Generating Change: The Power of AI in Government

July 31, 2023

Gartner®

Agenda

- Introduction
  - Why Gartner
  - Why Generative AI is so important to state government
- Briefing with Ben Kaner
  - What is Generative AI
  - Examples of how it could be used in Government
  - Risks
  - Guardrails

The Numbers Matter

Category	Gartner
Annual Revenues	**$5.5B**
Total Analysts (Excluding Consultants)	**2, 500**
CIO Membership	**10,000+**
Client Retention Rate	**86%**
Client Enterprises	**15,000+**
Client Inquiries/year	**490,000+**
Peer Community Members/LinkedIn Followers	**200,000/1.5M**
Vendor Briefings/year	**23,000+** (4,000+ unique vendors)
Contract Reviews/year	**11,000+**
Strategy Reviews/year	**20,000+**
Published Docs/year	**11,000+** (Library of 200,000+)
Destination Events/year	**41**
Total Destination Event Attendees	**60,000+**
Global Footprint (Offices)	**71 + 19 in U.S.**

Unparalleled Government Research Depth

- 2,500+ State/Local Government-Specific Research Notes
- Hundreds of Government-specific SMEs with expertise in:
  - Public Safety
  - Law Enforcement
  - Policy development
  - Testifying
  - Citizen Engagement
  - Modernization & Privacy
  - Public Services
  - Data Governance & Sharing

The 2 Sigma Problem: The Search for Methods of Group Instruction as Effective as One-to-One Tutoring

BENJAMIN S. BLOOM
University of Chicago and Northwestern University

FIGURE 1. *Achievement distribution for students under conventional, mastery learning, and tutorial instruction.*

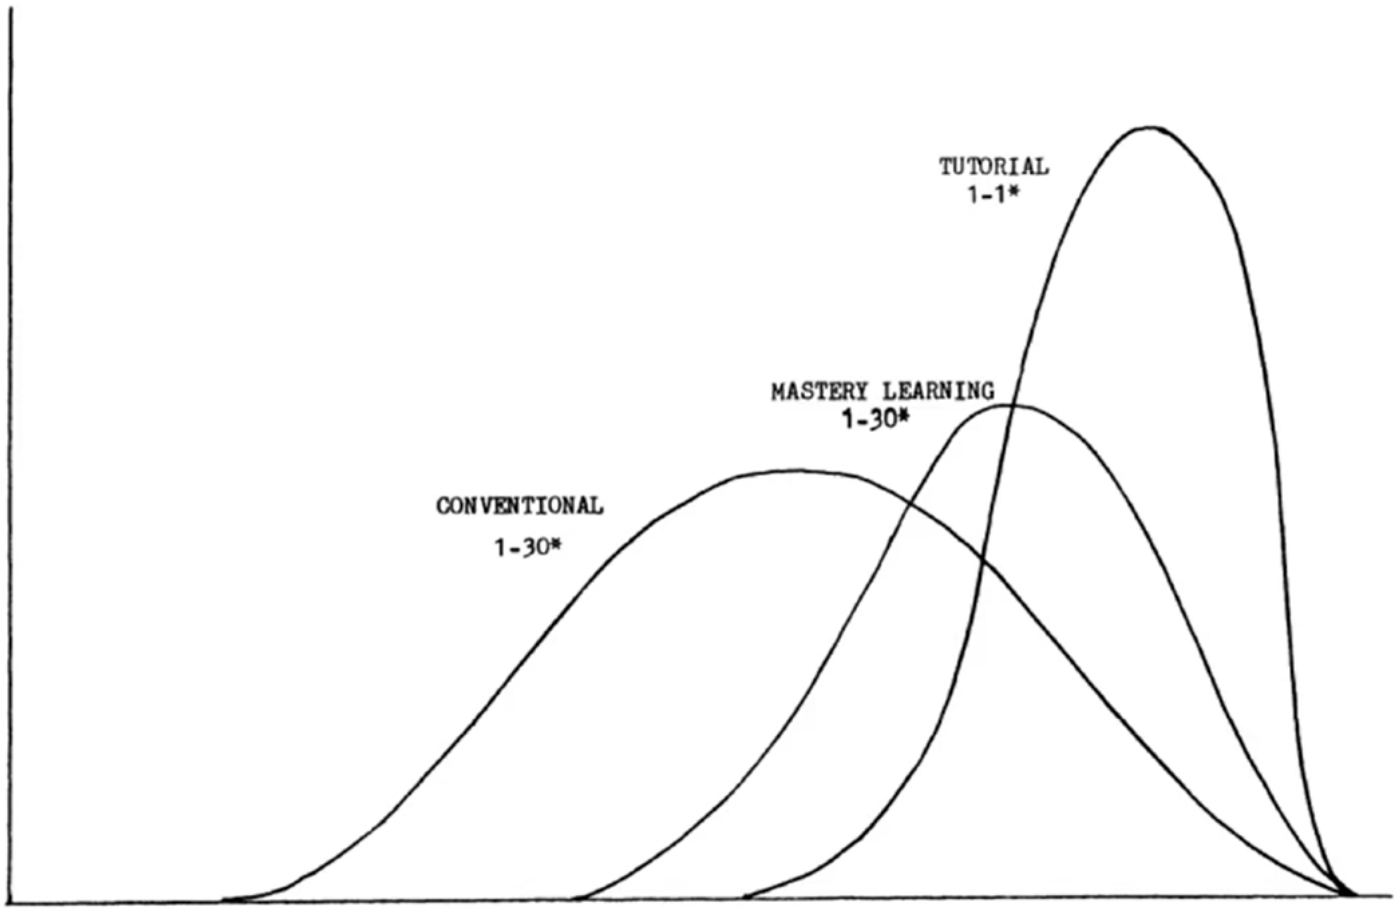

https://www.ted.com/talks/sal_khan_how_ai_could_save_not_destroy_education?language=en

Generative AI: A Government PoV

Ben Kaner

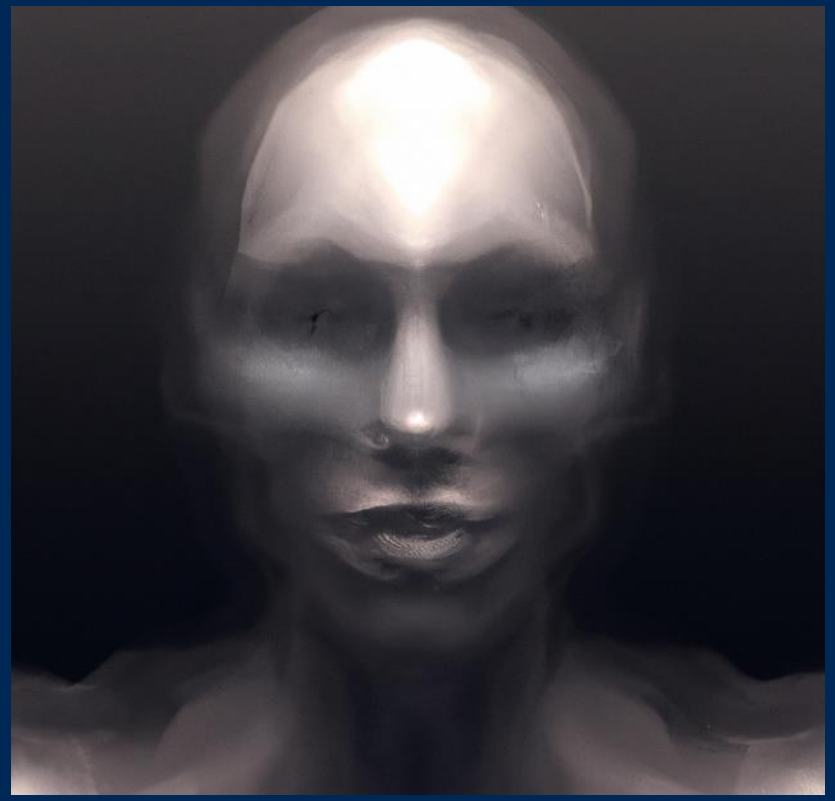

Gartner®

This is all about ChatGPT

AI – the Range

Capability

- ***Reactive***
  - Narrow single task, cannot learn ‘on the job
- ***Limited memory***
  - can predict a short distance into the future from current status, but limited in ability to adapt
    - Autonomous vehicles
    - Generative AI – they predict the ‘next word or phrase’ and adapts the responses based on previous information within a conversation (the prompt)
- ***Theory of mind***
  - Externally provided ‘model’ - a simulated human (in combination with previous techniques).
- ***Self Aware***
  - Artificial sentient.

Scope

- ***Narrow***
  - trained for a specific task.
- ***General***
  - learns and understands as well as a human
- ***SuperIntelligent***
  - Difficult to conceive of, and in the realms of SciFi (for now)

Simple view – uses of AI

- Predictive
  - Prediction of flow, modelling complex systems, simulations
- Discriminatory
  - Classification, soft-matching, risk assessments, anomalies
- Generative
  - Creates outputs that *look like* they came from an original dataset
  - May be text, voice, video or image

AI – the Power

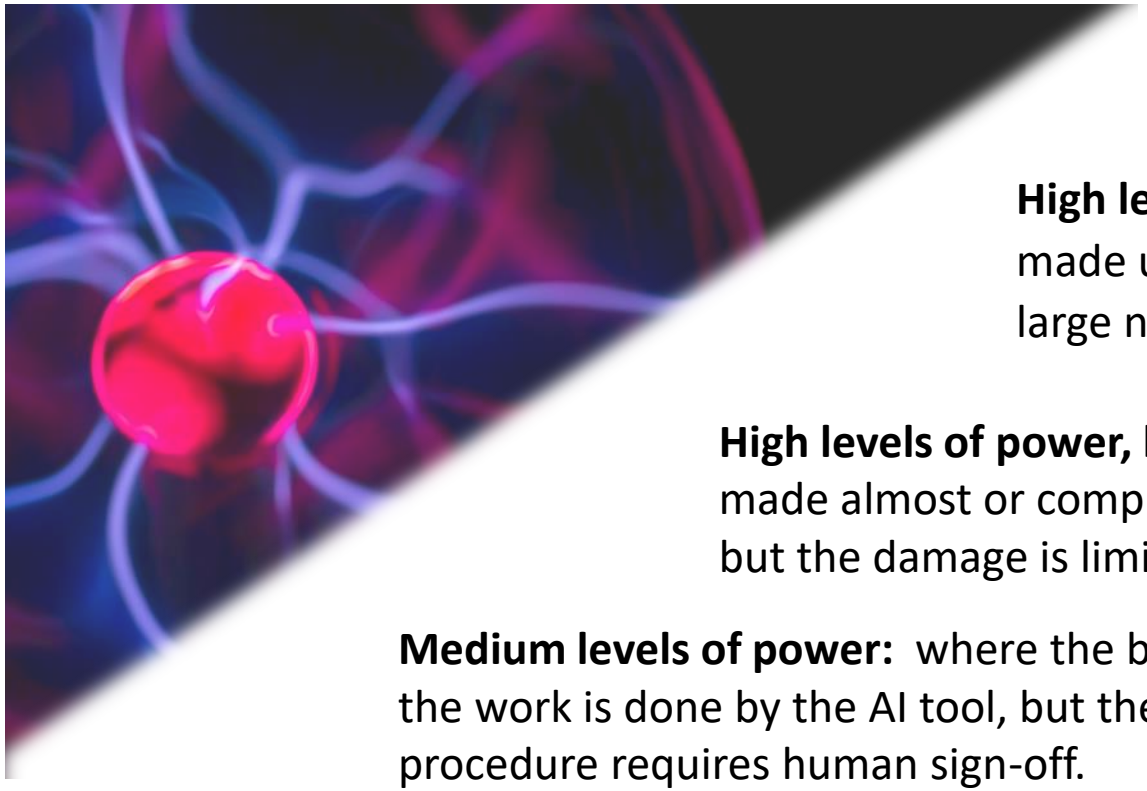

Unacceptable: Life changing decisions made automatically with insufficient human governance

High levels of power, high scale: decisions made using the same model but affecting large numbers of people.

Major scandal, loss of trust

High levels of power, low scale: decisions made almost or completely autonomously, but the damage is limited.

Autonomous driving, liability

Medium levels of power: where the bulk of the work is done by the AI tool, but the procedure requires human sign-off.

Human bias, habituation

Low levels of power: translation, proposal of initial wording – where there is *always* a human in the loop and the human unequivocally has the accountability and final say as to what is published.

Anchor bias, inadvertent bias

Gartner

Generative AI

Gartner

Generative AI – small set of examples

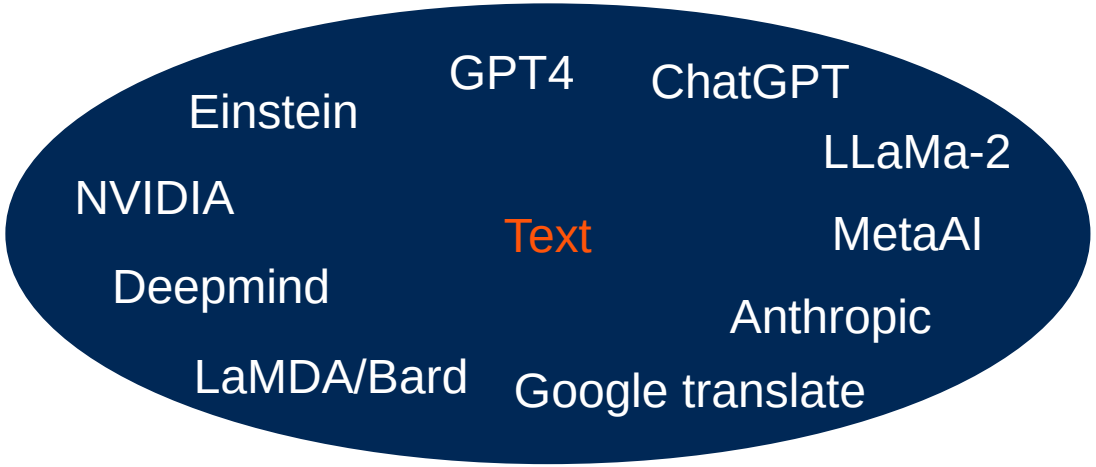

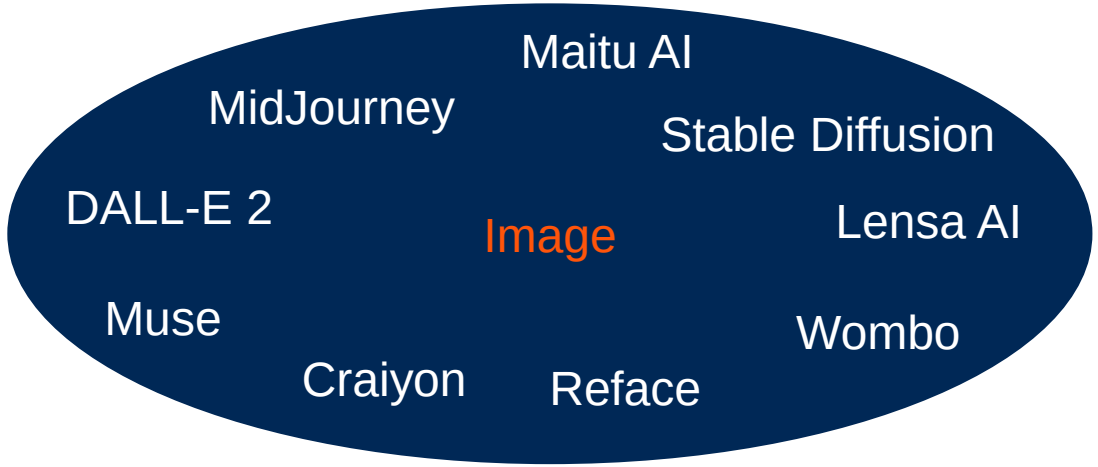

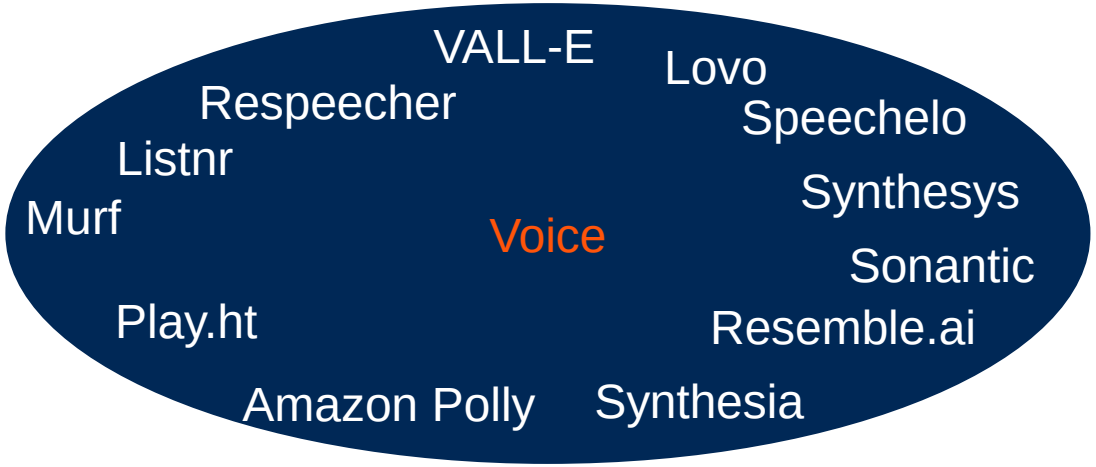

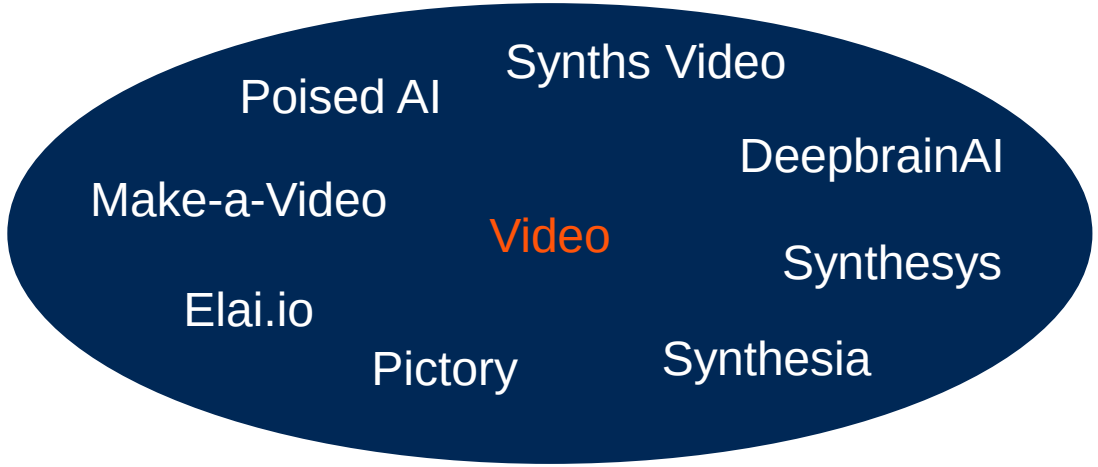

Gartner

What is a Generative AI Large Language Model?

- Neural network
- Many parameters
- Trained on large quantities of unlabeled text – or other data
- Data generally from public datasets – which introduces bias
- Self- or Semi-Supervised Learning

- ***Emergent*** behaviors with unexpected results

Gartner

What does a Large Language Model Do?

- Has read a lot
- Is very articulate
- Can speak multiple languages
- Can draft rapidly
- Has no real-world experience
- Is somewhat naïve
- May be able to produce interesting ideas
- Will learn from what they are given
- May or may not understand sensitivity and privacy

Gartner

…. And what might that remind you of?

- Has read a lot
- Is very articulate
- Can speak multiple languages
- Can draft rapidly
- Has no real-world experience
- Is somewhat naïve
- May be able to produce interesting ideas
- Will learn from what they are given
- May or may not understand sensitivity and privacy

Gartner

Use Cases in Government

Gartner

What can it be used for?

Use cases and techniques are proliferating rapidly

	Generating Text	Analysing Text	Manipulation and Translation	Question Answering
Manage Agency Mission	Reports, 1st draft	Sentiment Analysis – reviews of media Resolving Administrative Backlogs	Communications	Policy questions
Back Office	Process Guides, Reports	Recruitment Processes, Complaint handling	Reporting to multiple bodies	FOI requests Ramping on Staff
IT	Code Generation, User Guides	Discovery, functional analysis, code checking	Refactoring, legacy migration	User Support
Engaging Citizens	Guides, Communication	Sentiment Analysis, Empathic Communication	Targeting marginalized constituents	Help desks, Personalized Support and Information
Operations	Guides (internal and external), reporting, summarization	Triage, prioritization, classification, redaction	Multilingual support, outreach	Ramping on new staff, compliance, process adherence

What can it be used for?

Use cases and techniques are proliferating rapidly

	Generating Text	Analysing Text	Manipulation and Translation	Question Answering
Manage Agency Mission	Reports, 1st draft	Sentiment Analysis – reviews of media Resolving Administrative Backlogs	Communications	Policy questions
Back Office	Process Guides, Reports	Recruitment Processes, Complaint handling	Reporting to multiple bodies	FOI requests Ramping on Staff
IT	Code Generation, User Guides	Discovery, functional analysis, code checking	Refactoring, legacy migration	User Support
Engaging Citizens	Guides, Communication	Sentiment Analysis, Empathic Communication	Targeting marginalized constituents	Help desks, Personalized Support and Information
Operations	Guides (internal and external), reporting, summarization	Triage, prioritization, classification, redaction	Multilingual support, outreach	Ramping on new staff, compliance, process adherence

Specific Capabilities – Human Services *Protective Services*

Intake
- Inbound Referral Queues
- Triage
- Establish Case Information
- Update Case Information

Investigation Support
- Assemble Data
- Risk Assess Case Factors
- Workflow Management
- Evidence Management

Intervention Support
- Augmented Risk-Based Decisions
- Referral Orchestration
- Exception and Escalation

Case Monitoring
- Periodic Reporting
- Dynamic Case Monitoring
- Event Audit Logs
- Outbound Events

Placement Monitoring
- Allocating Placement
- Transition Handling
- Periodic Case Update

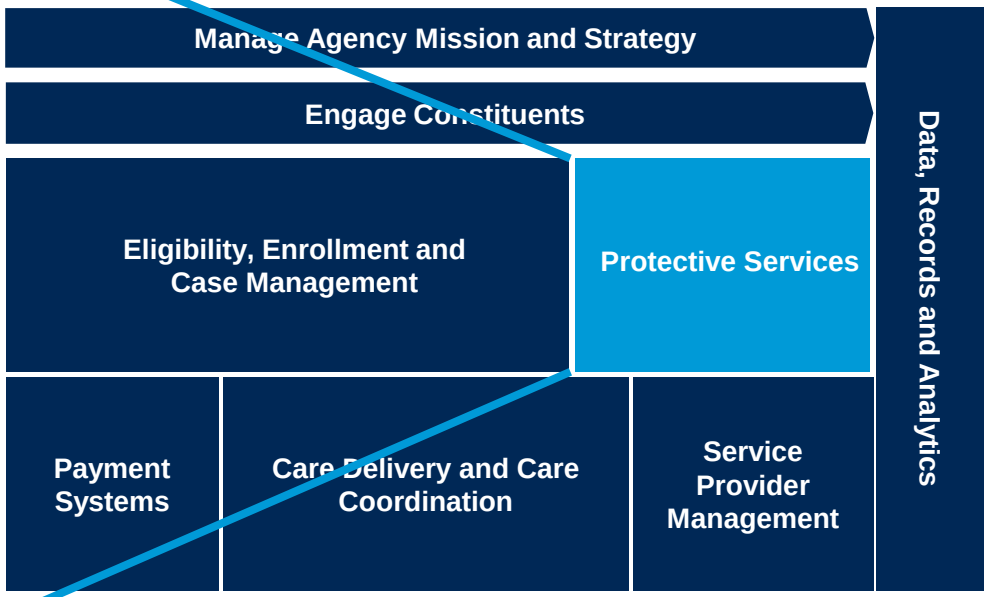

From 'A Business Capability Reference Model for Human Services' G00780325

Gartner

Specific Capabilities – Human Services *Protective Services*

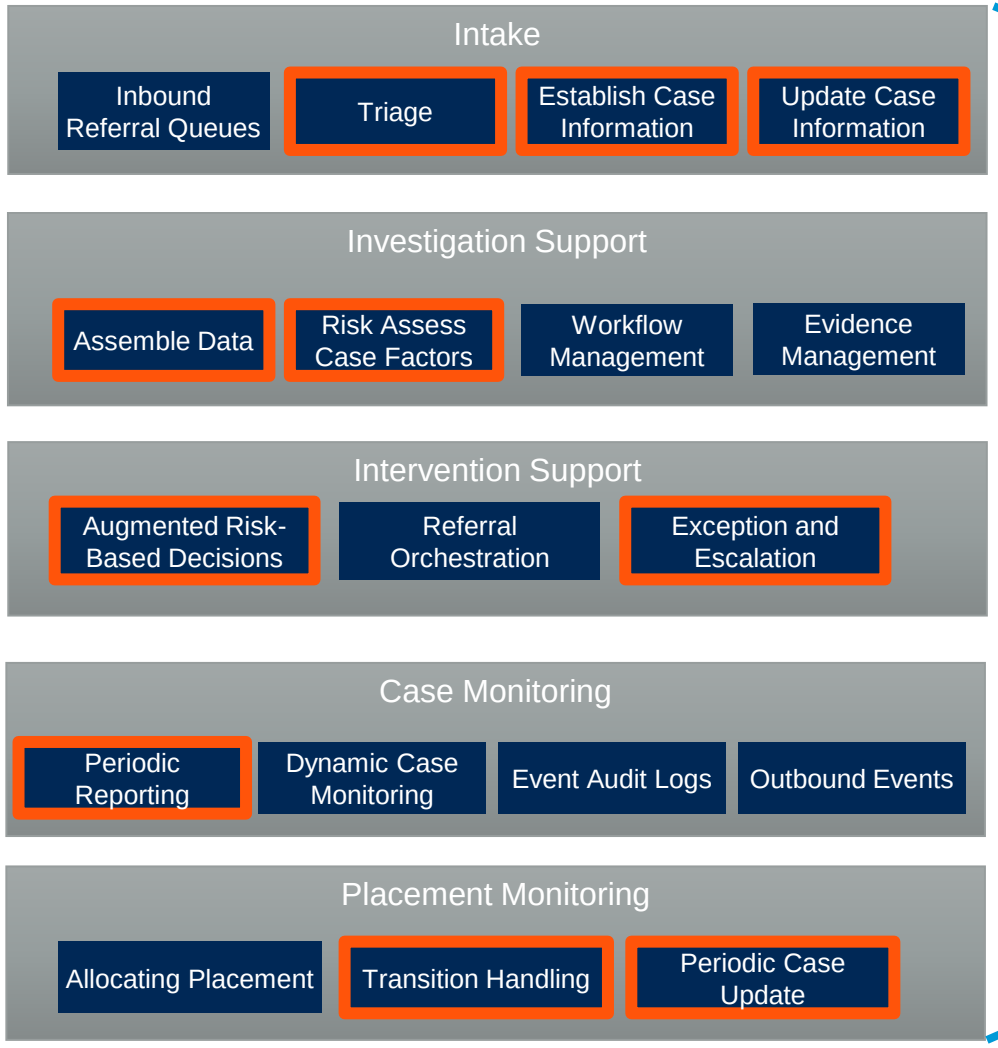

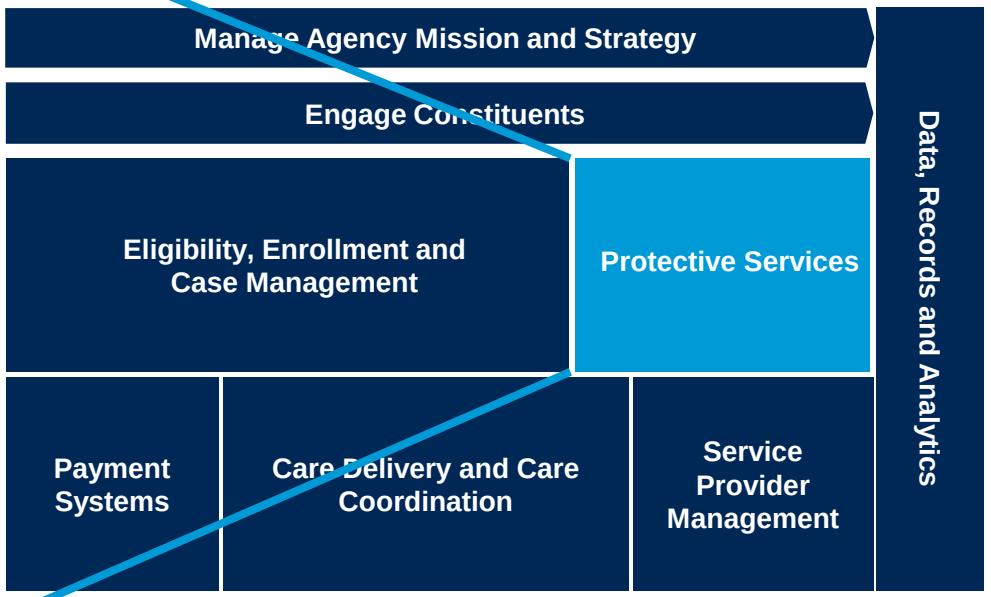

From 'A Business Capability Reference Model for Human Services' G00780325

Gartner

Prioritization

Gartner

For each proposed initiative:

- Score on :
  - Business Value
    - Policy/Societal Outcome
    - Operations Efficiencies
    - Risk to Service
  - Feasibility
    - Technology
    - Organizational (skills, governance, ability to adapt, operating model)
    - External (data availability, regulation, acceptability)

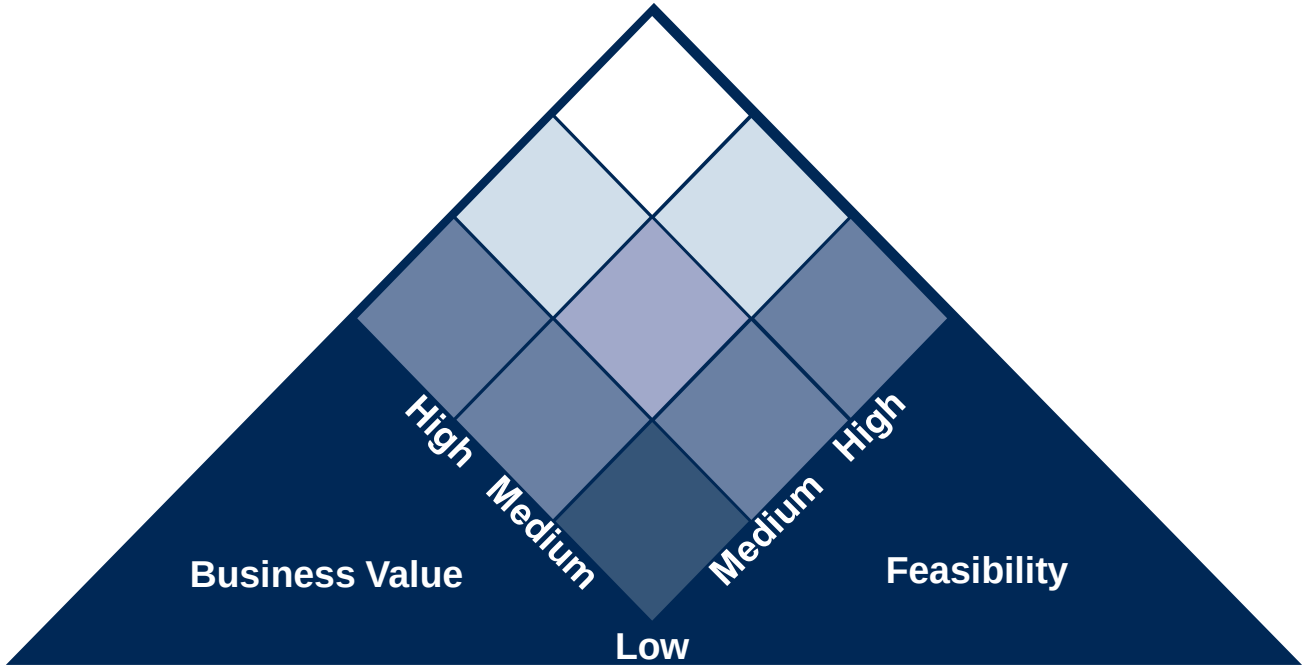

For each proposed initiative:

The Tech is only 1/6 of the challenge!

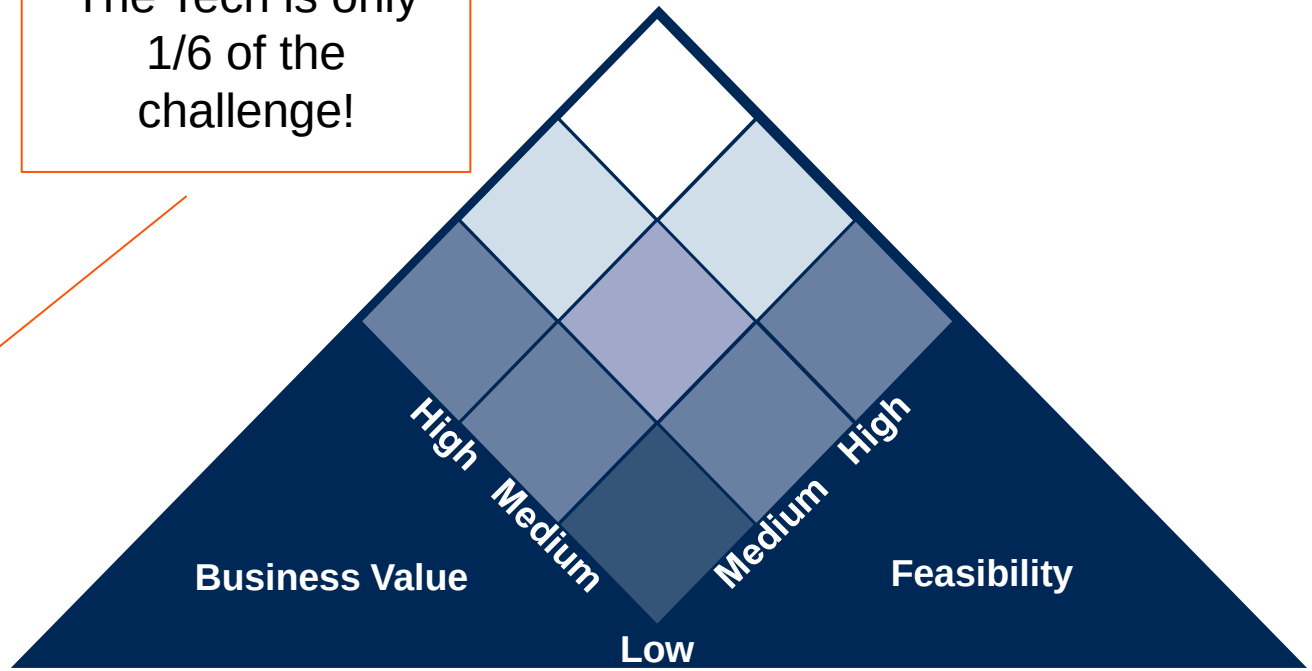

- Score on :
  - Business Value
    - Policy/Societal Outcome
    - Operations Efficiencies
    - Risk to Service
  - Feasibility
    - Technology
    - Organizational (skills, governance, ability to adapt, operating model)
    - External (data availability, regulation, acceptability)

Where things will end up

- Which are feasible?
- Which result in real return?
- Which look nice, but in reality....

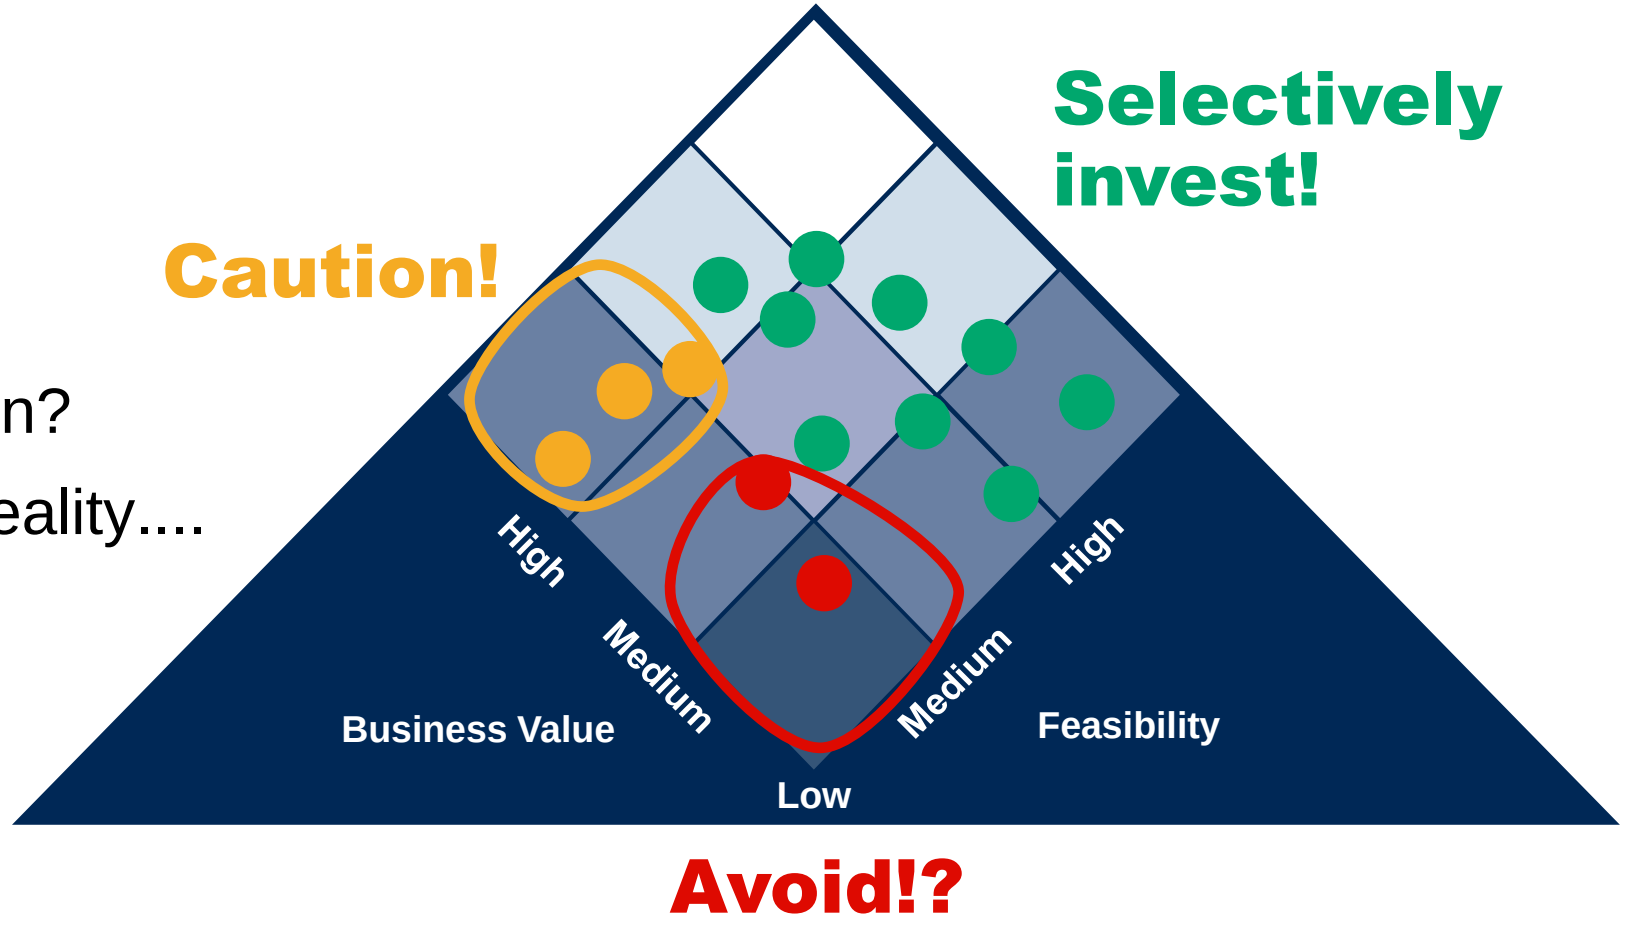

Recommendation: Prioritize Initiatives

- Appraise which *business capabilities* can benefit from which Generative AI capabilities
- Aggregate value
  - Value to mission
  - Reduction in operational resources
  - Reduction in risk of failure
- Ensure feasibility
  - Technical feasibility
  - Organizational Feasibility
  - Societal and Regulatory Acceptability

LLM Technology is immature and changing rapidly

The most common cause of failure in major change

AI is particularly sensitive to negative news

Challenges and Risks

Considerations

- Large Language Models learn from existing data
  - They will incorporate **biases** from their source data
  - The data is *not* live and may be **out-of-date**
  - They have no ‘conceptual model’ of the desired outcome
    - The data provided may be **seriously misleading**
  - Much of the material used may be copyrighted
    - The *products* of Generative AI cannot be copyrighted under US guidance
    - The implications on liability have not yet been legally tested – in **any jurisdiction**
- Unlike search results, A chat response usually gives ONE option
  - If that response is wrong or dangerous, **who is liable?**

Government are Not the Only Users

External users will not go direct to government sites

- **Malign Uses**
  - **Administrative Flooding**
    - Complaints, FOI
  - **DeepFake Attacks**
    - Social Engineering
    - Corruption of process, exposure of information, redirection of funds, provision of licences/permits
    - Impersonation of public figures
  - **Poisoning the Well**
    - Deliberate misinformation through the external models
  - **Code Generation**
    - Viral attacks generated at scale
- **Monitor developments and build plans**

Recommendation: Mitigate Risk

- Ensure that all staff understand the limits – Generative AI is not an infallible oracle
- Provide a safe space for experimentation
  - People are already experimenting!
  - Harvest good ideas
- Higher value/risk combinations should deliver in stages:
  - Start with ‘Human in the Loop’ – the *human is accountable*
  - Scale to automation *only* when the residual risk is low enough to be accepted by execs and legal
- **Prepare for new threats**

The Long View

The Long View: Is this a big change?

- Traditional engineering:
  - Precision and Repeatability = Quality
- Generative AI
  - Variability in outcome is a *feature* not a flaw. New outcomes continue to emerge
- **This is a behavioural system, not an engineered system**

- Who is making the decision?
  - For the ***first time*** humans may be subject to decisions made
    - By a non-human entity
    - Where the rules and reasons for the decision cannot be explicitly defined
  - The Future? Gen AI systems can invent code
    - Including the ability to invent a new AI… which means… evolution?
    - 'intent' can emerge and vary …. Which means… sentience?

Gartner

Recommendations

Gartner

Recommendations

1

Exploit the hype to drive policy

- Many are **excited**
- ChatGPT is **only one** large language model
- There are **many other uses of AI**
- This is an **opportunity**

2

Prioritize Initiatives

- **Value** to Purpose
  - Mission
  - Costs
  - Risk Reduction
- **Feasibility**
  - Technical
  - Organizational
  - AND Societal

3

Mitigate Risk

- Provide a **safe space** for experimentation
- Deliver in **stages**
- Start with human-in-the-loop
- **Scale** when residual risk is shown to be acceptable

4

Keep Abreast of Threats

- Involve **external parties** in discovery
- **Engage with providers** to control responses
- Build **mitigation plans**

Gartner

Questions?

Gartner