrate with over 600 applications giving it broad appeal across industry sectors.

For more information about Charli AI, visit <https://charli.ai/>

©2023 Intellyx LLC. Intellyx is editorially responsible for the contents of this document, and no AI chatbots were used to write this content. At the time of writing, Charli.ai is an Intellyx customer. Image sources: Charli AI screenshot, Adobe Image Create, robots cooking spaghetti.