FEEDBACK-GUIDED SEARCH

PART 7 00:00-01:11 / synthesis

Search Needs More Than More Samples

Stanford CS329A Part 7 - diversity, feedback, selection, and stopping

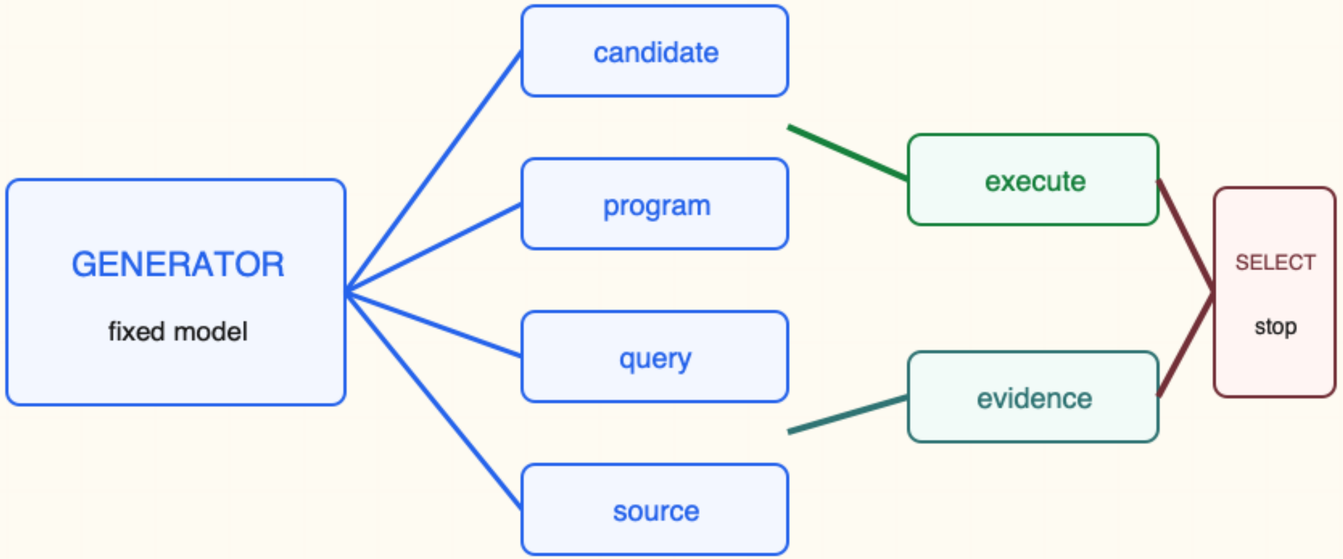

Diversity

Budget

TASK STRUCTURE

PART 7 01:11-04:30 + AlphaCode

Contest Code Requires a Whole Problem-Solving Chain

The model must interpret, plan, implement, and survive hidden tests.

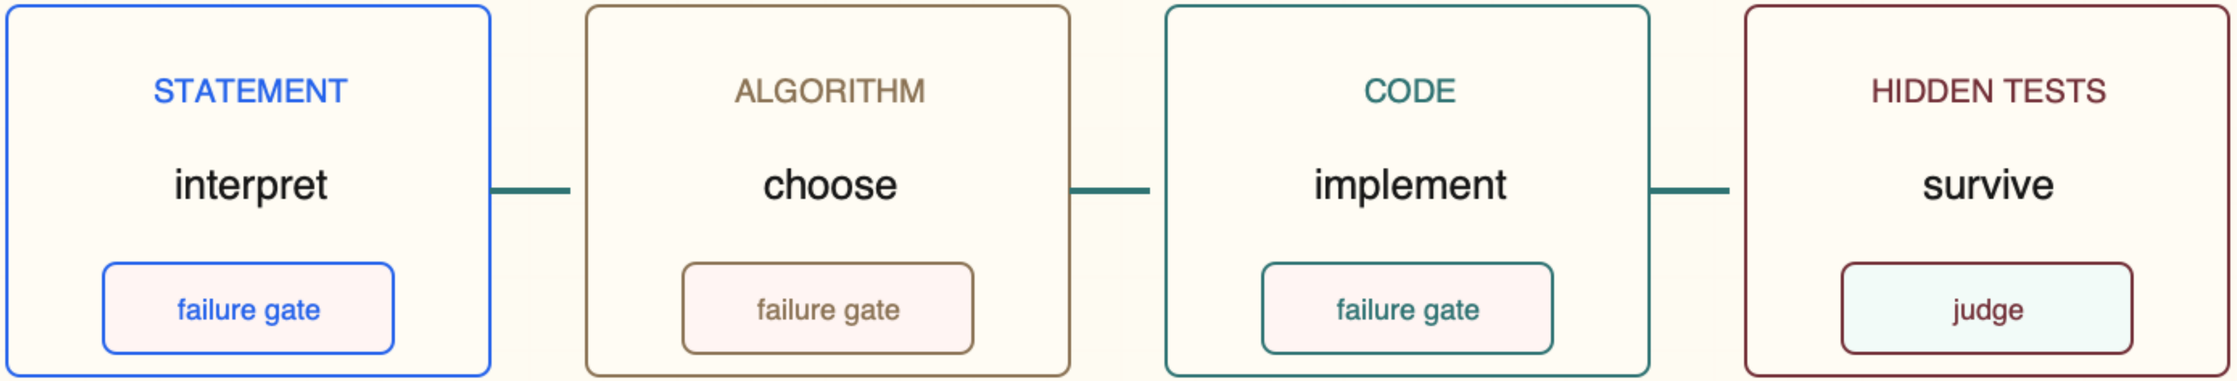

BOUNDARY

Competitive programming is structured and executable; it is not a complete proxy for professional software engineering.

CANDIDATE SEARCH

PART 7 04:30-09:40 + arXiv:2203.07814v1

AlphaCode Searches a Million-Scale Candidate Space

Sampling creates coverage; execution and clustering compress the pool.

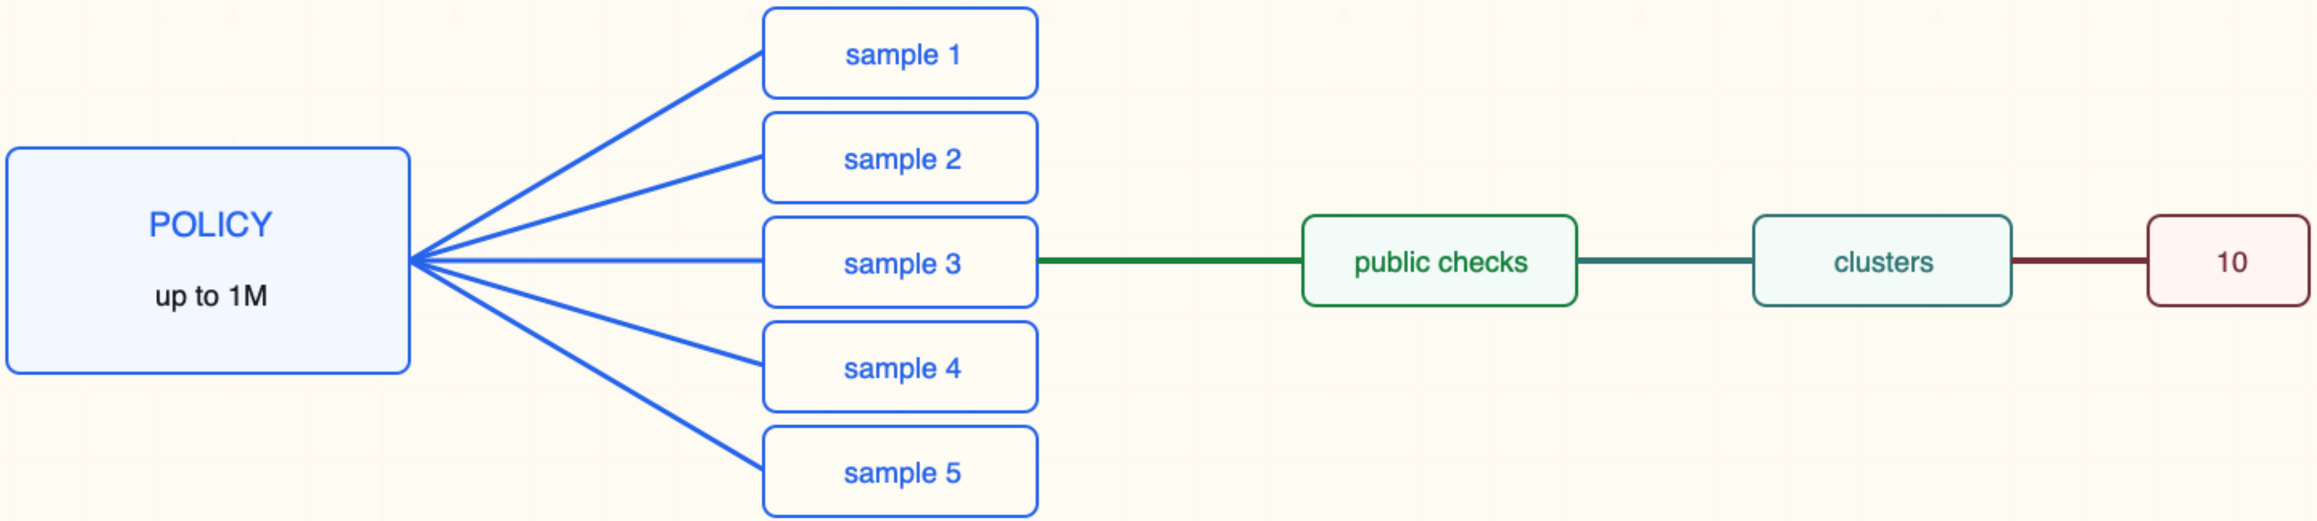

BOUNDARY

Public checks can reject failures; passing them does not prove correctness on hidden tests.

METRIC BOUNDARY

PART 7 09:40-16:35 + AlphaCode

Finding One Correct Program Is Not the Same as Selecting Ten

pass@k measures candidate coverage; 10@k also measures small-set selection.

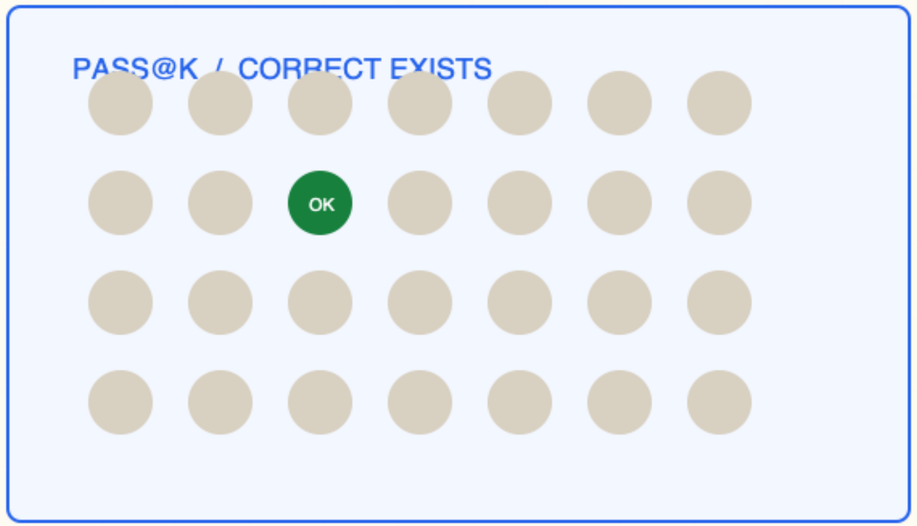

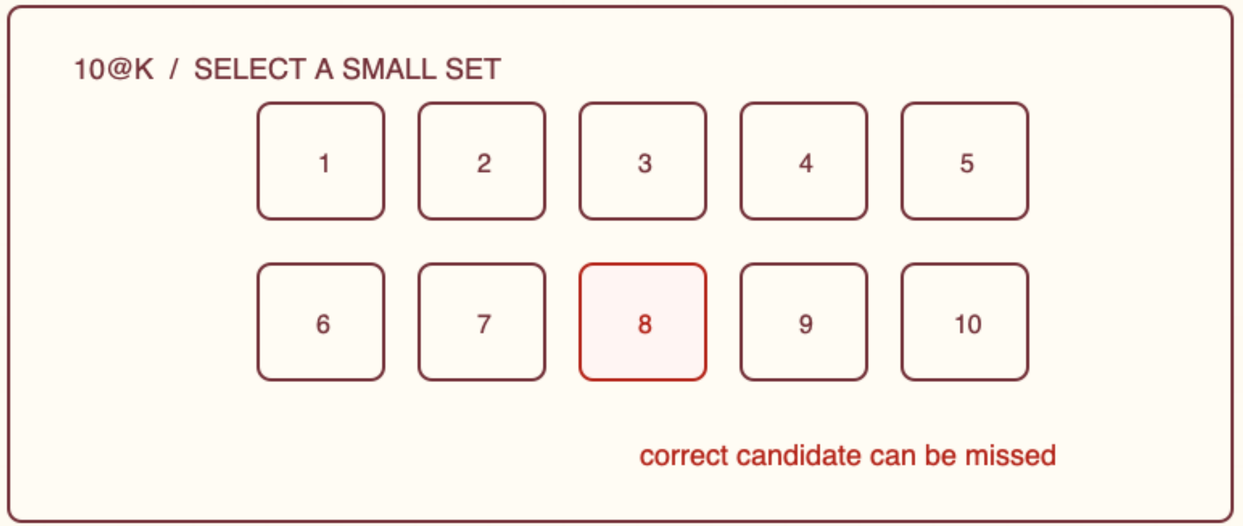

BOUNDARY

Candidate coverage and returned reliability belong to different parts of the system.

SCALING LIMIT

PART 7 16:35-24:38 + AlphaCode

Compute Keeps Rising After Useful Diversity Starts Flattening

Redundant candidates and weak ranking turn more samples into diminishing return.

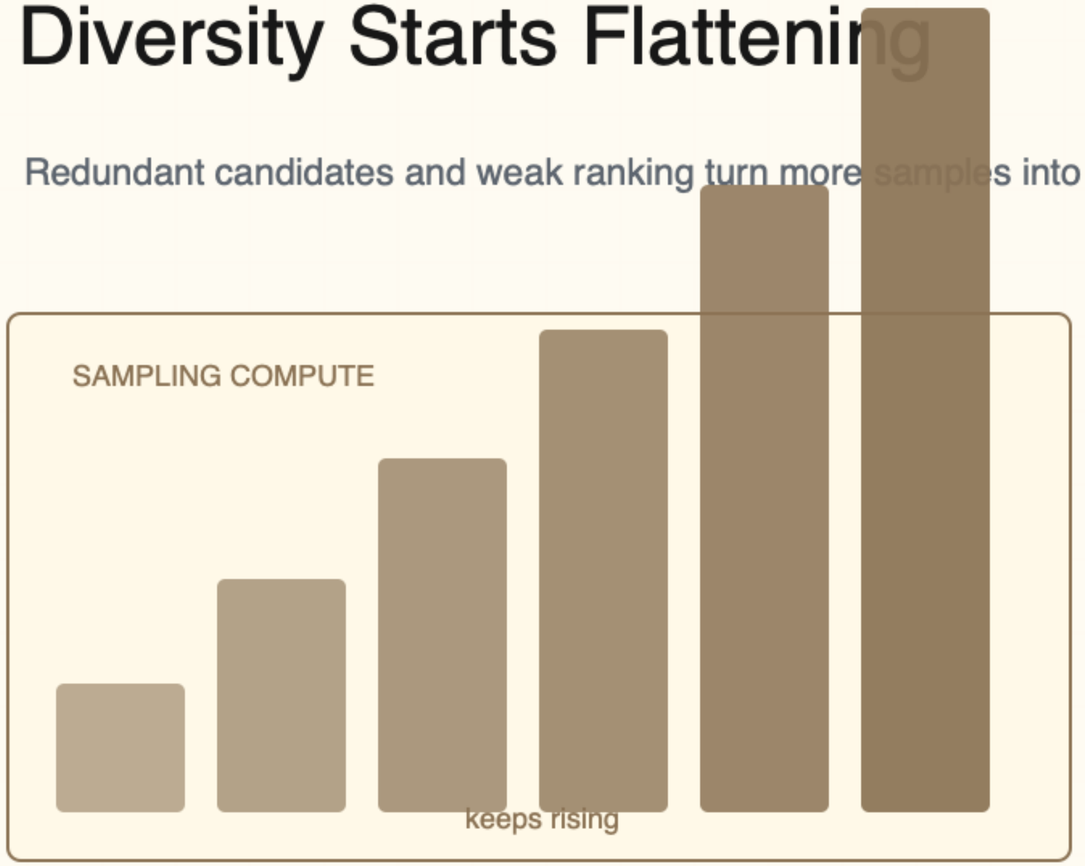

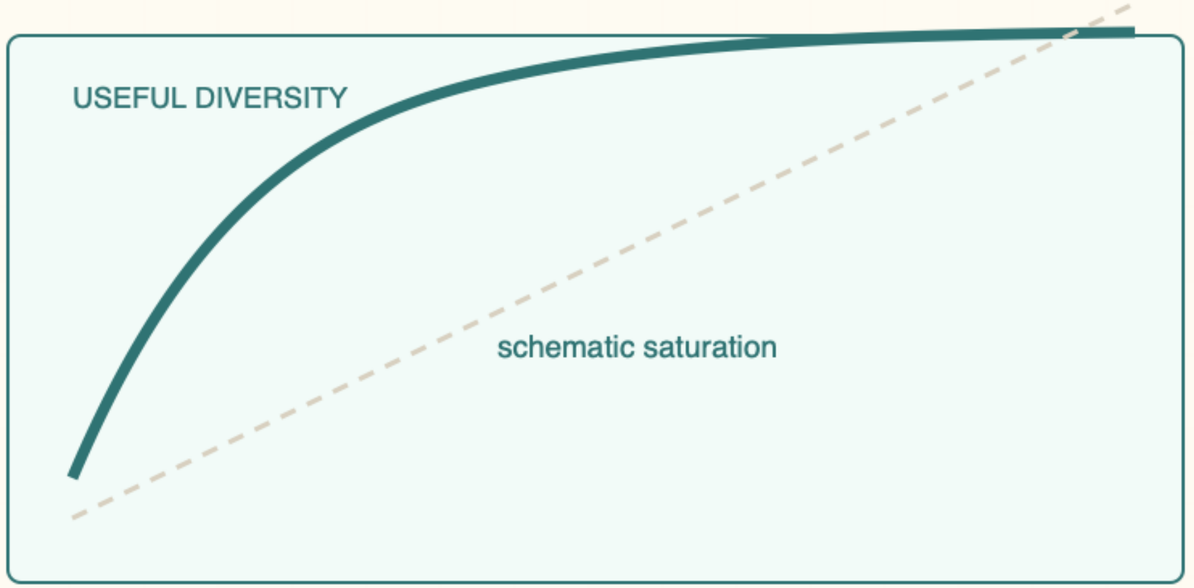

BOUNDARY

Schematic interpretation only: no empirical axis values or universal scaling law are claimed.

SYSTEM REDESIGN

PART 7 24:38-30:20 + AC2 report

AlphaCode 2 Improves the Candidate Supply and the Ranker

A policy family feeds execution, clustering, and a learned scoring model.

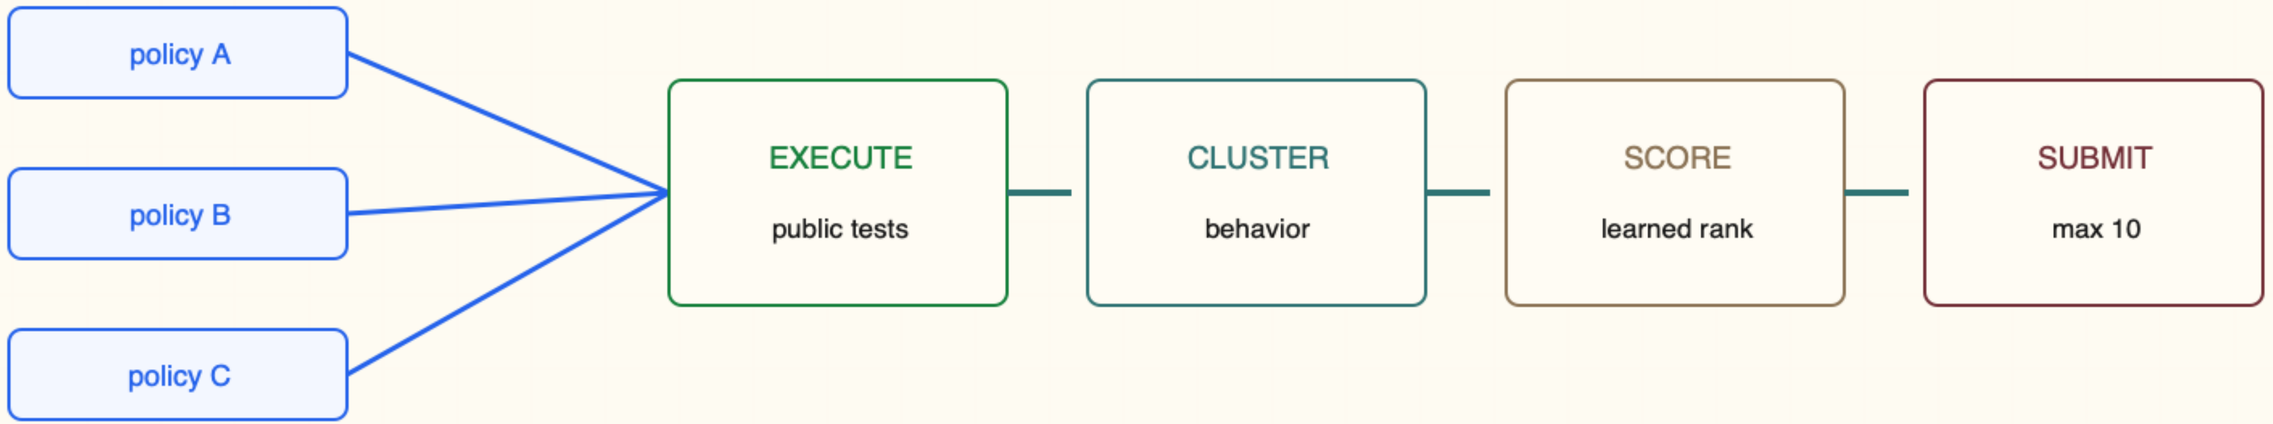

BOUNDARY

Every component exploits competitive programming's unusually strong executable feedback.

CLAIM CALIBRATION

PART 7 30:20-38:20 + AC2 report pp.1-3

The Headline Result Needs Its Denominator

Strong reported performance remains contest-, task-, and budget-specific.

43%	25%	85th
AlphaCode 2 solved	AlphaCode solved	estimated percentile

77 SELECTED PROBLEMS | UP TO 1M CANDIDATES EACH | MAX 10 SUBMISSIONS

Primary report text controls: 85th percentile (not the lecture graphic's apparent ~87%).

ALLOCATION

PART 7 39:25-46:33 / discussion

Harder Problems Need Adaptive Search, Not Just a Wider Fan

Route by difficulty toward sampling, decomposition, tree search, or human review.

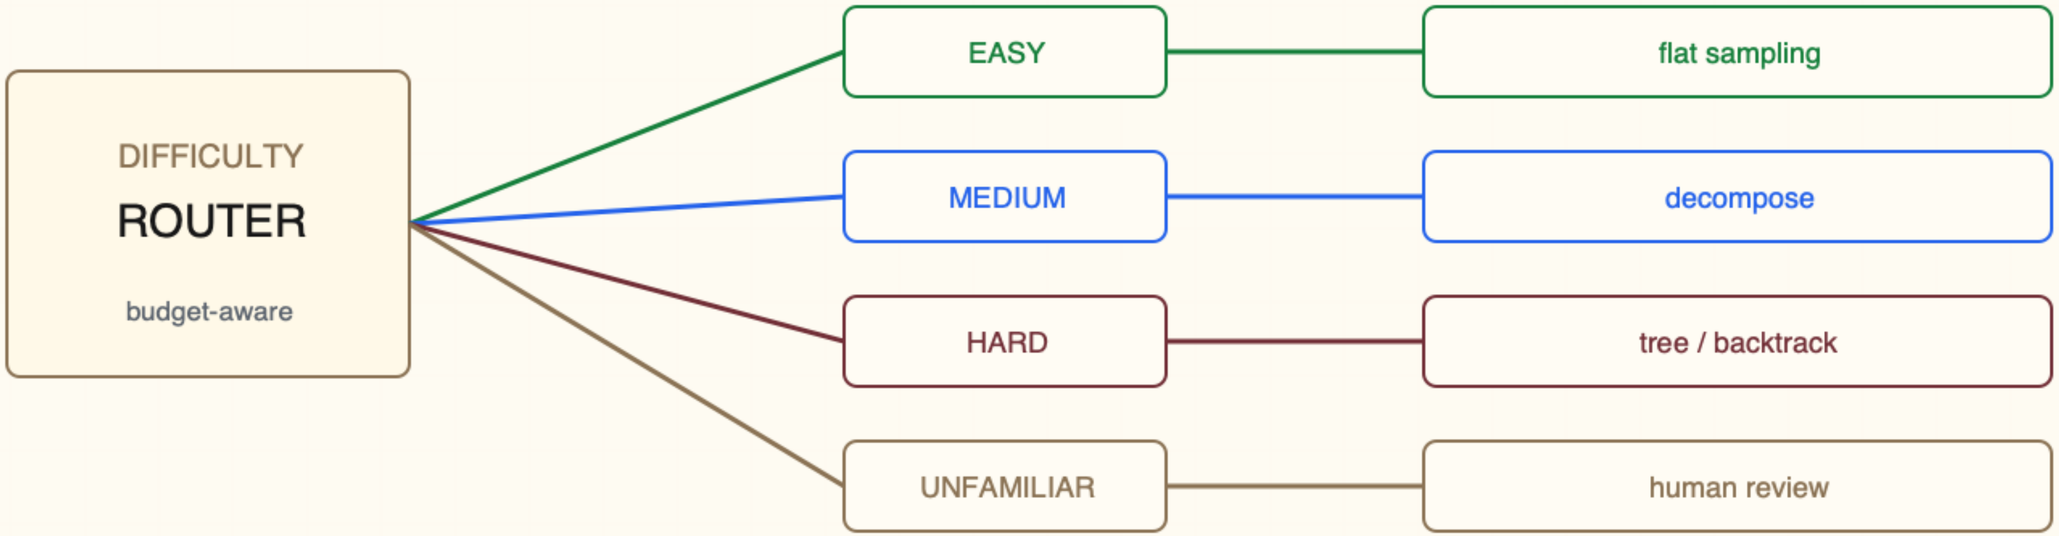

BOUNDARY

Lecture discussion directions, not reported AlphaCode 2 components or ablation results.

REASONING FAILURE

PART 7 46:33-50:40 + Search-o1

One Missing Fact Can Corrupt the Rest of a Reasoning Chain

Fluent continuation can hide the point where uncertainty became a guess.

Fluency continues after evidence stops.

BOUNDARY

The failure begins at the unsupported transition, not only at the final answer.

SEARCH ACTION

PART 7 50:40-55:50 + arXiv:2501.05366

Agentic Retrieval Searches When the Reasoning Process Needs It

Search becomes a recurrent action instead of a one-time preface.

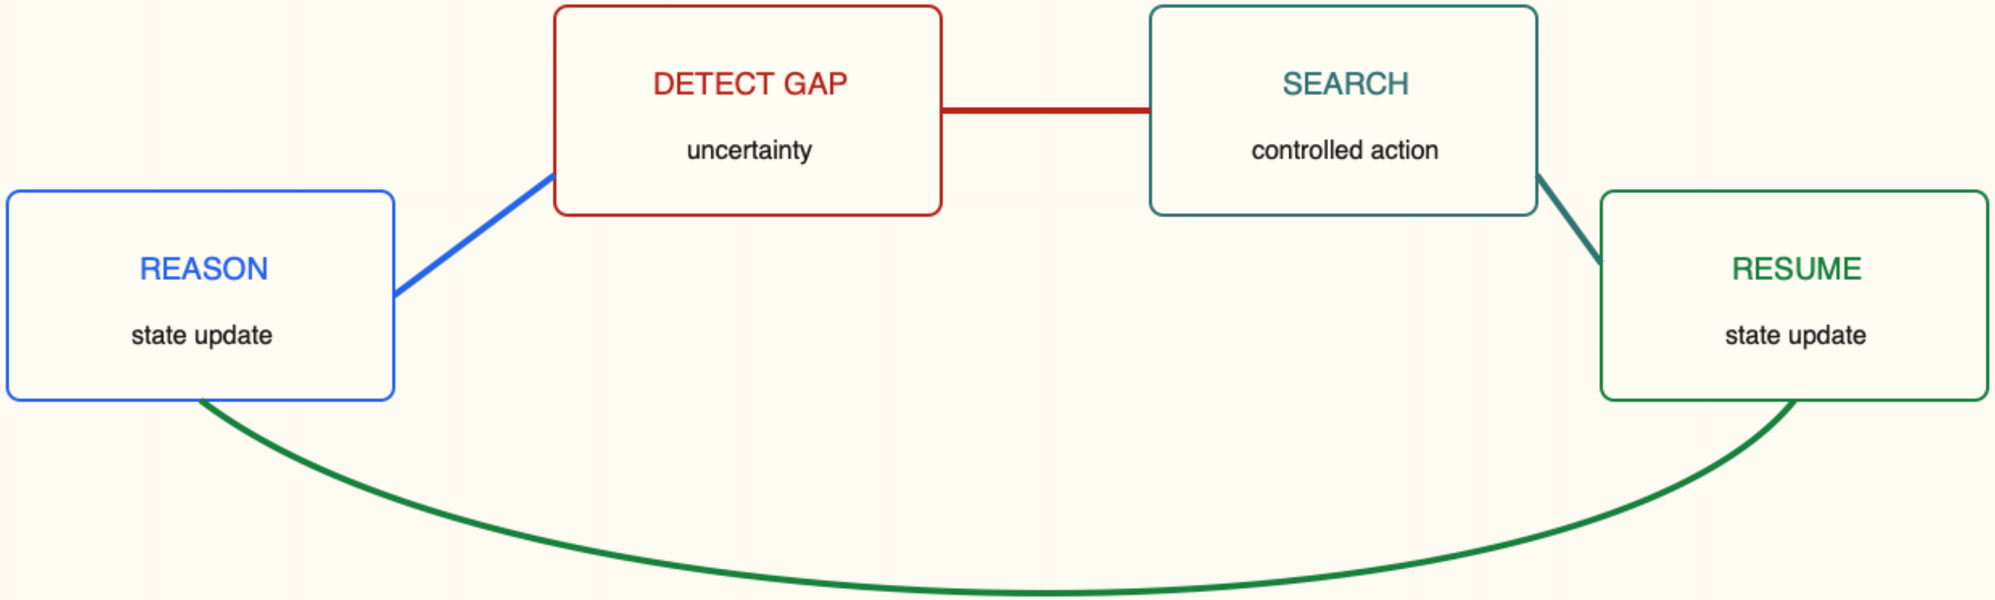

BOUNDARY

Agentic retrieval controls when to search; it does not certify the returned sources.

CONTEXT CONTROL

PART 7 55:50-68:15 + arXiv:2501.05366

Search-o1 Compresses Documents Into Focused Evidence

Retrieve sources, reason within them, extract evidence, then resume the main chain.

BOUNDARY

Focused extraction can omit key sources or preserve errors from retrieved documents.

TEACHING DIAGNOSTIC

PART 7 whole lecture / synthesis

Audit the Loop Before Trusting the Answer

Five questions expose where a search-enhanced agent can fail.

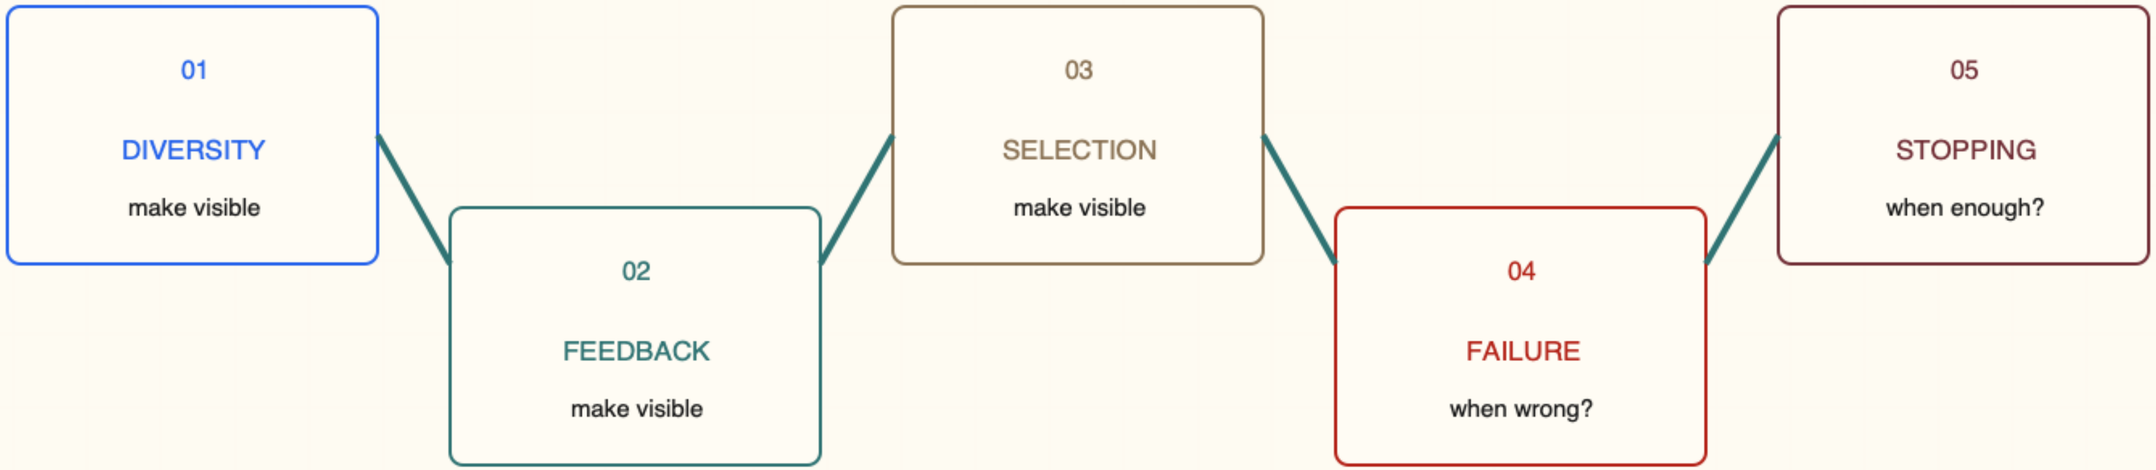

BOUNDED SEARCH != UNBOUNDED TRUST