Digital Tools You Wish You Implemented Earlier

Strategic Tech Enablers for CapEx

Cedric Berger, PhD, MBA

Date

4th Basel CAPEX SUMMIT - Innovations in Capex Project Delivery in Pharma & Biotech

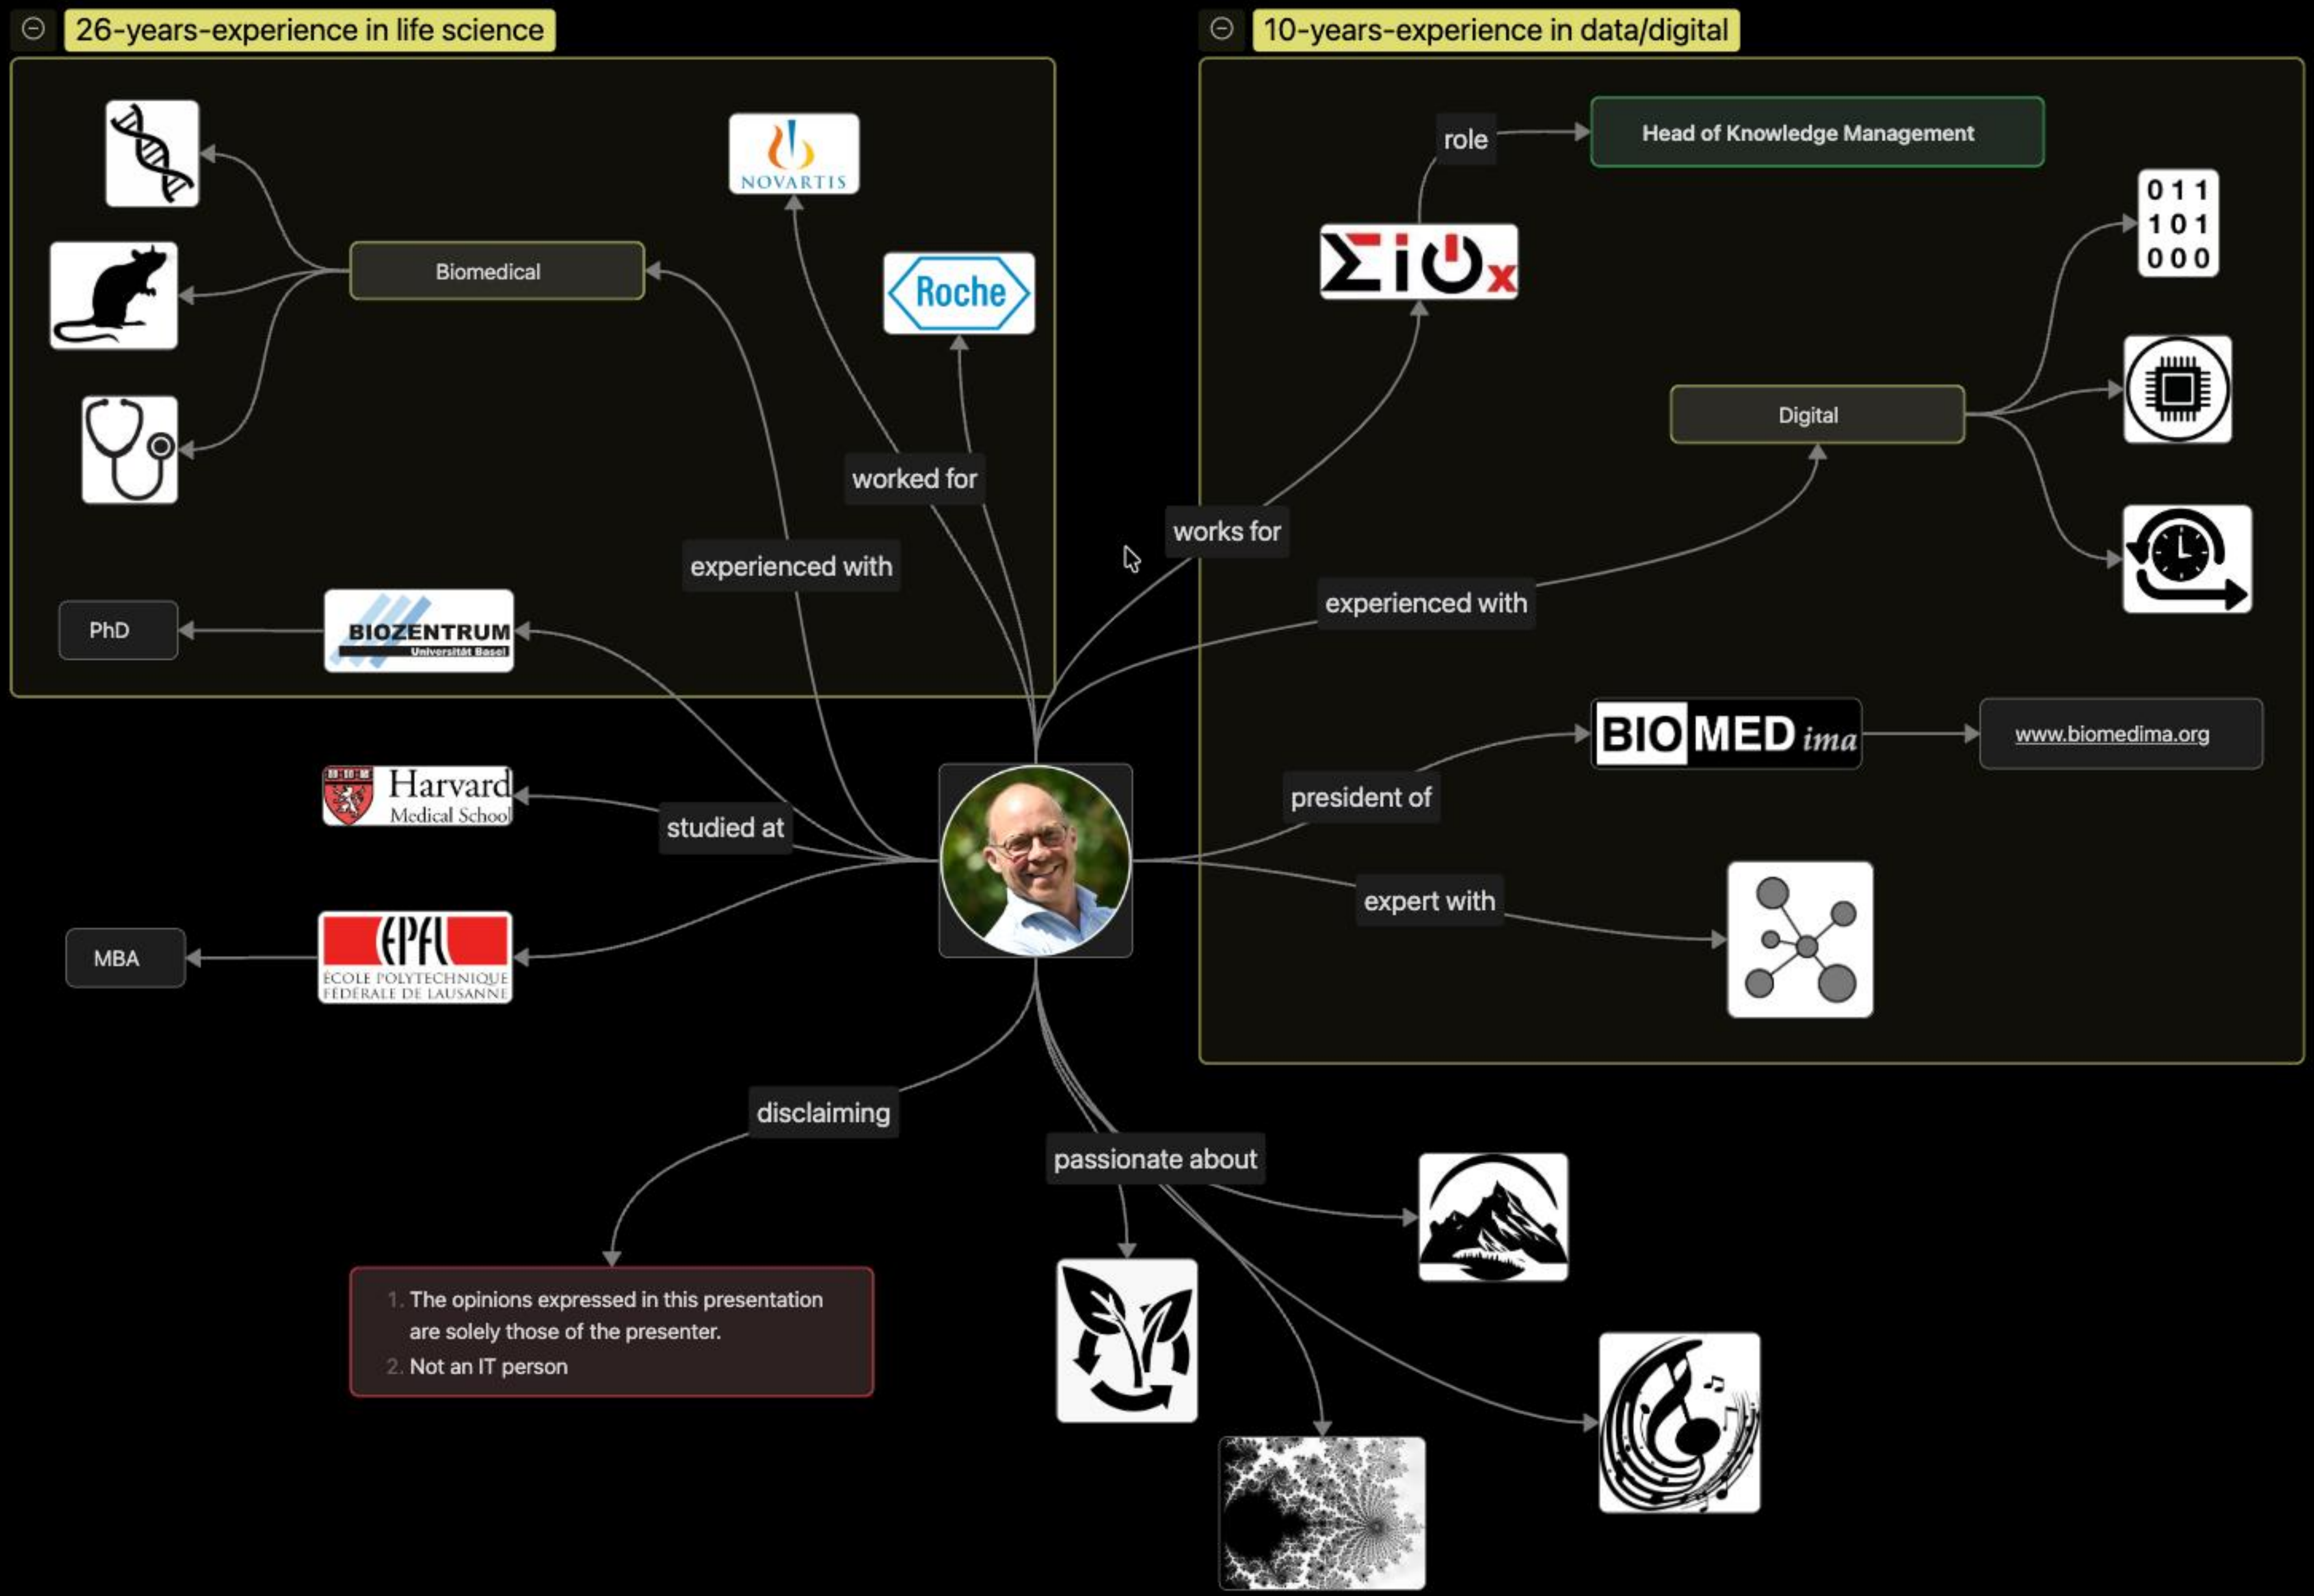
26-years-experience in life science
Biomedical
NOVARTIS
Roche
worked for
experienced with
PhD
BIOZENTRUM
Universität Basel
Harvard
Medical School
studied at
MBA
EPFL
ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE
10-years-experience in data/digital
role
Head of Knowledge Management
011
101
000
Digital
works for
experienced with
BIO MED ima
www.biomedima.org
president of
expert with
disclaiming
1. The opinions expressed in this presentation are solely those of the presenter.
2. Not an IT person
passionate about

The Age of Digital Complexity

The Explosion of Digital Complexity

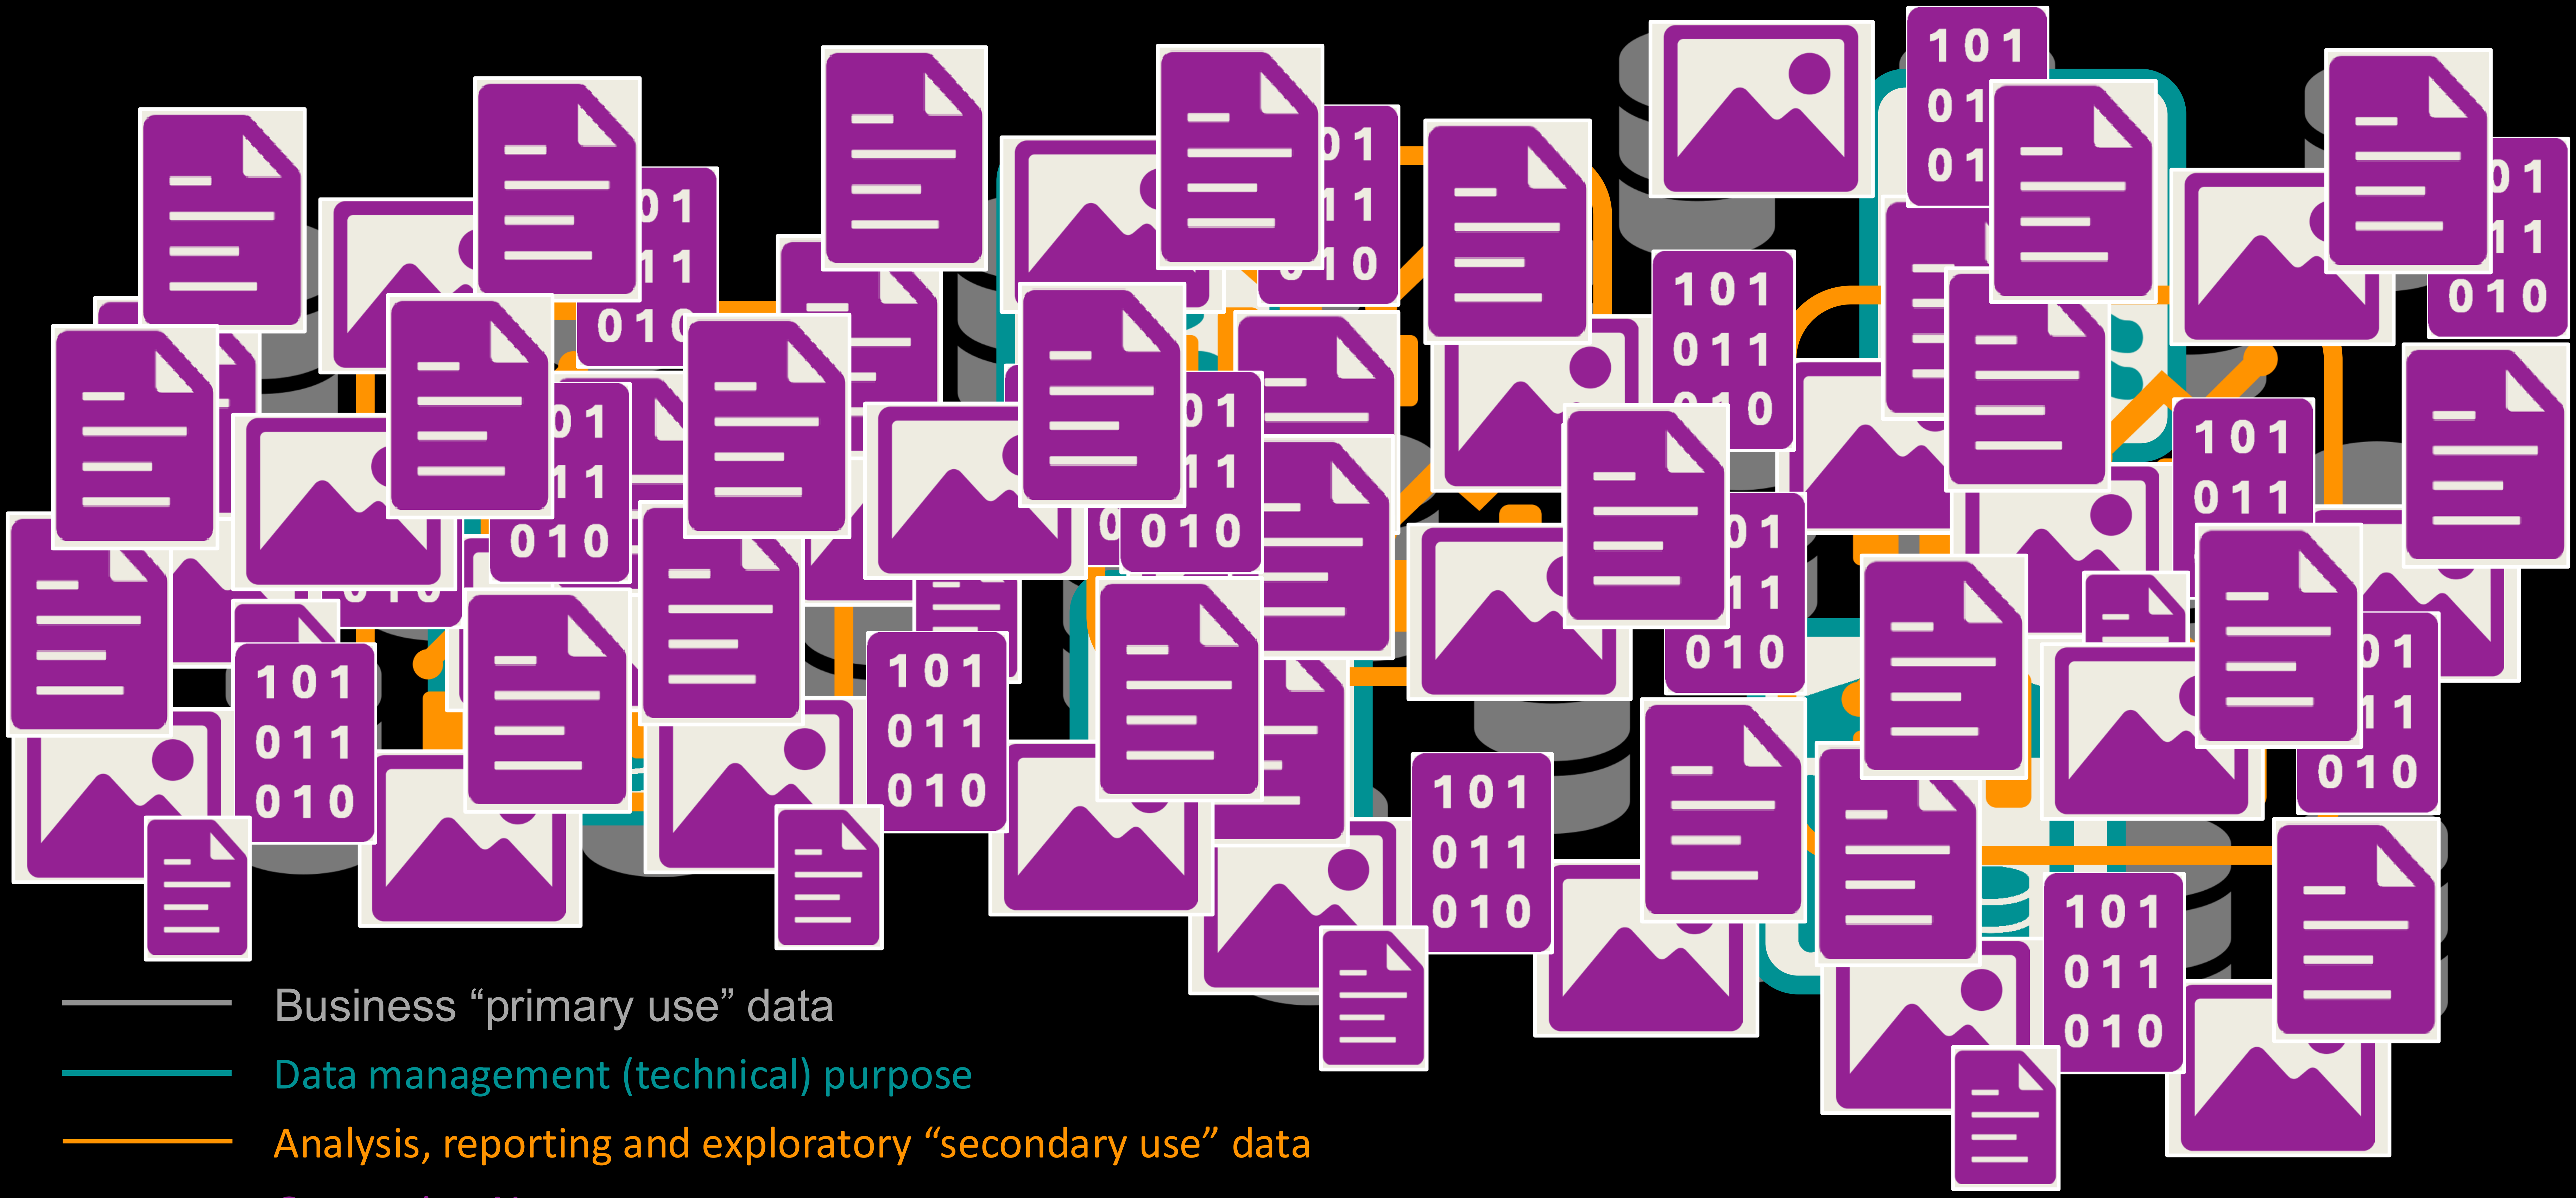

The proliferation of IT Systems…

… leads to the proliferation of siloed, disintegrated, contextless data

2012

2018

2024

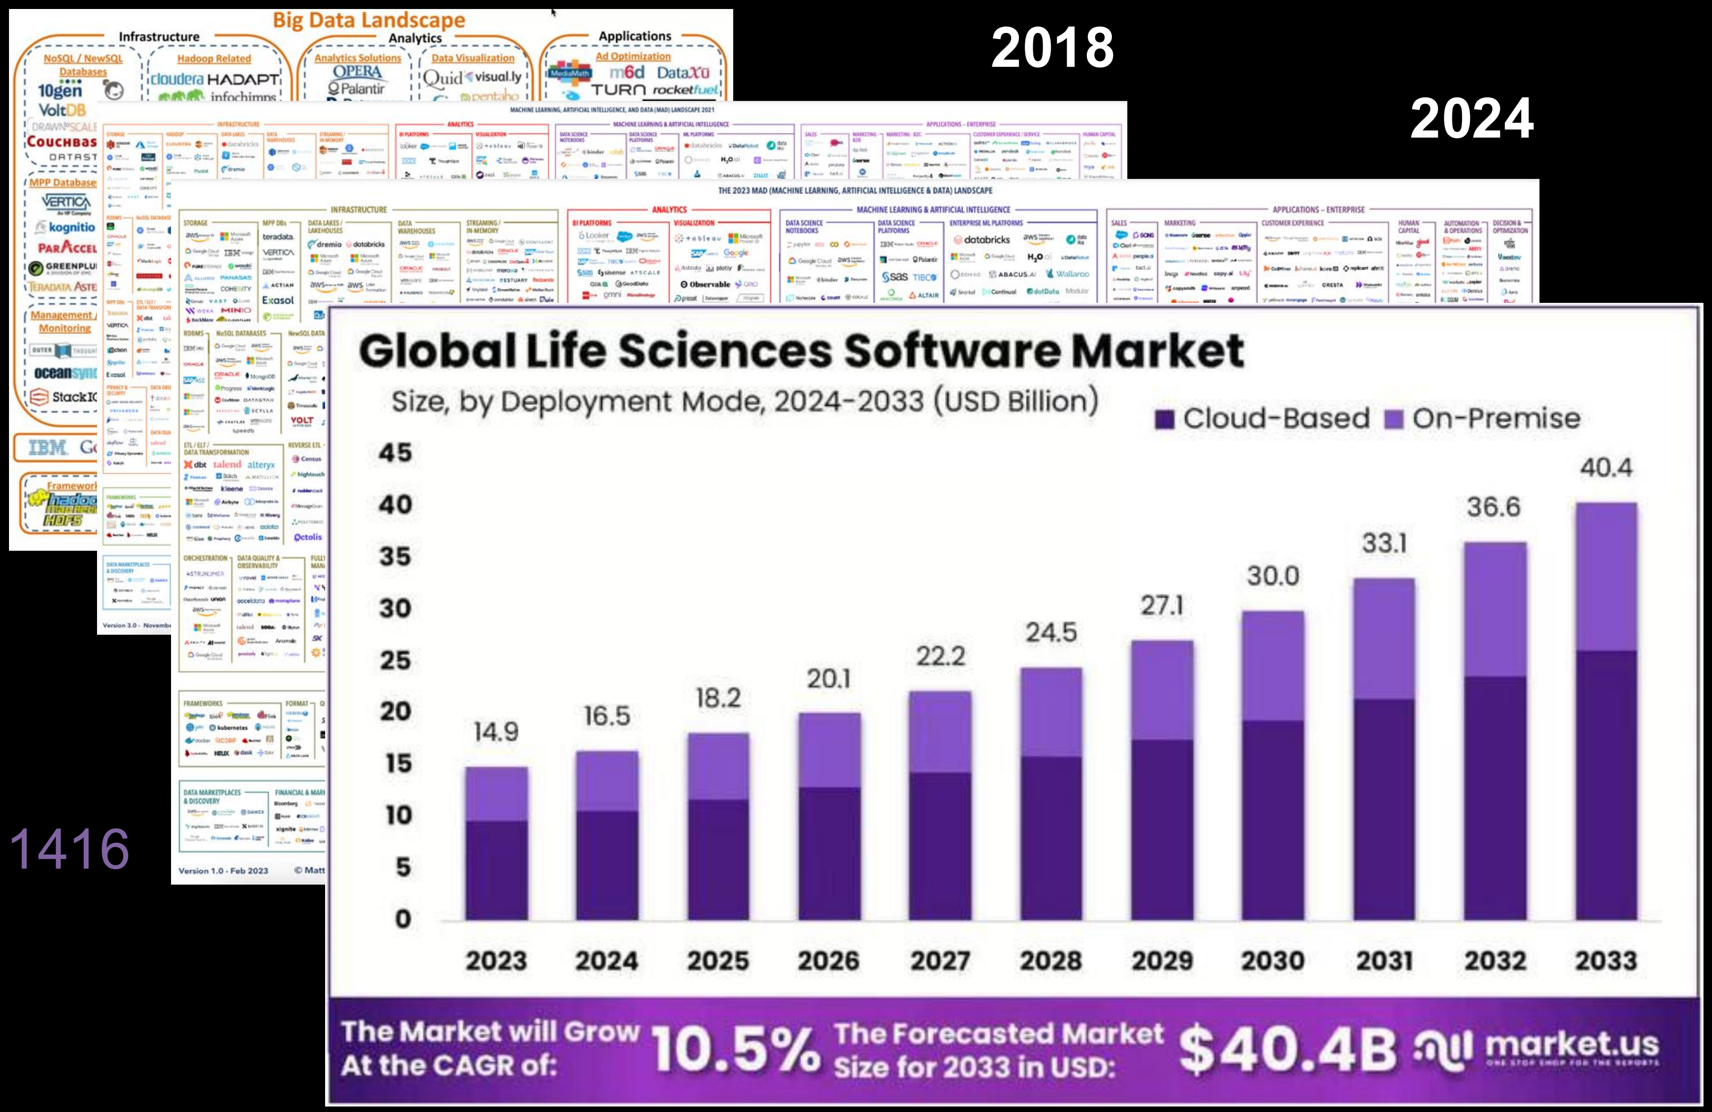

139

1416

https://mattturck.com

https://market.us/report/life-sciences-software-market/

Depending on definitions of "big" and "IT system", big pharma organizations *

- operate anywhere from 500 to over 6'000 IT systems
- handling 100 petabytes to several exabytes annually

100 petabytes = 25 million DVDs
1 exabyte = 1000 petabytes = 250 million DVDs

* Estimated from multiple sources: https://www.pharma-iq.com/manufacturing/articles/navigating-ai-integration-in-pharmas-legacy-systems; https://www2.deloitte.com/us/en/insights/industry/life-sciences/biopharma-digital-transformation.html; https://moldstud.com/articles/p-exploring-systems-engineering-practices-in-the-pharmaceutical-industry; https://www.pharmalex.com/thought-leadership/blogs/key-considerations-for-data-strategies-in-the-pharmaceutical-industry/

Lost in Digital Complexity

Business Expert

“Our patients' data spans over 3000 tables”

“HER systems give me incomplete, outdate, discrepant data about the same patients?”

Knowledge Expert

“Then, you will never know what is going on with your patients”

“These systems haven’t been designed to communicate together and exchange data.”

It is likely that nobody understands end-to-end business anymore

More about Digital Complexity

https://www.biomedima.org

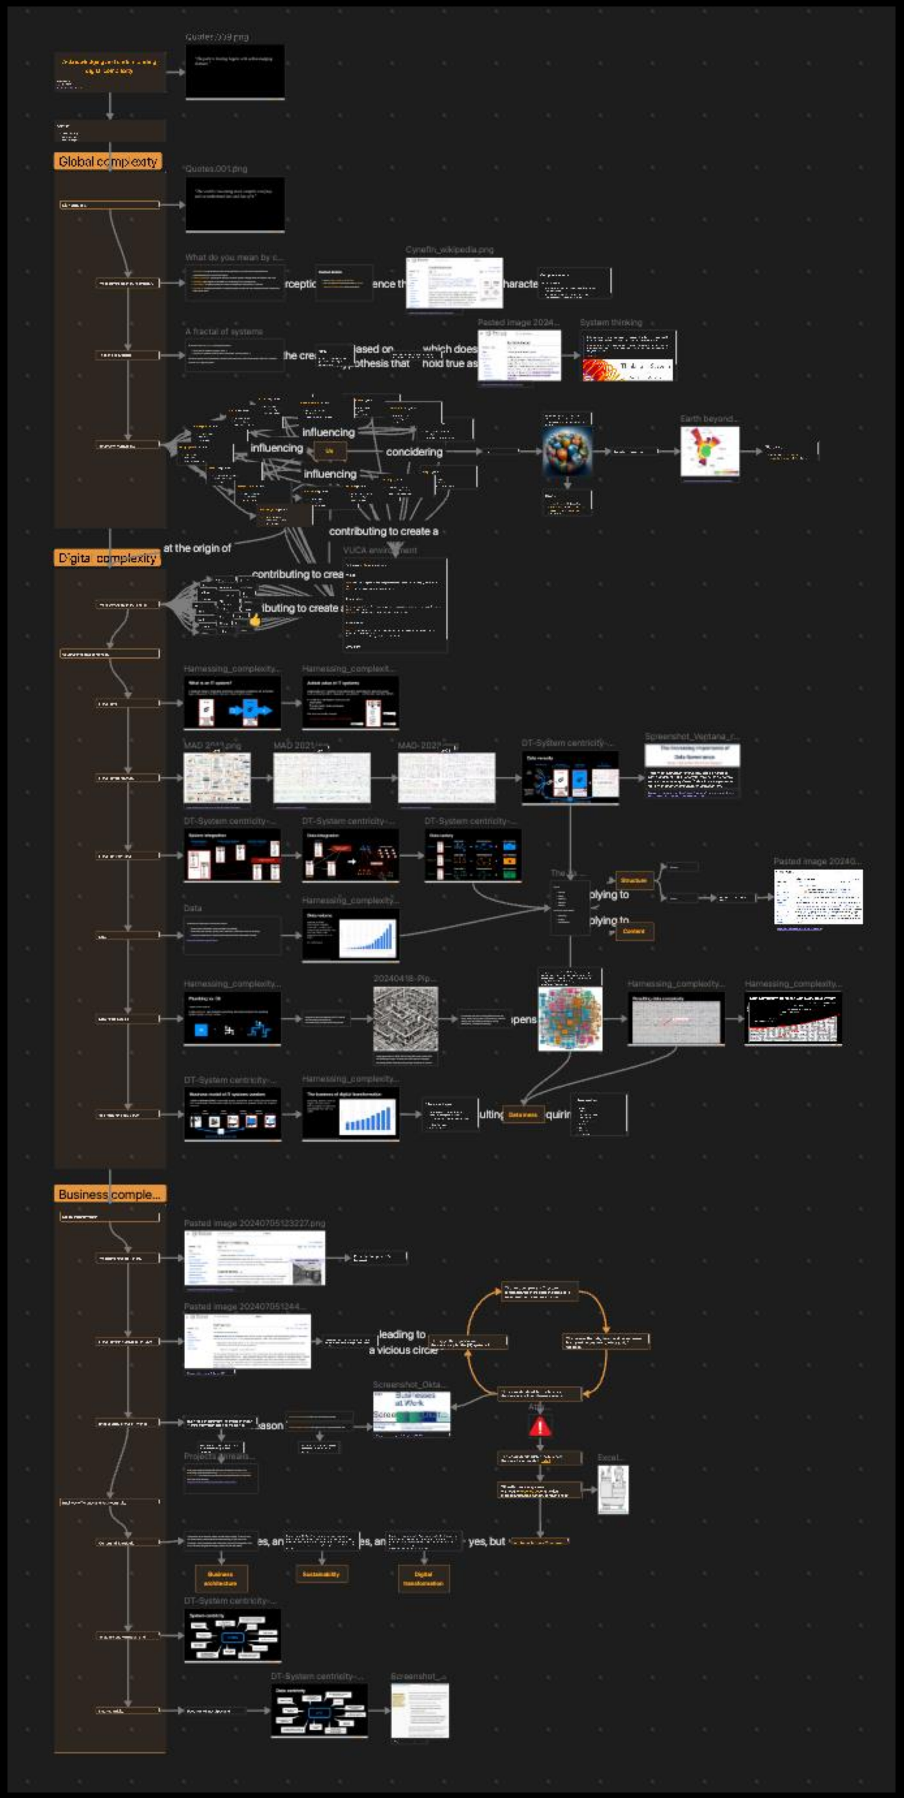

https://www.biomedima.org/project/acknowledging-and-understanding-digital-complexity/

Need #1: “Know thyself”

- What is your organization unique added value?
- In which current project are you using it and how?
- How do you quantitatively measure value?
- Who are the key contributing people using what skills?
- What are the key data sets supporting them?
- Where and how do you capture, manage and report on the above items?

Need: Diagnostic tools to capture a realistic, trustful, granular-enough and regular health snapshots of your company.

Risk of not doing it: Digital complexity will prevent you to have a good understanding of your organization which will prevent optimal decision-making.

Human

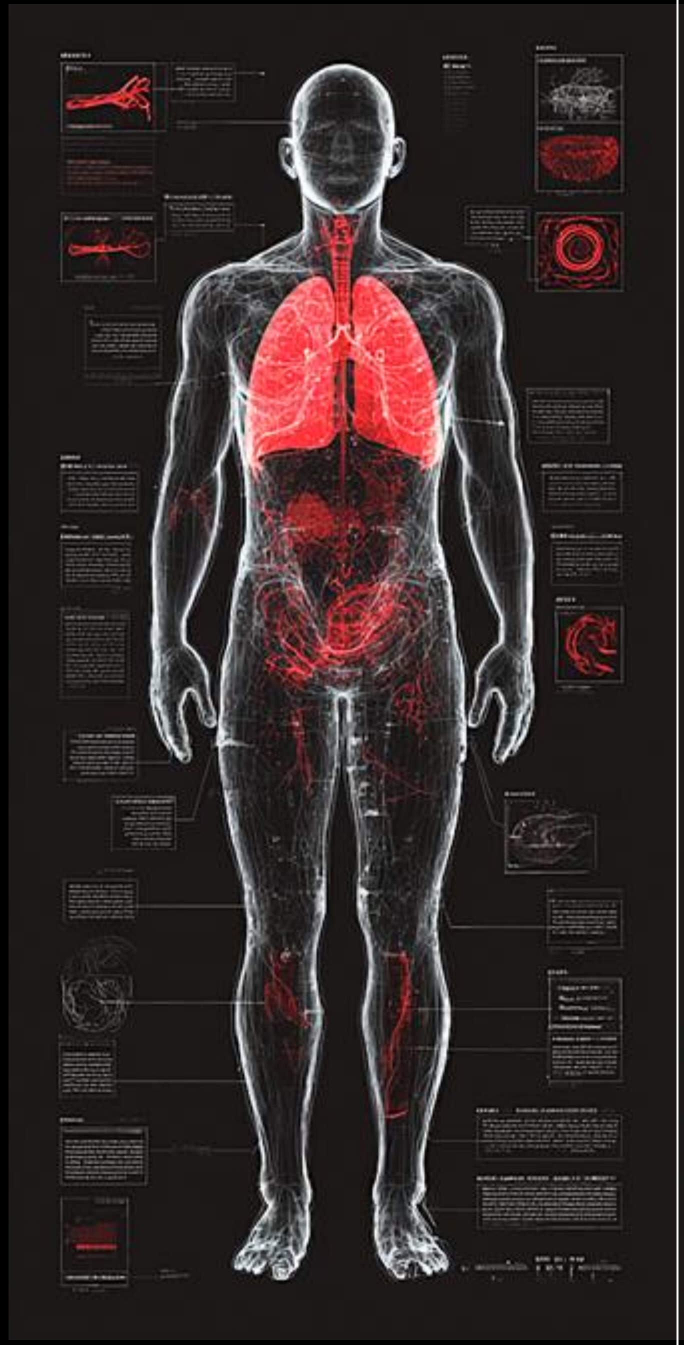

Enterprise

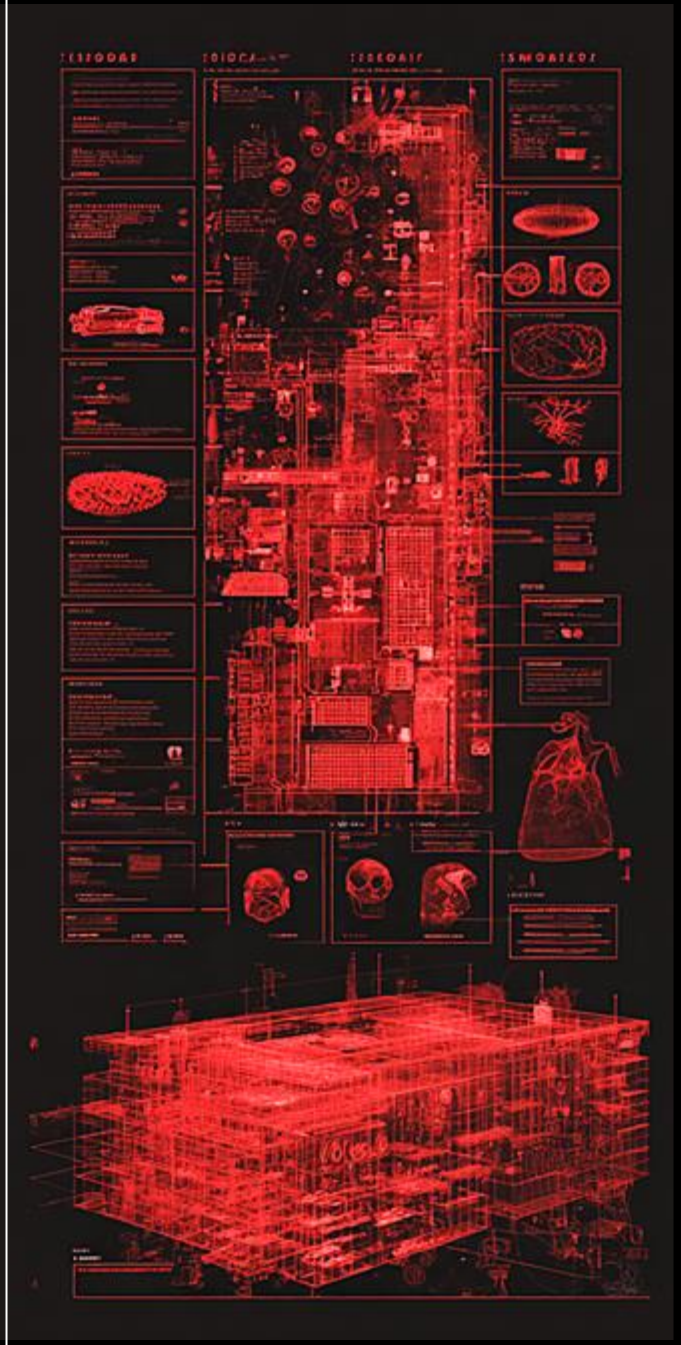

Chasing AI Value

Promises Only Bind People Believing in Them

MIT study shatters AI hype: 95% of generative AI projects are failing, sparking tech bubble jitters

By Paurush Omar, ET Online • Last Updated: Sep 09, 2025, 06:23:00 PM IST

Not Everyone will Benefit from AI

The Hype

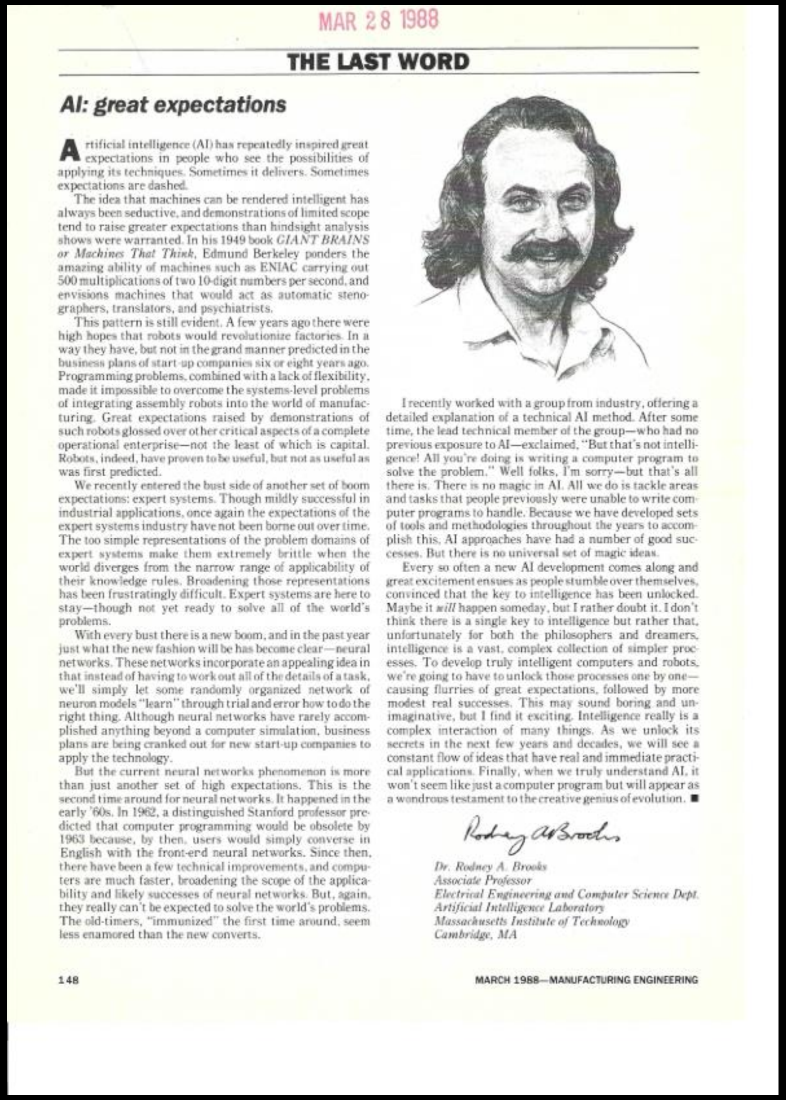

MAR 28 1988

THE LAST WORD

AI: great expectations

Artificial intelligence (AI) has repeatedly inspired great expectations in people who see the possibilities of applying its techniques. Sometimes it delivers. Sometimes expectations are dashed.

The idea that machines can be rendered intelligent has always been seductive, and demonstrations of limited scope tend to raise greater expectations than hindsight analysis shows were warranted. In his 1949 book *GIANT BRAINS or Machines That Think*, Edmund Berkeley ponders the amazing ability of machines such as ENIAC carrying out 500 multiplications of two 10-digit numbers per second, and envisions machines that would act as automatic stenographers, translators, and psychiatrists.

This pattern is still evident. A few years ago there were high hopes that robots would revolutionize factories. In a way they have, but not in the grand manner predicted in the business plans of start-up companies six or eight years ago. Programming problems, combined with a lack of flexibility, made it impossible to overcome the systems-level problems of integrating assembly robots into the world of manufacturing. Great expectations raised by demonstrations of such robots glossed over other critical aspects of a complete operational enterprise—not the least of which is capital. Robots, indeed, have proven to be useful, but not as useful as was first predicted.

We recently entered the bust side of another set of boom expectations: expert systems. Though mildly successful in industrial applications, once again the expectations of the expert systems industry have not been borne out over time. The too simple representations of the problem domains of expert systems make them extremely brittle when the world diverges from the narrow range of applicability of their knowledge rules. Broadening those representations has been frustratingly difficult. Expert systems are here to stay—though not yet ready to solve all of the world's problems.

With every bust there is a new boom, and in the past year just what the new fashion will be has become clear—neural networks. These networks incorporate an appealing idea in that instead of having to work out all of the details of a task, we'll simply let some randomly organized network of neuron models "learn" through trial and error how to do the right thing. Although neural networks have rarely accomplished anything beyond a computer simulation, business plans are being cranked out for new start-up companies to apply the technology.

But the current neural networks phenomenon is more than just another set of high expectations. This is the second time around for neural networks. It happened in the early '60s. In 1962, a distinguished Stanford professor predicted that computer programming would be obsolete by 1963 because, by then, users would simply converse in English with the front-end neural networks. Since then, there have been a few technical improvements, and computers are much faster, broadening the scope of the applicability and likely successes of neural networks. But, again, they really can't be expected to solve the world's problems. The old-timers, "immunized" the first time around, seem less enamored than the new converts.

I recently worked with a group from industry, offering a detailed explanation of a technical AI method. After some time, the lead technical member of the group—who had no previous exposure to AI—exclaimed, "But that's not intelligence! All you're doing is writing a computer program to solve the problem." Well folks, I'm sorry—but that's all there is. There is no magic in AI. All we do is tackle areas and tasks that people previously were unable to write computer programs to handle. Because we have developed sets of tools and methodologies throughout the years to accomplish this, AI approaches have had a number of good successes. But there is no universal set of magic ideas.

Every so often a new AI development comes along and great excitement ensues as people stumble over themselves, convinced that the key to intelligence has been unlocked. Maybe it *will* happen someday, but I rather doubt it. I don't think there is a single key to intelligence but rather that, unfortunately for both the philosophers and dreamers, intelligence is a vast, complex collection of simpler processes. To develop truly intelligent computers and robots, we're going to have to unlock those processes one by one—causing flurries of great expectations, followed by more modest real successes. This may sound boring and unimaginative, but I find it exciting. Intelligence really is a complex interaction of many things. As we unlock its secrets in the next few years and decades, we will see a constant flow of ideas that have real and immediate practical applications. Finally, when we truly understand AI, it won't seem like just a computer program but will appear as a wondrous testament to the creative genius of evolution. ■

Rodney A Brooks

Dr. Rodney A. Brooks
Associate Professor
Electrical Engineering and Computer Science Dept.
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA

148 MARCH 1988—MANUFACTURING ENGINEERING

Technology-centric mindset

Business Encoding

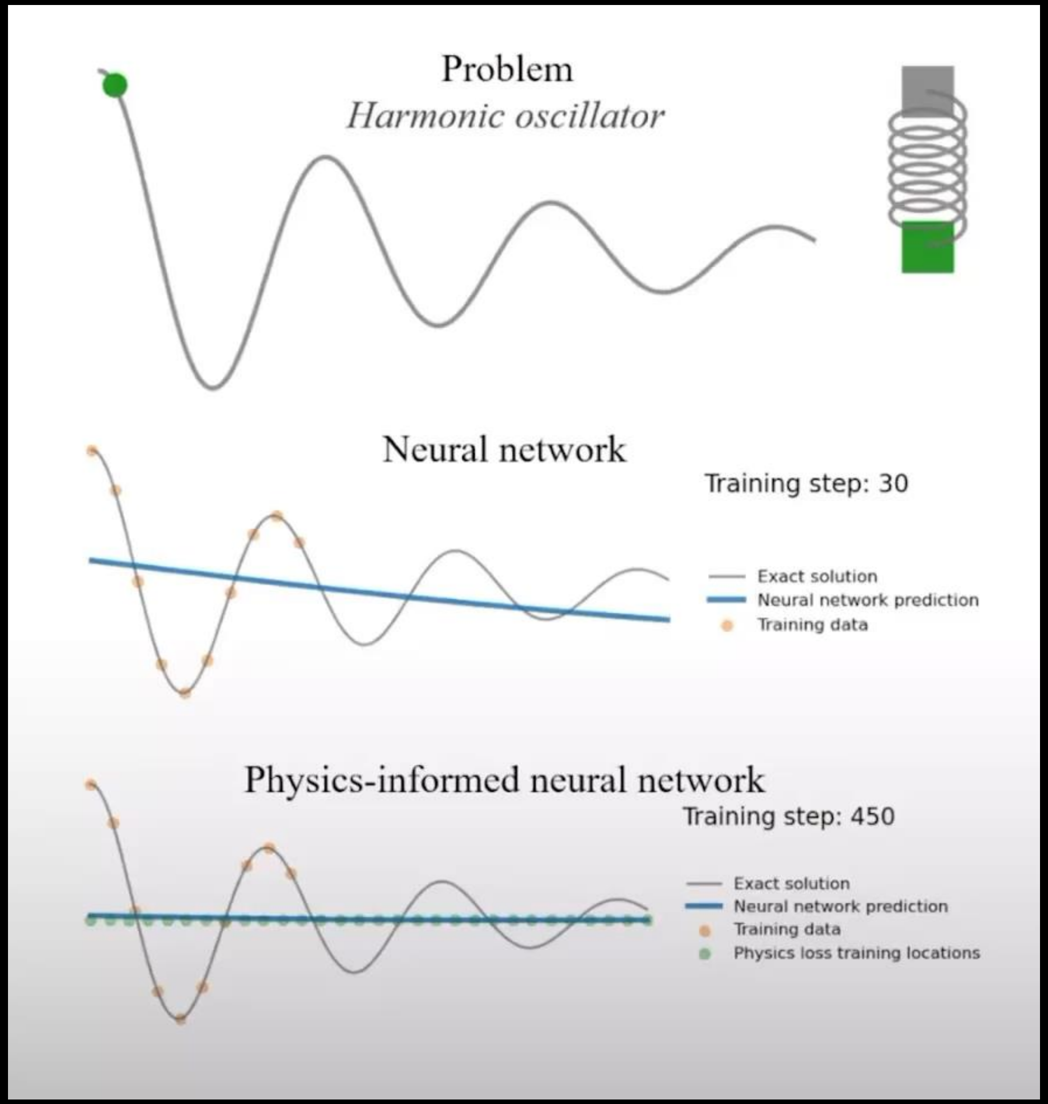

Data/Context Readiness

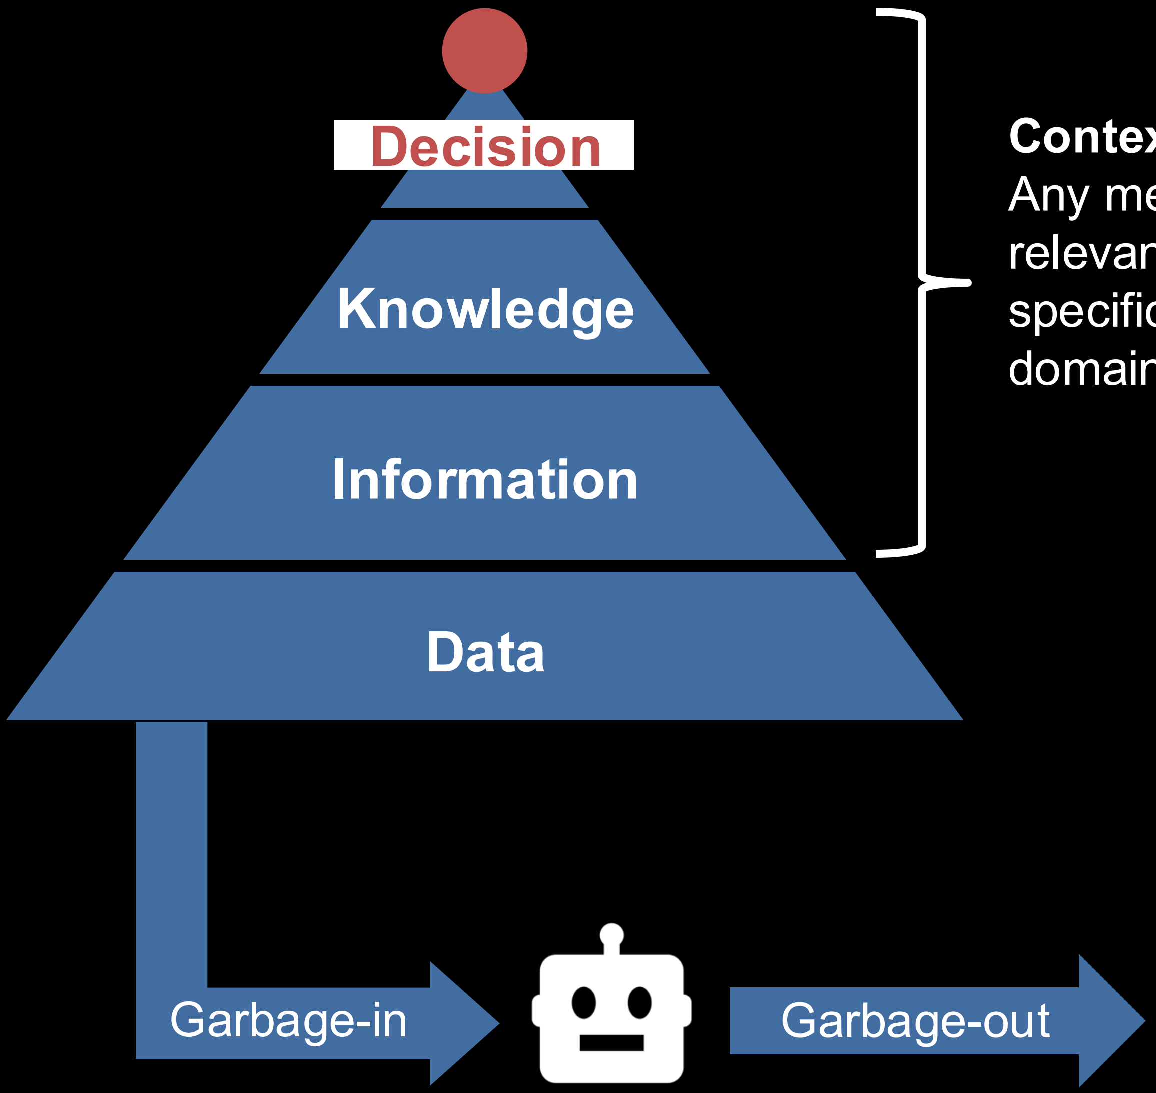

Context:
Any metadata relevant to one specific domain/use case

https://www.biomedima.org/project/messianism-in-ai-times-keynote/

https://en.wikipedia.org/wiki/DIKW_pyramid

“An IT system is as smart as the data available to it”

“Garbage In, Garbage Out”

<u>Facts</u>

- Models trained on inconsistent data (poor features preventing to identify patterns) = wrong predictions.
- Data issues discovered late multiply cost 10× (unplanned, single-shot, siloed, ungoverned)
- Reliable data shortens AI cycle time by 30–50 %.

<u>Common executive reaction:</u>

“Let’s buy a new tool !”

<u>Digital data/knowledge expert:</u>

“This is not about technology.
This is about 1) people/mindset, 2) processes and 3) (meta)data”

Image from LinkedIn; unknow author

Before any AI Strategy

More important is the data…

- Raw data → Features → Model → Deployment → Decisions
- Data quality affects every step
- Poor data = bias, drift, wasted spend
- 60–80 % of AI projects time spent cleaning data

…and metadata measuring its value/ROI

- No unified value-measurement framework
- Duplicate KPIs and conflicting over-complex dashboards
- Compliance and trust issues from unclear lineage.

No need to clean all legacy data !

- Legacy data: clean/integrate only what you need but with a consistent methodology (repeated iteratively)
- New data: apply the methodology systematically embedded in the business processes

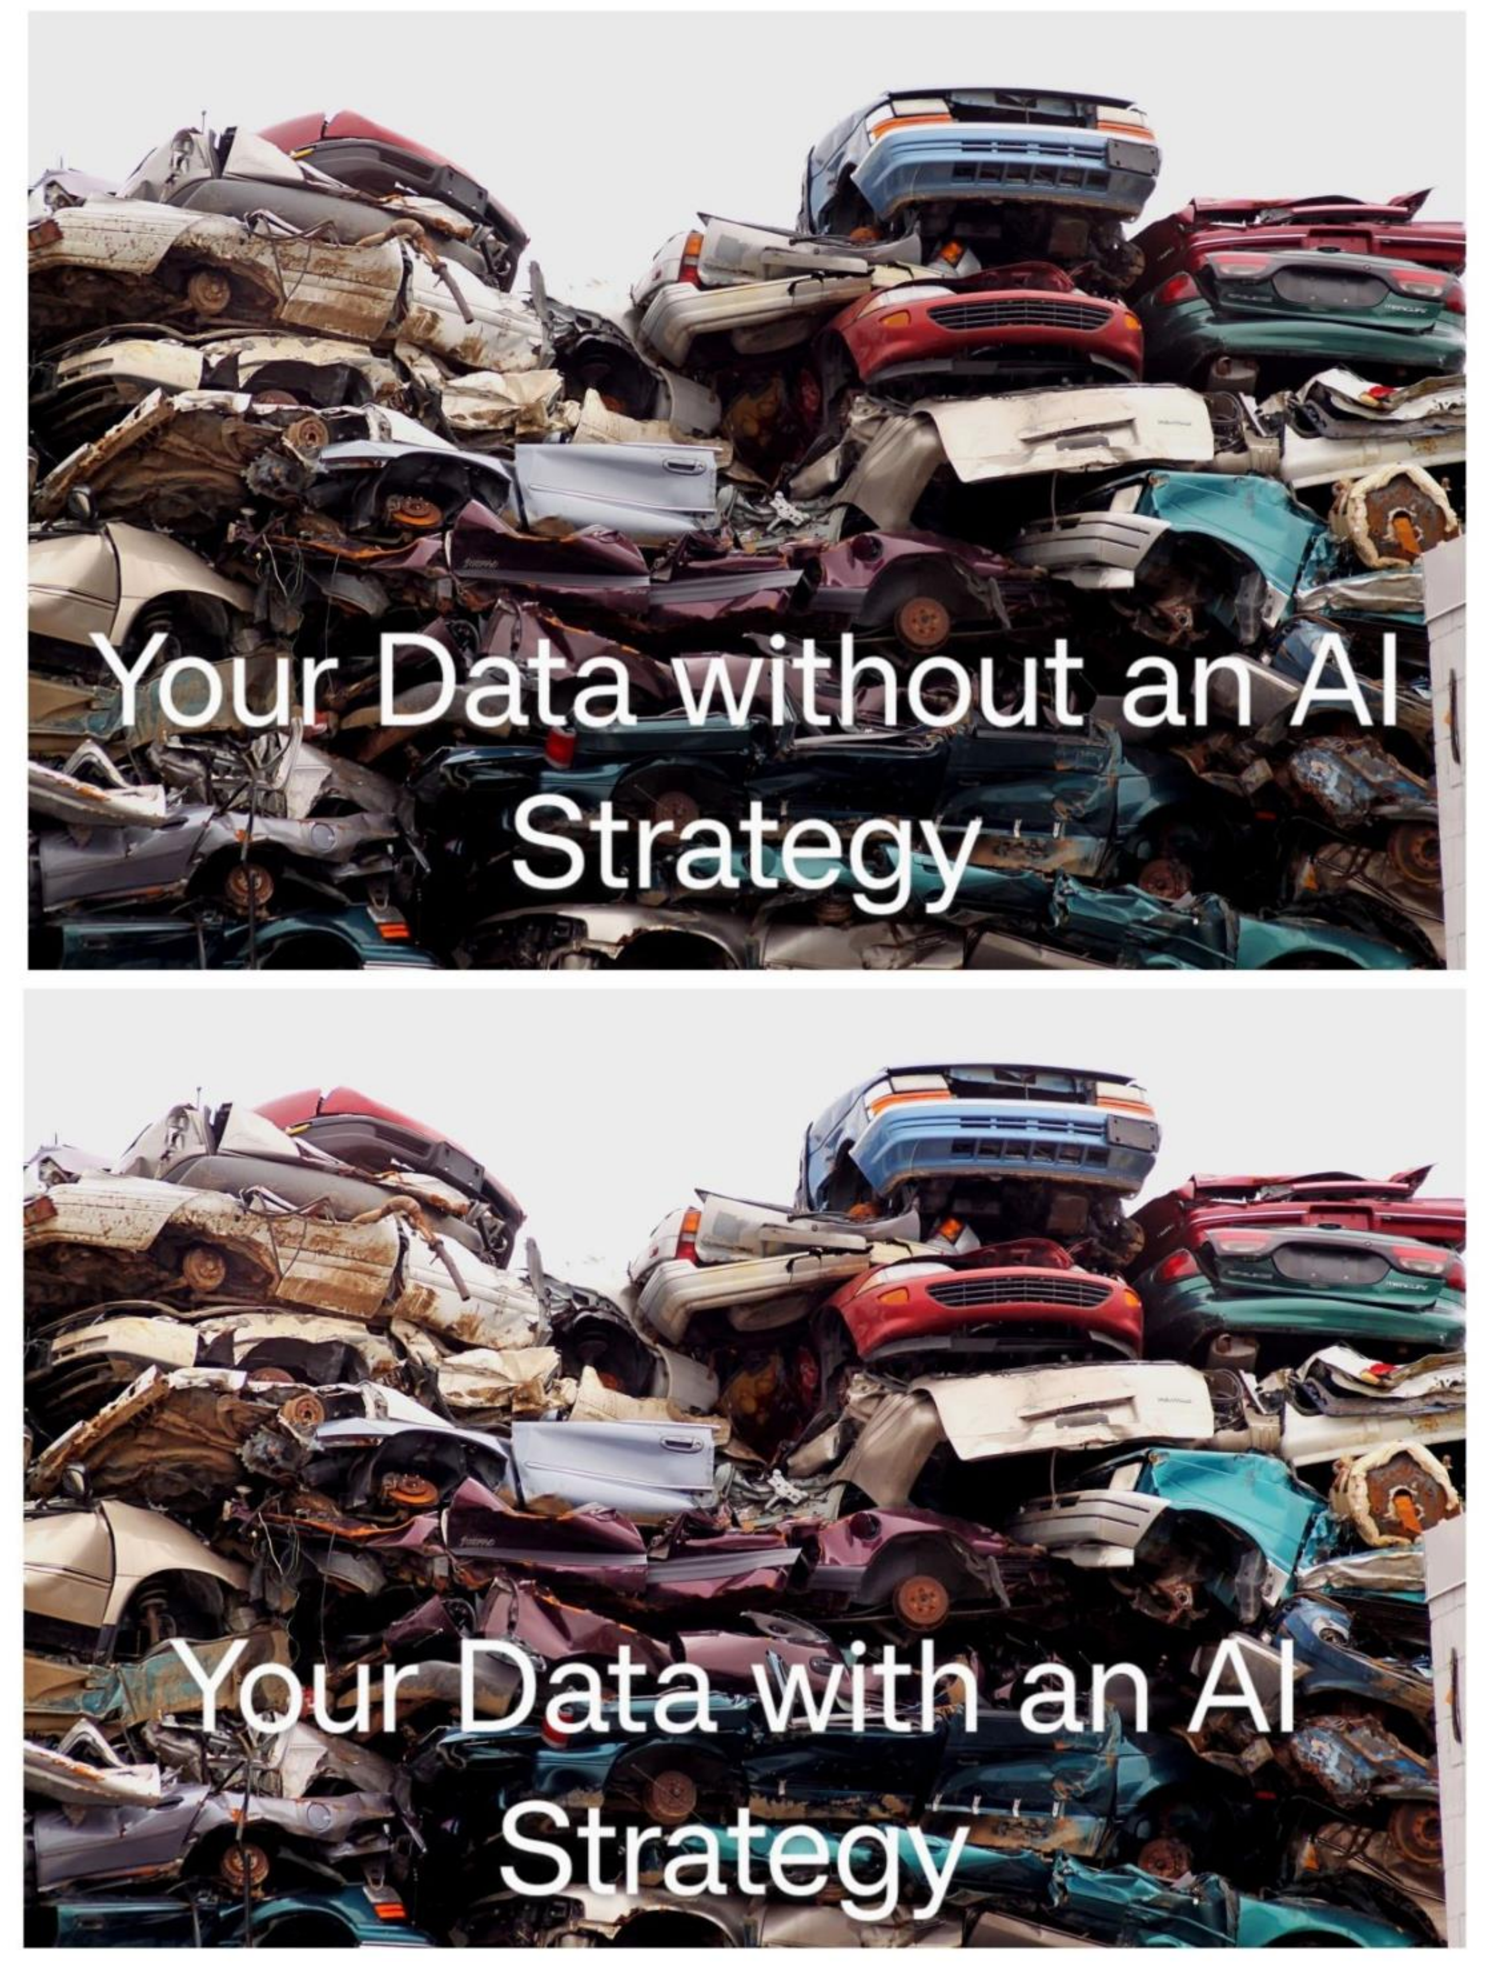

Need #2: Context = Connections

- How do you connect value measurement to projects? to contributors? to data sets? to data set quality?
- How do you connect structured (tabular) data with unstructured (text) data i.e. documents?
- How do you disambiguate and reconcile similar or identical entities across data sets?
- How do you integrate all the above to feed AI so that it brings value?

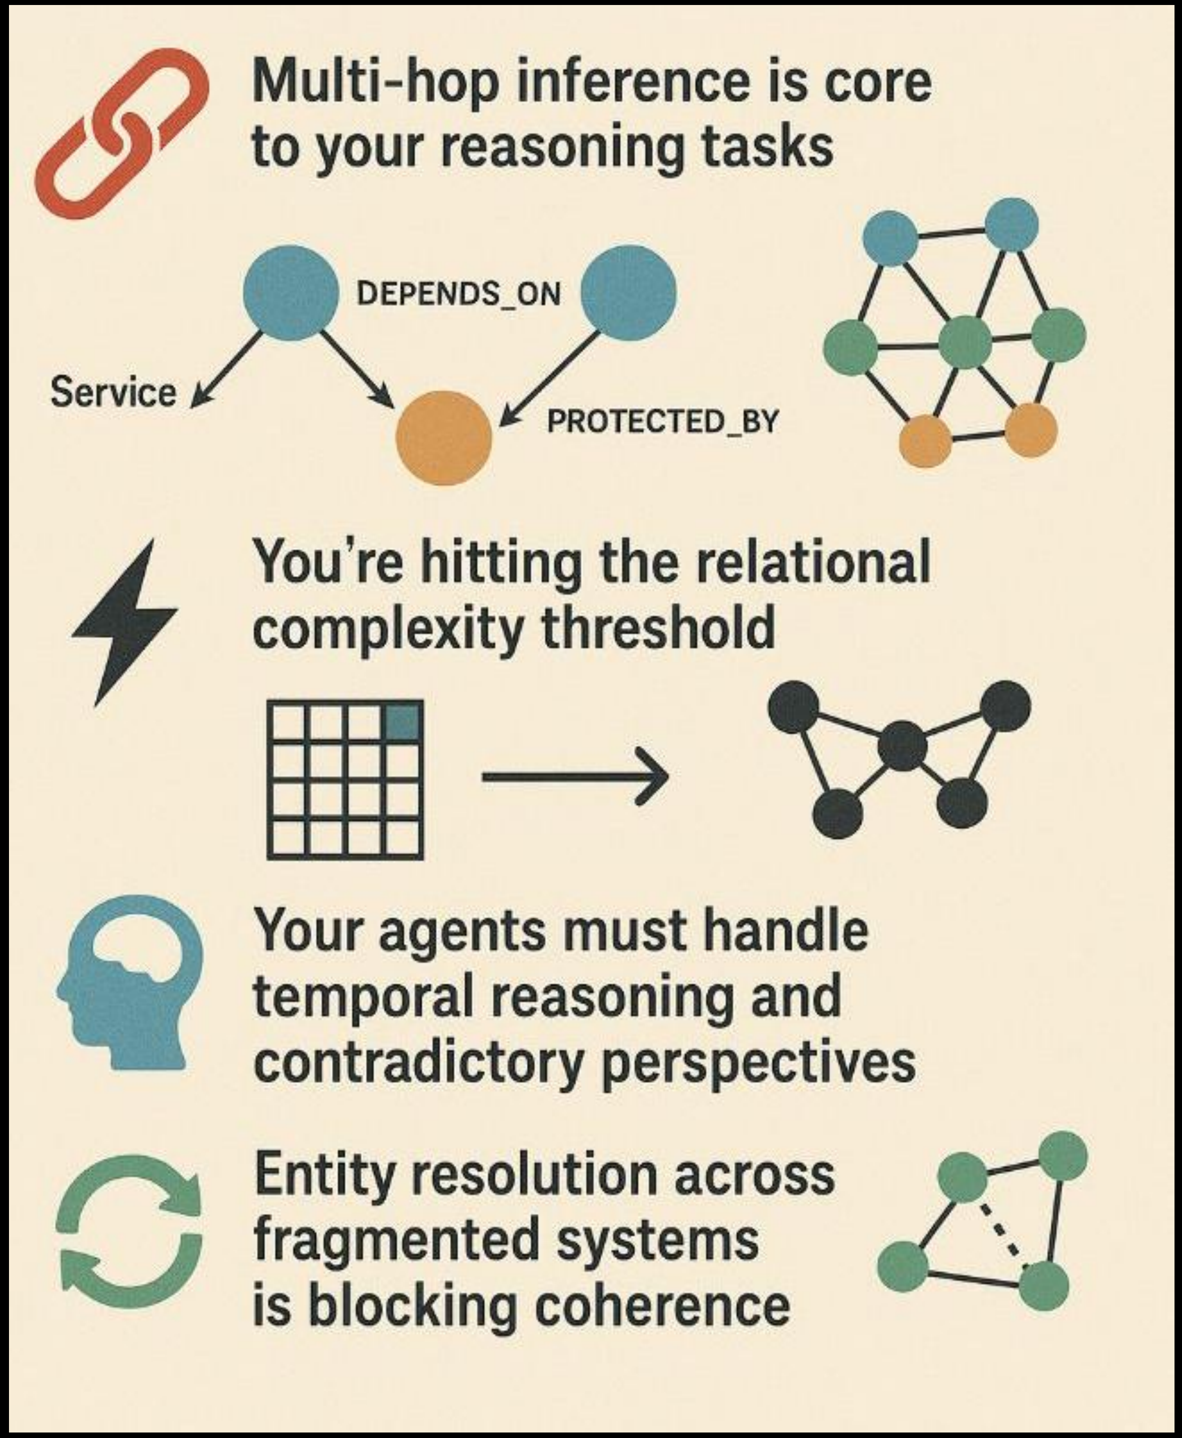

Image from LinkedIn; unknow author

Need: A data management tool that integrates/connects data and metadata (across systems and across data formats) to provide reliable context to AI

Risk of not doing it: Not getting (full) value out of ML/AI, especially generative AI that mostly relies on context.

DaaA

(Data as an Asset)

From OpEx to CapEx Thinking

Legacy OpEx thinking about data

- One-off data cleaning for individual projects
- Siloed (no share and reuse)
- End-of-process, quick and dirty fix
- Out of data governance scope
- Optional standards

Modern Capex thinking about data

- Durable data assets (catalogs, master data, lineage)
- Enables long-term productivity

Why?

- Curated training sets are intellectual capital
- Shared, governed data enables reuse across AI products
- Governance = faster AI iteration and safer deployment

A framework for measuring return on data investment (RODI)

Debarag Banerjee[†] Kamlesh Kumar[‡] Anik Paul[§]

The return of data investment on data asset Z is

$$RODI_Z = \frac{\Delta R_Z}{I_Z}$$

Where ΔR_Z is the total return, which is the sum of the change in expected credit loss and the expected return on loans.

$$\Delta R_Z = (ER_l(X,Z) - ER_l(X)) + (ECL_l(X) - ECL_l(X,Z))$$

The expected interest return is calculated as follows:

$$ER_l(X) = \sum_{i=1}^{N} M^i\Big(1 - P_d(D(X^i, M^i)|\hat{P}_d(D(X^i)) < \Delta)\Big)\left[\frac{rT}{(1+r)^T - 1} + (rT - 1)\right]$$

The expected credit loss is calculated as follows:

$$ECL_l(X) = \sum_{i=1}^{N} M P_d\Big(D(X^i, M^i)|\hat{P}_d(D(X^i)) < \Delta\Big)$$

AI Agents Relies on Semantic Connections

Knowledge Graph

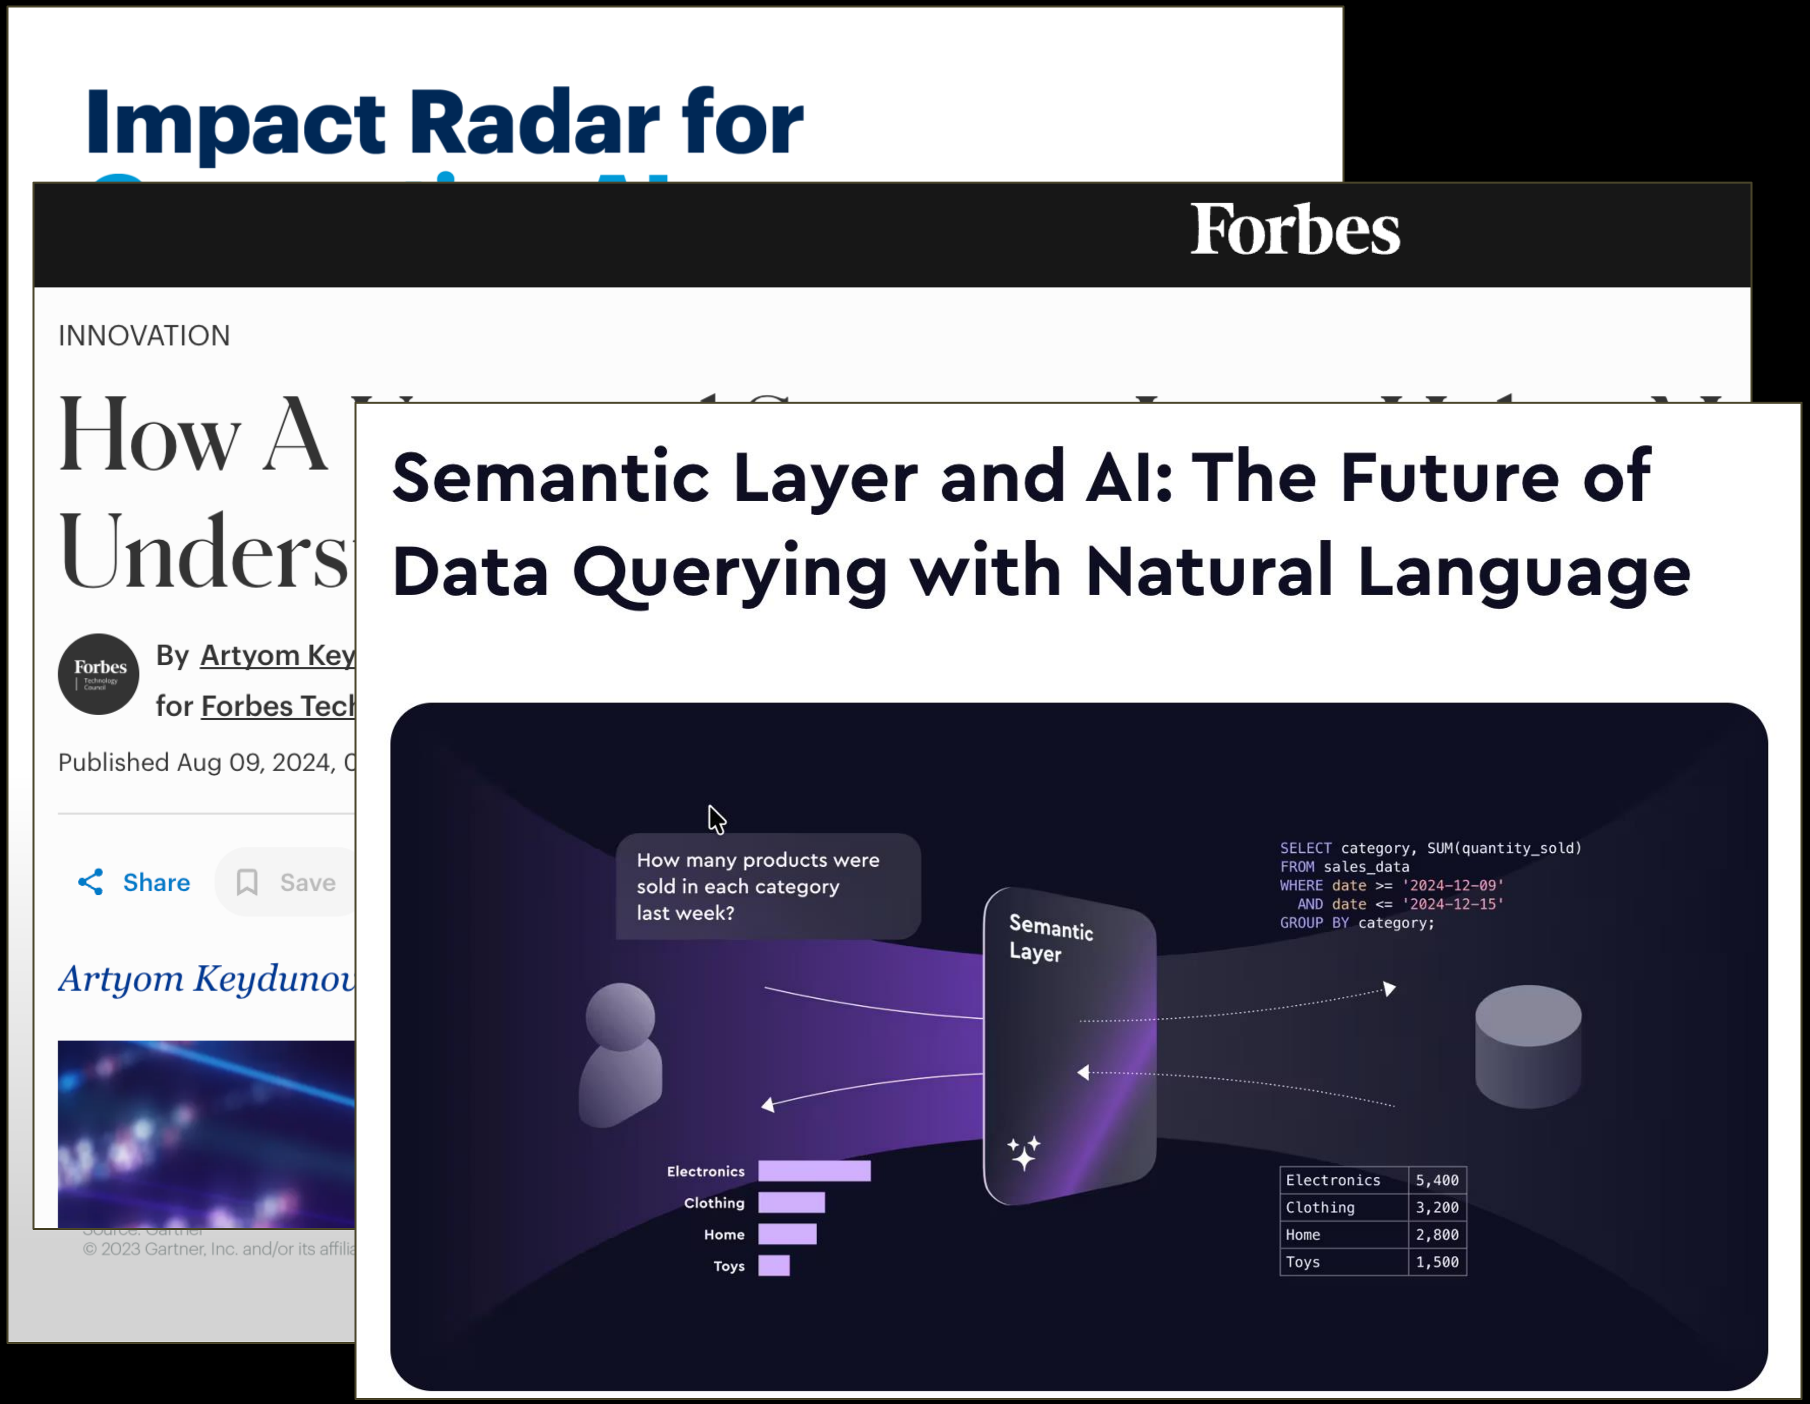

Semantic Layer

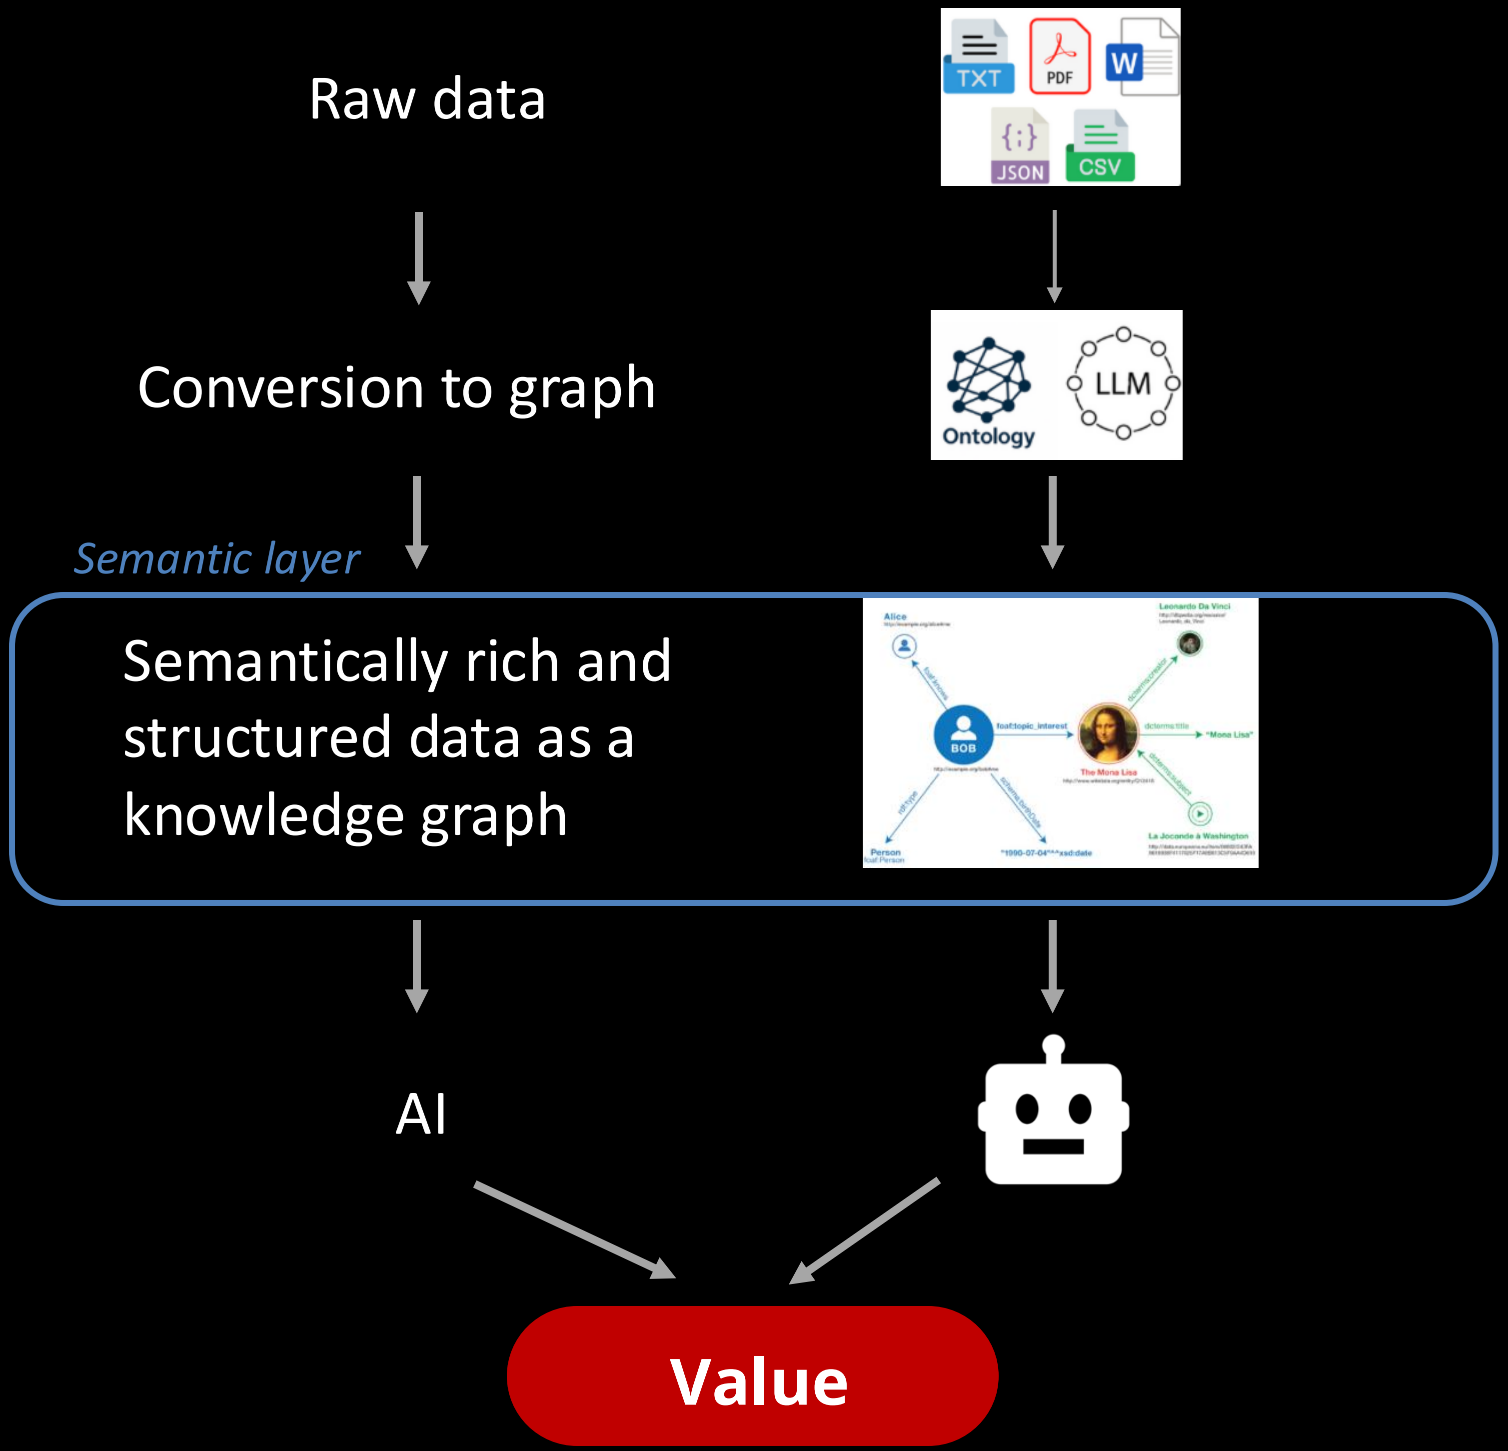

Semantic Layer

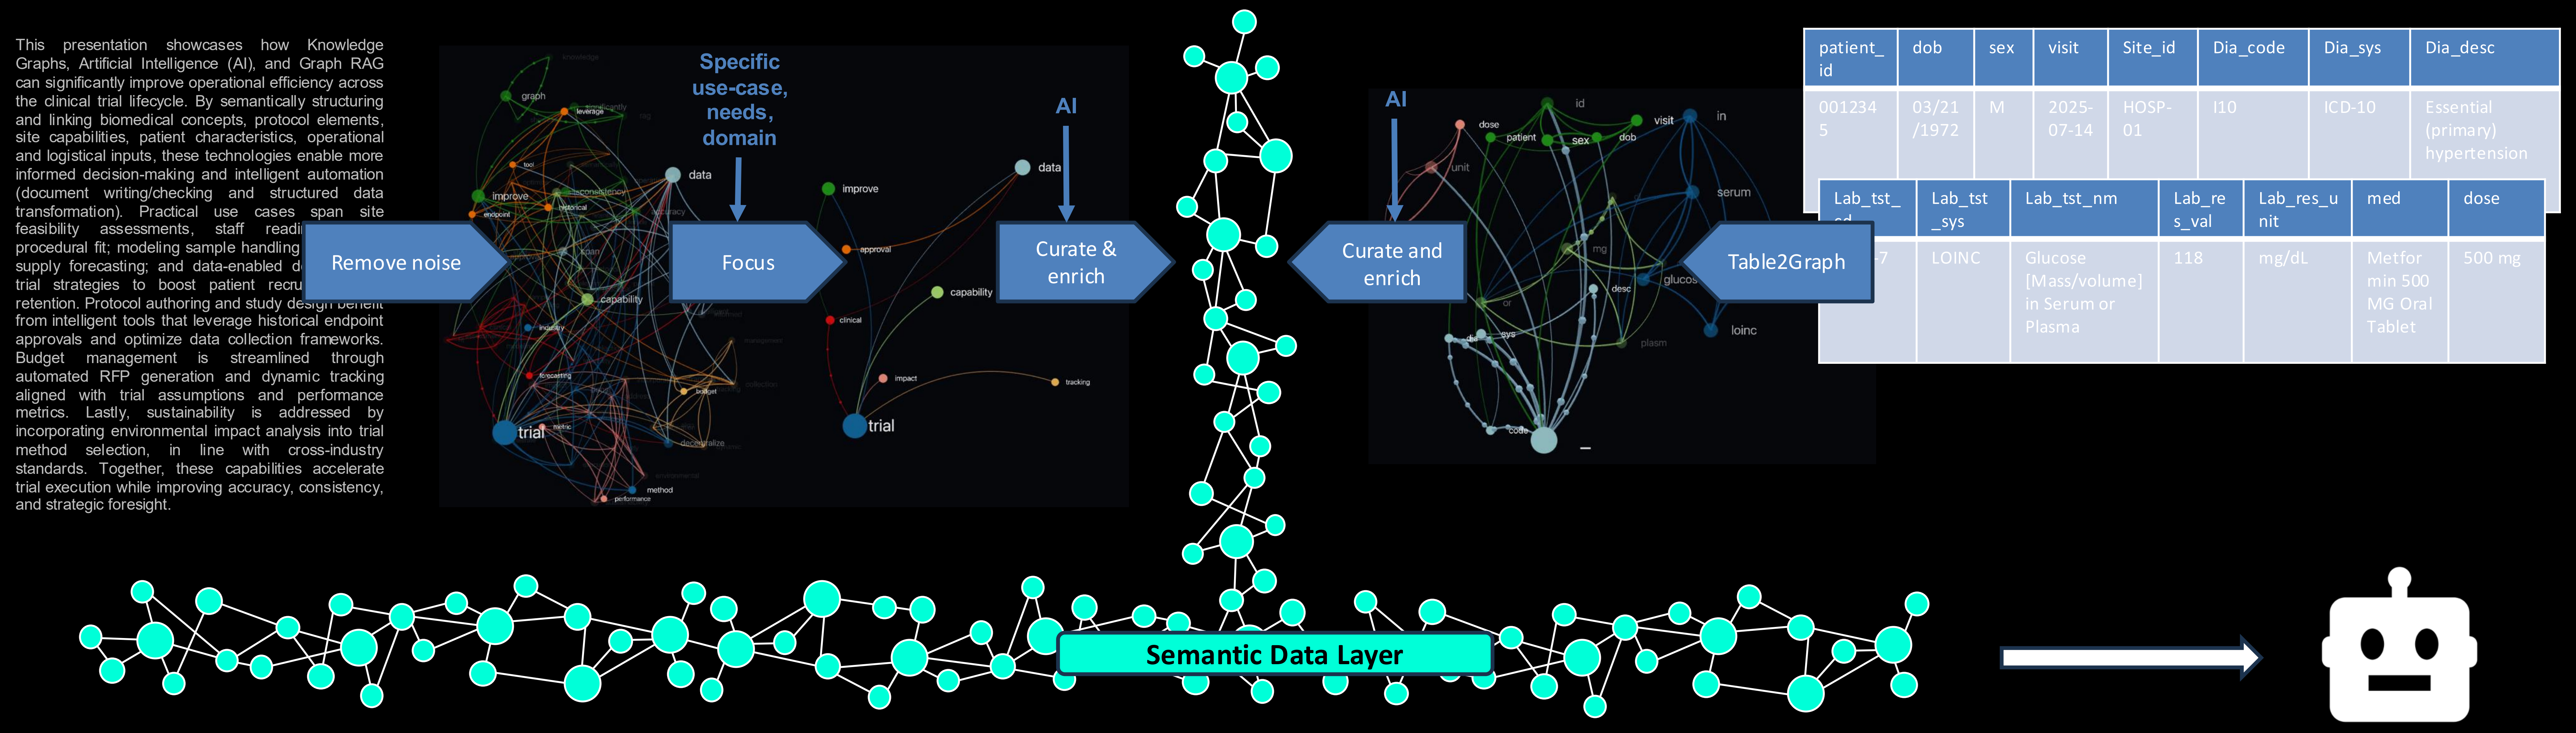

patient_id	dob	sex	visit	Site_id	Dia_code	Dia_sys	Dia_desc
0012345	03/21/1972	M	2025-07-14	HOSP-01	I10	ICD-10	Essential (primary) hypertension

Lab_tst_cd	Lab_tst_sys	Lab_tst_nm	Lab_res_val	Lab_res_unit	med	dose
-7	LOINC	Glucose [Mass/volume] in Serum or Plasma	118	mg/dL	Metformin 500 MG Oral Tablet	500 mg

Contains

- Models: schemas for internal and external data
- Data
  - Master: low-frequency changing IDs
  - Reference: terminologies

Enables

- Data Management incl integration
- Data Governance incl.
  - Traceability supports model explainability
  - Data lineage for audits and regulatory readiness

Adopters

Samsung

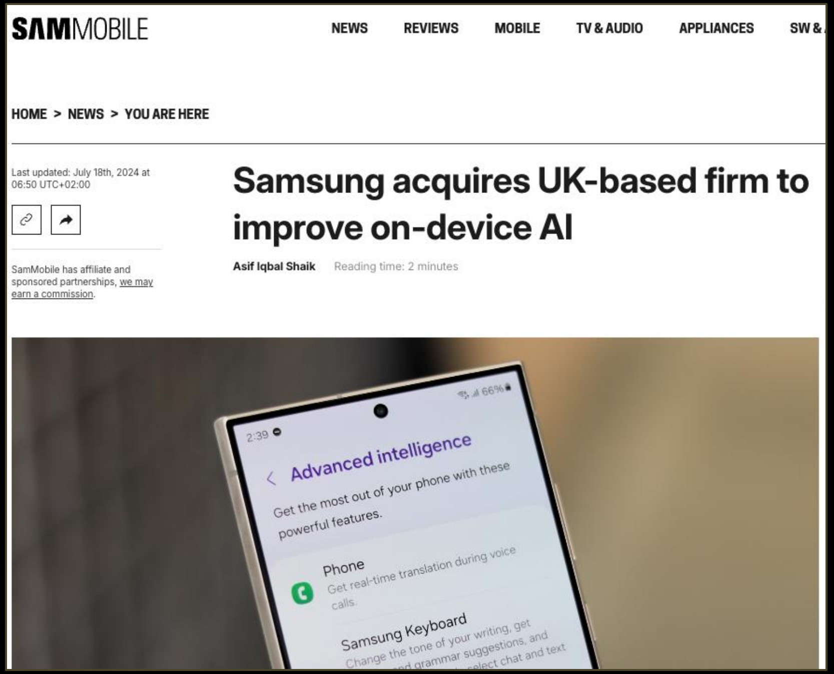

Netflix

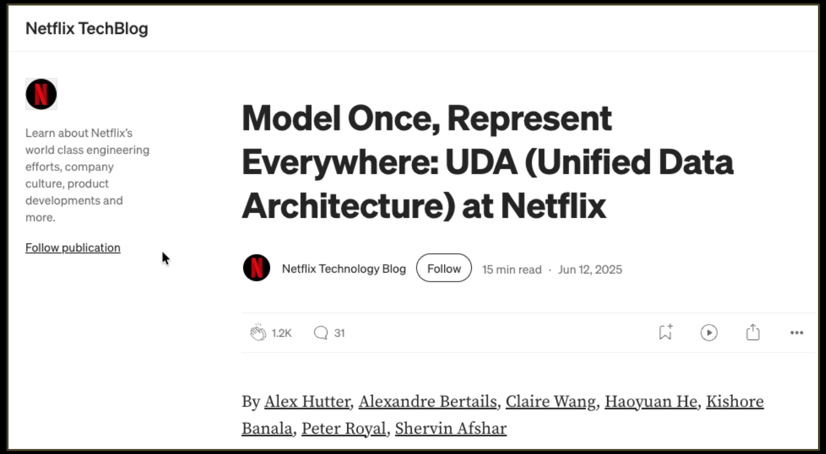

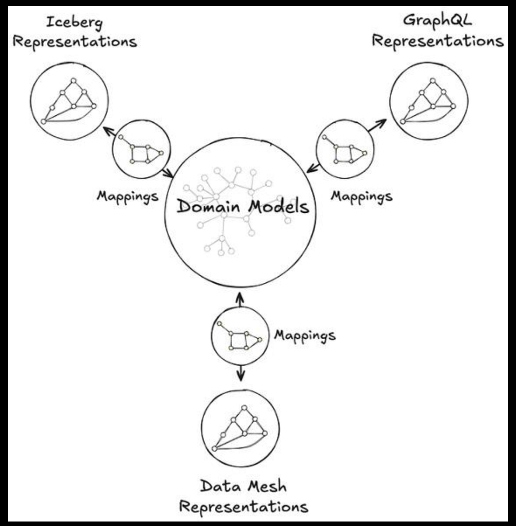

The same domain model can be connected to semantically equivalent data containers in the UDA knowledge graph.

ServiceNow

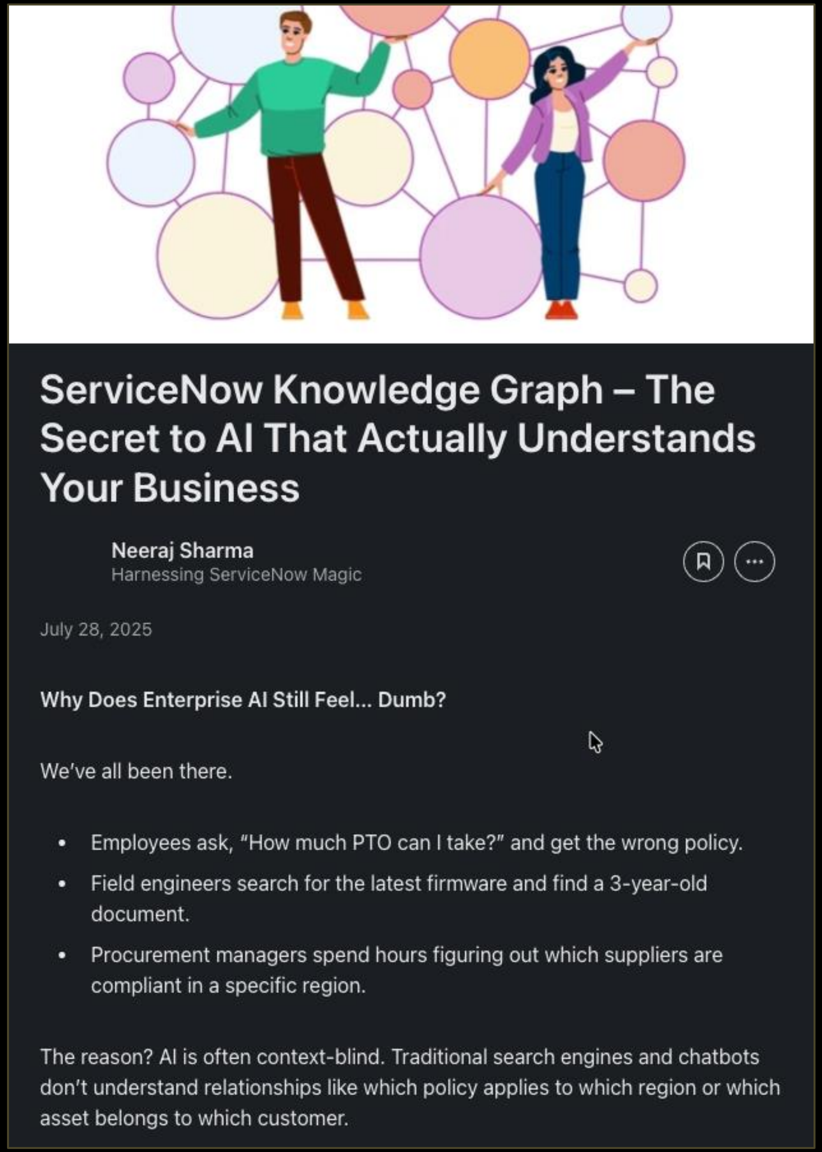

SAP

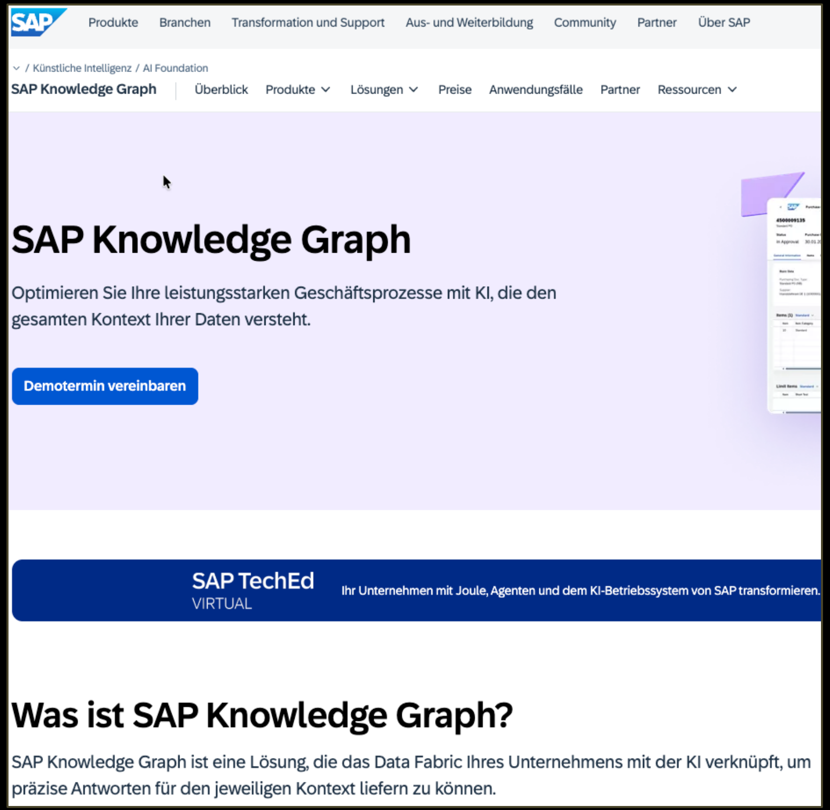

What will be the first Life Science solution adopting graph?

Need #3: Semantic Layer

- How do you feed AI with a machine-readable version of your core business/IP?
- How do you provide specific business context = connections between your value proposition/IP and externalities?
- How do you un-silo and integrate your knowledge?

Need: A framework (people, process, (meta)data and tools) to implement, manage and leverage a semantic layer:
- One or more knowledge graphs (containing models/ontologies and data)
- Pipeline to feed and retrieve graph data
- Feeding AI with it

Risk of not doing it: Not being AI-ready

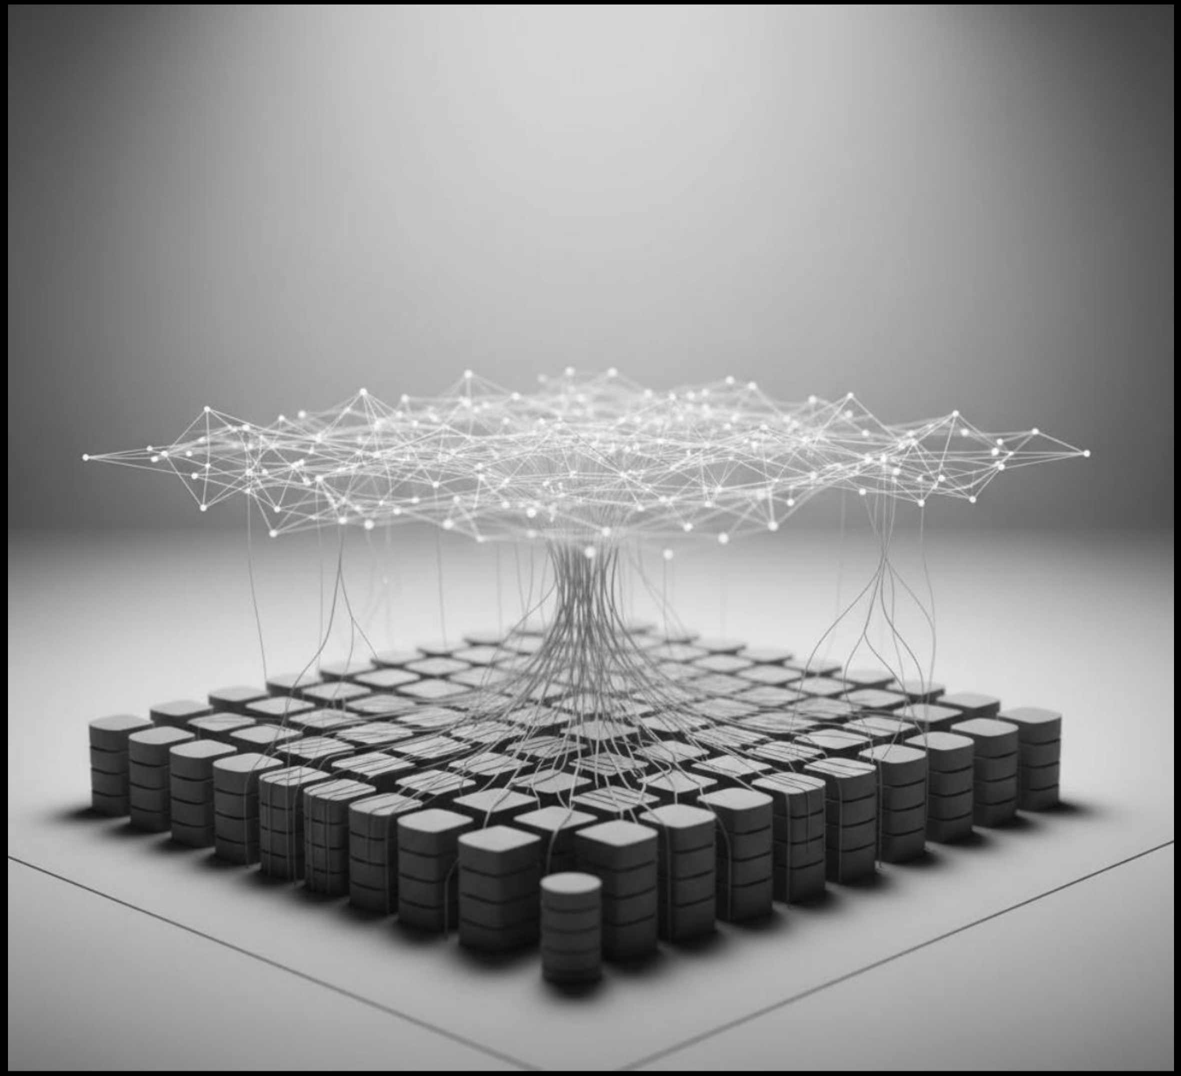

Image from LinkedIn; unknow author

More about Knowledge Graphs

https://www.biomedima.org

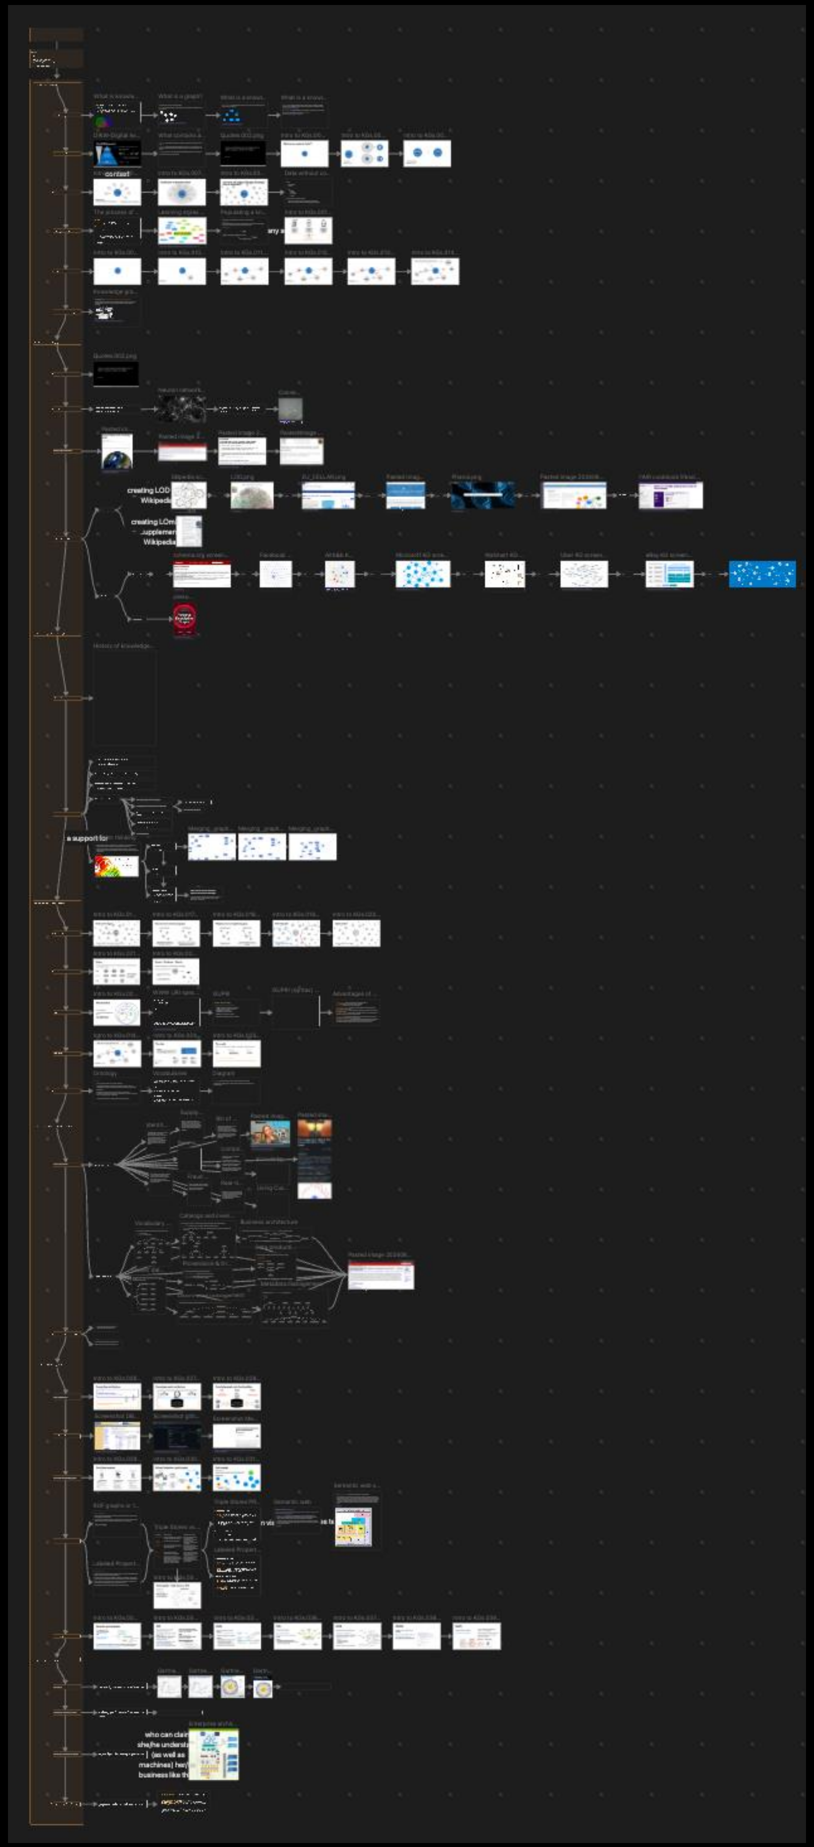

https://www.biomedima.org/project/introduction-to-knowledge-graphs/

Making the Data CapEx case

What Success Looks Like

One control room for your entire organization

AI-assisted

Relying on one single, multi-faceted, semantically annotated model

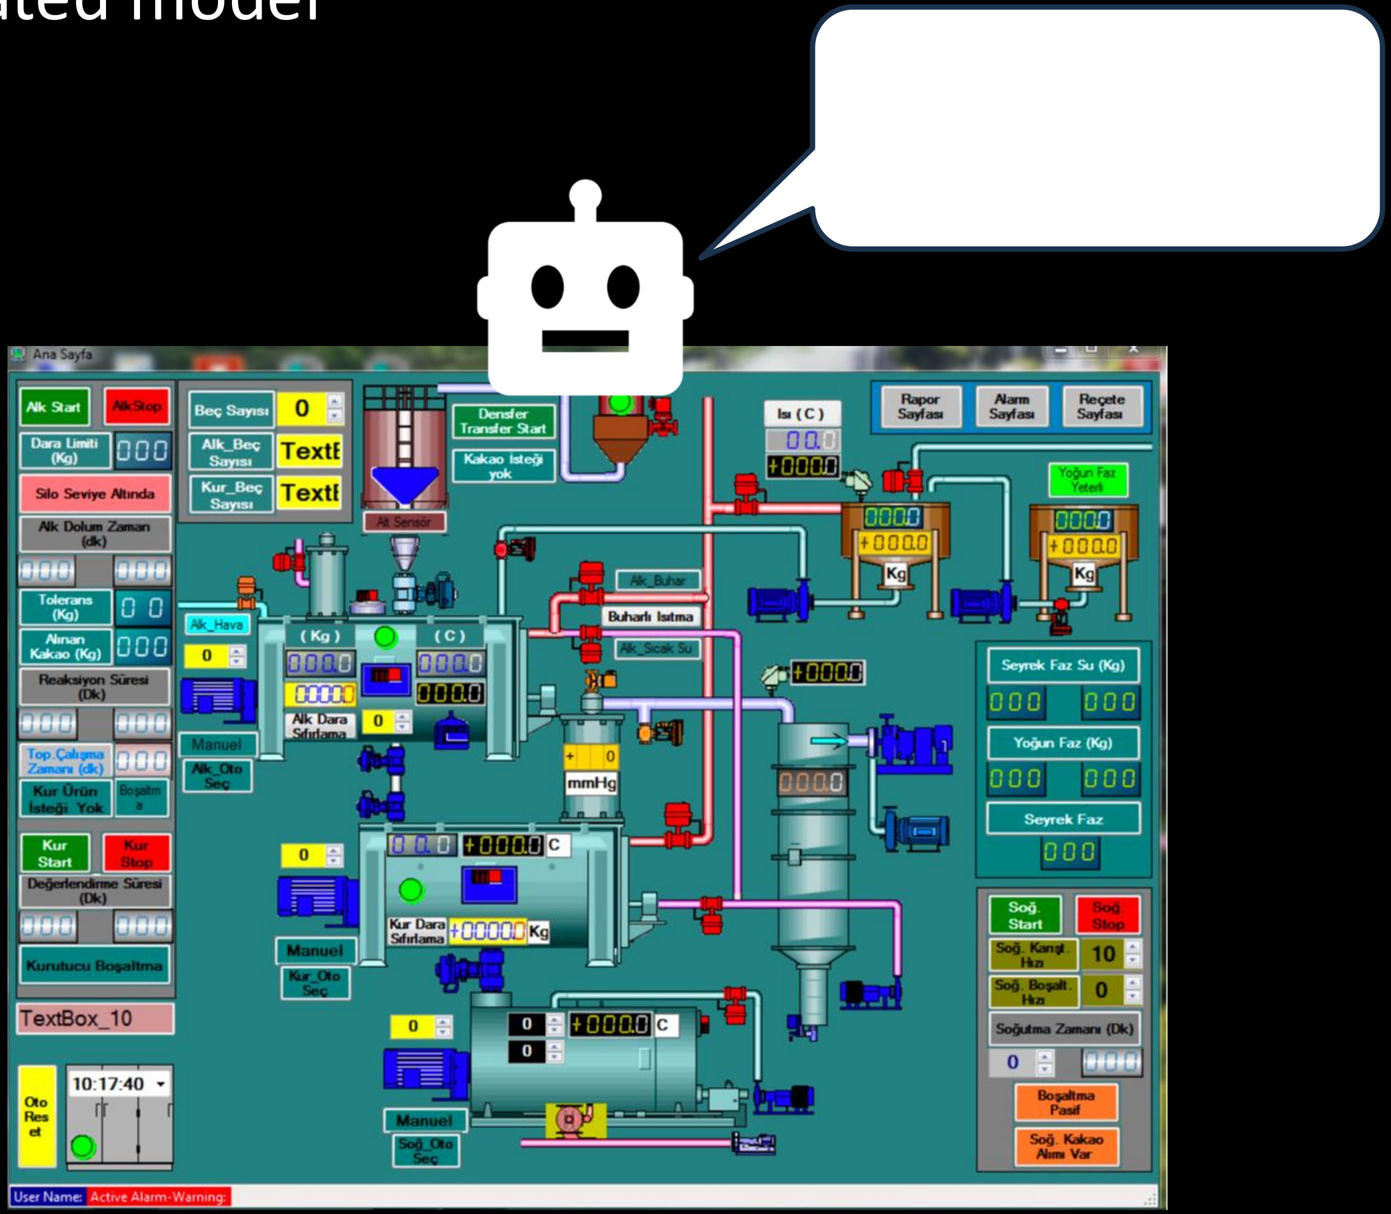

What Data-Centric AI Player are you?

Maturity Assessment

Motivation/ ability to change	Bronze - 50%	Silver – 40%	Gold – 10%
High	***Acceleration*** *10%* Fast-growing biotechs eager to modernize but still early-stage	***Industrialization*** *10%* Scaling companies investing heavily in automation and integration	***Transformation*** *2%* Digital leaders pushing AI/ML, real-time data, and next-gen architectures
Medium	***Standardization*** *25%* Organizations starting digitalization but lacking structure	***Expansion*** *20%* Most mid-cap pharma/biotech companies actively developing data capabilities	***Innovation*** *5%* Top-tier firms enhancing AI readiness and advanced analytics
Low	***Awareness*** *15%* Small biotechs, CROs, or growing firms with limited digital push	***Stabilization*** *10%* Mid-sized firms with decent tools but low adoption	***Optimization*** *3%* Large enterprises with strong foundations but low appetite for change

Recommendations

Based on a data/digital/AI maturity assessment:

Bronze

- Simple cloud storage, lightweight ETL, and basic BI tools start improving decision-making, traceability, and early ROI
- Failures often result from adopting tools too complex for immature data foundations
- Success requires selecting low-complexity, managed tools aligned with early-stage organizational readiness

Silver

- Data catalogs, lake-, warehouses, automated pipelines, and governed BI enhance traceability, deliver clearer ROI
- Implementations fail when integration effort or governance needs are underestimated
- Align tool complexity with maturing processes (especially DataOps and domain ownership) to scale effectively

Gold

- Data streaming, ML Ops, vector DBs, enterprise gov. transform decisions and create AI-driven ROI
- Risks occur when cutting-edge tech outpaces business readiness, talent, or governance capacity
- Value is maximized when advanced tools match a culture prepared for continuous change and high data literacy.

Need more information about maturity assessment and/or recommendations? *cedric@biomedima.org*

Break

Cedric Berger, PhD, MBA
2025-12-02

4th Basel CAPEX SUMMIT - Innovations in Capex Project Delivery in Pharma & Biotech

BIOMED ima

--- ROUND TABLE ---

DIGITIZING CAPEX GOVERNANCE

WHERE TO START, WHAT TO SKIP

- What to digitize first (and why)
- Governance templates vs bespoke processes
- How to scale from pilots to full rollout

4th Basel CAPEX SUMMIT - Innovations in Capex Project Delivery in Pharma & Biotech

Cedric Berger, PhD, MBA
2025-12-02

“You can only govern what you know”

- Is your organization having a consistent, coherent, unique process/framework to measure (project) business value delivery?

- Do you know this process/framework?
    - Success definition
    - KPIs

- How do you think this is going to evolve when using AI?

Key Concepts Alignment

<u>Purpose:</u> create a shared mental model before scoring our organizations

What does 'CapEx governance' mean in our context?

- "Is it about process compliance? Portfolio decisions? Accountability? Risk management?"
- "Where does governance start and stop? Intake → approval → execution → benefits?"
- "Which parts are currently consistent vs. variable?"

What do we mean by 'digitalization' ?

- "Are we digitalizing documents? Workflows? Data? Decision logic? Assumptions?"
- "Is the goal automation, transparency, auditability, or simplification?"
- "Are tools supporting processes, or are processes being reshaped for tools?"

What is 'good data' in CapEx governance ?

- "Which data must be standardized to make CapEx decisions (cost categories, risks, assumptions, KPIs)?"
- "Who defines these standards?"
- "What makes data trustworthy? Traceable? Complete? Contextual?"

What counts as 'value' in a CapEx project ?

- "Are we aligned on how we define expected value?
  → financial? strategic? compliance? sustainability? operational?"
- "How consistently do we capture assumptions and link them to value?"

What does 'integration' mean for us ?

- "Which systems must talk to each other to reduce manual work?"
- "Is integration about technical links or consistent metadata?"
- "Where does the flow break today (documents, spreadsheets, approvals)?"

What would 'AI-readiness' look like ?

- "What decisions could AI support if the foundations were in place?"
- "What quality and structure of data would be needed?"
- "What is out of scope for now?"

Diagnostic Questionnaire

Questions any business should be answering without any hesitation

Purpose & Identity

1. What core problem(s) does the organization exist to solve, and for whom? (customers)
2. What is the organization's mission, vision, and long-term ambition?
3. What values and principles guide decision-making and behavior?

Business Model

4. What products, services, or solutions does the organization offer?
5. Who are its target customers or stakeholders?
6. How does the organization generate revenue (primary and secondary revenue streams)?
7. What differentiates the organization from competitors (unique selling proposition or competitive advantage)?

Capabilities & Operations

8. What core capabilities and processes enable the organization to deliver its value?
9. What key technologies, systems, intellectual property, or data assets does the organization rely on?
10. What partnerships, suppliers, or ecosystems are essential to its operations?

Assets & Resources

11. What are the organization’s most valuable assets (people, data, IP, platforms, physical assets, brand)?
12. How is the organization structured (governance, teams, geographic presence)?

Performance & Strategy

13. What metrics define success for the organization (KPIs, OKRs, financial drivers)?
14. What major risks, constraints, or dependencies affect its ability to deliver value?
15. What strategic priorities, investments, or transformation initiatives are currently shaping its future?

All of this is knowledge, knowledge about your organization, your assets, your IP

All of this is about knowledge managment

Organization Core Blocks

Health check-up question are answered using the fundamental business architecture blocks that are all interlinked.

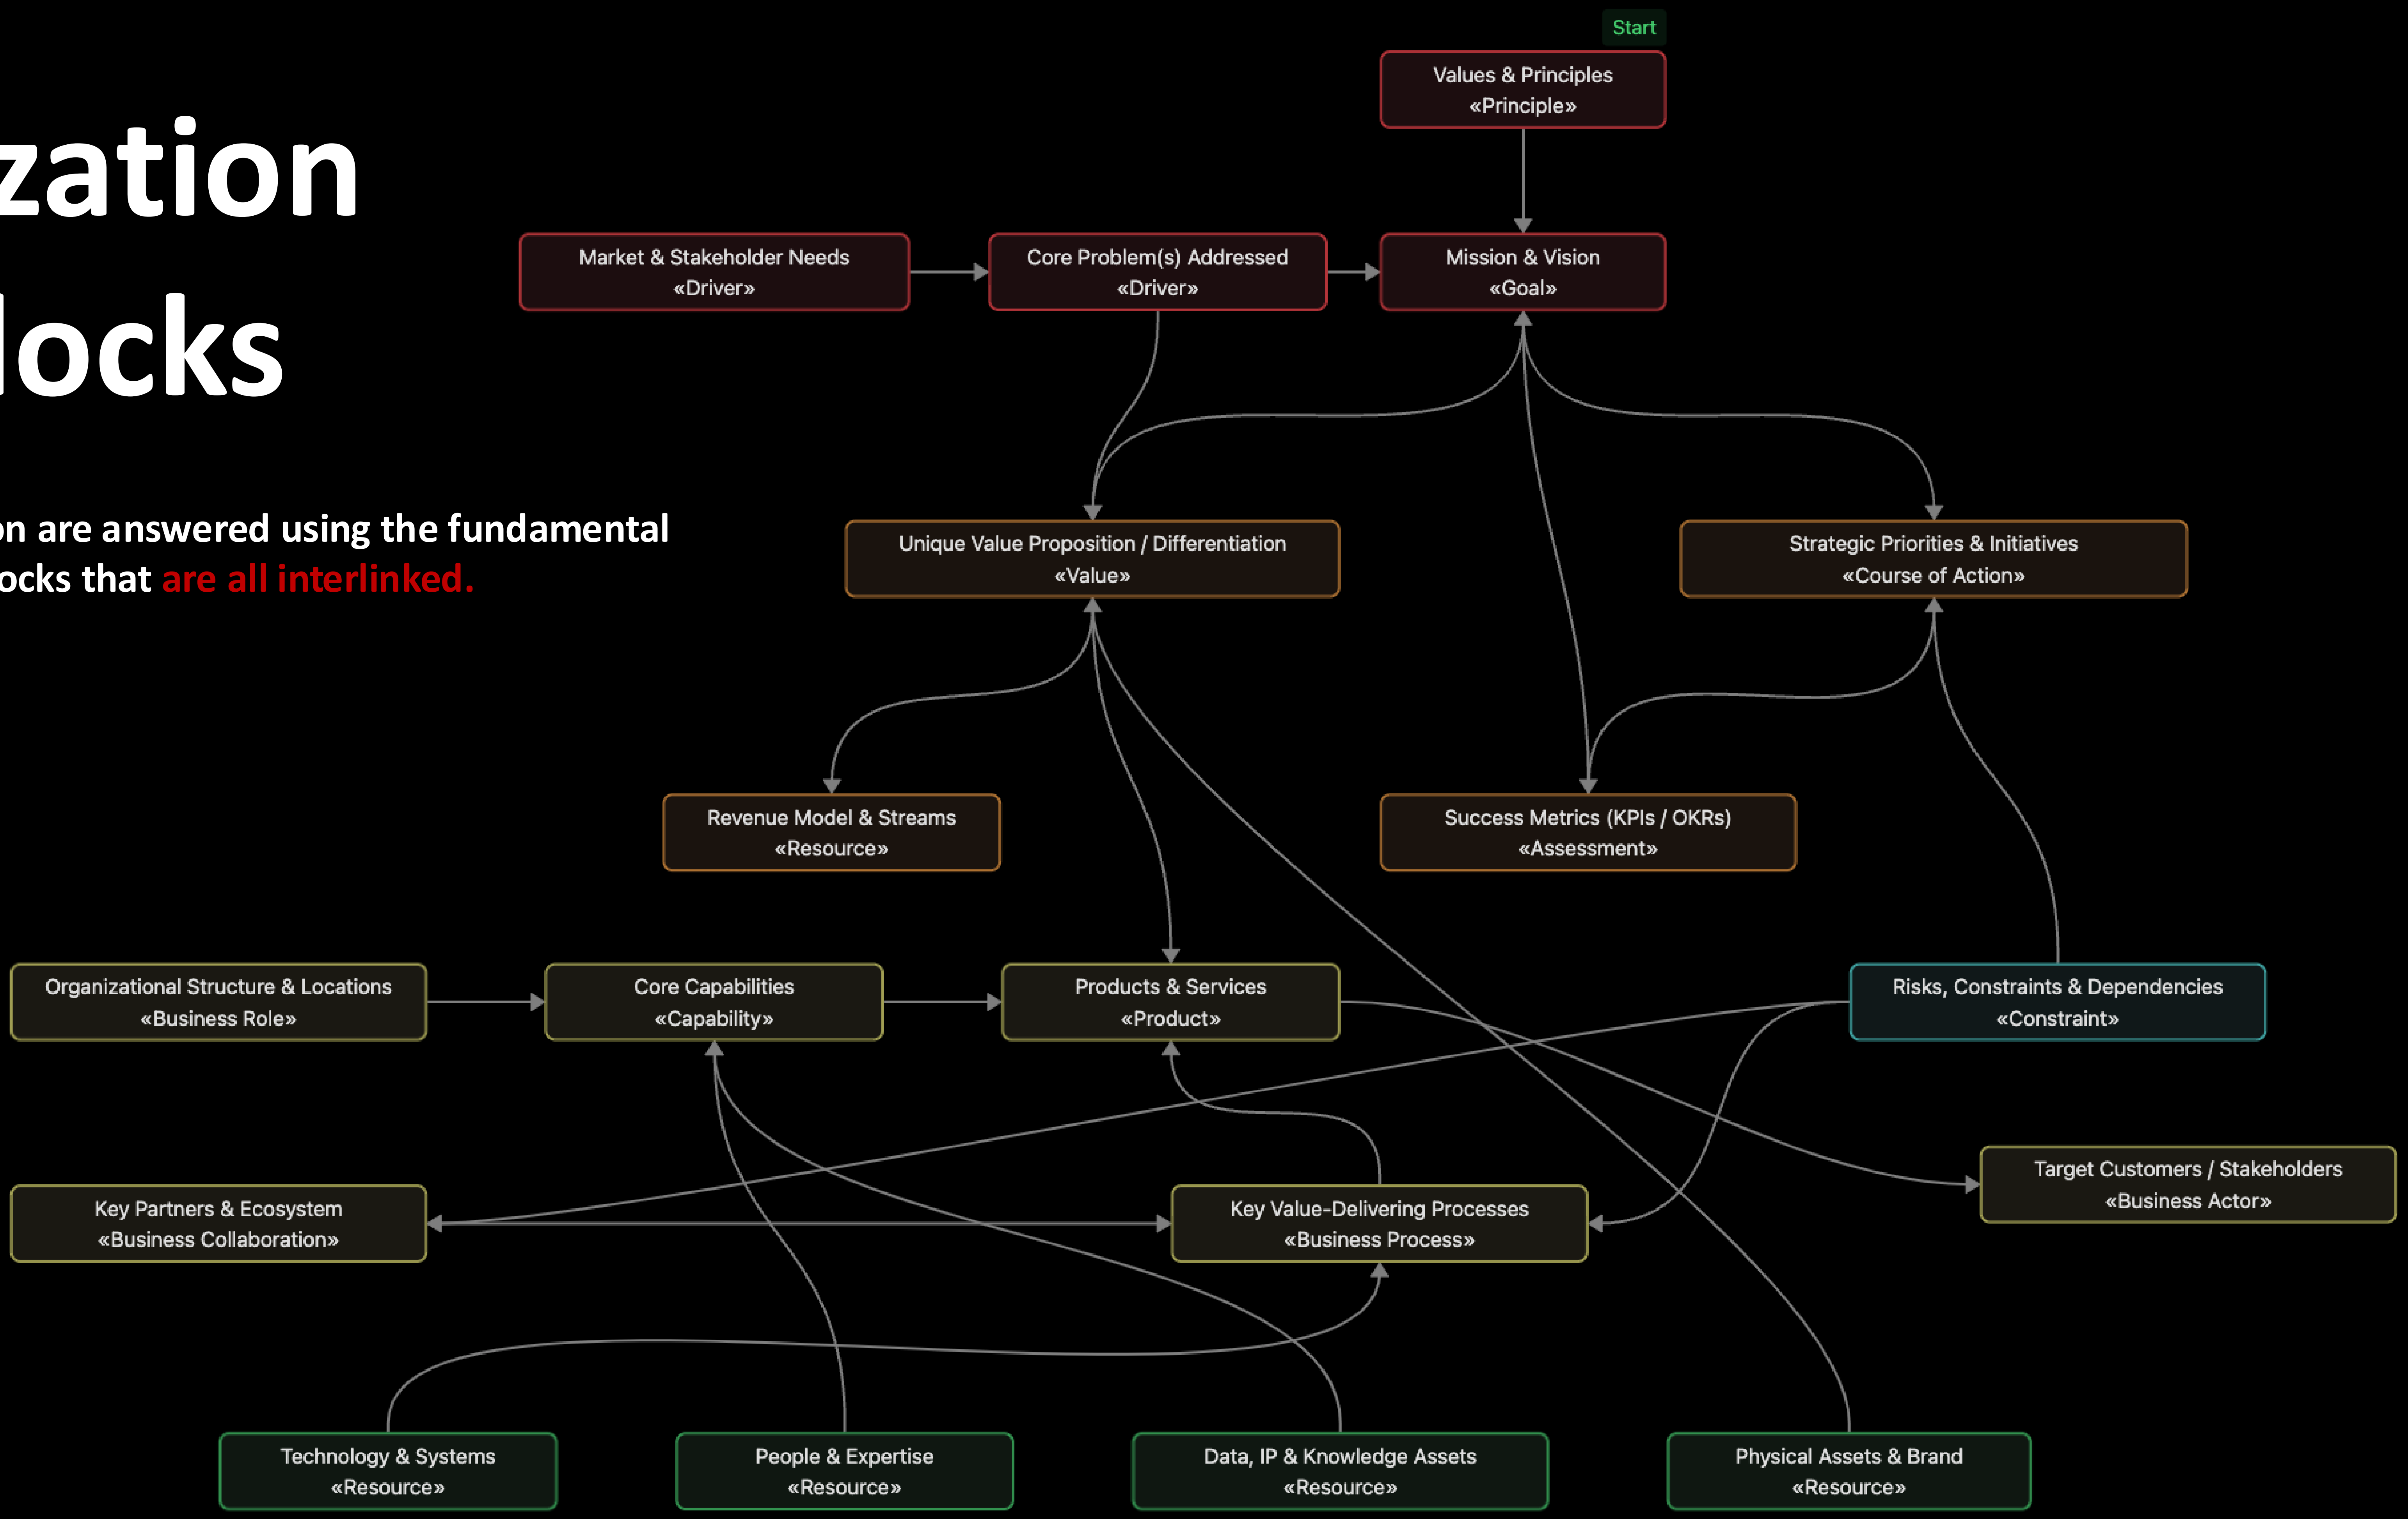

What Data-Centric AI Player are you?

Motivation/ ability to change	Bronze - 50%	Silver – 40%	Gold – 10%
High	***Acceleration*** *10%* Fast-growing biotechs eager to modernize but still early-stage	***Industrialization*** *10%* Scaling companies investing heavily in automation and integration	***Transformation*** *2%* Digital leaders pushing AI/ML, real-time data, and next-gen architectures
Medium	***Standardization*** *25%* Organizations starting digitalization but lacking structure	***Expansion*** *20%* Most mid-cap pharma/biotech companies actively developing data capabilities	***Innovation*** *5%* Top-tier firms enhancing AI readiness and advanced analytics
Low	***Awareness*** *15%* Small biotechs, CROs, or growing firms with limited digital push	***Stabilization*** *10%* Mid-sized firms with decent tools but low adoption	***Optimization*** *3%* Large enterprises with strong foundations but low appetite for change

A Strategic Roadmap to Data-Centric AI

Know thyself

- **Phase 1 (Foundation):** Inventory data assets, define ownership
- **Phase 2 (Enablement):** Deploy quality & master data tools
- **Phase 3 (Acceleration):** Integrate with AI/ML ops & data products.

Milestones & Deliverables

- **Year 1:** data catalog live, critical domains certified
- **Year 2:** data quality KPIs operational, stewardship active
- **Year 3:** AI models sourcing from certified datasets

Framing Governance as CapEx

- Governance tools, metadata guidelines, knowledge graph = capital assets
- One-time build enabling many future projects
- **Accounting alignment:** capitalize implementation, expense operation

Bronze (mainstream)

Bronze + Low Motivation

Profile: Minimal data foundations, unclear ownership, limited appetite for investment.

Recommendations:

1. People
- Create awareness sessions on "Why data matters"
- Assign a minimal part-time data steward in each domain

2. Processes & Pipelines
- Document basic data flows for 1–2 critical processes
- Introduce simple data-quality checks (completeness, duplicates)

3. Data Sets
- Identify top 5 data sets that are business-critical (e.g., customer, product, trial metadata)
- Create a small shared repository using existing tools (SharePoint, Confluence, etc.)

4. Technology
- Invest in starter tooling:
    - A lightweight cloud storage
    - A basic ETL/ELT builder
    - Begin experimenting with BI tools using shared dashboards
- Prefer OpEx-oriented managed services to reduce maintenance burden

Bronze + Medium Motivation

Profile: Some recognition of issues; willing to adopt structured improvements.

Recommendations:

1. People
- Designate part-time Data Champions in each function
- Start basic training in data literacy + dashboards

2. Processes & Pipelines
- Standardize naming conventions and introduce a "single source of truth"
- Establish a mini Data Governance Board meeting quarterly

3. Data Sets
- Start profiling key data sets for quality and lineage
- Launch a project to centralize at least one domain (e.g., HR, Sales, Clinical Operations)

4. Technology
Invest in:
- A data catalog
- A centralized data warehouse or lakehouse
- Begin implementing CI/CD for data pipelines

Bronze + High Motivation

Profile: Strong desire to modernize despite weak foundations.

Recommendations:

1. People
- Create an official Data Office (even if small)
- Train business teams in self-service analytics and low-code automation

2. Processes & Pipelines
- Define and enforce data policies (quality, access, retention)
- Build automated ingestion pipelines for critical systems (ERP, CRM, LIMS, etc.)

3. Data Sets
- Prioritize harmonization of master data domains
- Begin mapping data lineage across workflows

4. Technology
Invest aggressively:
- Deploy a modern ELT architecture
- Adopt a cloud-native lakehouse with scalable compute
- Introduce ML-powered data quality tools
- Set up centralized identity & access management

Silver (mainstream)

Silver + Low Motivation

Profile: Good core technology but suboptimal adoption.

Recommendations:

1. People
- Encourage better usage through recognition/reward programs
- Involve data-savvy talent to act as internal ambassadors

2. Processes & Pipelines
- Simplify existing data pipelines to reduce complexity
- Improve monitoring and alerting to reduce pipeline failures

3. Data Sets
- Focus on data cleanup and reducing redundancy
- Build initial KPI dashboards linked to validated data sources

4. Technology

Invest in:
- Automation-first pipeline orchestration
- Metadata harvesting tools to simplify discovery
- Improved data visualization tools with a governed semantic layer

Silver + Medium Motivation

Profile: Solid foundation and active commitment to scale data use.

Recommendations:

1. People
- Expand the Data Office with a dedicated Data Architect
- Regular training programs across functions

2. Processes & Pipelines
- Fully adopt DataOps practices
- Establish data product ownership per domain

3. Data Sets
- Accelerate consolidation of operational and analytical data sets
- Implement versioning for critical data

4. Technology

Invest in:
- A scalable lakehouse + warehouse hybrid
- Strong data governance extensibility (e.g., lineage graphs, policy APIs)
- Real-time streaming for operational use cases

Silver + High Motivation

Profile: Ready for transformation and modernization.

Recommendations:

1. People
- Launch a data competency center
- Hire or develop machine learning engineers

2. Processes & Pipelines
- Migrate toward fully automated ingestion + standardized transformations
- Build domain-based data products aligned with Data Mesh principles

3. Data Sets
- Expand enriched, curated data sets ready for analytics and AI
- Introduce synthetic data and active metadata management

4. Technology

Invest heavily:
- Implement real-time event streaming at scale
- Adopt MLOps platform
- Deploy enterprise-wide data catalog + governance suite
- Introduce API-driven data access for decentralized teams

Gold (mainstream)

Gold + Low Motivation

Profile: Highly capable tech but cultural fatigue or limited appetite for new initiatives.

Recommendations:

1. People
- Empower teams through incentives to maintain excellence
- Provide "innovation sabbaticals" for data teams

2. Processes & Pipelines
- Fine-tune operational efficiency; reduce reprocessing costs
- Audit pipelines for simplification opportunities

3. Data Sets
- Retire unused data products and optimize storage layers
- Improve accessibility without adding more tools

4. Technology
Invest in:
- Cost-optimization tooling (FinOps for data)
- Intelligent caching and data virtualization
- Next-gen performance accelerators (vectorized compute, query optimizers)

Gold + Medium Motivation

Profile: Superb data ecosystem, willing to push innovation.

Recommendations:

1. People
- Train employees on advanced AI literacy and data ethics
- Expand cross-functional innovation squads

2. Processes & Pipelines
- Adopt continuous metadata feedback loops (e.g., ranking, usage patterns)
- Expand federated governance models

3. Data Sets
- Create advanced domain-specific ontologies
- Enable data marketplaces inside the organization

4. Technology
Invest in:
- Vector databases for semantic search and RAG/AI applications
- Streaming analytics platforms
- Edge-to-cloud integration modules

Gold + High Motivation

Profile: Top-tier capability + strong executive push for innovation and AI-driven strategy.

Recommendations:

1. People
- Establish a Chief Data & AI Office with enterprise mandate
- Create AI Guilds bringing together engineering, science, and business

2. Processes & Pipelines
- Adopt Data Mesh 2.0 with fully automated governance, observability, and lineage
- Create closed-loop ML systems integrated directly into operations

3. Data Sets
- Develop rich, labeled corpora for AI, incl. multimodal data
- Implement enterprise-wide knowledge graphs

4. Technology
Invest strategically in:
- Enterprise-wide AI platform integrated with lakehouse
- Vector-native and graph-native architectures.
- Automated ML ecosystem (auto-feature engineering, auto-quality, auto-governance)
- High-performance GPU clusters or managed GPU compute
- Synthetic data generation frameworks for scenario testing

Thanks for your Attention

which is <u>not</u> all we need but is a rare, hence precious thing nowadays…

https://www.linkedin.com/in/cedricberg/