

# Park Slope Eye: When Visibility Isn't the Same as Understanding

An AI discoverability case study from an established independent optometry practice in Brooklyn - showing why search visibility, practice-level retrieval, and downstream AI understanding are different problems.

**1.1M+**

Google Search impressions in the final available 33-day pre-launch window.

**1,529**

Google reviews at the verified September 20 snapshot, with a 4.8-star rating.

**124,183**

Pre-launch property-level Generative-AI impressions reported by Google.

**13 / 22**

Non-branded local-intent retrieval tests where Park Slope Eye was missing.

*The content was visible. The practice itself was not always equally clear.*

# A strong search footprint did not guarantee practice-level discovery.

Park Slope Eye entered the rebuild with years of content, strong branded demand, a large review corpus, and measurable exposure in Google's generative-search reporting. That made it a useful test case because the problem was not simple obscurity.

## WHAT WAS ALREADY WORKING

- ✓ **1,119,893** pre-launch Google Search impressions.
- ✓ **1,914** organic clicks in the same 33-day window.
- ✓ **124,183** property-level Generative-AI impressions.
- ✓ **109** legacy educational articles worth preserving.

## WHAT THE SEPTEMBER 11 RETRIEVAL AUDIT EXPOSED

**13 missing · 4 confused · 3 poor · 2 prominent**

These are the 22 non-branded local-intent results reported separately from the broader governed 26-query audit.

Missing 13

Confused

Poor 3

Prom

*Visibility is not the same thing as practice-level discovery. And being found is not the same thing as being understood correctly.*

**Example:** best optometrist near 11215

**Example:** optometrist that accepts VSP Park Slope

**Example:** eye exam that takes EyeMed in Brooklyn

**Example:** independent optometrist Park Slope


**Evidence boundary:** the September 11 test was a controlled live-web retrieval audit, not 22 native ChatGPT or Gemini conversations. Native AI observations were captured separately after launch. The audit is retrieval evidence, not a universal ranking claim.

# The new site was accurate. Gemini still got two decision-critical insurance facts wrong.

The GSA-built Park Slope Eye website was already live and the first-party insurance facts were correct. Patient-intent testing that evening produced materially wrong conclusions about VSP and EyeMed.

9:51 PM ET · VSP

## In-network became "out-of-network."




Gemini said Park Slope Eye accepted out-of-network VSP benefits and could submit claims for reimbursement.

**Verified fact:** Park Slope Eye is in-network with VSP.

10:18 PM ET · EYEMED

## The same kind of error appeared again.




Gemini described the practice as operating primarily out-of-network and referenced out-of-network EyeMed claim submission.

**Verified fact:** Park Slope Eye is in-network with EyeMed.

These are dated native-AI observations from launch evening. They are not presented as a measure of every Gemini session, every user, or every AI product.

# Gemini offered a hypothesis - not forensic proof.



Gemini said broader out-of-network statements had overridden the more specific VSP information on the dedicated page.

**What this does prove:** an accurate first-party page did not instantly guarantee an accurate downstream answer.

**What this does not prove:** which source, retrieval path, weighting rule, or internal process caused the incorrect answer.

An AI system explaining its own prior output is another model-generated response. We preserve it as a hypothesis, not a verified causal explanation.

## THREE LESSONS

### Visibility is not retrieval

A large search footprint does not mean a practice will surface reliably for local patient intent.

### Retrieval is not understanding

A practice can appear and still be represented incorrectly on a decision-critical fact.

### Launch is not finished

Search, AI, directories, registries, and caches evolve on different schedules.

# The website is only one witness.

A practice can publish accurate first-party information while external directories, registries, legacy pages, and caches still carry conflicting context.

## CONTRADICTIONS DOCUMENTED DURING THE PROJECT

- Stale organization information in NPI/NPPES while a correction was pending.
- Birdeye pairing the practice's review corpus with incorrect hours before correction.
- Optix-now associating outdated services and languages.
- Yelp carrying inaccurate accessibility information before correction.
- Legacy and cached pages surfacing stale hours or service context.

These contradictions do not prove which source Gemini relied on. They show the kind of conflicting public evidence search and AI systems may encounter.

## WHAT GSA CHANGED

- ✓ Strengthened direct VSP, EyeMed, and UFT in-network answers.
- ✓ Separated those relationships more explicitly from Davis Vision and Spectera.
- ✓ Reinforced the distinction in patient-visible and machine-readable information.
- ✓ Redirected stale legacy routes that could continue supplying obsolete context.
- ✓ Continued reconciling external sources where correction was possible.

*Preserve what earned value. Clarify what was weak. Modernize the infrastructure underneath it.*

That approach also preserved **109 legacy educational articles** rather than discarding years of accumulated content simply to produce a cleaner rebuild.

# The failure is documented. The effect of the changes is still being measured.

Implementation evidence and outcome evidence are not the same thing. This dossier keeps those layers separate.

## VERIFIED NOW

- ✓ The current site clearly separates VSP, EyeMed, and UFT from Davis Vision and Spectera.
- ✓ The launch-evening Gemini failures are preserved as dated artifacts.
- ✓ Fresh first-party information and governed retrieval checks can be evaluated against the corrected facts.
- ✓ Production site, redirects, sitemap, and critical patient actions remain under ongoing checks.

## STILL BEING MEASURED

- Native AI retests.
- Broader search visibility and practice-level retrieval.
- AI citation and source-use patterns.
- Patient actions and confirmed appointments as separate evidence layers.
- Business outcomes.

*No ranking improvement, AI-recommendation gain, appointment increase, revenue gain, or ROI is inferred from structural work alone.*

**Evidence classes remain separate:** live-web retrieval, platform-reported search/Generative-AI exposure, native AI observations, external-source contradictions, site actions, confirmed appointments, and business outcomes. AI/search results can vary by product, model, date, location, personalization, and available web evidence.

# See what AI is currently telling patients about your practice.

GSA's complimentary AI Discoverability Report uses a controlled, evidence-preserving review across ChatGPT Search, Google AI Mode, Gemini, and Perplexity to show where answers are accurate, inconsistent, stale, or missing.

## COMPLIMENTARY DIAGNOSTIC

### Get your AI Discoverability Report

See how major AI/search systems currently surface and represent your practice across real patient questions, with dated evidence preserved for review.

## SEE THE DELIVERABLE FIRST

### Review the Park Slope Eye sample

The real two-page sample shows the format, evidence standard, and level of detail without turning one testing window into a universal score or ranking promise.

The report describes dated tested responses. AI/search results can vary by product, model, date, location, personalization, and available web evidence. A favorable or unfavorable response is an observation, not proof of permanent model behavior or a business outcome.

## Want to see what AI says about your practice?

[See the GSA AI Discoverability Report](#) · [Get Your Complimentary AI Discoverability Report](#)

# The same discovery test will be repeated.

At the **Day 30 checkpoint on October 18, 2026**, GSA intends to rerun the same controlled 26-query retrieval audit using the same classification method.

The result will be added to the public case study whether the practice performs **better, the same, or worse**.

| CORE EVIDENCE RECORD                                                                                         |                                                                                        | INTERPRETATION STANDARD                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Search baseline</b><br>Final available 33-day pre-launch Search Console window.                           | <b>Review snapshot</b><br>4.8 stars / 1,529 reviews, verified September 20, 2026.      | <b>Historical baseline evidence</b> remains historical. <b>Current implementation facts</b> describe what is live now. <b>Measured outcomes</b> are reported only when observed. Hypotheses stay labeled as hypotheses.<br><br>This is why implementation, deployment, impressions, referrals, clicks, or isolated AI/search tests are not treated as proof of improved search performance, conversion, revenue, or ROI. |
| <b>Retrieval audit</b><br>September 11, 2026; 22 non-branded results within broader governed 26-query audit. | <b>Native AI artifacts</b><br>September 18, 2026 at 9:51 PM, 9:53 PM, and 10:18 PM ET. |                                                                                                                                                                                                                                                                                                                                                                                                                          |

## Want to discuss what AI says about your practice?

[Discuss Your Practice](#) · [Gulfstream Strategic Advisory](#)

Park Slope Eye is GSA's founding flagship case study. Practice-specific operational facts remain governed in the Park Slope Eye client record; this dossier preserves the dated evidence used for the public case-study narrative.