A Peek at the Possible: AI in State Governing
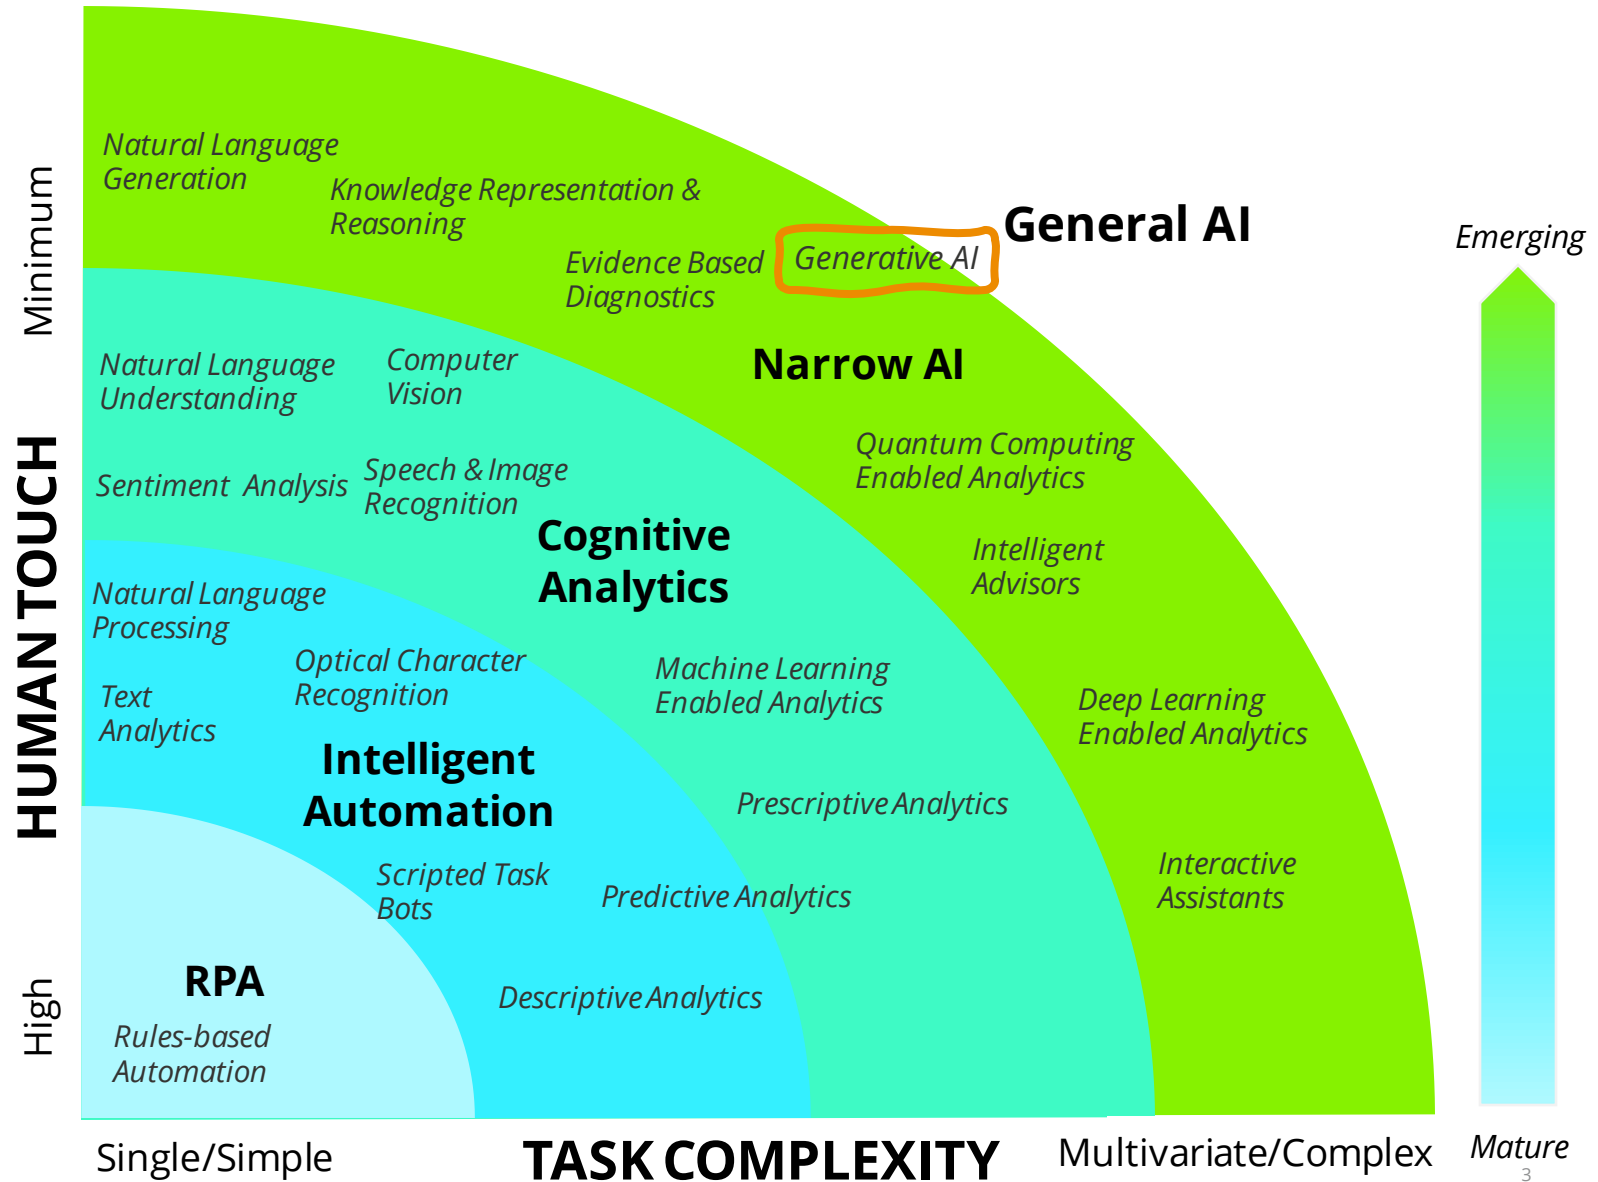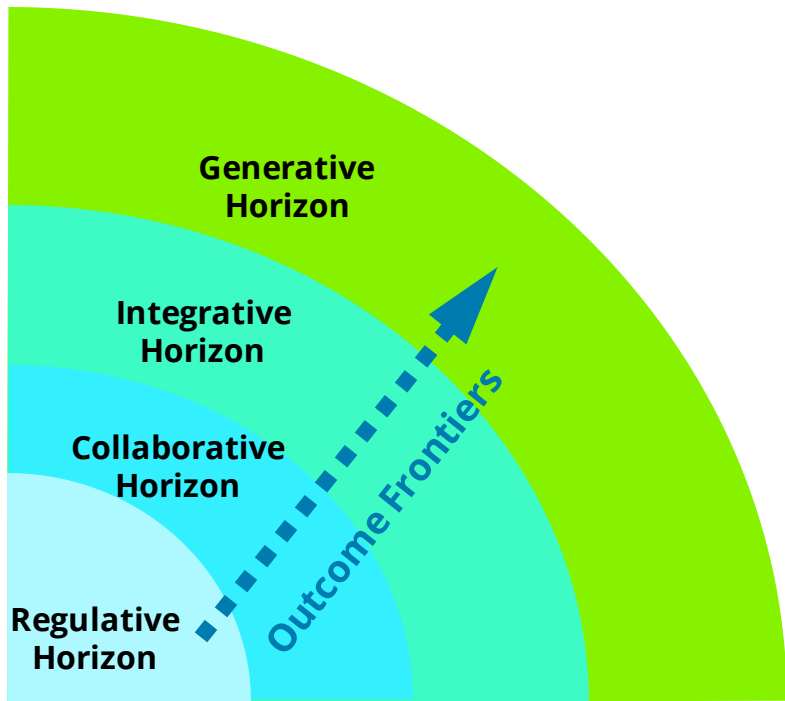
Thursday, November 16th, 2023

Miami Beach, FL

Insert 1st Video Here

The Human Services Value Curve and the Artificial Intelligence (AI) Spectrum

Charts the growth in efficiency of achieving outcomes in four horizons, each of which represent a progressive level of outcomes, impact, and social value.

AI encompasses a wide range of technologies that use inputs of varying complexities to generate human-like outcomes.

Generative Horizon

Integrative Horizon

Collaborative Horizon

Regulative Horizon

Outcome Frontiers

HUMAN TOUCH

Minimum

High

Natural Language Generation

Knowledge Representation & Reasoning

Evidence Based Diagnostics

Generative AI

General AI

Emerging

Narrow AI

Natural Language Understanding

Computer Vision

Quantum Computing Enabled Analytics

Sentiment Analysis

Speech & Image Recognition

Cognitive Analytics

Intelligent Advisors

Natural Language Processing

Optical Character Recognition

Machine Learning Enabled Analytics

Deep Learning Enabled Analytics

Text Analytics

Intelligent Automation

Prescriptive Analytics

Scripted Task Bots

Predictive Analytics

Interactive Assistants

RPA

Rules-based Automation

Descriptive Analytics

Single/Simple   **TASK COMPLEXITY**   Multivariate/Complex   *Mature*

What's Our Mission For Utilizing GenAI?

Act as a ***force multiplier*** *for our people and our clients*

*Reduce the **burden** of human interaction for many types of tasks*

*Use technologies to **guide decisions** and focus on critical tasks*

With Generative AI, we can work smarter, not harder

Summarization
Generate a concise and coherent summary of a long text or a collection of texts
E.g., meeting minutes, news summary, etc.
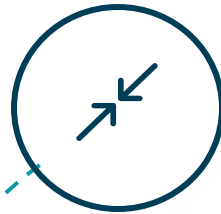
Textual Content Generation
Generate content based on a set of input examples, documents, data and/or a specific theme or topic
E.g., trends report generation, proposals/RFP generation, synthetic data generation

Q&A
Generate a natural language answer to a natural language question based on a given text or a knowledge base
E.g., policy document review, search, chatbot

Reasoning
Generate logical analysis (inference, deduction and/or explanation of relations) given a context or knowledge base
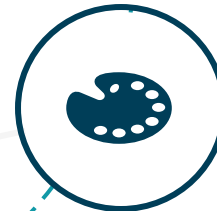*E.g., legislation impact analysis*

Content Transformation
Generate content by converting it into a new specified type, format or style
E.g., text to code, style transfer and personalization, text to table

Image Generation
Generate images from text prompts (descriptions, edits, etc.)
E.g., training data augmentation for computer vision models, Deloitte brand library

Classification
Generate a category or label to a given input
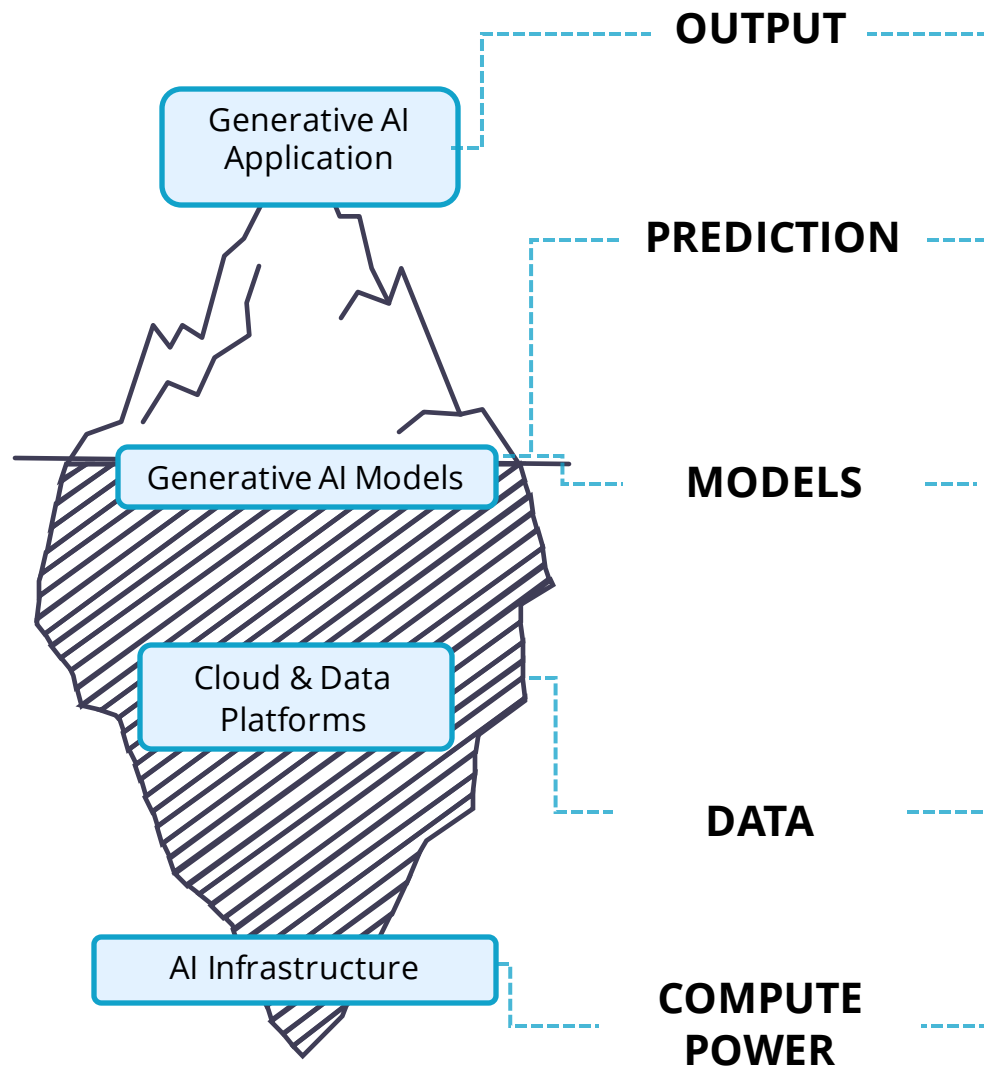*E.g., sentiment analysis, tech vendors classification*

Extraction
Generate specific information or entities from a given input
E.g., tax form key information extraction, data extraction from news for CI database

Mechanism of Generative AI

OUTPUT

Applications…what we see

Generative AI applications generate content across various modalities (e.g., text, image, video, audio) based on how the underlying model was trained.

PREDICTION

Why do these applications seem so human?

Like traditional AI, Foundation Models are models (1) that predict outputs based on inferences on the inputs it receives. However, through, fine-tuning (2), prompt engineering (3) and adversarial training (4) these models generate outputs based on its understanding of human communication.

MODELS

What are Foundation Models?

OpenAI's GPT-4, and NVIDIA's Megatron are two examples of Foundation Models, specifically large language models (LLM) which use deep learning to process huge amounts of data to form 'memories' on the input datasets through tokenization (5), thereby shaping the models' parameters (6). There are common Foundation Model architectures (e.g., Transformer (7), Diffusion (8)) which drive the modalities for each model.

DATA

Training on the world's knowledge

Foundation Models are trained on petabytes worth of global data to shape understanding, tone, and behavior while considering human communication styles.

COMPUTE POWER

Powering our journey to tomorrow

The scale of compute capacity required to train and process Foundation Models necessitates the usage of leading GPUs (e.g., A100 NVIDIA) and TPUs (e.g., Google TPU v4) on scalable infrastructure.

How Do Generative AI Applications Work?

Imagine GPT as an infant...

... learning how to talk. The infant first learns basic patterns of how to talk from their parents, in this case, from the Internet.

Although it is generally very smart, ChatGPT, like infants, would output comments that contain misinformation, biases, or information that is not required. Users then give feedback to course correct. It is a **constant loop** of users giving **feedback,** as the GPT improves the quality of output.

Judgement and edits are crucial steps for humans, as we address the **alignment problem**: how can human beings create a model that reduces the number of biases and that aligns with our human values of what we'd want the world to look like?

TRAINING

ChatGPT learns from any information on the web including Wikipedia, web pages, etc. and excluding most recent data after 2021
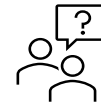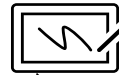
USER INPUT PROMPT

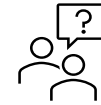Just as you would in Google Search, type in what you'd want to know / do for you

E.g., "How do I make pancakes?" "Write a poem for valentine's day in the tone of Woody Allen"

GPT OUTPUT

GPT will answer your question / request

E.g., "First you put 100g of flour..." "Roses are red...etc"

USER FEEDBACK

Specify or modify the prompt or give a thumbs up / down feedback to rate the output.
E.g. "How do I make pancakes for a calorie conscious mom with no eggs in the fridge?"
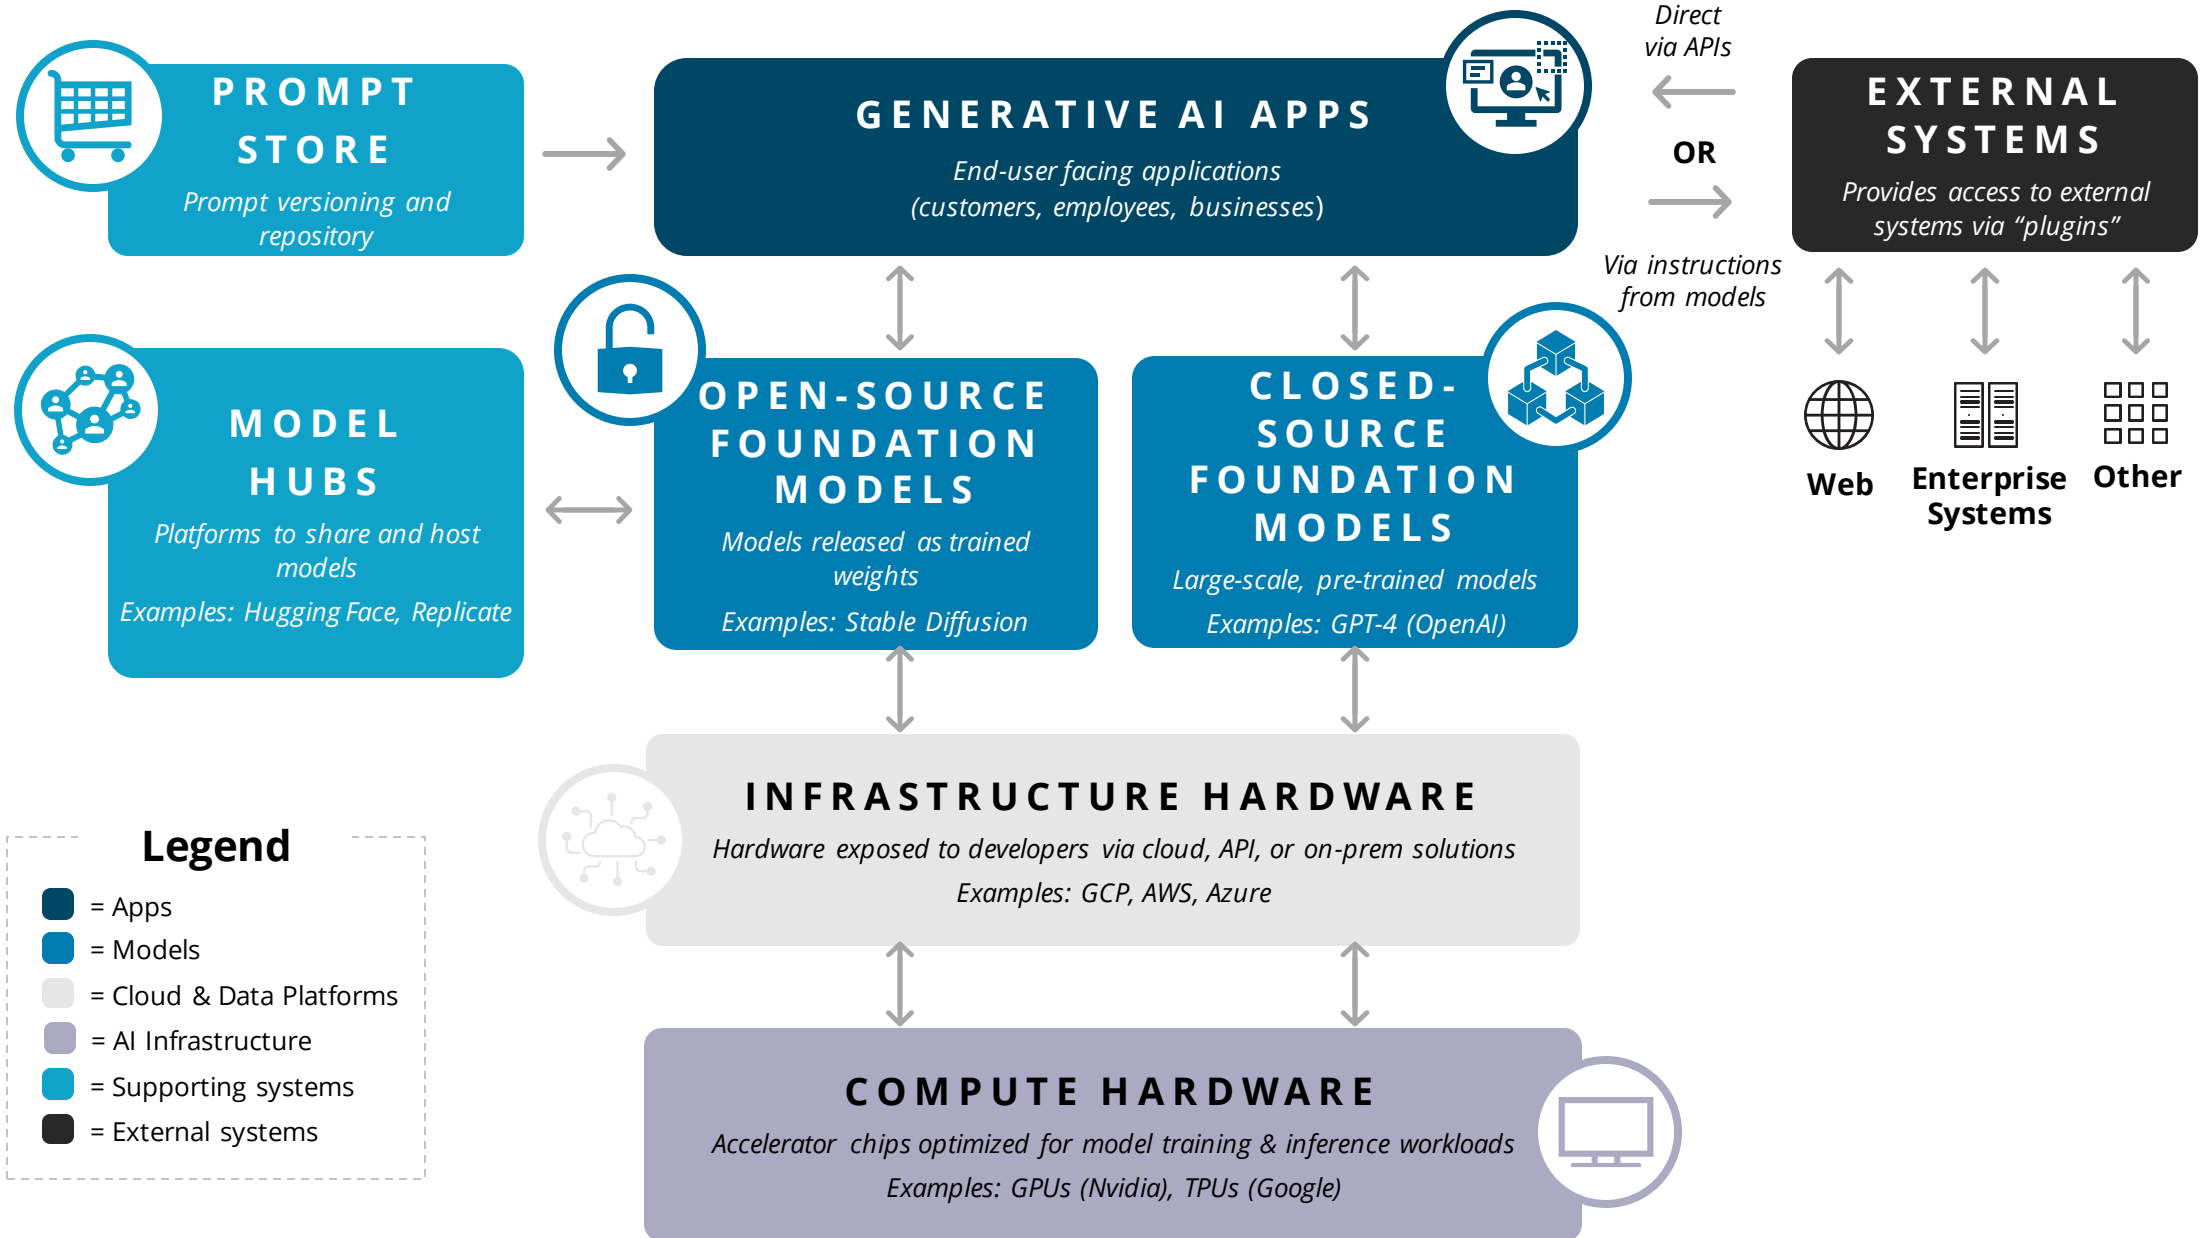
GPT UPDATED OUTPUT

GPT will provide you with, hopefully, an answer closer to what you want

USER EDIT / JUDGEMENT

Make the judgement on the quality of content (text, image etc.) and modify the prompt again if/as needed

High Level Architecture

PROMPT STORE
Prompt versioning and repository

GENERATIVE AI APPS
End-user facing applications (customers, employees, businesses)

Direct via APIs

OR

EXTERNAL SYSTEMS
Provides access to external systems via "plugins"

Via instructions from models

MODEL HUBS
Platforms to share and host models
Examples: Hugging Face, Replicate

OPEN-SOURCE FOUNDATION MODELS
Models released as trained weights
Examples: Stable Diffusion

CLOSED-SOURCE FOUNDATION MODELS
Large-scale, pre-trained models
Examples: GPT-4 (OpenAI)

Web **Enterprise Systems** **Other**

INFRASTRUCTURE HARDWARE
Hardware exposed to developers via cloud, API, or on-prem solutions
Examples: GCP, AWS, Azure

COMPUTE HARDWARE
Accelerator chips optimized for model training & inference workloads
Examples: GPUs (Nvidia), TPUs (Google)

Legend

■ = Apps
■ = Models
■ = Cloud & Data Platforms
■ = AI Infrastructure
■ = Supporting systems
■ = External systems

Ethical, Functional, Legal and Security Risks

Bias

Bias in; bias out. If the training data is biased (e.g., over/under-representation of a population cohort, sexism, racism), then the outputs generated could also exhibit biases. Bias reductions in the training data and/or human supervision during model training is needed.

Text Formatting

Models are good at understanding text, but they struggle when the data are in irregular formats or when the position of the text on the page (e.g., infographic, presentation slide) is relevant to the context and understanding. Other emphasis generators, such as bolded text, font color, etc., don't play a role yet.

Hallucination

Models might output statements that are factually false. Sources and citations are unavailable for most models. Users should be conscious that outputs could be inaccurate and perform due diligence to validate generated content.

IP Protection & Infringement

SaaS-AI companies may use prompt payloads to train future versions of the base model, potentially including confidential data that could expose the user to IP infringement claims – how could this affect your organization's competitiveness in the market?

Secure Infrastructure

External generative AI tools such as ChatGPT introduce an insecure extension of the environment, such as the inability to determine and how that suitable controls are in place for regulatory and policy compliance.

Ethical Use

Is the AI being used in a manner consistent with the purpose of the overall exercise? Is a human being brought into the loop to decide whether the AI's suggestion needs adjustment before actual use or whether the use of AI is ethical (e.g., submitting an AI-generated essay as your own)?

Model Performance

Foundation Models are comprised of billions of parameters (model size) and trained on petabytes of data. In theory, the larger the model, the better the output. Foundation Models take time to produce outputs, which may limit real-time use cases.

Cost

Foundation models generally offer a pay-as-you-go billing mechanism, and the cost per use of sophisticated models is materially significant. Fine- tuning the biggest model and running large documents through several times could quickly run up a bill of tens of US $1,000s.

Malicious Behavior

To maintain operations and customer trust, proactively minimizing risk from malicious behavior on the network is critical. For example, a customer service bot revealing confidential information to a hacker either by prompt or unintentionally.

Confidentiality & Privacy

Generative AI Models are built on data sharing which makes it challenging to maintain an individual's privacy rights. Consent for data used (confidential information, personally identifiable information) is necessary, but poses concerns around an individual's right to consent or be forgotten.

Approaches to Address Risks & Limitations

While not exhaustive, below are 4 key steps to take in response to GenAI risks and limitations

1

AI Literacy

Host informative sessions that focus on educating business and tech leaders, as well as data users, about the responsible utilization of generative AI. These sessions aim to enhance their understanding and knowledge of ethical practices in generative AI usage.

2

Continuously evaluate the changing regulatory landscape

Stay updated on the evolving regulatory requirements related to generative AI and ensure effective communication and integration of those requirements into policies, processes, and implementation approaches.

3

Establishing appropriate audit process for vendors

Implement an appropriate level of testing and validation, whether independent or internal, to allow for visibility and transparency of the impact of risks when generative AI models are implemented.

4

AI Governance Framework

Establish an operating model where experts can help and guide the organization to use generative AI solutions appropriately by accessing the implications for existing processes, proactive policy enhancements, and incorporating necessary safeguards and guidelines into those processes.

AI Federal and State Policy Landscape

AI Federal Policy Landscape

While there is broad agreement on the importance of AI issues across the Biden Administration and Congress, policymakers do not yet agree whether, when, and how AI should be regulated.

BIDEN ADMINISTRATION

- On October 30th, the Biden Administration issued a **long-awaited executive order on artificial intelligence**

- This EO identifies eight principles through which the administration intends to require — or incentivize — certain changes related to AI's development and use, including:
  - **AI safety and security**
  - **Innovation and competition**
  - **Equity and civil rights**

- A lengthy section on **federal government use of AI lays** out new directives to guide how agencies manage AI in their operations and missions, as well as how to increase AI talent in government

CONGRESS

- Most AI laws focus on **governmental uses of AI**

- Now, both the House and Senate are holding hearings and educational briefings[3] on **AI policy issues**, including:
  - **AI governance**
  - **Bias and discrimination concerns**
  - **National security impacts**
  - **Workforce effects**
  - **AI misuse**

- Senate Majority Leader Schumer (D-NY) announced a **regulatory framework and expert forums** to gather additional perspectives

- The House introduced bipartisan legislation that would create a **commission to study AI and make recommendations**

AI State Policy Landscape

------- B A C K G R O U N D -------

EDUCATION

Policymakers at state executive and legislative levels want to **understand how AI is being used in their own states** and educate themselves before legislating on the topic.

LEGISLATION

During the 2023 session, **artificial intelligence legislation** was introduced in **over 22 states and territories**.

WORKING GROUPS

Legislation enacted and executive orders issued in 2023 have pertained to creating **working groups and committees tasked with studying AI and producing policy recommendations.**

Example State Actions

California: Executive order issued to ensure deployment of ethical and responsible GenAI within state agencies.

Connecticut: Legislation established a working group to develop best practices for ethical and equitable use of AI in state government; required an inventory of state systems that employ AI.

Maine: MaineIT established a moratorium for at least 6 months on GenAI in order to conduct a risk assessment.

New Jersey: Executive order issued that establishes AI Task Force to study emerging AI technologies, to issue findings on their impacts, and offer recommendations on ethical use of them.

Oklahoma: Executive order issued to create the Governor's Task Force on Emerging Technologies to develop policy recommendations on responsible deployment of AI and GenAI.

Pennsylvania: Executive order issued to establish responsive and responsible standards for governance of Gen AI when used by state agencies.

Texas: Legislation established advisory council to study and monitor AI systems developed, employed, or procured by state agencies.

Virginia: Executive directive issued to ensure the responsible, ethical and transparent use of AI technology by state government.

Source: California, Connecticut, Maine, New Jersey, Oklahoma, Pennsylvania, Texas, Virginia

Insert 2nd Video Here

Solutions for State Government

GenAI Intelligent Solution Areas

DOCUMENT GENERATION

Application of AI to create documents and reports; saving manpower, reducing error and bias, and allowing Government employees to focus on what matters.

CASE MANAGEMENT

Using AI to parse Government case files for actionable details which are then sorted and prioritized for maximum impact. Guidance and Policy suggestions at the point of decision.

KNOWLEDGE MANAGEMENT

Natural language interface that allows users to cluster, search, and filter their data without coding or data science expertise. Includes image and video data sources.

BACK-OFFICE FUNCTIONS

Application of AI Engine modules and processes that augment the workforce to complete previously tedious back-office functions related to decisions, costs, finance, HR and procurement & logistics.

CUSTOMER ENGAGEMENT

Applying AI to enhance citizen and customer engagement, ultimately making Government services more user friendly, and resources easier to navigate and consume.

How is GenAI changing the conversation

DOCUMENT GENERATION

Application of AI to create documents and reports; saving manpower, reducing error and bias, and allowing Government employees to focus on what matters.

Documentation Engine
Generates documentation and summaries of information to reduce administrative burdens. This can range from developing training documents, system requirement documents, case note summaries, or detailed reports for decision-making.

User Stories and Test Case Generator
Automates the generation of user stories and test cases based on specific requirements. Includes domain-specific knowledge to further refine stories and cases.

Contract Generator and Analysis Engine
Generates contract frameworks or legal documents based on specified criteria. Conducts legal analysis by analyzing and parsing key pieces of information to generate analysis summaries.

Post-Interview Engine
Takes audio recordings and notes from an interview and populates data into system screens and conducts sentiment analysis for evaluation of interview and coaching support.

Notice Designer Engine
Assesses existing notices and recommends design and format changes to make them easier to understand for customers and to conform to defined readability levels.

AI-Assisted Appeals Decision Engine
For each case folder and hearing, documentation is submitted for a hearing which AI ingests to forecast decisions including applicable areas of decisions (e.g., fact findings, law, policy and conclusions).

How is GenAI changing the conversation

CASE MANAGEMENT

Using AI to parse Government case files for actionable details which are then sorted and prioritized for maximum impact. Guidance and Policy suggestions at the point of decision.

Interviewer Assist Engine

Assists caseworkers with benefit evaluations by nudging them with potential questions to ask during interviews and automatically extracts contextual information from the interview that can be prepopulated into screens, case notes, and summaries.

Intelligent Case Management

Assists caseworkers in processing cases by answering questions specific to the case or family they are serving. Connects with the case management system data / eligibility rules and identifies common data entry errors, provides detailed reasoning behind eligibility results, and predicts eligibility outcomes.

Case Management Co-Pilot

Interprets policy to answer questions specific to the case or family being served, or to flag areas for caseworkers to follow-up on. Connects the case and policy data to generate responses that can be validated.

Automated Fact-Finding Reminders

Automates outbound telephone dialing, SMS messaging, and/or emailing campaigns to benefit recipients who are approaching case action due dates.

Case Prep Engine

Provides summary information about a case that can be used for activities such as fair hearings, prep for customer interviews, or foster home placement.

How is GenAI changing the conversation

KNOWLEDGE MANAGEMENT

Natural language interface that allows users to cluster, search, and filter their data without coding or data science expertise. Includes image and video data sources.

Policy Knowledge Engine

Assists caseworkers, staff, and management by answering complex policy, operational procedure, and system questions. The engine takes data such as policy manuals, system documents, and process maps as inputs, and uses LLM models to provide responses and answers (with citations from relevant documents).

Dashboard Generator

Allows users to input natural language to generate and execute SQLs to perform data analysis, drill-down of reports, and synthetization of data/information through reports and dashboards.

Translation Engine

Provides multi-language text generation support for translating notices, including context of translation and quality of translation, from English to other languages.

Training Aid Content Engine

Develops initial versions of training support content such as FAQs and Quick Reference Guides as part of policy changes or enhancements to a system. This content can be leveraged by change management teams or training teams before, during, and after change situations.

How is GenAI changing the conversation

BACK-OFFICE FUNCTIONS

Application of AI Engine modules and processes that augment the workforce to complete previously tedious back-office functions related to decisions, costs, finance, HR and procurement & logistics.

Workload Optimization Engine
Assigns tasks to caseworkers based on unique skillsets, capabilities, and current/anticipated workload.

Funding Engine
Improves resource accuracy and can maximize federal funding by taking information from risk and safety assessments and providing the information to appropriate resources.

Rate-Setting Engine
Recommends provider rates based off network-wide data, best practices and generates analysis reports for decision-making.

Backlog Engine
Forecasts incoming workload on applications, renewals, etc. so that agencies can prepare for staffing and minimize the increase in their backlog.

Churn 2.0 Engine
Provides information and acts on cases that are anticipated to churn to minimize the number of customers who lose their benefits and need to reapply.

Reconciliation Generator
Generates third party liability output based on available documents/materials and conducts reconciliations on payments.

How is GenAI changing the conversation

CUSTOMER ENGAGEMENT

Applying AI to enhance citizen and customer engagement, ultimately making Government services more user friendly, and resources easier to navigate and consume.

Self Service Engine

Provides residents personalized guidance and responses to questions, by using information from existing policy/procedure documents. Assists with "completeness" of application, by nudging the client with appropriate information and proactively identifying data gaps.

Service Desk Assistant Engine

Assists help desk users in researching tickets, answering FAQs related to a tickets, resolving help desk issues, and/or tagging and summarizing help desk calls.

Notice Explanation Engine

Interprets language in customer-facing notices to explain eligibility determinations, changes, or requested information

Provider Search

Searches for and finds providers that meet specific criteria. Example: I am looking for a female provider that supports behavioral health needs within a 10-mile radius.

Application Engine

Assists customers with the application process by answering questions based on the family's situation. Connects with Social Security screens and policy manuals/procedure documents to provide responses.

Referral Engine

Provides automated referrals based on predictive analysis - (e.g. child welfare but also less dramatically remote family support classes if a move is detected)

State Government Focus for Generative AI Use Cases

COMMUNICATIONS & REPORTING

- Document summarization and categorization
- Chatbot / virtual assistant dialogue generation
- Strategic communication generation
- Policy and regulation change analysis and summary
- Report generation and impact analysis
- Video or image editing and generation

HUMAN RESOURCES

- Personal onboarding assistant
- Compensation analysis
- Workforce skill analysis
- Resume analysis
- Metaverse 3D workforce experience
- Metaverse 3D workforce upskilling

CONTRACTING & PROCUREMENT

- Language translation
- Supplier/Provider identification and evaluation
- Automated bidding and proposal evaluation
- Fraud detection and prevention
- FAR adherence and recommendations
- Contract adherence & anomaly detection

GOVERNANCE & OPERATIONS

- Intranet search (knowledge management)
- Process analysis
- Training for new team members
- Document inventory analysis
- News and media summaries
- Sentiment analysis for workforce

INFORMATION TECHNOLOGY

- Legacy code summarization & translation
- Peer review for optimized code writing
- Training on new technologies
- Test automation and test scenario creation
- Development lifecycle documentation
- Code generation across languages/frameworks/CSPs

FINANCE & ACCOUNTING

- Invoice processing and payment automation
- Risk management
- Forecasts and planning
- Financial report analysis
- Regulation and oversight analysis
- Fraud, waste, and abuse prevention

Insert 3rd Video Here

Q & A

Deloitte.

About Deloitte

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited, a UK private company limited by guarantee, and its network of member firms, each of which is a legally separate and independent entity. Please see http://www.deloitte.com/about for a detailed description of the legal structure of Deloitte Touche Tohmatsu Limited and its member firms. Please see http://www.deloitte.com/us/about for a detailed description of the legal structure of Deloitte LLP and its subsidiaries. Certain services may not be available to attest clients under the rules and regulations of public accounting.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor. Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.