


MVC


TURNING
TECHNOLOGY
INTO BUSINESS



**You and your company
already know where the
next big discovery is.**

**You own that information.
You just can't see it.**

This is not another paper about
how AI can help you find oil and gas.
It is about what you must do
to help AI find it for you.



MVC Paper

You Already Know Where the Next Discovery Is

Allow AI to Finish the Puzzle Before You Buy Another Box

Massimo Virgilio | MVC

Abstract

Your company believes there is more oil and gas to discover. It has to. If the remaining exploration potential were negligible, the rational response would not be to improve exploration; it would be to stop exploring. The continued existence of exploration portfolios, acreage positions, seismic programs and exploration wells therefore rests on a simple premise: material discoveries remain to be made.

At the same time, the basins in which those discoveries will be made are not blank spaces. They have been observed for decades. Seismic has been acquired and reacquired. Wells have been drilled. Gravity and magnetic data, electromagnetics, remote sensing, deformation measurements and geochemical observations have been acquired. Geological studies, petrophysical analyses, basin models, interpretations and previous decisions have accumulated. Operators have entered and exited. Licences have been relinquished. Partners and competitors have generated information that later became public, tradable or otherwise accessible. The quality, density and relevance vary enormously, but the important point remains: the exploration world is already heavily observed. Put those two statements together. There are discoveries left, and many of the places containing them have already been observed in one form or another. On the Norwegian continental shelf, for example, the Norwegian Offshore Directorate estimates that 22% of total resources remain undiscovered, with 60% of those undiscovered resources in areas already open for petroleum activity.^[1] The uncomfortable consequence is that some of the evidence associated with future discoveries may already have been acquired. Your company may own part of it. It may have inherited part of it. A former operator or competitor may have acquired another part that is now accessible. Some may sit in public repositories.

But the next discovery is probably not sitting intact in a forgotten folder waiting to be found. Think instead of a puzzle whose pieces have been accumulated for decades and scattered across thousands of boxes. One piece is a seismic observation. Another is an old well result. Another is a pressure anomaly, a regional study, a discarded interpretation, a relinquishment report, a production analogue or an assumption embedded in a presentation written fifteen years ago. Some pieces are internal. Some are external but reachable. Some were rejected because the picture imagined at the time was different. No individual piece contains the discovery. **The discovery is the picture that may emerge when the right pieces are finally assembled.**

And yet, when the picture remains incomplete, the exploration reflex is often to buy another box of pieces: another seismic survey, another dataset, another study, another measurement. Sometimes that is necessary. But buying more pieces before understanding what can already be assembled from the pieces you possess is not automatically better exploration. It can simply make the pile larger.

Artificial intelligence changes what is possible. It can search, retrieve, compare and correlate information at a scale no exploration team can reproduce manually. Increasingly, AI agents can move beyond retrieval: they can query databases, read technical documents, call APIs and deterministic scientific engines, execute defined workflows, test relationships, expose contradictions and iterate from one question to the next.

But this is not another paper about what AI can do for exploration. **It is about you and what your company must do before AI can do it.** AI cannot assemble pieces it cannot reach, understand or use. If data remain fragmented across databases and documents, if interpretation context remains trapped in specialist applications, and if scientific engines and workflows remain accessible only through human-operated interfaces, then a more capable AI simply becomes a better puzzle solver staring at locked boxes.

This is why digital transformation is the subject of this paper. In particular, the transformation of **data and application workflows**. Data must be liberated, abstracted and connected so that information can be discovered with context, semantics and provenance. Applications must stop being isolated destinations and expose the data, functions, calculations and scientific engines that AI agents and people need to invoke across workflows. Digital transformation puts the pieces on the table and makes the tools usable. Only then can AI help assemble the picture.

The provocation is therefore not “invest in AI.” It is more demanding: **transform the digital environment in which AI is expected to operate. Before you buy another box of data, make it possible for AI to work properly on the puzzle you have already paid for.**

1. You Believe There Is More Oil to Find

Every exploration strategy contains an implicit statement about the subsurface: something material remains unknown. Companies bid for acreage, acquire seismic, fund geological studies and drill exploration wells because they believe that today's understanding is incomplete and that tomorrow's interpretation can identify resources that have not yet been recognized.

This is almost too obvious to state, but it matters because it establishes the first premise. A company cannot simultaneously justify a substantial exploration program and behave as though all material opportunities have already been found. Exploration exists because uncertainty remains economically valuable.

The next discovery may be subtle, deeper, smaller, stratigraphically more complex or located in an area previously considered unattractive. It may require a different geological concept or a connection between observations that were previously treated independently. None of that changes the premise. If you are still exploring, you believe there is something left to find.

2. You Have Already Looked at Much of Where You Are Still Looking

The second premise is equally important. The world's prospective petroleum basins are not unknown territories. They have been investigated with different technologies, at different resolutions and at different moments in the industry's technical development. Some areas have dense modern 3D seismic and extensive well control; others may have sparse 2D lines, old wells, gravity and magnetic data, regional geology or limited public studies. "Explored" does not mean completely understood. It means observed.

For an individual company, the relevant information universe is also larger than the corporate archive. It includes information generated by partners, previous licence holders and competitors; national data repositories; relinquishment reports; released well and seismic data; academic work; purchased datasets; and other information that has become accessible over time. Corporate boundaries change faster than geology. Information that once sat behind somebody else's firewall may eventually become part of the evidence available to you. Norway provides a concrete example: in 2023, almost 30,000 datasets were made publicly accessible following changes in confidentiality periods.^[4]

Where genuinely underexplored acreage exists, new observation remains essential. But much modern exploration is difficult precisely because companies are also searching within areas that have already been examined, often repeatedly. The frontier is increasingly not simply geographic. It is interpretational and informational.

3. You Probably Already Have the Evidence

Now combine the premises. Material discoveries remain. Much of the acreage in which companies continue to explore has already accumulated substantial observation. It raises an important possibility: some of the evidence associated with future discoveries may already have been recorded somewhere.

This does not mean that an old seismic line contains a clearly visible field waiting for somebody to circle it. Nor does it mean that new acquisition is unnecessary. The proposition is more interesting: the evidence may be distributed. A structural observation may sit in one seismic interpretation; a

surprising pressure point in an old well; a source-rock implication in a regional study; an analogue in another asset; a failed prospect in a relinquishment report; and a key assumption in a presentation whose author left the company years ago. Individually, none provides the answer. Together, they may change the exploration hypothesis.

Johan Sverdrup is a useful reminder of how this can happen. The giant field was discovered in 2010 in a mature North Sea setting after roughly 50 years of regional exploration; wells drilled between 1967 and 2006 had already encountered oil shows.^[2] Rønnevik's account of the discovery process describes how the 2007 Edvard Grieg discovery changed the interpretation of the area and materially altered the perceived prospectivity of Johan Sverdrup.^[3] The pieces did not become valuable merely because they were old or new; their meaning changed when the picture changed.

The discovery is new. The evidence does not necessarily have to be.

4. The Problem Is That You Cannot See What You Know

A company can possess information without possessing usable knowledge. This is the central distinction.

Exploration information was created over decades by different disciplines, contractors, assets and generations of software. It was optimized for the task and application of the moment, not for future machine interrogation. Seismic interpretation lives in one environment, wells in another, reports in document stores, regional studies in shared drives, models in specialist applications and historical decisions in presentations, emails and the memories of people who may no longer be present.

Fragmentation also preserves old judgments. A prospect may have been rejected because of the seismic quality available at the time, an assumed charge model, an economic threshold or a regional concept that later changed. The data survives, but the path from observation to decision is rarely reconstructed systematically. Consequently, yesterday's conclusion can remain attached to today's evidence long after its assumptions have weakened.

No exploration team, however capable, can continuously reread decades of corporate and public information, compare every relevant well with every analogue, revisit every abandoned interpretation and follow every relationship across every application. The company therefore knows more in aggregate than any person or team can practically know. The problem is not the absence of information. It is the inability to assemble the company's distributed memory when a decision requires it.

5. And Yet You Keep Buying More Data

When uncertainty remains, the exploration reflex is understandable: acquire another survey, license another dataset, commission another study or drill another well. Sometimes that is exactly the correct decision. New measurements can resolve ambiguity that existing information cannot.

But there is a paradox. Companies can subject new-data investments to rigorous technical and economic scrutiny while devoting far less attention to whether the information already available to them can actually be exploited. The new survey is acquired, processed, interpreted and stored. The project finishes. Teams move. Applications change. Context decays. A few years later, the new information has joined the same fragmented estate that made the previous information difficult to use.

The organization has not necessarily reduced its information problem. It may simply have made the pile larger.

6. You Do Not Have a Data Pile. You Have an Unfinished Puzzle

The problem is often described as information being “hidden” in corporate data. That description understates it. A hidden answer suggests that somewhere there is one report, one seismic line or one forgotten interpretation containing the discovery, and that the task is simply to find it. Exploration knowledge is rarely that convenient.

A better metaphor is an unfinished puzzle. Over decades, your company and the organizations around it have accumulated pieces. Those pieces were created for different purposes, at different times and with different technologies. They sit in seismic projects, well databases, geological reports, interpretations, models, presentations, code, public repositories and specialist applications. Some are carefully catalogued. Some are poorly labelled. Some are mixed with irrelevant pieces. Some have become detached from the reasoning that once gave them meaning. Some sit in boxes owned by somebody else but are now available for you to open.

Most importantly, the picture you are looking for may not exist in any single piece. A seismic feature becomes interesting because of what an old well says about charge. That well becomes more relevant because of a regional source-rock study. A discarded prospect changes meaning when compared with an analogue from another basin. A public relinquishment report changes the interpretation of an internal dataset. The exploration insight is created by the relationship among the pieces.



This changes the question from “where is the answer?” to “which pieces belong together?” It also changes the logic of new-data acquisition. Buying another box may provide the missing piece. But it may also add thousands of pieces to a puzzle whose existing contents have never been systematically assembled.

The challenge is therefore not to stop acquiring data. It is to stop assuming that more pieces are the first answer to an incomplete picture.

7. AI Is Ready to Work on the Puzzle. Are You?

For decades, the scale of the puzzle imposed a practical limit. Even if every relevant piece could theoretically be found, no human exploration team could continuously read, compare and correlate all of it. People specialize. Teams change. Assets are reorganized. Knowledge leaves with employees. Attention is finite. The information estate grows faster than any organization can manually absorb it.

AI changes the scale at which that constraint can be addressed. It can search large bodies of structured and unstructured information, extract entities and relationships, compare technical descriptions, retrieve analogues, connect observations separated by discipline, geography or time, summarize

competing interpretations, expose inconsistencies and formulate new questions from what it finds. Agentic systems can increasingly query databases, invoke APIs, call deterministic scientific engines, execute defined workflows and return to the evidence with another hypothesis. Emerging subsurface research is already testing agent-to-agent and tool-interface patterns for geoscience and reservoir-simulation workflows.^[8,9]

Imagine an exploration question that causes an agent to retrieve relevant wells, inspect historical interpretations, search public relinquishment reports, compare pressure observations, identify geological analogues, call an existing basin-model or petrophysical engine, challenge an old assumption and return not with a magical answer but with a traceable set of relationships for geoscientists to evaluate. AI does not replace geoscience in this model. It increases the number of pieces that expert geoscience can practically consider.

This is no longer hypothetical. ExxonMobil reports that an AI model given its Guyana seismic data could find the oil the company had already discovered with a 90% success rate. Trained on its own Stabroek discoveries, drilling results and subsurface data, AI then identified four new exploration opportunities in parts of the block the company had not prioritised. Applied to data from 50,000 wells drilled across the industry, an analysis estimated at fifteen years of manual work took weeks and returned 150 exploration opportunities worldwide. The company’s own conclusion is that access to data is what makes AI useful in exploration.^[10,11] The pieces were already in the box.

But the important point of this paper is not that AI can do these things. That argument is already being made everywhere. The important point is that **your company may be preventing AI from doing them**. AI cannot liberate itself from a digital environment that was designed around disconnected applications, human-operated interfaces and fragmented stores of information.

Give AI selected documents and it knows selected documents. Give it a database without interpretation context and it knows the database. Leave a valuable scientific engine behind a manual interface and an agent cannot invoke it as part of a wider reasoning workflow. Keep the relationship between an interpretation, the data supporting it and the decision that followed it implicit in people’s heads, and AI cannot reconstruct that relationship reliably merely because the model became smarter.

The bottleneck is therefore increasingly not only AI capability. It is corporate digital readiness. **You have a machine capable of helping assemble the puzzle. The question is whether you have transformed your data and applications sufficiently to put the pieces and tools within its reach.**

8. Put the Pieces on the Table and the Tools Within Reach

This is the central argument. AI does not eliminate the need for digital transformation. It makes digital transformation more urgent.

For years, fragmented data and isolated applications could be tolerated because humans acted as the integration layer. An experienced geoscientist knew which database to search, which interpretation project to open, which specialist application to launch, which colleague to call, which report might contain the missing context and how to move an answer manually from one system into another. That model was inefficient, but it worked because people compensated for the architecture.

AI agents change the economics of that compromise. Their value comes from being able to work across far more informa-

tion and far more steps than a human can practically traverse. If the digital estate forces the agent back into the same narrow silos, much of that advantage disappears. A powerful reasoning capability connected to a fragmented environment remains fragmented intelligence.

The first requirement is therefore **data liberation**. The underlying principle is not new: FAIR data stewardship explicitly emphasizes machine findability, accessibility, interoperability and reuse, and extends those principles beyond data to algorithms, tools and workflows.^[6] Information must be reachable outside the application or repository in which it happened to be created. Heterogeneous sources need sufficient abstraction, semantics, context and provenance for people and machines to understand what they are looking at. Structured data, technical documents, interpretations, models and external information need to become connectable rather than merely stored.

The second requirement is **application and workflow liberation**. Valuable upstream applications contain more than screens. Behind them sit calculations, algorithms, scientific engines, quality controls and deterministic workflows developed over decades. Those capabilities should increasingly be exposed through interfaces that allow them to participate in broader workflows. This direction is consistent with OSDU's emphasis on open, standards-based access to subsurface data and with emerging work exposing deterministic reservoir-simulation capabilities to AI agents through explicit tool interfaces.^[7,8] An AI agent should not need to imitate a geoscientist clicking through every application. Where appropriate, it should be able to call the underlying capability, pass it the right information, receive a traceable result and use that result in the next step of the investigation.

This does not require one gigantic perfect database, nor does it mean cleaning every historical record before value can be created. It means designing an environment in which useful pieces can be discovered, understood, connected and acted upon, while preserving the scientific engines that remain authoritative for deterministic computation.

Digital transformation therefore has a much more consequential objective than modernizing IT. It creates the operating environment in which AI can express its potential against the company's highest-value technical questions. **Data abstraction puts the pieces on the table. Application liberation puts the tools within reach. AI helps determine which pieces might fit. Domain experts decide whether the emerging picture makes geological sense.**

Your company does not need AI to discover oil from nothing. It needs a digital environment that allows AI to help discover what your company already knows but has never been able to know collectively.

9. Before You Buy Another Box

None of this is an argument against new seismic, new wells or better measurements. Exploration will continue to require new data. Some puzzles genuinely have missing pieces. Existing measurements may lack the required resolution, geometry, coverage or physical sensitivity. A new play may demand observations that do not exist. A prospect may require a well.

The implication is not that your company should stop acquiring new data. New seismic, new wells and new measurements will remain essential wherever existing evidence is insufficient. The point is that acquisition should no longer be the automatic response to uncertainty.

Before approving the next data-acquisition budget, ask a different question first: *have we made it possible to discover what we already know?*

That requires more than deploying an AI model. It requires the digital transformation of the environment in which that model operates: data that can be found, understood and connected; applications whose scientific capabilities can be accessed beyond their traditional human interfaces; and workflows that can be interrogated, combined and orchestrated across organizational and technological boundaries.

AI may be capable of assembling relationships across a volume and diversity of evidence that no exploration team could examine manually. But it cannot assemble pieces it cannot reach, understand or use. The constraint is increasingly not the intelligence of the machine, but the digital environment your company gives it.

AI may help assemble the geological picture, but the subsurface still has the final vote: a prospect becomes a discovery only when physically demonstrated.

The next major discovery may still require new data. But before assuming that it does, consider a more uncomfortable possibility: much of the evidence may already be somewhere within your reach, fragmented across the information your company has accumulated or can access.

Before you buy another box, transform your company so that AI can finally assemble the puzzle already in front of you.

References

1. Norwegian Offshore Directorate (2025). *Resource Accounts as per 31 December 2024 – Undiscovered resources*. Estimates 22% of total NCS resources as undiscovered, 60% in areas open for petroleum activity.
2. Ottesen, S. et al. (2022). "Geology of the Johan Sverdrup field: A giant oil discovery and development project in a mature Norwegian North Sea basin." *AAPG Bulletin*, 106(4), 897–936. doi:10.1306/11042120037.
3. Rønnevik, H.C. (2020). "The Discovery Process behind the Giant Johan Sverdrup Field." In *AAPG Memoir 113*, Ch. 12.
4. Norwegian Petroleum Directorate (2023). "30,000 data sets are now publicly accessible." 19 June 2023.
5. Norwegian Offshore Directorate. *Diskos – National Data Repository for Petroleum Data*.
6. Wilkinson, M.D. et al. (2016). "The FAIR Guiding Principles for scientific data management and stewardship." *Scientific Data*, 3, 160018. doi:10.1038/sdata.2016.18.
7. The Open Group, OSDU Forum. *OSDU Data Platform: open, standards-based subsurface data platform and services*.
8. Novikov, A., Serrao Seabra, G., Agrawal, N., Garcia Alvarez, M. & Voskov, D. (2026). "Design and evaluation of MCP-server for agentic reservoir simulation workflows." *ECMOR 2026*. doi:10.3997/2214-4609.202637158.
9. "Dynamic GeoCopilot" (2026). *87th EAGE Annual Conference & Exhibition*. doi:10.3997/2214-4609.2026101472.
10. ExxonMobil (2026). Fireside chat, 42nd Annual Bernstein Strategic Decisions Conference, 28 May 2026 (N. Chapman, Senior Vice President). Transcript.
11. ExxonMobil (2026). Q2 2026 earnings call, 31 July 2026; reported in *World Oil*, 13 August 2026.