




Mapping the Enterprise:

Why Explainability in AI is More than a Requirement, It's a Transformative Benefit.



“Codifying judgment means that a person's decisions can have extraordinary scale. For various reasons we like to know who that person is. After all, without responsibility and its identification, how can someone be accountable for a decision?”



~ **Power and Prediction: The Disruptive Economics of Artificial Intelligence¹**



The Requirement of Explainability

The rise over the last couple of years of AI into every Enterprise and even mid-level organization's strategic focus, leads to questions, once theoretical, becoming now critical. Core among them is the question of "Explainability in AI". Massive Language Generating systems also known as "LLMs", the most famous being "ChatGPT" by OpenAI, present power and problem at the same time (at least with the traditional view). Now any user can converse directly with AI systems in native language. This in some ways, makes AI far easier for the average user to understand. At the same time, the size, scope and complexity of these foundation models (and applications built on top of them), present a great challenge: How can we explain why AI predicted, created or acted in the way it did? This is the often discussed "Black Box" problem. We as humans cannot explain our own intuition, much less the "intuition" or reasoning of massive AI systems.

This presents a great problem for the Enterprise: "How do you balance the benefit of leveraging this new age of AI with the risk of AI making a decision you would not agree with nor be able to explain?"

The answer for many today is "Explainability", defined by McKinsey as "...the capacity to express why an AI system reached a particular decision, recommendation or prediction."²




One approach to solving this problem, which is the one we here at Kognitos take, is to leverage the best of both worlds, in the same way that our brain leverages both logic and intuition to make decisions. Logic gives us the auditable reasoning and "explainability" of our decisions, while pattern recognition based heuristics provide greater efficiency and reduce the cost (in terms of mental labor) of making decisions. Kognitos intentionally builds both of these into its operating system, logging all of the inputs and outputs, the judgements and the exceptions. This provides an auditable record in English, where an Enterprise can see the decisions and actions of both people and AI systems.

The screenshot displays a workflow titled "create the billing reconciliation" with a list of 10 steps. Step 5, "get the first member table", is highlighted. To the right, a table titled "Table" lists member information.

MEMBER ID	MEMBER NAME
R00008563979700	ACOL, JACEY
R00008262543500	ACOSTA, RYAN
R00008387061800	ALING, DOMINADOR
R00008386754900	BAGUSO, EDE
R00008386753100	BAGUSO, WARREN
R00008542181600	BARLIE, HOLLY M
R00008386885100	BROOKS, DWIGHT
R00002500381500	BRUMAGHIM, WAYNE
R00008386830700	CORPUZ, ARNOLD
R00008401944700	DOMOGMA, ARAM
R00008386819000	DOWN, CARO

In many cases this solves the requirement of "Explainability" enabling the Enterprise to use AI in core business processes (instead of relegation to less risky endeavors in the way many LLMs are used today). But what if explainability offers more than a requirement or feature? What if it offers an opportunity to map the work of an Enterprise entirely?





Mapping the Enterprise: Why NLPA Offers a More Natural Form of Explaining and Understanding an Enterprise's Work

When discussing AI, most people (and to be fair we do too...) focus on the ability of AI to automate away manual tasks, yet Enterprises are spending an enormous amount of money trying to understand the information and activities of their own business. There is massive value in capturing this information, as well-honed organizations can derive insights for better decision making, and drive efficiency that provides a lasting competitive advantage. The Toyota production systems is perhaps the most famous example of the long-term competitive moat of "Economies of Learning", and started by understanding the actions of how their people "Work".

In an effort to capture the data, the global process mining market has grown to an estimated \$1.75B of license revenue alone in 2023.³ However, process mining has struggled to gain large scale adoption due to massive associated services costs, and limited ability to capture the day to day tasks of users. New applications of LLMs conceptually include the ability to quickly query information within the knowledge base and information systems of organizations. But all of this fails to capture the core business activity of the Enterprise, and leaves executives forced to pay consultants to manually attempt to map this out with users. But often in these studies only the standard, ideal process is captured, and we all know...business is rarely standard.

And here is where "Explainable AI", logic combined with pattern recognition, all recorded in English, moves from being a requirement to a core strategic benefit by enabling an organization to drive "Economies of Learning". AI Systems like Kognitos, capture and log in English the day to day work of people, every exception, every deviation, in a way that is fully explainable and can be "mined" for insights.

But before we dive deeper, let's take a quick departure on the road to the analogy of Tesla.






THE TESLA OF THE ENTERPRISE

Over the last two decades, another industry diligently worked to “Map” out its own environment: autonomous driving. Players in this industry took different approaches. Fully autonomous vehicles with no drivers, set out on the road to map roads and provide the data needed to train its systems to handle all the variations and “exceptions” that may cause problems. Unfortunately, the success of this approach faced limitations, as there are only so many “driverless” cars that society tolerates until accuracy and safety reach certain, admittedly undefined thresholds. Similar to “Process mining”, which is more akin to a GPS map of the city, scale and usability for the next step is limited.

Tesla on the other hand took a different approach. It sought to map how people drive, the actions they take, the “Exceptions” to the happy path of driving the speed limit with no obstructions etc. This provided immense scale immediately for Tesla to learn, and also did not require humans to change their behaviors. Now, Tesla has mapped far more “driving miles” than its competitors, giving it significant economies of learning advantage relative to others on which it is fine-tuning its “Autopilot”. Kognitos is doing the same, but with business users.





The Kognitos Approach:

Using A Previous Requirement to Map the Enterprise

At Kognitos, we sought to build "Explainability" into our system from the ground up to ensure that anyone in the Enterprise (technical or not), can review and audit the exact actions of both people and AI in their automated processes. Every action, every fact, every input and output is logged in English, in a process run. As business users converse with Kognitos and teach it how to handle exceptions, all of the variants to the standard process are captured as well. Because of this, some of the world's largest companies are using Kognitos to automate processes in their core business systems and mission critical functions.

What we discovered, or perhaps stumbled upon, is the value of this to the Enterprises.

The head of transformation for the financial shared services center at Wipro, a global services firm, recently outlined the value this brings to their organization, something they have never been able to achieve before in discussing a core financial process, (Collection Accounting), he noted:

"Where Kognitos has helped is capturing all of the logics that one of our team members knows. This includes how each client credits the account, the specific domain knowledge about each client...all of these manual learnings that my teams have accumulated over the years in working with these individual accounts, has now been inputted into Kognitos as they work with the system. This has huge value for our company, as previously the company would not have had this knowledge, and wouldn't know how to handle each client. Only the individual SME would have known that, how matching happens, how the credits come into the accounts, etc. All of that experiential learning is now saved in Kognitos and we can evaluate it for the future."

~Rajesh Sankaranarayanan

Kognitos is effectively becoming the "Tesla of the Enterprise". A feature, built originally to satisfy the explainability requirement of AI in the enterprise, is now recording in English, all of the actions as people work together with AI to better perform processes. Every exception, every unique form of domain knowledge previously only stored in people's brains, is now being "learned" by Kognitos in a way that organizations can query. In some of our customers, this is happening millions of times per year for a single process, all without the cost of deploying services heavy process mining. These English logs can then be turned into process flows showing the deviations from the happy path for further investigation and process optimization.

The screenshot displays a software interface for a ProService POC. The main window is titled "get the invoice summary" and contains a list of 20 numbered steps in natural language, such as "the file", "get the file as a scanned document", "split the scanned document into subdocuments where", "the confirmation texts are """, "the exclusion texts are """, "the subdocument identifiers are 'Group Number'", "imagine the invoicebatches", "the billing template", "get the billing template's table", "get the rows", "process each row as follows", "get the row's group number", "get the subdocument whose group number contains the group number", "add that to the invoicebatches", "process each subdocument as follows", "get the subdocument's first page", "add that to the invoicebatches", "get the subdocument's group number", "merge the invoicebatches into a single document where", and "the document name is 'firstpage_invoices.pdf'". To the right of the text is a "Process Flow" diagram. The flowchart starts with a "Start" node, followed by "Get the file as a scanned document", "Split the scanned document into subdocuments", "Imagine the invoice batches", and "Get the billing template's table". A "Process Row" section includes "Get the rows", which branches into "Get the row's group number" (orange) and "Process each row" (purple). "Get the row's group number" leads to "Get the subdocument with matching group number" (orange), then "Add to invoice batches" (orange), and "Process each subdocument" (purple). A "Process Subdocument" section includes "Get the subdocument's first page" (green), which feeds into "Process each subdocument". The flowchart ends with a "Finish" node.

Now, the underlying technology that ensures, "Explainability", helps make an organization more resilient (as they are less susceptible to losing domain knowledge in aging workforces or turnover), more efficient (as processes are mapped and can be optimized), and more strategic (as the insights into how their business "works" can lead to better decision making and economies of learning). What was once a "requirement" is now a transformative benefit as AI maps the enterprise.

REFERENCES

- ¹ Agarwal et. al, Power and Prediction: The Disruptive Economics of Artificial Intelligence, Harvard Business Review Press (2022)
- ² Grennan, Liz et. al, "Why Businesses Need Explainable AI and How to Deliver It", McKinsey, (2022)
- ³ Modi, Amerdeep et al, "The Evolution of Process Mining: Game Changing Innovations and Future Outlook", Everest Group (2022)