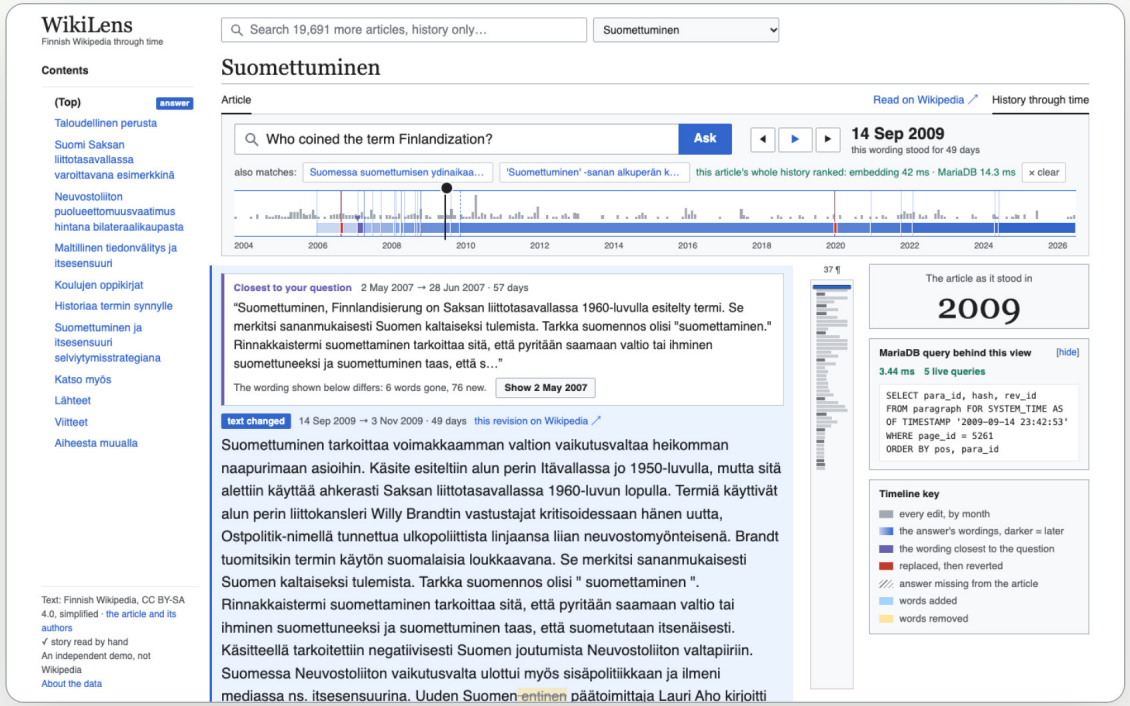
MARIADB CHALLENGE · JUNCTION 2026

WikiLens

Ask an article once.
Scrub the answer through time.

wikilens.pgf.lv



THE PROBLEM

A Wikipedia article is only the latest of thousands of versions.

To see how its answer to one question changed, you read diffs for an afternoon.

Talvisota, the Finnish article on the Winter War, has 2,143 revisions since 2003.

HOW IT FEELS

Ask once. Scrub the answer.

1 · Ask one question

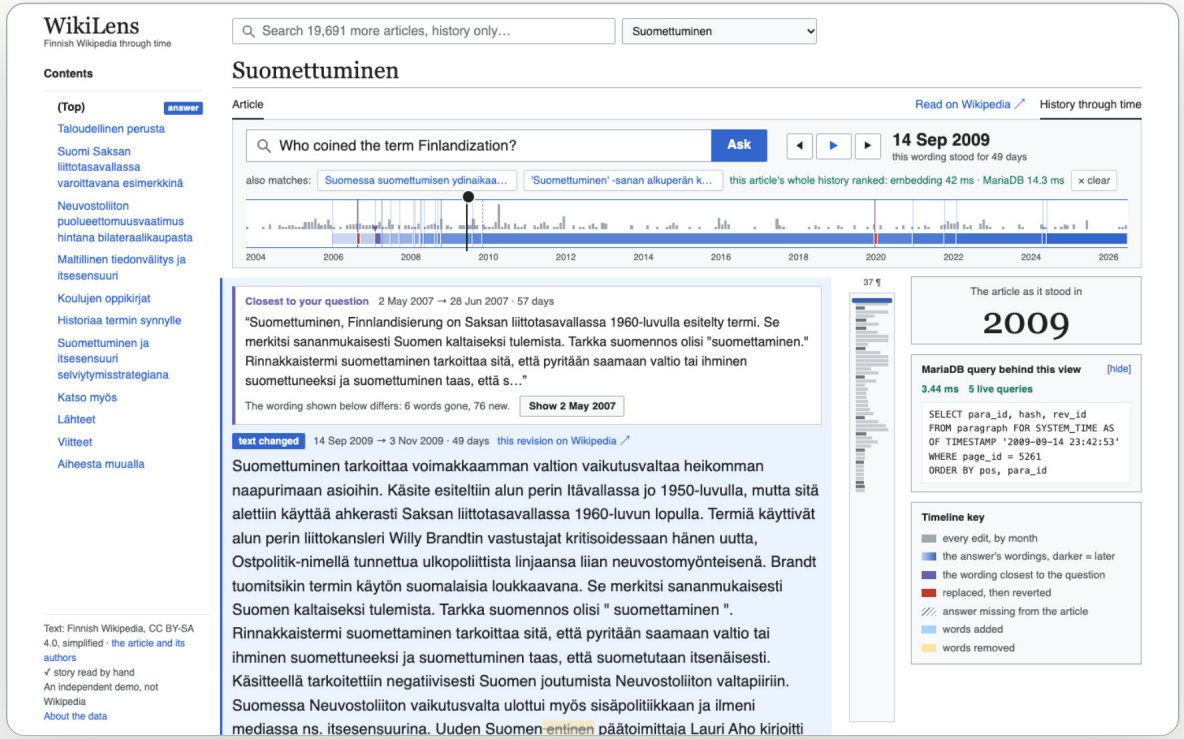
In Finnish or English. No dates to pick.

2 · The answer lights up

The paragraph that answers, right where it sits in the article.

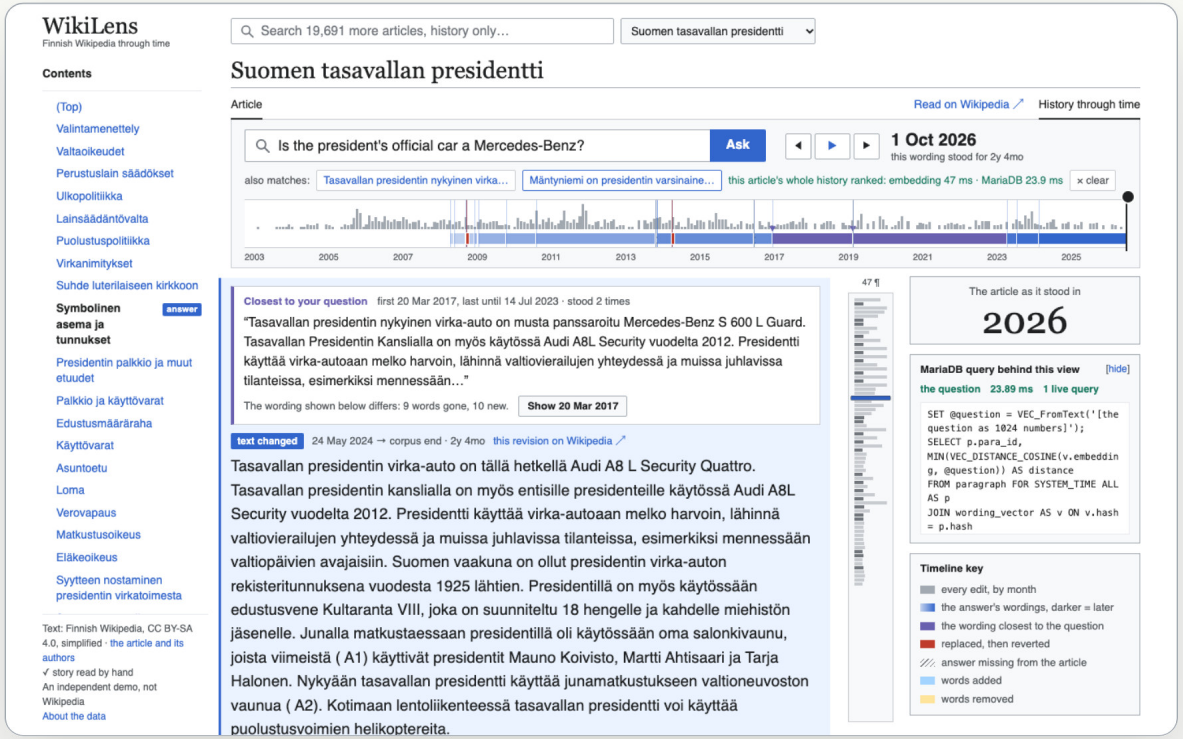
3 · Drag the timeline

It rewrites itself word by word, through twenty years of edits.



THEN AND NOW

It tells you when the answer was different.



"Is the president's official car a Mercedes-Benz?"

NOW · SINCE 24 MAY 2024

Audi A8 L Security Quattro

THEN · 2008 TO 2024

Mercedes-Benz S 600 L Guard

The question is ranked against every wording the paragraph ever had, so the page can say when the article read like that, and jump there.

HOW MARIADB IS USED

Time travel is the mechanism.

A paragraph is a row

in a system-versioned table.

An edit is a history row

stamped with the time it was saved on Wikipedia.

A wording has a vector

in a VECTOR column of the same database.

A SCRUB FRAME

```
SELECT para_id, hash, rev_id
FROM paragraph
  FOR SYSTEM_TIME AS OF TIMESTAMP '2013-06-01'
WHERE page_id = 1256 ORDER BY pos, para_id;
```

ASK: EVERY WORDING THE ARTICLE EVER HAD

```
SELECT p.para_id,
  MIN(VEC_DISTANCE_COSINE(v.embedding, @question)) AS distance
FROM paragraph FOR SYSTEM_TIME ALL AS p
JOIN wording_vector AS v ON v.hash = p.hash
WHERE p.page_id = 1256
GROUP BY p.para_id ORDER BY distance, p.para_id LIMIT 3;
```

MEASURED

The whole dataset, loaded and timed.

19,742

articles of Finnish Wikipedia, each checked against its source text

862,076

revisions, 2003 to October 2026, loaded in 8.5 minutes

1.2 M

paragraph versions, stored as history rows

0.6 ms

median for a frame, a FOR SYSTEM_TIME AS OF query

12 ms

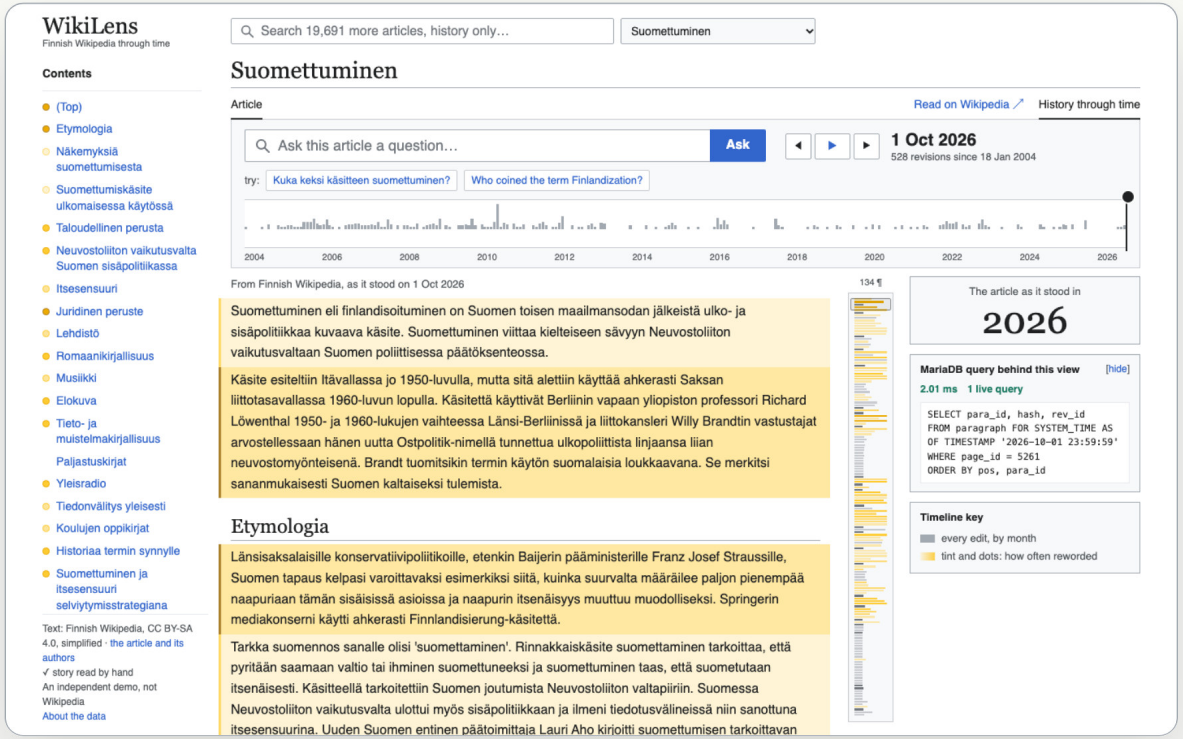
median to rank an article’s whole history against a question

12×

smaller: 11.6 GB of wikitext stored as 959 MB

DESIGN

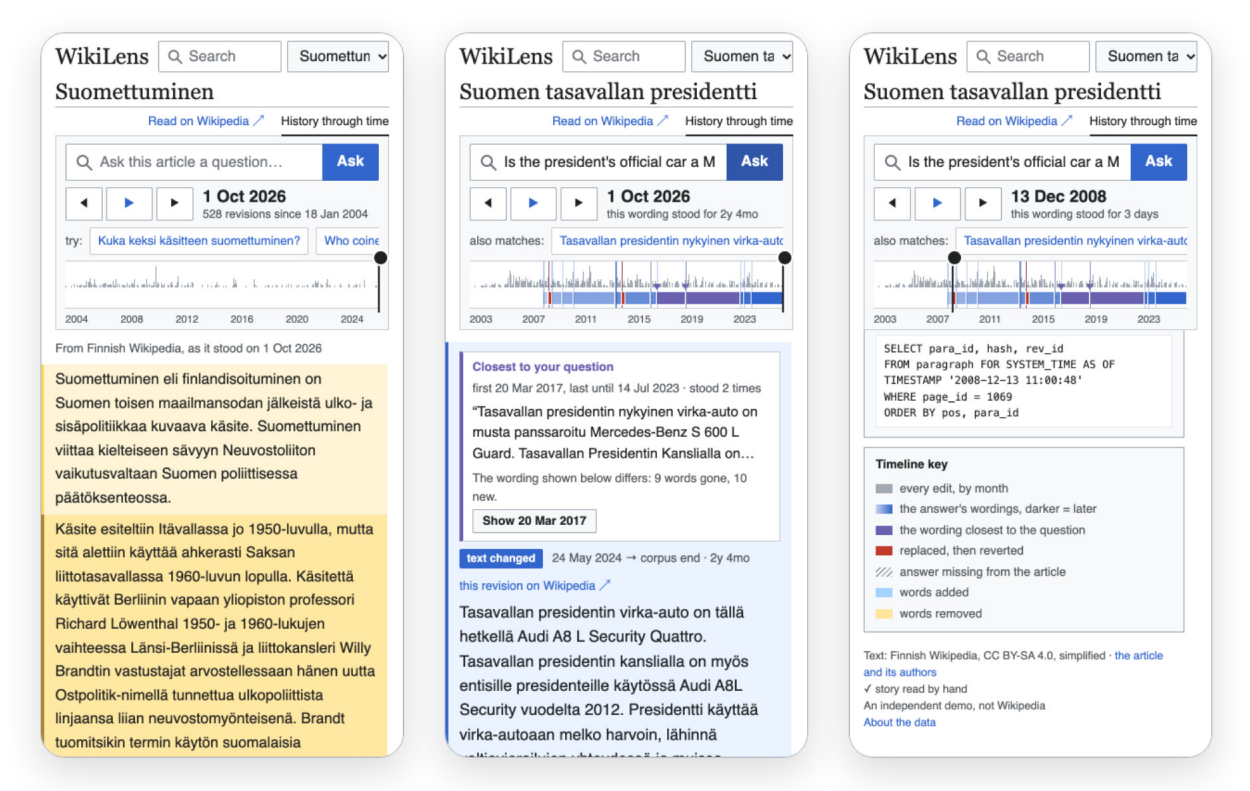
It reads like the article you know.



- Wikipedia’s own layout. The timeline is its history tab, opened.
- The contents list and the overview strip follow the playhead.
- A tint shows where the article was reworded most.
- Labels say what happened, never why.

MOBILE

On a phone, the timeline stays put.



At phone width the page becomes one column. The question and the timeline stay at the top while the article scrolls beneath them.

The side panels move into the page, and the same then-and-now note appears above the answer.

Covered by the browser tests at a 390 px touch viewport.

UNDER THE HOOD

How it is built.

DATASET

MariaDB’s Finnish Wikipedia set, size M, in Parquet.

→

INGEST

Clean the wikitext and match paragraphs between revisions.

→

MARIADB

System-versioned paragraphs and a VECTOR column.

→

PAGE

FastAPI and TypeScript. One live query per frame.

Multilingual embeddings

bge-m3 through Ollama, so English questions find Finnish text.

TRY IT

Wikipedia through time, one question at a time.

LIVE	wikilens.pgf.lv
FILM · 53 S	wikilens.pgf.lv/film
CODE	github.com/progresful/wikilens