

Chapter 1 - What SEO Is

How search exposure works, what SEO can influence and how evidence should be read

A useful webpage can exist in almost perfect silence.

It may contain the clearest answer, describe exactly the right service or offer the product someone needs. None of that guarantees that the page will be found. The people who would value it may not know it exists, while the organisation behind it may have no direct way to reach them at the moment their need arises.

Search engines help bridge that gap. They connect people who are trying to resolve something with information, organisations and services that may be able to help. SEO exists because this connection is neither automatic nor entirely **controlled** by either side.

This chapter follows that connection from beginning to end. We will see what must happen before a result can appear, what an organisation can legitimately change, where SEO ends and neighbouring practices begin, why visibility is not automatically valuable and how evidence limits what a responsible practitioner may claim.

The basic problem SEO addresses

Consider a heating company that provides emergency boiler repairs in Leeds. Its engineers are qualified, the telephone line is staffed through the night and the company has capacity to take another job. Yet when a nearby household searches for urgent help, the company is absent from the results.

The service has not suddenly become poor. The problem is that its usefulness has not become visible through search.

There are several ways this can happen. The search engine may not know that the company's emergency-service page exists. It may reach the page but fail to recognise the area the company serves. The page may describe boiler repair so vaguely that it appears no more suitable than dozens of alternatives. Important details may be missing, contradicted elsewhere or difficult to extract. The result may appear but communicate the service badly enough that the searcher passes over it.

Each failure produces the same outward complaint - "we are not getting enough business from search" - but the remedy depends on where the connection broke.

The company cannot solve the problem by announcing that it deserves to rank. From the search engine's point of view, the company must be represented by information the system can discover and interpret. From the searcher's point of view, the result must communicate a credible fit between the urgent need and the service being offered.

The company can improve those conditions. It can create a dedicated emergency-repair page, state the locations it covers, describe when help is available and organise its website so the page is easy to find. It can correct inconsistent business information and remove technical barriers that prevent search services from accessing important content.

Those changes do not compel a search engine to show the company. They make the company a clearer and more usable possibility for the search engine to consider.

That is the first boundary to understand. SEO works on a connection among a real need, a potentially useful resource and a search service that decides what to present. It can strengthen the information and access conditions on the represented organisation's side. It cannot manufacture an underlying fit that does not exist.

If the company serves only Leeds, no amount of optimisation makes it useful to someone who needs an engineer in Glasgow. If it provides routine maintenance but no emergency call-out, repeating the word "urgent" does not create that service. The result might attract more attention, but it would do so by creating a misleading connection rather than clarifying a genuine one.

We can therefore begin with a plain definition:

Search engine optimisation is the work of making suitable resources easier for public search services to find, understand and present to the people they can genuinely help.

The definition is deliberately modest. It does not promise first position, unlimited traffic or automatic commercial success. Before refining it, we need to understand why publishing a resource is only the beginning of the route into search.

Appearing in search is not one event

When the heating company publishes its emergency page, the page becomes available on the Web. It does not instantly become available as a result for every relevant search.

First, the search service must become aware of it. The service might follow a link from the company's homepage, read a submitted sitemap or encounter another page that refers to it. If nothing leads to the new page, the page can remain public while still being **unknown** to the search engine.

Awareness is not enough. The service must be able to obtain the page's contents. A server error may block access. A technical instruction may tell automated systems not to use the page. Essential text might depend on code that fails to load. In each case the company sees a working page under some conditions, while the search service receives nothing useful or receives only part of it.

The contents must then be interpreted. The service needs to distinguish an emergency boiler-repair page from a maintenance guide, a spare-parts shop or a general page about heating. It needs to recognise that the service operates in Leeds and that the page represents the same organisation described elsewhere on the site. If several addresses contain substantially the same information, it may also need to decide which representation should stand for them.


Only after those earlier conditions have been met can the page become a realistic possibility for a particular search. When someone asks for an emergency boiler engineer in Leeds, the service retrieves possible results, compares or combines them and decides what to present.

Ranking matters, but it happens late in a process that began with discovery and access. Treating every search problem as a ranking problem is like trying to improve the position of a runner who was never admitted to the race.

This explains why identical symptoms can require opposite remedies. Two pages may both receive no search traffic. The first has never been discovered. The second is known and available but does not clearly match any useful search. Rewriting the second page may improve its interpretation. Rewriting the first without creating a route by which it can be found leaves the original failure untouched.

The process continues after selection. A useful page may be shown with a title that hides its relevance. A clear result may lead to a slow page on which the telephone number is difficult to find. A technically successful visit may lead into a service process that does not answer calls.

The search part of the chain does not always end with a click. A result may give a direct answer, expose a call button, supply directions or open a booking tool. Search is still helping the person choose and reach the next step. Once the call must be answered, the booking processed or the repair performed, however, the underlying service has taken over. We will call that boundary the **functional handoff**: the point at which finding and choosing give way to carrying out the task.




A result passes through distinct stages before it can be shown. The functional handoff marks where search exposure ends and downstream action begins.

The distinction matters because organisations often try to solve downstream failures by increasing upstream visibility. If suitable customers already call but nobody answers, a higher ranking will send more people into the same broken process. Sensible diagnosis follows the route far enough to locate the actual constraint.

We can summarise the route as an **exposure chain**. A possible result must be discovered or admitted, made eligible, interpreted, retrieved for a relevant request, selected, presented and connected to whatever action the interface permits. These names do not claim that every search product implements one identical internal pipeline. They give us a practical model for asking which condition failed and which actor can affect it.

We must also keep the real pathway, the deliberate intervention and the recorded evidence separate. A person may search, see the company, telephone and book a repair. The company may have changed its title the previous day. Search reports, website analytics, call systems and customer records may each capture a different part of what happened.



A search pathway, the SEO change made within it and the evidence recorded about it are related, but they are not the same thing.

No single record automatically contains the whole event. An impression is not attention. A click is not a successfully loaded visit. A telephone event is not a completed repair. The records become useful when we know which part of the pathway each one describes.

The same separation protects us from easy causal stories. If the company changes a title on Monday and calls rise on Tuesday, we know the order of events. We do not yet know that the first event caused the second. Colder weather, a competitor's closure, an advertising campaign or a tracking change could also explain the increase.

Once the route is visible, we can ask a more precise question: which parts belong to the organisation, and which remain decisions made by the search service or the person searching?

What an SEO professional can influence

Return to the emergency page. The company controls whether the page exists, what it says and whether its claims are accurate. It can decide how the page is linked from the rest of the website, whether important information is present in the page's code and whether accidental instructions block access.

Control means that an actor can determine the state of a specific target. The company can change a heading on its own website because it governs that resource. A developer with the necessary access can deploy a corrected technical instruction because the deployment lies within that role's authority.

Influence is different. The same changes may affect how a search engine interprets or presents the page, but the search engine retains the final decision. The organisation can supply better evidence of its service area; it cannot order the engine to accept that evidence, choose the page or place it first.

The distinction continues beyond the result. The company can make its telephone number prominent. It cannot force a searcher to call. It can train staff and improve the booking process. It cannot guarantee that every enquiry becomes a completed repair.

This is not cautious wording for its own sake. It changes the quality of a recommendation.

"Make the page rank first" names a desired result without identifying an action the company controls. "Explain the emergency service and its actual coverage area clearly" names a change that can be implemented. It also permits a reasoned mechanism: clearer information may help searchers and search systems recognise when the service is relevant.

A professional recommendation goes further. It states how the change will be verified and what observation would weaken the explanation. The company might check that the correct text reached the live page, observe whether appropriate searches begin exposing it, compare the locations of resulting enquiries and watch for misleading demand from places it does not serve.

The recommendation now has a real structure. There is a defined problem, a change within someone's authority, a reason that change may affect one part of the exposure chain and evidence capable of supporting or challenging that reason. This is more useful than a confident prediction because it tells the organisation what can actually be done and learned.

We now have enough of the system to make the definition more precise without forcing it to carry the whole explanation by itself.

A precise working definition of SEO

In this academy, **SEO** means the professional work of deliberately managing how suitable resources, public entities and their representations may be discovered, understood, selected and presented through qualifying unpaid public-search routes.

This is a working convention for the curriculum, not a claim that every practitioner or company uses the term in exactly the same way. Its purpose is to keep difficult cases visible rather than letting a familiar label hide them.

The phrase **deliberately managing** matters because SEO does not always mean increasing exposure. A new service may need exposure established. A redesign may need existing exposure preserved. A technical failure may require it to be restored. Duplicate pages may need to be consolidated around a preferred version. An obsolete, private or dangerously misleading result may need its exposure limited or removed.

Direction alone therefore does not define the practice. Increasing and reducing exposure can both be sensible SEO when they deliberately manage a qualifying search route for a legitimate purpose.

The phrase **resources, public entities and their representations** prevents us from pretending that search results are always ordinary webpages. The heating company is an organisation. Its website contains resources it controls. A search service may also maintain a business profile that represents the company inside the search product. The company can provide or correct authorised information, while the search service controls the final representation and whether it is shown.

The phrase **public-search route** distinguishes the field from every system containing a search box. A retailer can optimise its internal product search. A seller can improve a listing inside a marketplace. A developer can work on discovery in an app store. These are all forms of **search optimisation**, the wider family. They may share techniques with SEO, but their eligibility rules, governing actors, commercial relationships and destinations differ.

The distinction is practical. A website owner can publish a public page without becoming a seller to the search engine. A marketplace merchant cannot make an ineligible product appear merely by publishing an independent page; participation in the marketplace helps constitute the route being optimised. An app developer depends on the store's distribution conditions. An internal-search team may control both the collection and the ranking system.

The phrase **unpaid** needs equal care, because it is often misunderstood.

Unpaid does not mean free

Organisations routinely spend money on SEO. They pay writers, developers, designers, analysts, consultants, hosting providers and software companies. They invest staff time and accept the cost and risk of changing a website. None of those expenses automatically turns the resulting search appearance into an advertisement.

The relevant question is not whether money was spent. It is what the payment purchased and who controlled the exposure.

When the heating company pays a writer to improve its service page, it purchases labour. The search service still decides whether the page will be admitted, selected or presented. When an advertiser pays the search service or its authorised seller for placement or an exposure advantage, the payment purchases something inside the exposure route itself. That route is advertising.

Some systems are less obvious. A seller may pay to join a commercial catalogue and then be ranked without paying for each individual position. The later ranking step may be unpaid, but the product became eligible only through paid

commercial participation. Calling the complete route simply "organic" would hide a dependency that materially changes how it works.

The same logic applies in reverse. A search service may pay a publisher to license information. That payment flows from the exposure controller to a source; it does not purchase placement for the publisher. A company may pay another site for a mention intended to **influence** search. That raises questions about deception, platform rules and harm, but it does not by itself mean the search engine sold the final exposure.

Where the recipient, purpose or effect of a payment is unknown, the commercial status remains unknown. Uncertainty is not a defect to be repaired with a convenient label.

The economic test therefore follows the complete route. Did payment purchase admission, placement or another advantage from the party controlling exposure? If so, the relevant chain is commercially conditioned. If not, the work may still be expensive, but the search exposure itself has not been purchased as advertising.

This definition works cleanly for a conventional webpage. Modern search screens, however, often place several different routes beside one another. To classify them sensibly, we have to stop treating the screen or brand as the unit of analysis.

The same screen can contain different kinds of search

Imagine a person planning a trip. One search screen may contain a hotel's website, a public business profile, a map result, a room offer supplied through a booking system and an advertisement. All refer to the same hotel, but they are not the same candidate and do not reach the traveller through the same route.

The hotel itself is an entity. Its website is a resource governed by the hotel. The business profile is a representation maintained inside the search service. A room offer for particular dates is a transactional candidate that may be supplied by the hotel or another seller. The advertisement is exposure purchased from the party controlling its placement.

Calling the whole screen "organic," "local," "shopping" or "AI search" collapses distinctions that determine what the hotel can change, what made each result eligible and what should be measured.

A reliable classification begins with the exact item being considered. What is the selected candidate: a webpage, a public entity, a product, an offer, a hosted profile or an application? What underlying resource or organisation does it represent? Who owns the representation? What had to happen before it became eligible? Who controls selection? Where does the person go to complete the important task?

These questions reveal several broad routes.

Search route	What makes the candidate eligible	Where the important task continues
Public Web search	An independently governed public resource or entity can be discovered or submitted without purchasing placement	An independent website, provider or public destination
Feed-mediated Web search	Structured data represents an independently available product, service or resource	The independent merchant or provider
Marketplace search	Commercial inventory participation is required and the operator controls an essential part of the transaction	Inside or materially through the marketplace relationship
App-store or platform-native search	The candidate exists or becomes eligible through distribution, hosting or profile creation in the platform	Inside the platform-controlled environment
Internal search	One publisher controls the collection being searched	Elsewhere in the same publisher's service

The marketplace distinction depends on more than the presence of a commission or a familiar brand. A route becomes materially marketplace-dependent when commercial seller participation is required and the operator controls or

performs something essential to the transaction, such as contractual offer terms, inventory allocation, checkout, payment or booking confirmation. Mere referral to an independent provider is not the same relationship.

Real products can combine these routes. A merchant's product may appear through ordinary Web search, a free feed-mediated listing, a paid advertisement and a marketplace offer on the same day. Each route should be recorded separately because its controls, economics and evidence differ.

The same discipline becomes even more important when one generated answer appears to be a single result.

AI-generated answers contain several different kinds of exposure

A generated answer is one visible composition, but the information inside it may have arrived through several paths.

The system may retrieve a webpage while preparing the response. Information from that page may help ground what is written. The interface may display the page as a citation, mention an organisation without linking to it, recommend a provider or expose an action the person can take. Licensed datasets may contribute under separate commercial agreements. Some wording may arise from patterns learned during model training rather than from a current identifiable source.

These events are related, but they are not interchangeable.

Retrieval means that a source was brought into the answer-making process for that request. **Citation** means that the interface displayed an attribution or link. **Mention** means that a name or entity appeared. **Recommendation** adds a judgement about suitability. **Action routing** provides a path by which the user or an agent can do something. A contribution from model memory cannot responsibly be converted into a specific live source merely because similar wording can be found elsewhere.

This is why "AI visibility" is too broad to serve as a complete objective. A citation tracker can count displayed citations under the conditions it observes, but it cannot automatically reveal every source used during generation. A conventional rank tracker may observe links while missing answers that influence people without producing a visit. Website analytics observes only what reaches the instrumented site; it cannot describe every direct answer, reformulation or abandoned task.

The organisation must first name the event it cares about. If it wants to know whether its pages are displayed as citations, it should measure citations. If it wants evidence that current content contributes to grounded answers, it needs evidence about retrieval and use. If it cares about recommendations, calls, bookings or agent actions, it must observe those events through the routes where they occur.

Some of this work fits the SEO convention developed here. Improving a public webpage so that a Web-grounded search service can discover, understand and cite it acts on a qualifying public-search route. Work intended to alter a closed platform's internal recommendation system or a model's parametric memory may sit elsewhere because the mechanism, actor controls and evidence are different.

Product labels will change. The durable method is to decompose the apparent answer into candidates, sources, representations, selection decisions and action routes. Once that has been done, the next question is not how much exposure can be obtained. It is what that exposure is supposed to accomplish.

Visibility is not the final objective

Suppose the heating company's page rises from position nine to position three and receives twice as many clicks. Has the SEO work succeeded?

The numbers establish that something changed under the reporting systems' definitions. They do not yet tell us whether the change was worthwhile.

More clicks may mean that households in Leeds are finding an appropriate emergency service. They may also come from people outside the coverage area, students seeking repair instructions, homeowners whose boilers need work the company does not provide or users attracted by a title that overstates availability. The same increase can contain useful demand, harmless curiosity and expensive mismatch.

Visibility is therefore an opportunity for another outcome, not the outcome's value. Its importance depends on what happens next, whom the result affects and what the organisation was trying to achieve.

Different organisations can reasonably want different things from search. A local service company may want suitable enquiries it has capacity to fulfil. A public authority may want people to find an answer without needing to telephone. A software company may want documentation to prevent avoidable support cases. A publisher may want an obsolete page to disappear because continued visibility creates confusion or harm.

The appropriate evidence changes with the objective. Rankings describe prominence under recorded conditions. Impressions show that a platform counted an eligible appearance. Clicks record a qualifying selection. Website analytics records events after arrival when measurement executes. Customer records may show that an enquiry became a completed job. Each observation can be useful, but each answers a different question.

A click does not prove that the page loaded. A visit does not prove that the person understood the offer. An enquiry does not prove that the person was suitable to serve. A sale does not prove that SEO caused the sale or that the same resources could not have created more value elsewhere.

The mistake is not using intermediate metrics. The mistake is asking them to stand in for outcomes they do not measure.

A useful objective can be stated in ordinary language before its metric is chosen: help more households inside the service area reach the correct emergency-repair route without increasing misleading enquiries or exceeding safe capacity. Rankings, impressions, clicks, calls and completed work can then be interpreted as different observations along that pathway.

This also reveals why reducing exposure can be successful. If a page offers outdated safety advice, attracts demand the company cannot fulfil or exposes information that should not be public, fewer appearances may improve the real system while making a visibility dashboard look worse.

To decide whether an intervention created value rather than merely accompanying it, we need a comparison.

Value must be judged against an alternative

Imagine that the company spends £20,000 improving a group of service pages and later attributes £80,000 of revenue to customers who arrived through unpaid search. The result sounds impressive, but the comparison is incomplete.

Some of those customers might have found the company anyway. The new pages may have shifted demand from another channel rather than creating it. The extra work may have required overtime, subcontracting or the rejection of

higher-value jobs. The benefit may continue for several years or disappear after a seasonal peak. Another use of the same people and money might have produced more value.

The missing comparison is the **counterfactual**: what would probably have happened without the intervention or under the next-best available decision?

We cannot watch the same organisation both make and not make the change under identical conditions. We estimate the difference through experiments, carefully chosen comparisons, models or bounded professional judgement. The method must match the importance of the decision and the evidence available.

This is why attributed revenue and incremental revenue are not synonyms. Attribution assigns credit for observed outcomes according to a rule or model. Incrementality asks how much of the outcome would not have occurred under the relevant alternative.

The distinction matters even where no money changes hands. A public-information page may reduce support calls, but some people would have found the answer elsewhere. A migration may show no growth and still be valuable because it prevented a large loss. Removing a harmful result may reduce traffic while improving safety and trust.

Value also depends on who benefits and who bears the cost. A page can increase revenue while misleading users. A change can reduce staff workload while making a service harder to use for people with particular access needs. A direct answer can help the searcher while weakening the publisher that funded the underlying information.

SEO metrics cannot resolve these conflicts by themselves. The organisation must state its objective, constraints, alternative, affected groups, material costs and uncertainty. A reversible wording change on a low-risk page may justify a modest evidence requirement. Advice affecting health, finance, legal rights or sensitive data requires stronger evidence, appropriate expertise and explicit authority.

The analysis has now moved beyond search mechanics. That is unavoidable: a recommendation can affect search while depending for its value on design, operations, finance, accessibility, privacy or safety. Clear responsibility begins by naming exactly who controls each decision.

Control, influence and verification are different

Complex SEO work often fails because responsibility is described too loosely. A team says that someone "owns the page" or "owns organic performance" without identifying which target that ownership covers.

Consider a proposed change to the heating company's service page. An SEO practitioner may recommend clearer wording. A writer may control the draft. A manager may decide whether the claim is commercially and legally acceptable. A developer may control deployment. The search service controls whether and how the page appears. The visitor decides whether to call. The call centre influences whether the enquiry becomes a booking.

No single actor controls the complete outcome.

It helps to separate four relationships. **Decision authority** concerns who may approve the change. **Execution control** concerns who can implement it. **Delivered-state control** concerns whether the intended version actually reached the live system. **Downstream influence** concerns effects on decisions retained by other people or systems.

Verification is separate again. A developer may correctly deploy a technical instruction without knowing whether a search service later processed it. A team may observe a page in search without confirming that every relevant user sees the same result. A customer record may confirm a completed repair while saying nothing about which earlier touchpoint caused the person to choose the company.

This separation prevents two opposite errors. A failed outcome does not prove that nobody controlled the implementation; the implementation may have matched a poor specification perfectly. A correct implementation does not prove that the desired external response occurred.

Clear responsibility therefore names an actor, a target and a relationship. Who may decide? Who can implement? Who verifies the delivered state? Who can only influence what follows? Once these answers are explicit, problems can be routed to someone capable of resolving them instead of being handed vaguely to "SEO."

The same clarity is needed when the possible harm of a decision changes.

Responsible SEO depends on the risk of the decision

The ease of changing a webpage tells us nothing about the seriousness of its consequences.

Changing a heading on a service page may take a minute. If the new wording makes a dangerous boiler fault sound safe, the simplicity of the edit does not reduce the potential harm.

Responsible SEO represents the method, evidence, uncertainty and likely outcome honestly. It respects applicable law and agreements, avoids deception and unauthorised interference, considers foreseeable harm and protects privacy, security, accessibility and information integrity where they matter. It also records who can approve the decision, what will be monitored and how the change can be corrected or reversed.

Different decisions require different levels of care. Clarifying the area served by the repair company is ordinarily a routine commercial change. Publishing emergency safety instructions is not. Search demand may show that people ask a question, but it cannot establish that an SEO practitioner is qualified to answer it. The content needs appropriate technical expertise and authorised review.

Platform rules form one part of this responsibility, not the whole of it. A tactic may comply with a search service's published policy and still be misleading, inaccessible, invasive or harmful. Conversely, a breach of platform rules is not made acceptable by a large expected commercial return.

The principle is proportionate review. Routine, reversible work should not be buried beneath needless ceremony. Material uncertainty should not be hidden behind confident language. When a credible unresolved fact could affect health, safety, rights, essential services, vulnerable groups or sensitive data, an irreversible public change may need to wait for the missing expertise or evidence.

This does not demand certainty before action. It demands that the strength of the decision, the safeguards and the authority match the consequence of being wrong.

That brings us to the final part of the chapter's central question. Search and business systems produce a great deal of data, but what does that data actually permit us to say?

Evidence determines what we are allowed to claim

An SEO report may contain rankings, impressions, clicks, sessions, calls, conversions, crawler findings, vendor scores and research results. The difficult part is not collecting another number. It is determining what each record describes and how far the resulting conclusion may travel.

A reliable claim can be built through five connected steps:

Claim > Evidence > Inference > Scope > Decision

The **claim** is the exact proposition we want to assert. The **evidence** is what the source, observation or study actually contains. The **inference** explains why that evidence bears on the claim. The **scope** states the engine, surface, population, jurisdiction, time, version and conditions in which the inference applies. The **decision** is the action justified at that evidence ceiling, including any safeguard or condition for reversal.

Suppose the company changes a page title and qualified calls rise during the following week. The records establish sequence if their timestamps and definitions are sound. They do not, by themselves, establish that the title caused the increase. Weather, competitor availability, advertising, changes in demand, tracking errors and ordinary variation remain plausible explanations.

The safe conclusion is not necessarily "we know nothing." The evidence may justify keeping the title while gathering more observations, testing similar pages or checking rival explanations. Precision narrows the claim to what the evidence supports and still permits a useful decision.

Source authority and methodological fitness answer different questions. A search platform is authoritative about the rule it publishes, but the rule alone does not prove the effect of one implementation. A vendor is authoritative about how it defines its metric, but not automatically about whether the metric represents search-engine quality. A controlled experiment can estimate an effect without revealing the proprietary mechanism that produced it.

Several familiar examples show how this discipline changes ordinary SEO language.

An automated accessibility check is not complete conformance

The Web Content Accessibility Guidelines distinguish the success criteria that define conformance from techniques that can help satisfy them.[\[1\]\[2\]](#) Passing one automated or technique-based check can support the statement that the named problem was not detected under the tested conditions. It does not prove that the complete page conforms to WCAG.

The useful correction is not to dismiss the test. It is to state its result accurately, retain it as evidence and complete the wider evaluation needed for the broader claim.

A Search Console click is not automatically a human visit

Google Search Console counts a click when a qualifying interaction in Google Search meets its reporting rules.[\[3\]](#) That event is not identical to a unique person, a fully loaded webpage, an analytics session or a completed task.

This is why Search Console clicks and analytics sessions need not match. The systems observe different events, use different collection rules and can lose or filter records at different points. Forcing the totals to agree would erase the information needed to explain the difference.

Current official documentation can conflict with itself

Official documentation is usually the strongest evidence of what a platform declares about its own product. It can still contain pages written at different times, updated by different teams or describing different product states.

In 2026, Google published information about generative-AI performance reporting while another current help surface described parts of that reporting differently.[\[4\]\[5\]](#) Choosing whichever statement supports a preferred conclusion would manufacture certainty.

The responsible response is to retain both records, preserve their dates, identify the exact conflict and narrow the current claim. Direct product observation may help, but it should be recorded as observation rather than silently replacing the documentation.

A vendor metric is not a search-engine score

Ahrefs defines Domain Rating as a measure based on the strength of a target's backlink profile relative to other sites in its database and method.[\[6\]](#) Ahrefs is authoritative about that definition.

The definition does not make Domain Rating a Google score, a universal measure of website quality or proof that a page will rank. Those are additional propositions requiring additional evidence.

Domain Rating can still support consistent comparisons inside the tool. The safe use names the vendor, the metric and the analytical purpose instead of borrowing authority from a search engine that did not produce it.

An AI search-volume field is not necessarily a raw prompt count

DataForSEO uses similar AI search-volume wording in products whose underlying measures differ. One endpoint documents modelled keyword data, while LLM Mentions materials describe another platform-specific measure constructed from external search-demand signals.[\[7\]](#)[\[8\]](#)

The familiar label does not reveal what was counted. The endpoint, platform, definition, method, date and unit are needed before the value can be interpreted.

The lesson extends beyond one supplier. A column name is not a measurement method. Two fields with the same label may represent different estimates and support different decisions.

Attribution credit is not automatically incremental effect

Google Analytics attribution distributes credit for observed conversions across touchpoints under a defined attribution method.[\[9\]](#) Conversion Lift and marketing-mix methods address a different question: what outcome would not have occurred under the relevant alternative.[\[10\]](#)[\[11\]](#)

An attributed conversion may be real and useful without being incremental. Attribution asks how credit is assigned within an observed path. Incrementality asks what difference an intervention caused. Confusing them can make several channels appear to have independently created the same outcome.

These examples do not require timid writing. Strong evidence can support a strong claim. Limited evidence supports a narrower one. The discipline is to expose the reasoning so that another person can inspect it rather than being asked to trust confidence, reputation or a proprietary score.

Evidence also changes with time, which means a once-careful claim can become careless when it is reused.

Evidence must be interpreted in time

Search products, interfaces, documentation, datasets and vendor methods change. A statement can be well supported and still stop describing the current system.

Important evidence therefore needs more than a source title. We may need to know when the behaviour occurred, when the data was produced, when the source was published or updated, when it became available to the decision-maker, when it was retrieved and when the claim was last verified.

Those dates answer different questions. A document published today may describe an earlier event. A study released later may reconstruct what happened, but it was not available to people making the original decision. A vendor page without visible version history may explain the current metric while leaving previous calculations uncertain.

The record should also state what would trigger another review. A platform announcement, endpoint change, revised methodology or contradictory observation may make the existing conclusion unsafe to reuse.

Three evidence states deserve plain names. **Unknown** means that the available evidence does not resolve the question. **Conflicted** means that credible evidence supports incompatible accounts that have not been reconciled. **Stale** means that the evidence may no longer represent the current system.

These states are not embarrassing gaps to hide inside a confidence score. They tell the next decision-maker what remains uncertain, why it matters and what kind of work could reduce it.

This distinction is especially important for an automated system. A model can produce fluent recommendations from poorly defined inputs. Fluency does not restore the meaning that was lost before the model received them.

Why this matters for Valinor

Valinor may eventually collect thousands of findings from crawlers, search services, analytics platforms, vendors and business systems. Producing the findings is not the hardest part. The harder task is deciding what they mean and what, if anything, should change.

Imagine that the system reports a low click-through rate for an important page. A crude decision layer could recommend rewriting the title because low click-through rate is commonly associated with weak presentation.

An expert would first ask what the metric actually contains. Which impressions and clicks were included? Which result types appeared? Did the mix of searches change? Was a website visit necessary for the page to fulfil its purpose? Could the title already be accurate while another result offers something more suitable? Would a more enticing title create a promise the organisation cannot keep?

The correct response might be a title change. It might instead be a content correction, a better service offer, a design improvement, a controlled test, further evidence or no change at all.

A future decision system must preserve the information that separates those outcomes. It needs the observation, event definition, source, scope, relevant exposure stage, competing explanations, objective, affected stakeholders, risks, proposed action and condition for reversal. A highly capable model can reason from that material in depth. A smaller model can also perform well when retrieval supplies a package in which the important distinctions have already been made explicit.

This is why the curriculum produces both human teaching and structured chapter knowledge. The human chapter develops judgement. The structured package preserves claims, evidence, uncertainty and qualifications for later machine use without pretending that the final graph, retrieval system or decision architecture is already known.

The Valinor commercial website and a neutral experimental site also answer different questions. The commercial site reveals whether a recommendation survives a real offer, audience, design system and operational constraint. The experimental site makes selected mechanisms easier to isolate. A controlled result may fail to transfer to a real business; a commercial result may contain too many simultaneous changes to explain causally.

Used together, the two environments can improve understanding. Neither removes the need for explicit reasoning. An observation becomes a recommendation only after the system has identified what happened, why a particular intervention might help, what else could explain the evidence and what would cause the decision to be reversed.

The same discipline now gives us a compact way to carry Chapter 1 into the rest of the curriculum.

Where the curriculum goes next

Chapter 1 has established the structure into which later technical knowledge will fit. It has not attempted to teach every mechanism named along the way.

Chapter 2 examines how human situations become searches and why a query reveals only part of a person's need. Chapter 3 develops query interpretation, inferred intent, reformulation and journeys.

Chapters 4 and 5 explain the Web, URLs, documents, links and rendering. Chapters 6 and 7 follow search services through crawling, indexing, canonicalisation, retrieval and ranking. Chapter 8 examines result types, vertical search and AI-mediated discovery.

Chapter 9 maps the SEO profession and its neighbouring disciplines. Chapters 10 to 12 address measurement, customer tracking, privacy, attribution and experimentation. Chapters 13 and 14 bring the foundation together through diagnosis, strategy, implementation and governance.

Later volumes develop technical SEO, content and information architecture, authority and digital PR, local and ecommerce search, international and enterprise work, AI-mediated discovery, analytics, agency operations, engine engineering and the detailed study of tools, vendors and APIs.

The separation is deliberate. A title rewrite, internal link, structured-data field, backlink campaign or tracking event should not be learned as an isolated trick. Later chapters will explain how those interventions work. The questions below determine whether any proposed intervention has been understood well enough to use responsibly.

Four questions every SEO recommendation must answer

Return once more to the heating company. "Rank first for emergency boiler repair" is not yet a professional recommendation. It names a desired outcome while leaving the action, mechanism, authority and evidence undefined.

The first question is: **what exactly are we changing?**

The answer should identify a resource, representation, configuration or process that someone has the authority and ability to alter. In this case it might be the wording and structure of the emergency-service page, the links that make it discoverable or inaccurate information in an authorised public profile.

The second question is: **why should that change affect search?**

The answer should locate the relevant part of the exposure chain. Clearer service and location information may improve interpretation. A sensible internal link may improve discovery. Removing an accidental blocking instruction may restore access. A slogan, correlation or tool score is not a mechanism.

The third question is: **who controls each important decision?**

The answer should separate approval, implementation and verification inside the organisation from selection by the search service, action by the searcher and fulfilment by the operational team. This prevents a recommendation from promising an outcome that no one involved can command.

The fourth question is: **what evidence would show whether the explanation is right?**

The answer should name the observation, its definition, its scope and its limits. It should also state what result would cause the organisation to reconsider the change. Appropriate impressions, qualified calls from the service area and

fewer unsuitable enquiries may support the explanation. Misleading demand, unchanged discovery or failure at another stage may point elsewhere.

Together, the questions turn an aspiration into a decision that can be understood, implemented, examined and reversed. They do not replace specialist expertise. They give later expertise a common structure.

SEO is therefore not a collection of tricks for pleasing a ranking system. It is disciplined work on the conditions through which suitable resources and representations can participate in public search. It operates with limited control, serves objectives beyond visibility and earns trust only when its claims remain within the evidence.

Source notes

1. W3C, Web Content Accessibility Guidelines (WCAG) 2.2, Recommendation dated 12 December 2024, [Open source](#).
2. W3C Web Accessibility Initiative, Understanding Techniques for WCAG 2.2 Success Criteria and Understanding Conformance, updated 26 July 2026, [Open source](#) and [Open source](#).
3. Google Search Console Help, What are impressions, position and clicks? and related Search Console data documentation, [Open source](#).
4. Google Search Central, Introducing Search Generative AI performance reports in Search Console, published 3 June 2026, [Open source](#).
5. Google Search Console Help, Generative AI performance report (Search), status observed 31 August 2026, [Open source](#).
6. Ahrefs, What is Domain Rating (DR)?, updated 31 October 2025, [Open source](#).
7. DataForSEO, Live Keyword Search Volume - AI Keyword Data, endpoint POST /v3/ai_optimization/ai_keyword_data/keywords_search_volume/live, [Open source](#).
8. DataForSEO, How the AI search volume metric works in LLM Mentions and LLM Mentions Target Metrics, [Open source](#) and [Open source](#).
9. Google Analytics Help, Get started with attribution, [Open source](#).
10. Google Ads Help, Understand your Conversion Lift based on users measurement data, [Open source](#).
11. Google Meridian, Incremental Outcome, ROI, mROI and Response Curves, updated 15 May 2026, [Open source](#).