

PwC Generative AI Knowledge Modernization

Harnessing the opportunity



To enable Generative AI, companies need to know it can be deployed responsibly

Generative AI (GenAI) is revolutionizing how firms operate: We've got the knowledge and capabilities and teams to be on the forefront. We are seeing 50% to 90% productivity improvements on use cases today.

How does GenAI work?

GenAI is a subset of deep learning that involves training a model to generate new data. PwC has built out solutions on Large Language Models (LLM) like GPT that are in live production, scaling this capability across our global workforce in a highly regulated industry.

How can it be governed effectively?


PwC has a pioneering factory model, with multidisciplinary agile pods to develop specific business and industry-led use cases built on a foundation of responsible AI to drive efficiency while ensuring governance.

Why PwC?

A recognized leader in responsible AI since 2017, PwC has 186+ use cases across industries ready to deploy. PwC has the knowledge and experience to help deliver end-to-end AI transformation, ranking as an [AI Consulting Leader by Forrester](#) for the past four years.

Has it been tested?

This model is being applied within PwC as client zero. We are building out solutions on GPT that are live in production, built for our regulated industry.



The solution and the scale

PwC Generative AI Knowledge Modernization



Securely and responsibly transform data and knowledge work to a competitive advantage

PwC and Microsoft are using Microsoft's OpenAI services to create the future of knowledge work by enabling Generative AI (GenAI) for specific industry and business use cases built on a foundation of trust.

We are reimagining the way we work with Azure OpenAI Service, internally and with our clients. GenAI is, today, amplifying knowledge work unprecedented efficiency, helping you scale further, work faster, reduce costs and enable new business models.

PwC is helping enterprises leverage this technology responsibly, with sound governance, security and data protection on Azure, from ideation to execution.

Key benefits

1

Safeguard data and securely deploy responsible AI with a proven governance model developed and is being tested with PwC as a first use case and replicated with our other clients.

2

Stand up and scale GenAI at any point in your journey, with offerings from assessment & discovery to integration, deployment & monitoring.

3







Transform knowledge work to drive value realization for enterprise businesses and their IT groups using deep industry use cases powered by Azure OpenAI Services.

4

Responsibly drive 50% to 90% productivity improvements in repetitive tasks such as knowledge capture, process automation and content generation.

To enable GenAI we work with clients from ideation to execution and operations



PHASE	 2-4 hours	 4 hours	 3-24 months	 3-24 months	 3-24 months	 Ongoing
OFFERINGS	Interest Intake	Demo Workshop	Scoping	Discovery	Delivery	Stewardship
CLIENT OUTCOMES	<p>Understand client's challenges and curiosity</p> <ul style="list-style-type: none"> Ideate on "Art of the Possible" Determine target functions & processes for demonstrations 	<p>Interactive session to collaboratively 'generate' real demos for client</p> <ul style="list-style-type: none"> Understand Gen AI & capabilities "in action" Collaborate on uses Define target functions & demonstrate sample use cases 	<p>Initial Assessment</p> <p>GenAI Solution Design</p> <p>GenAI Roadmap</p> <ul style="list-style-type: none"> Develop and prioritize use cases Assess Technology & Data Gaps Create GenAI Vision High level GenAI Roadmap 	<p>Data Prep & Annotation</p> <p>Prompt Engineering</p> <p>Model Selection/Adaptation</p> <ul style="list-style-type: none"> Model selection for use case Prompt database and solution Train / fine-tune model Model validation for use case 	<p>Integration & Deployment</p> <p>Transition & Execution</p> <p>AI Factory Enablement</p> <ul style="list-style-type: none"> Setup the "AI Factory" Build & Launch Tooling & Enhancements Implementation into business process and workflows; evangelization 	<p>Governance</p> <p>Monitoring</p> <p>Evaluation</p> <ul style="list-style-type: none"> Evaluation of insights and actions against business objectives Establish Controls & Governance Estimate Costs & ROI

Offerings Key: GenAI Offerings
Existing Offerings

New roles in the AI Factory Pod Model



Pod Leader

Is responsible for the overall direction of the use case and ensuring we are quickly getting to proof of value with GenAI for the sponsored use case. Makes overall decisions on the business objectives for the use case.

Business Analyst

Is responsible for the detailed breakdown of the use case objectives and ensuring the pod team is progressing on weekly deliverables for the sponsored use case.

Prompt Engineer(s)

Is responsible for designing and fine tuning prompts to meet the objectives of the use case. Works with model mechanic to validate LLM results and adjusts prompts accordingly.

Model Mechanic(s)

Is responsible for validating the results of the LLM responses and working with the Prompt Engineers to fine tune models to try and achieve more accurate and consistent results.



Data Engineer

Is responsible for accessing, preparing, organizing the unstructured data and embeddings for creating LLM-based solutions.

Data Scientist

Is responsible for ensuring the problem can be solved efficiently using an LLM, with data available and the business objectives. Informs the work of the prompt engineer & model mechanic to maximize accuracy & performance.

Dedicated roles per use case

Supporting roles across use cases

Real-world use cases



Here are some of the results we are currently seeing with our clients:

Regulatory Text Classification in financial services

Regulatory rules are correctly identified ~80-95% of the time and policy review time can be reduced by more than ~50%. Use cases include ESG, tax and audit reporting.

Negation Detection in life sciences

Generating templated clinical trial inclusion exclusion criteria to help drive patient enrollment and trial success, improving the precision of one model pipeline from ~73% to 94%.

Inclusion Exclusion criteria generation across industry

Our approach performs at an accuracy level of ~78% where no automation existed previously.

Insurance pre-authorization in healthcare


Effectively create pre-authorization letters with optimal approval rates for patient insurance claims, targeting more than ~70% time-savings.

Text Summarization and synthesis for Marketing Content

Reduction in human-hours that goes into enhancing search engine metadata on a monthly basis, as the AI generated outputs can serve as a starting point.

Safety Automated Narrative Generation in manufacturing

~90% touch-time savings for 250,000 cases of inventory a year.



Let's connect



Matt Hobbs

PwC Partner
US Microsoft Practice Leader
matthew.d.hobbs@pwc.com



Matt Labovich

PwC Principal, Data and AI
matt.labovich@pwc.com