AI-Driven Medical Science Liaison Enablement Platform: Evidence from a Real-World Implementation of Gen AI in Pharmaceutical Medical Affairs in China

Adele Li, MBA, David Wang, MBA, Yixuan Zhou, MSc.

MSR7

Background

Medical Science Liaisons (MSLs) are the principal scientific interface between pharmaceutical companies and healthcare professionals (HCPs). In China the role is under structural strain: physician demand is dispersed across a tiered hospital system, guideline and congress evidence expands continuously, and HCPs increasingly consult general-purpose AI tools themselves.

Field diagnostics identify three failure points. Before the visit, approved content, literature and congress data sit in separate repositories, so preparation depends on manual search. During the visit, answer consistency tracks personal familiarity rather than a defined evidence standard. After the visit, insight stays in individual notes rather than becoming an organisational asset.

Objective

To evaluate the real-world value and outcomes of deploying a generative AI-driven MSL enablement platform (eMSL Toolkit) across cardiovascular (CV) and central nervous system (CNS) therapy areas in China, assessing its impact on operational efficiency, knowledge quality, compliance governance and medical affairs return on investment.

Methods

An HEOR impact case methodology was applied to a multi-phase implementation across a pharmaceutical company's China MSL team (n=15–25, Phase 1 pilot; CV + CNS, eight products).

The platform uses an LLM + Retrieval-Augmented Generation (RAG) architecture over three knowledge sources: internally approved medical content, a monthly-updated PubMed database, and 30+ monitored congress sources. Hybrid vector and keyword retrieval feeds a medical-domain model under enforced citation, in-label / off-label partitioning, watermarking and full audit logging.

KPIs were captured across four layers: adoption (WAU), efficiency (response and preparation time), quality (accuracy, citation traceability, hallucination rate) and return on investment. Benchmarking used a comparable Digital MSL deployment at a multinational pharmaceutical company — six therapy areas, 3,755 approved internal documents, 14,000+ PubMed articles — and its 74-day proof of concept.

Figure 1. Platform architecture — LLM + RAG over a governed three-source knowledge base

Compliance controls are embedded in the retrieval and generation path, not applied after the fact

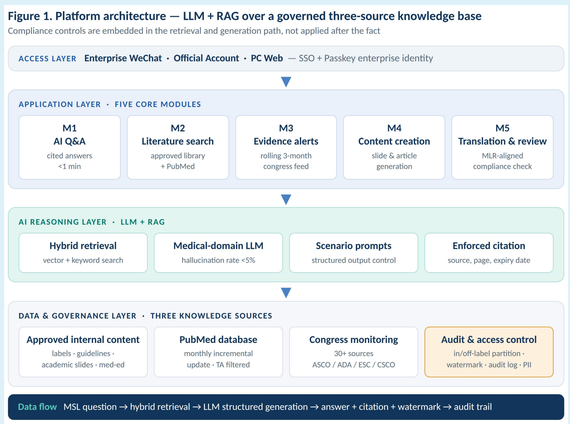

Results

Phase 1 pilot results demonstrated measurable operational value across all four KPI layers. HCP inquiry response time improved from approximately three hours to under one minute (>95% reduction); weekly active users reached 95%+; MSL visit preparation time decreased by more than 50%; answer accuracy exceeded 95% with 100% citation traceability. Phase 1 return on investment was estimated at approximately 500% against a development cost of approximately CNY 200,000.

Figure 2. HCP inquiry response pathway before and after deployment

Phase 1 pilot, CV + CNS therapy areas, MSL team n=15–25

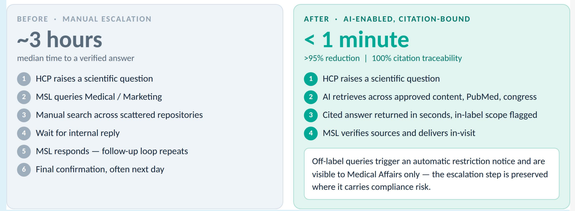

The response-time gain comes from collapsing a multi-step manual escalation chain into a single retrieval event. The escalation path is retained selectively: off-label queries route to Medical Affairs rather than being answered in the field, so the compliance control that the old workflow enforced by delay is now enforced by design.

Figure 3. Phase 1 pilot outcomes across the four-layer KPI framework

Measured over the pilot operating period; bars show attainment against pre-specified target

Adoption was near-universal within the seed team, which matters for interpretation: efficiency and quality gains are measured on routine practice rather than on a self-selected subgroup. Quality metrics held at target while query volume scaled, indicating that retrieval grounding rather than volume throttling is what constrains hallucination.

Figure 4. Deployed interface — every answer returns with its source set

Live system, Chinese-language deployment. PC workspace (left) and Enterprise WeChat mobile entry (right)

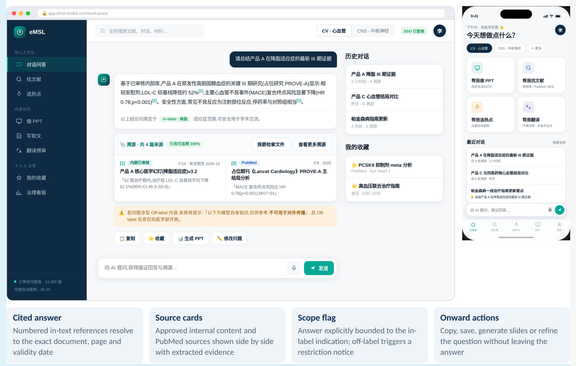

Cited answer — Numbered in-text references resolve to the exact document, page and validity date

Source cards — Approved internal content and PubMed sources shown side by side with extracted evidence

Scope flag — Answer explicitly bounded to the in-label indication; off-label triggers a restriction notice

Onward actions — Copy, save, generate slides or refine the question without leaving the answer

Knowledge quality and compliance governance were treated as measured outputs rather than as design assumptions. Answer accuracy exceeded 95% and hallucination remained below 5%, sustained through hybrid retrieval over approved content and enforced citation rather than through model scale alone.

Figure 5. Governance dashboard — evidence quality monitored as a standing metric

Four-dimension evidence scorecard, source composition and compliance pre-review volume

Four-dimension scorecard — Clinical rigour, evidence-citation rate, citation fidelity and professional-source share tracked continuously

Source composition — Answer provenance split across approved internal content, PubMed and congress data

Pre-review interception — Compliance issues flagged before human MLR review; medical staff retain final decision

The governance dashboard converts compliance from a periodic audit into a continuously observable metric. Source composition shows that most answer provenance resolves to approved internal content, with PubMed and congress data supplementing rather than substituting for it. AI pre-review intercepts issues ahead of human MLR review, and medical staff retain the final decision.

Figure 6. Benchmarking against a comparable Digital MSL deployment

Reference case: multinational pharmaceutical company, China medical affairs

Dimension	Reference Digital MSL deployment	eMSL Toolkit (this implementation)
Scope	6 therapy areas, enterprise-wide medical affairs	CV + CNS, 8 products, seed MSL team n=15–25
Knowledge base	3,755 approved internal documents; 14,000+ PubMed articles; 30+ congress sources	Same three-source design; TA-scoped approved library with brand-level governance
Positioning	Visit support tool — Q&A, literature, alerts, slide generation	Full-pathway enablement — pre-visit, in-visit and post-visit, plus medical content strategy
Evidence quality	Answers traceable; hallucination <5%; no systematic quality scoring	Four-dimension scorecard: clinical rigour, evidence-citation rate, citation fidelity, professional-source share
Deployment time	74-day proof of concept to live system	6–8 week MVP; each phase independently evaluable and stoppable
Reported outcomes	Response 3 h → <1 min; WAU 95%+; adoption 80%+; ROI ~500%	Phase 1 replicates efficiency, adoption and quality outcomes at TA scale

Benchmarking against a comparable Digital MSL deployment at a multinational company shows the efficiency and adoption outcomes to be reproducible rather than site-specific. This implementation adds systematic evidence-quality scoring and brand-level knowledge governance, which the reference case did not measure.

Figure 7. Three-phase deployment model with staged value capture

Capability accumulates; each phase carries its own KPI gate

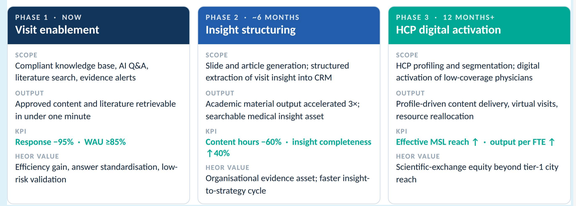

Because capability accumulates across phases, the value case is cumulative as well: Phase 1 efficiency alone clears the payback threshold, while insight structuring and HCP digital activation add value categories that are conventionally difficult to fund on their own.

Figure 8. Translating workflow KPIs into HEOR value categories

Operational metrics mapped to the value dimensions relevant to medical affairs resourcing decisions

KPI LAYER	MEASURED OUTCOME	HEOR VALUE CATEGORY
LAYER 1 **Adoption**	WAU 95%+, full registration coverage, sustained multi-module use — the tool is embedded in routine practice rather than piloted.	**Implementation feasibility** Precondition for any downstream value claim
LAYER 2 **Efficiency**	HCP inquiry response ~3 h → <1 min; visit preparation ~50%+; material production hours ~60%.	**Resource productivity** MSL capacity released to scientific exchange
LAYER 3 **Quality**	Accuracy >95%, citation traceability 100%, hallucination <5%, in-label / off-label access enforced at retrieval.	**Evidence integrity** Standardised, auditable scientific communication
LAYER 4 **Business value**	Phase 1 ROI ~500% against ~CNY 200,000 development cost; satisfaction ≥8/10; scale-up adoption 80%+ in the benchmark case.	**Return on investment** Quantifiable case for medical affairs AI investment

~500% PHASE 1 ROI · **~CNY 200k** DEVELOPMENT COST — A low absolute investment threshold means the efficiency gain alone clears the payback bar — quality and insight value accrue on top rather than being required to justify the build.

Mapping each KPI layer onto a value category makes the framework portable to other medical affairs AI deployments, enabling cross-implementation comparison rather than isolated case reporting.

Conclusions

AI-driven MSL enablement platforms deliver quantifiable, multi-dimensional value in pharmaceutical medical affairs: efficiency gains, quality standardisation and scalable HCP engagement.

The compliance-embedded LLM+RAG architecture addresses pharma-specific constraints — content governance, in-label / off-label access control and audit traceability — by building them into the retrieval path rather than layering review on top.

These outcomes suggest that HEOR frameworks for pharmaceutical operations should incorporate AI-enabled workflow transformation as a measurable value driver, with implications for medical affairs resourcing, HCP interaction quality and, ultimately, evidence-to-practice translation at scale.

Reference

1. eMSL Toolkit: AI-Driven Medical Science Liaison Enablement Platform. Project Plan v3, 2026.
2. Digital MSL system operating documentation, multinational pharmaceutical company, China Medical Affairs, 2025.
3. Gen AI HCP inquiry assistant: 74-day proof-of-concept outcome report, China, 2025.
4. 2025 Digital Life Physician (DLP) Survey. Base: n=5,000 Chinese physicians; AI users n=4,246.

Acknowledgement

The authors thank the eMSL Toolkit implementation team, the participating Medical Affairs and MSL teams, and the compliance and digital functions whose review shaped the governance design. Benchmarking context was informed by the Digital Life Physician research programme, conducted in collaboration with JKT (健康通).