

Designing the AI-Enabled CMC Ecosystem

For Drug Development and Manufacturing in a 10× Era



Pablo Rolandi* (PhD)

Executive Director**

Transformative Digital Capabilities • Process Development • Operations

* This presentation was produced by **the author**, augmented by **ChatGPT** tokens through **26** prompts and **4953** seconds of inference time (and **163** references).

** At Amgen, Pablo is a power user of **ChatGPT** Enterprise & MS **Copilot**.

** At home, Pablo is an avid user of **Claude** & **Gemini**, and an occasional user of **DeepSeek**, **Grok**, **Manus**, & **Perplexity**.

The Amgen logo, consisting of the word "AMGEN" in a bold, white, sans-serif font. The letters are closely spaced and have a slight shadow effect against the blue background.

AMGEN

Perception vs Reality

Calibrating mental models on LLMs and foundation models

Stochastic parrot

- ✓ In-context learner ¹
(approximating Bayesian updating)

Memorization oracle

- ✓ Emergent reasoning ²
(language mastery and advanced reasoning patterns)

Search engine

- ✓ Thought partner / Research assistant ³
(with tools; eg., “deep research”)

Question-answer bot

- ✓ Stepwise reasoning ⁴
(prompt scaffolding / chain-of-thought)

Text-only inputs

- ✓ Multi-modal reasoning
(text + images + data)

Uncontrollable hallucinations

- ✓ Grounded, verifiable responses ⁵
context + retrieval (RAG) to reduce hallucinations

1) <https://arxiv.org/pdf/2402.03175>; 2) <https://arxiv.org/pdf/2501.12948>; 3) <https://arxiv.org/pdf/2506.12594>

4) <https://arxiv.org/pdf/2201.11903>; 5) <https://arxiv.org/pdf/2005.11401>

Guiding Principles

Lens for Case Studies



Good AI practice in drug development

1) Human-centric by design

Use aligns with ethical and human-centric values.

2) Risk-based approach

Validation and controls are proportional to risk, based on context of use.

3) Adherence to standards

Adhere to all relevant standards, including technical, scientific, and regulatory.

4) Clear context of use

Have a well-defined context of use.

5) Multidisciplinary expertise

Integrate multidisciplinary expertise across the life cycle.

6) Data governance & documentation

Ensure data provenance, traceability, privacy, and documentation.

7) Model design & development practices

Use fit-for-purpose data and best practices to promote "trustworthiness".

8) Risk-based performance assessment

Performance evaluation of the complete (human-AI) system is proportional to risk.

9) Life cycle management

Adopt risk-based quality management across the life cycle.

10) Clear, essential information

Communicate to the intended audience all relevant information in plain language.

Case Study #1

AI4x Suite – Grounded Knowledge Work at Scale

AI4x Suite: consists of multiple LLM-based systems that tap into >2.0m controlled documents and >1.5m curated experiments in the Amgen's corporate knowledge base.

AI solution characteristics

RAG¹ grounded responses →

reduced hallucinations, full traceability... (hyperlinks to references).

Not a “vanilla” ChatGPT →

multimodal reasoning, agents, and various optimizations increase accuracy.

FAIR² data + automated data threads →

heavy use of metadata, daily data refresh for scalability.

FDA principles anchored

7) *Model design & development practices* → “promotes transparency, reliability, (...) and robustness”

5) *Multidisciplinary expertise* → “integrated throughout the technology’s life cycle”

6) *Data governance & documentation* → “[data provenance, processing are traceable and verifiable]”

1) RAG: retrieval-augmented generation; <https://arxiv.org/pdf/2005.11401>

2) FAIR: findable, accessible, interoperable, reusable; <https://www.nature.com/articles/sdata201618.pdf>

Case Study #2

Virtual Clone – Automated AI-Enabled Clone Selection

Virtual Clone: performs 10% better than SMEs on a subset of drug modalities.

AI solution characteristics

FAIR¹ data + automated data threads → automated data flows were essential to success...

(previous unsuccessful attempts had manual workflow steps).

Multi-modal² feature set →

images, categorical data, time series, and analytical measurements contribute to prediction accuracy.

Rigorous model qualification →

governed by two SOPs: 1) AI SOP (based on NIST AI RMF³), 2) ML SOP (internal extension of ASME VVUQ 40 standard⁴).

FDA principles anchored

6) *Data governance & documentation* → “data (...) processing steps, and analytical decisions are documented in a detailed (...) manner”

7) *Model design & development practices* → “best practices in model and system design and software engineering and leverages data that is fit-for-use”

2) *Risk-based approach* / 8) *Risk-based performance assessment*

3) *Adherence to standards* → “AI technologies adhere to relevant (...) technical, scientific (...) standards”

1) FAIR: findable, accessible, interoperable, reusable; <https://www.nature.com/articles/sdata201618.pdf>

2) <https://arxiv.org/pdf/2309.10020>

3) <https://www.nist.gov/itl/ai-risk-management-framework>

4) <https://www.asme.org/codes-standards/publications-information/verification-validation-uncertainty>

What We Should Debate Next



A) Determinism is the wrong lens

Manage **variability statistically** (ensembles, repeated runs) and **operationally** (observability, monitoring).



B) Interpretability over explainability

Interpretability (why) is more important than **explainability** (how) → inspect AI's **reasoning traces** and **decision rubrics**.



C) FAIR data is limiting AI value

Trustworthy AI operates on **FAIR data**, which faces significant cultural, infrastructure, and technology challenges.



D) Embrace context engineering

Prompt engineering is obsolete. Effective solutions implement **context engineering** and **systems thinking**.



E) Benchmarks for our industry

Open benchmarks for our **specific domain** are needed → *FrontierScience* for expert scientific reasoning tasks as a source of inspiration.

Conclusion

Takeaway Messages

AI in CMC is an **engineered system**, not a model

There is “*no silver bullet*”, “*no free lunch*”, and “*the bitter lesson*” still applies.

Context engineering and **systems thinking** (plus AI **organizational agility**) will distinguish *winners* from *followers*.

1

Trustworthy AI is achievable **now**

Grounding (RAG), “**evals**” (evaluation harnesses), **lifecycle** management, and **change controls** work *today*.

Organizational inertia and **FAIR data** threads are *the limiting step*.

2

Move from pilots to **intelligence at scale**

ROI scales when *strategy, data, people, process, and technology* turn **point tools** into durable **integrated solutions**.

3