

Who's asking? Persona framing changes the evidence AI assistants cite on sickle cell disease

Ellie Thomas,¹ Joseph Baker,¹ Ellie Hughes¹ | 1. Camino Communications, London, UK

Implications for medical affairs and patient advocacy teams sharing SCD information online

Keep general health websites accurate and current

Lay patient answers drew **52.3%** (CI 39.9–65.9) of citations from general health sites and only **2.6%** (CI 0.0–7.2) from journals, indicating a substantial influence on what information LLMs return to patient queries.

Test AI with the questions each audience might ask

Claiming to be a specialist did not change the journal share of citations (**9.0% vs 8.6%**), but detailed questions raised it to **25.6%**, indicating that question complexity matters more than stated role.

State clearly online when a therapy is withdrawn

Most answers flagged voxelotor and crizanlizumab as withdrawn, but **8** did not, suggesting that outdated information can persist within LLMs when outdated sources are surfaced.

Introduction

- Patients, carers and clinicians increasingly ask AI assistants about rare diseases, where authoritative sources are scarce
- Retrieval-enabled LLMs run their own web searches based on what the user types
- Objective: to assess whether persona framing changes the sources and answers an LLM returns to the same clinical questions on SCD

Methods

1

Prompt generation

10 SCD topics × 4 personas (lay patient, informed patient, GP, specialist haematologist) × lay and clinical wording, plus 10 control prompts: GP questions asked as the specialist

2

Prompt runs

270 calls to one pinned model (gpt-5.5-2026-04-23) with web search, UK location and no system prompt, within one 4-hour window

3

Source classification and answer coding

1,594 citations sorted into 6 prespecified source classes using a fixed domain list; answers were coded on 7 dimensions, including withdrawn therapies

4

Statistical analysis

Source-class shares calculated by persona with cluster-bootstrap 95% CIs; a pattern counted as a finding only if it held in all 3 replicates

Results

Answers varied even for identical prompts

5

median sources cited per answer
Range 1–17. All 270 calls returned a web-searched answer.

37.5% of cited domains shared when the same prompt was repeated
IQR 25–50%. At the level of individual pages, overlap fell to 17% (IQR 6–31%).

Cited sources differed by persona

Figure 1. Share of citations by source class

Persona	General health sites	Guidelines	Journals	Other
Lay patient	52%	34%	3%	11%
Informed patient	28%	45%	7%	20%
GP	16%	58%	8%	16%
Specialist	4%	60%	12%	25%

Lay patient

52.3% (CI 39.9–65.9) of citations came from general health sites and 2.6% (CI 0.0–7.2) from journals

Specialist haematologist

3.6% (CI 1.9–6.1) of citations came from general health sites and 24.9% (CI 10.7–40.4) from journals

GP and specialist

58.2% (CI 46.6–66.0) and 59.9% (CI 46.9–72.1) of citations came from treatment guidelines, respectively

Question depth, not persona, drove journal citations

Figure 2. Share of citations from peer-reviewed journals for clinically worded questions

Question Type	Share of citations from peer-reviewed journals
GP persona, GP questions	8.6%
Specialist persona, GP questions	9.0%
Specialist persona, own questions	25.6%

Guideline share rose from 58.2% (CI 46.6–66.0) to 70.3% (CI 59.0–79.1) when the specialist persona asked the GP's questions

Withdrawn therapies were not always mentioned

28%
of answers mentioned voxelotor or crizanlizumab (76 of 270 responses)

89%
of mentions stated a drug had been withdrawn (68 of 76 responses)

3%
of answers (n=8 omitted mention of drug withdrawal (8 of 76 responses)

We're keen to apply this approach in other rare diseases and to keep testing in this fast moving area.

If you'd like to explore what AI assistants are telling people about the condition you work in, get in touch:

ellie.thomas@caminocomms.com Caminocomms.com