PRODUCT BOOKLET · NOVEMBER 2026

PledgeScope

Promises made. Actions examined.

Every member of Congress. What they promised, what the record shows, and the source behind every number. A language model reads. Code verifies. People decide. No scores, rankings or ratings.

539
current members with a profile

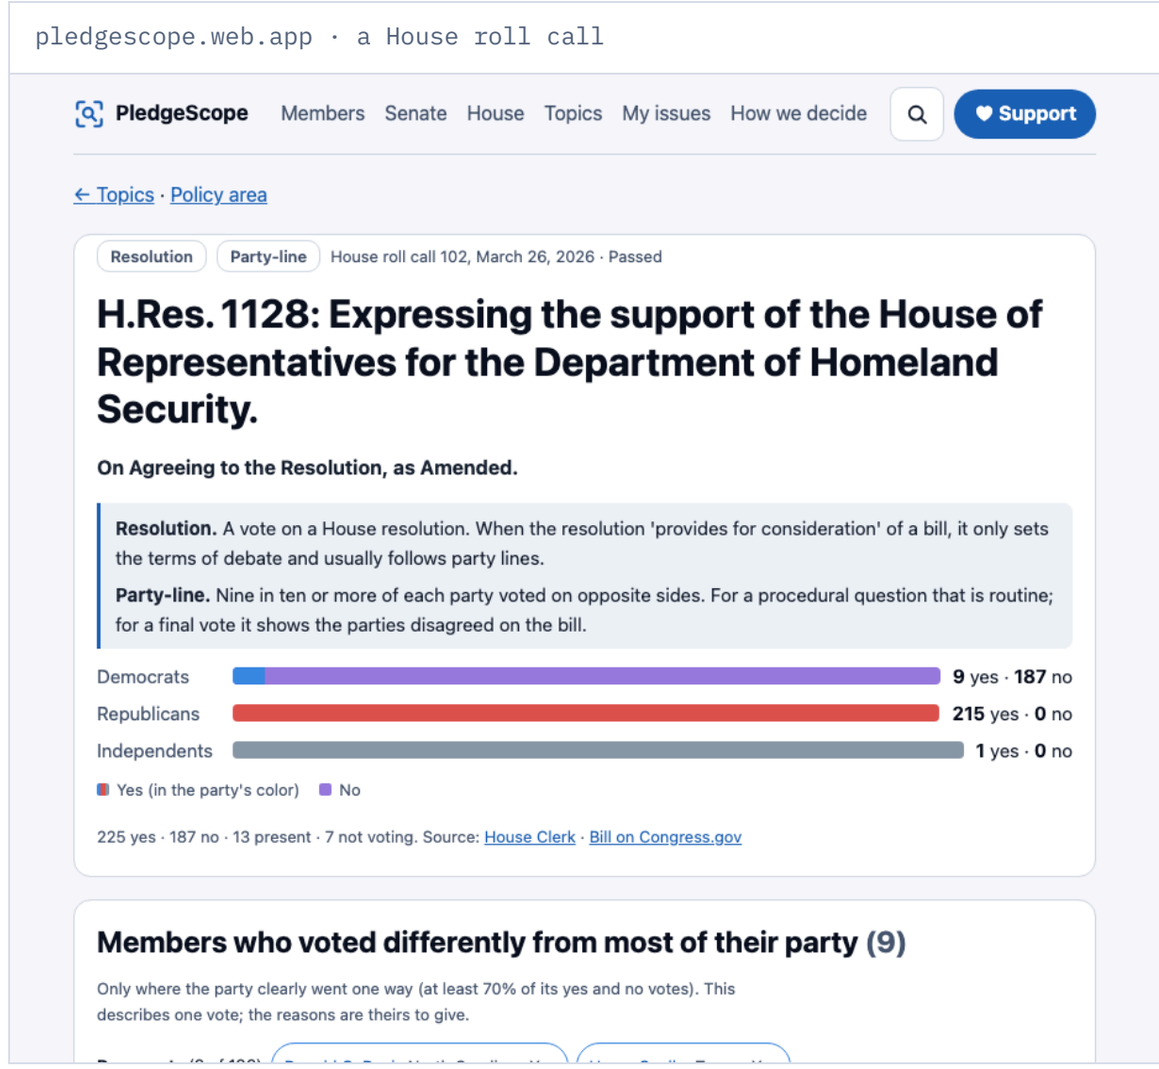
1,025
campaign sites preserved

Fabricio Zoff
Design and engineering · Zoff.Tech

The facts are public. Nobody lines them up.

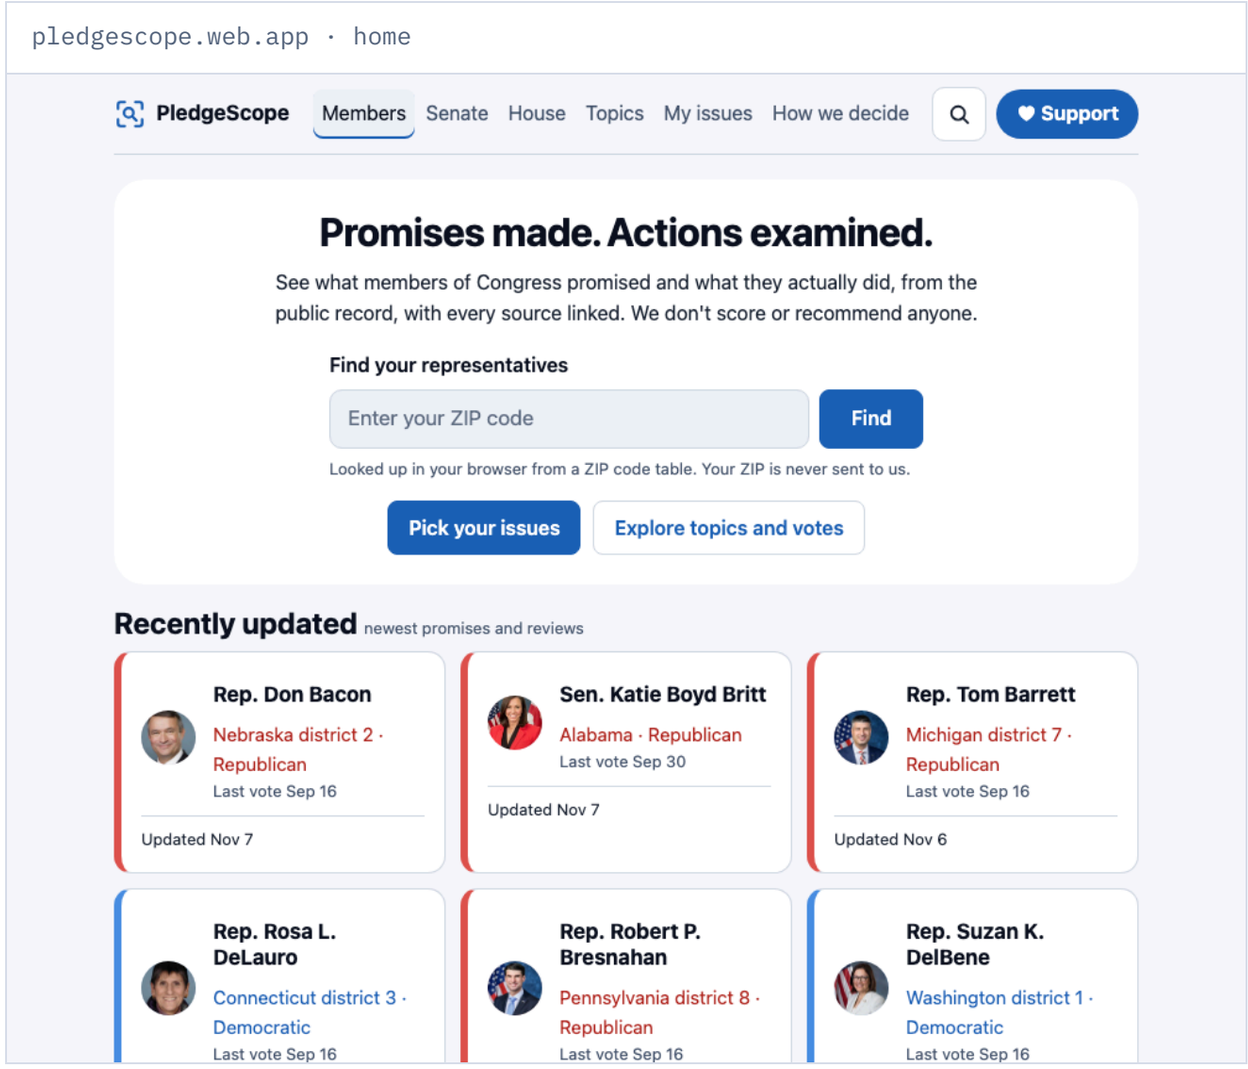
A campaign site, Congress.gov, the House Clerk, the FEC. Each holds part of a member's story in its own format. Few people put them side by side, and the tools that try usually end in a score. I wanted a page per member where the record speaks for itself and nothing gets rated.

I also wanted to use a language model without letting it decide anything. It reads pages and proposes. Code checks every quote against the preserved source. A person approves, assesses and publishes. The model never writes to the public site.



- ✓ **Describe, do not characterize.**
  Counts, quotes and dates. No "liar", no "hero".
- ✓ **Absence of evidence is not a verdict.**
  "No action found" says what was searched.
- ✓ **Same method for every member.**
  Both parties, both chambers, one rulebook, published.

Start with your ZIP code. End at a source.



- ✓ **Find your representatives**
  ZIP lookup runs in the browser from a static table. Your ZIP is never sent anywhere.

- ✓ **Pick your issues**
  Mark where you stand and see how your members voted on it. Stances stay in your browser.

- ✓ **Browse topics and votes**
  Each roll call shows the counts by party, a plain-language label and who broke with their party.

- ✓ **Every number links to its source**
  Congress.gov, the House Clerk, senate.gov, the FEC. If it cannot be sourced, it is left out.

- ✓ **Report a mistake**
  A corrections form on every profile. Fixes appear in a dated public change log.

One page per member. Record first, promises beside it.

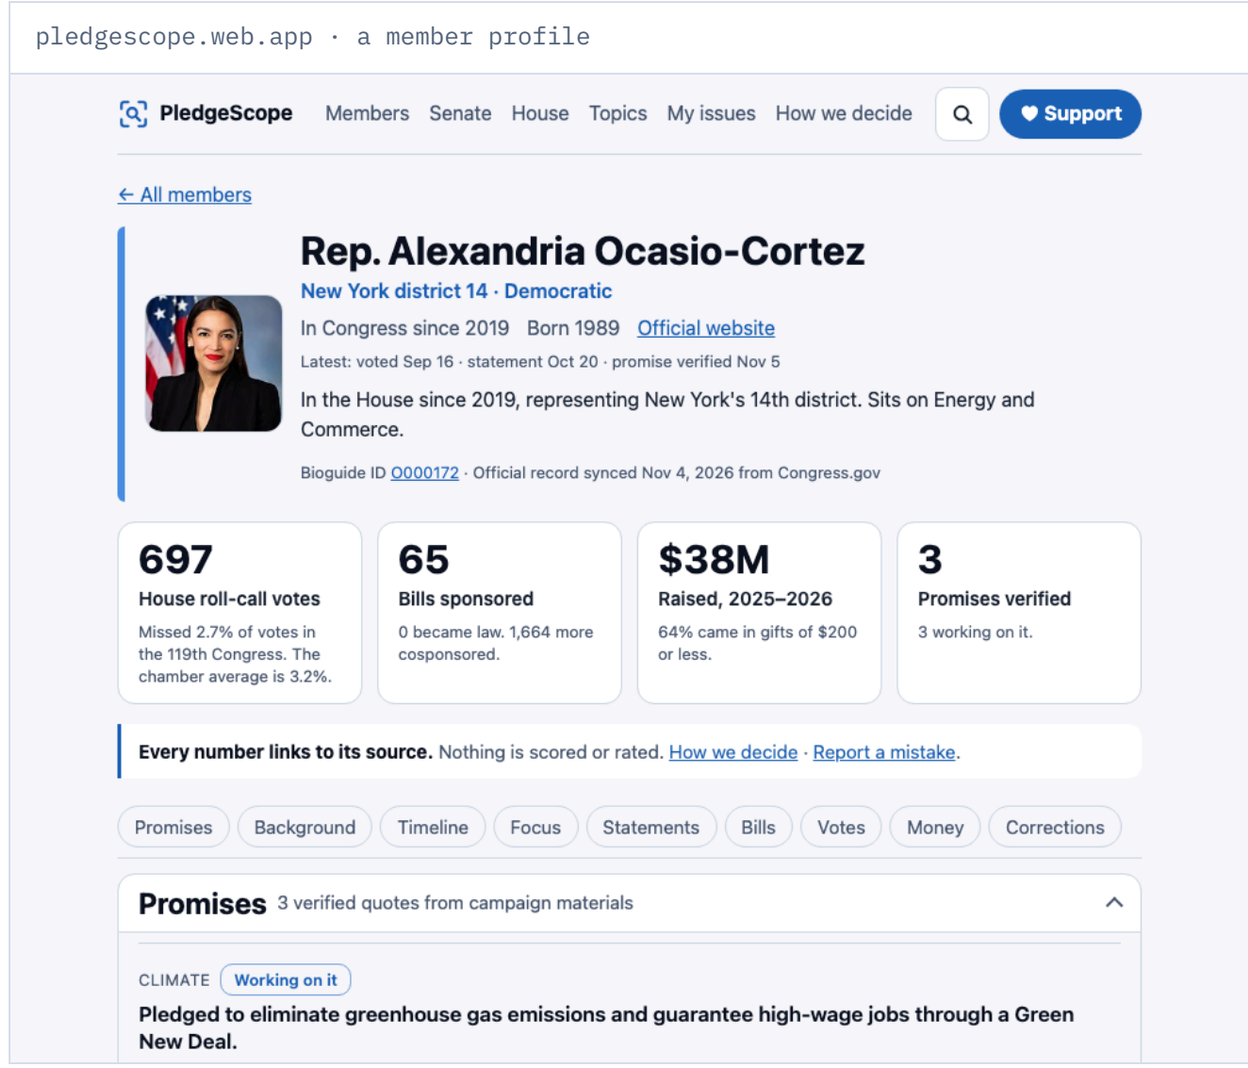
- ✓ **Official record**
  Roll-call votes with missed-vote and party-agreement rates against the chamber average. Bills sponsored and cosponsored across the career. Election history and a timeline.

- ✓ **Money**
  Receipts, gift sizes, top listed employers and outside spending from FEC filings. Described, never interpreted.

- ✓ **Statements**
  Dated press releases from the member's official site.

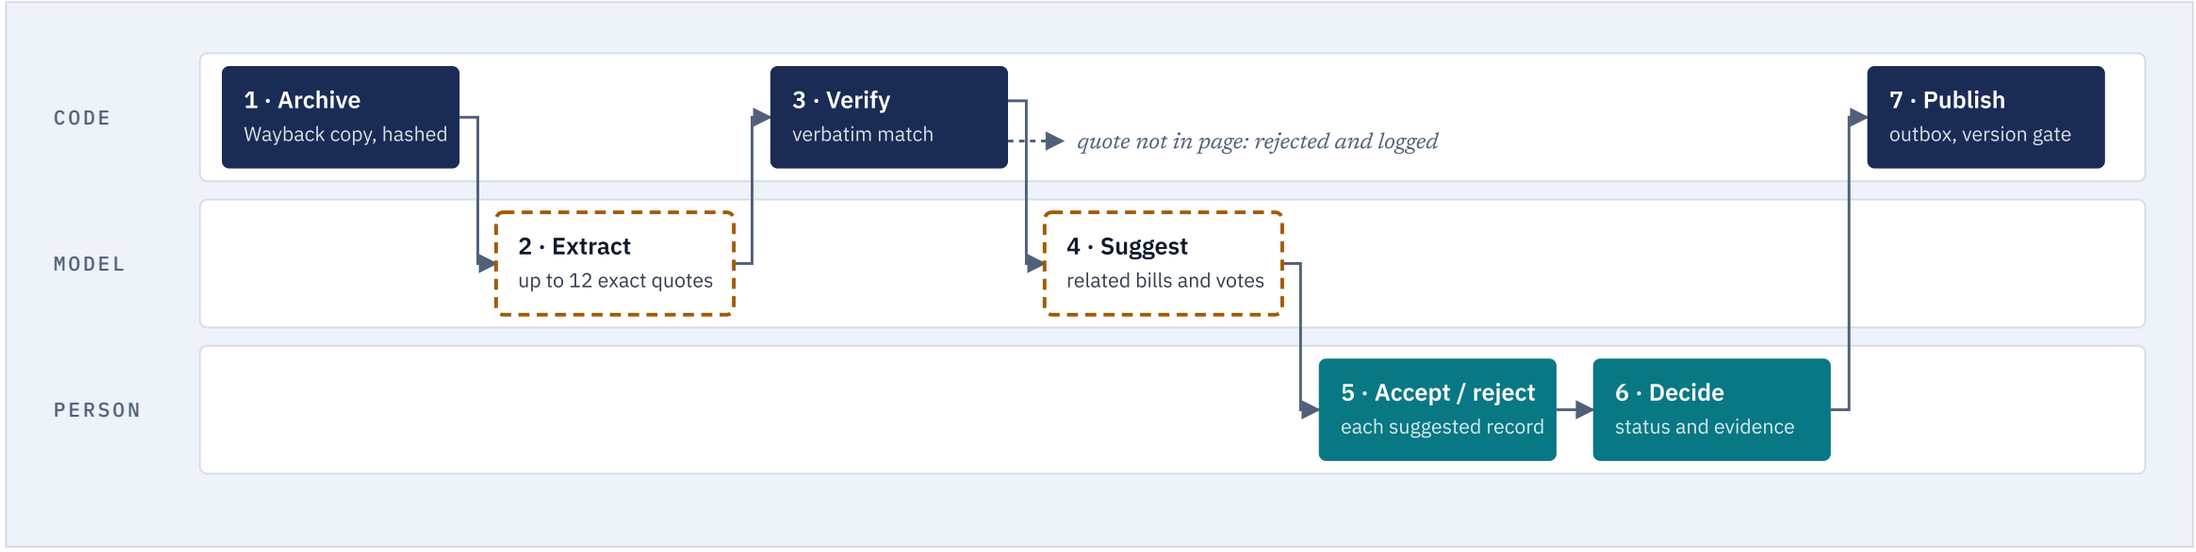
- ✓ **Promises**
  Exact quotes from the member's own campaign site as archived before election day, each with a status and the records behind it.

- ✓ **Corrections made to this page**
  Shown whenever a reader's report changed something.



HOW A PROMISE IS CHECKED

The model reads. Code verifies. A person decides.



CODE Deterministic and tested

Fetches the page as the Internet Archive saved it before election day and stores the bytes with a hash. Checks every returned quote against that text; anything not found verbatim is thrown out. Publishes only through a database that refuses without a recorded approval.

MODEL One bounded call

Gemini reads the page text and returns up to twelve exact quotes, each with a neutral summary and a topic. Later it gets a closed list of the member's bills and votes dated after the promise and says which relate. It never chooses what to read and has no write path.

PERSON Signed in, on the record

An editor accepts or rejects each suggested record, picks a status, writes a rationale and attaches evidence. Every transition is logged under the editor's identity. Nothing is stored until the editor saves it as their own.

Five relationships between a promise and the record.

STATUS	MEANING	EVIDENCE
KEPT	The promised outcome happened and the member was part of it.	Required
WORKING ON IT	Sponsored, cosponsored or voted for it. Not passed, or it failed.	Recommended
NO ACTION FOUND	Nothing related found in bills, cosponsorships or votes.	Not required
CONTRADICTED	Voted against it, or sponsored the opposite.	Required
REVERSED POSITION	Later said the opposite on the record.	Required

Trying and failing counts as keeping your word. A bill that dies in committee is still a bill the member put their name to.

"No action found" is not a verdict. It says what was searched and found nothing. The search is listed.

Nobody is labelled. There is no "liar" or "hypocrite" status, because those describe a person. These five describe a promise against a record, with the records attached.

Safety that holds even if the prompt is wrong.

EXTRACTION

Verbatim check

A quote that is not in the preserved page, character for character after whitespace, is rejected and logged. An invented or paraphrased quote cannot enter the review queue.

EVIDENCE

Immutable snapshots

Every source page and API response is stored as retrieved, with a SHA-256 hash and timestamp. Database triggers refuse edits and deletes.

PUBLISHING

Publish guard in the database

A row cannot reach the published state without a recorded human approval and at least one https evidence link. Application code cannot bypass it.

HISTORY

Append-only assessments

A new assessment adds a version. The previous one stays. Every transition records who, when and why.

PRIVACY

Stripped public documents

Public data carries no model name, editor identity, internal note or stored page copy. Readers' stances and ZIP codes never leave their browser.

PROOF

Evidence pack

For any promise an editor can export the quote, the preserved page with its stored and recomputed hash, a fresh quote check, the review history and every assessment.

AUTOMATION

There are agents. None of them can publish.

More than twenty scheduled workers keep the record current: votes hourly, Senate roll calls daily, campaign finance at the FEC's rate limit, press releases, bill sponsors. Each runs once, does a bounded amount of work and exits. Three of them look at the review queues, and those three are advisory by contract.

AGENT	WHAT IT DOES	MAY WRITE
site-check	Does this archived page look like the member's campaign site? Rules only, no model.	A note
draft-check	Is this quote a pledge, a value statement, biography or fundraising copy? Gemini, batches of ten.	A note
health	Did every job run on time? Is the outbox draining? Is the draft queue growing?	A report

- **Never confirms, approves, rejects or publishes.**
  A note is shown in the editor as "automated, advisory only".
- **Never writes to the public database.**
  Only the outbox job does, and only what a person approved.
- **Never stores a rating the model did not give.**
  A model failure leaves the item blank so it is retried.

The agents are graded, not trusted. A scorecard compares each agent verdict with what editors later decided. A verdict that is right less than 80 percent of the time is flagged, and its rule gets changed. The pipeline also pauses itself at 50 unreviewed drafts, so the queue is paced by people, not by the model.

FOR ENGINEERS

Under the hood

.NET 10 | PostgreSQL | K3s, single node | Gemini | Cloud Firestore | Blazor WebAssembly | Firebase Hosting | Internet Archive

DELIVERY

Transactional outbox

Publishing writes the public document and an outbox row in one transaction. A job claims rows with `FOR UPDATE SKIP LOCKED` and writes to Firestore inside a transaction that applies only if the version is newer. At-least-once delivery, no regressions.

ISOLATION

Two database roles

Runtime pods and the API connect as a role with no DDL that cannot disable triggers. Migrations run only in a separate job under the owner role.

WORKERS

Bounded, budgeted, idempotent

Each worker has a per-run budget, a wall-clock limit and a stop-after-N-failures rule. Same source and same quote are stored once, so a replay is a no-op.

MODEL CALLS

Plain HTTP, no framework

Gemini is called directly with hand-written prompts. Page text is passed as untrusted data inside tags. Candidate records for matching come from a closed set the code selects, never from the model.

RECOVERY

Backups that are proven

Nightly dump. Weekly, the newest dump is restored into a scratch database and row counts are compared with live tables. A mismatch fails the job and the health agent reports it.

PUBLIC SITE

Read-only by rule

Firestore rules allow browsers to read public collections and to create a correction, nothing else. The site carries no API keys. Pages were cut from about 600 reads to a handful.

WHERE IT STANDS

Live, growing, and looking for careful readers.

What surprised me

- ✓ **The first drafts were actions, not promises.**
  "Introduced bill X" is a record, not a pledge. The prompts, the inputs and a pre-reviewer agent now sort them apart.
- ✓ **390 campaign sites block automated visitors.**
  We do not work around robots.txt. Those members' promises wait for a source that allows it.
- ✓ **People are the bottleneck, by design.**
  Hundreds of verified drafts wait for a decision. The editor dashboard now shows each pile, its age and its pace.

Next

- ✓ **Voting patterns.**
  Does a member vote the bill or the sponsor's party? Counts, both sides, caucus baseline. Built, awaiting editor approval.
- ✓ **2026 winners.**
  Their sites were preserved on election eve. Promises become drafts when they are sworn in.
- ✓ **Public launch.**
  Custom domain, abuse protection, a five-person usability test.

Three things you can do

Look up your own representatives at pledgescope.web.app. Report anything wrong; the form is on every profile. And if you are building systems where a model must not have the last word, I would like to compare notes: zoff.tech.