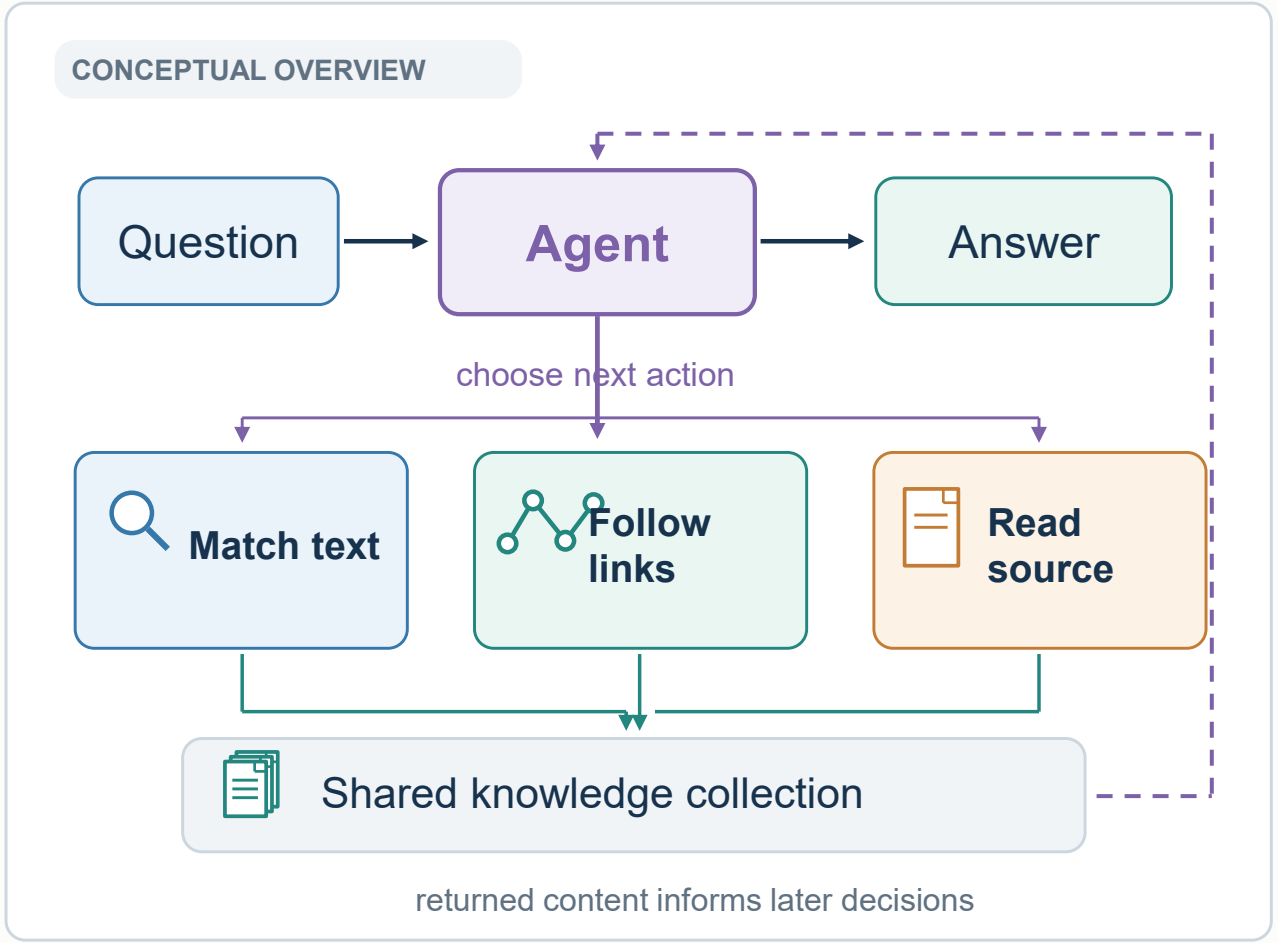
Exploring Agentic RAG

How Can We Find Sufficient Evidence at Lower Cost?



LM-RtA

Team 5

Explore the process before committing to an architecture.

Retrieval offers different ways to access knowledge

Matching content and organizing knowledge are related, but different, design choices.

(a) Sparse matching

Query- who is the director of Odessey

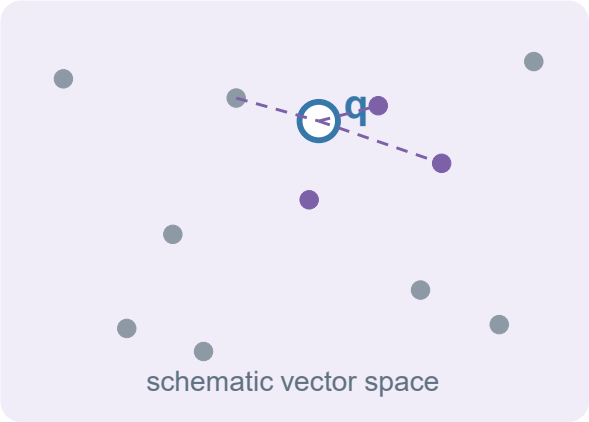
D1 Odessey is written by Homer

D2 Nolan directed a new movie

D3 Orion launch notes

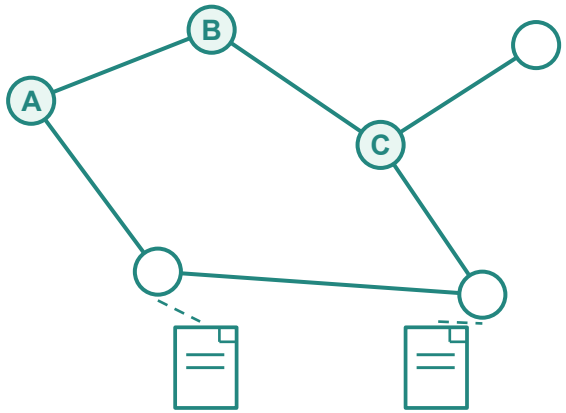
Match lexical features against an index

(b) Dense matching



Match learned query and document representations

(c) Graph-based access



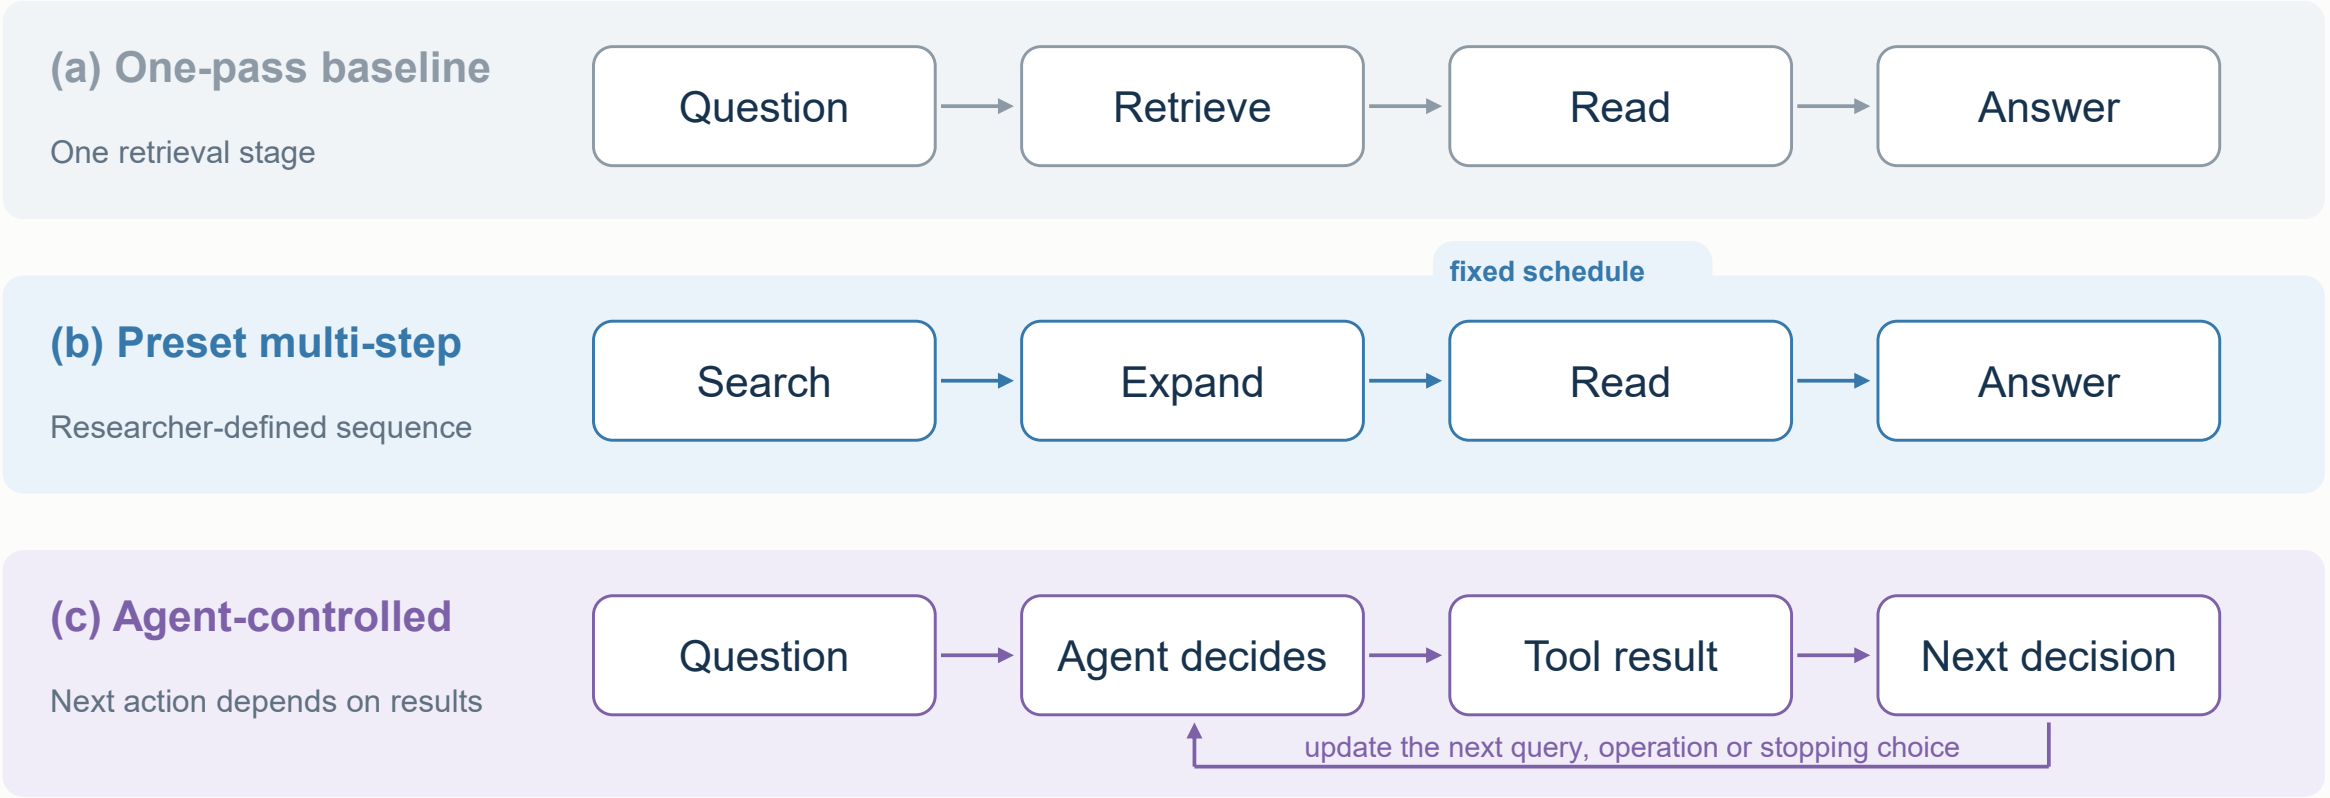
Use links, neighborhoods or graph-derived summaries

These mechanisms can be combined; a graph can also use sparse or dense search.

Our interest: what happens when an agent repeatedly selects and combines these operations?

The key distinction is who controls the next step

Repeated retrieval is not automatically agent-controlled retrieval.

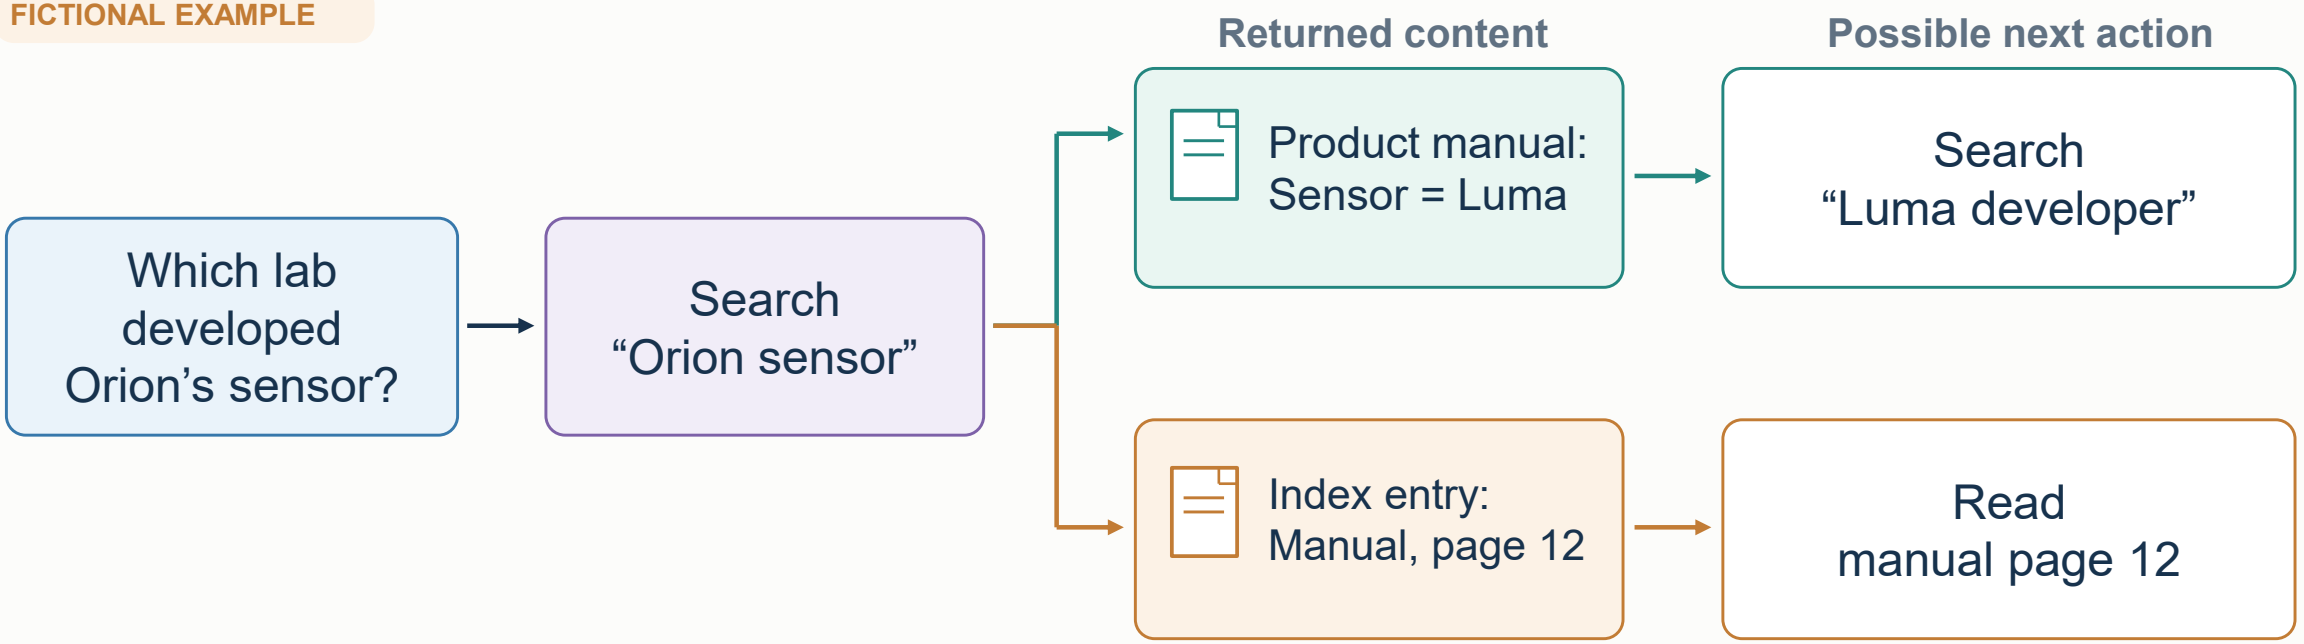


Study adaptation to returned information, not simply the number of retrieval steps.

A retrieved result can change what happens next

The same initial search can lead to different follow-up actions.

FICTIONAL EXAMPLE



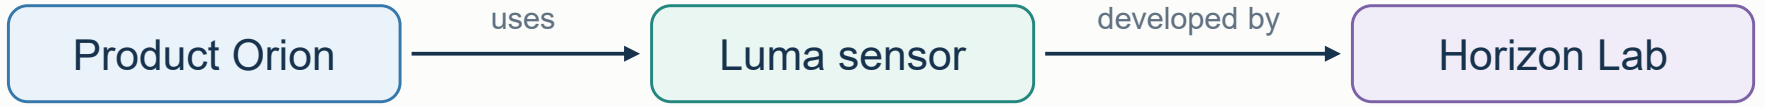
The route is shaped by what becomes available along the way.

A retrieval operation changes the context in which later operations are chosen.

One question can lead to several search routes

Which lab developed the sensor used in Product Orion?

FICTIONAL KNOWLEDGE



(a) Search route

Search Orion sensor → Read product manual → Search Luma developer → Read lab profile

(b) Link route

Locate Orion node → Follow uses relation → Follow developed-by link → Read linked source

(c) Mixed route

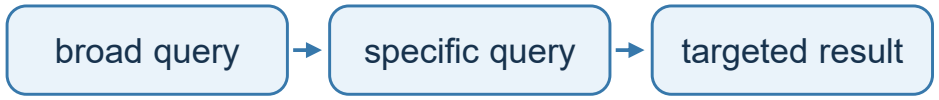
Search Orion sensor → Read product manual → Locate Luma node → Read linked lab source

Compare the process that produces an answer, not only the final answer.

Which search patterns are worth investigating?

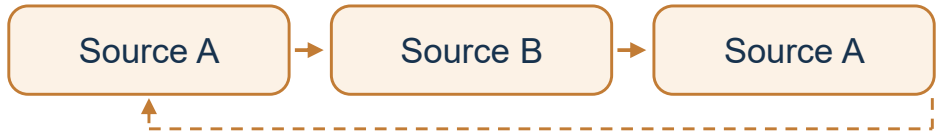
These are candidate observations, not findings or a settled taxonomy.

(a) Progressive refinement



Does each step make the search more targeted?

(b) Revisiting



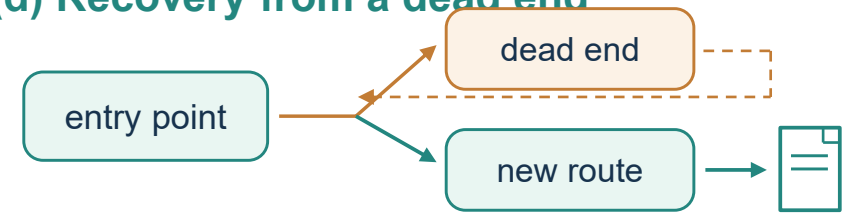
Useful confirmation or repeated effort?

(c) Method switching



Complementary operations or an unnecessary detour?

(d) Recovery from a dead end



Can the process return to a productive path?

Start by identifying recurring patterns, then ask what explains them.

Our question: how does adaptive search change retrieval?

The project aims to characterize a process before proposing an improvement.

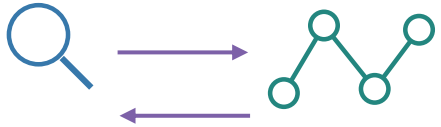
How does agent-controlled search reshape the strengths, limitations and trade-offs of retrieval mechanisms?

(a) Characteristics



What patterns recur?

(b) Interactions



How do methods work together?

(c) Conditions



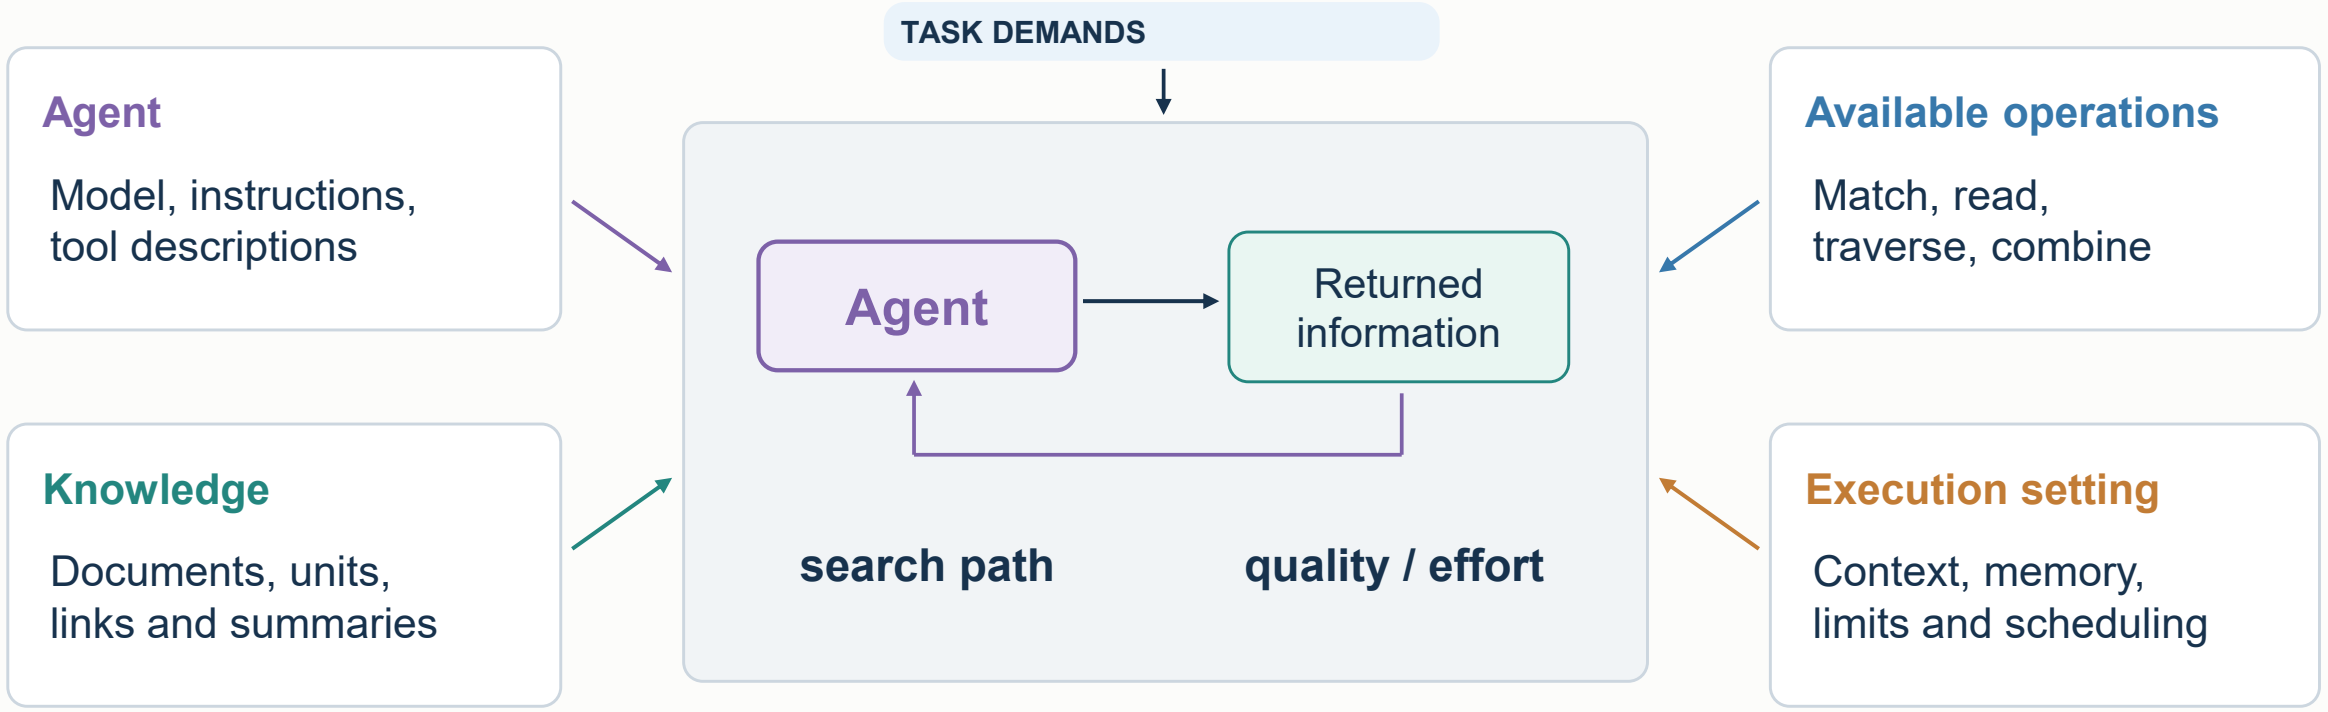
model × task × knowledge

When does the pattern change?

Use exploration to formulate sharper questions, not to preselect a winning method.

Observed behavior has more than one possible explanation

Candidate factors to map; not a claim that they can be varied independently.

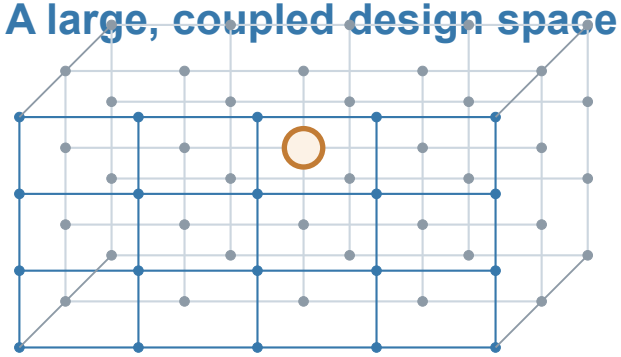


A useful research question needs an explicit boundary around what may explain the behavior.

The challenge: many choices, unclear attribution

The design space is open-ended, and many choices are coupled.

(a) A large, coupled design space



… more choices and interactions

Model	Operations	Knowledge
Instructions	Memory	Budget

A change in one component may require another.

(b) Bundled system comparisons

System A		System B
Model	vs.	Model
Search tools		Search tools
Knowledge		Knowledge
Context rule		Context rule

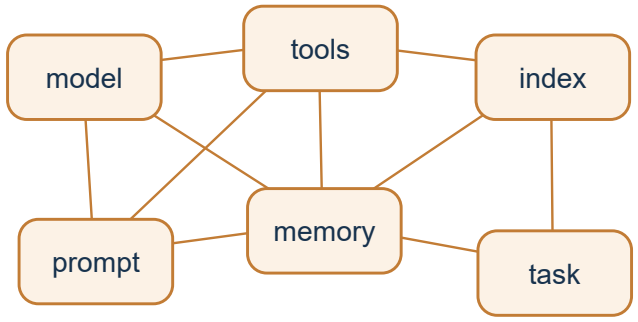
Different outcome.
Which change explains it?

The first challenge is choosing an interpretable question, not enumerating every configuration.

Minimal must not mean removing the phenomenon

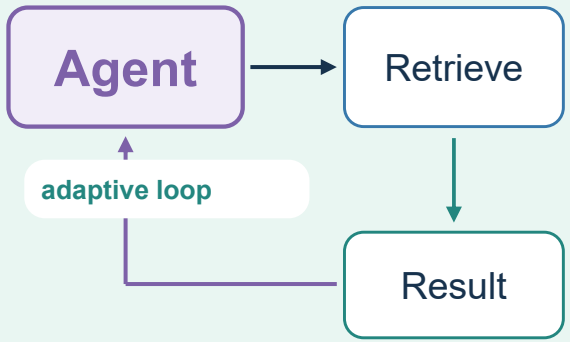
A small setting is useful only if the target behavior can still occur.

(a) Too broad



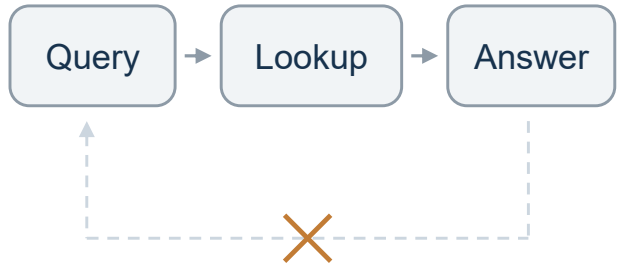
Many moving parts
No clear attribution

(b) Minimal + meaningful



Retain adaptation
Isolate a useful contrast

(c) Too reduced



No consequential choices
Target behavior disappears

Design the smallest setting that preserves the question, not merely the fewest components.

Milestone 1: turn the survey into a design brief

The survey should inform a decision about what to study and what to hold fixed.

Representative systems + studies



Read design choices and implementations

Extract a common description

- Who chooses the next action?
- What operations are available?
- What knowledge can they access?
- What changes together?
- What evidence supports a claim?

Design brief

- Design-space + dependency map
- Shortlist of research questions
- Minimal-setting rationale
- Known limits and open gaps

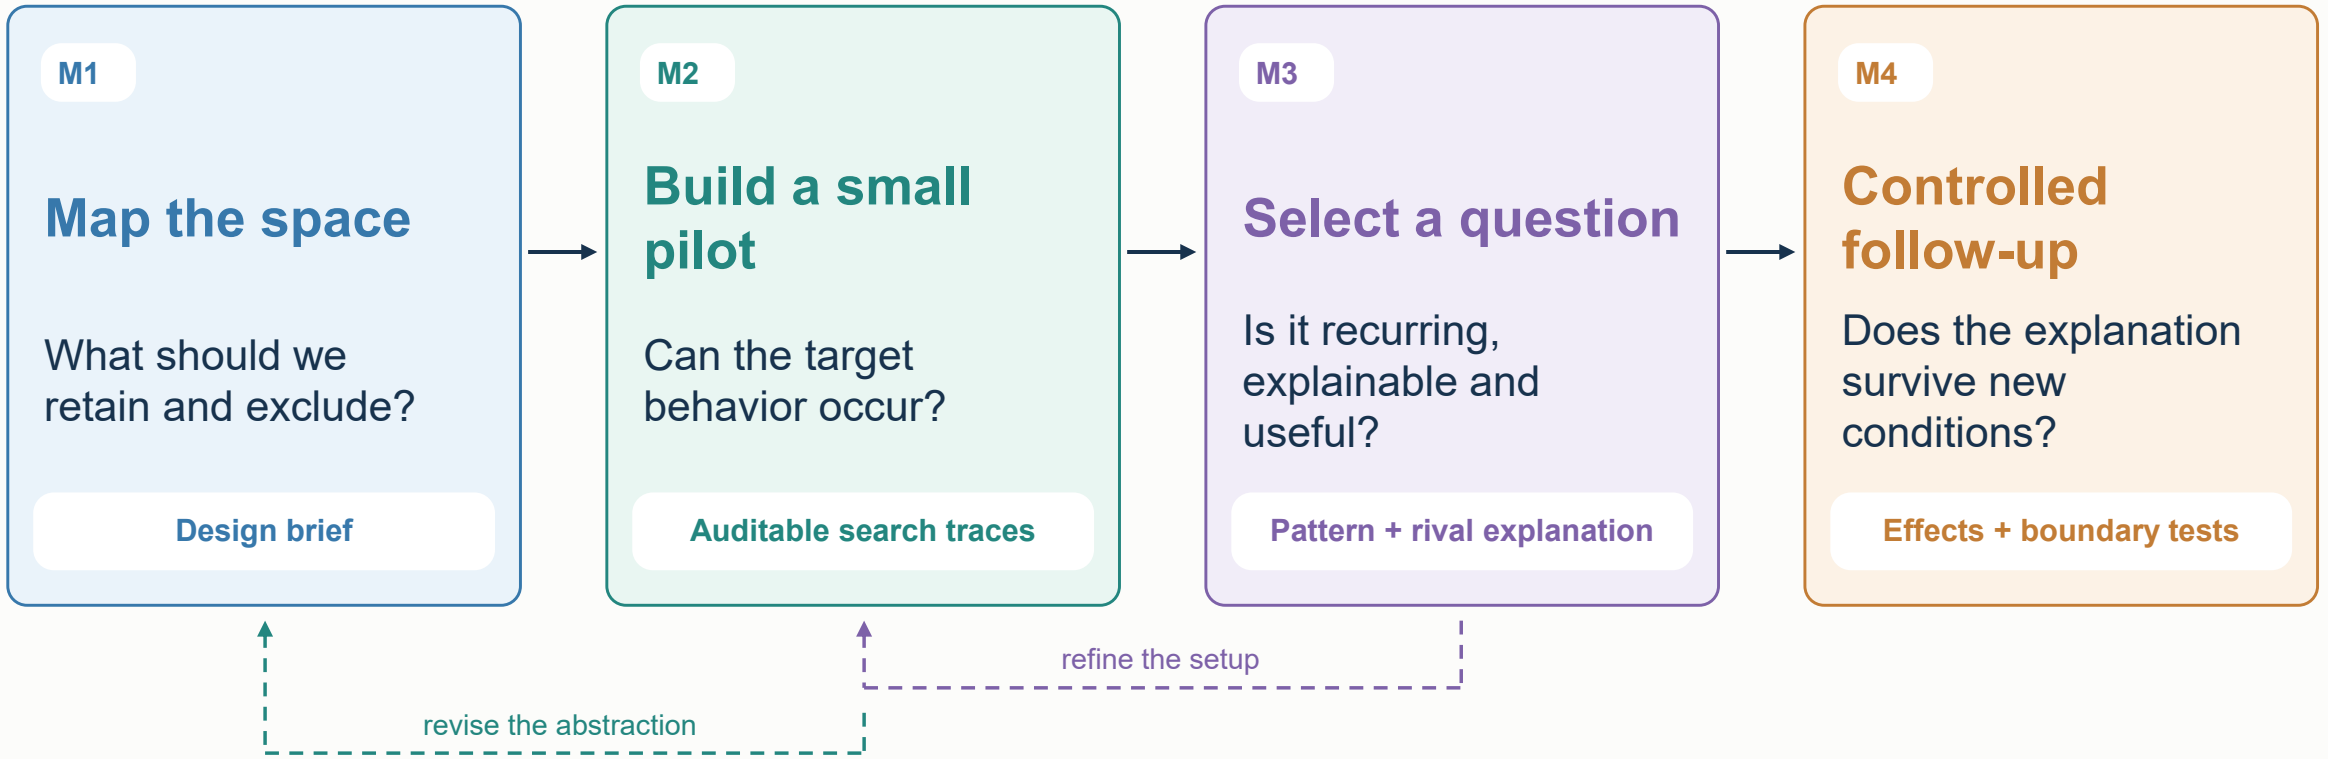
COMPLETION CRITERION

Justify what to retain, vary, hold fixed and defer.

Survey is a research-design milestone, not just a bibliography-building task.

Discover first, then test the explanation

The initial setup is a tool for discovery, not a commitment to the final research question.



Long-term aim: reproducible, generalizable insight into how agentic retrieval behaves.

Today’s decision: agree on the first exploration

Leave with a survey scope and ownership, not a full experiment grid.

1

Scope

Which aspects of agent-controlled retrieval belong in this project?

2

Starting question

Which possible characteristic is important enough to investigate first?

3

First deliverable

Who owns the design-space map and the minimal-setting rationale?

FIRST MILESTONE



Survey-driven design brief

A justified scope.
A tractable next question.

The first contribution is learning how to ask a useful, answerable question.