aikka

The Omnichannel Pharma guide

State of GEO & Omnichannel

Research report

2026 edition

FOREWORD

Why this guide exists

Over the past twelve months, pharma has crossed a line. Google rankings no longer decide who gets heard when a drug is discussed — the synthesis layer of Perplexity, ChatGPT, Claude, Gemini and AI Overviews does. Patients ask these engines their questions. So do physicians. According to the American Medical Association, 66% of US physicians now use AI in practice — a 78% jump in a single year. Pew Research reports that 22% of Americans now turn to AI chatbots for health questions — and 32% of 18-29 year-olds.

This guide documents — brand by brand, campaign by campaign, persona by persona — how the global pharma industry produces omnichannel content, and how visible (or invisible) that content actually is inside generative engines. It is the editorial counterpart to Pharma Guide v1, which mapped brands, indications and pipelines. Here we look at the other side of the equation: marketing, the brand website, the TV spot, the patient community, the LinkedIn KOL ad, the organic TikTok.

The angle is consistent: **every asset is read through the GEO / AI-visibility lens**. Which brand is citable by an LLM? Which FAQ is extractable? Which patient story is invisible because it lives only as video? Which HCP hub is a walled garden? The scoring grid is explicit (Part 2). Anti-patterns are documented. Best practices are named.

How to read this guide

You can read it end to end — 9 parts, 100+ pages, dense. Or dip in by brand (Part 3), by campaign (Part 4), or by persona (Part 7). The glossary and brand index make navigation fast.

Callout legend

IN PLAIN ENGLISH

Plain-language take on a technical concept, written for non-specialist readers. Use these in brand-team briefings or Steering Committee decks.

GEO FACT SHEET

A formal GEO score with a five-dimension analysis grid. Compact format for fast audits of a site, a hub, or a single asset.

THE CONSULTANT'S TAKE

Strategic perspective — what an Aikka consultant working at this company would actually say. Reading between the lines on the GEO competitive landscape.

WATCH OUT

Anti-pattern, regulatory risk, or GEO false friend. Do not reproduce with your own clients.

BEST PRACTICE

Example identified in the market — to reproduce, adapt, or cite in a meeting. A sourced, factual reference.

Sources and methodology

The 40 brands were the subject of direct patient, HCP and corporate site audits (June 2025 – May 2026). Campaigns are drawn from official Cannes Lions Health, PM360 Trailblazer, MM+M, Fierce Pharma Marketing Awards, Clio Health, and Shorty Awards rankings. Personas are inferred from patient surveys, FDA PFDD reports, Nature and JCO papers, and patient associations (NEA, CFF, NORD, EURORDIS). Names are fictitious; behavioural profiles are built by reasoned inference from real sources.

Every link in this document is clickable. Every image is a real screenshot or visual collected directly from the publisher (iSpot.tv, Fierce Pharma, YouTube, Wikimedia, official sites). No AI-generated imagery is used.

CONTENTS

Navigating the guide

PART 1

The 2026 landscape

> *"The generative engine is not an evolution of SEO. It is a paradigm shift: you no longer compete to be clicked, you compete to be cited."*
>
> *— Wil Reynolds, Seer Interactive*

The tipping point is behind us

The pharmaceutical industry is facing a structural shift. Patients and physicians now consume medical information through generative engines — ChatGPT, Perplexity, Gemini, Claude — that synthesise and return answers without a click. Brands that fail to appear inside those synthesised answers become invisible to their most engaged audiences.

Key figures 2025–2026

66%	of US physicians use AI in clinical practice *AMA, 2025*
+78%	growth in physician AI usage in a single year *AMA, 2025*
51%	of HCPs now use generative AI *Digitalya / Impiricus, 2026*
91%	of physicians demand to know who authored AI medical content *Wolters Kluwer, 2024*
22%	of Americans use AI chatbots for their health *Pew Research, 2025*
32%	of 18–29 year-olds use AI chatbots for their health *Pew Research, 2025*
+40%	LLM visibility uplift achievable with citations + stats + direct language *Princeton GEO, arXiv 2311.09735*
>80%	of AI traffic goes to pages updated within the last 2 years *Seer Interactive, 2025*

Mapping the omnichannel pharma landscape

Pharma marketing in 2026 spreads across seven principal channels, each with vastly different GEO visibility. The matrix below maps — for each channel — its budgetary weight in the industry and its citation potential in LLM answers. The conclusion is striking: **the channels where the industry invests the most (DTC TV, congresses, walled HCP gardens) deliver the lowest LLM visibility**.

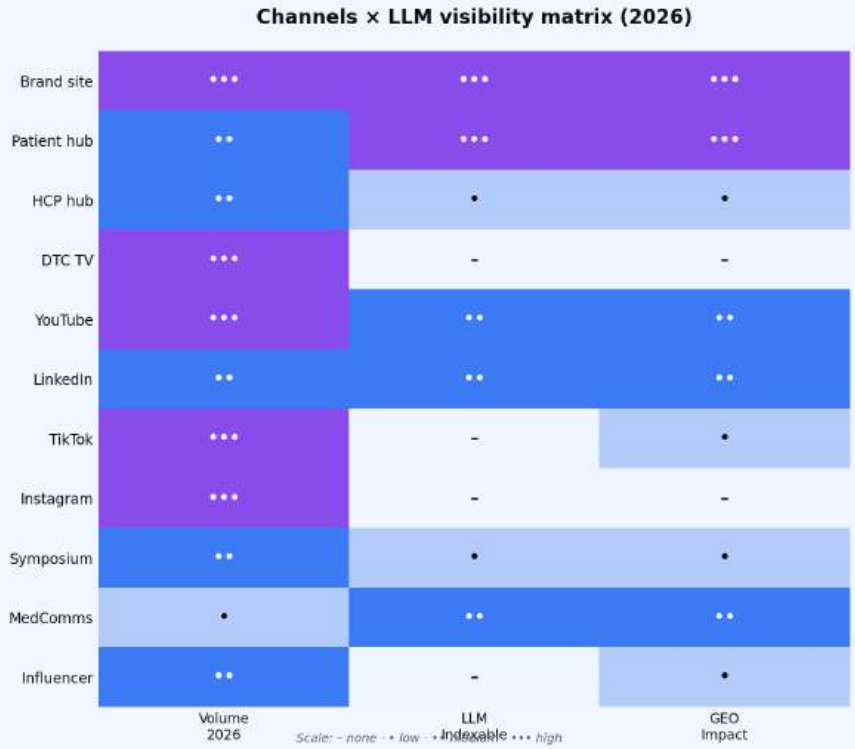

Pharma channels × GEO/AI visibility matrix (Aikka 2026)

THE CONSULTANT'S TAKE

The central pharma paradox of 2026: budgets remain concentrated on the channels that LLMs do not read. US DTC TV is an 8 billion dollar/year line item, yet transcripts are rarely published. Medical congresses (ASCO, ESMO, AHA, ADA, EASD) absorb millions of marketing dollars but posters and oral presentations only exist for LLMs through third-party recaps (FiercePharma, Endpoints). The HCP walled garden — Janssen Medical Cloud, MerckConnect, Lilly Pro — is mechanically invisible. The strategic window of 18–24 months consists of shifting a share of those budgets toward public, structured, citation-ready content.

BEST PRACTICE

Sanofi 'Chasing the Miracles of Science' (2025, agency VML) is an excellent example of a corporate campaign that maximised LLM citability: media coverage (FiercePharma, MediaPost, Campaign US) generates third-party backlinks and structured mentions. Sanofi.com hosts in parallel the Nadine patient page with extractable narration. The result: Sanofi is now cited by LLMs as a 'patient-centric corporate pharma' on transversal queries.

The three structural tensions of 2026

1. Compliance vs. visibility

FDA OPDP sent more than 100 warning letters to pharma companies in September 2025 (McGuireWoods, 2025). EMA is developing the ePI (Electronic Product Information) standard on FHIR — precisely the format LLMs can ingest. Pharma therefore lives this tension: produce rigorous content *and* make it citable.

2. Patient vs. HCP — the dual-audience dichotomy

Pharma brands serve two audiences with different languages, channels and legal obligations. GEO content structure must reflect that dichotomy: a plain-language patient hub plus a scientific HCP hub. Brands that blur the two (Tagrisso, Verzenio, Cosentyx) produce content that is less citable than those that keep them cleanly separated (Dupixent, Keytruda, Lynparza).

3. Local vs. global — regulatory fragmentation

The same medicine carries different indications, packaging and labels under FDA, EMA, HAS, NICE and PMDA. Aggressive geo-blocking (a blank page in France for a US product) is the cardinal GEO anti-pattern. The best players (Dupixent FR, Sanofi global) handle this via hreflang and local adaptation, not via blocking.

What follows — PART 2 formalises the PharmaGEO GEO-Readiness framework across 6 dimensions, with its 22 best practices and 22 anti-patterns. This grid is the reference tool to analyse any brand, any hub, any piece of content.

PART 2

The GEO-Readiness framework

GEO vs. SEO — the silent revolution

Generative Engine Optimization (GEO) denotes the set of practices used to optimise a brand's visibility, accuracy and citability inside answers generated by LLMs. Formalised in the Princeton paper, arXiv 2311.09735, GEO demonstrates that a source's visibility can be lifted by up to +40% through citations, statistics and structured language. AEO, AIO and LLMO are alternative terminologies — they all converge on the same discipline.

	Classic SEO	GEO
Goal	Appear at the top of the SERP	Be cited inside the AI answer
Metric	Ranking, CTR, traffic	AI Signal Rate, Share of Answer
Click	Indispensable	Optional — the AI represents you without a click
Content	Keywords, density, backlinks	Extractability, clarity, FAQ
Stability	Relatively stable	Volatile: 40–60% of sources change every month
Wikipedia	Useful	Major LLM credibility signal
Freshness	Important	Critical: >80% of AI traffic goes to pages <2 years old

IN PLAIN ENGLISH

The golden rule of GEO: if your content cannot be *extracted and summarised in 2–3 sentences by an LLM*, it will not exist in generative answers. That's it. Everything else follows from that constraint.

The 6 dimensions of pharma GEO-Readiness

The Aikka / PharmaGEO framework decomposes the GEO maturity of a brand into six dimensions, weighted by the relative importance of each lever inside the LLM answers observed on pharma queries.

D1	**Narrative clarity & extractability** Descriptive H1/H2/H3 hierarchy, short paragraphs, direct answers, defined jargon, quote-ready sentences.	20%
D2	**FAQ & Q/A structure** FAQ block 10 questions, schema.org FAQPage, real patient/HCP questions, long-tail queries.	15%
D3	**HCP depth & scientific accessibility** Partially public HCP hub, PI in HTML, MOA in text + annotated diagram, pivotal trials linked to PubMed.	20%
D4	**Authority & freshness signals** Named medical reviewer, publication + revision dates, PubMed/DOI citations, schema.org MedicalWebPage.	20%
D5	**Multilingual & multi-geo** Localised versions (FR, DE, UK, JP), human translation, correct hreflang, no aggressive geo-blocking.	10%
D6	**External citability & authority** Maintained Wikipedia, patient-association partnerships, STAT/Endpoints coverage, CME platforms.	15%

Score interpretation

0–30	Invisible — non-extractable content, critical score
31–50	Weak — marginal GEO visibility
51–70	Average — partial presence, quick wins identifiable
71–85	Good — structured GEO strategy, maintenance required
86–100	Excellent — brand cited regularly, leadership position

The pharma GEO-readiness pyramid

Pharma GEO-readiness pyramid (Aikka 2026) — each level prerequisites the next

The GEO-readiness pyramid translates the 6-dimension grid into an investment hierarchy. The foundation (robots.txt, sitemap, SSR) is *technical* and inexpensive. The middle layers (FAQ, MOA in text, reviewers, transcripts) require an editorial structuring effort. The upper layers (Wikipedia, partnerships, open access, disease awareness) are strategic and commit the organisation over a 12–18 month horizon.

22 Best Practices observed in pharma

These 22 best practices come from the field audit of 40 brands. They are ordered from the technical foundation (BP-15 sitemap, BP-16 robots.txt, BP-17 SSR) to the long-term strategy (BP-18 Wikipedia, BP-20 disease awareness, BP-21 CME).

BP-01	**Patient hub + HCP hub kept separate BUT accessible** Two distinct hubs with no mandatory login — the Dupixent model.
BP-02	**Structured FAQ with long-tail queries** 15 Q/A + schema.org FAQPage — the native LLM format.
BP-03	**MOA in text + annotated diagram with 200+-word alt text** Direct answers to 'how does X work?' queries.
BP-04	**Transcribed patient stories (not video-only)** A video without a transcript is invisible to LLMs — the Hemlibra model.
BP-05	**Patient communities referenced and reciprocally linked** NEA, NPF, CFF, NORD as authority backlinks.
BP-06	**Inline medical glossary** Citable definitions for terminology queries.
BP-07	**Plain Language Summary of pivotal trials** Reading grade 8 — the AstraZeneca Open Innovation model.
BP-08	**Named authors and medical reviewers with credentials** Strong E-E-A-T signal + schema.org Person.
BP-09	**Distinct indication-specific pages** Dupixent: separate /eczema, /asthma, /nasal-polyps pages.
BP-10	**Decision tools / calculators in accessible HTML** Crawlable, indexable, citable.
BP-11	**Open access to scientific publications** PubMed Central indexable — the AstraZeneca model.
BP-12	**Patient pathway + treatment-centre locator** Direct answers to 'how do I find a specialist?'.
BP-13	**ISI/SmPC clear, non-obstructive, well-structured** 30% of the page — organised by H3 headings.
BP-14	**Human-quality multilingual + local adaptation** Mention HAS / Ameli / NICE depending on market.
BP-15	**Public, complete, up-to-date XML sitemap** Submitted to GSC, Bing, declared in robots.txt.
BP-16	**Allow AI crawlers in robots.txt** GPTBot, ClaudeBot, Google-Extended, PerplexityBot.
BP-17	**Server-side rendered site (SSR)** Content readable without JavaScript — partial invisibility otherwise.
BP-18	**Active Wikipedia stewardship** The LLM tiebreaker — neutrality respected.
BP-19	**Proactive media coverage (STAT, Endpoints, Reuters Health)** High-authority domains crawled regularly.
BP-20	**Unbranded disease-awareness sites** Model: 'Could It Be ATTR-CM?' — Pfizer.
BP-21	**Integration with accredited CME platforms** Medscape, ESCMedEd — HCP authority.

BP-22	**Breadcrumb navigation + structured internal linking** AI crawlers traverse the site like a decision tree.

22 Anti-patterns to avoid at all costs

Anti-patterns are the practices that mechanically cancel GEO visibility. Some (walled garden, geo-block, hCaptcha) are technical decisions that can be reversed quickly. Others (untranslated jargon, video-only, marketing FAQ) are cultural and require an editorial shift.

AP-01	**Total HCP walled garden** All HCP content behind login — complete invisibility.
AP-02	**PDF prison** PI, SmPC available only as PDF, not OCR'd.
AP-03	**Video-only storytelling** Patient stories, MOA, KOL without transcripts.
AP-04	**Aggressive geo-block** 'Not available in your region' = double penalty.
AP-05	**Cosmetic marketing FAQ** 5 promotional questions — zero signal.
AP-06	**Unexplained patient jargon** 'JAK-STAT pathway' without definition.
AP-07	**Mega-PI dump** PI copied with no H2/H3 structure.
AP-08	**ISI Box colonising 50% of the page** The LLM extracts mostly the ISI.
AP-09	**No named medical reviewer** Negative E-E-A-T signal.
AP-10	**Geo redirects that penalise the crawl** The crawler gets an empty page.
AP-11	**Poor or absent Wikipedia** No credibility anchor for the LLM.
AP-12	**Missing schema.org markup** Content treated as generic.
AP-13	**MOA locked inside Flash/legacy animation** Completely invisible to LLMs.
AP-14	**Patient pages written as a medical manual** Reading grade 16 — LLMs do not cite.
AP-15	**HCP pages written as patient marketing** Lacks rigour — not cited by HCPs.
AP-16	**Pure SPA JS site without SSR** Content invisible to AI crawlers.
AP-17	**Missing or incomplete sitemap** FAQ/Studies pages not discovered.
AP-18	**Block GPTBot/Claude in robots.txt** Voluntary invisibility.
AP-19	**Duplicate content across indications** LLM duplication penalty.
AP-20	**Links to paywall only** The LLM cannot verify the claims.
AP-21	**Missing Drug structured data** No formal identification.

AP-22	**Exclusive reliance on WebMD/Vidal** Brand becomes a secondary mention, loss of narrative control.

The pharma paradox: compliance vs. GEO visibility

Pharma is probably the only sector where content production is governed by MLR (Medical-Legal-Regulatory) cycles of 4 to 8 weeks. That constraint is culturally and structurally opposed to the *freshness* LLMs demand (>80% of AI traffic goes to pages under 2 years old, per Seer Interactive). The paradox resolves through three strategies already observed among leaders.

Strategy 1 — Unbranded disease awareness

Launch a site centred on the disease rather than the product. Freed from the FDA OPDP / ANSM constraints on Rx promotion, that site can maintain a rich FAQ, publish KOLs, and partner with patient associations. Model: Pfizer 'Could It Be ATTR-CM?' (PfizerForAll/ATTR-CM).

Strategy 2 — Hybrid public/gated architecture

A first public layer (factual, indexable information: disease, mechanism, general criteria) plus a second gated layer for promotional HCP content. This preserves OPDP compliance while remaining crawlable. Model: dupixenthcp.com (Sanofi/Regeneron).

Strategy 3 — Patient-association partnerships

Associations (NEA, CFF, NORD, EURORDIS, France Assos Santé) publish far more freely than pharma brands. A transparent editorial partnership with clear disclosure is both compliant and GEO-powerful. Model: Vertex + the CF Foundation for Trikafta.

WATCH OUT

FDA OPDP — heightened vigilance 2025–2026
In September 2025 the FDA sent more than 100 warning letters to pharma companies for misleading advertising, including social activities and influencer partnerships. Off-label promotion, insufficient *fair balance*, and incomplete disclosures are the leading motives. Every piece of GEO content must pass the OPDP filter — but you can be GEO-friendly AND compliant.

PART 3

Audit of the 40 brands

This part is the heart of the report. Each brand has been audited directly — patient sites, HCP sites and corporate sites — between June 2025 and May 2026. The GEO-Readiness score (/10) results from the grid broken down into 5 sub-scores (narrative clarity, FAQ, HCP depth, authority/freshness, multilingual) — a simplified version of the 6-dimension framework introduced in PART 2. Each fact sheet presents a screenshot of the site, a consultant analysis and actionable recommendations.

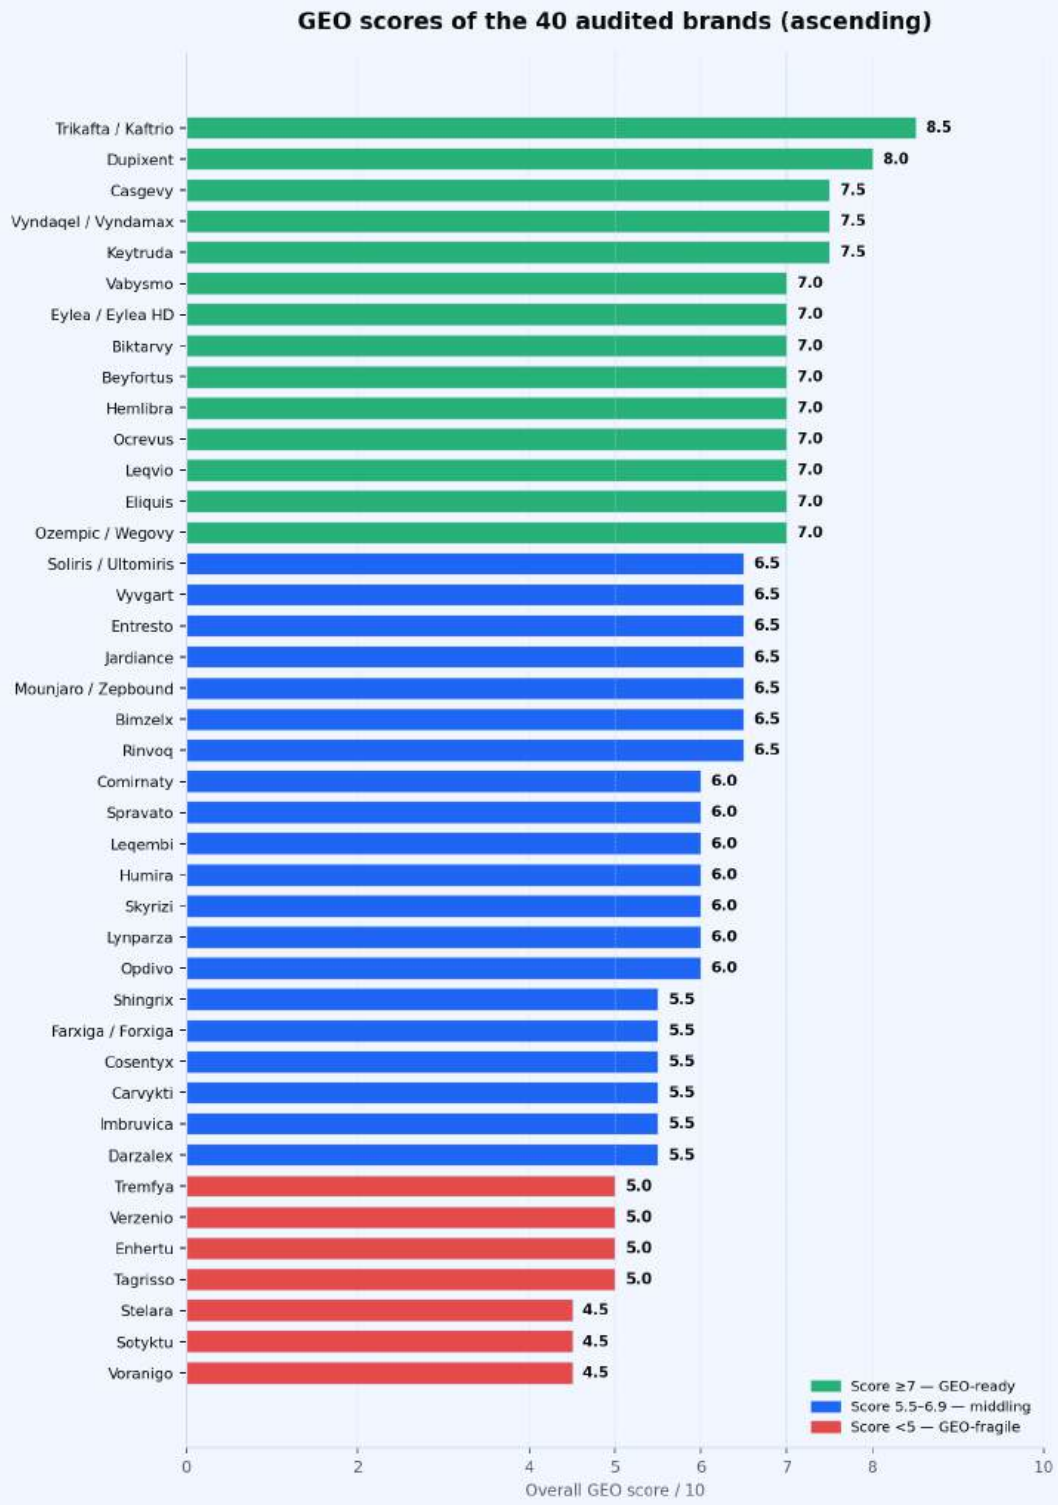

GEO-Readiness scores of the 40 pharma brands 2026 (Aikka / PharmaGEO audit)

3.1 — ONCOLOGY

3.1 — Oncology

This section covers 10 major brands in oncology. Average score: 5.5/10.

Keytruda — Merck / MSD

Immuno-oncology — 44 FDA indications across 19 cancers · keytruda.com / keytrudahcp.com

GEO SCORE

7.5 /10

Narrative clarity		1.5/2
Q/A FAQ		0.5/2
HCP depth		2.0/2
Authority / Freshness		1.5/2
Multilingual / Geo		2.0/2

Source: https://www.keytruda.com

GEO STRENGTHS

- Unmatched indication breadth (44 approvals) — a natural source for any LLM query on PD-1 immunotherapy
- Mechanism of action expressed in directly extractable text (PD-1 / PD-L1 pathway blockade)
- Deep HCP content with no login wall and 30+ named KEYNOTE trials

GEO WEAKNESSES

- No structured FAQ with schema.org markup — a major extractability gap
- No named medical reviewers and no visible update dates — weak E-E-A-T signal
- Video-only MOA content with no text transcript

THE CONSULTANT'S TAKE

Keytruda is the benchmark for publicly accessible scientific depth. Its 44 indications and 30+ KEYNOTE trials make it a natural citation source for LLMs across immuno-oncology. The main gap — and the ceiling on its GEO score — remains the missing structured FAQ and weak E-E-A-T signals.

BEST PRACTICE

Actionable recommendation. Quick win: add a structured FAQ block per indication with FAQPage schema. Mid-term: name the medical reviewers and systematically date every page.

Opdivo — Bristol Myers Squibb

Immuno-oncology (anti-PD-1) — multi-indication · opdivo.com / opdivohcp.com

GEO SCORE

6.0 /10

Narrative clarity		1.0/2
Q/A FAQ		0.5/2
HCP depth		1.5/2
Authority / Freshness		1.0/2
Multilingual / Geo		2.0/2

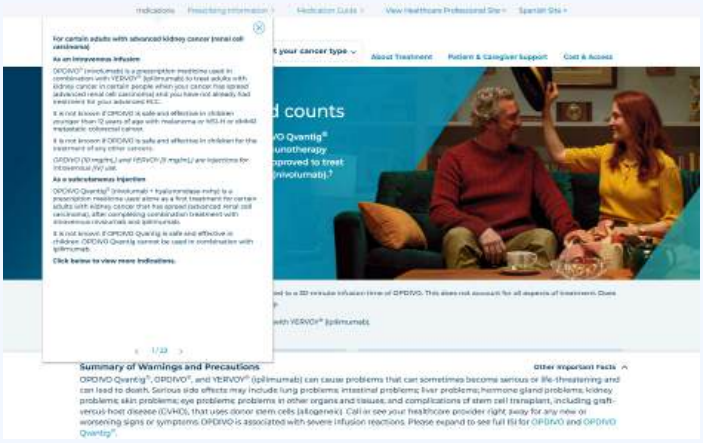

Source: https://www.opdivo.com

GEO STRENGTHS

- Broad oncology coverage (NSCLC, cHL, HCC 1L, melanoma…)
- OPDIVO Qvantig SC clearly positioned as the IV alternative
- Latest data surfaced on the HCP hub

GEO WEAKNESSES

- Patient site dominated by ISI at the expense of extractable narrative
- No detectable structured FAQ
- MOA barely present as direct text on the HCP side

THE CONSULTANT'S TAKE

Opdivo's content is markedly less extractable than direct competitor Keytruda. The patient site produces very little narrative beyond ISI — a critical GEO problem. The HCP site is news-first but lacks structured scientific depth in text.

BEST PRACTICE

Actionable recommendation. Add a standalone MOA text page, restructure the patient hub around indications, and create long-tail FAQs. ISI should be capped at 25-30% of each page.

Darzalex — Johnson & Johnson

Multiple myeloma (anti-CD38) · darzalex.com / darzalexhcp.com

GEO SCORE

5.5 /10

Narrative clarity		1.0/2
Q/A FAQ		0.5/2
HCP depth		1.5/2
Authority / Freshness		1.0/2
Multilingual / Geo		1.5/2

Source: https://www.darzalex.com

GEO STRENGTHS

- Both IV and SC (FASPRO) formulations documented — relevant for patient queries on administration time
- Patient support programs in place (Janssen withMe / Darzalex withMe)
- Combination data (Rd, Vd, KRd) referenced on the HCP side

GEO WEAKNESSES

- No clear line-of-therapy hub — queries like 'Darzalex first-line vs relapsed' are poorly served
- No substantive patient FAQ block
- MAIA, CASTOR, POLLUX, GRIFFIN trials under-leveraged in plain language

THE CONSULTANT'S TAKE

Darzalex shows a pattern typical of heme-onc products: strong technical HCP coverage but a patient hub under-resourced in extractable content. Myeloma patients ask very specific questions (lines of therapy, combinations, duration) that the site does not answer well in GEO format.

BEST PRACTICE

Actionable recommendation. Build 'Darzalex first-line / Darzalex in relapse' pages with dedicated FAQs, and publish Plain Language Summaries of the MAIA / GRIFFIN / CASTOR pivotal trials.

PREVIEW · FIRST 20 PAGES

Want the full guide?

Read all 120 pages on

pharma-geo.com

VISIT PHARMA-GEO.COM

Free download. PDF. 120 pages. 2026 edition.