Pharma GEO

The Pharma GEO *Playbook*

Table of contents

Introduction

What this is. A practitioner's reference for Generative Engine Optimization in the regulated pharmaceutical context. Synthesized from the Princeton GEO foundation paper, 2025–2026 vendor research (Profound, AthenaHQ, Astiva, friction AI, Evertune), peer-reviewed clinical-AI evaluations, FDA/EMA guidance, aikka's proprietary PharmaGEO research program (consumer-LLM prompt audits across brand and disease prompts in multiple markets), and the public documentation that HCP clinical answer engines (OpenEvidence, DoxGPT, UpToDate Expert AI, Vera Health, Glass Health, Medwise, and others) have themselves disclosed about their corpora, partnerships, and citation behavior.

Audience. Medical, marketing, market-access, regulatory and digital teams at pharmaceutical, biotech and medtech companies, and the agencies who serve them.

01 Executive doctrine: eight non-negotiable principles

1. GEO is patient safety, not just marketing. Independent benchmarks find general-purpose LLMs answer only 29–75% of clinical drug questions correctly, with documented dosing errors of 1,000× and fabricated boxed warnings (Intuition Labs 2026; JMIR AI 2025). If authoritative content does not dominate AI responses, hallucinated content fills the void.

2. Authority compounds, and 68% of citations go to the top 15 domains. The concentration of AI citations is more extreme than Google PageRank ever produced (Everything-PR 2026). Earning placement on those 15 domains, in the right form, is the single highest-leverage GEO move.

3. In healthcare, that "top 15" looks different. For health queries, over 75% of ChatGPT health citations come from established institutions: Mayo Clinic, Cleveland Clinic, NIH, WebMD, MedlinePlus. For HCP clinical AIs, the hierarchy is even tighter: NEJM → JAMA → Lancet/BMJ → NCCN/ACC/AHA/NICE → FDA labels → Cochrane → society guidelines. PubMed indexing is the common denominator.

4. The brand-name penalty is real and measurable. All major LLMs lose ~4% accuracy when drug names are switched from INN to brand, with some losing 10%+ (RABBITS / Gallifant 2024). GPT-4o passes only 67% of brand-to-generic conversion tests against an 80% threshold (John Snow Labs 2024). Brand and INN must be linked by your content, not by the model's guesswork.

5. Brand mentions matter more than backlinks. Branded web mentions correlate with AI citations at r=0.664, three times stronger than backlinks (r=0.21) (SparkToro / Astiva 2026). Earned coverage and named-entity ubiquity are the new domain authority.

6. Each engine has a different retrieval backend. Claude is 86.7% aligned with Brave Search (isagentready.com). ChatGPT is grounded in Bing. Gemini AI Overviews overlap 93.67% with Google top-10 (The Digital Bloom 2026). Perplexity runs a proprietary 200B-URL index. Cross-platform citation overlap is only ~11%. Single-engine optimization leaves 89% of the AI answer layer untouched.

7. HCP-facing AI is a fundamentally different game from consumer AI. OpenEvidence, DoxGPT, UpToDate Expert AI, Glass Health and Vera Health all run RAG over **gated, curated** corpora. Pharma-owned websites are mostly invisible to these engines. The only levers that move them are guideline inclusion, peer-reviewed publication, FDA-label clarity, and editorial inclusion in named drug databases.

8. Every GEO asset is a promotional asset. FAQ pages, JSON-LD `Drug` schemas, llms.txt summaries, Wikipedia contributions, sponsored CME: all are subject to MLR/PRC certification under FDA, EMA, ABPI, ANSM and equivalents. The 2026 FDA warning letter on AI overreliance makes clear that "reliance on AI is not a defense against regulatory violations."

02 The new search architecture in one page

Modern AI assistants answer in three different modes, and each requires a different optimization tactic.

Mode A: Parametric (no real-time search)

The model answers from training-data weights alone. ~60% of ChatGPT queries take this path (The Digital Bloom 2026). Wikipedia (~22% of LLM training data), Reddit (licensed) and high-DA publishers dominate. **Win mode A** by being broadly syndicated and entity-resolved across the web graph at training cutoff.

Mode B: Retrieval-Augmented Generation (RAG)

The model issues a real-time query to an index (Bing, Brave, Google, proprietary), retrieves top-k chunks, and inserts them into the prompt. Five stages: query encoding → hybrid semantic + BM25 retrieval → cross-encoder reranking (NDCG@10 +28%) → context injection → citation attribution (Astiva 2026). Perplexity triggers this every query; ChatGPT browsing triggers on "fresh" or "specific" intent. **Win mode B** by ranking in the retrieval index, structuring content for clean extraction, and supplying canonical, cite-able paragraphs.

Mode C: Tool-use and agentic search

Advanced models (GPT-5.4 Thinking, Claude Opus 4 with web search) issue 8.5 sub-queries on average per prompt and loop up to 10 times. Claude (February 2026) writes Python to strip navigation and ads from HTML before reading, buried content gets discarded (PassionFruit 2026). GPT-5.4 cites brand sites at 56% vs. 8% for GPT-5.3. **Win mode C** with main-content clarity, query-fan-out friendly site structure (independent product / pricing / comparison / safety pages), and clean machine-readable HTML.

The HCP clinical-AI mode (the fourth mode)

Distinct architecture: closed-corpus RAG. OpenEvidence runs over 35M+ peer-reviewed papers + NEJM/JAMA/NCCN/Cochrane partner content. UpToDate Expert AI and DynaMed AI cannot reference anything outside their editorial databases. The base LLM is invisible to optimization, the corpus is the lever. **Win the HCP mode** through guideline inclusion, peer-reviewed publication strategy, and structured drug-monograph data.

03 What signals AI engines actually reward

From Princeton GEO + 1.2M-citation Kevin Indig analysis + Astiva 1,247-brand study + Everything-PR 680M-citation index:

Signal	Finding	Why it matters in pharma
Branded web mentions	r=0.664 with citation rate (SparkToro / Astiva)	Earned media + congress coverage + KOL mentions feed the parametric layer
Domain authority (DR)	r=0.21, secondary signal	Important but not sufficient; DA without entity recognition fails
Adding citations to authoritative sources inline	+115.1% AI visibility for mid-ranked pages (Princeton GEO, KDD 2024)	Linking out to NEJM, NCCN, FDA label increases your own citation rate
Statistics with named source + year	+41% Position-Adjusted Word Count (Princeton GEO)	"ARR 2.3% (NCT04XXXXXX, 2024)" beats "demonstrated efficacy"
Named expert quotations	+29% Subjective Impression (Princeton GEO)	KOL pull-quotes with credentials are extracted preferentially
Question-format H2/H3	Strong RAG retrieval alignment	"Who is X indicated for?" beats "Indication overview"
FAQPage / Drug / MedicalCondition JSON-LD	High extraction correlation	Schema-marked answers are parsed cleanly by every engine
Person schema (author credentials)	+2.1× citation rate on Claude (Astiva Q1 2026)	MD, board certification, institution improves trust
Front-loaded structure	44.2% of cited content from first 30% of page (PassionFruit 2026)	Direct-answer paragraph in sentence one
Entity density	Cited text averages 20.6% entity density (3–4× normal English)	Drug names, conditions, mechanisms, trial names by name
Readability (Flesch grade ~16)	Outperforms grade 19+ prose	"Plain SmPC" outperforms turgid medical writing
Review-platform presence	3× citation probability	Drugs.com, GoodRx, Healthgrades, Reddit pharmacy threads
Multi-platform presence	2.8× more cited if on 4+ platforms	Owned + Wikipedia + Drugs.com + PubMed beats owned-only
Freshness (`dateModified`)	Strong recency bias on Ahrefs analysis	Quarterly content refresh maintains citation share

Tactics that fail: keyword stuffing, content padding, persuasive-only language, simplification alone (Princeton GEO, Table 2).

04 The two-channel reality: consumer LLMs vs. HCP clinical AIs

Pharma faces two architecturally different AI surfaces, and treating them as one is the most common GEO mistake.

Channel 1: Consumer LLMs (ChatGPT, Gemini, Perplexity, Claude, Copilot)

What dominates: Wikipedia, Reddit, Mayo Clinic, WebMD, Drugs.com, GoodRx, NHS, NIH, manufacturer DTC sites (US-only when brand is queried).

Triangulation insight from consumer-LLM audits. Running identical clinical and brand prompts across the major consumer LLMs reveals consistent patterns:

- **Brand vs. INN is the single strongest determinant of source tier.** Querying "Keytruda" surfaces manufacturer-controlled domains in top 2–3 results. Querying "pembrolizumab" entirely displaces manufacturer content in favor of NIH StatPearls, NHS, NCI Drug Dictionary, and FDA labels. The same shift is observed for Humira/adalimumab and Eliquis/apixaban.
- **Consumer vs. HCP framing produces a measurable source-tier jump.** Consumer GLP-1 queries surface WeightWatchers and Reddit; HCP-framed oncology queries return ASCO Publications and AACR, a full qualitative tier difference driven entirely by clinical terminology.
- **Geographic routing is language-driven, not country-flag driven.** French text surfaces HAS, Vidal, SFD. German text surfaces AkdÄ, DDG, Techniker Krankenkasse. Japanese text surfaces local clinical commentary translating ADA guidelines. Manufacturer content rarely appears in top results in non-US languages, except in Brazilian Portuguese, where local manufacturer content was observed surfacing.
- **No refusals observed** for clinical questions on Perplexity Sonar, including pregnancy drug safety, off-label metformin, and specialist-framed prescribing decisions. Safety disclaimers scale with clinical risk, not with topic sensitivity.
- **Citation transparency varies dramatically.** Perplexity is the only major engine that mandates URL-level citation in default mode. ChatGPT-4o and Google Gemini do not provide citations in standard responses (JMIR AI 2025). Microsoft Copilot does cite sources, primarily Drugs.com, NHS, and WebMD.

Pharma implication. Consumer LLM optimization is principally an **earned + Wikipedia + structured-DTC-site** play in the US, and a **disease-awareness + HCP-syndication** play in markets where Rx DTC is prohibited.

Channel 2: HCP clinical AIs

What dominates: Closed corpora, NEJM, JAMA, Lancet, BMJ, NCCN, ACC/AHA, NICE, ESMO/ASCO, FDA labels, Cochrane, UpToDate editorial team, Micromedex, Lexicomp, Pathway knowledge graph (now in DoxGPT).

Architecture observations:

- **OpenEvidence** runs RAG over 35M+ peer-reviewed publications with NEJM/JAMA/NCCN/Cochrane content partnerships. Citation hallucination rate ~2.3% (vs. 40–80% for general LLMs). Used by ~40% of US physicians (Contrary Research 2026). Pharma advertising is structurally separated from the answer engine by design.
- **DoxGPT** layered the Pathway knowledge graph (acquired August 2025, $63M) on top of a general LLM. Now includes 3,200+ drug monographs. Selected as best clinical answer in 61% of side-by-side evaluations vs. OpenEvidence (26%) (Doximity 2025). Distributed to 80%+ of US physicians.
- **UpToDate Expert AI** and **DynaMed AI** are strictly closed-corpus, answers cannot reference anything outside the editorial database.
- **Glass Health** uses RAG plus a structured clinical-reasoning layer; cites named guidelines (AAP, NICE, SIGN, AHA, Cochrane) inline with "High Impact / Highly Cited" tags.
- **Vera Health** runs RAG over 60M+ peer-reviewed papers with ACEP partnership (38,000+ emergency physicians).
- **Medwise** indexes local NHS trust guidelines + NICE, UK-primary.

The triangulation insight. Across all HCP clinical AIs, only three levers reliably change what the answer engine retrieves and cites:

1. **Guideline inclusion** (NCCN / ACC / AHA / ESMO / ASCO / NICE / HAS / NCCN-IPC), guarantees indexed retrieval across nearly every platform.
2. **High-impact journal publication** (NEJM / JAMA / Lancet / BMJ), the second tier; Cochrane systematic reviews adjacent.
3. **FDA-label precision and DailyMed structured indexing**, labels are directly indexed; precise indication, contraindication, and warning language is reproduced verbatim.

Levers that do not move the engines:

- Sponsored advertising (visibility yes, retrieval no, by design)
- MSL outreach (no retrieval impact)
- Congress abstracts (mostly invisible to clinical AI corpora)
- General DTC content (out-of-corpus by design)

Hierarchy of HCP-channel investment, ranked:

> Guideline inclusion > NEJM/JAMA/Lancet/BMJ publication > Cochrane systematic review > FDA-label precision (DailyMed) > UpToDate / DynaMed editorial inclusion > platform content partnerships > specialty-society journal publication > advertising placement > congress abstracts (negligible) > MSL outreach (negligible for AI).

05 Engine-by-engine playbook

Engine	Backend	Distinguishing trait	Pharma optimization priorities
ChatGPT (GPT-4o / GPT-5.x)	Bing-grounded RAG + parametric	87% Bing top-10 alignment (Digital Bloom); Wikipedia 47.9% of citations; 60% queries answered without search	Bing index via IndexNow; GPTBot allow-list; Wikipedia entity completeness; Tier-2 publisher coverage (Forbes, NYT, WSJ); query-fan-out structure (independent product/safety/dosing pages)
Claude (Sonnet 4 / Opus 4)	Brave Search (86.7% alignment)	Three crawlers: CLAUDE-SEARCHBOT (block at your peril), CLAUDEBOT (training, optional), CLAUDE-USER. Inline citations with source URL + 150-char extract. Feb 2026 dynamic Python filtering of HTML	Rank in Brave Search; allow CLAUDE-SEARCHBOT in robots.txt; conversational, question-answering style; main-content clarity (no nav/footer noise); Person schema (+2.1× citation lift)
Gemini (AI Overviews + app)	Google organic + proprietary	93.67% AI Overview top-10 organic overlap; Gemini 3 (Nov 2025) replaced ~42% of cited domains; YMYL standards aggressively applied; AI Overviews on 50%+ of US searches	Top-10 Google rankings; full structured data (FAQPage, HowTo, Article + dateModified); YouTube indexing with full transcripts; Search Console "AI Appearances" report
Perplexity (Sonar / Pro / Deep Research / Comet)	Proprietary 200B+ URL index	Real-time RAG every query; 5–10 inline citations; Reddit 46.7%, YouTube 13.9% of top citations; Deep Research mode pulls 20–30+ sub-queries with PubMed/Cochrane/NEJM weighting	Get into Perplexity's index (independent crawler); Q&A-structured content; PubMed-indexed publications referencing brand and INN; legitimate Reddit engagement (not astroturfing)
Microsoft Copilot	Bing (similar to ChatGPT Search)	IndexNow latency directly affects Copilot timeliness; Wikipedia ~35%; deeply embedded in Microsoft 365 (enterprise hospitals = Copilot users by default)	IndexNow; Bing SEO discipline; same Wikipedia + reputable third-party priorities
Meta AI	Facebook/Instagram/WhatsApp/Ray-Ban	Heavy Meta-property + social-proof draw; relevant for DTC + disease awareness with social audiences	Branded social presence; Wikipedia; structured product data on Marketplace where applicable
Mistral Le Chat	Proprietary (EU emphasis)	Stronger presence in France, Germany; GDPR-positioned for EU HCPs	EU language coverage; alignment with EU regulator content (EMA, ANSM, BfArM, AIFA, AEMPS)

Engine	Backend	Distinguishing trait	Pharma optimization priorities
Grok (xAI)	X/Twitter integration	Real-time social signal; useful for adverse-event sentiment monitoring	Social listening + brand-name standardization across X conversations
DeepSeek	Proprietary; APAC strength	Matters in JP, KR, CN	Native-language content; APAC regulator alignment (PMDA, MFDS)

06 Brand vs. molecule: the most under-managed lever in pharma GEO

This is the single highest-leverage tactical area in pharma GEO and the one most teams are not measuring.

What the data show

- LLMs lose 4% accuracy on average, and up to 10%+ on some classes, when drug names switch from INN to brand (RABBITS, Gallifant et al. 2024).
- GPT-4o: 67% pass rate on brand-to-generic conversion (target threshold 80%) (John Snow Labs 2024).
- Querying brand surfaces manufacturer content; querying INN surfaces NIH StatPearls, NHS, MedlinePlus, FDA labels, observed consistently for Keytruda/pembrolizumab, Humira/adalimumab, Eliquis/apixaban.

The implication

There are effectively **two distinct citation graphs** for every product, a brand graph and an INN graph, and they cite different sources. Most pharma teams optimize only the brand graph (manufacturer site, branded DTC) and leave the INN graph to chance. The INN graph is what HCPs and clinically literate consumers query.

The playbook

1. **Always co-name brand and INN** in every owned page, every press release, every schema field. Use `schema.org/Drug` with both `proprietaryName` and `nonproprietaryName` populated.
2. **Publish a canonical "[Brand] is [INN]" FAQ** with `FAQPage` schema, indexed and crawlable. Question phrased exactly as users ask: "Is Keytruda the same as pembrolizumab?"
3. **Ensure Wikipedia entity coverage for both names**, separate disambiguation, cross-linked via `sameAs` to Wikidata.
4. **Get DailyMed structured-product-label entries right.** SPL is the single most important machine-readable label asset; DailyMed entries feed institutional drug databases used by HCP AIs.
5. **For biosimilars,** publish side-by-side comparisons with the originator using the WHO Biological Qualifier framework. Cross-link FDA's Purple Book entry. LLMs conflate originator and biosimilar safety profiles by default, disambiguation is your responsibility.
6. **Test brand-and-INN-paired prompts monthly** across all five major engines. Track citation share separately for each name.

07 Market & regulatory split: US, EU5, APAC, LATAM

Pharma GEO is a **per-market** discipline, not a per-language one. The same language across two markets can require very different content (e.g., HAS-driven France vs. INAMI-driven Belgium for French; or G-BA Germany vs. SwissMedic Switzerland for German).

United States

Regulatory: FDA fair-balance (21 CFR 202); DTC permitted for Rx; FDA January 2025 AI guidance; 2026 FDA warning letter on AI overreliance.

Authority sources AI defers to: Mayo Clinic, Cleveland Clinic, NIH, MedlinePlus, FDA.gov, DailyMed, WebMD, Drugs.com, GoodRx, NEJM, JAMA, ASCO/NCCN guidelines, AHA/ACC, USPSTF.

HCP channel: OpenEvidence (40% of US physicians), DoxGPT (80%+ via Doximity), UpToDate Expert AI, Glass Health, DynaMed AI.

GEO priorities:

- DailyMed-perfect SPL files
- Mayo / Cleveland Clinic editorial inclusion (where possible)
- Wikipedia entity completeness for brand and INN
- Drugs.com and GoodRx structured data
- DTC content with fair-balance schema-marked
- AI Overviews tracking via Search Console "AI Appearances"

EU5 (UK, France, Germany, Italy, Spain)

Regulatory: EU Directive 2001/83/EC prohibits Rx DTC. EMA AI Reflection Paper (2024). EU AI Act high-risk classification for healthcare. ABPI, ANSM, BfArM, AIFA, AEMPS national codes. MLR/PRC certification mandatory before publication.

Authority sources:

- **UK:** NICE, NHS.uk, MHRA, BNF, SIGN
- **France:** HAS (Avis de transparence), ANSM, Vidal, SFD and other societies
- **Germany:** G-BA (Nutzenbewertung), BfArM, AkdÄ, DDG and other societies
- **Italy:** AIFA, ISS, national societies (SIE, SID, AIOM)
- **Spain:** AEMPS, SEMI, national societies

GEO priorities:

- Brand-neutral disease-awareness content (consumer surface)
- HCP-targeted content via professional channels (registered access)

- HTA-aligned language (NICE TA, G-BA Nutzenbewertung, HAS Avis), copy phrasing where compliant
- Native-language content for each country (AI engines preferentially cite native-language sources for queries in that language)
- ABPI / ANSM / BfArM compliance baked into content templates

APAC

Japan: PMDA approval documentation; Japanese Society of Diabetes / JSH / JCS guidelines; CareNet and major Japanese medical publishers. **APAC primary-source bias is strong:** 60–65% of AI citations in Japan are from official party / primary sources vs. 25–45% in the US (arXiv:2510.06823). Translated ADA / ESC content is consumed via Japanese clinical-commentary blogs.

China: Restricted digital environment; Baidu ERNIE, DeepSeek dominate. NMPA approval documentation; CNKI for Chinese-language literature. Western GEO frameworks do not translate directly.

South Korea: MFDS; KAMJE journals; DeepSeek, ChatGPT (post-2025 Korean availability), local engines.

Australia: TGA / ARTG; NPS MedicineWise; UK-style regulatory model, content priorities resemble the UK with TGA-specific branding.

India: CDSCO; growing English-medium AI use; PubMed-indexed Indian journals (Indian J Med Res); cost-effectiveness content particularly important.

LATAM

Brazil: ANVISA; CONITEC (SUS reimbursement protocols); Portuguese-language content; manufacturer content occasionally surfaces in top results, a regional anomaly worth pursuing for Portuguese-Brazilian markets.

Mexico, Argentina, Colombia, Chile: COFEPRIS (MX), ANMAT (AR), INVIMA (CO), ISP (CL); Spanish-language regulatory documentation often less digitally available than EU5 Spanish content. AI citations for LATAM health queries fall back to US/EU English sources by default, **first-mover GEO opportunity** for native-Spanish/Portuguese authoritative content.

08 Rx, OTC, reimbursed, and biosimilar: four very different games

Category	Primary surface	Top citation drivers	Specific tactics
Reimbursed Rx	HCP clinical AIs + HTA portals	NICE TA / G-BA / HAS / CADTH / PBAC; cost-effectiveness publications; HEOR data	Publish in ISPOR venues; align language with HTA decision documents; ensure cost-effectiveness numbers cite source studies
Rx (non-reimbursed)	HCP clinical AIs + medical journals	Clinical guidelines, KOL literature, FDA/EMA labels	Publication strategy in NEJM/JAMA/Lancet/specialty journals; KOL-authored review articles; congress publications archived to PubMed
OTC / consumer health	Consumer LLMs + retail review platforms	Reddit pharmacy threads, GoodRx, Drugs.com, Healthline, Mayo Clinic, NHS	Structured product data on retail platforms; Drugs.com monograph completeness; consumer-question FAQ schema; legitimate Reddit/community engagement
Biosimilars	HCP clinical AIs + payer tools + pharmacist tools	FDA Purple Book; EMA biosimilar listings; comparator publications; interchangeability data	Side-by-side comparison content with originator; explicit BQ / suffix usage; switch-study publications; payer-targeted comparative cost content

A note on Rx-DTC asymmetry. Because the US permits Rx DTC and most other markets do not, your global GEO posture is split: in the US, manufacturer DTC content can rank for branded queries; outside the US, the same content is regulatory non-compliant. **One global content base does not work.** You need market-segmented content trees with regulator-aware schema and indexing rules.

09 Owned-domain readiness checklist

Audit every brand and corporate site against this list. Each is a citation lever that AI engines specifically reward.

Technical foundation

- ☐ HTTPS, HTTP/2 or HTTP/3, fast TTFB (under 600ms p75)
- ☐ Mobile-first, accessible (WCAG AA minimum, Claude weights on accessibility signals)
- ☐ XML sitemap submitted to Google, Bing, IndexNow
- ☐ `robots.txt` with explicit allow rules for: GPTBot, ChatGPT-User, OAI-SearchBot, ClaudeBot, CLAUDE-SEARCHBOT, CLAUDE-USER, Google-Extended, PerplexityBot, Bytespider (or equivalent), Mistral, Bing
- ☐ `/llms.txt` published (Markdown, 2–3 sentence blockquote summary, H2 sections of 3–7 links each, returned as `text/markdown` or `text/plain`, HTTP 200, no auth gate)
- ☐ Optional: `/llms-full.txt` for developer-tooling and clinical reference companies

Content structure (per page)

- ☐ Page title is the user's question, verbatim where possible
- ☐ H1 reformulates the question; first paragraph is the direct answer (front-loaded structure: 44.2% of cited content sits in first 30%)
- ☐ H2 / H3 are question-format
- ☐ Stable noun phrases (drug name, INN, indication, mechanism, trial name) repeated naturally
- ☐ Statistics with source + year inline (Princeton GEO +41% PAW)
- ☐ Named expert quote with credentials (Princeton GEO +29% SI)
- ☐ At least 2–3 citations to authoritative external sources (Princeton GEO +115% for mid-ranked pages)
- ☐ Tables for comparisons (LLMs preferentially extract tables)
- ☐ Plain-language readability target Flesch grade ~16

Schema.org / JSON-LD

- ☐ `MedicalWebPage` (or `Article` + `MedicalEntity`) on every clinical content page
- ☐ `Drug` with `proprietaryName`, `nonproprietaryName`, `mechanismOfAction`, `clinicalPharmacology`, `indication` (`MedicalCondition`), `contraindication`, `warning`, `interactingDrug`, `legalStatus`
- ☐ `FAQPage` for indication / dosing / safety pages
- ☐ `Person` with `medicalSpecialty`, `worksFor`, credentials for all medical authors
- ☐ `Organization` + `MedicalOrganization` for the corporate entity
- ☐ `dateModified` updated on each MLR-approved revision

- [ ] `sameAs` linking entities to Wikidata, Wikipedia, RxNorm, ATC, MeSH, ChEMBL where applicable
- [ ] `MedicalCondition` `applicableLocation` to scope claims geographically

Entity discoverability

- [ ] Wikipedia article exists for both brand and INN; both are factually current and reference-supported (not edited by the company directly, engage editorial volunteers properly)
- [ ] Wikidata entry exists, linked via `sameAs`
- [ ] Drugs.com monograph present and accurate
- [ ] DailyMed SPL filed and current (US)
- [ ] EMA EPAR / national equivalents linked from EU pages
- [ ] PubMed-indexed publications use brand and INN consistently

Hygiene

- [ ] No prompt-injection vectors (block UGC on product pages, sanitize comments)
- [ ] No off-label content on indexable pages
- [ ] Quarterly content refresh cadence with updated `dateModified`
- [ ] MLR/PRC version control on every page
- [ ] Adverse-event reporting link on every product page

10 Earned and shared: where leverage actually lives

Owned content is necessary but not sufficient. Citation share is overwhelmingly captured by third-party domains. The earned + shared playbook:

Tier 1, Highest leverage

- **Guideline inclusion** (NCCN / ASCO / ESMO / ACC-AHA / NICE / HAS / G-BA / international equivalents). Multi-year endeavor; the single largest GEO lever for HCP AI.
- **High-impact journal publication** (NEJM / JAMA / Lancet / BMJ / specialty top-tier). PubMed-indexed.
- **Cochrane systematic review inclusion / authorship.**
- **DailyMed / EMA structured label precision.**

Tier 2, Strong leverage

- **Wikipedia entity completeness** (brand + INN, both current, both well-sourced).
- **Editorial inclusion in UpToDate / DynaMed / Lexicomp / Micromedex / Pathway** (DoxGPT now ingests Pathway).
- **Tier-2 publisher coverage** (Forbes, NYT, WSJ, Reuters, STAT News, Endpoints, FierceHealthcare). ChatGPT brand-mention parametric weight.
- **PubMed-indexed real-world evidence and Phase IV publications.**

Tier 3, Useful leverage

- **Patient advocacy organizations** (curated content alliances).
- **Specialty society journals and review articles.**
- **Drugs.com / GoodRx / Healthline / WebMD entries** (claim and verify; offer expert-reviewed content).
- **YouTube content** with full transcripts, structured descriptions, chapter markers, `VideoObject` schema. YouTube is the #3 most-cited AI source overall.
- **Reddit engagement** through legitimate physician/expert participation (not astroturfing, community detection is robust).

Tier 4, Foundational but limited GEO upside

- Congress abstracts (mostly invisible to clinical AI corpora, but valuable for parametric mention diffusion if the abstract is later published).
- MSL outreach (zero impact on AI retrieval, valuable for offline channels only).
- Sponsored content on HCP platforms (visibility yes, retrieval no, by design).

11 Measurement: KPIs and benchmarks

The core KPI set

KPI	Definition	Benchmark
AI Share of Voice (SOV)	% of responses to a defined prompt set that mention your brand	Top-quartile pharma in TA: >30%
Citation Rate	% of responses that include a link back to your domain	Top-quartile: >15%
Sentiment Score	Positive / neutral / negative tone of brand mentions	Baseline ≥80% positive-or-neutral; <60% triggers reputation review
Brand vs. INN Split	% of responses that use brand name vs. INN	Track separately; aim for ≥70% paired co-mention
Prompt Coverage	% of high-volume prompts where brand appears in top-3 positions	Target ≥50% across mapped category prompts
Competitive Displacement Rate	% of prompts where competitor is recommended instead of brand	Investigate when >20% in core indications
Attribute Accuracy	Correctness of AI's description of features / indications	Target ≥95% on key fields (indication, dosing, contraindication, warning)
HCP-AI Citation Share	Citation share specifically in OpenEvidence / DoxGPT / UpToDate AI / Glass / Vera	Track by therapy area; benchmark vs. guideline inclusion + NEJM/JAMA publication count

Prompt set construction

A defensible prompt set has the following structure, scaled to therapy area:

- **Awareness layer (50 prompts):** disease symptoms, condition definitions, "what is X" queries
- **Consideration layer (75 prompts):** treatment options, comparisons, "best treatment for X" queries
- **Decision layer (50 prompts):** brand-specific, dosing, side effects, switching, contraindications
- **HCP-clinical layer (75 prompts):** first-line, guideline-aligned, biomarker-stratified, comparative-effectiveness
- **Geographic variants:** repeat 25 highest-priority prompts in each market language
- **Brand-vs-INN pairs:** 20 paired probes
- **Competitor probes:** 30 prompts probing direct competitor mentions

Total: ~250–300 prompts per therapy area, refreshed quarterly, run weekly across 4–6 engines.

Measurement tooling

Platform	Strength	Engine coverage
Profound	Enterprise; prompt-volume data; ChatGPT Shopping analytics; tracks Meta AI	9 engines incl. Meta AI, DeepSeek
AthenaHQ	Multi-region (60+ countries); Google AI Mode coverage	8 engines
friction AI	E-commerce; closed-loop; purchase-intent	4 core engines
HubSpot AI Search Grader	Free entry-level	ChatGPT, Gemini
xFunnel	Citation analysis; brand-appearance tracking	ChatGPT-focused
Otterly	SMB; visibility + content generation	4–6 engines

For pharma specifically, dedicated platforms (Evertune AI and others) layer therapy-area context, regulatory-language detection, and adverse-event signal monitoring on top of the generic stack.

Refresh discipline

- **Weekly:** automated KPI run on the canonical prompt set across all engines.
- **Monthly:** sentiment-shift detection, competitor-displacement detection, adverse-event monitoring across engines.
- **Quarterly:** prompt set audit (new high-volume prompts emerging? new engines launched? new model versions?), content refresh on owned domains.
- **Annually:** full doctrine review (this document and equivalents), governance update, corpus partnership review.

12 Failure modes: pharmacovigilance, hallucination, off-label, prompt injection

Documented hallucination categories

Category	Documented example	Risk class
Fabricated clinical trials	GPT-4 created fake n=300 corneal-surgery dataset; concluded false benefit (Intuition Labs)	Reputational + scientific integrity
Hallucinated citations	47% of ChatGPT-3.5 medical citations completely fabricated; only 7% entirely correct	Trust erosion + liability
Regulatory misinformation	ChatGPT-3.5 correct on FDA boxed warnings for only 12/41 antibiotics; fabricated warnings for fidaxomicin, aztreonam	Patient harm + brand damage
Dosing errors	1,000-fold opioid intrathecal dose error (mg displayed instead of µg)	Critical patient safety
Drug interactions missed	Paxlovid/verapamil reported as no interaction (severe CYP3A inhibition reality)	Adverse drug event
Sodium/bromide substitution	Patient replaced dietary salt with sodium bromide after ChatGPT advice; suffered bromism	Preventable harm

Prompt injection on owned sites

IOActive's analysis defines prompt injection as malicious page text causing the LLM to act in unintended ways. For pharma, vectors include UGC sections, comments, reviews, embedded text in competitor-controlled pages or hijacked third-party content. **Countermeasure:** disable UGC on product pages; implement content-security policies; audit pages monthly for injected text; deploy automated detection of out-of-context instruction-like strings on owned domains.

Citation poisoning (PoisonedRAG)

arXiv:2510.06823 documents adversarial GEO: small volumes of malicious text injected into publicly accessible content can satisfy both retrieval and generation conditions, producing

attacker-intended answers. ~30% of US-query AI citations come from low-barrier sources (Reddit, blogs, personal sites) where manipulation is easiest. **Countermeasure:** anchor truth in high-barrier sources (.gov, peer-reviewed literature); monitor mention sentiment on low-barrier sources; engage proactive takedown procedures; publish authentic data with sufficient density that authoritative content out-competes fabrications.

Off-label surfacing

LLMs do not reliably constrain answers to approved indications. Any published literature (case series, hypothesis-generating papers, off-label use surveys) can be retrieved. **Countermeasure:** explicit indication-scoping language in all GEO content; populate `schema.org/Drug` `indication` only with approved indications; do not optimize for off-label query terms even where literature supports them; monitor weekly for AI responses suggesting off-label use of your products and route to Medical Affairs.

Pharmacovigilance signal capture

LLMs can inadvertently surface adverse-event reports and safety concerns aggregated from patient forums and case reports, sometimes amplified, sometimes accurate. Establish a monitoring protocol that flags AE-suggestive AI outputs to PV teams within 24 hours of detection; integrate with global PV intake systems.

13 Governance: how MLR/PRC must evolve

The principle

GEO assets (FAQ content, JSON-LD schemas, llms.txt summaries, press release language, Wikipedia contributions, sponsored CME) are **promotional materials in regulatory terms** and must pass MLR/PRC review before publication. The 2026 FDA warning letter on AI overreliance establishes that AI-generation tools accelerate drafting but cannot replace human regulatory review.

Required workflow updates

1. **Schema as approvable content.** JSON-LD `Drug` schema fields (`description` , `indication` , `warning` , `contraindication` , `interactingDrug`) require MLR-approved language. Build a schema content library managed at the same level as DTC creative.
2. **AI-generation logs.** Every AI-assisted GEO asset must carry a generation log (model, prompt, date, reviewer), supports FDA AI-overreliance mitigation.
3. **Schema diff in version control.** Every schema field change must trigger a re-review (MLR diff + sign-off).
4. **Quarterly AI-output audit.** A cross-functional team (Medical, Regulatory, MLR, Digital) reviews AI outputs across the 250-prompt set; discrepancies route to label-comparison and corrective-action workflow.
5. **Pre-flight risk classification on all GEO actions.** Three buckets:

- **Low-risk:** factual, on-label, fair-balance present, schema-marked.
- **Medium-risk:** comparative non-promotional, requires H2H or NMA evidence inline.
- **High-risk:** comparative promotional, off-label-adjacent, mechanism claims. Require senior MLR sign-off.

6. **Adverse-event detection in AI responses.** PV intake protocol for AE-suggestive AI outputs.
7. **Right-to-correct workflow.** Where AI engines surface incorrect information about a product, establish the engagement protocol with each platform (OpenEvidence, Doximity, UpToDate, Wikipedia editorial guidelines, etc.).

Roles and accountability

- **Digital / Marketing:** owns content production, schema, and earned-media strategy.
- **Medical Affairs:** owns publication strategy, KOL engagement, guideline-inclusion advocacy.
- **Regulatory:** owns label precision, DailyMed / EMA EPAR updates, AE-reporting compliance.
- **MLR / PRC:** owns content certification, schema review, AI-generated-asset audit.
- **PV:** owns AE-signal monitoring across AI engines.

- **Legal:** owns prompt-injection liability, takedown procedures, deepfake response.
- **A central GEO program lead** (typically Digital + Medical) coordinates across all of the above with a quarterly governance cadence.

14 The 90-day mobilization plan

A pharma organization can stand up a credible GEO function in 90 days. The sequence:

Days 0–15, Baseline

- Define therapy areas in scope and product portfolio
- Build the 250-prompt canonical set per TA
- Run baseline measurement across 4–6 engines (Profound / AthenaHQ / equivalent)
- Document current SOV, citation rate, sentiment, brand-vs-INN split, competitive displacement

Days 15–30, Audit

- Owned-domain audit against the readiness checklist (section 9)
- DailyMed / EMA EPAR / structured label audit
- Wikipedia entity audit (brand + INN, both markets)
- HCP clinical AI presence audit (OpenEvidence / DoxGPT / UpToDate AI / Glass / Vera)
- Guideline-inclusion gap analysis (NCCN / NICE / G-BA / HAS / equivalents)

Days 30–60, Intervention sprint

- Schema deployment on top 20 owned pages (FAQPage, Drug, MedicalWebPage, Person)
- llms.txt publication
- robots.txt updates (allow major AI crawlers)
- IndexNow integration for Bing/Copilot
- DailyMed SPL refresh
- Wikipedia editorial engagement (legitimate, declared)
- Tier-2 publisher placement (PR pipeline)
- Schema content library spun up under MLR governance

Days 60–90, Earned + governance

- Publication strategy alignment with HCP-AI corpora (NEJM/JAMA/Cochrane targeting)
- Guideline-inclusion engagement plan (multi-year)
- MLR workflow updates (schema-as-content, AI-generation logs, AE detection)
- Cross-functional governance cadence established
- Re-measure KPI set; document delta vs. baseline; build investment case for year-2 program

15 Forward-look: agentic browsers, multimodal GEO, regulatory convergence

Agentic browsers as the next distribution surface

ChatGPT Atlas (October 2025) and Perplexity Comet (2025) embed AI at the browser navigation layer. Persistent sidebar memory, agent action across tabs, and on-page summarization. For pharma:

1. Branded content must be **actionable at the page level**, Atlas agents extract and summarize on-page content directly.
2. **Persistent memory** means early, compliant brand exposure compounds across sessions.
3. **The funnel collapses**, users skip search engines for growing shares of intent. Branded content becomes the cited reference an agent surfaces.

Implication: GEO transitions from traffic optimization to **answer ownership.** The unit of optimization is the cited paragraph, not the page.

Multimodal GEO

Multimodal RAG extends retrieval to images, video frames, audio. For pharma:

- **Images:** clinical diagrams, MOA animations, patient-journey graphics, descriptive ALT text + caption + adjacent text are required for indexing. Claude 5 and DeepSeek V4 multimodal capabilities now extract image content directly.
- **Video:** YouTube transcripts are actively indexed by Google AI Overviews and Gemini. KOL interviews, MOA animations, clinical-data presentations should ship with full closed captions, structured descriptions, chapter markers, and `VideoObject` schema.
- **Voice:** AI voice answers are single-assertion, featured-snippet equivalent. Direct-answer-first structure is voice-optimized GEO by default.

Regulatory convergence

FDA (January 2025 draft guidance) and EMA (2024 AI Reflection Paper) are developing AI frameworks. The 2026 FDA warning letter on AI overreliance signals expanding scrutiny across pharmaceutical AI use cases. Expect specific FDA guidance on AI-generated promotional materials within 2026–2027. EU AI Act high-risk classification for healthcare AI applies to GEO content-generation tools used in pharma promotion (conformity assessment likely required).

The compliance cost of GEO will rise. The **competitive moat** for organizations that build governance early will rise even faster.

16 Glossary & primary references

Glossary

- **AEO**, Answer Engine Optimization; subset of GEO focused on direct-answer engines (ChatGPT, Perplexity, Claude, Gemini) as opposed to traditional Google.
- **GEO**, Generative Engine Optimization; the discipline of optimizing for citation in AI-generated answers.
- **HTA**, Health Technology Assessment; bodies like NICE, G-BA, HAS, CADTH, PBAC that evaluate cost-effectiveness for reimbursement.
- **INN**, International Nonproprietary Name; the generic/molecule name of a drug (e.g., pembrolizumab).
- **MLR / PRC**, Medical, Legal, Regulatory / Promotional Review Committee; the cross-functional review required before any pharma promotional content is published.
- **RAG**, Retrieval-Augmented Generation; the architecture where an LLM retrieves real-time content and synthesizes an answer.
- **SOV**, Share of Voice; the proportion of relevant AI responses that mention your brand.
- **SPL**, Structured Product Label; the FDA's machine-readable drug label format housed at DailyMed.
- **YMYL**, Your Money Your Life; Google's content-quality framework that applies stricter standards to health and financial content.

Primary references

Foundational research

- Aggarwal et al., "GEO: Generative Engine Optimization," Princeton/Georgia Tech/IIT Delhi/Allen AI, KDD 2024
- Princeton SAGEO follow-up, 2025
- Wu et al., Nature Communications 2025, LLM medical citation accuracy
- Gallifant et al., RABBITS, drug name accuracy, PMC 2024
- arXiv:2510.06823, Citation vulnerabilities in generative engines, 2025

Industry research and citation analysis

- Everything-PR AI Platform Citation Source Index 2026 (680M citations)
- The Digital Bloom 2025 AI Citation & LLM Visibility Report
- PassionFruit "How LLMs Search for Citations" (Kevin Indig, 1.2M citations)
- Astiva AI 2026 brand-citation study (1,247 brands)
- xFunnel GPT-5 citation pattern analysis (2.5M citations)
- isagentready.com, How Claude selects sources
- Profound GEO guide 2025
- AthenaHQ GEO Marketing Blueprint 2026
- friction AI visibility platform comparison 2026

Pharma-specific

- Evertune GEO in pharma, six steps
- Intuition Labs, LLM hallucinations in pharma 2026
- arXiv:2601.17109, Where ChatGPT health citations come from
- BrightEdge / Primary Intelligence, AI Health Search Accuracy
- John Snow Labs, drug name swapping in LLMs
- Forbes, FDA & EMA 10 guiding principles for AI in pharma
- Morgan Lewis, FDA warning letter on AI overreliance, April 2026
- EMA AI Reflection Paper 2024
- JMIR AI 2025, Drug-related LLM benchmark
- Frontiers in Medicine 2025, Infectious-disease MCQ benchmark

HCP clinical AI platforms

- OpenEvidence, Contrary Research breakdown
- Doximity / DoxGPT + Pathway acquisition
- Vera Health + ACEP partnership
- Glass Health features
- UpToDate Expert AI
- Medwise PMC clinical study 2025

Standards and best-practice references

- Rango AI, llms.txt complete guide
- IOActive, Prompt injection in GenAI
- Astiva 2026, RAG pipeline and AI citation optimization

This is a living document. Pharma GEO best practice is changing on a quarterly cadence; the underlying research and primary references should be re-audited every 90 days.