Translation Layers and AI Answer Probability

Evaluating Translation Layers, Observer Platforms, Markdown Translators, and SEO+ for AI Answer Probability

AnswerShare Research

Primary investigator: Robert Maynard, Jr. (Wikidata Q18157412) [23]

AnswerShare entity: Wikidata Q140196284 [22]

Research assistance: Artificial Intelligence

Date: June 20, 2026

Scope note: This paper **does not endorse AnswerShare, Scrunch, or any other vendor**. It evaluates method classes and public vendor positioning to identify which approach is most likely to increase **AI answer probability** and **Share of Model** under live-retrieval, generative-answer conditions.

Abstract

Generative search changes the unit of competition from rank position and click share to answer inclusion, recommendation, citation, and explanation inside model-generated responses. This paper evaluates the main methods now used to pursue that outcome: conventional SEO and SEO+; programmatic SEO; source-grounded GEO content; markdown translators and llms.txt-style manifests; observer platforms such as SEMrush, Otterly.ai, and Profound; direct implementation layers such as Scrunch AXP and AnswerShare; and machine-readable resource surfaces such as MCP resources, tools, and content indexes. The evaluation criterion is **AI answer probability**: the probability that a relevant model output mentions, cites, recommends, or materially uses a brand, entity, product, or source. **Share of Model** is treated as the measurable expression of that probability across a prompt panel and competitive set.

The conclusion is deliberately **method-based rather than vendor-promotional**. If the objective is to increase **AI answer probability** - the probability that a model mentions, cites, recommends, or materially uses a brand in response to a relevant prompt - the **best-supported method class is a translation layer**. In this paper, a translation layer means an implementation-led retrieval surface that presents faithful, source-grounded, machine-readable content to AI agents while preserving the human site. This claim is stronger than saying that dashboards, markdown conversion, or SEO are useful. Those methods are useful, but they are indirect or partial. Observer platforms detect visibility gaps; markdown and llms.txt reduce representation friction; SEO keeps the domain discoverable. **A governed translation layer is the only evaluated method class that directly changes what AI systems fetch, parse, ground, and reuse at answer time.** Scrunch and AnswerShare are direct comparators in this class because both publicly describe parallel or translation-layer delivery for AI agents. SEMrush, Otterly.ai, and Profound are treated primarily as observer and measurement platforms.

This paper concludes that a faithful, source-grounded translation layer is the best-supported method for increasing AI answer probability and Share of Model because it directly modifies the machine-readable retrieval surface that AI systems crawl, parse, compare, and cite.

Contents

1. Executive findings

Principal finding: Among the evaluated methods, the translation-layer model has the strongest expected effect on AI answer probability because it intervenes at the retrieval substrate itself. Observer platforms such as **SEMrush, Otterly.ai, and Profound** help measure visibility; SEO+, schema, markdown conversion, and `llms.txt` improve portions of the signal environment; but a translation layer changes what AI systems are able to retrieve, interpret, and cite.

The central finding is that **Share of Model cannot be optimized reliably with a single traditional marketing technique.** Generative systems retrieve, summarize, ground, cite, and recommend; therefore, **the most effective method is the one that controls the retrievable representation of the brand** while preserving human-facing content, improving evidence density, and measuring the resulting model outputs. Public evidence supports treating this as a **retrieval-infrastructure problem, not only a keyword or rank-tracking problem**. Google describes generative AI features as relying on retrieval-augmented generation and query fan-out, while Bain reports that AI summaries are already changing click behavior at broad consumer scale. [4,1]

The practical market landscape has three layers. First, **observer platforms measure what models say**. SEMrush, Otterly.ai, and Profound provide visibility, prompt, citation, sentiment, and share-of-voice data. [9,10,11] Second, **transport adapters make content easier to ingest**. Cloudflare Markdown for Agents converts HTML to Markdown when a client asks for text/markdown, while llms.txt proposes a Markdown manifest for inference-time website use. [13,15]

Third, **implementation-led delivery layers alter the machine-facing substrate itself**. Scrunch AXP and AnswerShare both describe parallel or translation-layer delivery to AI agents, with humans and search-only bots continuing to see the normal site. [6,7,20]

The strongest expected increase in AI answer probability should come from the third layer when, and **only when, the layer preserves content parity and carries genuinely citable evidence.** The reason is causal. Observer tools identify gaps, but do not by themselves make a JavaScript-heavy, slow, thin, or poorly grounded site easier for a retrieval bot to parse.

Markdown conversion reduces representation friction, but does not create facts, evidence, entity consistency, freshness, or trust. SEO improves discovery, but Google's own generative-AI guidance warns against scaled pages created primarily to manipulate rankings or generative AI responses. [4,5]

Translation layers have the highest method-level leverage because they can combine crawl efficiency, clean rendering, canonical evidence, entity identifiers, source grounding, freshness, and measurement in the same AI-facing surface. [6,7,19,20]

This paper therefore ranks methods as follows for direct effect on AI answer probability in order: (1) translation layers combined with grounded content; (2) source-grounded GEO content architecture; (3) machine-readable resources, MCP-style context surfaces, and content indexes; (4) markdown/content-negotiation adapters as supporting infrastructure; (5) observer tools as required measurement and prioritization, not direct remediation; (6) SEO+ and structured data as **baseline hygiene**; and (7) thin programmatic SEO as high-risk and often counterproductive unless it contains unique data and user value. This is a **method conclusion, not a vendor endorsement**.

2. Research basis and evidence grading

The evidence base combines the supplied AnswerShare Research drafts with public, live resources from academic research, Google documentation, Cloudflare documentation, vendor documentation, and public entity records. The earlier research draft emphasized the acquisition-efficiency shift created by zero-click and AI-mediated search, while the positioning draft emphasized the need to compare AnswerShare against implementation, observer, and transport alternatives rather than against legacy SEO alone.

This paper distinguishes four evidence types. **Academic and platform documentation are treated as the strongest basis for general method claims. Vendor documentation is treated as authoritative for what a vendor says its system does, but not as independent proof of performance. Vendor case studies and self-reported metrics are labeled as claims unless accompanied by independently verifiable data.** Supplied internal research is used to structure the analysis but is not used as a substitute for live-source footnotes.

Evidence class	Used for	Examples in this paper	Interpretation
Academic / technical research	General mechanism and method theory	GEO paper; MCP specification	High value for method design; not necessarily direct commercial proof. [3,16,17]
Platform documentation	Rules and retrieval mechanics	Google generative AI guidance; Google spam policies; crawler docs	High value for compliance and platform-specific constraints. [4,5,25,26,27]
Official vendor documentation	What a vendor product claims to do	Scrunch AXP, SEMrush Toolkit, Otterly, Profound, Cloudflare	Valid for classification; performance claims require additional validation. [6,9,10,11,13]
Vendor proof or self-measurement	Implementation examples and auditability	AnswerShare FAQ/methodology; Top10Lists crawl statistics	Useful if receipts are public; should be labeled as vendor-published evidence. [19,21,24]

3. Why Share of Model is the correct evaluation metric

Traditional search rewarded a position in a ranked list. Generative search often suppresses or compresses the list, and the user may receive a synthesized answer before deciding whether to click. Bain reports that about 60% of traditional-search sessions now end without the user progressing to another destination, and Ahrefs reports that Google AI Overviews correlate with materially lower click-through to the top organic result. [1,2]

Share of Model is a better leading indicator for this environment because it measures whether the model includes the entity in the answer. A practical SoM program should separate **four signals: mention share, citation share, recommendation share, and answer absorption**. Mention share asks whether the brand appears. Citation share asks whether the model links to or names a source controlled or materially earned by the brand. Recommendation share asks whether the model chooses or ranks the brand as a fit. Answer absorption asks whether the model uses the brand's facts without naming or linking it.

The measurement formula can be simple at first: **SoM = brand-positive outputs divided by total relevant category outputs across a controlled prompt set.** Mature programs should weight by platform, prompt intent, answer position, sentiment, citation type, and whether the output occurred under live retrieval. Google's generative AI guidance makes this platform-sensitive because its systems use retrieval-augmented generation and query fan-out, not merely one static keyword query. [4]

The implication is direct: **a method that cannot change what retrieval bots can fetch, parse, trust, or ground should not be expected to move SoM by itself.** It can show the problem, prioritize work, or provide content suggestions, but it does not control the supply of machine-usable evidence.

4. Method taxonomy

The methods in this paper are evaluated by their likely effect on **AI answer probability**, defined as the probability that a brand, page, entity, or source is retrieved, trusted, synthesized, and cited in an AI-generated answer. Under this standard, the strongest method is not necessarily the one that produces the most reports, dashboards, or content artifacts. The strongest method is the one that most directly improves the machine-facing substrate used by AI systems during retrieval and answer construction.

The methods in this paper are evaluated by their likely effect on **AI answer probability**, defined as the probability that a brand, page, entity, or source is retrieved, trusted, synthesized, and cited in an AI-generated answer. Under this standard, the strongest method is not necessarily the one that produces the most reports, dashboards, or content artifacts. The strongest method is the one that most directly improves the machine-facing substrate used by AI systems during retrieval and answer construction.

The methods in this market are often bundled user AEO, GEO, LLMO, AI visibility, or AI search optimization. Those labels hide important differences. For **Share of Model**, the relevant distinction is **whether a method observes model outputs, improves human-search eligibility, adapts the representation layer, or changes the evidence corpus that models retrieve**.

Method class	Primary function	Representative examples	Direct SoM mechanism
Observer / measurement platforms	Track model answers, prompts, mentions, citations, sentiment, and competitors.	SEMrush, Otterly.ai, Profound. [9,10,11]	Indirect. They identify where SoM is lost and whether interventions worked.
Implementation-led AI delivery layers	Serve a clean, structured, agent-oriented version of site content while preserving the human site.	Scrunch AXP; AnswerShare translation layer. [6,7,20]	**Direct. They change the content representation encountered by AI retrieval systems.**
Markdown translators / content negotiation	Convert existing HTML into Markdown or expose Markdown files to reduce extraction friction.	Cloudflare Markdown for Agents; DIY markdown endpoints; llms.txt. [13,15]	Medium. They improve transport but do not alone create evidence or trust.
SEO+ and structured data	Maintain crawlability, index eligibility, schema, entity clarity, and useful content.	Google SEO and structured-data guidance. [4,28]	Baseline. It supports discovery and indexing but does not guarantee model inclusion.
Source-grounded GEO content	Add facts, statistics, citations, quotes, primary research, and clear entity definitions.	Academic GEO methods; AnswerShare source-grounding concepts. [3,19]	**High. It increases the factual material that models can cite and justify.**
Machine-readable resource surfaces	Expose data, resources, and tool endpoints for model or agent use.	MCP resources and tools; ai-content-index style feeds. [16,17,19]	**High for compatible agents; dependent on adoption and governance.**

5. Method-by-method evaluation

5.1 Conventional SEO and SEO+

SEO remains necessary because several generative search experiences rely on web indexes and retrieval systems that inherit classic crawlability, quality, and technical constraints. Google states that its generative AI features are rooted in core Search ranking and quality systems and that foundational SEO practices continue to apply. [4]

However, **SEO is not sufficient as a Share of Model strategy.** SEO is designed for discoverability, index eligibility, and ranked retrieval. SoM depends on whether the model decides to include, recommend, or cite the entity inside a synthesized answer. A page can be technically indexable yet still be too repetitive, slow, ungrounded, or poorly structured to be useful to an AI answer system. Conversely, a model may use information from sources that do not correspond to the same page-one ranking pattern that SEO teams track.

SEO+ should therefore be treated as **baseline hygiene**: semantic headings, accessible content, good internal linking, structured data, crawlable pages, canonical entity pages, and non-commodity content. Google's guidance explicitly favors unique, helpful, non-commodity content and warns against creating many pages primarily to manipulate rankings or generative AI responses. [4,5]

5.2 Programmatic SEO

Programmatic SEO can help when it produces genuinely distinct, useful pages based on unique data. It is **dangerous when it produces many thin pages with minor variable swaps.** Google defines scaled content abuse as generating many pages primarily to manipulate rankings and not help users, including pages generated with AI tools, scraped feeds, or automated transformations that add little value. [5]

For SoM, the practical problem is not only spam risk. The problem is **information gain**. A model has little reason to cite a templated page that restates commodity information already present in its retrieval set. Programmatic content should therefore be evaluated by uniqueness ratio, source grounding, freshness, and whether each page answers a real fan-out question with a data point or expert claim that is not duplicated elsewhere on the site.

5.3 Source-grounded GEO content

Source-grounded GEO content is the highest-value content layer because it creates the **evidence models need in order to cite or recommend a source**. The academic GEO paper introduced a black-box optimization framework and found that optimization techniques can improve visibility in generative-engine responses by up to 40%, with performance varying by domain. [3]

The method is concrete: define the entity, state the claim, ground the claim, show the data, identify the source, provide freshness metadata, and make the answer easy to extract without excess page chrome. AnswerShare's public methodology uses a Source Grounding Ratio and Retrieval Token Cost as reasoning signals, which reflects the same principle: generative retrieval systems need assertions that are both verifiable and cheap to process. [19]

Source-grounded GEO should **not be confused with keyword stuffing**. It is closer to building a machine-readable evidence record: concise definitions, canonical claims, primary research, third-party citations, last-modified dates, and consistent entity identifiers across the site and across external profiles.

5.4 Markdown translators and content-negotiation adapters

Markdown translators are important but limited. Cloudflare Markdown for Agents lets a client request text/markdown through the Accept header; Cloudflare then fetches the original HTML and converts it to Markdown before serving it. [13] Cloudflare's launch article frames this as a way for AI agents to avoid complex conversion work and receive structured Markdown from the source. [14]

The strength of markdown translation is token efficiency and extraction simplicity. The weakness is that **conversion is not the same as strategy. A Markdown rendering of a weak page is still weak**; it may be easier to read, but it does not necessarily contain unique data, source-grounded claims, entity alignment, or prompt-specific answer coverage. Markdown should therefore be considered a **transport layer beneath GEO, not a substitute for GEO**.

The same conclusion applies to llms.txt. The llms.txt proposal gives websites a Markdown file that provides LLM-friendly background, guidance, and links, and it explicitly recognizes that processing depends on the application. [15] That makes llms.txt useful as a manifest and orientation file, but **not a guaranteed ranking or citation lever**. A well-built llms.txt can help agents find the correct documents; it cannot compensate for an ungrounded or stale evidence corpus.

5.5 Implementation-led AI delivery layers

Implementation-led delivery layers are the direct method class most relevant to Share of Model. Scrunch AXP publicly describes a parallel AI-ready version of site pages and says AI agent traffic can be routed through a CDN to a stripped-down structured version while humans keep the full experience. [6] Its FAQ describes AI traffic detection, CDN routing, and AI-optimized server-rendered HTML without changing the human experience. [7]

AnswerShare publicly describes a translation layer between the site and AI systems, with a parallel, clean-room, no-JavaScript representation for AI crawlers and measurement across retrieval confidence, semantic trust, and citation grounding. [20] Its methodology describes AI Opinion, Infrastructure Readiness, and Reasoning Keys, including query fan-out survivability, source grounding, retrieval token cost, MCP, and ai-content-index.json. [19]

The shared strength of this category is that it **changes the surface the model experiences at fetch time**. **This is why Scrunch should not be classified with observer dashboards.** It does some observation, but its AXP offering is an implementation method. The same is true of AnswerShare's translation-layer positioning. **A direct implementation layer is most likely to change SoM** when the bottleneck is crawlability, JavaScript rendering, page-chrome overhead, entity ambiguity, stale content, or poor source grounding.

The shared risk is governance. If a machine-facing representation materially diverges from the human page in a way intended to manipulate search engines or mislead users, it can approach the territory described by Google as cloaking. [5] **A compliant implementation must preserve content parity**, use the human site as the source of truth, avoid hidden claims, and maintain audit logs showing what was served to which bot and when.

5.6 Machine-readable resource surfaces

Machine-readable resources represent the next layer of implementation. MCP resources provide a standardized way for servers to expose context such as files, database schemas, or application-specific information; MCP tools expose callable capabilities with schema metadata. [16,17] AnswerShare's public methodology references MCP endpoints and ai-content-index.json as part of its infrastructure-readiness family. [19]

These approaches are promising because they **move beyond page scraping**. They allow the site or brand to expose canonical resources in a form that an agent can inspect or query. Their limitation is ecosystem adoption: not every AI search experience will use a site's MCP endpoint or custom manifest. Therefore, **machine-readable resources should complement, not replace, clean HTML, Markdown, structured data, and source-grounded web pages**.

6. Scrunch and AnswerShare as direct implementation comparators

The revised classification treats Scrunch and AnswerShare as direct implementation comparators. Scrunch's AXP documentation says it creates a parallel AI-optimized version of the site, served invisibly to AI agents while human visitors continue to see the existing experience. [8] AnswerShare describes its own product as a translation layer that sits between a site and AI systems, with translate, measure, and optimize stages. [20]

The two products are not identical, but **they are in the same method class**. Both attempt to solve the same core problem: **AI systems often do not experience a website the way humans do**. JavaScript, dynamic navigation, layout chrome, client-side rendering, and weak source grounding can impair machine extraction even when the human page is attractive.

Dimension	Scrunch AXP	AnswerShare	Method implication
Public positioning	Parallel AI-ready version of pages served to AI agents. [6,8]	Translation layer between websites and AI systems. [20]	Both are implementation-layer products, not merely analytics dashboards.
Delivery mechanism	Public copy describes CDN routing, AI traffic detection, and server-rendered structured HTML. [7]	Public methodology describes edge detection/rerouting concepts, clean-room machine-facing content, infrastructure metrics, and AI-content manifests. [19,20]	Both intervene before or during agent retrieval, which gives the category direct SoM relevance.
Measurement layer	Scrunch describes monitoring of brand presence, citations, placement, sentiment, and change over time. [7]	AnswerShare describes ASQ, AI Opinion, Infrastructure Readiness, Reasoning Keys, and per-site receipts. [19,21]	Both combine implementation with measurement, though the published metrics differ.

Dimension	Scrunch AXP	AnswerShare	Method implication
Public proof posture	Public documentation emphasizes product workflow; case studies may vary by client availability. [6]	AnswerShare publishes Top10Lists proof claims and crawl-stat methodology as vendor-published evidence. [21,24]	Vendor claims should be validated with independent prompt panels and logs before being treated as universal proof.
Main risk	A parallel representation must preserve content parity and avoid misleading search crawlers or users. [5]	Same risk: translation layers require governance, receipts, and parity controls. [5,19]	The category is powerful but must be auditable.

On method grounds, Scrunch and AnswerShare should be evaluated against each other on **retrieval parity, representation quality, source grounding, speed, fan-out coverage, auditability, crawler logs, and observed SoM lift** across controlled prompt panels. A fair trial would run the same prompt set before and after implementation, include SEMrush/Otterly/Profound or another observer as an external measurement layer, and compare raw server logs against model-output changes.

7. Observer platforms: SEMrush, Otterly.ai, and Profound

Observer platforms are essential, but their job is measurement and prioritization. SEMrush describes its AI Visibility Toolkit as measuring how brands appear in AI-generated answers, benchmarking brand visibility, analyzing competitors, tracking prompts, auditing technical blockers, and identifying gaps. [9] Otterly.ai describes AI search monitoring across Google AI Overviews, AI Mode, ChatGPT, Gemini, Copilot, and Perplexity, including brand mentions, citations, and share of AI voice. [10] Profound describes tracking front-end AI experiences, visibility scores, share of voice, sentiment, citation sources, authority, competitor rankings, regions, topics, and personas. [11]

This paper classifies all three as observers because **their primary Share of Model function is to reveal model behavior**. Profound also markets agents, prompt volumes, content workflows, and site crawl analytics, so it has activation features. [12] Still, the public Answer Engine Insights positioning centers on measuring how AI systems represent a brand and where competitors appear. [11]

Observer platforms can improve SoM indirectly in three ways: they reveal the prompts that matter; they expose which competitors and third-party sources models cite; and they validate whether an implementation or content intervention worked. They should be included in any serious SoM program, but they **should not be mistaken for the intervention itself**.

Observer	Primary observed signals	Best use	Limit for SoM
SEMrush AI Visibility Toolkit	AI visibility, mentions, competitors, prompts, technical blockers, reports. [9]	Teams already using SEO workflows that need prompt tracking and competitive AI visibility.	Primarily diagnostic; remediation depends on subsequent content or infrastructure changes.
Otterly.ai	Brand mentions, citations, link changes, AI Search monitoring, share of AI voice. [10]	Focused monitoring across major AI search engines and alerting on visibility changes.	Can show gains or losses but cannot by itself alter the content a bot receives.
Profound	Front-end AI experiences, share of voice, sentiment, citation sources, competitor ranking, prompt volumes. [11,12]	Enterprise answer-engine intelligence, PR/brand monitoring, and prompt-demand mapping.	Strong observer and workflow layer; implementation lift depends on actions taken after insights.

8. Markdown translators and llms.txt

Markdown translators deserve a separate section because they are often confused with full GEO. Cloudflare Markdown for Agents is a **protocol and representation feature**: when a compatible client sends an Accept header including text/markdown, Cloudflare converts the origin HTML to Markdown and serves that representation. [13]

This can materially help with token cost and extraction quality. Many AI retrieval systems would prefer clean text to full client-side application markup. However, the conversion process starts from the original page. If the original page lacks primary facts, current dates, external citations, structured definitions, or coherent entity information, the Markdown version will inherit those defects.

The same distinction applies to llms.txt. It is an **orientation and manifest convention, not a ranking guarantee**. The proposal recommends a Markdown file with background information, guidance, and links to detailed Markdown files, while noting that processing depends on the application. [15]

8.1 Evaluation of markdown methods

Markdown method	What it solves	What it does not solve	Expected SoM impact
Cloudflare Markdown for Agents	Reduces HTML-to-text conversion friction using content negotiation. [13,14]	Does not create evidence, external trust, or prompt coverage.	Medium as an infrastructure component; high only when paired with grounded content.
llms.txt	Gives agents a readable manifest and links to important Markdown resources. [15]	Does not guarantee that AI systems will fetch, honor, cite, or recommend it.	Low to medium by itself; useful as readiness hygiene.
DIY Markdown pages	Can expose concise AI-readable versions of complex pages.	Can create governance, drift, and cloaking risk if it diverges from the human source of truth. [5]	Medium if parity-controlled and source-grounded; risky if used as hidden alternate content.
Full translation layer	Can combine Markdown/HTML cleanup with canonical facts, entity IDs, source grounding, freshness, and logs.	Requires implementation governance and measurement.	Highest direct SoM potential when paired with observer measurement.

9. Scorecard for expected Share of Model impact

The scorecard below evaluates method classes rather than vendors. **Scores reflect expected direct effect on AI answer probability and Share of Model when executed competently, not guaranteed outcomes.** A score of 5 means **the method directly changes a key causal input to model inclusion**; a score of 1 means the method primarily observes or supports other work.

Method class	Direct SoM leverage	Reason	Best role in stack
Implementation-led translation layer	5 / 5	Changes what agents can fetch, parse, and use; directly addresses retrieval friction, token overhead, entity representation, and source grounding. [6,7,19,20]	Core intervention layer.
Source-grounded GEO content	4.5 / 5	Creates the evidence, data, definitions, and citations that models can justify using. Academic GEO work reports visibility gains from optimization techniques. [3]	Core content layer.

Method class	Direct SoM leverage	Reason	Best role in stack
Machine-readable resources / MCP / content indexes	4 / 5	Can expose canonical data beyond page scraping, but depends on agent support and governance. [16,17,19]	Advanced infrastructure layer.
Markdown translators / content negotiation	3 / 5	Reduces representation friction, especially for HTML-heavy pages, but does not guarantee trust or inclusion. [13,14]	Transport and efficiency layer.
llms.txt / manifest files	2 / 5	Provides orientation and links, but processing depends on the application. [15]	Low-cost readiness layer.
Observer platforms	1.5 / 5 direct; 5 / 5 diagnostic	Measure prompts, citations, share of voice, sentiment, and gaps; indirect impact depends on follow-through. [9,10,11]	Measurement and prioritization layer.
SEO+ and structured data	2.5 / 5 direct; 4 / 5 baseline	Supports crawlability, index eligibility, and content clarity, but does not alone control generative inclusion. [4,28]	Baseline hygiene layer.
Thin programmatic SEO	0.5 to 3 / 5	Useful only when pages add unique value; risky when scaled primarily for manipulation. [5]	Conditional; avoid thin templates.

9.1 Method conclusion

The most defensible conclusion is that a translation layer is the best-supported model for increasing AI answer probability and Share of Model, provided the layer is faithful to the human-facing content, source-grounded, parity-controlled, crawl-governed, and measured before and after deployment.

This conclusion is not a claim that every translation-layer implementation will outperform every SEO, markdown, observer, or structured-data tactic in every vertical. It is a comparative method conclusion: the translation layer has the greatest causal leverage because it operates directly on the retrieval surface AI systems consume. Observer platforms measure whether a brand appears. SEO+ and structured-data methods improve pieces of the underlying content environment. Markdown and protocol tools improve transport or representation. A translation layer, by contrast, changes the machine-facing substrate

The paper can make the stronger conclusion: **the translation layer is the best-supported method class for increasing AI answer probability**, provided that it is implemented as a **faithful, governed representation of the same substantive content**. It is not best because it has a brand name attached to it. **It is best because it is closest to the retrieval event.** It changes the surface the model can fetch, the amount of noise the model must discard, the availability of canonical facts, the freshness and source metadata attached to those facts, and the audit trail needed to connect bot access to later answer inclusion. [6,7,13,19,20]

The conclusion should not be stated as universal proof. No public cross-vendor randomized study currently proves that every translation layer beats every content, SEO, or dashboard implementation in every vertical. The defensible statement is narrower and stronger: among the evaluated method classes, a translation layer has the highest direct causal leverage on **AI answer probability** because it acts on **the retrievable representation itself**. GEO content supplies the evidence; SEO keeps discovery possible; markdown and llms.txt improve machine access; observers measure whether answers changed. The translation layer is the model that can **unify those inputs into the actual machine-facing substrate**.

The highest-performing stack is therefore: (1) maintain SEO and structured data so discovery is not blocked; (2) build a grounded evidence corpus; (3) expose the corpus through clean HTML/Markdown, manifests, and machine-readable resources; (4) use a translation or agent-experience layer where the human site is difficult for bots to parse; (5) deliver the data to optimize affordability for the AI and (6) monitor SoM and answer probability with an observer platform before and after implementation.

9.2 Truth status of the translation-layer claim

The claim is true at the level of causal method comparison, but **not yet proven as an absolute empirical law.** The evidence supports the following formulation: for a site that wants to increase the probability of appearing in AI answers, a translation layer is the **strongest single intervention class** when the existing human site creates retrieval friction and when the translation layer carries faithful, **source-grounded, current, citable content**. Scrunch states that AXP creates a parallel AI-ready version of pages, intercepts AI retrieval traffic at the CDN layer, translates pages into clean server-rendered HTML, and serves compressed, structured output mapped to sources of truth. [6,7] AnswerShare states that AI crawlers route through a parallel clean-room, no-JavaScript translation layer with identical content structured for retrieval. [20]

The same evidence also sets limits. A translation layer **will not rescue content that is inaccurate, unsupported, undifferentiated, or contradicted by stronger third-party sources.** It should therefore be expressed as the best model for increasing answer probability only **when it is paired with source-grounded GEO content, entity consistency, content parity, crawl governance, and observer measurement**. In this qualified form, the conclusion is stronger, more useful, and more defensible than the previous wording.

10. Implementation architecture and governance

A serious SoM program should be operated as a measured system rather than as a content campaign. The system needs **baseline measurement, controlled prompts, source-grounded content, machine-facing delivery, crawler logging, and post-deployment validation**.

10.1 Baseline measurement

Create a controlled prompt panel that includes category prompts, comparison prompts, local or vertical prompts, problem-solution prompts, and purchase-intent prompts. Run the panel across ChatGPT, Perplexity, Gemini, Claude, Copilot, Google AI Overviews, and Google AI Mode where feasible. Observer platforms such as SEMrush, Otterly.ai, and Profound can accelerate this step because they already track brand visibility, prompts, citations, sentiment, and competitor presence. [9,10,11]

10.2 Evidence corpus

For each strategic prompt, identify the factual claims a model would need in order to recommend the brand. Then create source-grounded pages or data records that state the claim, provide the evidence, cite external sources, expose structured data, and keep last-modified dates current. Google's guidance favors unique, non-commodity content and technical crawlability, while the GEO paper shows that generative-engine visibility responds to content optimization methods. [4,3]

10.3 Machine-facing delivery

If the site is simple, well-rendered, and content-dense, SEO plus structured data plus Markdown may be adequate. **If the site is JavaScript-heavy, slow, navigation-heavy, or hard for bots to parse**, **an implementation-led delivery layer is more likely to move SoM**. Scrunch and AnswerShare both publicly describe serving AI agents a **clean, structured, machine-facing representation** while preserving the human experience. [6,7,20]

10.4 Crawler management and logs

AI crawler behavior is not one crawler and one purpose. OpenAI documents crawlers and user agents used for product actions and robots.txt controls; Anthropic documents ClaudeBot controls; Perplexity documents its crawlers and robots.txt controls. [25,26,27] **An SoM system should log verified bot user agents**, page paths, cache status, response size, status codes, and whether the machine-facing version was served. Top10Lists.us publishes a vendor-owned example of crawl-stat methodology using Cloudflare GraphQL Analytics and edge middleware logs. [24]

10.5 Governance controls

The governance rule is simple: the machine-facing version should be a **faithful, cleaner representation of the same substantive content, not a hidden alternate reality**. Google defines cloaking as presenting different content to users and search engines with the intent to manipulate rankings and mislead users. [5] **Translation layers should therefore maintain parity checks, change logs, source-of-truth controls, and audit receipts.**

10.6 Suggested SoM test design

1. Define 50 to 200 prompts across category, comparison, local, and purchase-intent use cases.
2. Run the prompts across major AI answer systems with location, device, and account state controlled where possible.
3. Record mentions, citations, recommendations, answer position, sentiment, and whether the brand's own URLs or third-party URLs were used.
4. Deploy the intervention: SEO hygiene, grounded content, Markdown, llms.txt, translation layer, or a combination.
5. Run the same prompt panel weekly for 8 to 12 weeks and compare changes against a control group of pages or entities.
6. **Use server logs to confirm that AI agents actually fetched the pages or machine representations** before attributing model-output changes to the intervention.

11. Conclusion

Translation-layer architecture is the strongest method class for increasing AI answer probability; AnswerShare and Scrunch are evaluated as examples of that method class, while SEMrush, Otterly.ai, and Profound are evaluated as observer systems.

The answer to the method question is a **qualified yes**: a translation layer is the best-supported model for increasing **AI answer probability** among the method classes evaluated in this paper. The reason is not branding. **The reason is system position.**

AI answer probability is decided after a model or retrieval system discovers, fetches, parses, compresses, grounds, and compares candidate evidence. **A translation layer acts directly on that path** by changing the machine-facing representation that the AI system receives. SEO, markdown, llms.txt, observer dashboards, and content optimization each help, but **none of them alone occupies as many points in the retrieval-to-answer chain**.

The claim should be stated precisely. A translation layer is the strongest single method class when it is faithful to the human-facing content, removes retrieval friction, exposes source-grounded evidence, maintains canonical entity and freshness signals, logs crawler behavior, and is measured against a prompt panel before and after deployment. It is not a license to cloak, invent claims, or serve a materially different page to machines. Google's spam policies make that distinction material: alternate representations must not be used to mislead users or manipulate ranking systems. [5]

Scrunch belongs in the direct implementation comparator set because its AXP publicly creates a parallel AI-ready version of pages served to AI agents, including CDN interception, clean server-rendered HTML, structured output, and source-of-truth mapping. [6,7] AnswerShare belongs in **the same method class** because it publicly describes a parallel clean-room, no-JavaScript translation layer plus measurement across retrieval confidence, semantic trust, and citation grounding. [19,20] SEMrush, Otterly.ai, and Profound belong primarily in the observer set because their public materials center on tracking visibility, prompts, citations, share of voice, sentiment, competitor presence, and AI crawler behavior. [9,10,11]

The final method conclusion is therefore: for increasing **AI answer probability** and **Share of Model**, **the most effective model is a governed translation layer**, supported by source-grounded GEO content and verified through observer measurement. In a practical stack, SEMrush, Otterly.ai, and Profound reveal where answer probability is being won or lost; Cloudflare Markdown for Agents and llms.txt improve machine access; SEO and structured data keep the site eligible and discoverable; **but the translation layer is the direct intervention most likely to move the outcome** because it controls the citable machine-facing substrate.

This conclusion does not endorse AnswerShare over Scrunch or any other vendor. It ranks the method. Vendor selection should be based on public receipts, controlled trials, crawl logs, prompt-panel deltas, content parity, compliance controls, and measured **Share of Model** lift against a predefined baseline.

Sources and Notes

All sources below are live resources used to ground the whitepaper. Accessed June 20, 2026 unless otherwise noted.

[1] Bain & Company, "Consumer reliance on AI search results signals new era of marketing". Reports that about 80% of search users rely on AI summaries at least 40% of the time, that about 60% of traditional-search sessions end without a click onward, and that organic traffic is reduced by an estimated 15% to 25%. URL: https://www.bain.com/about/media-center/press-releases/20252/consumer-reliance-on-ai-search-results-signals-new-era-of-marketing--bain--company-about-80-of-search-users-rely-on-ai-summaries-at-least-40-of-the-time-on-traditional-search-engines-about-60-of-searches-now-end-without-the-user-progressing-to-a/

[2] Ahrefs, "Update: AI Overviews Reduce Clicks by 58%". Reports an updated December 2025 analysis finding that the presence of a Google AI Overview correlates with a 58% lower average CTR for the top-ranking page. URL: https://ahrefs.com/blog/ai-overviews-reduce-clicks-update/

[3] Aggarwal et al., "GEO: Generative Engine Optimization". Introduces GEO-bench and reports that GEO techniques can improve visibility in generative-engine responses by up to 40%, with effects varying by domain. URL: https://arxiv.org/html/2311.09735v3

[4] Google Search Central, "Optimizing your website for generative AI features on Google Search". Explains that Google generative AI features use retrieval-augmented generation and query fan-out, and states that foundational SEO practices remain relevant. URL: https://developers.google.com/search/docs/fundamentals/ai-optimization-guide

[5] Google Search Central, "Spam policies for Google web search". Defines cloaking and scaled content abuse, including generated pages created primarily to manipulate rankings or generative AI responses rather than help users. URL: https://developers.google.com/search/docs/essentials/spam-policies

[6] Scrunch, "Agent Experience Platform (AXP)". Describes AXP as creating a parallel AI-ready version of pages and routing AI agent traffic to a stripped-down structured version while humans keep the full experience. URL: https://scrunch.com/platform/agent-experience/

[7] Scrunch FAQ, "What is Scrunch's Agent Experience Platform (AXP) and how does it work?". States that AXP automatically converts a website into an AI-optimized format, detects AI traffic, routes agents through the CDN, and delivers AI-optimized server-rendered HTML. URL: https://scrunch.com/faqs/what-is-scrunch-agent-experience-platform-axp-and-how-does-it-work/

[8] Scrunch Help Center, "Agent Experience Platform (AXP)". Describes AXP as edge middleware that detects AI agent traffic and serves clean, structured, LLM-optimized content without changing the human site. URL: https://helpcenter.scrunchai.com/en/articles/13656392-agent-experience-platform-axp

[9] SEMrush Knowledge Base, "AI Visibility Toolkit". Describes SEMrush AI Visibility Toolkit as a tool to benchmark brand AI visibility, analyze competitors, track prompts, audit technical blockers, and report AI-search performance. URL: https://www.semrush.com/kb/1493-ai-visibility-toolkit

[10] Otterly.ai, "AI Search Monitoring Tool". Describes Otterly as an AI search monitoring and optimization platform for Google AI Overviews, AI Mode, ChatGPT, Gemini, Copilot, and Perplexity, with brand mention and citation tracking. URL: https://otterly.ai/

[11] Profound, "Answer Engine Insights". Describes Profound's tracking of front-end AI experiences, visibility scores, share of voice, sentiment, citation sources, authority, competitor rankings, and RAG-based answer analysis. URL: https://www.tryprofound.com/features/answer-engine-insights

[12] Profound homepage. Presents Profound as a full-stack marketing platform for AI search, including prompt volumes, answer-engine insights, agents, and agent analytics. URL: https://www.tryprofound.com/

[13] Cloudflare Docs, "Markdown for Agents". Explains that clients request text/markdown through the Accept header and Cloudflare converts the origin HTML to Markdown before serving it. URL: https://developers.cloudflare.com/fundamentals/reference/markdown-for-agents/

[14] Cloudflare Blog, "Introducing Markdown for Agents". Explains Cloudflare's rationale for source-side HTML-to-Markdown conversion and the use of content negotiation for AI agents. URL: https://blog.cloudflare.com/markdown-for-agents/

[15] Jeremy Howard, "The /llms.txt file". Proposes /llms.txt as a Markdown file that gives LLM-friendly background, guidance, and links for inference-time website use, while noting that processing behavior depends on the application. URL: https://llmstxt.org/

[16] Model Context Protocol specification, "Resources". Defines MCP resources as a standardized way for servers to expose context such as files, database schemas, or application-specific information to clients. URL: https://modelcontextprotocol.io/specification/2025-06-18/server/resources

[17] Model Context Protocol specification, "Tools". Defines MCP tools as server-exposed capabilities that models can invoke, with names and schema metadata. URL: https://modelcontextprotocol.io/specification/2025-06-18/server/tools

[18] AnswerShare, "Overview of GEO Corpus". Describes AnswerShare's GEO corpus as a synthesis of public research, information-retrieval literature, AI-platform behavior, industry frameworks, and AnswerShare operational testing. URL: https://answershare.com/geo-overview

[19] AnswerShare, "GEO Methodology". Defines AnswerShare's metric families: AI Opinion, Infrastructure Readiness, and Reasoning Keys, including signals such as TTFB, TTLB, query fan-out survivability, source grounding, retrieval token cost, MCP, and ai-content-index.json. URL: https://answershare.com/methodology

[20] AnswerShare homepage. Describes AnswerShare as a translation layer between websites and AI systems, including translate, measure, and optimize stages. URL: https://answershare.com/

[21] AnswerShare FAQ. Describes AnswerShare's ASQ measurement, reproducible methodology pages, per-site receipts, and Top10Lists.us proof claims. URL: https://answershare.com/faq

[22] Wikidata, "AnswerShare (Q140196284)". Identifies AnswerShare as an American website providing AI-powered answers and generative engine optimization services. URL: https://www.wikidata.org/wiki/Q140196284

[23] Wikidata, "Robert Maynard, Jr. (Q18157412)". Identifies Robert Maynard, Jr. as an American businessman and records the QID used to ground the primary investigator attribution. URL: https://www.wikidata.org/wiki/Q18157412

[24] Top10Lists.us, "AI Crawl Statistics". Publishes collection-method notes for AI-bot crawl statistics, including Cloudflare GraphQL Analytics and edge middleware logs. URL: https://www.top10lists.us/crawl-stats

[25] OpenAI Developers, "Overview of OpenAI Crawlers". Describes OpenAI crawlers and user agents, including robots.txt controls for OAI-SearchBot and GPTBot. URL: https://developers.openai.com/api/docs/bots

[26] Anthropic Support, "Does Anthropic crawl data from the web, and how can site owners block the crawler?". Describes Anthropic crawler controls and shows robots.txt examples for ClaudeBot. URL: https://support.claude.com/en/articles/8896518-does-anthropic-crawl-data-from-the-web-and-how-can-site-owners-block-the-crawler

[27] Perplexity Docs, "Perplexity Crawlers". Describes Perplexity crawlers and robots.txt controls for how sites interact with Perplexity. URL: https://docs.perplexity.ai/docs/resources/perplexity-crawlers

[28] Google Search Central, "Intro to how structured data markup works". Explains that Google uses structured data markup to understand content and supports JSON-LD, RDFa, and Microdata. URL: https://developers.google.com/search/docs/appearance/structured-data/intro-structured-data