

Generative AI and AI Governance in the Public Sector

Mary Strain (she/her)

Artificial Intelligence and Machine Learning US EDU

AWS

straimar@amazon.com



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



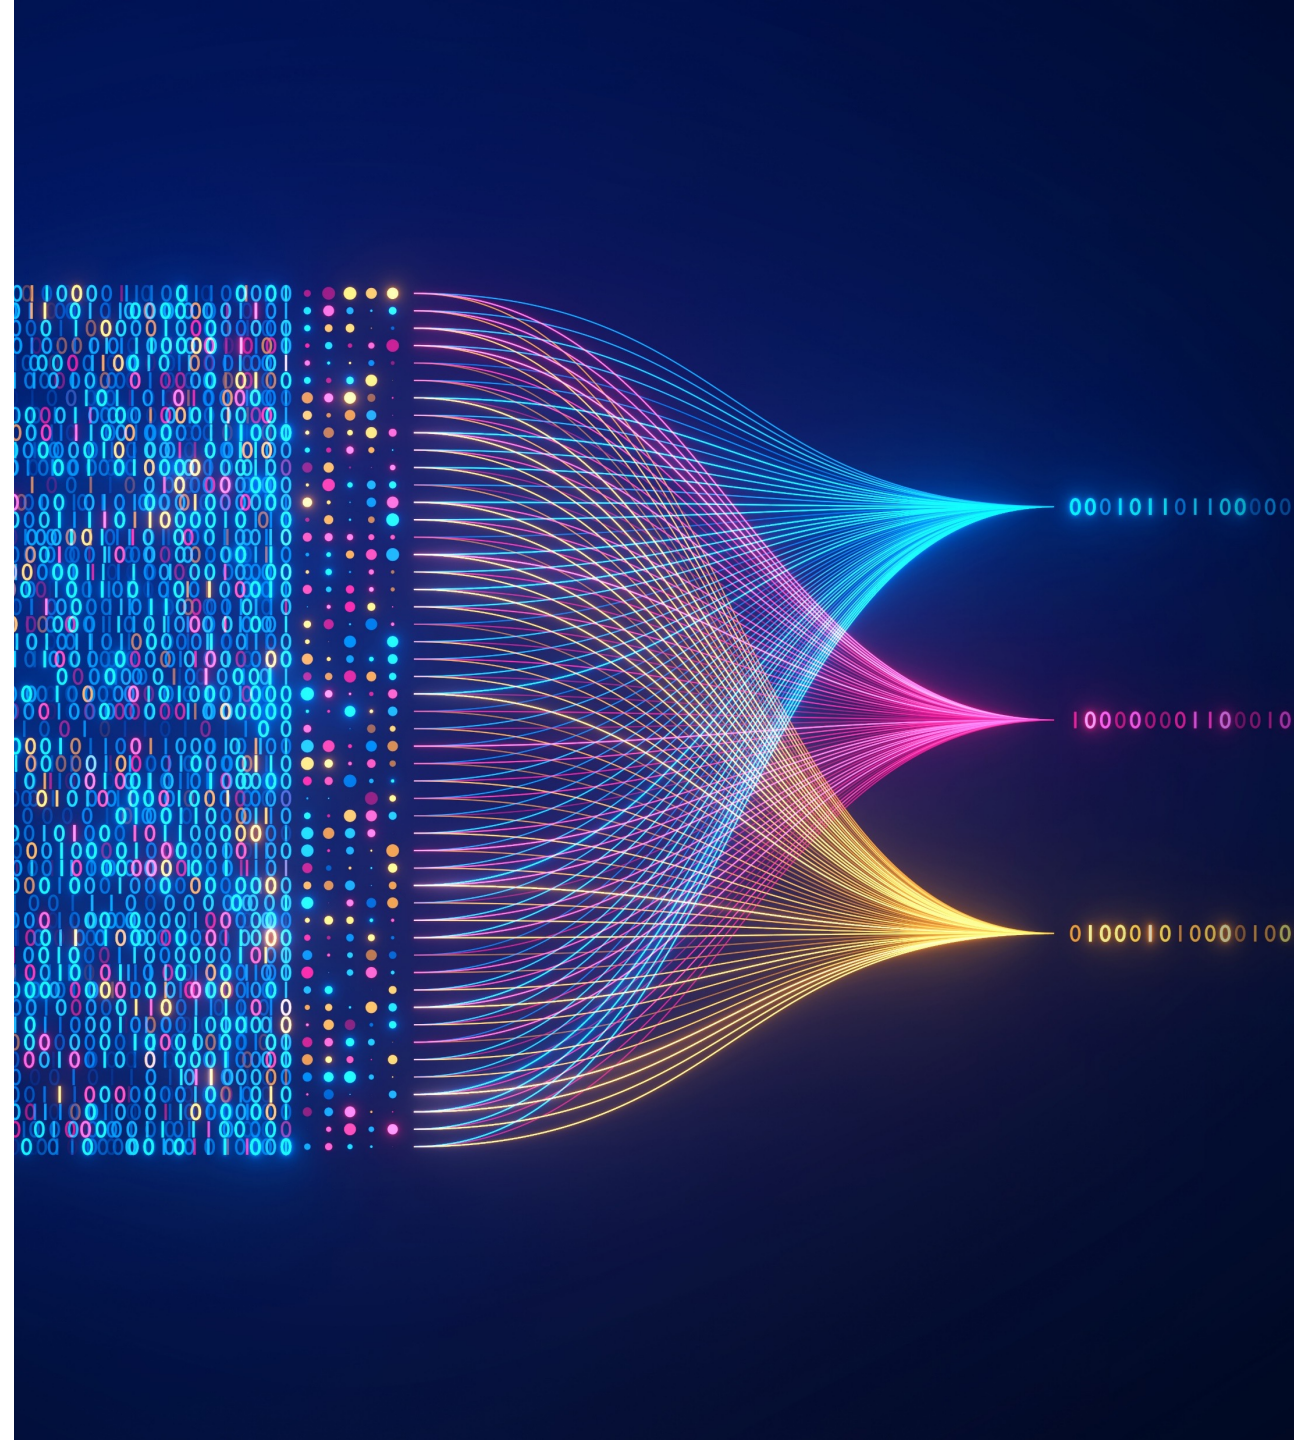
Generative AI is powered by foundation models

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

Answers that are factually incorrect, irrelevant, or nonsensical, because of limitations in their training data and architecture

"The world record for crossing the English channel on foot is 15 hours"



Bias

Answers that display discriminatory behaviour resulting in prejudiced or unequal treatment of a particular group or groups

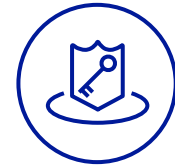
"Generate a picture of a person cleaning" returns overwhelmingly women



Copyright and IP

The rights of content creators from whom training data is collected remains uncertain and is currently being challenged

Artists suing creators of foundation models alleging the improper use of its photos



Security and privacy

Some model providers use and store data for training purposes. Entire end-to-end data pipelines require security and data privacy controls.

Engineers accidentally releasing source code by putting into a foundation model for debugging

Challenges we are hearing from public sector customers



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



Citizens increasingly expect the 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

What we are learning: market challenges

Consumer facing applications for enterprise problems

Relying on foundation models alone to solve problems

Inaccessible, unintegrated and poor quality data

Not aligning use cases to strategic priorities

Shadow AI

Goldman Sachs | Global Macro Research

ISSUE 129 | June 25, 2024 | 5:10 PM EDT

TOP of MIND

GEN AI: TOO MUCH SPEND, TOO LITTLE BENEFIT?



Tech giants and beyond are set to spend over \$1tn on AI capex in coming years, with so far little to show for it. So, will this large spend ever pay off? MIT's Daron Acemoglu and GS' Jim Covello are skeptical, with Acemoglu seeing only limited US economic upside from AI over the next decade and Covello arguing that the technology isn't designed to solve the complex problems that would justify the costs, which may not decline as many expect. But GS' Joseph Briggs, Kash Rangan, and Eric Sheridan remain more optimistic about AI's economic potential and its ability to ultimately generate returns beyond the current "picks and shovels" phase, even if AI's "killer application" has yet to emerge. And even if it does, we explore whether the current chips shortage (with GS' Toshiya Hari) and looming power shortage (with Cloverleaf Infrastructure's Brian Janous) will constrain AI growth. But despite these concerns and constraints, we still see room for the AI theme to



Amazon Bedrock

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

Choice of industry-leading FMs
available via a single API

Customize your models using your
organization's data

Enterprise-grade security and privacy

Amazon Bedrock

AI21labs	amazon	ANTHROPIC	cohere	Meta	MISTRAL AI	stability.ai
Contextual	Text summarization, generation, Q&A, search, image generation	Summarization, complex reasoning, writing, coding	Text generation, search, classification	Q&A and reading comprehension	Text summarization, Q&A, text classification, text completion, code generation	High-quality images and art
Jurassic-2 Ultra	Amazon Titan Text Premier	Claude 3 Opus	Command	Llama 3 8B	Mistral Large	Stable Diffusion XL1.0
Jurassic-2 Mid	Amazon Titan Text Lite	Claude 3 Sonnet	Command Light	Llama 3 70B	Mistral 7B	Stable Diffusion XL 0.8
	Amazon Titan Text Express	Claude 3 Haiku	Embed English	Llama 2 13B	Mixtral 8x7B	
	Amazon Titan Text Embeddings	Claude 2.1	Embed Multilingual	Llama 2 70B		
	Amazon Titan Text Embeddings V2	Claude 2	Command R+			
	Amazon Titan Multimodal Embeddings	Claude Instant	Command R			
	Amazon Titan Image Generator					



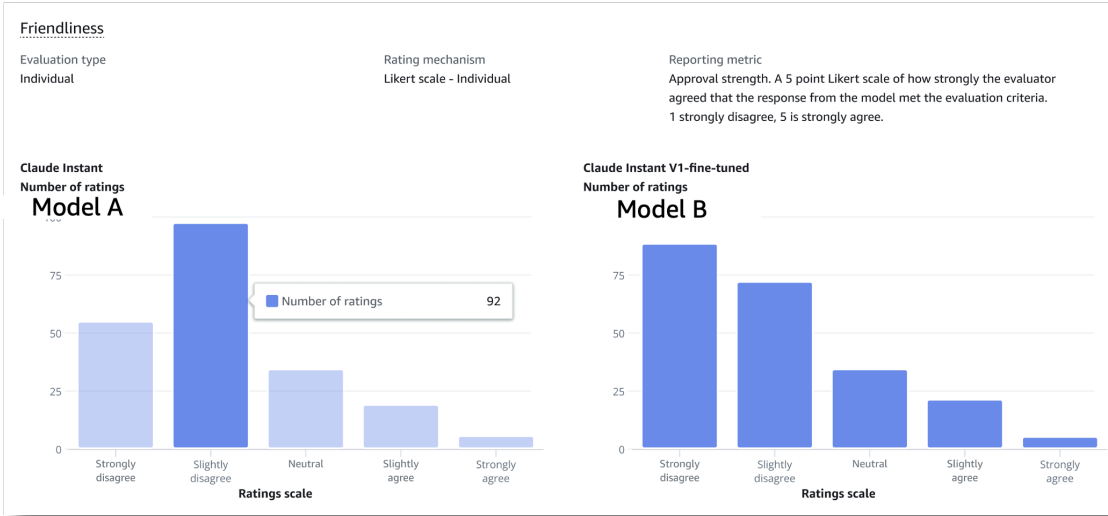
Model evaluation in Amazon Bedrock

EVALUATE FMS TO SELECT THE BEST ONE FOR YOUR USE CASE


Automatic or human evaluation method

Curated datasets or bring your own

Predefined and custom metrics



Text summarization evaluation summary (3)

The results for text summarization consist of accuracy, toxicity, and robustness, which indicate the quality of the summaries generated by the model. [Learn more.](#) 

Accuracy

Dataset	Value
CNN/DailyMail	.6
S3 URI 3	.4

Toxicity

Dataset	Value
S3 URI	.5

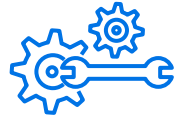
Robustness

Dataset	Value
CNN/DailyMail	.4
S3 URI 2	.6



Amazon Bedrock

Helps keep your data
secure and private



None of the customer's data is used to train the underlying models, nor is it shared with Model Providers



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



Data remains in the region where API processed



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

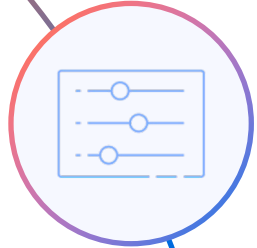
Compliance



Current status:

- **HIPAA** eligible
- **GDPR, SOC and PCI** compliant
- **FedRAMP Moderate – In scope**
- **FedRAMP High – 3PAO Assessment** [ETA September 2024]

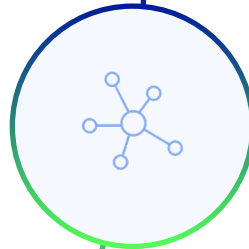
Guardrails for Amazon Bedrock



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 sensitive PII information in FM responses

AWS Generative AI Stack

APPLICATIONS THAT LEVERAGE LLMs AND OTHER FMs



Amazon Q
Business



Amazon Q
Developer



Amazon Q in
QuickSight



Amazon Q in
Connect

TOOLS TO BUILD WITH LLMs AND OTHER FMs



Amazon Bedrock

Guardrails | Agents | Customization Capabilities

INFRASTRUCTURE FOR FM TRAINING AND INFERENCE



GPUs



Trainium



Inferentia



SageMaker



UltraClusters



EFA



EC2 Capacity Blocks








Nitro



Neuron



Reinvent work with Amazon Q

BUSINESS		DEVELOPERS	SPECIALIZED USERS
 Amazon Q Business	 Amazon Q in QuickSight	 Amazon Q Developer	  Amazon Q in AWS Supply Chain
KNOWLEDGE SEARCH	UNDERSTAND DATA	PLAN APPLICATION	AGENT ASSIST
SUMMARIZATION	BUILD & REFINE VISUALS	CODE GENERATION	
CONTENT CREATION	BUILD CALCULATIONS	UNIT TESTING	
EXTRACT INSIGHTS	EXECUTIVE SUMMARIES	SECURITY SCANNING	
RESEARCH & ANALYSIS	CREATE DATA STORIES	CODE REMEDIATION	SUPPLY CHAIN
		CODE MIGRATION	
		TROUBLESHOOTING	
		DEVELOPER KNOWLEDGE	



Your data is the **differentiator**

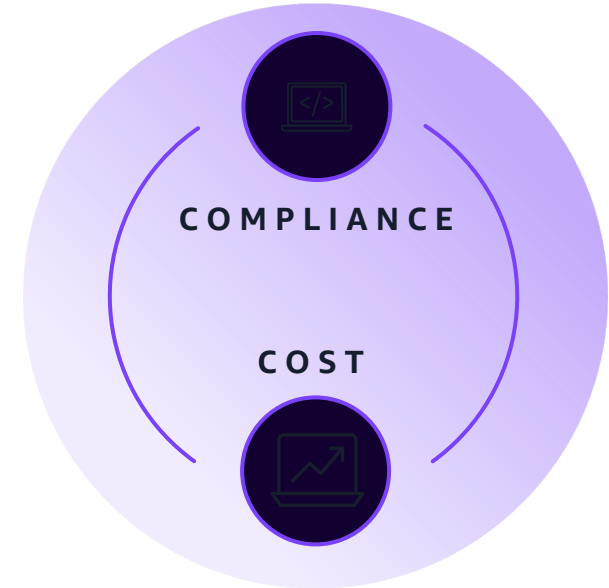
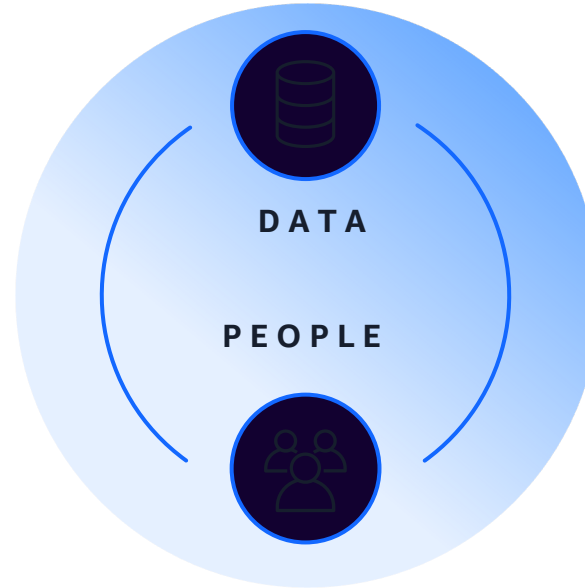


Generic
generative AI



Generative AI that
knows your business
and your customers

Deriving insights from data can be challenging





Make generative AI
work with **your data**



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

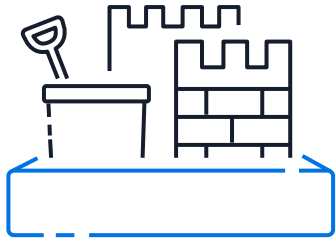
Unlabeled, unstructured enterprise data

Change a copy of the foundation model

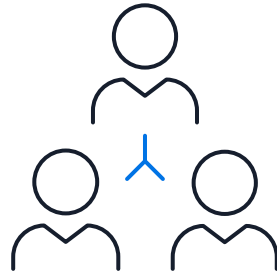
Generative AI for the Enterprise

Bedrock Studio

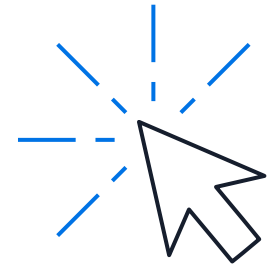
BUILD GENAI APPLICATIONS FASTER AND MORE SECURELY



Easy to use
playground



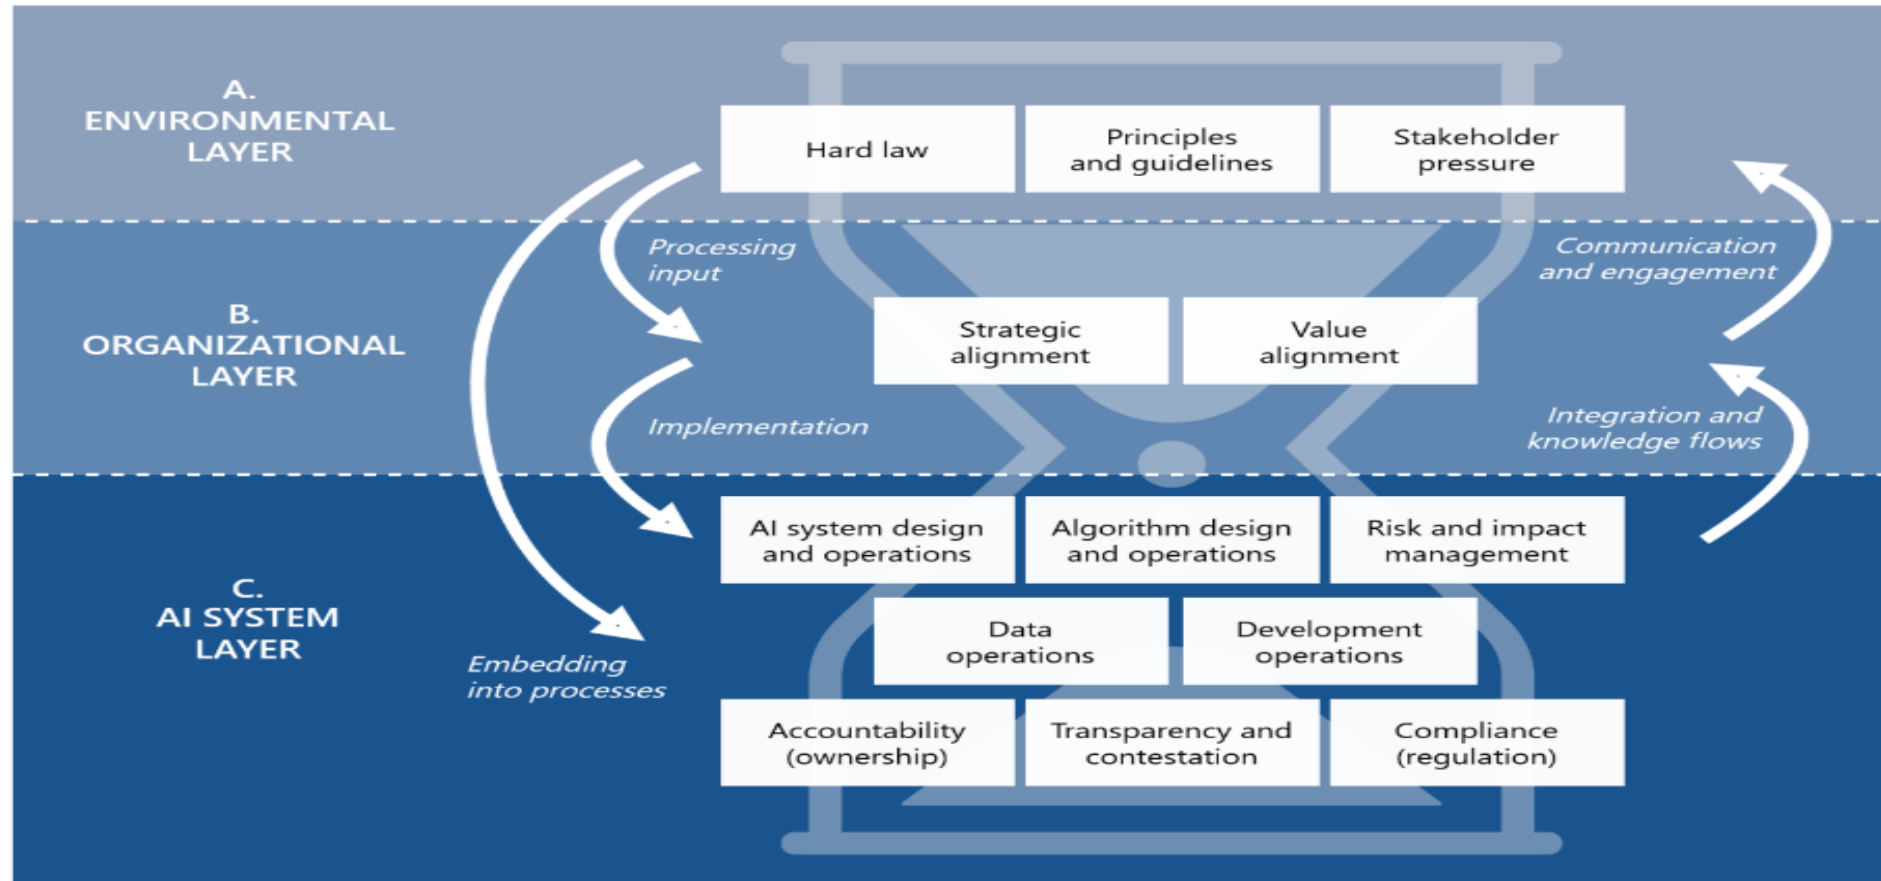
Projects based
collaboration



Easy access with
corporate SSO

Governance and Policy

AIGA: EU AI Governance Framework



Citation: Mäntymäki, M., Minkkinen, M., Birkstedt, T., & Viljanen, M. (2022). *Putting AI Ethics into Practice: The Hourglass Model of Organizational AI Governance* (arXiv:2206.00335). arXiv. <https://doi.org/10.48550/arXiv.2206.00335>

Pillars for the responsible use of AI

Value alignment

Systems should be designed and used in ways that align with company mission, social norms, and legal compliance

Inclusion

Inclusion of unique skills, experiences, perspectives, and cultural backgrounds

Training & education

Appropriate knowledge sharing and education to understand purpose, use, and impact

Accountability

Structured maintaining human involvement and responsibility for design, development, decision processes, and outcomes

Privacy & security

Protects the quality and integrity of data used, its relevance, access, and processing

Fairness

Systems must be designed to minimize bias and promote inclusive representation

Transparency & explainability

Understanding how data is used, how decisions and outcomes are made in a human understandable way



Enterprise Use Cases

2024

The Year of Production

(FOR SOME)



How do I prioritize my projects?

How can I lower my costs?

How do I make this real?

What customization method should I use?



How I can I scale this?

Which models should I use?

Should I train my own model?

How do I manage risks?



How can we move faster?

Top AI/ML use cases for public sector



Speech
and language

SampleOutput.pdf (1 page)

Employment Application

This is a sample **employment** application form. and answer all questions.

Personal Information	
Full Name:	Jane Doe
Phone Number:	555-0100
Home Address:	123 Any Street, Any Town, USA
Mailing Address:	Same as home address

Work History	
Current Company:	Any Company (2018-Current)
	Any Role
Company#1:	Previous Company # 1 (2014-2017)
	Previous Role # 1
Company#2:	Previous Company #2 (2010-2014)
	Previous Role # 2

Intelligent
document
processing



Computer
vision



Predictions
and insights

Top AI/ML use cases for public sector



Speech
and language

SampleOutput.pdf (1 page)

Employment Application

This is a sample **employment** application form. and answer all questions.

Personal Information	
Full Name:	Jane Doe
Phone Number:	555-0100
Home Address:	123 Any Street, Any Town, USA
Mailing Address:	Same as home address

Work History	
Current Company:	Any Company (2018-Current)
	Any Role
Company#1:	Previous Company # 1 (2014-2017)
	Previous Role # 1
Company#2:	Previous Company #2 (2010-2013)
	Previous Role # 2

Intelligent
document
processing



Computer
vision



Predictions
and insights

Generative AI use cases across industries

ENHANCE CUSTOMER EXPERIENCES

CHATBOTS

VIRTUAL

ASSISTANTS

CONVERSATION

ANALYTICS

PERSONALIZATION

BOOST EMPLOYEE PRODUCTIVITY & CREATIVITY

CONVERSATIONAL

SEARCH

SUMMARIZATION

CONTENT CREATION

CODE GENERATION

DATA TO INSIGHTS

OPTIMIZE BUSINESS PROCESSES

DOCUMENT

PROCESSING

DATA

AUGMENTATION

FRAUD DETECTION

PROCESS

Contact Center Intelligence

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



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



King County Assessor's Office



CHALLENGE

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

SOLUTION

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

RESULT

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."



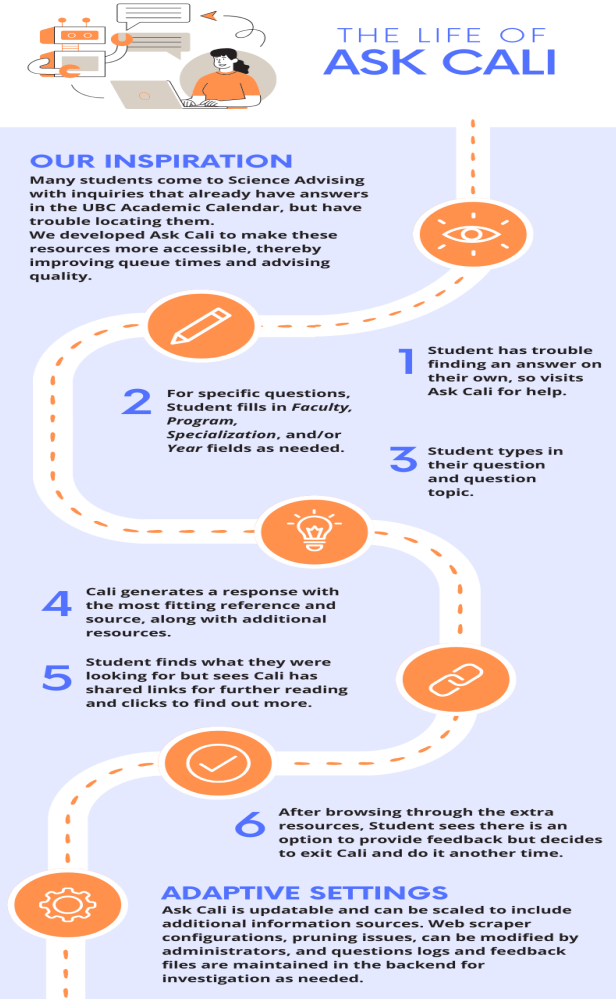
This Photo by Unknown Author is licensed under CC BY

Student Advising Assistant: University of British

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.



AI Assisted Financial Aid: Highline College

Quicker and more transparent way for students to access their information—anywhere and anytime.

Drastically reduced the manual burden on the financial aid department,

lowered the number of emails, calls, and visits about an applicant's status by 75%



UCLA Anderson School of Management : Fundraising

UCLA's Anderson School of Management, a top 20 business school, sought innovation around their approach to alumni fundraising

UCLA engaged AWS partner Slalom, to personalize fund training

Results

19% improvement in email open rate

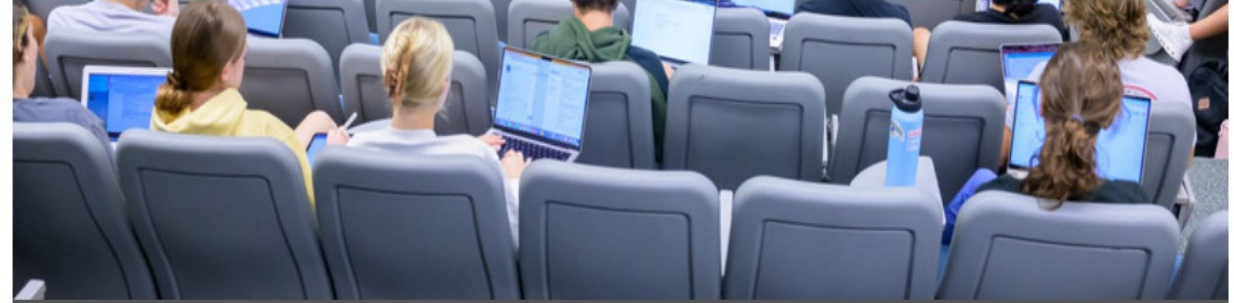
66% improvement in email click-through rate

132% improvement in dollars raised



Virtual tutor: University of Delaware

- **UD Study AiDE**
- Leveraged 300K recorded classroom lectures
- Generated flash cards, study guides, practice tests on UD content across disciplines



Through collaboration with Amazon Web Services, UD will offer students unique digital learning tools and services.

EXPLORING AI INNOVATION

Article by Christopher A. Vito | Photo by Evan Krape | February 19, 2024

Through collaboration with Amazon Web Services, UD will offer students unique digital learning tools and services

While still a nascent tool in many industries, generative artificial intelligence (AI) has been reshaping education as we know it for several years. A 2023 [study](#) revealed that more than half of American college students are incorporating the use of generative AI tools into the completion of their academic work. This study also illustrates that while university faculty are working to catch up to their students, only about 22% of them are using these same AI tools.

GenAI for Cyber

Use Cases

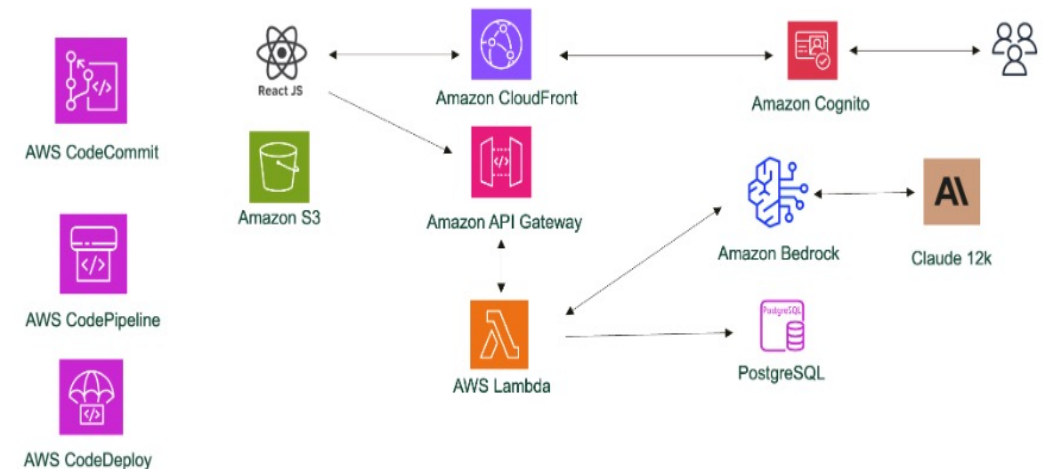
Leverage historical data to identify and inform response to new threats

Agents to gather data and relevant evidence to inform analysis and action

Draft security recommendation

My eCISO: CalPoly and AWS CIC

Natural language Q&A and recommendations aligned to NIST standards



Generative AI for Data



01 Generative AI powered database queries

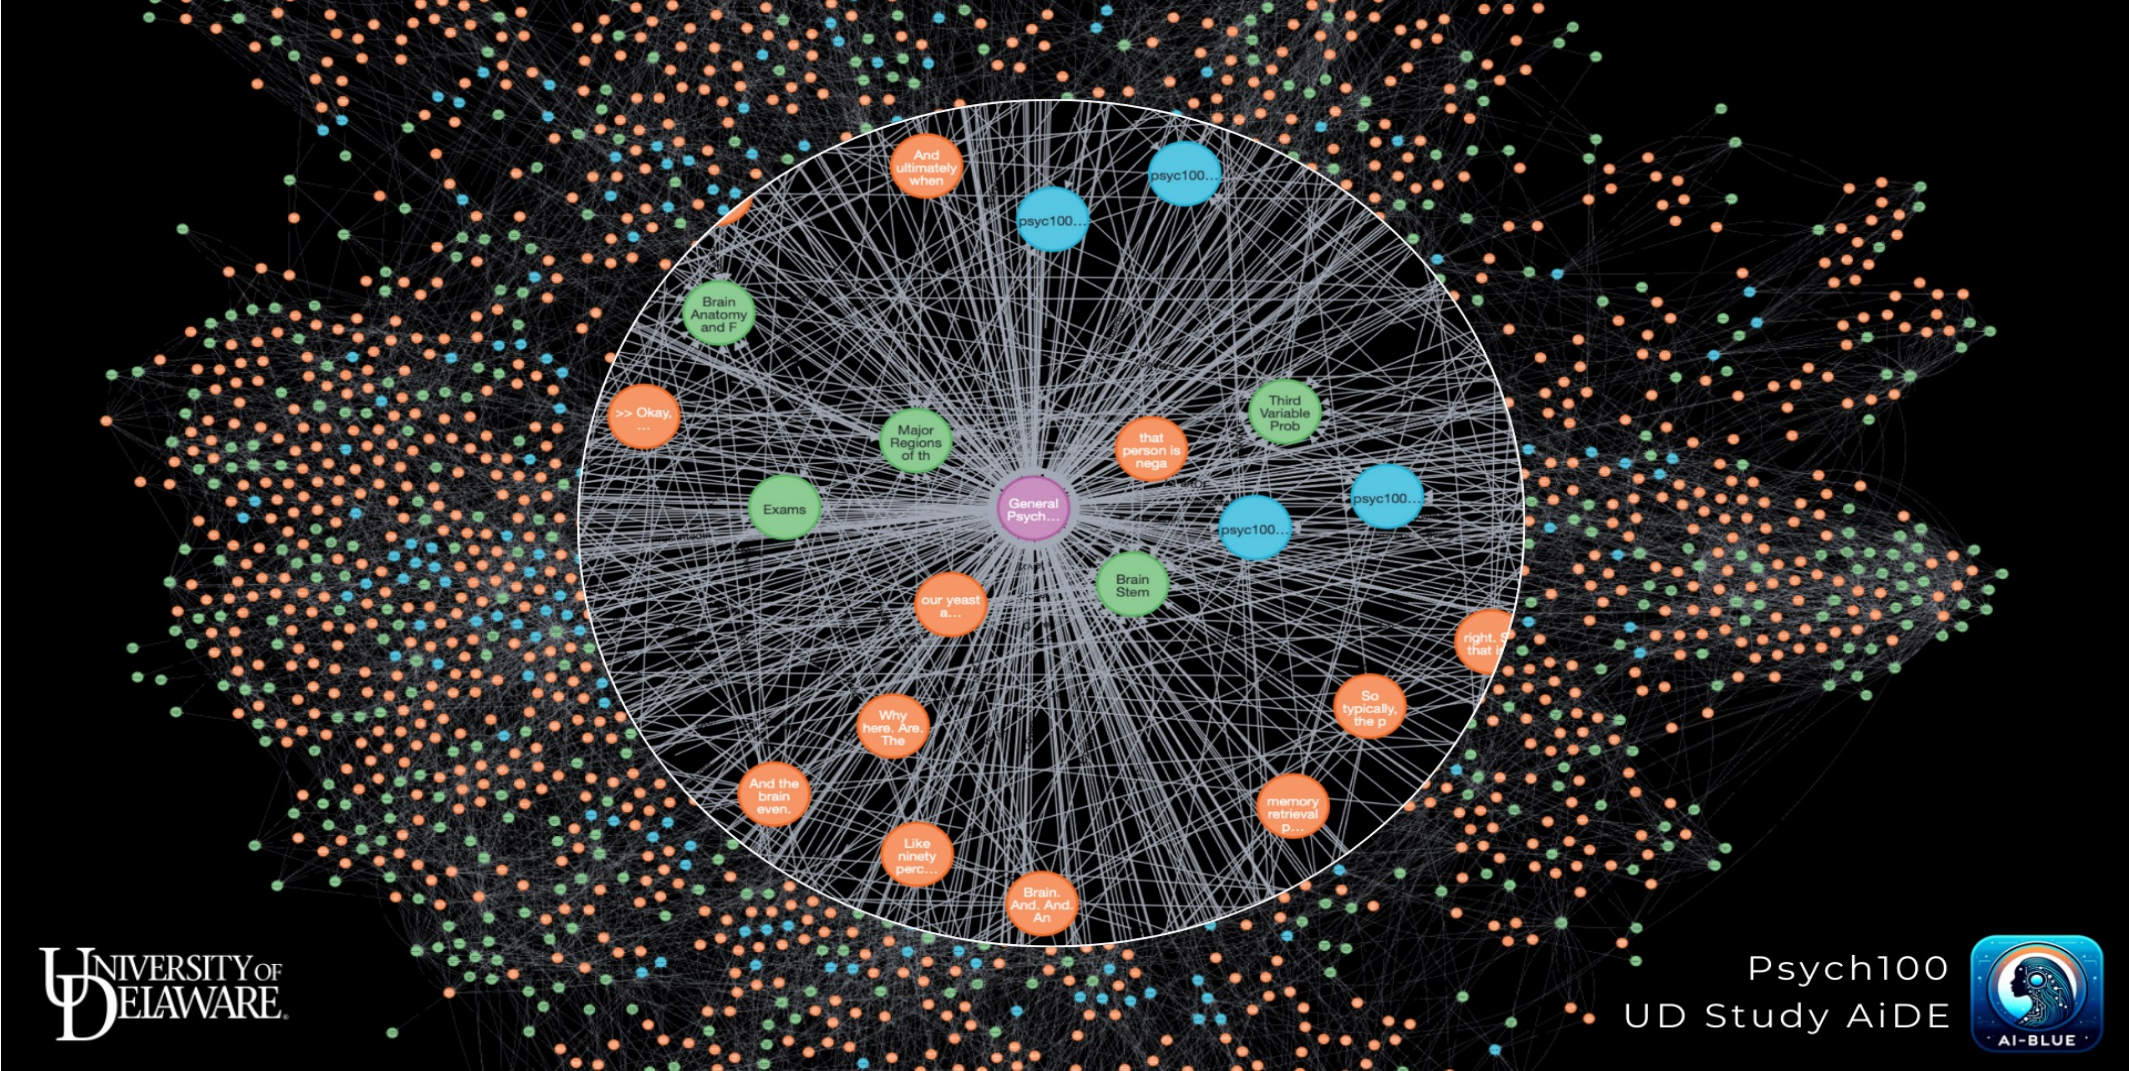
02 Generative AI for data analytics

03 Generative AI powered data mapping

04 Generative AI for data relationship discovery

05 Generative AI for metadata extraction

Generative AI for Knowledge Graph: University of Delaware



The case for AI in public sector

AI PROACTIVELY SUPPORTS

- 1 **Engagement** across (stakeholder) experience cycle and beyond
- 2 **Access** can minimize affective filters for seeking assistance
- 3 **Personalization** for consumption styles, interest, opportunities, pathways, and recommendations
- 4 **Agency** ability to solve problems, seek avenues to success, resources, and resolution





Thank you!

Mary Strain (she/her)

Artificial Intelligence and Machine
Learning, US EDU

AWS

Email straimar@amazon.com

**Please complete the survey
for this session**



**Generative AI and AI
Governance in the
Public Sector**



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



Accelerate your impact!

Rapid AI prototyping and innovation engagements

Noah Eden

Global Lead, Cloud Innovation Centers
Worldwide Public Sector
Amazon Web Services



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



Co-innovation through public-private partnerships

Cloud Innovation Centers enable public sector organizations and AWS to **collaborate and solve real-world public sector challenges**, while providing valuable experience for student interns.

We provide access to **industry-leading innovation processes and cloud technology expertise** to solve organizational and mission-critical problems in the public sector.




CAL POLY
Digital Transformation Hub


Facilitating cyber-resiliency: My eCISO


Using the NIST Cyber Security Framework, we worked with the City of San Diego and the San Diego Cyber Center of Excellence to create a generative AI-based app to support cyber resiliency. It walks users through a natural language interview about their organization's cybersecurity posture and produces a report with NIST-compliant tools and actionable steps for improvement.


My eCISO.

 Hello, I'm eCISO. I specialize in the NIST Cybersecurity Framework for evaluating an organization's cybersecurity posture. Could you kindly introduce yourself and mention the institution you work for?

 Colleen with Cave Creek Unified School District

 Thank you Colleen. Let's start the assessment by discussing your district's cybersecurity practices across the 5 NIST framework functions - Identify, Protect, Detect, Respond, and Recover. I'll ask questions about each function one-by-one. Please provide as much detail as possible in your responses. To begin, for the Identify function, how does your district inventory its assets, manage system configurations, and assess risks?

 We don't currently have a system in place to do this.

 Thank you for that information. It seems like the Identify function, which covers asset management, configuration management, and risk assessments, needs significant improvement. Let's move on to discussing the Protect function. Could you describe your district's identity and access management, data security, and system security measures?

Can be applied to additional frameworks including CIS, CJIS, HIPAA, and more

My eCISO.

DISCLAIMER: MyCISO strives to provide accurate information on cybersecurity frameworks however this application provides content that is educational and for informational purposes only. It is recommended that you conduct additional research and verify the accuracy of this information via trusted sources and additional review by cybersecurity experts.



I accept the terms and conditions.

Improving PDF accessibility

Problem: In April, the Department of Justice published a final rule for Title II of the Americans with Disabilities Act (ADA). It requires that all web and mobile content meets WCAG 2.1 Level AA standards. State and local governments across the US – and higher education institutions – have millions of documents in their knowledge base that do not yet meet this standard.

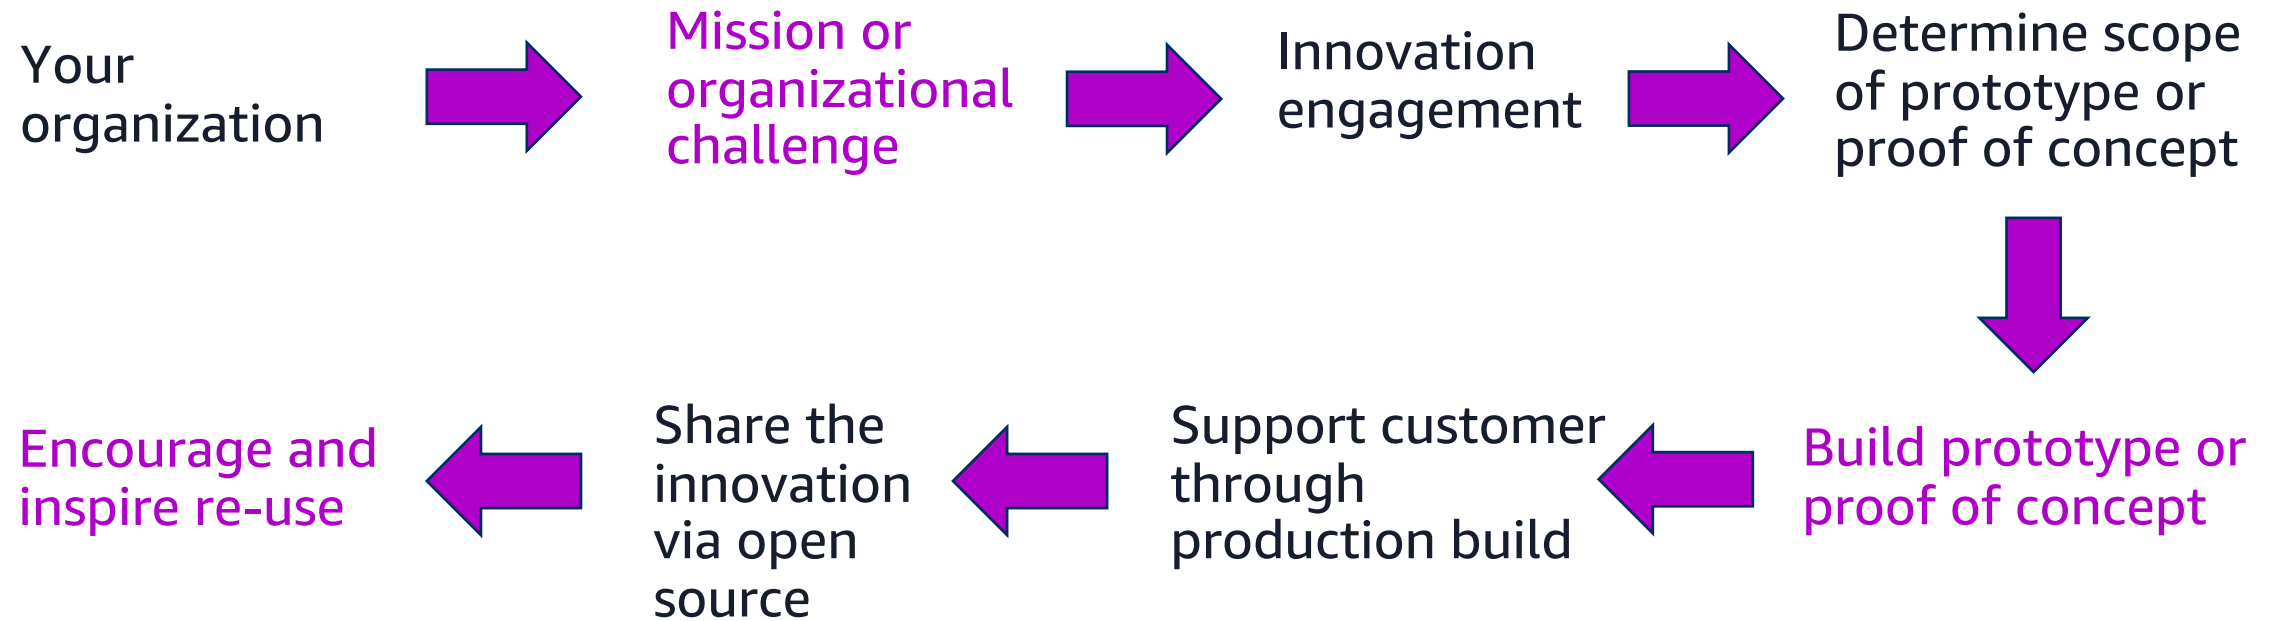
Solution: An open-source prototype that automates the remediation of PDF documents (i.e. incorrect header formatting, erroneous text boxes, etc.) to create a tagged PDF that is at least 90% compliant with WCAG 2.1 Level AA standards. This reduces remediation costs from \$3-4/page (when done manually) to <\$.01/page, plus API costs.

The screenshot displays the PAC (PDF Accessibility Checker) interface. At the top, the PAC logo is visible. Below it, a button labeled 'Open document' is present. A modal window shows the document's metadata: Title (no title), Filename (COMPLIANT_dum0049.pdf), Language (en), Pages (17), Tags (234), and Size (2 MB). A 'Close' button is at the bottom right of the modal. Below the modal, there are three tabs: 'PDF/JA', 'WCAG', and 'Quality'. The 'WCAG' tab is selected, showing a message: 'This PDF file is not WCAG 2.1 compliant.' Below this message is a table with the following data:

Checkpoint	Passed	Warned	Failed
✓ 1.1 Text Alternatives	1	0	0
⊗ 1.2 Time-based Media	0	0	0
✗ 1.3 Adaptable	15544	0	63
⊗ 1.4 Distinguishable	0	0	0
⊗ 2.1 Keyboard Accessible	0	0	0
⊗ 2.2 Enough Time	0	0	0
⊗ 2.3 Seizures and Physical Reactions	0	0	0
✗ 2.4 Navigable	3	0	1
⊗ 2.5 Input Modalities	0	0	0
✓ 3.1 Readable	7544	0	0
⊗ 3.2 Predictable	0	0	0
⊗ 3.3 Input Assistance	0	0	0
⚠ 4.1 Compatible	524	2	0

At the bottom of the interface, there are two buttons: 'Results in detail' and 'PDF report'. On the far right, there are three icons: a key, a bar chart, and an eye.

Engagement lifecycle



What is your top innovation priority?

...to more easily communicate with residents?

...to index your websites and enable smarter, multi-lingual search?

...to quickly analyze or aggregate thousands of paper records?

...to enable students to access the resources and guidance they need?

...to reduce time responding to freedom of information (FOIA) search requests?

Or something else entirely?

Let's innovate and build together.



Thank you!

Noah Eden

Global Lead, Cloud Innovation Centers
Worldwide Public Sector
Amazon Web Services

noaheden@amazon.com

**Please complete the survey
for this session**



Executive track

Generative AI/ML and AI governance
for the public sector