

Generative AI/ML and AI governance for the public sector

Mary Strain

Al and ML Services, EDU Amazon Web Services

© 2024, Amazon Web Services, Inc. or its affiliates.

Challenges we are hearing from state and local government customers

Demand for government services is rising while resources and capacity to deliver them **aren't keeping pace**

Citizens increasingly expect government to **provide modern digital experiences** for conducting online transactions

Aging infrastructure for data capture, storage, and management creates friction for leveraging data for analytics and machine learning

Complex security, privacy, and compliance requirements create barriers to change and block adoption of many SaaS solutions

Risk averse culture and institutional inertia slow innovation

Ŀ

1 MÜ

 $\ge \land \land$

₹@?́^

Generative AI is powered by foundation models (FMs)

Pretrained on vast amounts of unstructured data

Contain large number of parameters that make them capable of learning complex concepts

Can be applied in a wide range of contexts

Customize FMs using your data for domain specific tasks



What could go wrong?

Inaccuracies	Bias	Copyright and IP	Security and privacy
Answers that are factually incorrect, irrelevant, or nonsensical, because of limitations in their training data and architecture	Answers that display discriminatory behaviour resulting in prejudiced or unequal treatment of a particular group or groups	The rights of content creators from whom training data is collected remains uncertain and is currently being challenged	Some model providers use and store data for training purposes. Entire end-to-end data pipelines require security and data privacy controls
"The world record for crossing the English channel on foot is 15 hours"	"Generate a picture of a person cleaning" returns overwhelmingly women	Artists suing creators of foundation models alleging the improper use of its photos	Engineers accidentally releasing source code by putting into a foundation model for debugging

AWS generative Al strategy



Secure customization



The easiest way to build with FMs



Generative AIpowered solutions

Generative AI Stack

APPLICATIONS THAT LEVERAGE LLMs AND OTHER FMs



Amazon Q Business_____



Amazon Q in QuickSight



TOOLS TO BUILD WITH LLMs AND OTHER FMs

— မို္င္က်ိဳ Amazon Bedrock –

Guardrails | Agents | Customization capabilities

INFRASTRUCTURE FOR FM TRAINING AND INFERENCE

🛱 GPUs 🏥 Trainium 🏥 Inferentia 🚯 SageMaker

🖸 UltraClusters 🐵 EFA 🖸 EC2 Capacity Blocks 🎽 Nitro 🛞 Neuron

Amazon Bedrock

The easiest way to build and scale generative AI applications with foundation models (FMs)

Choice of leading FMs through a single API Model customization Retrieval Augmented Generation (RAG) Agents that execute multistep tasks

Security, privacy, and safety



Amazon Bedrock Helps keep your data secure and private

None of the customer's data is used to train the underlying model

All data is encrypted in transit and at rest; data used for customization is securely transferred through customer's virtual private cloud (VPC)

Data remains in the Region where the API is processed

Support for GDPR, SOC, ISO, CSA compliance, and HIPAA eligibility

Guardrails for Amazon Bedrock

Implement safeguards customized to your application requirements and responsible AI policies Apply guardrails to multiple foundation models and agents for Amazon Bedrock

Configure harmful content filtering based on your responsible AI policies

Define and disallow denied topics with short natural language descriptions

Redact or block sensitive information such as PIIs, and custom Regex

NEW

Amazon Q

Your generative AI assistant designed for work that can be tailored to your business, data, code, and operations Provides interactive answers, solves problems, generates content, and takes action

Understands your organizations information, code, and systems

Personalizes interactions based on your role and permissions

Built to be secure and private



Generative AI application

Data foundation

S T O R A G E

GOVERNANCE & COMPLIANCE

DATABASES, ANALYTICS, & DATA LAKES

DATA INTEGRATION

Your data is the **differentiator**



Personalized virtual assistant

THE OBJECTIVE

An agency wants to generate a more personalized virtual assistant

The data

Constituent DATA



Past benefit history



Current benefits



aws





Agency DATA

Confirmed eligibility records

Current regulations



Other public services



Similar constitutent services



RETRIEVAL-AUGMENTED GENERATION (RAG)

Specialized knowledge through prompt augmentation Enterprise knowledge corpus No change to the foundation model



FINE-TUNING

Specialized knowledge for specific tasks Small number of labeled examples Change a copy of the foundation model

CONTINUED PRE-TRAINING

Generalized and specialized knowledge for your domain Unlabeled, unstructured enterprise data Change a copy of the foundation model



Make generative AI work with **your data**



Bedrock Studio – Now in public preview



SSO web interface

Rapid prototyping environment and streamlines access to multiple Foundation Models (FMs) and developer tools in Bedrock

Configure one or more workspaces for their organization in the AWS Management Console for Bedrock, and grant permissions to individuals or groups to use the workspace

Generative AI public sector application examples



Constituent communications

Citizen engagement and feedback, transparency



Public health

Personalized care, population health assessments



Public safety

Public safety and crime prevention, emergency response and disaster management

Transportation

Traffic optimization, autonomous vehicle control, personalized transportation experiences



Finance

Budget optimization, fraud detection, risk assessment, and mitigation



Constituent services

Urban planning, personalized urban services



Energy and utilities

Energy management, waste management, smart grid optimization



Research and engagement

Environmental monitoring



Generative AI use cases across public sector

TEACHING AND LEARNING	ADMINISTRATION AND OPERATIONS	LIBRARIES, RESEARCH, AND ATHLETICS
CHATBOTS	VIRTUAL AGENTS	META DATA CREATION
VIRTUAL TUTORS	PROCESS OPTIMIZATION	NEW INSIGHTS XXL DATA
GUIDED PATHWAYS	CONTENT CREATION	REPORTING
HYPER- PERSONALIZATION	CODE GENERATION	COMPLIANCE
COURSE CONTENT	DATA TO INSIGHTS	RECRUITMENT
		MEDIA GENERATION

Generative AI-Enabled Citizen Engagement

Engage citizens and drive improvements in customer satisfaction

- Improve contact center agent effectiveness with real-time translation and decision support using Amazon Connect and Contact Center Intelligence
- Analyze call and text interactions with citizens to spot issues and trends and drive improvement
- Improve self service









Add AI/ML to existing contact centers





Using AI to improve agent efficiency

"During peak hours, previously you're 45-50 minutes on hold, and now that's has been reduced to about three and a half minutes. One of the other benefits we've gotten from Amazon Connect is sentiment analysis. On a call, we get real-time feedback on whether or not the customer was happy, frustrated, or angry..."

—Benny Chacko, Deputy General – LA County Internal Services Department





Extract insights from unstructured content

Extract insights from unstructured documents and forms, like images, PDFs, and audio

- Analyze text with natural language processing (NLP) to identify topics, extract entities, understand sentiment, and classify documents with Amazon Textract, Amazon Rekognition, and Amazon Comprehend
- Translate content at scale with Amazon
 Translate



King County Assessor's Office



CHALLENGE

Reduce data entry, eliminate data errors, and improve data time lines.

SOLUTION

Intelligent documents processing for documents and electronic files, streamlining & unlock data and information from paper documents and electronic files

RESULTS

King County employees will focus on higher value, more satisfying work, and ultimately help the County realize its vision for connected communities, connected data, and connected government."



This Photo by Unknown Author is licensed under CC BY-SA



Student advising assistant: University of British Columbia

- Science advising office and the UBC CIC designed a solution that:
 - Generates a conversational response that draws on relevant information from the UBC Calendar and web content by the Faculty of Science
 - Enhanced accessibility and readability of the information



2023 The year of proof of concepts





Get moving with generative AI



Select the right use case

© 2024, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Assessing the fit of generative AI use cases



Selecting the right use cases



Feasibility: Data, technology, skills, risk/compliance

Value: productivity,

Get moving with generative Al





Select the right use case

Empower your teams to innovate with tools and trainings



Comprehensive, customized training for every stage of your cloud journey



Skills assessment and personalized training plans

Our **Learning Needs Analysis** helps you understand where your gaps are and design a plan to target the Al and ML skills that are most important to your business goals.



Blended training designed by AWS experts

Digital courses, instructor-led classes, and game-based learning meet teams' diverse training needs.

AWS Skill Builder: 100+ learning resources on AI and ML, including generative AI. Interactive challenges in a secure sandbox environment.



Industry-recognized credentials

Validate expertise with certifications like **AWS Certified Machine Learning – Specialty.** Retain top talent by funding exam fees.



Get moving with generative AI







Select the right use case

Empower your teams to innovate with tools and trainings

Get started on your top use cases



Find the right **AWS Partners** to fit your business needs



<u>aws</u>

Strategic recommendations on generative AI

EYES WIDE OPEN

- Culture of continuous experimentation
- Prevent early dependencies

FLEXIBILITY IS KEY

aws

- Innovation requires flexibility
 free from technical or contractual lock-ins
- Infrastructure supporting third-party Generative Al integration
- A breadth of services ensures long-term flexibility and business value



NO ONE SIZE FITS ALL

- Generative AI as an enhancement not a replacement
- Requirements in enterprises differ based on financial resources, security needs, governance, and skills.
- Evaluate buy-vs-build

LISTEN AND ENABLE

- Work backward from your constituents
- Listen to your domain experts
- Enable employees with the right set of tooling

Securing the use of generative AI in your organization

DON'T • Outright ban generative AI technologies

- Empower users by creating policies for the use of generative AI technologies
- Refer to and reinforce your existing data policies
- DO'S
- Implement controls to remove harmful/inappropriate/incorrect content from inputs and outputs
 - Threat model your generative AI applications
 - Track model versions as part of your software bill of materials (SBOM)



Responsible AI dimensions



What does **the future** hold for generative AI?



Amazon Al ready initiative

- Eight new and complimentary AI and generative AI courses
- AWS generative AI scholarship, providing more than 50,000 high school and university students globally with access to a new generative AI course on Udacity
- New collaboration with Code.org, designed to help students learn about generative AI



The case for AI in public sector



Engagement across the student (stakeholder) experience cycle and beyond



Access can minimize affective filters for seeking assistance



Personalization for learning styles, interest, opportunities, pathways, and course recommendations



aws

Agency ability to solve problems, seek avenues to success, resources, and networks





Thank you!

Mary Strain straimar@amazon.com