WORLD EPA CONGRESS 2026

AI & Digital Transformation Track

Five verdicts from two days of sessions — and why they matter now

Ryan Bishop | March 2026 | ryanbishop.co.uk

Why I Was in the Room

Amsterdam, 3–4 March 2026. I attended the World EPA Congress as the launch platform for ryanbishop.co.uk — a consultancy built on a deliberate positioning: twelve years of pharmaceutical experience in HEOR, market access, and HTA, combined with AI strategy capability developed through formal study and practice.

The AI & Digital Transformation track was where I spent both days. Not because it was the only relevant track, but because it was the room where the questions I've been working through — about implementation, domain expertise, governance, and what AI actually changes in this field — were being debated by the people doing the work.

This is not a conference summary. A summary would tell you what speakers said. This is an analysis of what the sessions revealed — five verdicts drawn from two days of evidence. The conference validated the launch. Here is why.

> ***AI in pharma market access has crossed the line from concept to delivery. The tools exist. The use cases are live. What is missing is the implementation layer — the people and structures to make it work at scale.***

TAKEAWAY 1

Generic AI is failing HTA. Domain-specific architecture is the answer — and most organisations haven't built it yet.

Multiple sessions converged on the same finding: generic large language models do not work for HTA. The reasons are consistent and structural. Hallucinated citations. No source traceability. Outdated clinical trial data. No access to unpublished proprietary material. Outcome inconsistency. The list appeared session after session, from BMS to Pharmacoevidence to the Harvard presentation, not as caveats but as the central problem framing each talk.

The proposed solution — Retrieval Augmented Generation, or RAG — was equally consistent. RAG constrains model output to a curated evidence corpus rather than a general training dataset. For HTA, where a hallucinated citation in a dossier is not an inconvenience but a regulatory and scientific failure, this distinction is fundamental.

But the d-fine session introduced a critical caveat that most of the room glossed over. Basic RAG operates on text similarity — vector matching without structured connections. It reduces

hallucinations, but it does not solve provenance or complex relational reasoning. Knowledge Graph RAG, which layers structured domain knowledge onto retrieval, is the architecture that actually addresses the problem. It is also significantly more complex and expensive to build.

The MIT figure cited by Cube RM — 95% of companies seeing zero return on GenAI investment despite $30–40 billion in enterprise spend — is not a coincidence. It is the direct consequence of organisations deploying generic tools at a specialist problem and measuring the gap in ROI rather than architecture.

> **Verdict:** Generic AI in HTA is not a version problem — it will not be solved by upgrading to the next model release. It requires domain-specific architecture. Most organisations are not there yet.

TAKEAWAY 2

HITL is non-negotiable — but it is being misunderstood.

Human-in-the-loop appeared in almost every session. By the end of day two, it had the quality of a conference consensus — something everyone agreed on, which often means something important is being flattened into a slogan.

There are two versions of HITL operating in this space, and they produce very different outcomes. The first is compliance HITL: a human reviews AI output before it goes out the door, providing sign-off and accountability. This reduces risk. It does not unlock value.

The second is expert HITL — what BMS articulated clearly when they said that upskilled subject matter experts are the differentiator. The distinction matters because a compliance reviewer catches errors. An expert HITL shapes the workflow, identifies what the AI cannot see, and interprets output in clinical and commercial context. Jag Chhatwal from Harvard put it directly: HITL is not optional, it is absolutely necessary — but the quality of the human input determines the quality of the output.

The Cube RM human-AI collaboration matrix made the hierarchy explicit. Human direction plus AI with human optimisation: excellent results. Human direction plus AI alone: acceptable but limited. AI only: limited impact with hallucination risk. Human only: limited impact. The peak performance is not AI replacing human judgment. It is expert human judgment directing and evaluating AI capability.

Astellas described their team culture as one of inquisitive attitude, thoughtful evaluation, and willingness to experiment — but always with expert oversight. Sanofi framed responsible AI as HITL across the entire value chain, not at the end of it.

> **Verdict:** Compliance HITL protects you. Expert HITL is where the competitive advantage lives. The industry is investing heavily in the former and underinvesting in the latter.

TAKEAWAY 3

The tools are not the barrier. The thinking is.

This was the most direct statement of the conference, and it came from BMS: the LLMs are not the barrier — how you think and how you use the system are. It was not a throwaway line. It was the conclusion of a presentation documenting BMS's progression from crude experiments to sophisticated end-to-end tools across a single year.

The implementation failures described across sessions were not technical. They were structural. Organisations deploying AI without decomposing workflows into tasks and sub-tasks. Teams adopting tools without understanding why. Vendors being paid for pilots when organisations needed solutions. The Harvard session named it explicitly: AI literacy is essential. Not AI use — AI literacy. The ability to understand what the system can and cannot do, and to make strategic decisions accordingly.

Sanofi's Agile methodology — test, learn, fail fast, iterate, with stage gates at every step to justify value — is not a complicated framework. It is disciplined thinking applied to implementation. The chair's remark on day two that AI is not a concept but a component of the value chain for better decision-making captures the shift: the question is no longer whether to adopt AI, but how to embed it into decisions that already happen every day.

The Daniel Moreira session raised an important question about vendor framing — whether 'agentic AI' is genuinely new architecture or RAG functionality repackaged in more marketable language. The answer matters for procurement decisions. A knowledge base with a query layer is not an autonomous agent, regardless of how it is presented. Organisations need enough AI literacy to ask the right questions of their vendors.

> **Verdict:** The organisations pulling ahead are not the ones with the biggest AI budgets. They are the ones that understand the problem clearly enough to deploy the right tool at the right step.

TAKEAWAY 4

The regulatory and HTA layer is about to be redefined by AI — whether the bodies are ready or not.

The Blueprint Medicines session was the most forward-looking of the two days, and arguably the most consequential for anyone working in market access strategy. The core argument: AI has the capability to solve the two structural gaps that cause most rare disease submissions to fail — the absence of a comparator and fragmented evidence.

AI-generated natural history models and synthetic control arms are not hypothetical. They are being developed and submitted. The implications are significant. If AI can turn fifty patients' worth of data into an evidence base that powers a dossier — by synthesising digital monitoring data, ML digital twins, and continuous multi-modal records into a single coherent dataset — then the boundary between AI tool and regulatory evidence is already moving.

Oliver Van Zon raised the consequence directly: what NICE and G-BA accept as standard of care may soon be defined by an AI model. That is a paradigm shift in how dossier strategy is built. The HTA bodies are aware of it — NICE has a positioning statement, HTAi has a policy lab, the FDA and EMA are developing frameworks — but awareness is not the same as readiness.

The panel on the European Health Data Space surfaced a parallel tension. Real-world evidence and digital behaviour change data are increasingly central to HTA submissions, particularly for medical devices and digital health interventions. The infrastructure to collect, link, and use that data is uneven across countries. Norway is ahead. The UK is behind. That asymmetry will affect which markets get prioritised for AI-enabled evidence generation — and which patient populations wait longer.

> **Verdict:** The regulatory environment for AI-generated evidence is in active formation. Organisations that engage with HTA bodies now — rather than waiting for frameworks to be finalised — will shape the standards they eventually have to meet.

TAKEAWAY 5

Data quality is the unglamorous foundation. It will determine who wins.

Every session that described a working AI system eventually arrived at the same prerequisite: the data layer has to be right. Sanofi named it as the first enabler: AI-ready data, structured and in the right place. Crinetics described poor data quality as the single greatest barrier to AI effectiveness in pricing and market access. The d-fine knowledge graph presentation framed structured data as the foundation of traceability, explainability, and hallucination reduction.

This is not a new insight. But its repetition across sessions by practitioners who have already built and deployed systems carries a different weight than the same observation from a vendor pitch. These are organisations that have invested in AI, run pilots, scaled operational models — and arrived at the same conclusion: the AI is only as good as what you feed it.

The European Health Data Space panel added a geopolitical dimension. The UK's fragmented approach to data collection and linkage — explicitly contrasted with Norway's integrated electronic health record and disease registry infrastructure — means that AI capabilities in real-world evidence generation will not be uniformly available across markets. Countries with richer, better-linked data assets will generate better AI-driven evidence. The gap will compound over time.

For organisations assessing their AI readiness, the data question is often treated as an infrastructure problem to be solved later, after the tools are chosen. The evidence from this Congress suggests the opposite: data infrastructure is the strategic decision. Everything else is conditional on it.

> **Verdict:** Investment in AI tools without investment in data infrastructure is investment in the wrong layer. The organisations that will lead AI-driven market access are the ones that treat data governance as a strategic priority, not a technical afterthought.

The Verdict from Two Days

The World EPA Congress AI & Digital Transformation track was not a room discussing the future. It was a room of practitioners comparing notes on what is already working, what is failing, and what needs to be built next. The use cases are live. The architecture debates are real. The governance gaps are documented.

The gap that ran through every session was not technological. It was structural: the implementation expertise to deploy AI correctly, the domain knowledge to evaluate its output, and the governance frameworks to embed it responsibly into decisions that affect patients and organisations alike.

That gap is the work. It is also the opportunity — for organisations willing to invest in the right layer, and for the people who can operate at the intersection of deep domain expertise and practical AI capability.

> ***AI in pharma market access is no longer a concept. The question is no longer whether. It is who has the expertise to make it work — and who is still waiting for certainty that will never come.***

I launched ryanbishop.co.uk at the World EPA Congress, March 2026. I work at the intersection of pharmaceutical experience in HEOR, market access, and HTA, and AI strategy consulting. Analysis, tools, and further writing at ryanbishop.co.uk.

Source

AI & Digital Transformation Track. World EPA Congress, Amsterdam, 3-4 March 2026.

Available at: https://www.terrapinn.com/conference/pharma-pricing/agenda.stm

[Accessed 7 March 2026]

Quick Reference: Five Verdicts

1. Generic AI fails HTA. Domain-specific architecture — RAG, Knowledge Graphs, guardrailed workflows — is the answer. Most organisations haven't built it yet.
2. HITL is non-negotiable. But compliance HITL and expert HITL produce very different outcomes. The competitive advantage is in the latter.
3. The tools are not the barrier. The thinking is. AI literacy — understanding what AI can and cannot do — is now a core professional competency in market access.
4. The HTA regulatory layer is in active formation around AI-generated evidence. Organisations that engage now will shape the standards they will eventually have to meet.
5. Data quality is the foundation. Every working AI system in this space depends on structured, governed, accessible data. Everything else is conditional on getting this right.

World EPA Congress 2026 | Amsterdam | AI & Digital Transformation Track | Ryan Bishop | ryanbishop.co.uk