

FedGPT

Generative AI models
tailored to federal requirements.

Accenture Federal Services FedGPT delivers generative artificial intelligence (GenAI) capabilities, including chat, search, summarization, and content creation, to fit federal requirements and use cases.

FedGPT adapts state-of-the-art large language models – commercial, open-source, cloud-based, and on-prem – to generate explainable, accurate, and traceable outputs. Built for Federal agencies who want to experiment with and deploy GenAI, FedGPT enables natural language question-and-answering against client datasets through a ChatGPT-like interface. The platform complements and strengthens existing enterprise search capabilities and is deployable in the cloud and on-prem.

Flexible and modular, allowing for build-to-fit customization.

Cumulative power of combined skillsets

FedGPT's modular *skills* are combined and customized to fit client tasks and data sources. The skills are flexible and model-agnostic to meet enterprise and mission needs.

FedGPT Chat – interactive chat through a ChatGPT-like interface

FedGPT Search – retrieve relevant documents and answer questions based on client data

FedGPT Draft – populate templates with structured data for documents such as job descriptions and solicitations

FedGPT Translate – translate across languages and software code bases

FedGPT Query – generate database queries from natural language input

FedGPT API – integrate skills into third-party applications through standard Python interfaces

FedGPT Summarize – ingest client knowledgebases and draft informative summaries

FedGPT CodeQA – answer questions about code and generate snippets in various coding languages

Approach to Federalized LLMs




FedGPT carefully federalizes LLMs for security, safety, and reliability, using guardrails to monitor model prompts and outputs. We apply several techniques to mitigate the risk of hallucinations that are an artifact of Generative AI models.

- Citations are provided for each query, linking model outputs back to ground truth documents
- Confidence scores show the alignment of model outputs with source content, flagging potential hallucinations
- LLMs are tested against a gold-standard question repository to measure performance against factual, multi-part, and single-part questions
- An LLM firewall blocks undesirable queries based on keywords, patterns, and client-specific classifiers

Cost-effective



The combination of models and skills can be run in CPU or GPU environments, bringing cutting-edge AI/ML capability with reasonable resource investment.



Multiple FedGPT models combine for a seamless user experience, while guardrails limit unsafe use and unwanted outputs.

The new generation of generative models we see in today's market reflect a step-change in capability and usability, but the underlying technology can be outdated and inflexible.

In developing FedGPT, we have been able to build from foundations in NLP, large language models, and generative imagery to ensure our platform is secure, reliable, and adaptable.

For more info

Michael Thieme
AFS National Security Portfolio
Technical Account Lead
michael.thieme@afs.com

AccentureFederal.com