

Foundations of Information Retrieval

Translating human intent into machine action.




Based on Search Semantically
by Dr. Mahdi Habibi.

The Lossy Bridge

Information Need

The Translation Gap

Query




What the user actually wants and understands perfectly in their own mind.

The user struggles to state their need perfectly. Information is lost.

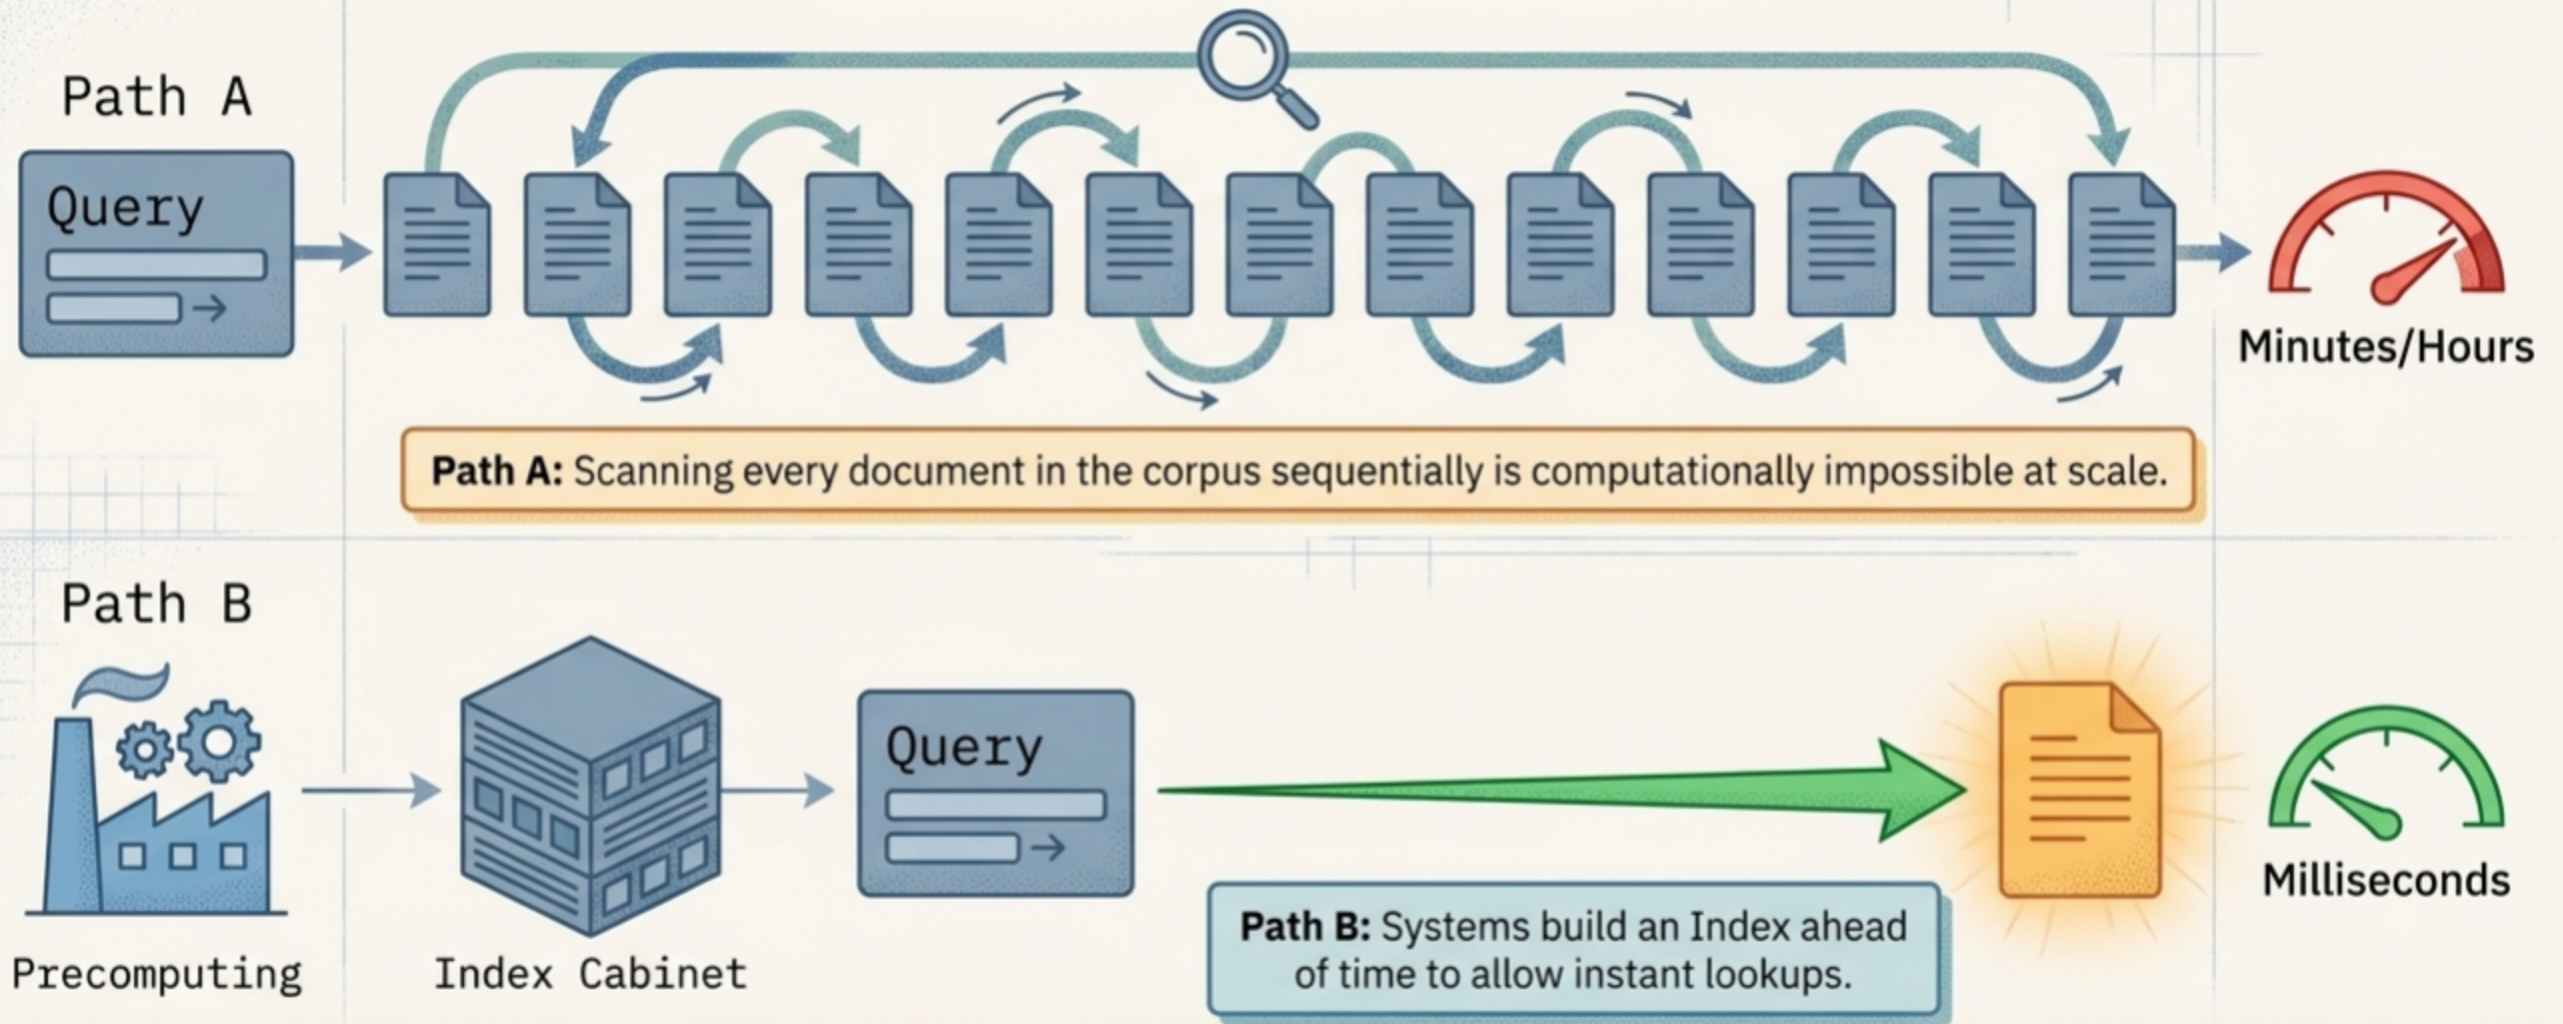
The imperfect, lossy stand-in the user actually types into the search bar.

Our job as search engineers is to bridge this lossy query back to the documents that truly answer the original, unspoken need.

The Anatomy of Search



The Necessity of Indexing




The index must be built before the query ever arrives so answers can be retrieved via instant lookup rather than exhaustive reading.

Stage 1: Retrieval

Selects a broad set of candidate documents from the entire index.

Stage 2: Ranking

Precisely re-orders the candidates so the most relevant document appears first.



	Stage 1: Retrieval	Stage 2: Re-Ranking
Primary Goal	Cast a wide net to find potentials.	Sort the top potentials flawlessly.
Action Metaphor	Fishing with a massive net. 	Using precise tweezers. 
Data Volume	Scans the entire Corpus.	Only scans the narrowed Candidate Set.
Computational Cost	Cheap and Fast.	Expensive and Heavy.
Precision Level	Coarse (Accepts false positives).	Exacting (Prioritizes true relevance).

Retrieve wide and cheap; re-rank narrow and precise.

The Lexical Miss



The Mechanical Blind Spot


Simple search engines count shared words. This works perfectly for simple facts (e.g., "what color is the grass").



The Semantic Failure

When the user and the author use completely different words to describe the exact same meaning, the system fails to connect them.

The Vocabulary-Mismatch Problem



Blind to Meaning

Traditional search systems are highly optimized for speed and sorting, but they are entirely blind to actual meaning.

Literal Overlap

They rely strictly on literal, mechanical word overlap to calculate a match between a query and a document.

Relevance Lost

If the query doesn't perfectly match the document's exact vocabulary, relevance is lost—even if the underlying information need matches perfectly.

To fix this, search must evolve from matching words to matching meaning.

Knowledge Check

Q: What is the difference between an information need and a query?

A: The need is what the user actually wants; the query is the imperfect, lossy text they type to express it.

Q: Why build an index before any query arrives?

A: To allow queries to be answered in milliseconds via instant lookup, rather than linearly scanning every document.

Q: What are the two classic search stages and what does each optimize?

A: Retrieval (optimizes for fast, wide, cheap candidate selection) and Re-ranking (optimizes for precise, narrow ordering).