



# AI use case inventory

Use case	Intended purpose and expected benefits	Stage of development	Rights or safety impact?
Solicitation review tool	The SRT intakes SAM.gov data for all ICT solicitations. The system then compiles the data into a database to be used by machine learning algorithms. The first of these is a natural language processing model that determines if a solicitation contains compliance language. If a solicitation does not have compliance language, then it is marked as non-compliant. Each agency is asked to review their data and validate the SRT predictions. GSA also conducts random manual reviews monthly.	Implementation and assessment	Neither
Acquisition analytics	Takes detailed data on transactions and classifies each transaction within the governmentwide category management taxonomy.	Implementation and assessment	Neither
City Pairs Program ticket forecast and scenario analysis tools	Takes segment-level City Pair Program air travel purchase data and creates near-term forecasts for the current and upcoming fiscal year by month and at various levels of granularity including DOD versus civilian, agency, and region.	Implementation and assessment	Neither
Category taxonomy classifier	Classification of obligation transactions in to GWAS subcategories when the existing FPDS coding scheme, PSC/NAICS, is insufficient for positive classification of a transaction.	Implementation and assessment	Neither
CPRM contract breach forecast	Applies obligation forecasting to ITC contracts to provide situational awareness of risk of contract ceiling breach.	Implementation and assessment	Neither
Category taxonomy refinement using NLP	Uses token extraction from product descriptions to more accurately shape intended markets for PSCs.	Retired	Neither
Key KPI forecasts for GWCM	Takes monthly historical data for underlying components used to calculate KPIs, and creates near-term forecasts for the upcoming fiscal year. Pilot effort focuses on total agency and category spend — the denominator in multiple KPIs. If the pilot program is successful, the same methodology can be extended to other KPIs.	Retired	Neither
Enterprise Brain	A document repository by Tanjo (tanjo.ai) that leverages AI to improve document discovery.	Retired	Neither
AWS CAPTCHA solver	Part of a project to automate actions in AWS console using Selenium. We developed an AI model that can solve AWS CAPTCHA so our automated system won't be blocked for being detected as a bot. The model is used in an AWS Lambda that uses a headless Selenium driver. The model was trained from scratch using Tensorflow. We have	Initiated	Neither

Use case	Intended purpose and expected benefits	Stage of development	Rights or safety impact?
	discussed making a more general form of this solution that could be applied to other AI and ML problems.		
ServiceNow generic ticket classification	We are building a model to take generic Service Now tickets and classify them so that they can be automatically re-routed to the correct team that handles these types of tickets. The process of re-routing generic tickets is currently done manually, so the model will allow us to automate it. The initial model will target the top five most common ticket types.	Initiated	Neither
ServiceNow virtual agent (Curie)	Virtual agent that uses ML to provide predictive results for chat entries. A natural language chatbot (virtual assistant), we named Curie, as part of a multi-model customer service experience for employee's IT service requests leveraging knowledge-based articles.	Initiated	Neither
G-REX AI document classification	Office of Leasing and IT Modernization partnered in a pilot program to test the feasibility of using artificial intelligence and machine learning using the AWS cloud service from FAS to classify PDF documents emailed into G-REX. The pilot built the pipeline to train individual document types that are typically uploaded in PDF format. It was successful at illustrating how time could be saved by completing this project in production, improving accuracy of document type assignment. The pilot utilized AWS FAS service, which would require additional funding. Appian announced they will provide the functionality as an out of the box feature in late 2024 or early 2025, to be able to scan PDF document types if G-REX can move to the cloud. The cloud assessment is ongoing as of submission of the need. This was an approved FY 2023 need.	Initiated	Neither
Contract acquisition lifecycle intelligence	CALI tool is an automated machine learning evaluation tool built to streamline the evaluation of vendor proposals against the solicitation requirements to support the source selection process. CALI is offered by Octo Consulting as infrastructure as a service, or IaaS, and is implemented and hosted within Amazon Web Services public cloud US East and US West regions. Once the contracting officer has received vendor proposals for a solicitation and is ready to perform the evaluation process, the CO will initiate evaluation by sending solicitation documents along with all associated vendor proposal documents to Bizagi's source selection module, which will pass all documents to CALI. CALI will process the documents, associated metadata and begin analyzing the proposals in four key areas: format compliance, forms validation, reps and certs compliance, and requirements compliance. The designated evaluation members can review the evaluation results in CALI and submit finalized evaluation results back to source selection in Bizagi. CALI is currently being trained with sample data from the EULAs under the Multiple Award Schedule program.	Retired	Neither
Survey comment ham/spam tester	The website USA.gov receives a lot of survey comments. We seek a way to determine which comments are worth the time of analysts to	Retired	Neither

Use case	Intended purpose and expected benefits	Stage of development	Rights or safety impact?
	read and which are not.		
Classifying qualitative data with Medallia	USA.gov has a lot of customer-driven qualitative data that we would like classified into topics: survey comments from multiple surveys, web chat transcripts, contact center agent case notes, and search box queries. We want to know what the topics of the text chunks are about. This is very challenging because of the massive diversity of USA.gov covered content and topics. Medallia offers natural language processing in which users can create rules based on words and their relationships with other words to tag qualitative data with our topics, such as passports, tax refunds, and more. We are focusing on this. Medallia also uses AI to do classifying on its own. However, we have not derived value from this functionality and therefore do not use it.	Acquisition and/or development	Neither
Elastic machine learning threat detection	Use AI to further analyze security logs and alert on anomalous data patterns for a human SOC analyst to review. The goal is for a human and AI partnership that can more effectively and quickly search through the data.	Acquisition and/or development	Neither
Elastic machine learning threat detection, phase 2	Leverage an LLM that has been given access to our security logs, to provide a more efficient search and discovery tool for SOC analysts. The benefit would be helping the SOC analyst find records of interest and associated logs more quickly. It will also analyze logs automatically and suggest records of interest.	Initiated	Neither
Document workflow and intelligent data capture and extraction	GSA is driving towards a more accurate and scalable document workflow platform. GSA seeks to intelligently capture, classify, and transfer critical data from unstructured and structured documents, namely PDF files, to the right process, workflow, or decision engine.	Retired	Neither
OAS Kudos chatbot	The Office of Administrative Services would like to utilize a ChatBot to easily capture employee peer to peer recognitions. The OAS Kudos Chatbot leverages Natural Language Processing, inherent to Google Dialogflow, to facilitate a conversation between the AI software and the end user. This allows the chatbot to interpret and apply branching logic to what the user enters into the chat window.	Retired	Neither
NCMMS AI chatbot	Provide inline tech support for NCMMS and IBM Maximo via a chatbot. This is an off-the-shelf product that provides natural language-driven, contextual help using Maximo.	Initiated	Neither

Use case	Intended purpose and expected benefits	Stage of development	Rights or safety impact?
Login.gov	<p>The purpose of Login.gov’s identity verification service is to establish a link between the claimed identity and the subject creating a user account. The facial matching AI component of this service provides enhanced controls to protect users from identity fraud and the government from fraudulent access. The benefit to the user is to provide a fast, convenient way to prove their identity remotely and protect the public from identity theft. The benefit to the government is that government partners may use this tool to verify the identities of people accessing government websites in a remote, unattended session and protect Government systems from fraudulent access.</p>	Implementation and assessment	Rights-impacting
Gemini for Workspace	<p>The purpose of the Gemini for Workspace pilot study is to evaluate how advanced AI capabilities can enhance productivity, collaboration, and efficiency within GSA. Over a 90-day period, 200 users will interact with Gemini to assess its functionality and impact on daily operations. The system offers intelligent assistance for tasks like drafting documents and composing emails, enhances communication through smart scheduling and automated meeting summaries, provides data analysis and insights from large datasets, and automates routine administrative tasks to reduce manual workloads. All features available with Gemini Enterprise feature set will be made available within the test profile instance for pilot users. These features aim to help users accomplish productivity tasks more quickly and accurately, and purport to increase productivity across the agency. Enhanced communication tools are designed to facilitate better teamwork and coordination among staff, while access to advanced analytics could support data-driven decision-making. By automating routine tasks, employees can focus on higher-value activities, contributing further to operational efficiency and innovation. Access to the latest Gemini models can also support code development and debugging practices previously identified in earlier studies. The pilot will help us understand user engagement with AI tools and assess the practical benefits of integrating Gemini into the Google workspace. Users will only be permitted to use the features for non-sensitive use cases that involve public information and datasets.</p>	Acquisition and/or development	Neither
QBIQ test fit layouts pilot	<p>The purpose of the QBIQ pilot is to test the capabilities to rapidly test-fit spaces in consideration to evaluate whether AI can sufficiently be used to speed up vetting of space to fit requirements as well as utilize visualizations as a communication bridging tool with tenant agencies</p>	Acquisition and/or development	Neither

**Questions?** Email us at [caio@gsa.gov](mailto:caio@gsa.gov).