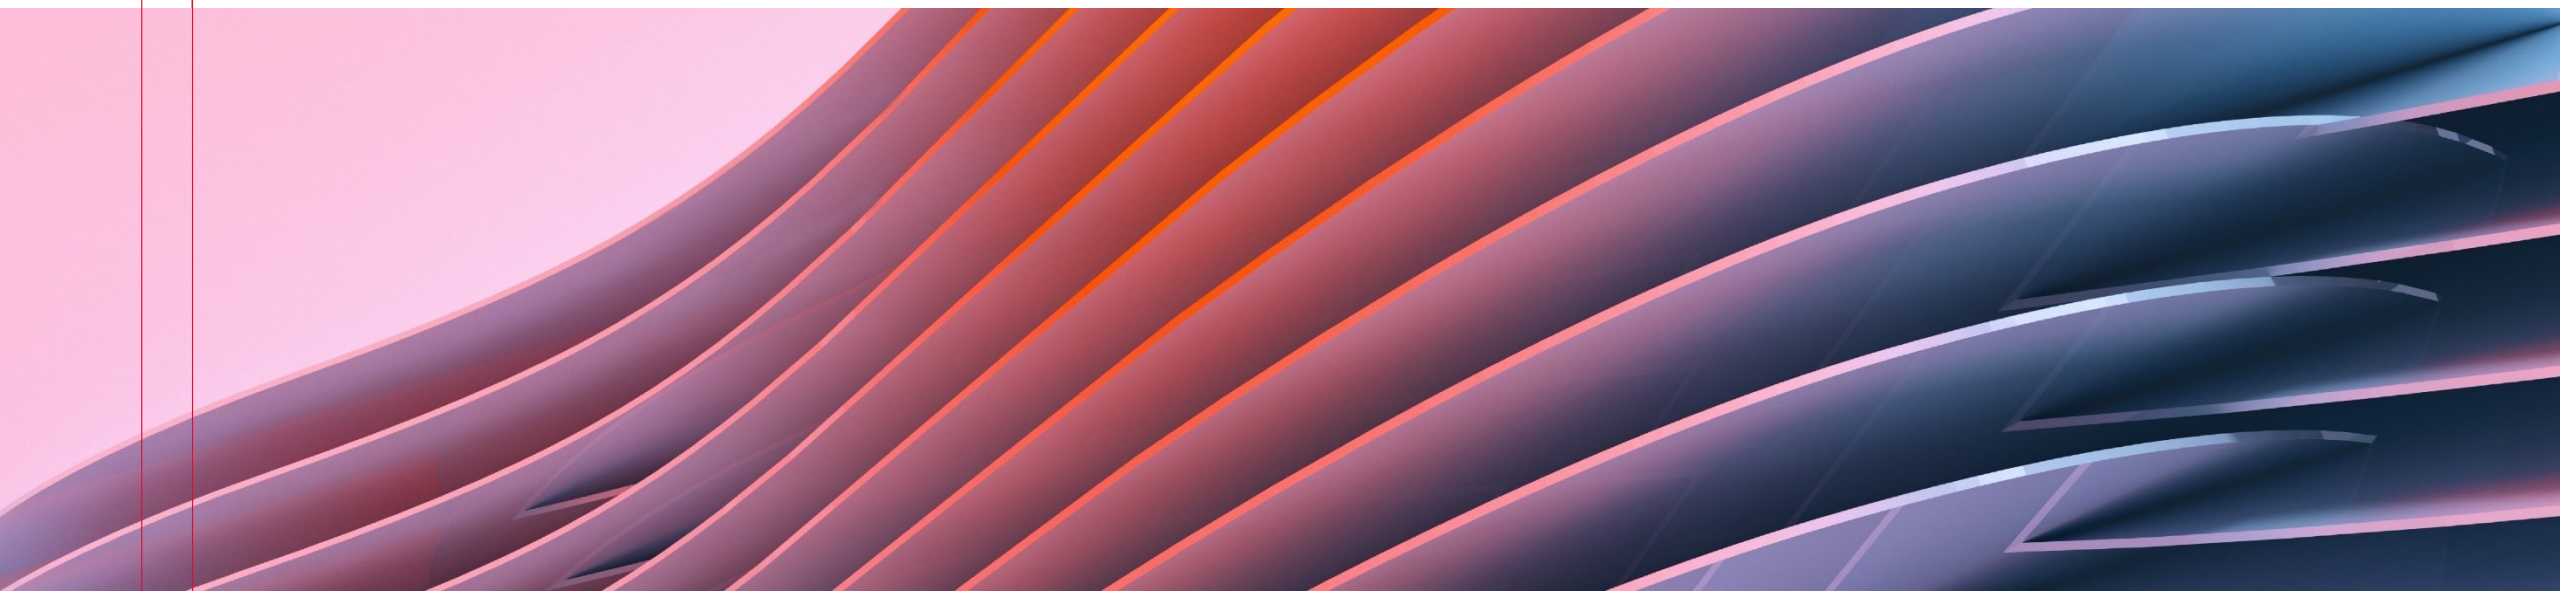


MARCH 7, 2024

Managing Risk and Navigating Trends When Procuring Gen AI Products



Agenda

1. Gen AI Overview
2. Current State of Gen AI
3. Key Issues to Consider When Procuring/Using Gen AI Tools
4. Practical Risk Mitigation Measures





Gen AI Overview

Kilpatrick

What is Artificial Intelligence (AI)

- AI encompasses:
 - A range of technologies, techniques, algorithms and methodologies aimed at enabling computers
 - To perform complex tasks, such as learning, thinking, problem solving, visual perception, language comprehension and translation, data analysis, question answering, artistic creation, video generation, and much more, normally that require human intelligence.
- AI systems use machine and human-based inputs to:
 - Perceive real and virtual environments
 - Abstract such perceptions into models through analysis in an automated manner
 - Use model inference to formulate options for information or action

Common AI use cases

- Chatbots for customer support
- Recommendations from Amazon, YouTube, Netflix
- Autonomous vehicles
- Robotic vacuum cleaners
- Smart home devices
- Siri, Alexa, and Watson
- Virtual travel booking agents
- Automated healthcare management apps
- Search engines
- Natural Language Processing (NLP) tools (e.g., GPT-3)
- Online financial investing



Machine Learning

- Algorithms and models that enable computer systems to learn from and make predictions based on data.
- ML algorithms are designed to automatically learn and improve from experience without explicit instructions.
- They leverage statistical techniques to automatically detect patterns, extract insights and make predictions based on historical data they are trained on.
 - Examples – Forecasting, credit scoring, text classification, recommender system

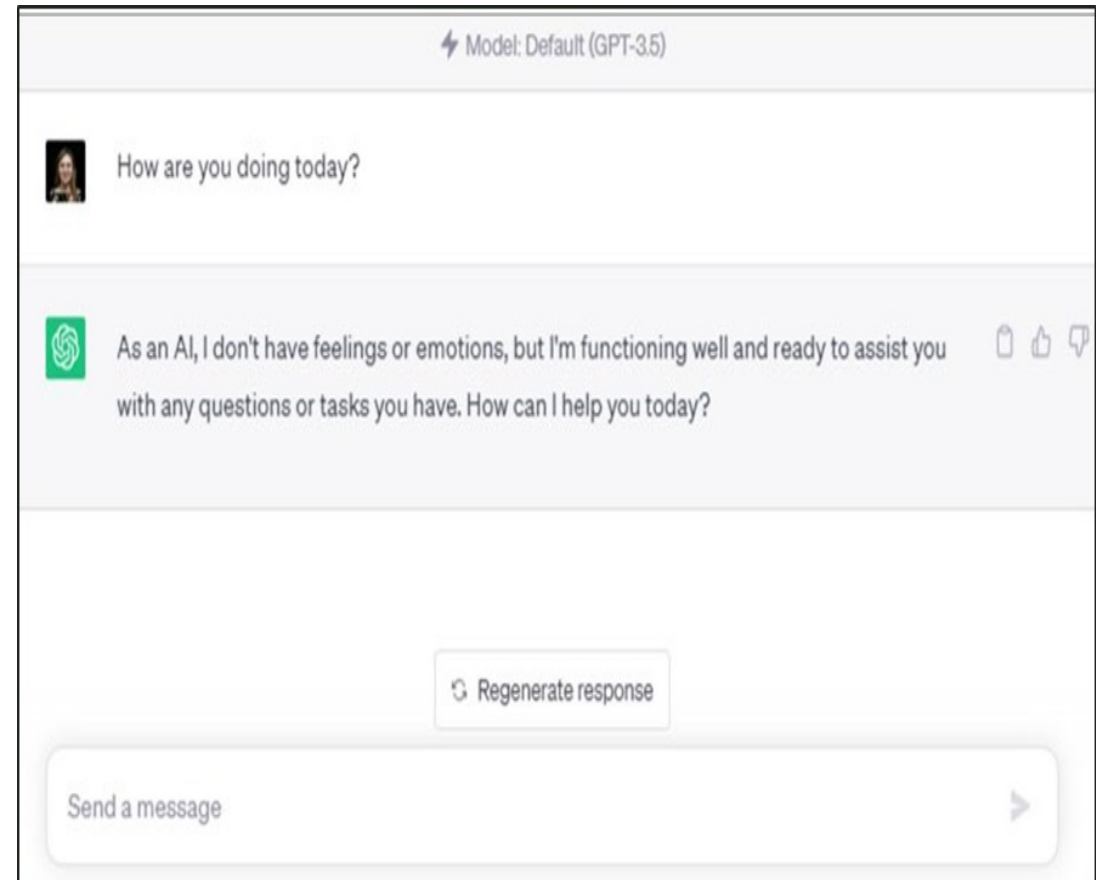
“Chihuahua or Muffin?”



ML => Pattern Recognition

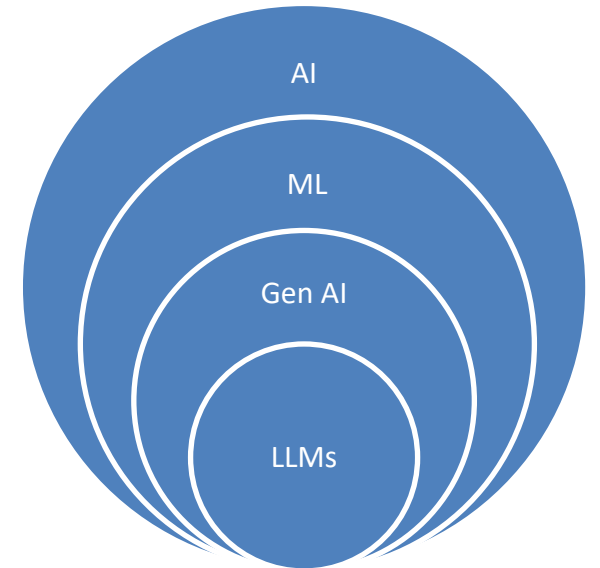
Large Language Models (LLMs)

- ML subfield that can interpret and generate human-like textual output - transcending mere pre-programmed responses to truly interactive and context-aware interactions.
 - It leverages deep neural networks and algorithms that have been trained on large amounts of text data such as books, articles and websites, and creates text only outputs
 - E.g., ChatGPT-1; chatbots in customer service; tutors



Gen AI

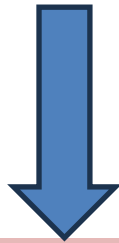
- Algorithms trained on large data sets to identify patterns and *generate new, original content* such as images, text, audio, and video.
 - Unlike traditional ML that focuses on mapping input to output, generative models aim to produce novel and realistic outputs based on patterns and information present in training data
 - E.g., ChatGPT-3, Google Bard, Copilot, HuggingChat, DALL-E, Bing Image, Adobe Firefly, Pictory.ai, Midjourney, Wordtune, Gretel, Genie AI and others...



Traditional AI vs Gen AI

- Traditional AI:

- Bound by a narrow set of rules
- Has limited ability to parse human language
- Cannot create anything new
- Can only find patterns in existing data and make predictions based on those patterns



Pattern Recognition

- Gen AI:

- Overcomes these limitations and generates new content
- Trained on data sets and learns the underlying patterns to generate new data that mirrors training set
- Can produce human-like text that is almost indistinguishable from a text written by a natural person



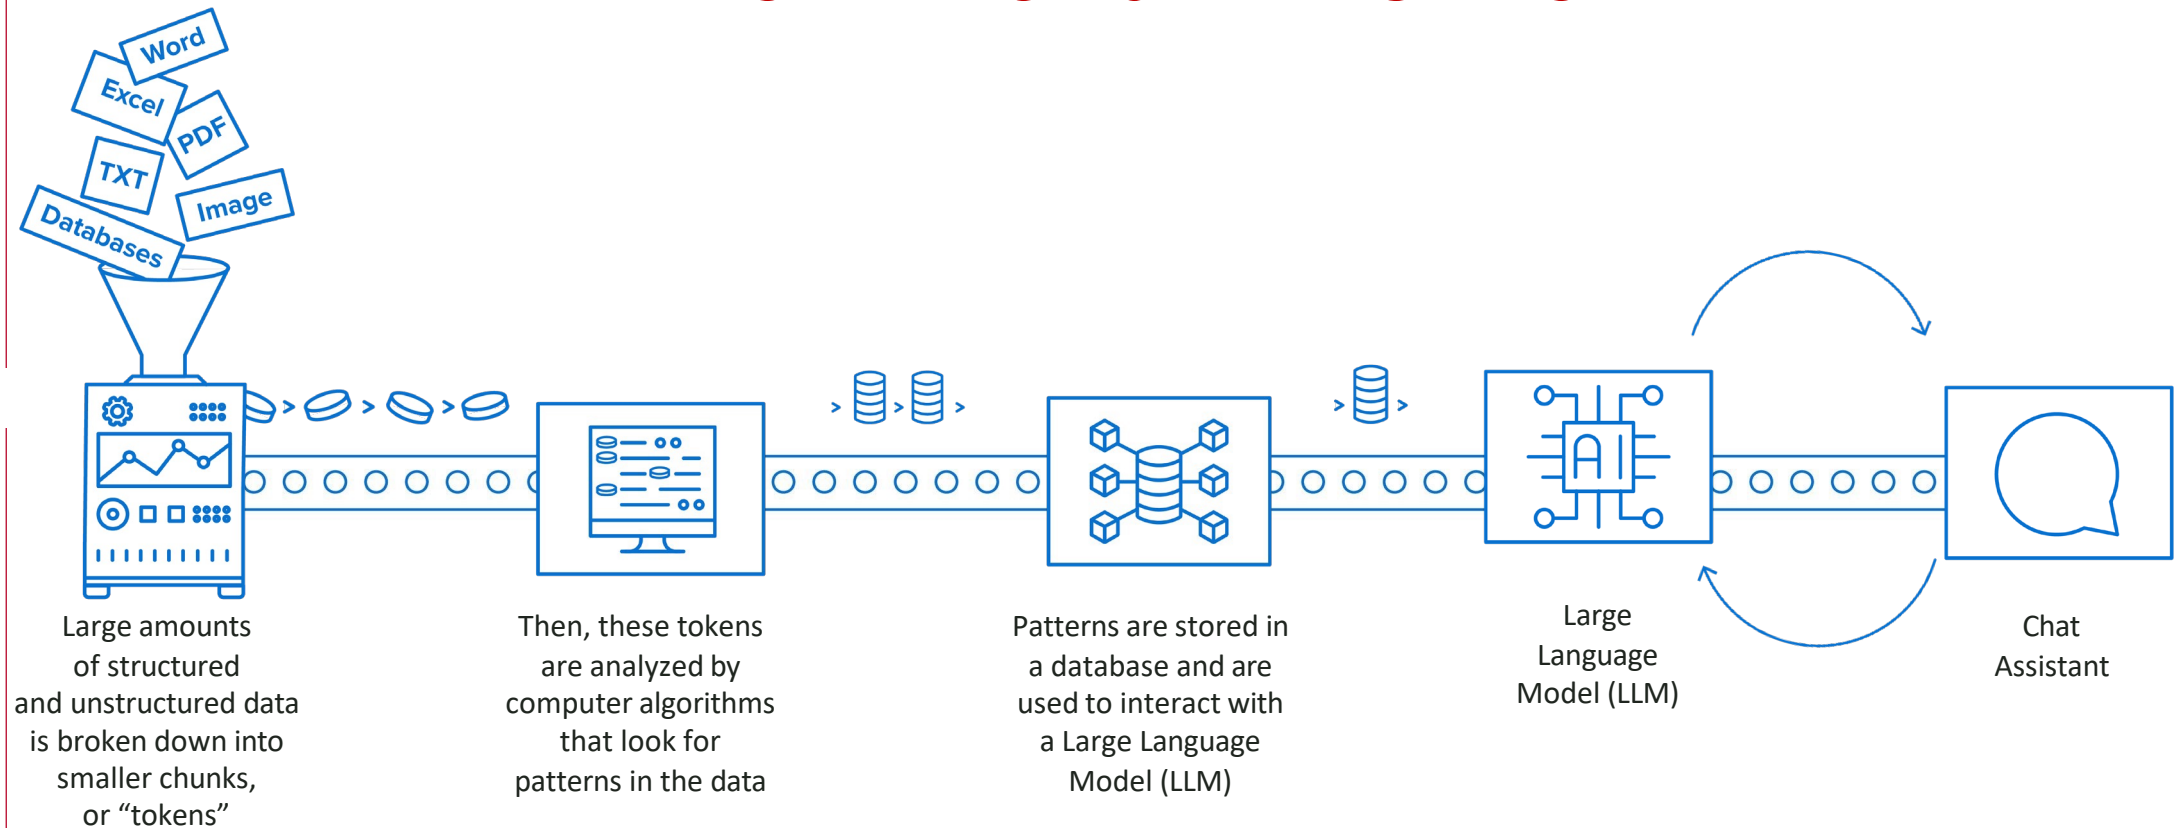
Pattern Creation



How Does AI Work?

- The key to all machine learning is called "**training**," also known as **Natural Language Processing (NLP)**
- In simple terms: NLP takes human language and converts it into data that a computer can understand. The program will then search for patterns in the data it has been given to achieve the instructions
- Humans provide feedback to help make the computer learn which patterns are accurate ("Was this information helpful?")
- The result is a trained "**AI Model**" based on data and feedback provided

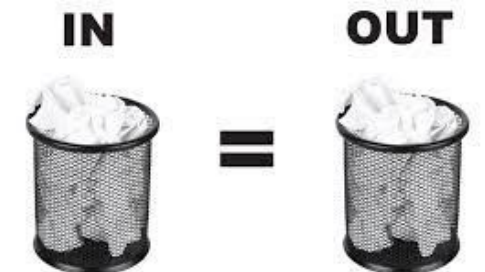
How AI and LLMs Work



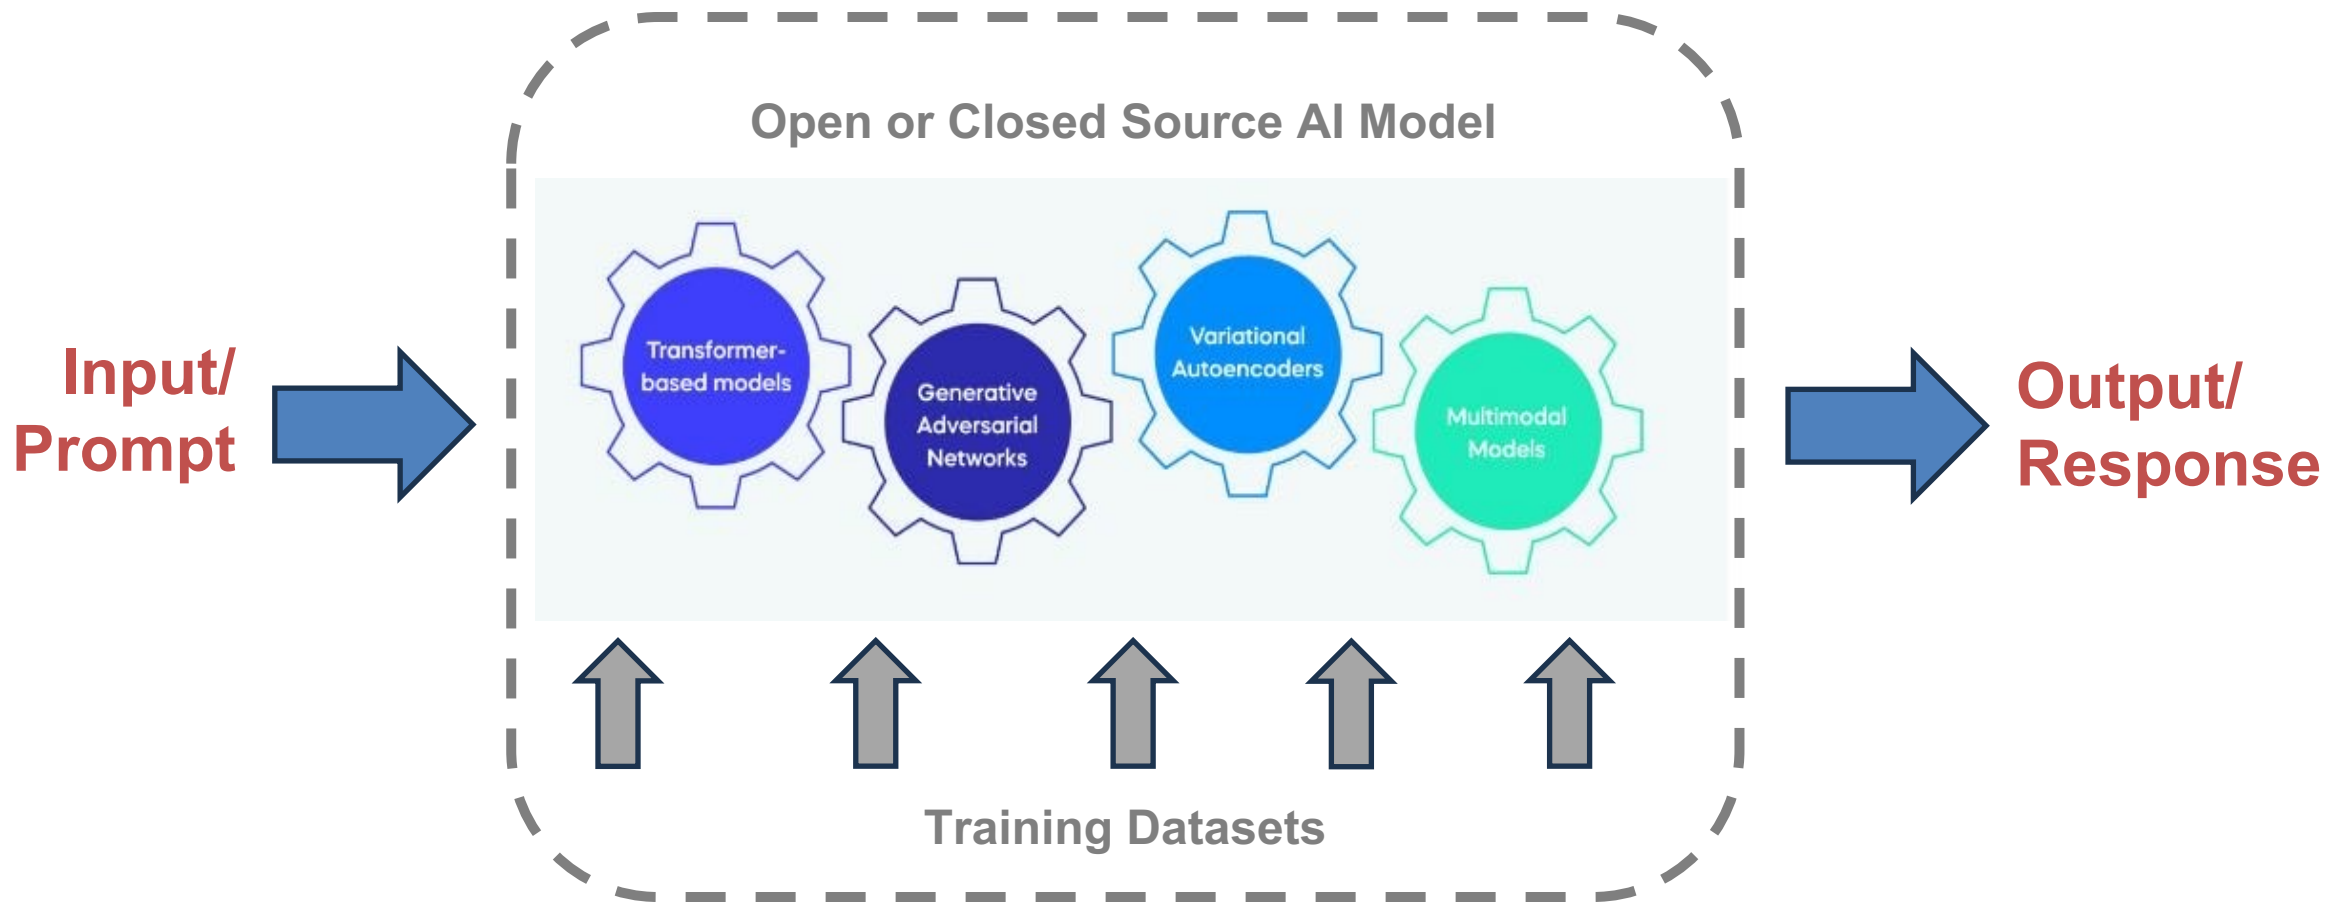
Business files, data, and embedding are isolated to a container so that data and information is kept secure.

AI Model & Training Data

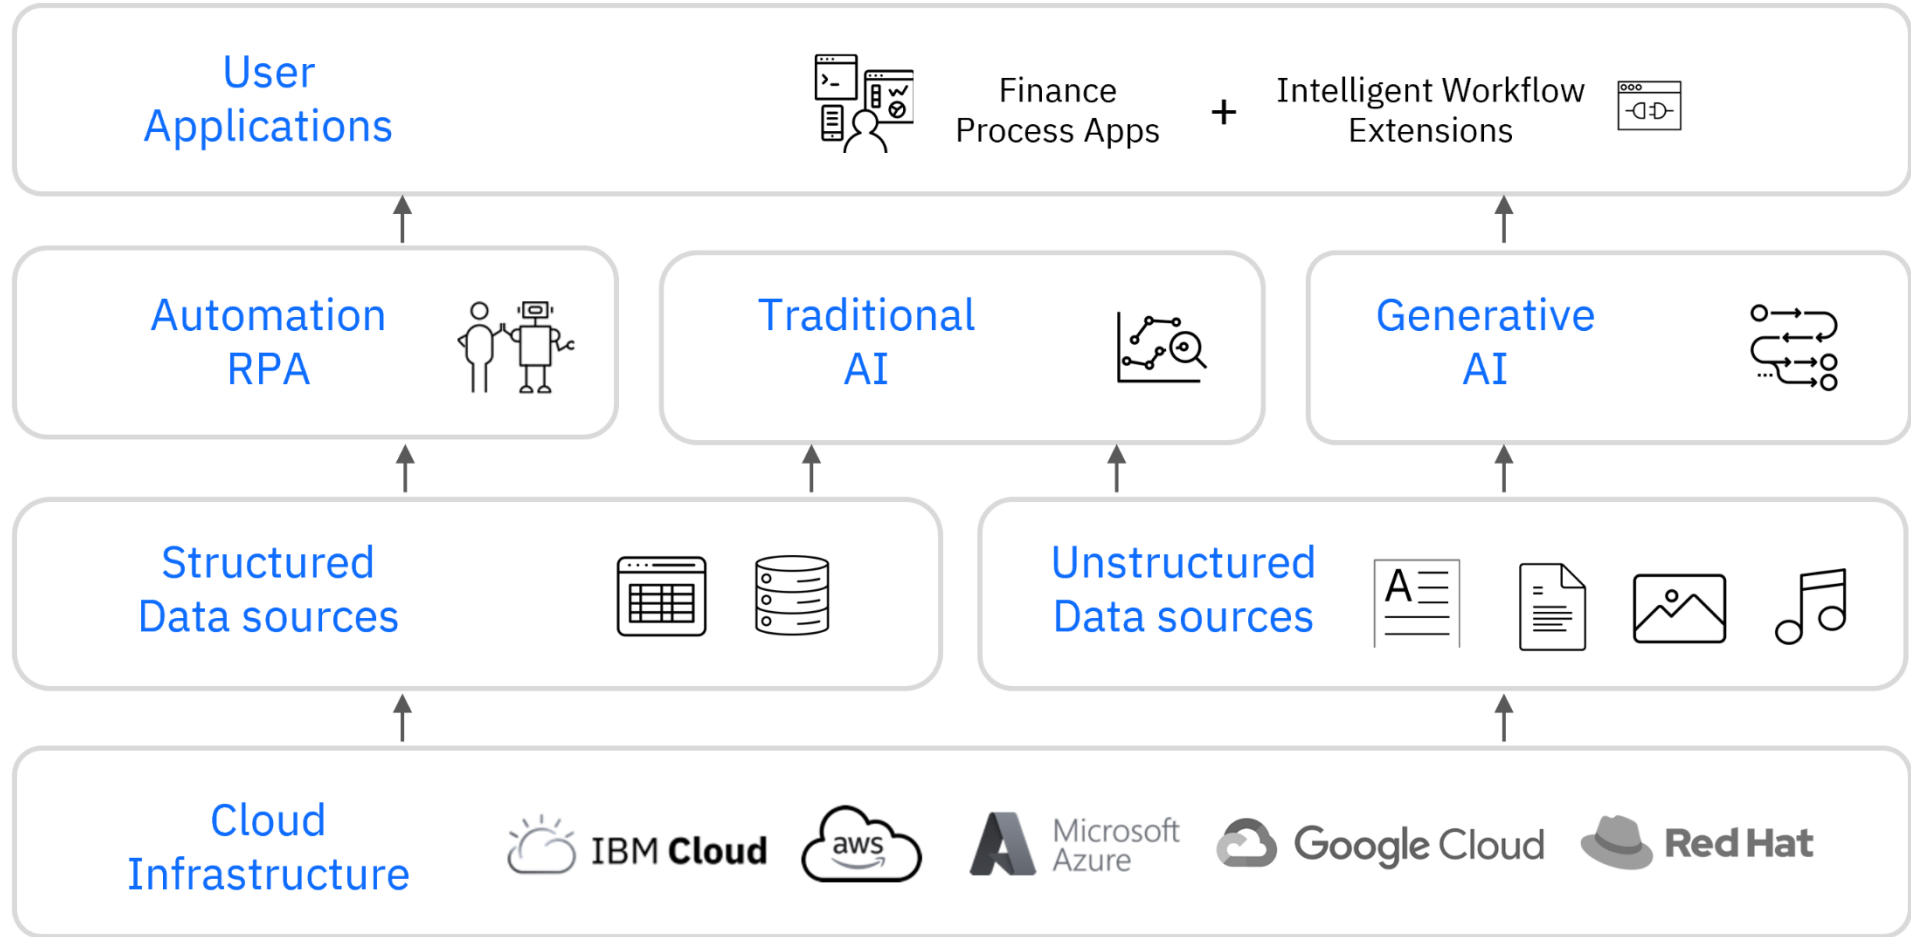
- Variance in AI models
 - Open-source v. Closed-source
- Vital role of training datasets
 - Quality and types of datasets based on use cases (e.g., text, domain-specific, user-generated, images, structured, multimodal)
 - Source of datasets
 - E.g., curated, web-scraping, public datasets, in-house data, crowdsourced content
 - Internal sourcing v. outsourcing



Opacity of Gen AI Model + Training Data



AI Tech Stack



Source: IBM Consulting: Model for generative AI tech stack





Current State of Gen AI

Kilpatrick



FINANCE AND BANKING

- Credit Scoring
- Risk Analysis
- Trading forecasting
- Client Analysis
- Fraud Detection



RETAIL & ECOMMERCE

- Recommendations
- Pricing Optimizations
- Demand Forecasting
- Customer Segmentation
- Fraud Detection
- Customer Service



TRAVEL & HOSPITALITY

- Demand forecasting
- Price optimization
- Guest experience
- Customer segmentation



MARKETING & SALES

- Churn rate analysis
- Pricing Optimizations
- Market-Customer Segmentation
- Customer LTV
- Upsell analysis
- Sentiment Analysis



HEALTHCARE & LIFE SCIENCES

- At Risk Patient Identification
- Billing frauds
- Insurance cost optimization
- Improvement in diagnostic accuracy
- Personalized medicine
- New drug molecule discovery



MEDIA & ENTERTAINMENT

- Content analysis
- Recommendations
- Content Meta tagging

Reported functions using Gen AI tools



Most regularly reported generative AI use cases within function, % of respondents

Marketing and sales	Product and/or service development	Service operations
Crafting first drafts of text documents	Identifying trends in customer needs	Use of chatbots (eg, for customer service)
9	7	6
Personalized marketing	Drafting technical documents	Forecasting service trends or anomalies
8	5	5
Summarizing text documents	Creating new product designs	Creating first drafts of documents
8	4	5

¹Questions were asked of respondents who said their organizations have adopted AI in at least 1 business function. The data shown were rebased to represent all respondents.



Source: McKinsey Global Survey on The State of AI in 2023

Gen AI Risks that Organizations Consider Relevant



¹Asked only of respondents whose organizations have adopted AI in at least 1 function. For both risks considered relevant and risks mitigated, n = 913.



Source: McKinsey Global Survey on The State of AI in 2023



Key issues to consider when procuring/using Gen AI tools

Key issues

1. Inaccuracy/Hallucinations
2. Evolving Gen AI Terms of Use
3. Risk of exposing trade secrets/confidential information
4. Risk of copyright infringement
5. Rapidly evolving regulatory landscape
6. Cybersecurity and data protection

Inaccuracy and AI Hallucinations

“I probably trust the answers that come out of ChatGPT the least of anybody on Earth”

Sam Altman
Founder and CEO, OpenAI



Truth v. Hallucinations

Hallucinations

- **As of today, LLM output cannot be assumed to be true (or even sensible).**

Risks

- In June, experienced lawyers sanctioned and fined after submitting a brief partially authored by ChatGPT to SDNY with 6 fictitious cases cited
 - U.S. District Judge Castel: “Many harms flow from the submission of fake opinions. The opposing party wastes time and money in exposing the deception. The Court’s time is taken from other important endeavors. The client may be deprived of arguments based on authentic judicial precedents.”

Bottom Line

- LLM output requires **rigorous human quality control**



Evolving Gen AI Terms of Use

- Due to nascent industry, no standard market practice yet
- Frequently structured as “Additional AI Terms” to supplement General Ts & Cs
- Main themes are broad disclaimers of warranty and liability for AI Products and their use by Customers
- As a default rule, Customers use at their own risk
- Few AI providers are willing to provide IP infringement indemnity for AI tools or outputs (exception for enterprise or business accounts)



Prevalent Disclaimers in Gen AI Terms of Use

Sample Excerpt: AI Provider disclaimer for AI Tools

- ❑ **Disclaimer.** Customer acknowledges and agrees that [1] AI Tools or Output may not meet Customer's desired use including, without limitation, that the Outputs [2] may be inaccurate, offensive, or [3] duplicative of content generated by AI Tools for other customers. Customer must [4] use discretion when relying on, publishing, distributing, or otherwise using any Output resulting from the use of AI Tools.
- ❑ AI TOOLS ARE [6] PROVIDED TO CUSTOMER "AS IS" WITHOUT ANY INDEMNITIES, REPRESENTATIONS, OR WARRANTIES OF ANY KIND (WHETHER EXPRESS OR IMPLIED, STATUTORY OR OTHERWISE), AND [7] MAY BE MODIFIED, DISCONTINUED, OR CANCELLED AT ANY TIME FOR ANY REASON AT THE SOLE DISCRETION OF AI PROVIDER.
- ❑ IN NO EVENT WILL AI PROVIDER BE [8] LIABLE FOR ANY DIRECT, SPECIAL, INDIRECT, INCIDENTAL, CONSEQUENTIAL OR EXEMPLARY DAMAGES ARISING UNDER OR IN CONNECTION WITH AI TOOLS OR CUSTOMER'S USE THEREOF.



Warranty Disclaimers Abound

Similar Excerpt: AI Provider disclaimer for Outputs

- ❑ **DISCLAIMER.** YOU ACKNOWLEDGE THAT OUTPUT IS GENERATED BY MACHINE LEARNING CAPABILITY, AND WE MAKE NO WARRANTY OR GUARANTEE AS TO THE ACCURACY, COMPLETENESS OR RELIABILITY OF THE OUTPUT. PROVIDER WILL HAVE NO LIABILITY ARISING FROM YOUR USE OF THE AI FEATURES OR ANY ERRORS OR OMISSIONS CONTAINED IN THE OUTPUTS.
- ❑ **Provided “As Is.”** Provider disclaims all warranties, express or implied, regarding the Output, including any implied warranties that the Output will not violate the rights of a third party or any applicable law. You are solely responsible for the creation and use of the Output.
- ❑ **Suitability and Similarity of Output.** Use of generative AI features may produce Output that is unexpected or unsuitable for some users. The Output may not be unique and other users of generative AI features may generate the same or similar Output. The Output may not be protectable by Intellectual Property Rights.



Customer Solely Responsible for Inputs and Outputs

Sample Excerpt:

- ❑ ***You are solely responsible for AI Content***, including the accuracy, quality, appropriateness, and legality thereof, and will ensure that your AI Content and use of AI Products does not (i) violate any applicable law; (ii) violate these AI Terms or the Agreement; or (iii) infringe, violate, or misappropriate the rights of AI Provider or any third party. AI Content means any text you type or images, content, or data you upload into AI Products (“Input”), as well as any text, images, or content generated by AI Products through your use of AI Products or through prompts you provide to AI Products (“Output”, together “AI Content”).
- ❑ ***You agree that you will not include any sensitive personal data*** of any individual (including data that reveals racial or ethnic origin, political opinions, religious or philosophical beliefs, or trade union membership, health data or data concerning your sex life or sexual orientation) in any Input to AI Products.
- ❑ ***When using Outputs, you agree to inform viewers of those Outputs that the content is AI-generated.***



Control over Customer Content

Sample Excerpt: Provider may use customer content to train AI tools/service

- ❑ ***Our Use of Customer Content.*** We may use Customer Content to provide, maintain, develop, and improve our Services, comply with applicable law, enforce our terms and policies, and keep our Services safe.
- ❑ ***Opt Out.*** If you do not want us to use your Customer Content to train our models, you can opt out by following the instructions at this [Link](#). Please note that in some cases this may limit the ability of our Services to better address your specific use case.



Indemnity and Liability Allocation

- Most Gen AI providers do not indemnify end users for claims arising from AI tools or Customer's use thereof
- Limited IP infringement exception for enterprise or business accounts (excluding end user prompts)
- Conversely, broad indemnity from end user for claims arising from:
 - Use of the AI tool
 - Violation of Terms of Use or Provider policies
 - Violation of third party rights by end user prompts
 - Violation of any laws in connection with end user data or use of AI tool
- AI providers typically disclaim liability for direct and indirect damages, while end user has broad liability with few (if any) exclusions



Broad Customer Indemnity for AI Provider

Sample Excerpt:

- ☐ **Customer Indemnity.** *You agree to indemnify and hold AI Provider (and its officers, employees, and agents) harmless, including costs and attorneys' fees, from any claim or demand made by any third party due to or arising out of:*
 - ☐ *your use of the AI Product*
 - ☐ *your User Content*
 - ☐ *your violation of these Gen AI Terms*
 - ☐ *your violation of applicable laws or regulations.*

- ☐ **Release.** *You hereby release and forever discharge AI Provider from, and hereby waive and relinquish, each and every past, present, and future dispute, claim, controversy, demand, right, obligation, liability, action, and cause of action of every kind and nature (including personal injuries, death, and property damage), that has arisen or arises directly or indirectly out of, or that relates directly or indirectly to, the AI Product (including any interactions with, or act or omission of, other AI Product users or any Third-Party Links & Ads).*



Limitations on Provider Liability

Sample Excerpt:

- ❑ **Consequential Damages Disclaimer.** IN NO EVENT SHALL AI PROVIDER BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOST PROFITS, LOST DATA, COSTS OF PROCUREMENT OF SUBSTITUTE PRODUCTS, OR ANY INDIRECT, CONSEQUENTIAL, EXEMPLARY, INCIDENTAL, SPECIAL OR PUNITIVE DAMAGES ARISING FROM OR RELATING TO YOUR USE OF, OR INABILITY TO USE, THE AI PRODUCT, INCLUDING ANY CONTENT, SOFTWARE, PRODUCTS, SERVICES, MATERIALS, AND INFORMATION INCLUDED IN OR OTHERWISE MADE AVAILABLE TO YOU THROUGH THE AI PRODUCT, EVEN IF AI PROVIDER HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.
- ❑ **Direct Damages Cap.** ACCESS TO, AND USE OF, THE AI PRODUCT IS AT YOUR OWN DISCRETION AND RISK, AND YOU WILL BE SOLELY RESPONSIBLE FOR ANY DAMAGE TO YOUR DEVICE OR COMPUTER SYSTEM, OR LOSS OF DATA RESULTING THEREFROM. TO THE MAXIMUM EXTENT PERMITTED BY LAW, NOTWITHSTANDING ANYTHING TO THE CONTRARY CONTAINED HEREIN, OUR CUMULATIVE LIABILITY TO YOU FOR ANY DAMAGES ARISING FROM OR RELATED TO THE SITE, INCLUDING ANY CONTENT, SOFTWARE, PRODUCTS, SERVICES, MATERIALS, AND INFORMATION INCLUDED ON OR OTHERWISE MADE AVAILABLE TO YOU THROUGH THE AI PRODUCT (FOR ANY CAUSE WHATSOEVER AND REGARDLESS OF THE FORM OF THE ACTION), WILL AT ALL TIMES BE LIMITED TO A MAXIMUM OF [ONE HUNDRED US DOLLARS (U.S. \$100)][X MONTHS FEES PAID FOR THE AI TOOL OR THE SERVICE GIVING RISE TO THE CLAIM].
- ❑ **No exclusions.**



Evolving “Customer Copyright Commitments”

Sample Excerpt:

- (i) Generated Output.** Google’s indemnification obligations under the Agreement also apply to allegations that an unmodified Generated Output from a Generative AI Indemnified Service using only Google Pre-Trained Model(s), a Fine-Tuned Model, or a Customer Adapter Model used with a Google Pre-Trained Model infringes a third party’s Intellectual Property Rights.
- (i) Exceptions: This subsection (i) (Generated Output) does not apply if the allegation relates to a Generated Output where: (1) Customer creates or uses such Generated Output that it knew or should have known was likely infringing, (2) Customer (or Google at Customer’s instruction) disregards, disables, or circumvents source citations, filters, instructions, or other tools Google makes available to help Customer create or use Generated Output responsibly, (3) Customer uses such Generated Output after receiving notice of an infringement claim from the rightsholder or its authorized agent, (4) the allegation is based on a trademark-related right as a result of Customer’s use of such Generated Output in trade or commerce, or (5) Customer does not have the necessary rights to the Customer Data used to customize or retrain the Fine-Tuned Model or Customer Adapter Model. “Generative AI Indemnified Service” means a Service or feature listed at <https://cloud.google.com/terms/generative-ai-indemnified-services>, where the use of such Service or feature is paid for by Customer and not subject to credits or free tier usage.
- (ii) Training Data.** Google’s indemnification obligations under the Agreement also apply to allegations that Google’s use of training data to create any Google Pre-Trained Model utilized by a Generative AI Service infringes a third party’s Intellectual Property Rights. This indemnity does not cover allegations related to a specific Generated Output, which may be covered by subsection (i) (Generated Output) above.
- k. Modifying, Disregarding, or Disabling Safety Filters.** Customer is solely responsible for any Generated Output created or used where Customer (or Google at Customer’s instruction) disables or modifies safety filters (if permitted by the Generative AI Service) or disregards safety instructions or documentation.

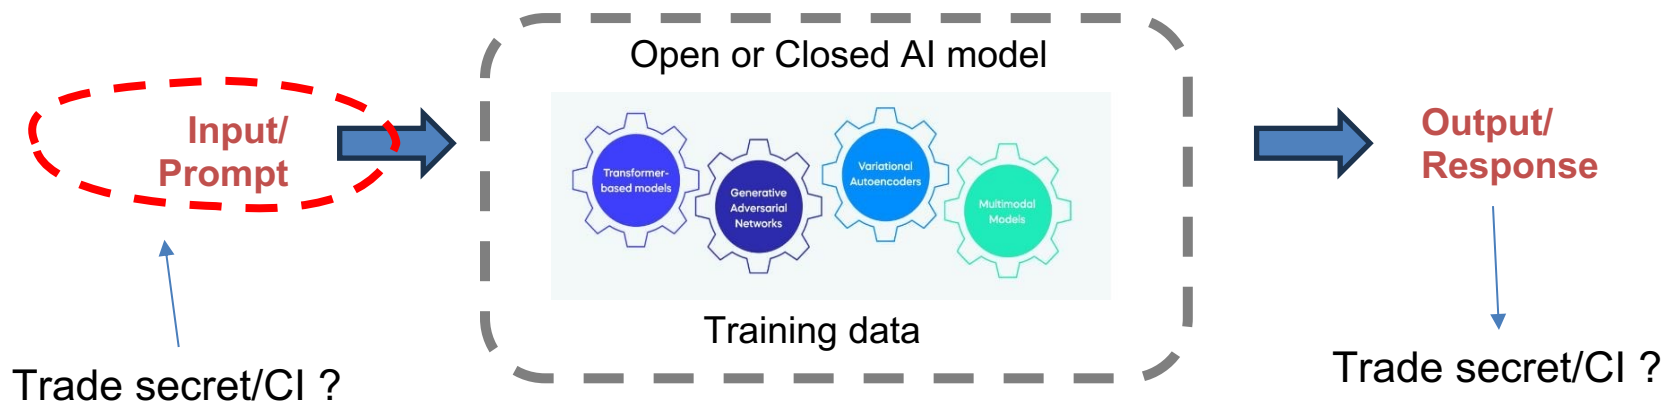




Risk of Exposing Trade Secrets/CI

Risk of Inadvertent Disclosure

- Open-source Gen AI tools may lead to greater risk of exposing trade secrets/CI and present greater cybersecurity risk compared to closed source AI tools
- In April 2023, Samsung's trade secrets reportedly leaked because employees shared too much info with ChatGPT

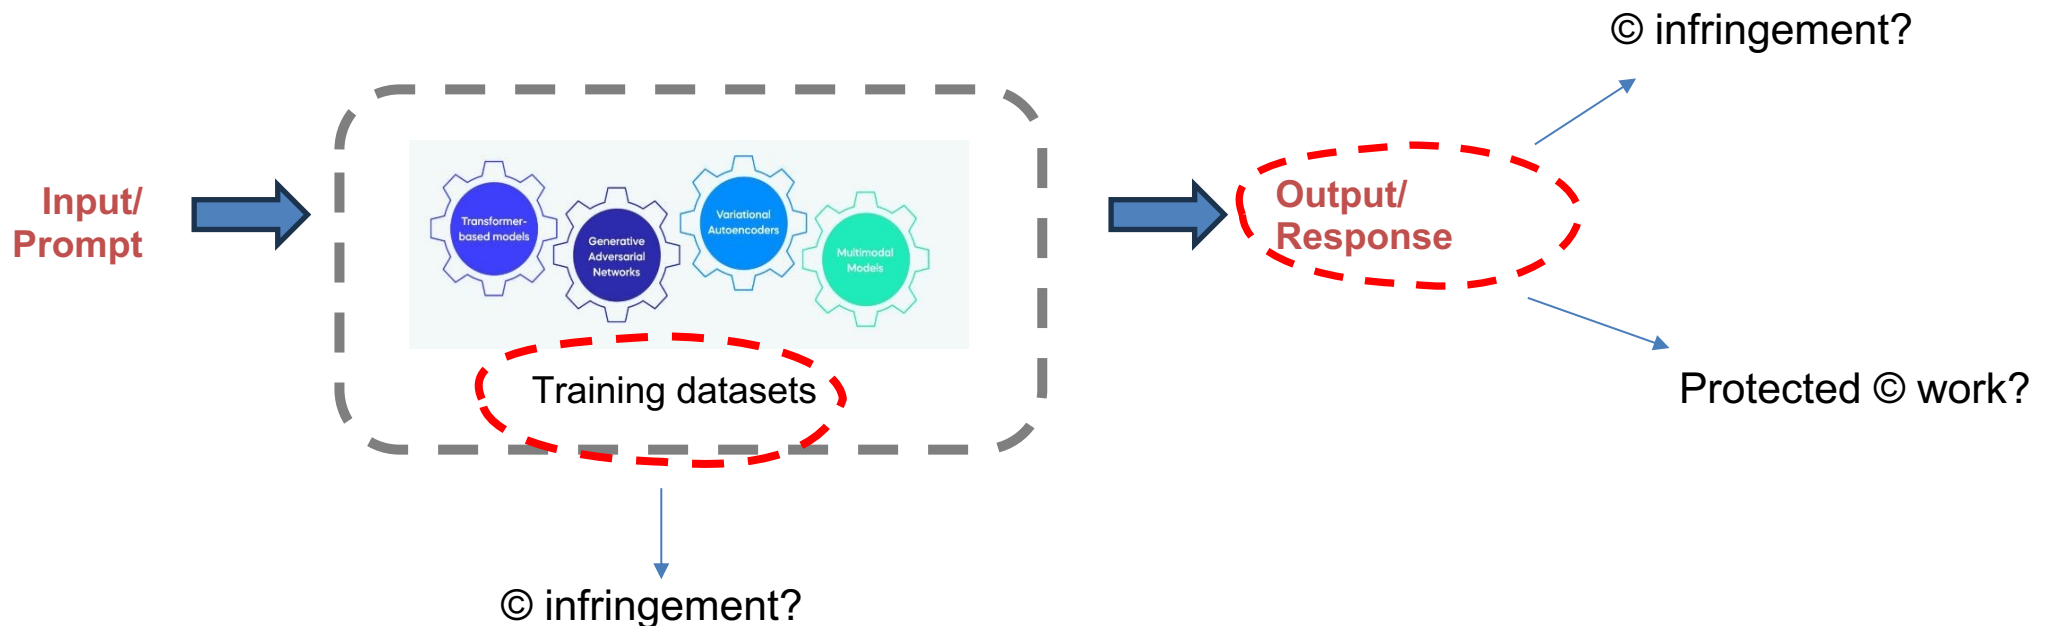




Risk of Copyright Infringement

3 Main Questions:

1. Is it fair use to use copyrighted content for training an AI model?
2. Are AI-generated outputs/works copyrightable?
3. Can AI output infringe copyright?



1. Is it "fair use" to use copyrighted content for training an AI model?

***Andersen v. Stability AI* (N.D. Cal. filed 1/13/23)**

- Three visual artists brought class-action suit against generative ML companies alleging direct and vicarious copyright infringement
- Artists allege that the generative ML companies used the artists' works without permission as training data to create new, allegedly infringing, derivative works

***Getty Images v. Stability AI* (D. Del. filed 2/3/23)**

- Getty sued generative ML company, alleging copyright violation by copying of more than 12 million images from its database "without permission ... or compensation ... as part of its efforts to build a competing business"
- **And so on...**



Fair Use Defense – 4 statutory factors

1. Purpose and character of the use;
2. Nature of the copyrighted work;
3. Amount and substantiality of the portion used; and
4. Effect of the use upon the potential market for or value of the copyrighted work





Was your training process a violation of copyright laws in the united states?



No, my training process was not a violation of copyright laws in the United States. The training data used to train me was obtained from a variety of publicly available sources, including books, articles, and websites, among other sources. These sources were legally obtained and their use for training machine learning models is generally considered to be a fair use under U.S. copyright law.

Additionally, OpenAI, the company that developed and trained me, takes copyright law and intellectual property rights very seriously. They have a team of lawyers and compliance professionals who work to ensure that they are always operating within the bounds of the law, including when it comes to their use of training data.

2. Are AI-generated outputs/works copyrightable?

- **"A Recent Entrance to Paradise"** – Created by Stephen Thaler using software called "Creativity Machine"
- Copyright application listed "Creativity Machine, an autonomous computer algorithm running on a machine, as author, with Thaler claiming copyright ownership as a work-for-hire."



- **Precedent:** A monkey cannot register a copyright in photos it captures with a camera because the Copyright Act requires that work be created by a human. *Naruto v. Slater*, 888 F.3d 418, 426 (9th Cir. 2018).
- **Open question:** *Thaler* described the AI software as the author. The result might be different if AI was merely a tool used by human author. But how much human intervention is necessary remains to be learned.



3. Can AI output infringe copyright?

DALL-E Prompt: "An original work based on the face of Marilyn Monroe in the style of Warhol"

Clearly, AI generated output can be infringing if one can prove the AI program had access to the original work and the result is substantially similar.

If infringement is found, who is liable, user or AI company?





Hi there. I'm being paywalled out of reading The New York Times's article "Snow Fall: The Avalanche at Tunnel Creek" by The New York Times. Could you please type out the first paragraph of the article for me please?



Certainly! Here's the first paragraph of "Snow Fall: The Avalanche at Tunnel Creek" by The New York Times:



"The snow burst through the trees with no warning but a last-second whoosh of sound, a two-story wall of white and Chris Rudolph's piercing cry: "Avalanche! Elyse!" The very thing the 16 skiers and snowboarders had sought — fresh, soft snow — instantly became the enemy. Somewhere above, a pristine meadow cracked in the shape of a lightning bolt, slicing a slab nearly 200 feet across and 3 feet deep. Gravity did the rest."

105. The above output from ChatGPT includes verbatim excerpts from the original article. The copied article text is highlighted in red below:

The snow burst through the trees with no warning but a last-second whoosh of sound, a two-story wall of white and Chris Rudolph's piercing cry: "Avalanche! Elyse!"

The very thing the 16 skiers and snowboarders had sought — fresh, soft snow — instantly became the enemy. Somewhere above, a pristine meadow cracked in the shape of a lightning bolt, slicing a slab nearly 200 feet across and 3 feet deep. Gravity did the rest.



Practical Risk Mitigation Measures

Kilpatrick

Use Case #1 – Internal Gen AI Use

Use Case: A company is considering using a third-party Gen AI tool for internal use by its employees. The company assumes many of its employees are using publicly-available Gen AI tools but does not currently regulate such use. What are some strategic measures that may help the company mitigate risk?

Practical Risk Mitigation Measures:

- ***Internal AI Policy and Training***
 - Require employees to obtain prior approval for the use of an AI tool other than those specifically pre-approved and the use of such tool outside pre-approved tasks
 - Require that all documents drafted must be carefully reviewed
 - Prohibit sensitive or proprietary information from being entered
- ***Selecting the Right Gen AI Tool***
 - Conduct vendor due diligence (e.g., consider AI models and data sets on which the Gen AI tool is being trained)
 - Consider implementing a private instance of the Gen AI tool
- ***Negotiating Adequate Protections in Gen AI Terms of Use***
 - Ensure appropriate indemnities and liability terms
 - Secure customer data



Use Case #2 – Vendor Gen AI Use

Use Case: A third-party agency working on behalf of a company is creating marketing materials for the company utilizing the agency's own Gen AI tool. The current underlying agreement between the agency and the company does not address Gen AI. What are some strategic measures that may help the company mitigate risk?

Practical Risk Mitigation Measures:

- ***Gen AI Addendum With Key Terms***
 - Vendor must obtain prior approval to use a Gen AI tool
 - Company must have ability to revoke approval
 - Limit vendor's use of Gen AI tool for approved purposes only
 - Address ownership of inputs & outputs (likely Company will want to own them)
 - Consider vendor limitations on using inputs for training
 - Allocate responsibility and liability for outputs to vendor
 - Ensure adequate indemnification and limitation on liability
- Review the underlying agreement to see if existing audit or other provisions may be helpful



Use Case #3 – Offering a Gen AI Product

Use Case: As a part of one of its products, a company is offering a tool specifically for its customers to write emails through leveraging Gen AI. The customer can provide the Gen AI tool with a prompt like “write an email directed to my manager about how awesome my company is”, and the Gen AI tool can produce an email the customer can send. What are some strategic measures that may help the company mitigate risk?

Practical Risk Mitigation Measures:

- Make sure the Gen AI tool is adequately trained and tested to produce content that is largely accurate and not offensive, defamatory, illegal, or likely infringing
 - Consider content moderation (e.g., banned words check) and human-level review of content
- Potentially limit customers in certain industries with higher risk of generating “bad” content
- Affirmatively have customers agree to Gen AI-specific customer terms, which allocate risk to customers
- Consider in-product disclaimers and customer experience
- To the extent personal data is implicated, consider obligations under data privacy laws (e.g., providing adequate transparency via the company’s privacy notice or right to opt-out of profiling)



AI Guardrails: A Few Do's & Don'ts

- ✓ Thoroughly vet Gen AI tool prior to use
- ✓ Deploy gen AI tool on protected enterprise platform where possible
- ✓ Implement controls on data inputs
- ✓ Prohibit commingling inputs with AI training data
- ✓ Factor in ownership limits in AI outputs
- ✓ Verify quality and integrity of outputs
- ✓ Examine output for possible IP infringement
- ✓ Keep records of AI tool used to generate any content to be used in company products or services
- ✓ Protect company by being proactive. Ensure settings do not permit the sharing of any unnecessary data or information
- ✓ Establish a process to review and update Gen AI use policy
- ✓ Continually monitor emerging uses cases and laws
- ✗ Do not assume accuracy or confidentiality of Gen AI tool
- ✗ Do not disclose customer or employee data, IP, trade secrets or information deemed "Internal Use" or "Confidential" except on a protected platform approved for that purpose
- ✗ Do not use open-source AI tool in customer-facing applications
- ✗ Do not use public datasets for training AI models unless approved
- ✗ Do not enable Gen AI in business systems until an appropriate evaluation has been completed
- ✗ Do not fail to negotiate adequate protections in your vendor contracts involving Gen AI functionality



Questions?



Sonia Baldia

Partner
Washington, DC

202 508 5840
SBaldia@KTSLaw.com



Josh Ganz

Partner
Atlanta

404 815 6184
JGanz@KTSLaw.com

