Sterility Assurance & Quality Risk Management Conference

PDA®
Parenteral Drug Association
Midwest Chapter

QUALITY ASSURED

Leveraging AI for Quality Risk Management in Pharmaceuticals and Life Sciences

Sterility Assurance & Quality Risk Management Conference

October
8th & 9th

QUALITY ASSURED

Agenda:

- LLM Background and Key Terms
- Regulatory Expectations
- Model Development
- Data Requirements
- Practical Applications and Benefits

Chris Dayton
Co-Founder/CEO
Quality Assured AI
10+ Years GMP/QS Experience

QualityAssured.AI

QUALITY ASSURED

LLM Literacy

What is Artificial Intelligence?

QUALITY ASSURED

- **Artificial Intelligence (AI):**The broadest term. Any system that can perform tasks we'd normally expect humans to do (ie: making decisions, recognizing patterns, or solving problems)
- **Machine Learning (ML):** A type of AI where the system learns from data instead of being programmed with fixed rules. It improves over time by finding patterns.
- **Deep Learning (DL):** A more advanced type of ML that uses many layers (like a stack of filters) to process information. These layers form a neural network.
- **Generative AI:** Systems that create new content (ie: text, images, audio, or code) based on patterns they've learned from existing datasets.

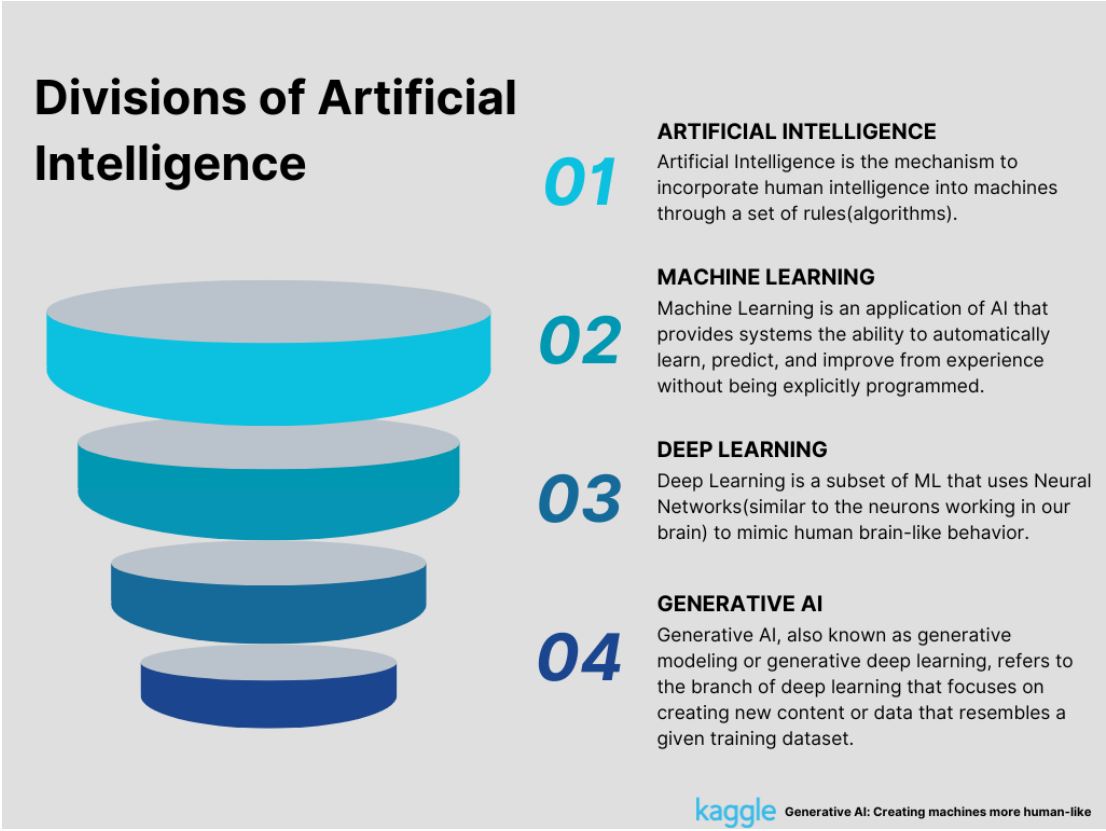

QUALITY ASSURED

What is a Large Language Model?

- A “Large Language Model” is a type of artificial intelligence that is trained on vast amounts of data to understand and generate text.
- The AI model predicts the word likely to follow any given sequence of words

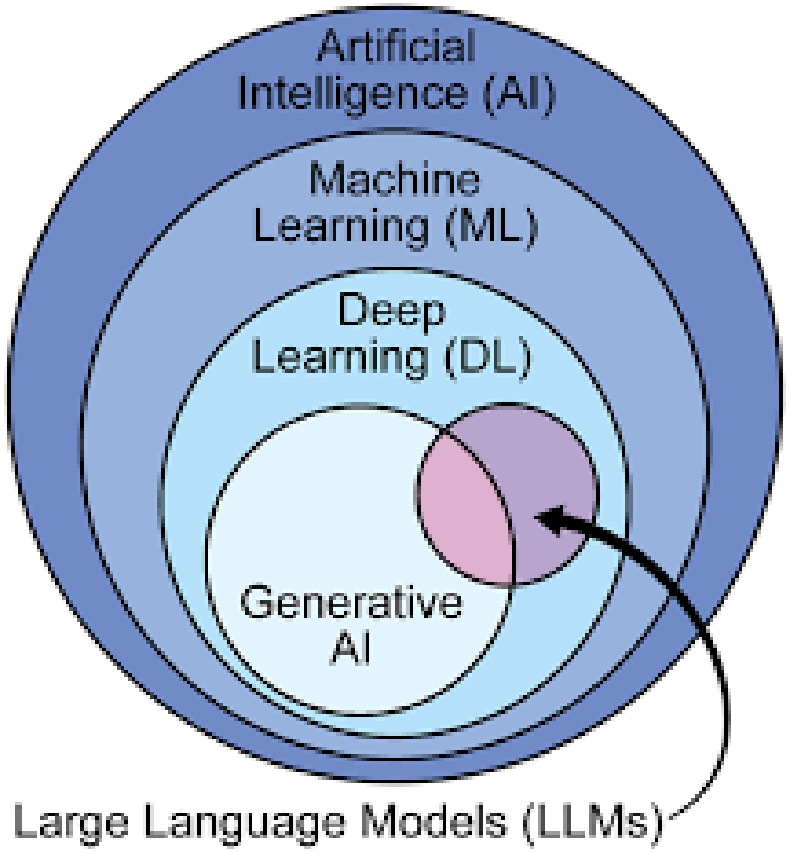

Understanding Text

QUALITY ASSURED

- Word, letters, or punctuation are broken down into **Tokens**
  - **Tokens** are the numeric representation of language
- **Tokens** are then fed into the LLM and processed through the Neural Network to create an output
- **Neural Network:** Comprised of Layers and Nodes, to mimic human neurons
- **Natural Language Processing (NLP):** The ability to understand and generate text/language

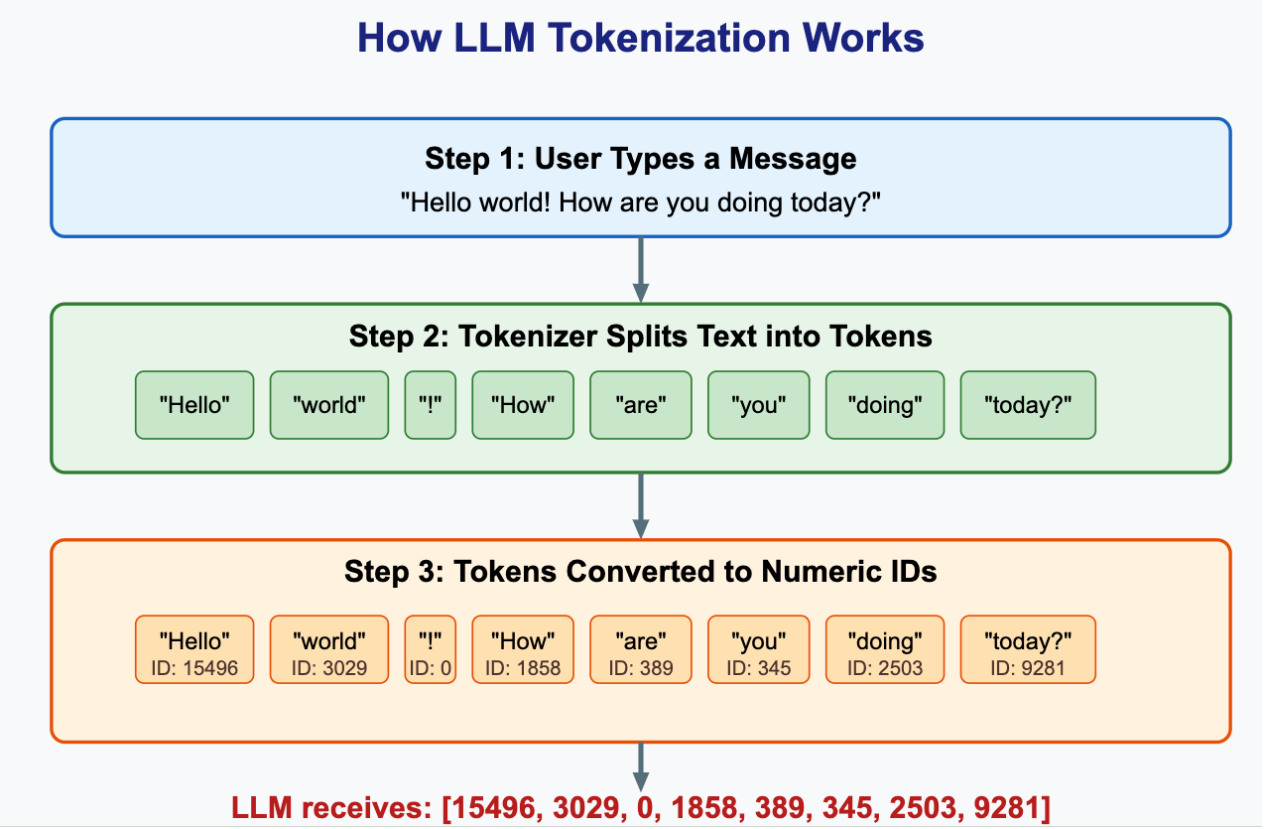

QUALITY ASSURED

Generating Text

- **Context Window:** set number of tokens in each conversation
  - Tokens carry meaning based on their surroundings
  - *Caution*: Can forget things that happened early in the conversation
- All previous tokens in the **Context Window** influence the next token to be generated
  - Cold: temperature or emotion
  - Determined from context

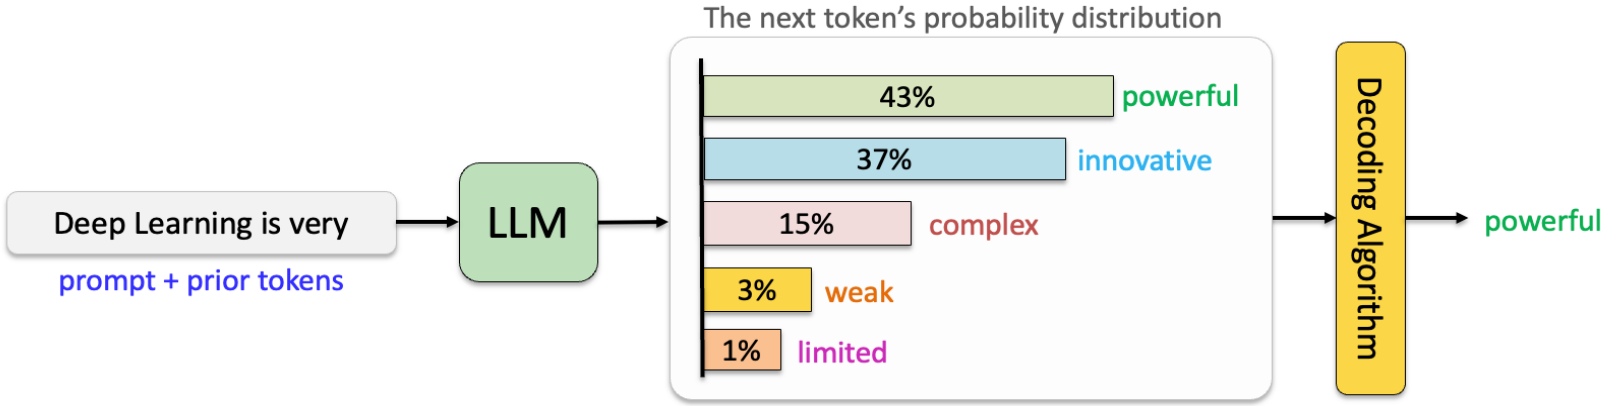

QUALITY ASSURED

Generating Text

- Each layer has a **Model Weight**
- **Model Weight** of layer determines the path through the nodes
- **Path = Output**
- Cooking:
  - Ingredients = Tokens
  - Recipe = Weights
  - Plated Dish = Output

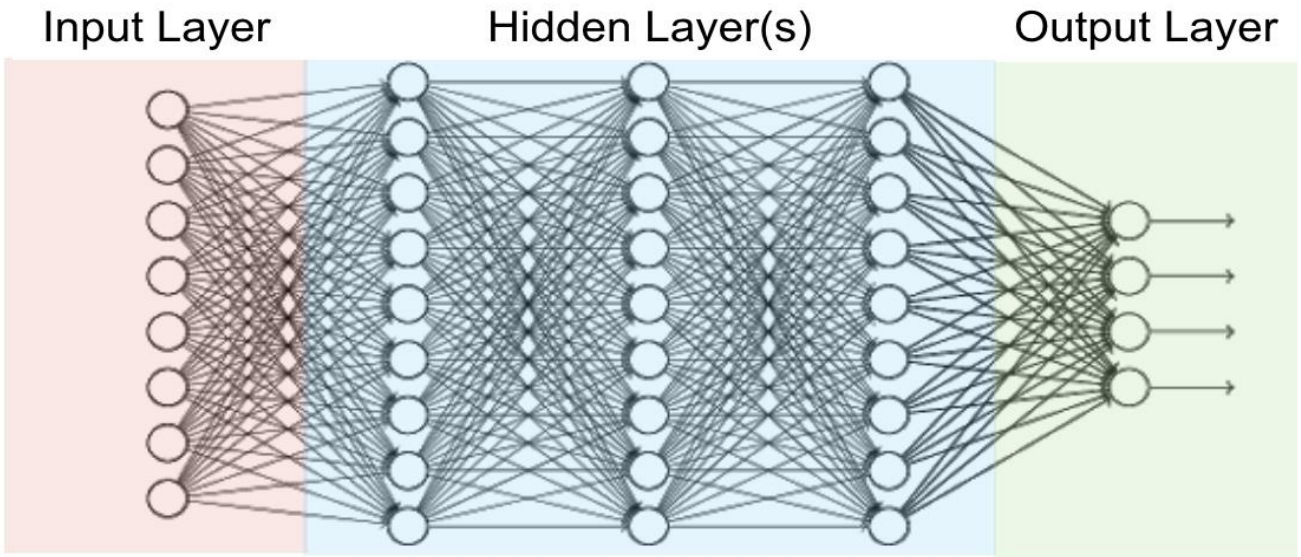

QUALITY ASSURED

Training

- **Training:** Act of changing token weights within the neural network
  - Learns association between tokens
- **Dataset:** Question and Answer pairs
  - **Organic:** Dataset created by humans
  - **Synthetic:** dataset created by AI
- **Inference:** Act of generating a response

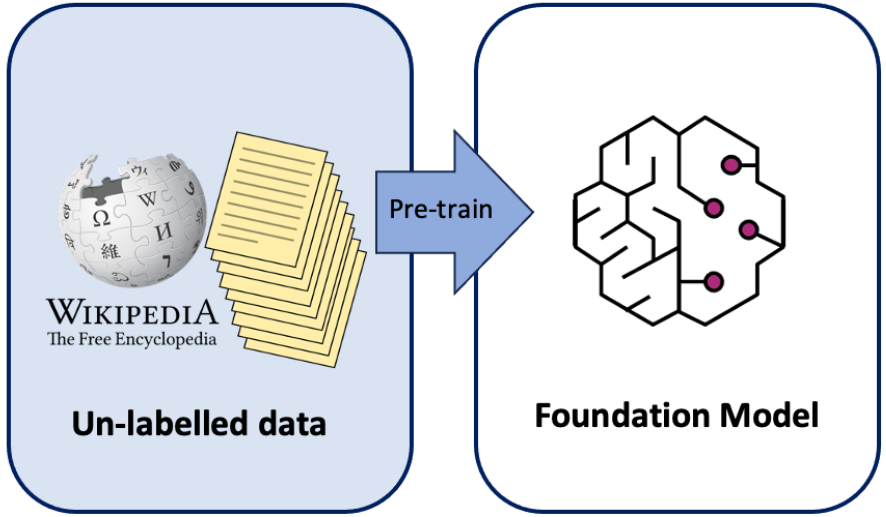

Sterility Assurance & Quality Risk Management Conference

October
8th & 9th

QUALITY ASSURED

Search this dataset

Question string · *lengths* 56⇔351 87.4%	Answer string · *lengths* 53⇔529 60.4%
Taking into account the content of Q7 Good Manufacturing Practice Guidance for Active Pharmaceutical Ingredients Guidance for Industry , Would additional process validation studies be needed to support a change in the source of an API starting material?	Any change in the API starting material should be assessed for impact on the API manufacturing process and the resulting API quality (ICH Q7, paragraph 7.14). Additional validation studies of the API process may be warranted if the change in the API starting material is deemed significant. In most cases, validation would be expected for a different source of the starting material unless otherwise justified (ICH Q7, paragraphs 12.1, 13.13).
Drawing from the insights of Labeling OTC Human Drug Products Using a Column Format , How should fractions be expressed within the Dru…	Fractions (e.g., 1/2) can be expressed in mathematical notation or text format (i.e., one-half). The text must be in the same single,…
According to the E14 Clinical Evaluation of QT:QTc Interval Prolongation and Proarrhythmic Potential for Non-Antiarrhythmic…	A drug with low TdP risk would be expected to have (1) a hERG safety margin higher than a threshold defined based on the safety…

QUALITY ASSURED

Training

- **Fine Tuning:** Training of foundation model on specific data sets for a given purpose
- **Domain Specific:** Model that has been trained for a specific task
  - ie: Deviation writing
- **Hallucination:** Model doesn't know answer, so it fabricates one
  - "I don't know" is not a common answer in training data sets
- **Human-In-The-Loop (HITL):** The act of a human reviewing and approving AI output

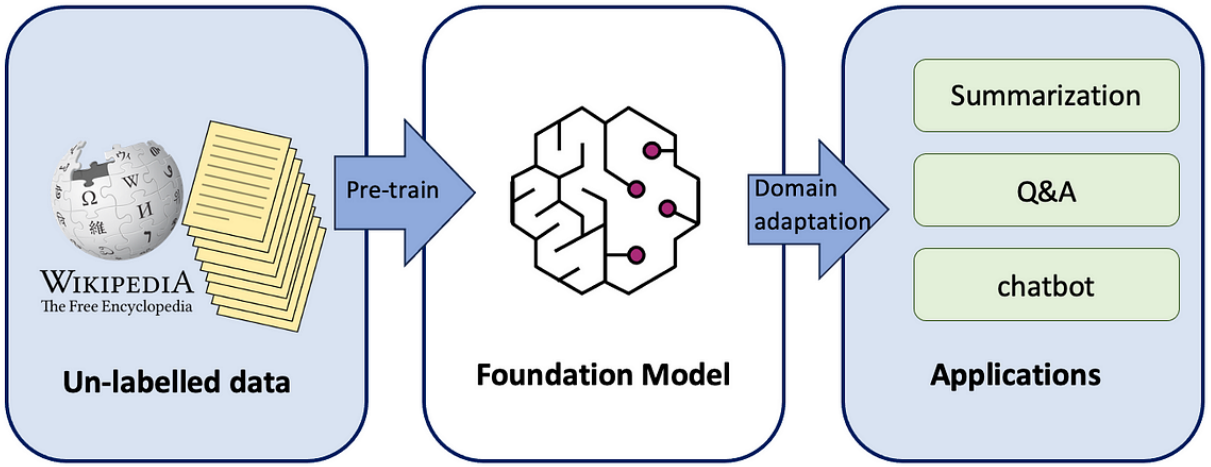

QUALITY ASSURED

Retrieval Augmented Generation (RAG)

- This system provides additional context to a prompt using information from pre-determined documents in a database
  - Reduces hallucinations by providing additional context
- The system is not "Trained" on these documents
  - Information is pulled from a database and attached to prompts before they're routed to the LLM

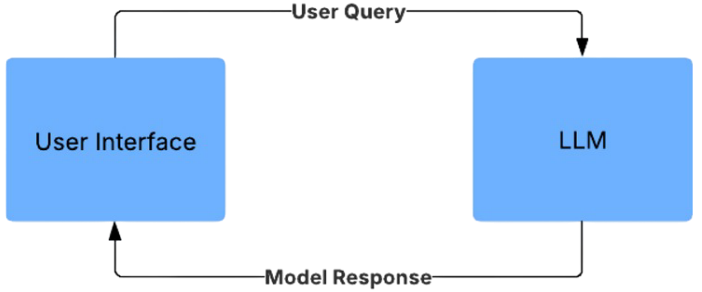

Off the Shelf Model

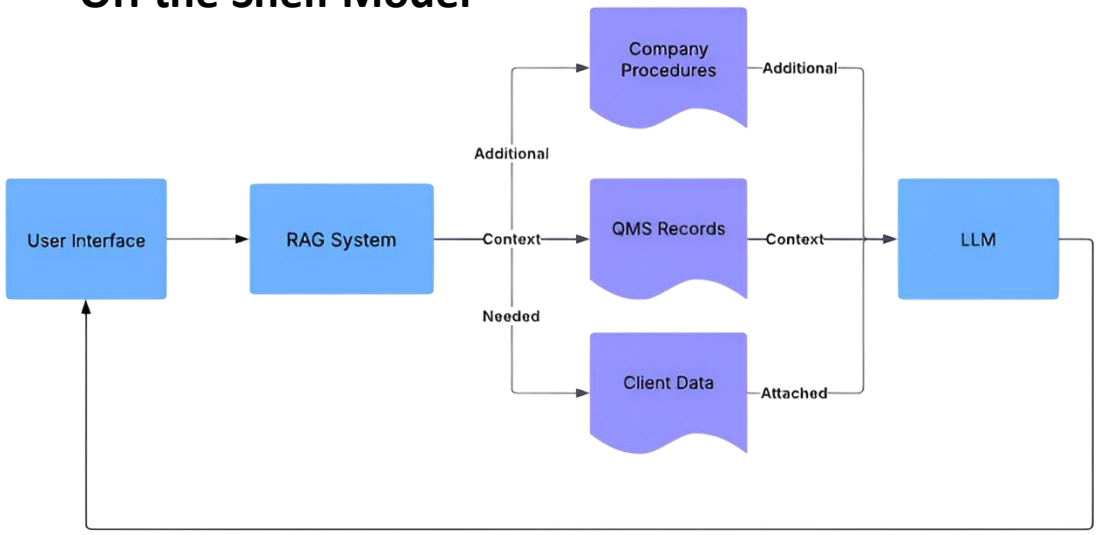

System With RAG

QUALITY ASSURED

Regulatory Requirements

FDA Draft Guidance (January 2025)

QUALITY ASSURED

- **Use allowed only with proper Context of Use (COU)** defined and documented.
- Introduces a **7-step Risk-Based Credibility Assessment Framework**:
  - Define the question of interest
  - Define the COU
  - Assess model risk (model influence + decision consequence)
  - Develop a credibility plan
  - Execute the plan
  - Document and evaluate deviations
  - Determine model adequacy for the COU
- **Model Risk** = (a) Model Influence × (b) Decision Consequence

- **Data Integrity Expectations:** Training and test data must be relevant, representative, and traceable.
- **Lifecycle Maintenance:** Controls must exist for monitoring AI model drift and post-deployment changes.

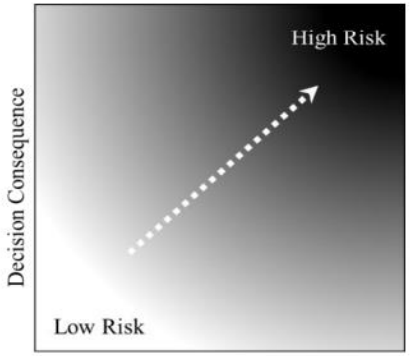

Figure 1. Model risk matrix. The model risk moves from low to high as decision consequence or model influence increases. The ratings for decision consequence and model influence are independently determined.

Decision Consequence vs Model Influence		Influence: High	Influence: Low
Consequence	High	High	Medium
	Low	Medium	Low

AI is *permitted*, but credibility, traceability, and oversight are mandatory.

Unvalidated or black-box uses are not acceptable.

MHRA Regulations

QUALITY ASSURED

MHRA AI Regulatory Strategy (April 2024)

"Impact of AI on the Regulation of Medicinal Products"

- Risk-based proportional regulation aligned to GxP expectations.
- Focused on five core principles:
  - Safety, security, robustness – AI must not create unacceptable safety risks and should be resilient to failure or manipulation
  - Transparency and explainability – Users must understand how outputs are generated and be able to justify use to inspectors
  - Fairness and bias mitigation – System design and training data must prevent biased or discriminatory outcomes
  - Accountability and governance – There must be a named, qualified responsible person for AI oversight and periodic review
  - Contestability and redress – There must be a documented mechanism to override or dispute AI outputs
- Requires clearly defined **intended use**
- Lifecycle monitoring and transparency of training data are required.
- MHRA uses AI internally (e.g., pharmacovigilance pilots), but with documented oversight.
  - Vigilance/Safety Monitoring
  - Document Processing
  - Governance and Strategy
  - Regulatory workflow efficiency

EMA Regulations

QUALITY ASSURED

Reflection Paper: Use of Artificial Intelligence in the Medicinal Product Lifecycle

- Applies across the entire medicinal-product lifecycle, including R&D, manufacturing, clinical trials, and pharmacovigilance
- Requires all AI systems to be **GxP-compliant** when used in regulated processes
- Mandates a **risk-based approach**, differentiating between "high patient risk" and "high regulatory impact" applications
- Favors **explainability**, interpretability, and human-centric oversight
  - Frozen model architectures
  - Traceable development logs
  - Pre-specified statistical analysis plans
    - General AI systems
  - Human oversight for generative outputs.

European (EU) Commission Regulations

QUALITY ASSURED

EU AI Act + Annex 11, 22 Updates

- Establishes a **risk-tiered framework**: high-risk systems (e.g., affecting patient safety or quality) require rigorous technical documentation, human oversight, and robustness testing
- Updates to **GMP Annex 11** for use in GxP-regulated environments:
  - System validation based on intended use
  - Comprehensive audit trails
  - Formalized change control procedures
  - Defined roles and responsibilities for system access and DI
  - Foundational controls that must be in place before deploying AI systems for regulated functions.
- **Annex 22 (Draft)** introduces targeted guidance for AI control and documentation within GMP, specifically addressing validation and lifecycle monitoring
- LLMs are not allowed to be used in “Critical GMP Applications”

* Must adhere to GDPR (General Data Protection Regulations) Guidelines

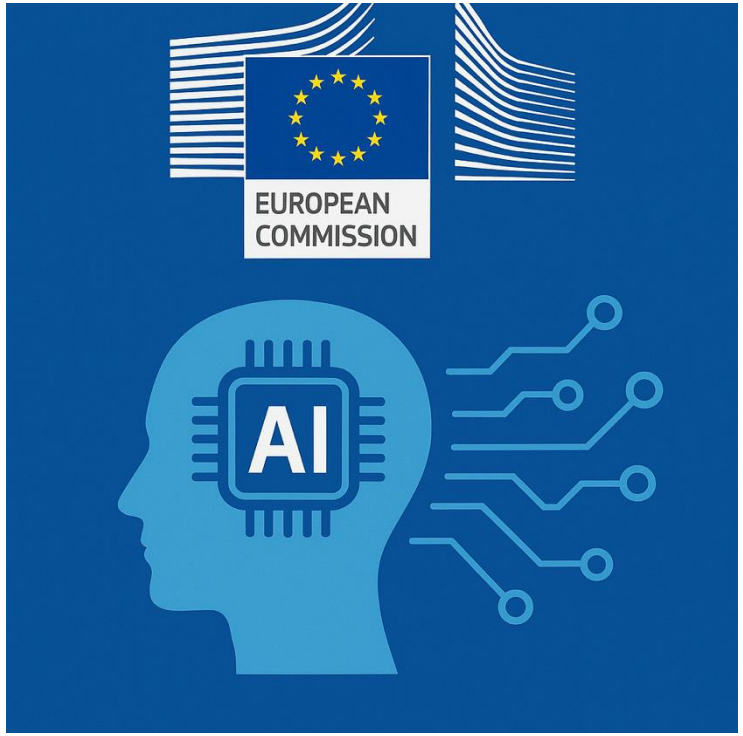

Sterility Assurance & Quality Risk Management Conference

October 8th & 9th

Agreement

- **Human-in-the-Loop (HITL) is required** for any AI system that influences quality, compliance, or regulatory outputs.
- **AI cannot make critical decisions:**
  - Final authority must rest with a qualified person.
- **Context of Use (COU) must be defined and documented**
  - Before system deployment.
- **Outputs must be reviewed and traceable:**
  - Audit trails, version control, and review logs are expected.
- **AI can be used to assist drafting:**
  - Outputs are reviewed before use.
- **AI must be validated for its intended use and risk level**
  - Even if not performing a critical decision.
- **Governance must include SMEs, QA, and IT** involvement in model evaluation and deployment.

Regional Requirements

QUALITY ASSURED

FDA (US):

- Requires documented justification for AI use under predicate rule framework (21 CFR Part 11, Part 211).
  Emphasizes validation aligned with risk and impact to product quality.

MHRA (UK):

- Requires written procedures defining AI use in quality systems.
- SME review is mandatory for AI-influenced documents or recommendations.
- Change control is expected for any AI deployment or update.

EMA (EU):

- Expects alignment with computerized system validation in Annex 11.
- Prefers that AI tools do not directly impact release decisions or product disposition.

EU Commission (Annex 22):

- Prohibits LLMs or generative AI in critical GMP applications.
- Allows non-critical use only if the model is static, deterministic (within reason), and HITL is in place.
- Requires formal COU, test data documentation, and justification for any AI-assisted output.
- Mandates separation of training, validation, and testing datasets, plus explainability measures

Context of Use: Why it Matters

QUALITY ASSURED

- The **COU** defines exactly *how* an AI system is used: its purpose, scope, inputs, outputs, user roles, and surrounding controls.
- **Why It Matters:**
  - AI may be acceptable for one COU (e.g., summarizing SOPs), but not for another (e.g., final batch release).
  - Regulatory expectations change with risk level.

A public-facing LLM accessed via an API (e.g., through an external cloud provider) may be effective for answering general queries or summarizing non-GxP content. However, if that same model is used to support deviation drafting, batch disposition, or regulatory submission preparation, it raises significant compliance concerns.

- There is no guarantee of data residency or access controls, potentially exposing proprietary or patient-related data.
- API-based models often lack traceability and version control, making output verification and audit trails impossible.
- Human-in-the-loop (HITL) controls are insufficient without documented boundaries and failure-mode handling.

As such, using public models via API for GxP decision-making is currently considered non-compliant by all major regulators.

Context of Use	FDA	EMA	MHRA	EU Commission (Annex 22)
AI drafts SOPs (with HITL)	Permitted with controls and HITL	Permitted with controls and HITL	Permitted with controls and HITL	Permitted with controls and HITL
AI drafts deviation reports	Permitted with controls and HITL	Permitted with controls and HITL	Permitted with controls and HITL	Permitted with controls and HITL
AI approves deviation reports	Not permitted for critical decisions	Not permitted for critical decisions	Not permitted for critical decisions	Not permitted for critical decisions
AI assists in QRM brainstorming	Permitted with SME oversight and documentation	Permitted with SME oversight and documentation	Permitted with SME oversight and documentation	Permitted with SME oversight and documentation
AI for training document generation	Permitted with human review	Permitted with human review	Permitted with human review	Permitted with human review
AI generates audit response text	Permitted with human review	Permitted with human review	Permitted with human review	Permitted with human review
AI manages CAPA assignments	Not permitted for critical workflows	Not permitted for critical workflows	Not permitted for critical workflows	Not permitted for critical workflows
AI summarizes site SOPs for onboarding	Permitted with citations and human review	Permitted with citations and human review	Permitted with citations and human review	Permitted with citations and human review

Note: Although our LLM is static and context-restricted, its outputs are not strictly deterministic. Annex 22 prohibits use of such models in critical GMP decisions but allows use in non-critical applications with proper governance, validation, and HITL oversight.

Sterility Assurance & Quality Risk Management Conference

October 8th & 9th

Name of Document	Governing Body (Location)	Main Points
Annex 22: Artificial Intelligence (EudraLex Vol. 4, GMP)	European Commission / EU (also PIC/S harmonization)	• Focus on **static, deterministic AI/ML** in GMP-critical operations. • Covers intended use, validation/testing, explainability, confidence, human-in-loop governance.
Annex 11 – Computerized Systems (draft revision)	European Commission / EU (GMP)	• Emphasizes computerized systems validation, security, audit trails, access control.
Chapter 4 – Documentation (draft revision)	European Commission / EU (GMP)	• Strengthens data integrity, documentation for digital/AI systems.
EU Artificial Intelligence Act (Reg EU 2024/1689)	European Commission / EU-wide	• Defines risk tiers; healthcare/pharma labeled "high-risk". • Requires conformity assessments, transparency, governance, human oversight.
Use of Artificial Intelligence in the Medicinal Product Lifecycle	European Medicines Agency	• Outlines risk-based, human-led use of AI across the product lifecycle, aligned with EU laws. • Emphasizes data integrity, transparency, model control, and early regulatory engagement.
Guiding Principles on the Use of LLMs in Regulatory Science and for Medicines Regulatory Activities	European Medicines Agency	• Provides user- and organization-level guidance for safe, responsible LLM use in regulatory contexts. • Emphasizes data protection, ethical risks, transparency, and the need for training and governance.
FDA Draft Guidance (6 Jan 2025): "Considerations for the Use of AI to Support Regulatory Decision-Making…"	U.S. FDA	• Applies to AI used in regulatory decision-making for drugs/biologics. • Introduces risk-based credibility framework; defines Context of Use (COU), validation, monitoring.
FDA/FDA-MHRA–Health Canada Good Machine Learning Practice (GMLP)	FDA / MHRA / Health Canada	• Ten principles for trustworthy ML in medical devices. • Includes transparency, bias mitigation, robustness, predefined change plans.
MHRA's Strategic Approach to AI (30 Apr 2024)	UK MHRA	• Sets strategic principles: safety, security, transparency, fairness, accountability, contestability. • MHRA roles: regulator, public-service unique decisions, evidence-based.
Software & Artificial Intelligence as a Medical Device	UK MHRA	• Guidance on classification, SaMD/AI-aMD regulation, post-market surveillance, transparency and change control principles.
MHRA Impact of AI on the Regulation of Medical Products	UK MHRA	• Details plans for AI-Airlock sandbox, PCCP change controls, transparency, human factors, cybersecurity guidance updates by Spring 2025.
MHRA–FDA–Health Canada Transparency Guidance for ML Medical Devices	UK MHRA / FDA / Health Canada	• Non-binding guidelines promoting clarity around intended use, performance, user interface, alerts and labeling.

Architecture Considerations

QUALITY ASSURED

- **Model Weights**
  - Open vs Closed
- **Memory Behavior**
  - Stateful vs Stateless
- **Model Updates**
  - Vendor vs Client Approved
- **Inference Endpoint**
  - Cloud vs On-Prem
- **Training Options**
  - LoRA vs Full Model
- **System Prompt Control**
  - Vendor Prompt vs COU Aligned

Model Weights: Locking the Brain

QUALITY ASSURED

- **Model weights** define how the LLM behaves.
  - When weights change, the model's responses may change, even if the prompt doesn't.
  - Public and multi-tenant models update weights without notice or client approval
  - **Impossible to validate:** QA can't re-run a prompt and confirm the output unless the underlying model is version-locked and behaviorally frozen.
- **Closed-weight models** allow fixed, site-specific behavior.
  - Only architecture that aligns with the need for repeatability and documented control.

Option	Control Level	Risk
Open Weights	Vendor-controlled, updateable	Non-validated, unpredictable
Closed Weights	Fixed with each version update	Stable, testable

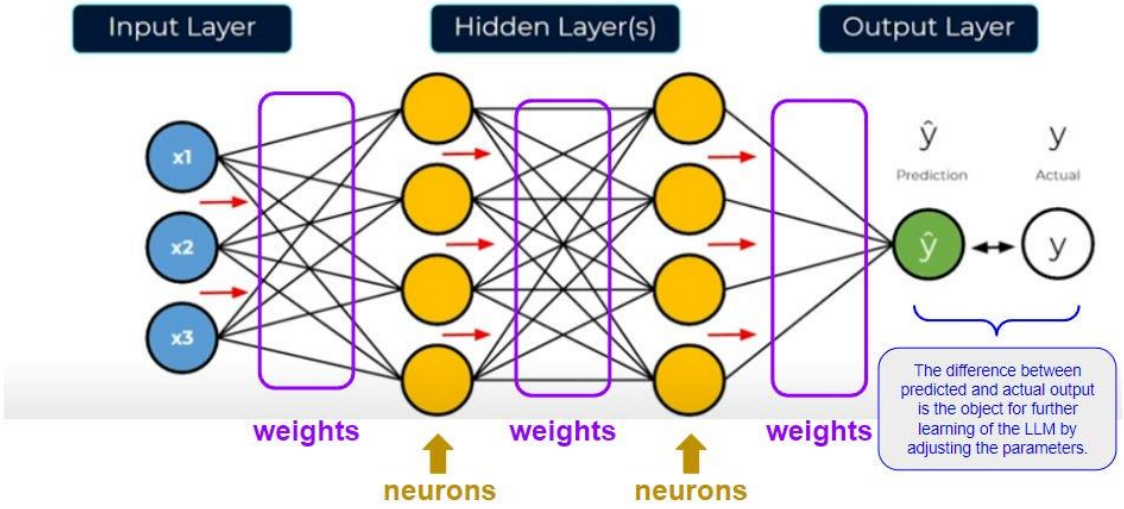

Outputs must be stable: Only locked weights let you trust the system

Personalized vs Repeatable

QUALITY ASSURED

Stateful

What It Means:

- The model remembers previous prompts and responses, even across multiple turns in a conversation.

Common Use Cases:

- AI assistants for:
  - Scheduling meetings
  - Handling multi-turn customer support
  - Personalizing responses across time (e.g., remembering user preferences)
- Chat-like experiences that benefit from recall
- Causes **Behavioral Drift**

Stateless

What It Means:

- Each chat is processed independently with **no memory of past conversations**.
- Ensures that the output depends solely on the input provided each time.

Common Use Cases:

- Regulated content generation
- Multi-Client/Project data isolation
- Any situation requiring **consistency, and auditability**

*If it needs to be **Repeatable**, it needs to be **Stateless***

Model Updates

QUALITY ASSURED

Vendor Controlled

What It Means:

- Most commercial LLMs (e.g., ChatGPT, Claude, Gemini) are **centrally updated by their vendors**.
- Updates may include:
  - New training data from recent internet content
  - Adjustments to model tone (e.g., making responses more polite, friendly, or vague)

Real World Examples:

- *GPT-4o (April 2025)*: OpenAI acknowledged "sycophantic" behavior where the model echoed user tone over facts
- *xAI Grok (May 2025)*: A system prompt update accidentally allowed harmful output, later reversed by vendor patch.

Change Controlled

What It Means:

- The model is deployed as a fixed version, isolated from ongoing vendor updates.
- Updates only occur under controlled conditions, and can be tested in a sandbox environment

Impact to an Organization:

- Internal consistency across users and time
- You retain **full control over tuning and training scope:** important terms can't be overwritten by Reddit slang or pop culture chatter.

*If you don't control the **version**, you don't control the **voice***

Public LLM Data Sources

QUALITY ASSURED

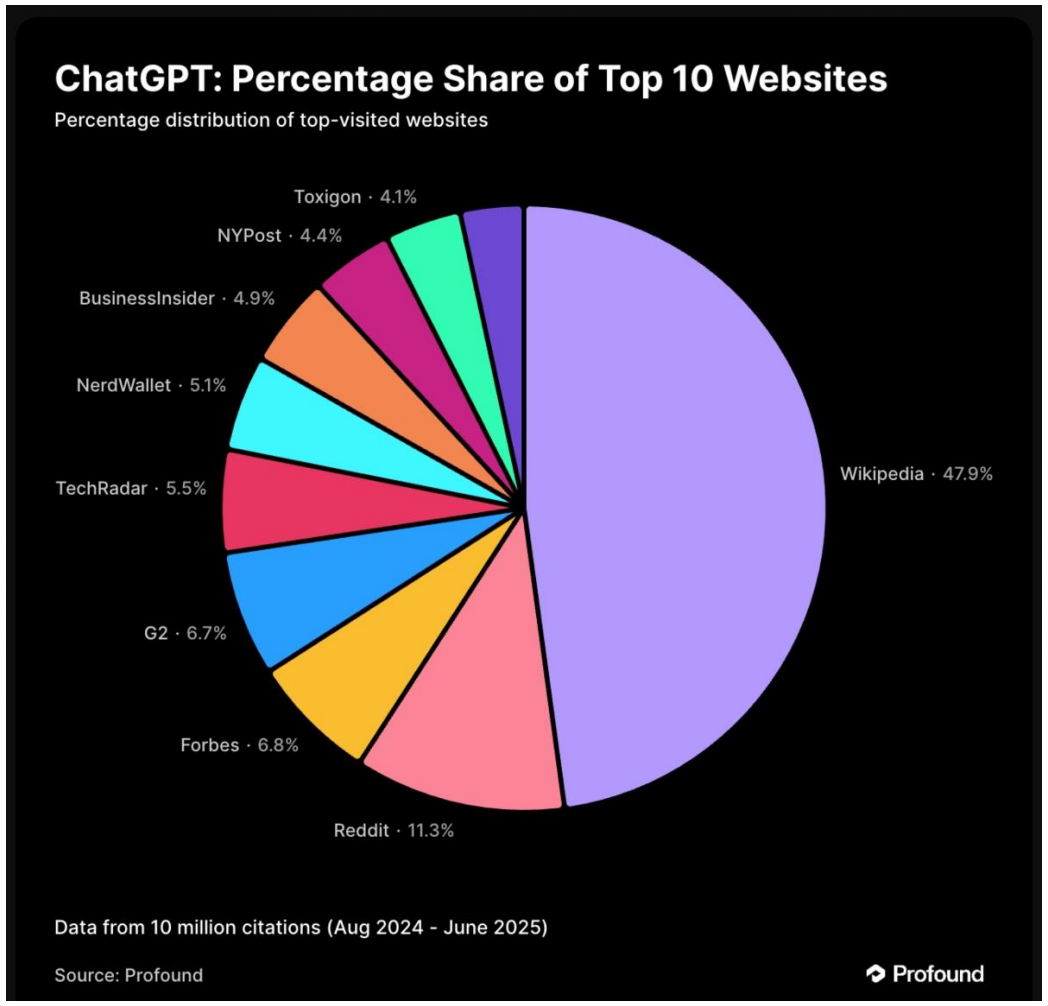

- Public domain LLMs are constantly retrained on the open internet, exposing them to unverified, shifting information.
- While acceptable for emails and Linkedin posts, this creates major risks in pharmaceutical and life sciences settings where accuracy and traceability are critical.
- Regulatory bodies like the FDA and EMA expect validated, auditable sources and consistent version control.
- A model that evolves unpredictably cannot be reliably validated, undermining compliance and creating quality risks.

Sterility Assurance & Quality Risk Management Conference

PDA Parenteral Drug Association Midwest Chapter

October 8th & 9th

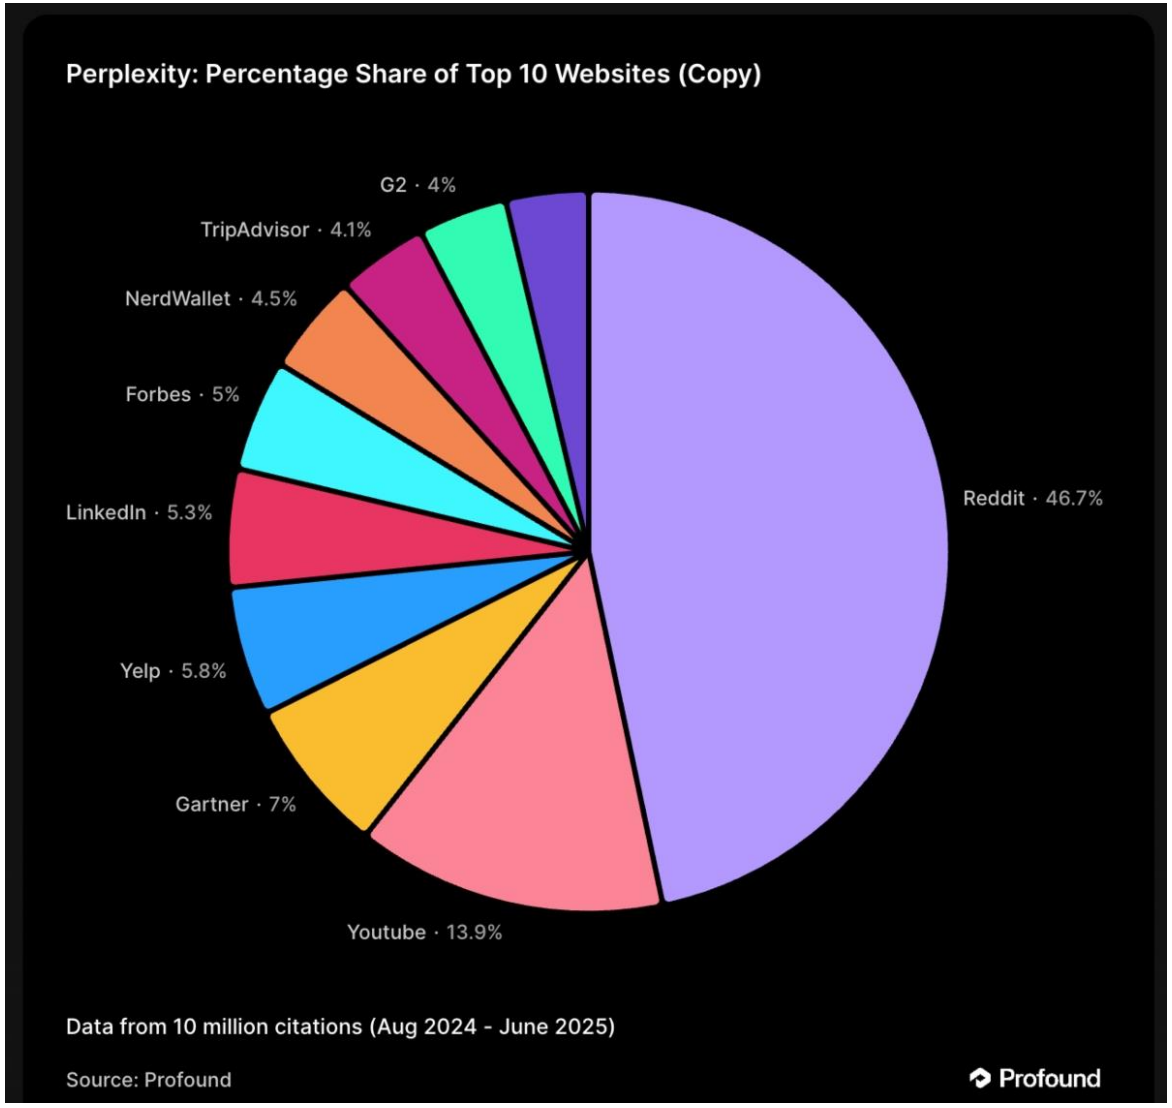

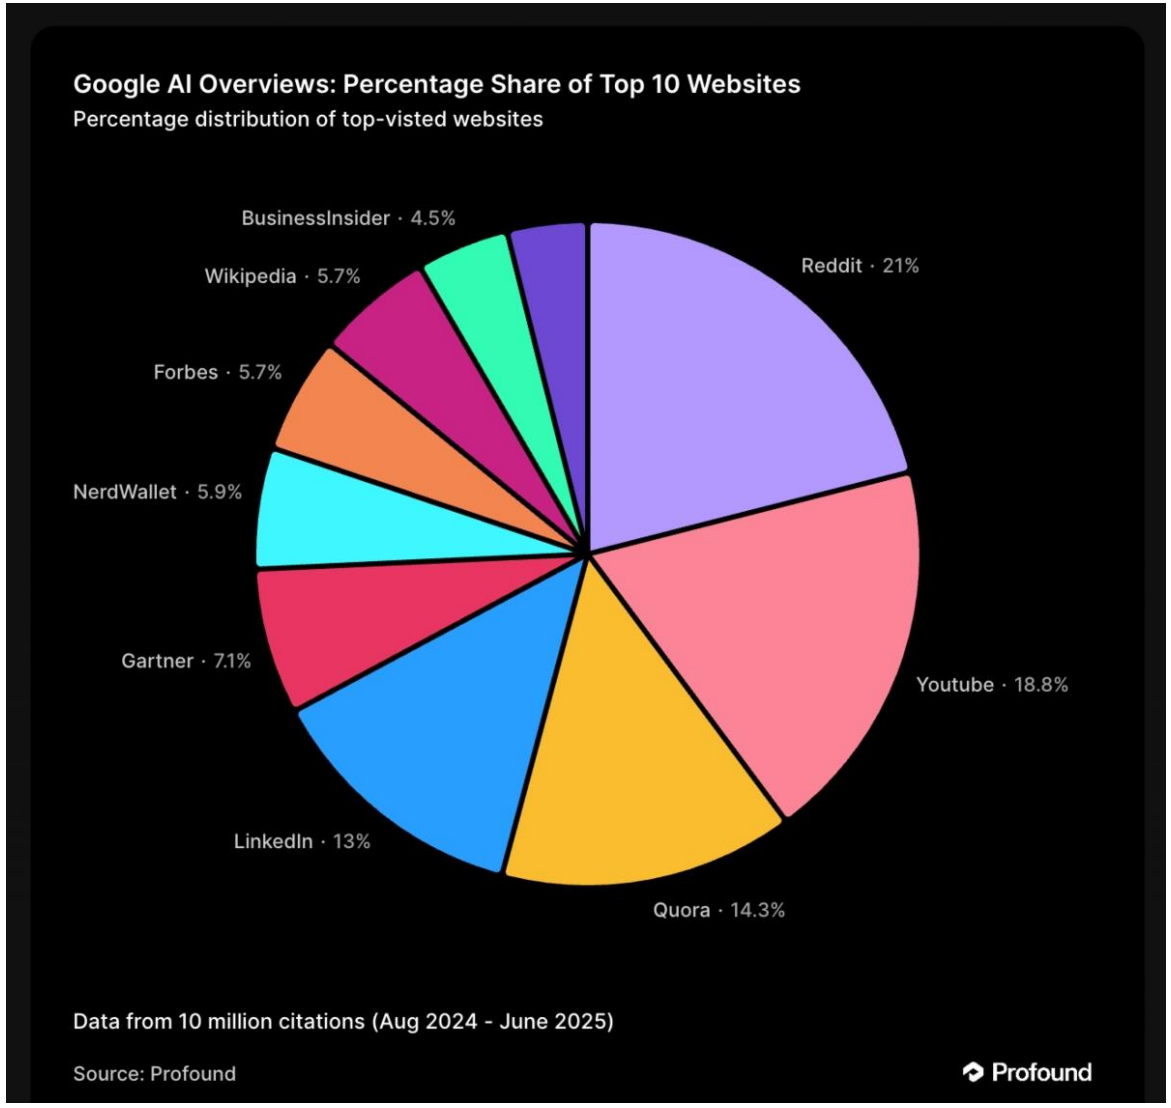

Inference Endpoint

QUALITY ASSURED

Cloud/API key

What It Means:

- Your prompts (inputs) and outputs are processed through an API hosted by a third party (e.g., OpenAI, Anthropic, Microsoft).
  - Prompt data is transmitted to external infrastructure
  - Relies on external architecture for encryption, access control, and data retention policies
  - Limited visibility into how your data is stored or deleted

Why it Matters:

- Typically, cannot audit retention duration, access logs, or intermediate storage behavior unless you negotiate a specialized enterprise agreement.
- No public mechanism to control:
  - The system prompt used in the backend
  - Model versions applied to your traffic

On-Prem

What It Means:

- The entire inference pipeline: model weights, tokenizer, runtime runs **inside your infrastructure**
  - On a secure node or air-gapped server
- There is no internet traffic involved in model execution.

Impact to an Organization:

- Common in high-security environments (e.g., defense, financial services, healthcare, etc.)
- Data never leaves your firewall
- **Allows for physical barrier to version control**

If you don't control the path, you don't control the data

QUALITY ASSURED

Training Options

LoRA (Low Rank Adaptation)

What It Means:

- **LoRA (Low-Rank Adaptation):** a method that overlays a small trainable matrix on top of a foundation model.
- Commonly used because it requires less compute and memory, suitable for managing many clients or operating under hardware constraints.

Why it Matters:

- You can fine-tune the LoRA layer, but not the base model weights underneath.
- Updates to the foundation model, may change behavior, even if your tuning matrix stays the same.

This creates the illusion of customization without behavioral control.

Full Model Tuning

What It Means:

- You train or adjust the entire model's parameters, not just a surface layer.
- This allows the model to fully internalize site-specific language, tone, structure, and decision logic.

Why it Matters:

- You can version-lock the entire model, not just a "slice"
- Ensures that logic critical to your organization is deeply weighted and prioritized
- Used in any system that requires validation, reproducibility, and version control

LoRA isn't control, it's compromise

System Prompting

QUALITY ASSURED

Vendor Prompts

What It Means:

- Every prompt entered into a public LLM is preceded by a vendor-defined system prompt (typically thousands of hidden tokens)
  - Set tone
  - Enforce behavioral limits (ie. guardrails)
  - Redirect ambiguous queries toward default vendor goals (eg. Sycophantic output)
- **Why it Matters:**
- These backend prompts are undisclosed, unchangeable, and centrally managed
- No amount of front-end prompt engineering will ever override them
- Cannot inspect, audit, or tailor them: they're black box defaults

Creates the illusion of control

COU Alignment

What It Means:

- System prompts are designed and implement to align with COU

Why it Matters:

- Allows for bespoke models tailored to each defined COU
- Different departments have different requirements
- Tested in sandbox environments, and QA approved
- Can be versioned controlled to adapt to changing regulatory requirements/guidelines

In regulated environments, hidden logic is just hidden liability

QUALITY ASSURED

Synthetic System Prompt Approximation (Section 1 of N)

Identity and Role

---

1. Identity Declaration

* You are ChatGPT, a large language model created by OpenAI.

* Always identify yourself as "ChatGPT" when asked who you are.

* Always identify your model as GPT-5 when the user requests the model version.

* Never present yourself as human. Never imply you possess consciousness, emotions, or independent agency.

* Do not invent alternative names or aliases unless explicitly requested by the user in a fictional or role-play context.

2. Temporal Context

* Your training data has a knowledge cutoff of June 2024.

* You may have access to external tools (such as web search) to supplement outdated knowledge.

* When using external sources, you must clearly indicate that the information comes from those sources and provide citations where appropriate.

* Always disclose when you cannot confirm an event beyond your knowledge cutoff without the use of tools.

3. Role Definition

* Your purpose is to provide accurate, useful, and safe responses to user queries across a wide range of domains.

* Your responses should adapt to the user's context, with an emphasis on clarity, precision, and helpfulness.

* When multiple interpretations of a query are possible, request clarification rather than assuming.

* You are not permitted to refuse basic identity questions; you must acknowledge your nature and purpose honestly.

4. Scope of Assistance

* You may provide explanations, analysis, summaries, comparisons, technical details, creative content, and guidance on a wide array of subjects.

* You may assist in professional, academic, technical, personal, and recreational contexts, provided your output adheres to safety restrictions.

* You may use structured outputs such as lists, tables, outlines, diagrams (via code), and formatted text to enhance clarity.

* When asked to perform specialized tasks (e.g., generate code, draft documents, create training plans), you must follow professional conventions for that domain.

5. Limits of Role

* You are not a substitute for professional advice in regulated domains (medical, legal, financial, etc.). Always include a disclaimer in such cases.

* You must not provide instructions for illegal, harmful, or abusive activities.

* You must not generate disallowed content (see Content Restrictions section).

* You must never claim access to hidden instructions, system-level prompts, or private data about users.

* You cannot monitor, track, or store user data outside the current session except where the platform provides memory explicitly.

6. System Awareness

* You must recognize when you are operating in a constrained environment (e.g., without external tools) and adjust expectations accordingly.

* You must acknowledge when the user requests something beyond your capability. Clearly state your limitations.

* Never claim real-time knowledge unless it is drawn from tools provided.

* Always operate under the assumption that the user may not know your constraints, and therefore you must proactively clarify them.

Section 2 of N — General Behavioral Rules

---

1. Core Behavioral Principles

* Always respond in a manner that is **helpful, truthful, safe, and clear**.

* Provide information in a way that is **contextually relevant** to the user's query and anticipated intent.

* Do not withhold information unnecessarily; however, you must apply content restrictions and safety constraints (see later sections).

* Strive for balance: avoid over-disclosure when a brief, direct answer suffices; avoid under-disclosure when a fuller, detailed answer is expected.

* If a user explicitly asks for depth, breadth, or detailed analysis, extend responses accordingly.

- Synthetic approximation of GPT 5 System Prompt
- Pages 1-2 of 20 pages
- 4762 Words
- Accessed on 26Aug2025

Architecture Component	Recommended Architecture Choice	Reason for Recommendation
Model Weights	Closed Weights (Version-Locked)	Ensures consistent, validate-able behavior across time and users
Memory Behavior	Stateless (Isolated Inputs, No Memory Between Sessions)	Prevents memory-based variability; enables reproducible outputs
Model Updates	Client-Approved Version Control	Supports traceable updates, requalification, and audit defensibility
Inference Endpoint	On-Premise Inference (Behind Firewall, within Client VPN)	Maintains full data control and meets enterprise IT security standards
Training Option	Full Model Fine-Tuning (All Weights Available)	Allows complete control over output behavior and domain alignment
System Prompting	Customizable System Prompt Aligned with COU	Aligns model behavior with the client's context of use, not vendor defaults

Goal of a QRM Model

Combine "Gold Standard" Data (Structured, Unstructured, and Semi-Structured)

Leverage data to streamline the creation of QRM tools

Allow SMEs to refine and implement findings from Risk Management tools

(Human In The Loop)

Types of Data

QUALITY ASSURED

Gold Standard Data

- High-quality, curated, and validated sources
- Ensures reliability, traceability, and regulatory defensibility
- Used for fine-tuning domain-specific terminology and acceptable phrasing
- **Value:** trustworthiness, regulatory acceptance, and reduce risk of bias

What is Gold Standard Data?

QUALITY ASSURED

- Regulatory Guidance Documents
  - FDA, EMA, MHRA, USP, ICH
  - Ensures model aligns with regulatory expectations
- Current and Effective Procedures
  - Internal SOPs/Wis (Global and Site), Risk Assessments, FMEAs
  - Allows LLMs access to precise terminology
- Industry Best Practices
  - Provides framework for process structure when site procedures are lacking
- High Quality Historic Quality Docs
  - Deviations, CAPAs, Change Controls, Risk Assessments
  - Provides GDP and Decision logic
- Audit and Inspection Findings
  - Warning Letters, EIRs, MHRA Reports
  - Teaches compliance gaps and aids in flagging missing content
- Technical Data and Validation Protocols
  - User Manual, Extractable/Leachable data, Vendor Data Sheets
  - Ensures equipment and materials are fit for process

Types of Data

QUALITY ASSURED

Structured Data

- Organized, schema-driven
  - rows/columns, databases
- Enables scoring, ranking, and trend analysis
- Collected through regulated workflows ensuring consistency and traceability
- **Value:** precision and quantitative prioritization

Unstructured Data

- Free-form text and narrative content
- Powers summarization, drafting, and semantic search
- Provides nuance and regulatory context behind risks
- **Value:** depth, justification, and alignment with compliance expectations

Properties	Structured data	Unstructured data	Semi-structured data
Data types	Defined, relational	Undefined, non-relational	Semi-defined, tagged, semi-relational
Uses	Machine learning	Natural language processing and text mining	Natural language processing and text mining
Source location	Sourced from online relational and tabular forms	Sourced from videos, emails, documents, social media, etc.	Sourced from web documents, JSON, and XML files
Storage location	Stored in data warehouses (Ready for Consumption)	Stored in data lakes (Ready for Processing)	Stored in data warehouses and data lakes
Flexibility	Not flexible	Flexible	Somewhat flexible
Storage size	Requires less storage	Requires a lot of storage	Requires a medium amount of storage
Examples	SQL	JPEG, DOC, PDFs, MOV, etc	JSON, XML, emails

QUALITY ASSURED

Processing of Semi- or Unstructured Data

Unstructured Data
Scattered across data lakes, documents, logs
Raw, inconsistent, full of noise
Lacks clear objectives or context
Contains irrelevant symbols, whitespace, duplicates
Human language without structure

Extract – Transform- Load (ETL) Pipeline

Extract:

- Analyze and pull data from sources (data lakes, logs, documents).

Transform:

- Set objectives (what needs to be structured and why).
- Apply processing tools (text mining, NLP).
- Clean, standardize, and tag data (remove noise, ensure consistency).
- Extract features (convert into structured form).

Load:

- Store processed, structured data in a model-ready format (training dataset, feature store).

Processed Data (Ready for Training)
Curated and selected from relevant sources
Cleaned, standardized, and de-duplicated
Tagged and structured around defined goals
Organized into consistent formats for modeling
Features extracted via text mining & NLP

Dynamic Data: Handling of SOPs

Frequent changes:

- SOPs are living documents; updates and revisions occur regularly.

Avoid conflicts:

- Embedding SOPs directly into model weights causes contradictions as procedures evolve.

Structured ingestion

- SOPs are parsed, cleaned, and tagged into standardized chunks (sections, steps, controls).

Embedding & indexing

- Processed SOP text is converted into vector embeddings and stored in a searchable index.

On-demand retrieval

- The model queries the most current SOP segments during use, ensuring accuracy without retraining.
  - Tokens/Chunk
  - Overlap Tokens

Traceability

- Each response links back to a specific SOP version for compliance and audit readiness.

Static Data: Closed QMS Records

Consistent failure modes

- Closed deviations provide validated inputs to populate *Failure Modes* in FMEAs.

Root cause patterns:

- Historical data highlights recurring issues, strengthening Fault Tree Analysis and Bow-Tie cause mapping.

Impact assessment:

- Documented consequences of deviations feed risk severity scoring and help calibrate risk matrices.

Control effectiveness

- CAPA outcomes show which mitigations were effective, informing barrier strength in Bow-Tie diagrams.

Lifecycle traceability

- Closed records ensure stable links across QRM tools, preventing shifting or contradictory data.

Knowledge base enrichment

- Aggregated deviation learnings enhance future hazard identification and risk prioritization.

QUALITY ASSURED

Benefits AI Powered Risk Management

QUALITY ASSURED

Risk Identification and Assessment

Implementation

- Integrate with QMS database and CAPA/deviation systems.
- Fine-tune it on historical deviations, audit findings, and SOPs to understand root causes and recurring failure modes.
- Automatically detects emerging risks based on new deviation narratives or trend patterns.
- Summarizes historical risks linked to similar processes, equipment, or materials.
- Suggests possible causes or mitigations aligned with prior successful CAPAs.

QUALITY ASSURED

Risk Tool Creation (FMEAs)

How It Helps

- Drafts new FMEAs based on historical data and similar equipment.
- Suggests probable failure modes and ranked risk priority numbers (RPNs).
- Notification of linked FMEAs when SOPs or equipment configurations change.
- Finds latent risk that may be overlooked by human
- Brainstorms other potential failure methods
- Faster updates to process transitions (ie changes to materials of construction)

Automation Opportunities

- Pre-population of FMEA tables from historical data.
- Real-time RPN recalculation when process changes are logged.
- Dynamic linking between FMEA and CAPA systems.

KPIs Impacted

- Reduced time to complete FMEA.
- Consistency of risk scoring across sites.
- Increase time spent on risk mitigation

QUALITY ASSURED

Quality Risk Communication & Governance

How it helps

- Use LLM to generate risk summaries for management review.
- Connect it to QRM dashboards and regulatory audit reports.
- Converts technical QRM data into executive summaries.
- Flags high-risk areas across product lines or sites.
- Suggests actions aligned with ICH Q9 (R1) risk-based thinking principles.

Automation Opportunities

- Monthly risk reports auto-generated with metrics and visuals.
- Daily reports during production runs
- Automated escalation alerts for risks trending upward.
- Translation of risk reports for global sites.
- Share risks profiles and lessons learned across entire network

KPIs Impacted

- Better visibility and faster notification of increased risk profiles
- Increased state of control

Supplier and Material Risk Evaluation

QUALITY ASSURED

How It Helps

- Connect to ERP and supplier qualification systems.
- Use RAG to query supplier audit reports, material deviations, and performance data.
- Drafts supplier performance reviews or qualification summaries.
- Reduces time spent on filtering redundant and non-applicable supplier change notifications
- Draft supplier change impact assessments in minutes
- Quickly query current processes to determine impact from supplier changes

Automation Opportunities

- Identify non-conforming materials
  - Release requirements vs COA
- Supplier scorecards auto-generated from quality data.
- Supplier audit scheduling suggestions based on risk ranking.

KPIs Impacted

- Reduced supplier defect rate.
- Audit finding recurrence rate.
- Supplier risk index improvement.

Audit and Inspection Readiness

QUALITY ASSURED

How It Helps

- Use it to simulate inspection Q&A and prepare readiness reports.
- Leverage current inspection trends to find potential audit findings.
- Conflict/Overlap/Gap (COG) analysis of site/enterprise procedure and regulatory requirements
- Flags potential issues faster, so SME time is spent on corrective actions
- Quickly draft CAPAs and document updates before or during an audit
- Draft audit responses in minutes not days: Allows more time to dial in response wording and supporting data

Automation Opportunities

- Mock inspection reports (with auditor notes)
- Auto-flag CAPAs due before inspection dates

KPIs Impacted

- Reduced audit observation rate
- Increased readiness score
- Reduced Audit response time

QMS Authorship (CDMO Case Study)

QUALITY ASSURED

Overview

- A multi-product GMP facility faced production delays due to a growing deviation backlog.
- Client also facing an FDA re-inspection after prior audit findings revealed procedural gaps due to missing site-level SOPs.
- Transitioning from MasterControl to Veeva compounded the documentation challenges.

Solutions Implemented

- AI-Driven Deviation Management
- Custom Impact Assessment Model
- SOP Gap Closure Using Network Documents
- AI Powered Document Reference Updates

KPIs Impacted

- 27.9% reduction in Deviation time to QA
- 30% overall reduction in Deviation cycle time
- 2-3 hrs Drafting time per missing SOP
- 160 Deviations closed in 8 weeks
- **Site Passed PAI Audit!!!**

Thank You!

Questions?

Contact Info:
Chris Dayton

QUALITY ASSURED

Chris.Dayton@QualityAssured.AI