



W H I T E P A P E R

Intent Driven **Search & Keyword** Optimization

A practical framework for intent-driven query lifecycle governance
and closed-loop optimization in paid search campaigns



Table of Contents

01 Problem Statement

02 Landscape Review - Existing Tools

03 The Capability Gap

04 Intent Classification - The Core Layer

05 System Architecture

06 Decision Framework & Actions

07 Business Impact

08 Pilot Approach & Next Steps

09 Implementation Prerequisites

01. PROBLEM STATEMENT

After working across paid search accounts of varying scale, one structural inefficiency keeps surfacing. The platforms (**Google Ads, Microsoft Ads, Optmyzr, Adalysis**) are capable, and others surface rich query data and automate bids effectively. But when a new search query enters the account, the decision of what to do with it still requires a human to review, judge, and act.

Promote it? Negate it? Monitor it? That judgement happens through **Search Query Report (SQR) analysis**, manually performed by campaign analysts on a weekly review cadence at best. **Wasteful spend accumulates quietly between cycles and high-intent terms go unnoticed in the query report.** This is a solvable problem. The gap is not in the data or the tooling - it is in the **last mile decisioning**. No platform today automates the semantic decision of what a search query actually means and what to do about it.



The Current Reality

In most PPC programmes, search query reviews happen weekly or bi-weekly. A practitioner downloads a search term report, filters by spend or clicks, makes a judgement about each query, and either promotes it to a keyword, adds it as a negative, or leaves it for another cycle. This approach has three structural problems:

It is reactive.

Inefficient queries continue to spend until they are caught in the next manual review cycle. A query that generates clicks but no conversions may run undetected for weeks.

It is inconsistent.

Decisions about what to promote, negate, or monitor are made differently by different team members, at different times, with different thresholds. There is no enforced decision framework.

It does not scale.

As account complexity grows - more campaigns, more match types, more geographies - the volume of queries requiring review outpaces the capacity of manual processes.

The core problem is not a lack of data. Platforms like Google Ads and Microsoft Ads surface extensive search term data. The problem is the absence of a consistent, automated system for interpreting the data and acting on it at speed.

7–14
days

Typical lag from a wasteful query first appearing in an account to a human reviewer negating it – for teams on weekly cycles. Monthly reviewers commonly see 15–30 days. Budget drains throughout this window.

Source: Improvado PPC Analysis Guide, 2026; Negator.io, Negative Keyword Velocity research, 2025

10–20%

Of total PPC spend in broad and phrase match campaigns is typically tied to irrelevant or low-intent queries that bypass initial negative keyword filters – spend that intent-level classification would contain earlier.

Source: Improvado PPC Analysis Guide, 2026; corroborated by Seer Interactive search term analysis

Weekly
cycles

The review cadence most PPC teams operate on – meaning new wasteful queries can run for 7 or more days before being caught, and new high-intent terms can go unidentified for the same window.

Based on standard PPC management practice; the proposed system replaces periodic cycles with continuous query processing

What This Creates

The consequence is measurable waste. Budget is consumed by queries that will never convert. High-performing new search terms are slow to be identified and promoted. Negative keyword lists grow stale. And the effort required to keep up with query management crowds out higher-value strategic work. **The gap is not primarily about tool capability – it is about the absence of a decision framework that operates continuously and consistently.**

Where the Domain Expertise Sits

Building this well requires more than technical capability. It requires a deep understanding of how paid search campaigns behave – query patterns, match type dynamics, bid interactions, account structure, and how automation errors compound. The framework documented here is grounded in that domain knowledge.

Search intent mechanics

- Understanding the spectrum from navigational to transactional – and the grey areas that rule-based systems routinely misclassify
- Recognizing how query construction signals intent beyond just keyword matching

Campaign architecture

- How query decisions interact with match types, ad group structure, and bid strategy
- Optimization actions that ignore campaign architecture cause downstream regression

Threshold calibration

- Knowing which performance signals justify automated action, and at what confidence levels
- The difference between a useful system and a noisy one is where the thresholds are set

Feedback loop design

- Structuring the closed-loop so outcome signals flow back to improve decisions – not just report on them after the fact
- Ensuring the system learns from its own actions over time

02. LANDSCAPE REVIEW – EXISTING TOOLS

This framework was developed following structured research across the leading PPC optimization platforms and native advertising systems. The research aimed to understand what existing tools offer, where they perform well, and where gaps remain.

Tools Evaluated

Tool	Primary Strength	Approach to Query Management
Optmyzr	Rule-based automation, scripts, reporting	Rule triggers on performance thresholds; no semantic query understanding
Adalysis	Search term analysis, negative keyword management	Performance-based filtering; analyst reviews and approves changes
Adspert	Algorithmic bid optimization	Focused on bid adjustments; limited query-level governance
Opteo	Suggestions and alerts for campaign health	Surfaces recommendations; requires manual approval and execution
SEMrush	Keyword research, competitive intelligence	Research and discovery tool; limited live campaign integration
Google Ads	Native automation, Smart Bidding, Performance Max	Platform-level automation; limited and non-transparent query intent grouping; no actionable cross-account framework
Microsoft Ads	Native campaign management	Similar to Google; rule-based automation; manual query review
SA360	Cross-engine campaign management at scale	Centralized management; workflow automation; no intent layer

What the Research Found

Across all tools evaluated, a consistent pattern emerged: optimization is driven by performance metrics (clicks, conversions, cost) and rule-based thresholds, not by semantic understanding of what a query means or what a user intended when they typed it.

Tools are good at surfacing queries that are already performing badly. They are not designed to classify queries by intent before performance data accumulates, or to enforce a consistent decision framework about what to do with each category of query.

03. THE CAPABILITY GAP

We identified three specific capability gaps that existing tools do not address. These gaps are what the proposed system is designed to fill.



Gap 1

No Intent-Level Query Understanding

Existing tools evaluate queries based on what they have done (clicks, spend, conversions). They do not evaluate queries based on what they mean. A query like "how to fix a pipe" and a query like "emergency plumber near me" may both generate clicks on a plumbing company's ads - but they signal very different user intents and warrant different responses. Current tools treat them the same until performance data separates them.

Gap 2

No Consistent Decision Framework

When a query is flagged for review, the decision about what to do with it is made by a human, in the moment, based on their individual judgement. There is no enforced framework that says: if a query is classified as informational and has spent above a threshold with no conversions, add it as a negative. Rules can approximate this, but they operate on metrics, not meaning - and they must be configured and maintained manually per campaign.

Gap 3

No Closed-Loop Feedback

Query optimization in existing tools is largely one-directional. A query is reviewed; an action is taken, and the outcome of that action is not systematically fed back to improve future decisions. There is no mechanism by which the system learns that certain query patterns reliably lead to waste and proactively applies that learning to new queries matching the same pattern.

Existing Tools Provide

- Rule-based campaign automation
- Search term performance reporting
- Bid optimization algorithms
- Negative keyword suggestions
- Keyword discovery and research

Campaign architecture

- Semantic understanding of query intent
- Consistent automated decision framework
- Proactive classification before spend accumulates
- Closed-loop learning from optimization outcomes
- Cross-account query governance at scale

04. INTENT CLASSIFICATION - THE CORE LAYER

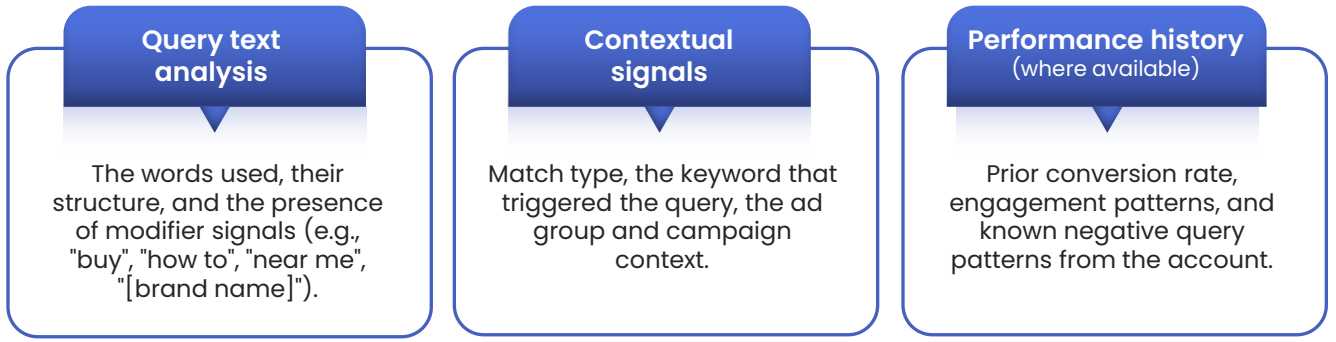
Intent classification is the differentiating capability at the heart of this system. It is the layer that existing tools do not have, and it is what makes the rest of the framework possible.

The premise is straightforward: every search query encodes something about what the user is trying to do. Classifying that intent - reliably and consistently, at scale, is what enables the system to make better decisions faster than a human review process can.

What Intent Classification Does

The classifier takes a raw search query as input and assigns it to one of a defined set of intent categories. This is not keyword matching - it is semantic analysis of what the query means in context. The same word or phrase can belong to different intent categories depending on construction, modifier terms, and surrounding signals.

Classification is performed using a combination of:



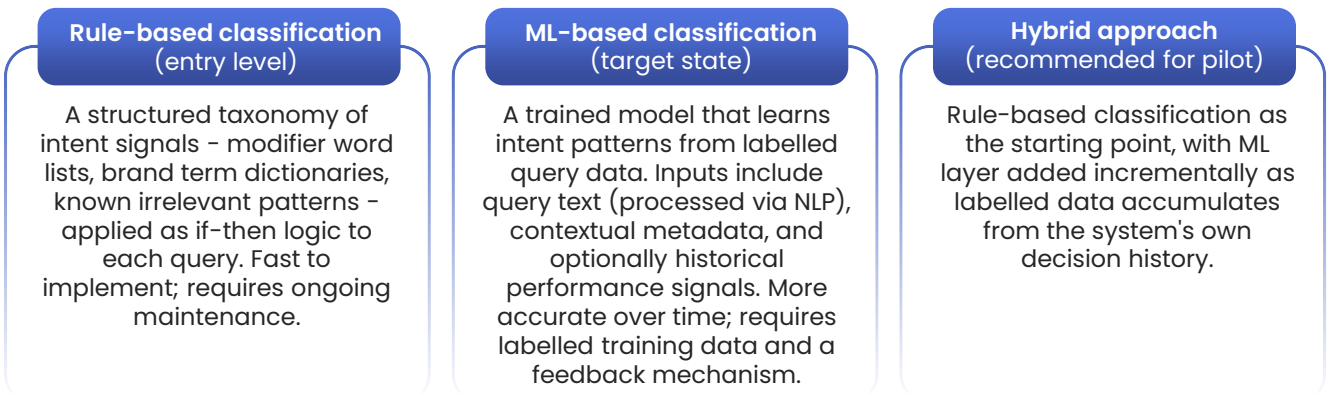
Intent Categories

The system classifies queries into the following categories. Each category maps directly to a defined action, which is described in Section 06.

Intent Category	What It Signals	Example Queries	Default Action
Transactional	User is ready to act - purchase, book, enquire	"buy running shoes online", "emergency plumber quote"	Auto-promote to keyword
Informational	User is researching - not yet in decision mode	"how to choose running shoes", "pipe types explained"	Monitor; promote if it converts
Navigational	User is looking for a specific brand or site	"Nike store near me", "[Competitor] login"	Route to strategy review
Competitor	Query references a competitor brand or product	"Adidas running shoes sale", "[Competitor] prices"	Filter or flag per account strategy
Irrelevant	Query is off-topic - unrelated to the advertiser's offer	"free shoe lace patterns", "tanker room film"	Auto-add as negative keyword
Ambiguous	Intent is unclear - insufficient signals to classify	"running help", "heating problem"	Monitor; await performance data

How the Classifier Works

At a technical level, intent classification for search queries can be implemented at varying degrees of sophistication, depending on available data and tooling. The approach can range from structured rule sets to ML-based models:



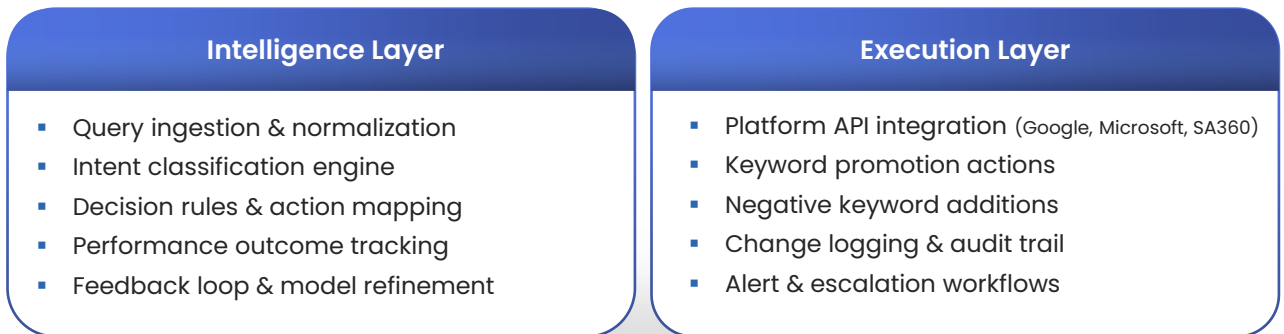
Intent classification is not a one-time process. As new queries enter the account, the classifier runs continuously.

As the system takes actions and observes outcomes, those outcomes feed back to improve classification accuracy over time. This is what makes it a closed-loop system.

05. SYSTEM ARCHITECTURE

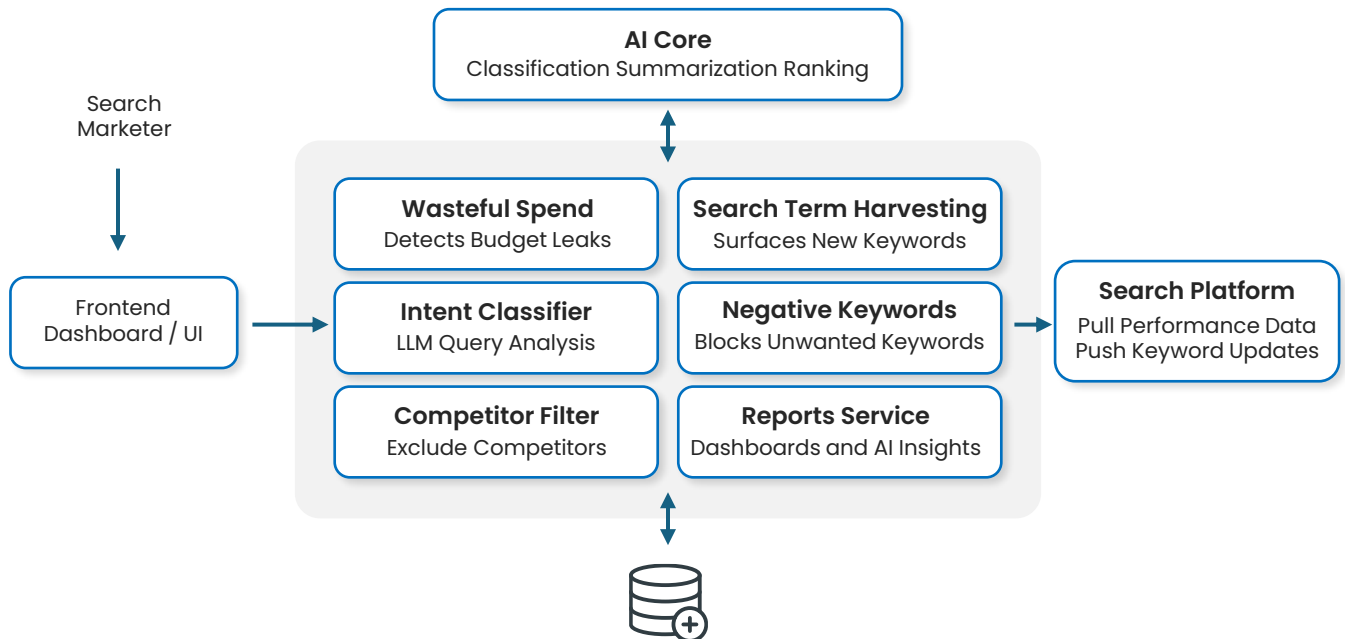
The system operates as a layer above existing ad platforms - it does not replace *Google Ads*, *Microsoft Ads*, or *SA360*, but sits between the platform's query data and the decisions that govern it. The architecture has two distinct layers: an intelligence layer that classifies and decides, and an execution layer that acts and logs.

Intelligence Layer vs. Execution Layer



The diagram below illustrates how these layers connect in practice. The LLM Core sits above all decision services, providing semantic classification that feeds into each functional module. Data flows bi-directionally, the Search Platform pushes query data in, and the system pushes keyword actions back out.

Keyword Intent Classifier - System Architecture

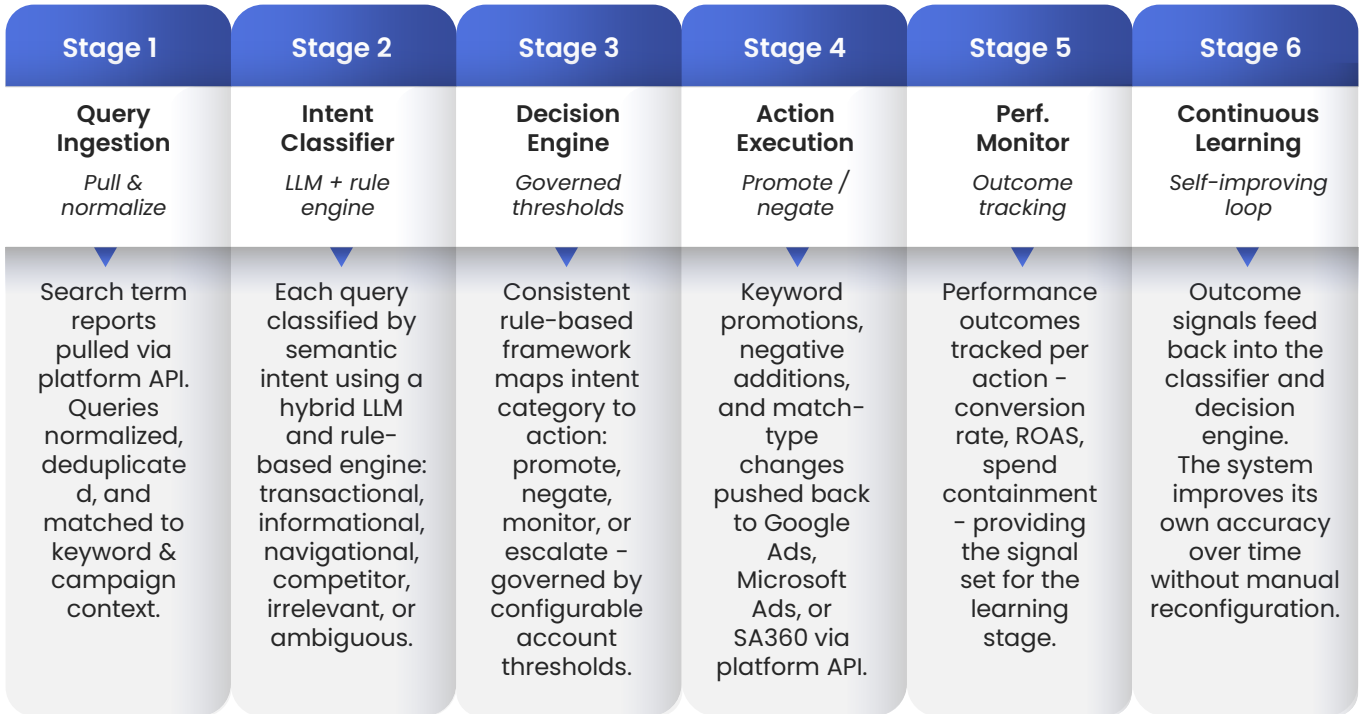


The LLM Core orchestrates classification across all services; the Datastore provides the feedback loop that enables continuous learning.

Each of the six stages described below maps to one or more of these components. The Frontend Dashboard gives practitioners visibility into decisions being made, the Datastore is what transforms the system from a one-directional automation tool into a closed-loop learning system.

End-to-End Flow

The system operates as a continuous cycle across six stages:



The continuous optimization loop - performance outcomes from Stage 5 feed back into the classifier and decision engine at Stage 2, making each cycle more accurate than the last.

Stage Detail

Stage 1

Query Ingestion

Search term reports are pulled from the ad platform on a scheduled basis (typically daily). Raw query data is normalized: duplicates removed, queries matched to their triggering keyword and campaign context, and performance metrics attached where available.

Stage 2

Intent Classifier

Each query is passed through a hybrid classification engine combining an LLM for semantic understanding with a rule-based layer for speed and precision. The classifier assigns a category (transactional, informational, navigational, competitor, irrelevant, ambiguous) based on query text and contextual signals. Confidence scores are generated; low-confidence classifications are routed to the monitor queue rather than actioned immediately.

Stage 3

Decision Engine

The decision engine applies the action framework (see Section 06) to each classified query. It determines whether to promote, negate, monitor, or escalate for human review, based on intent category, configurable performance thresholds, and account-level governance rules.

Stage 4

Action Execution

Approved actions are pushed back to the ad platform via API. Keyword promotions, negative keyword additions, and match type adjustments are applied programmatically. All actions are logged with timestamp, intent classification, and decision rationale.

Stage 5

Performance Monitor

The system tracks the performance outcome of each actioned query. For promoted keywords, it monitors conversion rate and ROAS. For negated queries, it confirms that spend on those patterns has ceased. These outcomes form the input to the learning stage.

Stage 6

Continuous Learning

Outcome signals from Stage 5 are fed back into the classifier and decision engine. The system refines its intent models and threshold calibrations based on observed results – progressively reducing misclassification and improving action quality without requiring manual reconfiguration.

06. DECISION FRAMEWORK & ACTIONS

The decision engine applies a consistent set of rules to each classified query. These rules encode what a skilled practitioner would do – but applied automatically, at scale, and without variation. The framework below defines the default action for each intent category, with thresholds that can be adjusted per account.

Intent Category	Default Action	Conditions / Overrides	Execution
Transactional	Promote to keyword	Minimum impression threshold met; not already a keyword	Add as exact or phrase match keyword
Informational	Add as negative keyword	Campaign type is conversion-focused; query has spent without converting	Add to negative keyword list at ad group level
Navigational	Add as negative keyword	Query contains a brand name that is not the advertiser's own	Add to negative keyword list at campaign level
Competitor	Flag for review or negate	If competitor bidding is not a stated strategy: negate. If it is: monitor.	Configurable per account strategy
Irrelevant	Add as negative keyword	No override; immediate action	Add to shared negative list
Ambiguous	Monitor	Re-evaluate after defined impression or spend threshold is reached	No action; added to monitoring queue

Four Applications of the Framework

The decision framework supports four practical use cases that represent the most common query management tasks in paid search campaigns:

Negative keyword automation

- Irrelevant queries identified and negated at first encounter - not at the next review cycle
- Applied at the correct scope: keyword, ad group, or campaign level
- No missed exclusions due to analyst oversight or review cadence

Search term harvesting

- High-intent queries meeting configurable thresholds auto-promoted to keywords
- No missed long-tail opportunities - the system catches what manual review skips
- Promotion scope and match type governed by the decision engine, not individual judgement

Competitor query filtering

- Competitor-branded queries handled consistently per your strategy - not variably depending on the reviewer
- Strategy-level rules (negate vs. monitor) enforced uniformly across every campaign

Informational query routing

- Upper-funnel queries monitored over time rather than immediately negated or ignored
- Promoted if they convert; negated if they keep burning spend - handled consistently, not case by case

Governance Rules

The decision engine operates within a set of account-level governance rules that define boundaries for automated action. These rules ensure that automated decisions remain aligned with campaign strategy:



Action scope

Define which campaigns and ad groups are within scope for automated action versus human review.



Spend thresholds

Set minimum spend or impression levels before a query can be negated, to avoid premature exclusion of low-volume queries.



Promotion criteria

Define the conditions under which a query can be promoted to a keyword (e.g., must have generated at least one conversion, or must exceed a minimum CTR threshold).



Escalation triggers

Queries that exceed a defined spend level without a clear classification are escalated to a human reviewer rather than actioned automatically.



Audit log

All actions taken by the system are logged with full decision context - intent classification, confidence score, threshold applied, and outcome. This provides transparency and supports governance review.

07. BUSINESS IMPACT

Teams that move to automated query governance stop spending analyst time on routine triage. The optimization cycle that used to run weekly runs continuously. Note that the scale of impact will vary by account – outcomes below are directional based on research observations.

Faster Spend Containment	Broader Keyword Coverage	Consistent Optimization Logic	Analyst Time Redirected
<p>Wasteful queries blocked before the next impression – not after the next review cycle.</p> <p>The 7–14 day lag between a bad query appearing and a human negating it is eliminated.</p>	<p>Long-tail high-intent terms captured systematically, not caught opportunistically.</p> <p>The system promotes qualifying queries without waiting for a practitioner to spot them in a report.</p>	<p>Decisions apply uniformly across every campaign, every account, every day – not based on who did the review or when.</p> <p>The same threshold, the same framework, applied at scale.</p>	<p>Teams move from routine query to strategic work – the decisions that actually need human judgement.</p> <p>High-confidence actions are automated; edge cases are escalated.</p>
<p>Continuous</p> <p><i>optimization – not weekly</i></p>	<p>Systematic</p> <p><i>discovery – not manual triage</i></p>	<p>Zero drift</p> <p><i>across campaigns and reviewers</i></p>	<p>Human effort</p> <p><i>reserved for what matters</i></p>

Additional Operational Benefits

<p>Full audit trail</p> <p>Every query decision logged with intent classification, confidence score, and rationale – supporting governance review and accountability.</p>	<p>Self-improving accuracy</p> <p>Outcome signals feed back to the classifier continuously – the system gets more accurate over time without manual reconfiguration.</p>
<p>Scales with account complexity</p> <p>As campaigns grow – more match types, more geographies, more ad groups – query volume is handled without proportional increases in team capacity.</p>	<p>Strategy-aligned execution</p> <p>Competitor and navigational handling configured per account strategy – not variably applied depending on the reviewer or review cadence.</p>

08. PILOT APPROACH & NEXT STEPS

Rather than describing the framework in abstract terms, this section outlines a concrete pilot approach. – The minimum viable implementation that would allow the system to be tested, validated, and refined in a real account environment.

Pilot Scope

A pilot should be scoped to a single advertiser account, with a defined subset of campaigns. Recommended pilot criteria:

Account type

An account with meaningful search query volume - sufficient data to evaluate classification accuracy within a short observation window.

Campaign scope

Start with broad and phrase match campaigns in one or two verticals. These generate the highest query diversity and are where the classification layer adds the most value.

Duration

A minimum of four to six weeks is recommended to allow query volume to accumulate and to observe the performance impact of actions taken.

Pilot Phases

Observe (Weeks 1-2)

Ingest query data, run the classifier, and generate intent labels - but take no automated actions. Review the classification output manually to validate accuracy and calibrate the rule set.

Recommend (Weeks 3-4)

Enable the decision engine to generate recommended actions (promote, negate, monitor) for each query. Present recommendations to a human reviewer who approves or rejects each one. Track approval rate and correction patterns.

Automate (Weeks 5-6)

Enable automated execution for high-confidence, low-risk actions (irrelevant query negation, clear transactional keyword promotions). Maintain human review for ambiguous and competitor categories. Monitor performance impacts.

Evaluate and Refine

At the end of the pilot, review classification accuracy, action approval rates, performance outcomes, and spend impact. Use findings to refine the decision framework and set parameters for broader rollout.

Decision Criteria for Progression

The decision to progress from pilot to broader implementation should be based on:

Classification accuracy

Intent labels reviewed by a practitioner should be correct in the substantial majority of cases. A defined accuracy threshold (agreed at pilot design) should be met before automation is expanded.

Action approval rate

If the human reviewer is approving the majority of recommended actions with few corrections, that indicates the decision framework is well-calibrated for the account.

Performance outcome

Actions taken during the pilot should demonstrably improve campaign efficiency - reduced spend on non-converting queries, increased share from promoted keywords - within the observation window.

The pilot is designed to generate evidence, not just demonstrate capability. The output should be a data-backed recommendation on whether and how to proceed - not a generic endorsement of the framework.

09. IMPLEMENTATION PREREQUISITES

Full implementation of the system requires access to a set of technical and operational resources. These are framed here as prerequisites for planning purposes - they define what needs to be in place before a pilot can begin, not as blockers to the framework's validity.

Technical Prerequisites

Ad Platform API Access

The system requires API-level access to pull search query data and push keyword actions. For Google Ads, this requires a Developer Token with at least Basic access level. For Microsoft Ads and SA360, equivalent API credentials are needed. These should be scoped and requested as part of pilot preparation.

Classification Infrastructure

The classifier and decision engine need access to account-level performance data - search term reports, keyword lists, negative keyword lists, and campaign structure. This data should be available in read/write format for the accounts included in the pilot.

Ad Platform API Access

A compute environment capable of running the classification pipeline on a scheduled basis. For a rule-based pilot, this is lightweight. For an ML-based system, model training and inference environment is required.

Logging and Audit Storage

A structured data store for capturing all classification decisions, actions taken, and outcome metrics. This is essential for the feedback loop and for governance review.

Operational Prerequisites

Nominated Account

A specific advertiser account agreed for the pilot, with appropriate permissions and stakeholder sign-off.

Governance Review Process

A defined process for human review of automated actions during the recommend phase. This should include a named reviewer, a review cadence, and a mechanism for capturing approvals and rejections.

Decision Framework Calibration

Account-specific parameters for the decision engine - spend thresholds, escalation rules, and scope boundaries - agreed with the account team before the pilot begins.