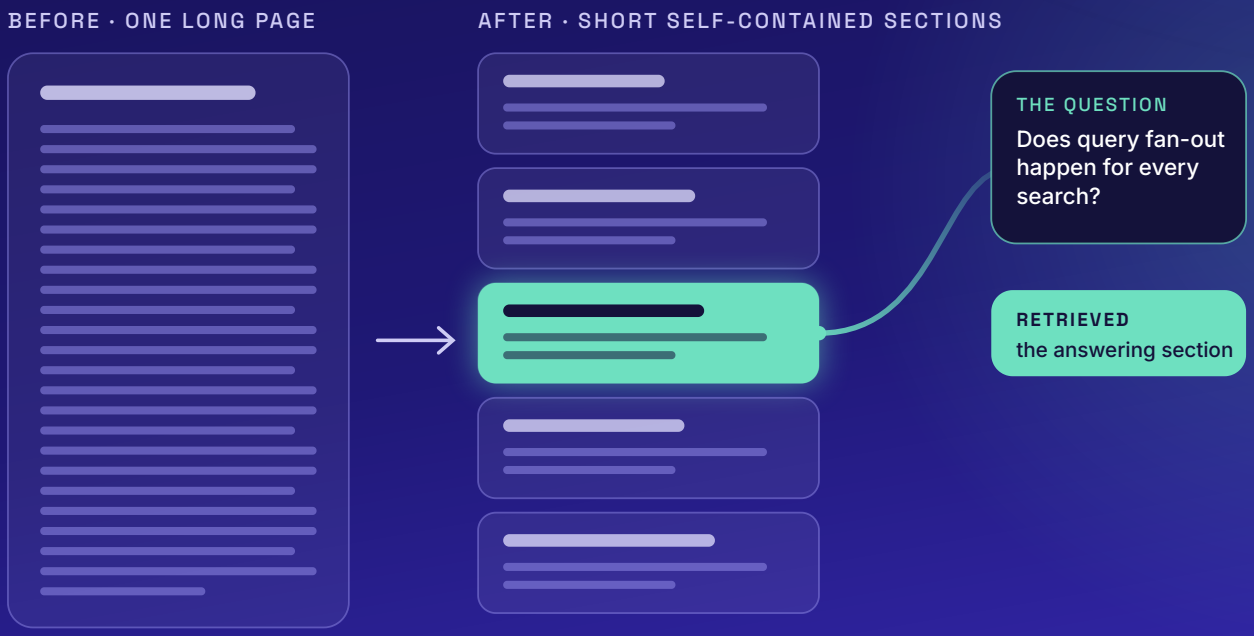
Built to Be Retrieved

AI search doesn't read your page. It lifts one passage. I rewrote ten pages five ways to find out which structure gets the right passage picked.



84–92%
of questions: short sections beat a plain rewrite

4 of 4
retrieval models agree

+20–29 pts
more often, the passage with the answer wins

Ameneh Saeednia
Senior SEO and Generative Search Strategist

10 pages · 50 rewrites · 49 questions
Run October 28, 2026

The short version

AI search retrieves passages, not pages. I rewrote ten real pages in controlled ways, kept every fact identical, and measured which rewrites made the right passage easier to find.

84–92%

Short sections win. Splitting pages into self-contained sections under question headings beat a plain paraphrase on most questions, in all four models.

+20 to +29

The right passage gets picked. Short sections raised how often the top passage held the actual answer, in points, against the control.

61–78%

Naming entities helps a little. Replacing "it" and "this" with real names gave a smaller but consistent lift.

51–61%

Answer first did almost nothing. A summary on top and answers moved up was close to a coin flip in every model.

+20 to +44

Conversational, human-first posts gained most. Blog posts written for people, with stories and examples, gained more than Wikipedia articles from being split into sections.

What it means: a strong page is not enough if its answers are buried in long, connected paragraphs. Each answer needs its own short, self-contained section that still makes sense when it is lifted out alone.

Why passages matter

Before an AI writes an answer, a retrieval step picks a handful of passages from many pages. Your page competes passage by passage, not as a whole.



Page-level thinking

Rank the whole page for a keyword. A strong page wins even if its best answer sits in paragraph nine and depends on the paragraphs above it.

Passage-level thinking

Each chunk of a few hundred words is scored on its own against the question. A chunk that needs context it does not contain loses, however good the page is.

Why test structure, not content

Every rewrite kept the facts identical. Every number in the originals survives in every variant. Only structure and wording changed, so any difference in retrieval comes from how the page is built.

Why a control

Any rewrite could change scores just by changing words. So each page also got a plain paraphrase with no structural improvements. A rewrite only counts if it beats that control, not just the original.

How the study worked

Each question was scored against its own page's passages and against the original passages of the nine other pages.

1

Ten pages

Five of my seogirl.ca posts and five Wikipedia articles on the same topics, so each page competes with close rivals.

2

Five rewrites

Answer first, short sections, explicit entities, all three combined, and a plain-paraphrase control.

3

Four models

Three open embedding models (MiniLM, bge-small, e5-small) and BM25 keyword search, on 49 natural questions.

Version	Rewrite rule	Length
Original	Unchanged page text	781 to 2,008 words
Answer first	A 2 to 3 sentence summary on top; the direct answer moved to the start of each section	+3 to +6%
Short sections	100 to 250 word sections under question headings, each understandable on its own	-2 to +11%
Explicit entities	Pronouns and vague references replaced with names; entities named in headings	+4 to +8%
All combined	All three rewrites together	+9 to +23%
Control	A plain paraphrase, same length and structure	-5 to +2%

Pages were split into passages at headings, capped at 200 words. Each question has an answer phrase found in exactly one passage of the original page, which checks whether the right passage was picked. Effects were tested with a paired Wilcoxon signed-rank test. Run on October 28, 2026.

Finding 1

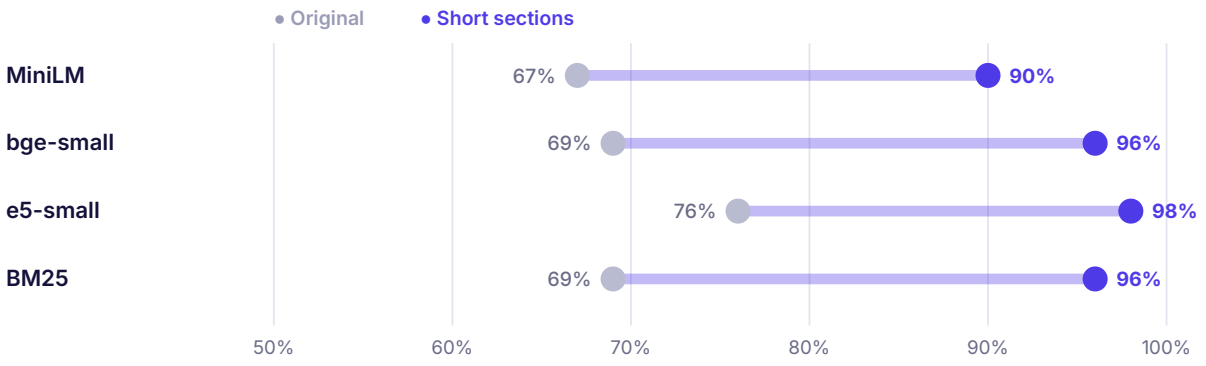
Short sections win in every model

Splitting pages into self-contained, question-headed sections was the only rewrite that clearly beat the control everywhere.

	MiniLM	bge-small	e5-small	BM25
Short sections	84%	90%	90%	92%
All combined	88%	90%	90%	94%
Explicit entities	61%	73%	65%	78%
Answer first	51%	57%	61%	55%

Share of 49 questions where the rewrite beat the plain-paraphrase control. 50% = no effect. Darker = stronger.

How often the top passage held the answer: original vs short sections



The control stayed flat

The plain paraphrase beat the originals on only 47 to 55% of questions. Rewording alone did not move retrieval, so the gain comes from structure.

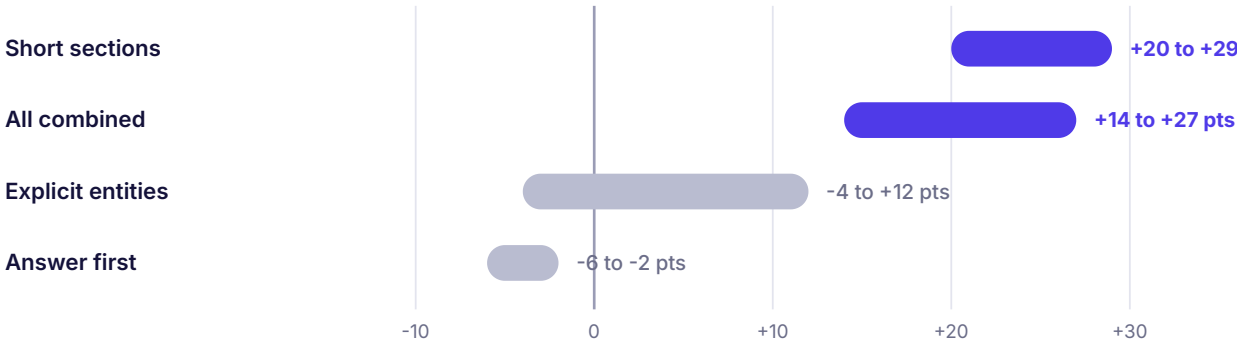
Stacking added nothing

Combining all three rewrites beat the control on 88 to 94% of questions, about the same as short sections alone.

Finding 2

Names help a little, answer-first barely moves

Two popular rewrite rules gave very different results once they were compared with the control.



Change in how often the top passage held the answer, vs the control, in points. Bar = range across the 4 models.

Entities: real but small

Scores rose in every model ($p < 0.02$), but the right passage was picked only -4 to +12 points more often. Names make a passage clearer to a model without changing which passage wins.

Answer first: no effect

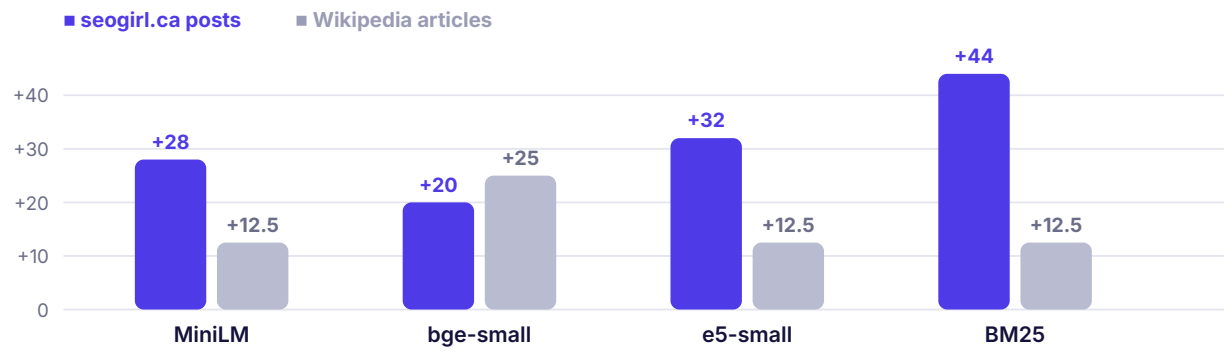
A summary on top and answers moved up made no significant difference (p 0.18 to 0.77). Moving an answer within a long section does not help if the section itself is too long.

Finding 3

Blog posts gain more than encyclopedia pages

The pattern held for both sources, but conversational, human-first posts gained more when they were split into sections.

Source	Questions	Short sections beat control	Right passage picked	Entities beat control
seogirl.ca posts	25	84 to 88%	+20 to +44 points	68 to 80%
Wikipedia articles	24	79 to 96%	+12 to +25 points	54 to 75%



Short sections: points gained in picking the passage with the answer, vs the control.

Why posts gained more

Blog posts spread answers across stories, examples and FAQ items. Splitting them put each answer into its own passage. Wikipedia prose is already dense with facts.

Read it as a hint

With about 25 questions per group, the difference between sources is a direction to test further, not a finding to quote.

What content teams can do

If retrieval happens passage by passage, structure becomes part of the content, not decoration.

Give each answer its own section

Keep sections around 100 to 250 words, one question each, under a heading that names it.

The evidence: Short sections beat the control on 84 to 92% of questions in every model.

Make sections stand alone

Cut "as mentioned above" and repeat the context a lifted passage needs.

The evidence: The right passage was picked 20 to 29 points more often.

Name things, not ”it”

Use the entity's real name at the start of sections and in headings.

The evidence: Explicit entities beat the control on 61 to 78% of questions.

Restructure before you rewrite

Change the shape of the page first; rewording alone did not help.

The evidence: The plain paraphrase moved results by about a coin flip.

Start with long, conversational pages

Blog posts and guides with buried answers have the most to gain.

The evidence: Conversational posts gained +20 to +44 points vs +12 to +25 for Wikipedia.

Test before you roll out

Run your own pages and questions through a passage test with a control.

The evidence: Answer-first, a common rule, showed no measurable effect.

Limits, and the series

Read this as directional

- **Heading bias.** I wrote the questions before the rewrites, and some question headings echo them closely. This likely inflates the short-section and combined results.
- **Near the ceiling.** The originals already ranked in the top 3 for 98 to 100% of questions, so ranking gains were capped.
- **Small set.** 49 questions, 10 pages, one writer for all rewrites.
- **Open models only.** This measures retrieval similarity in small open models, not how ChatGPT, Perplexity or Google AI Mode rank or cite pages.
- **Closed competition.** Each page competed with nine other pages, not the open web. The combined rewrite also ran 9 to 23% longer.

The complete series

Study 1, The AI Shelf: which brands AI recommends. **Study 2, Who Are You, Really?:** brand consistency and the knowledge graph. **Study 3, Does Size Buy Trust?:** E-E-A-T by brand size. **Study 4, Absorbed, Not Credited:** when AI credits its sources. **Study 5, this report:** how structure changes passage retrieval.

Sources

Five seogirl.ca posts by the author: query fan-out, RAG, entity SEO, how Google mines Wikipedia, and hybrid retrieval. Five Wikipedia articles under CC BY-SA 4.0 (Wikipedia contributors): Query expansion, Retrieval-augmented generation, Knowledge graph, Information retrieval and Wikidata, retrieved October 28, 2026. Full URLs, revision IDs and data are in the lab files.

Independent research

This is an independent analysis using open-source models and openly licensed content. I have not worked with, or been endorsed by, any company mentioned.

SEO Girl

Where AI meets search, I make brands visible.

Ameneh Saeednia is a Senior SEO and Generative Search Strategist with more than nine years of experience across e-commerce, international education and travel. She specializes in entity SEO, E-E-A-T and AI visibility.

seogirl.ca