Session 2: Solving Planning Problems

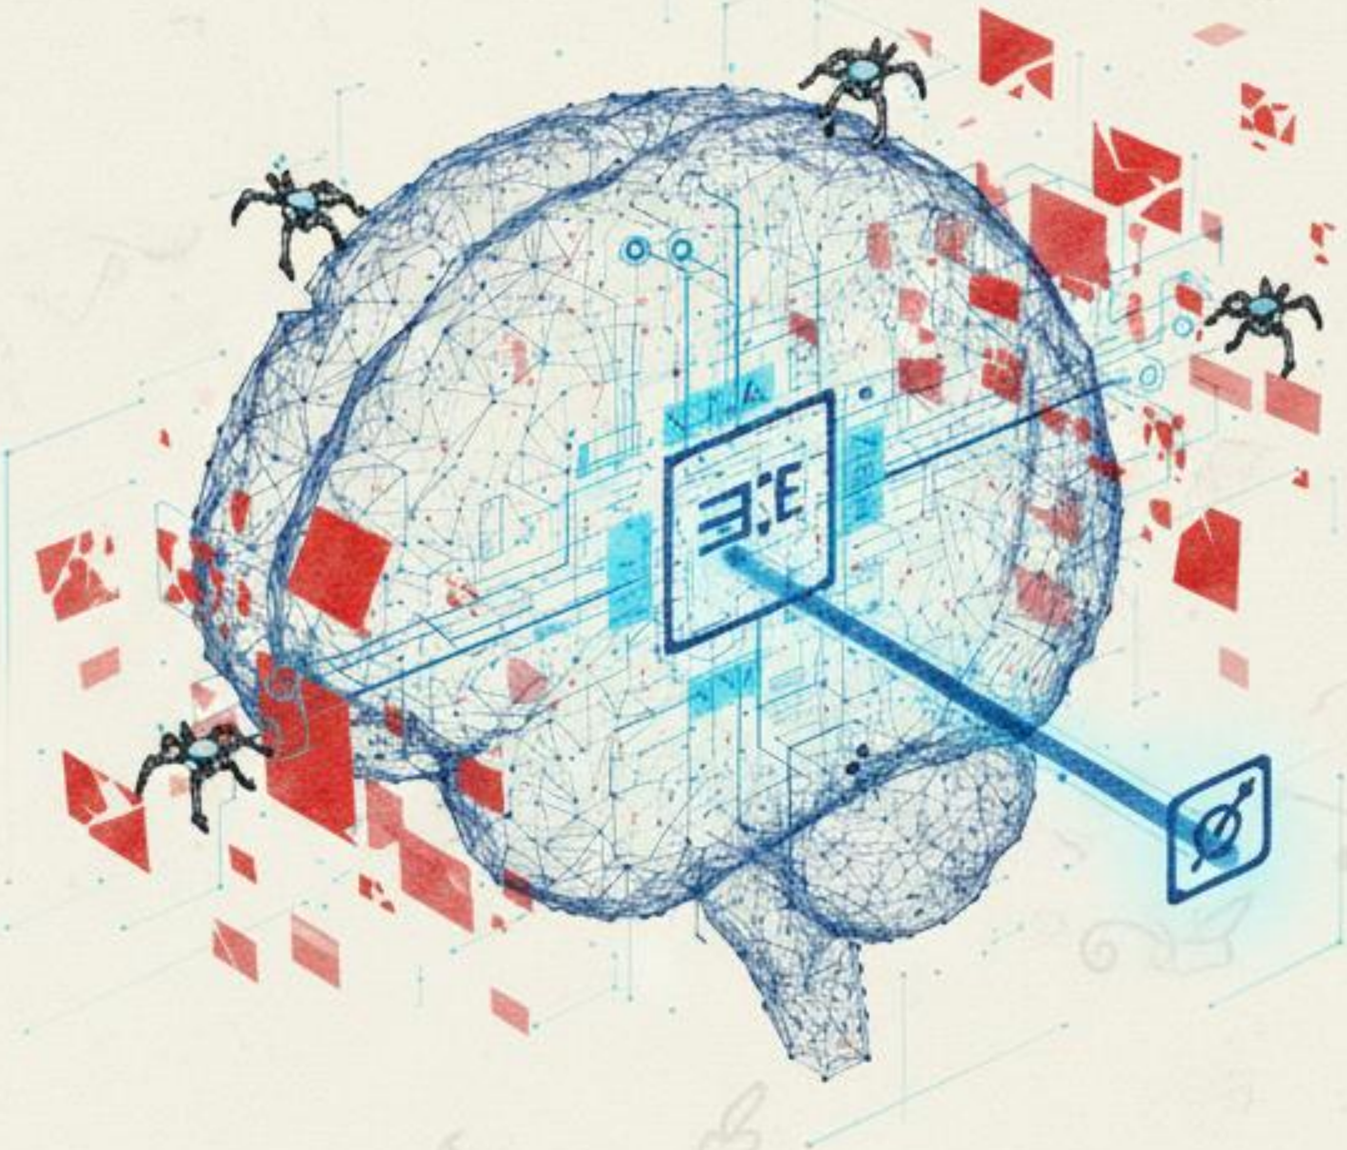

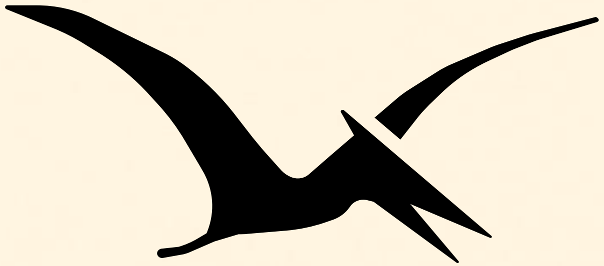

PLANNING IN THE PRE-LLM ERA

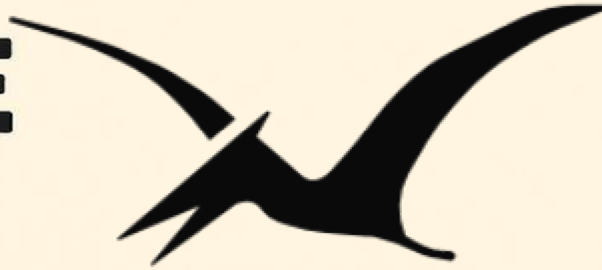

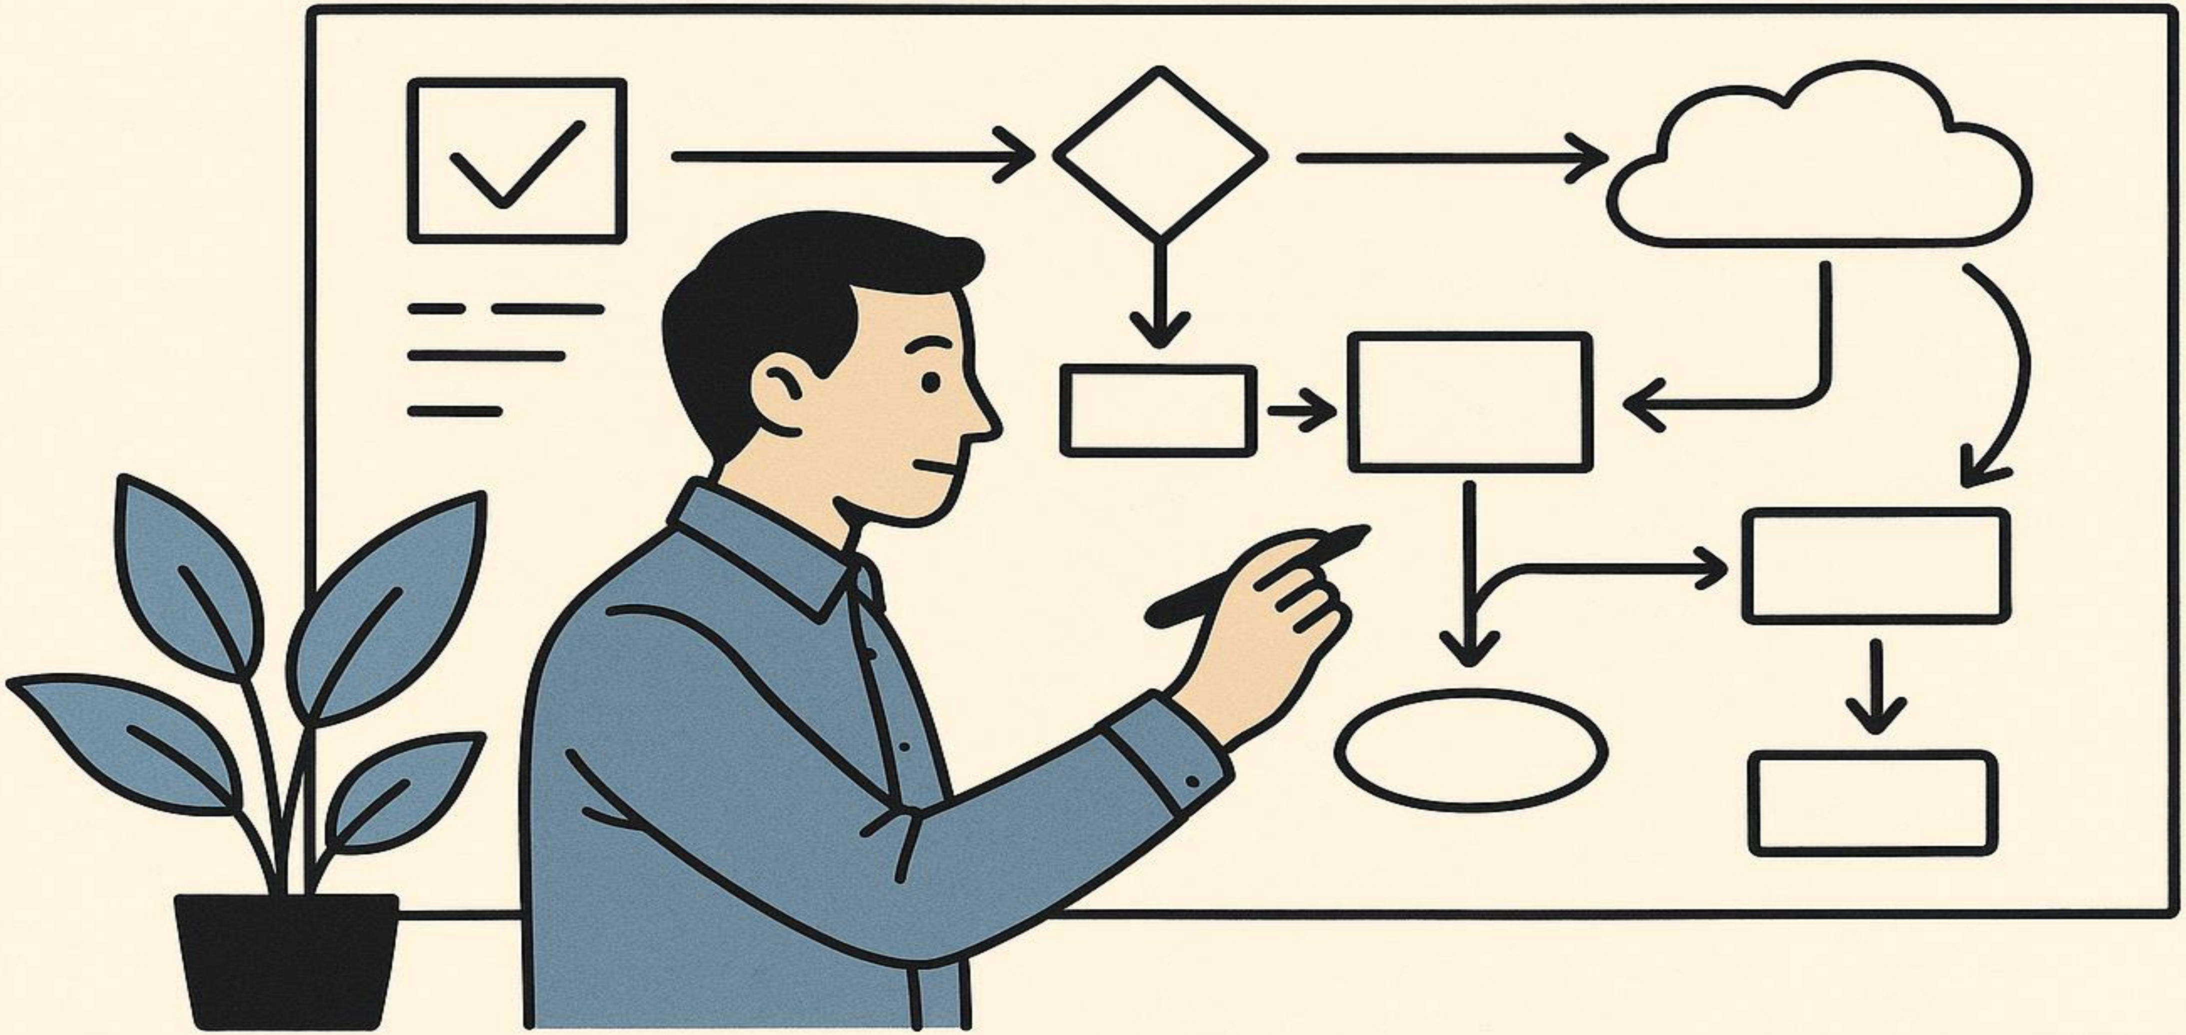

Solving Classical PDDL Planning

PDDL → **Planner** → plans

STRIPS	Fikes & Nilsson, AIJ 1971
GraphPlan	Blum & Furst, IJCAI 1995
SATPlan	Kautz & Selman, AAAI 1996
POCL	Penberthy & Weld, KR 1992
HSP	Bonet & Geffner, AIJ 2001
MCP	Cimatti et. al., ECP 1997

Solving Classical PDDL Planning

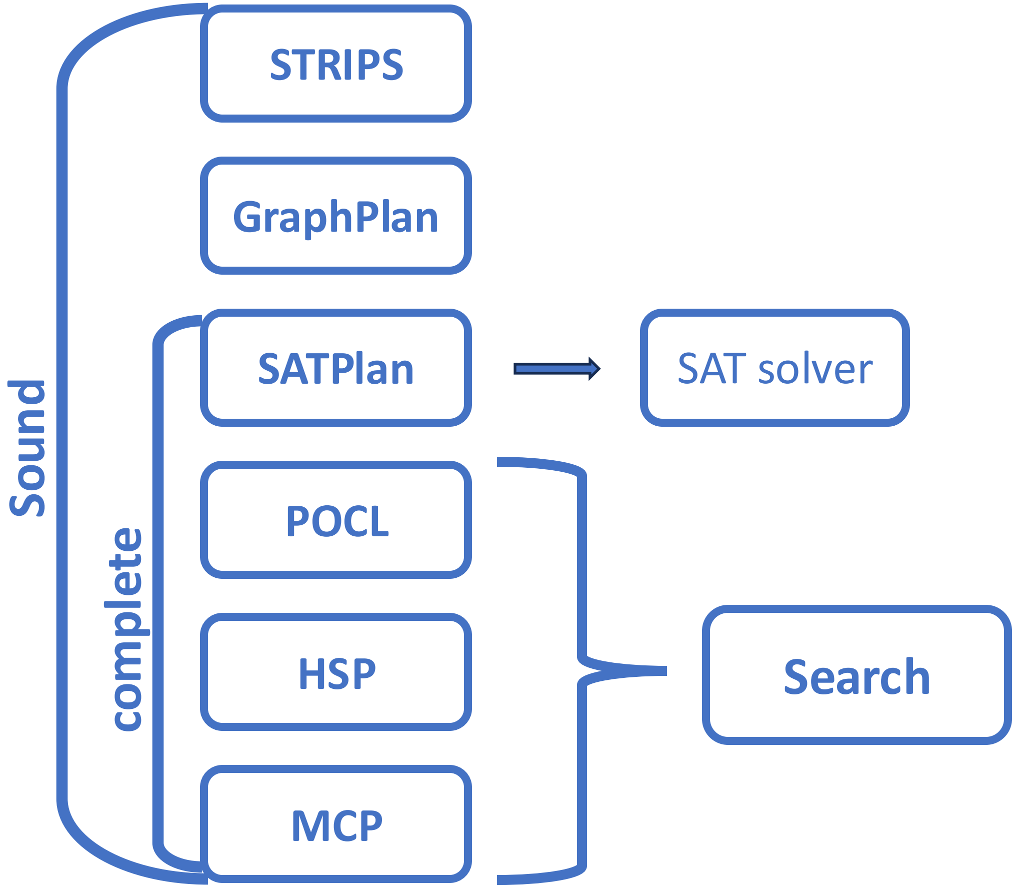

Solving Classical PDDL Planning

PDDL → **Planner** → plans

Heuristic Search

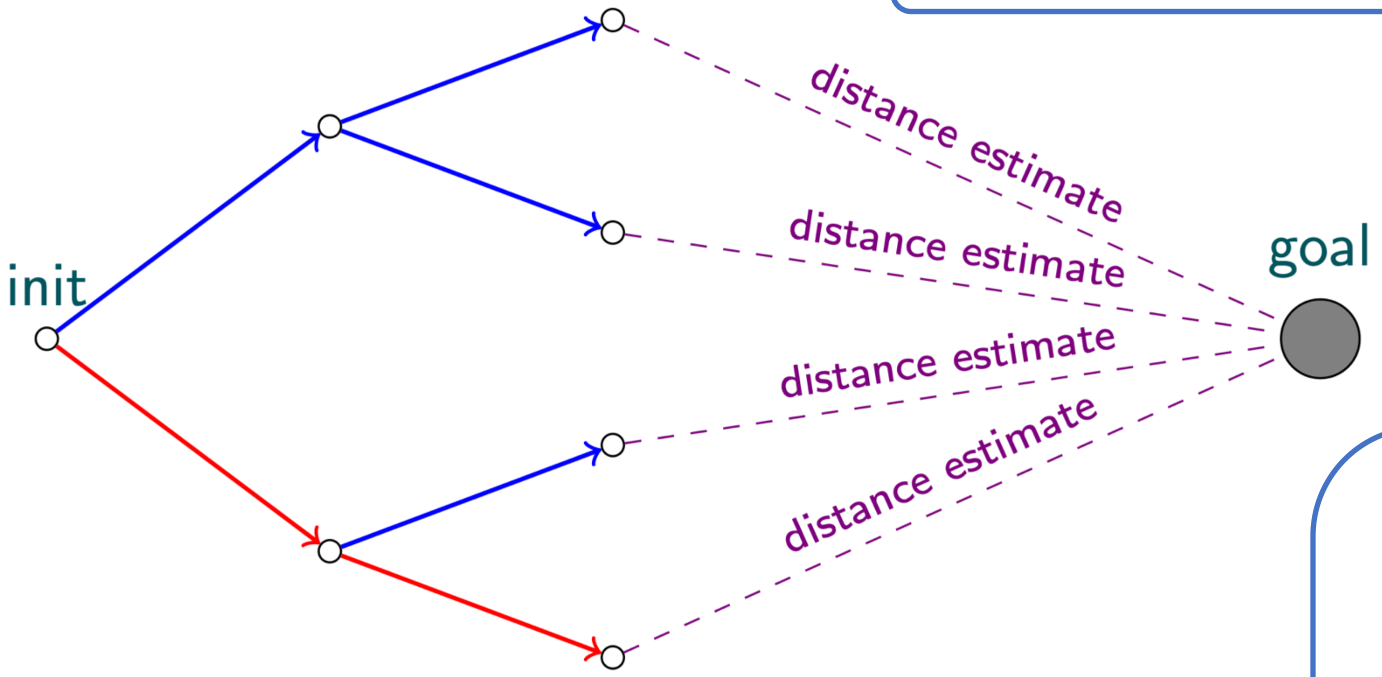

Best-First Search

- Distance estimate: **h**
- Distance from init: **g**
- Explores nodes in the order of $\mathbf{f = g + w \cdot h}$
- **w**=1 — A*
- **w**>1 — wA*
- **w** → ∞ — GBFS

Solving Classical PDDL Planning

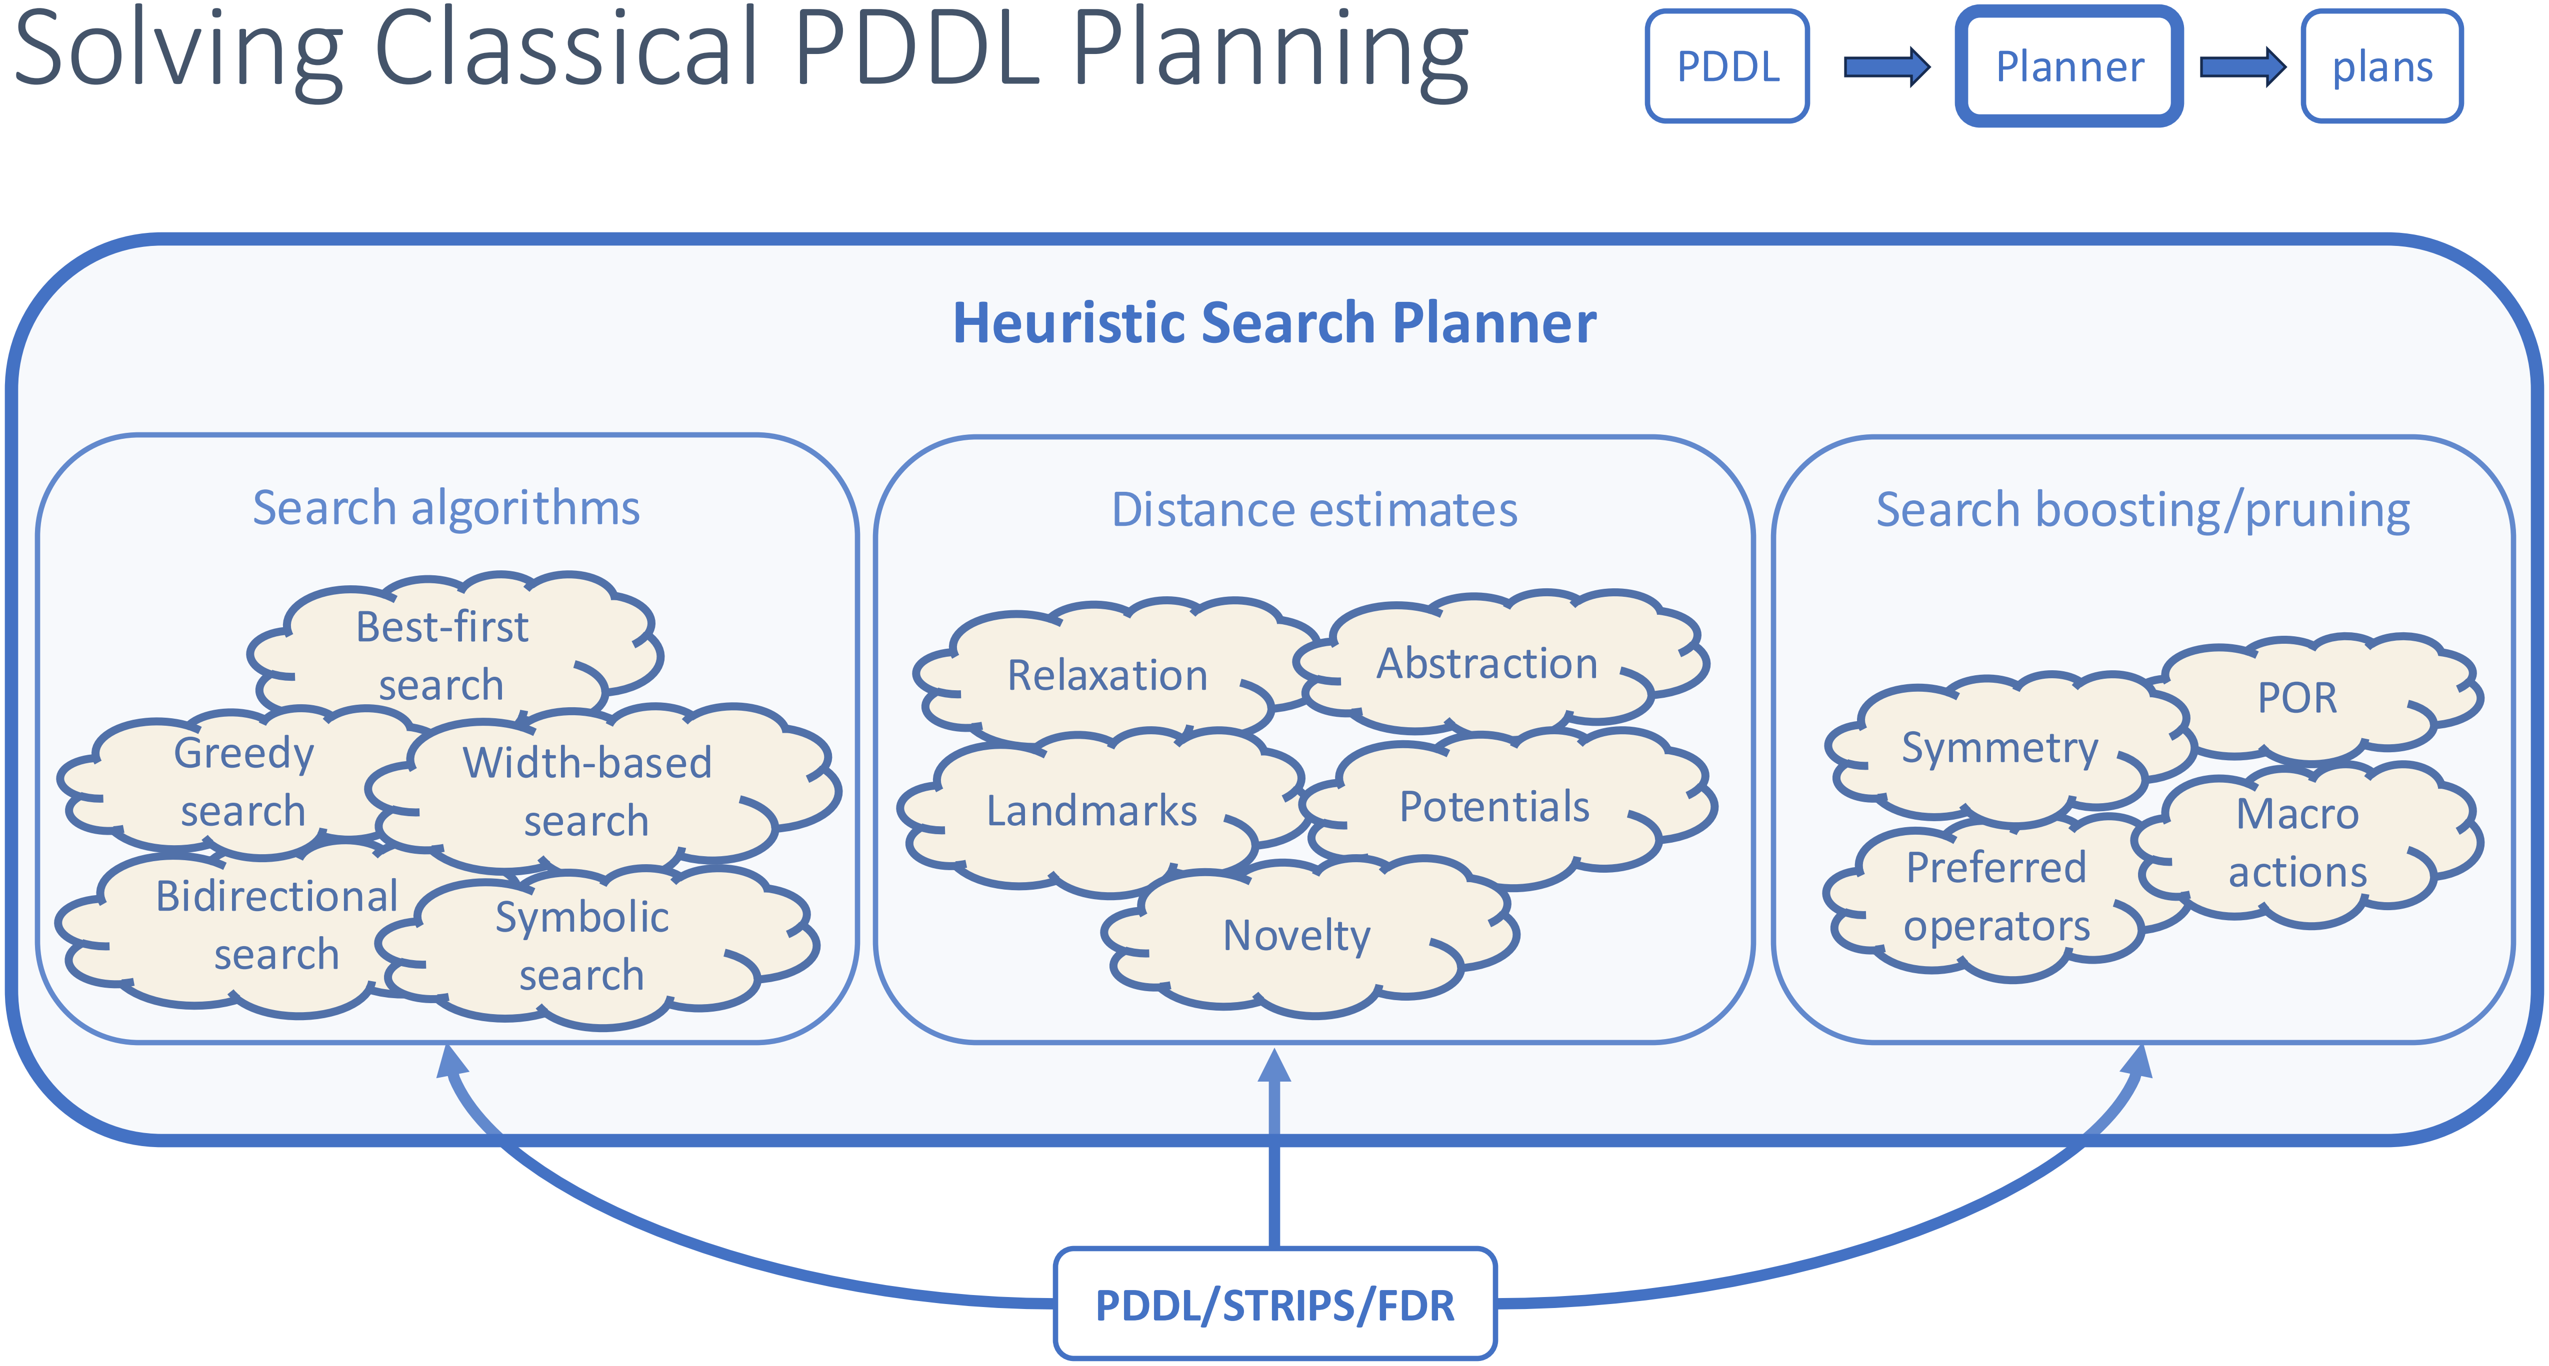

PLANNING IN THE EARLY LLM ERA

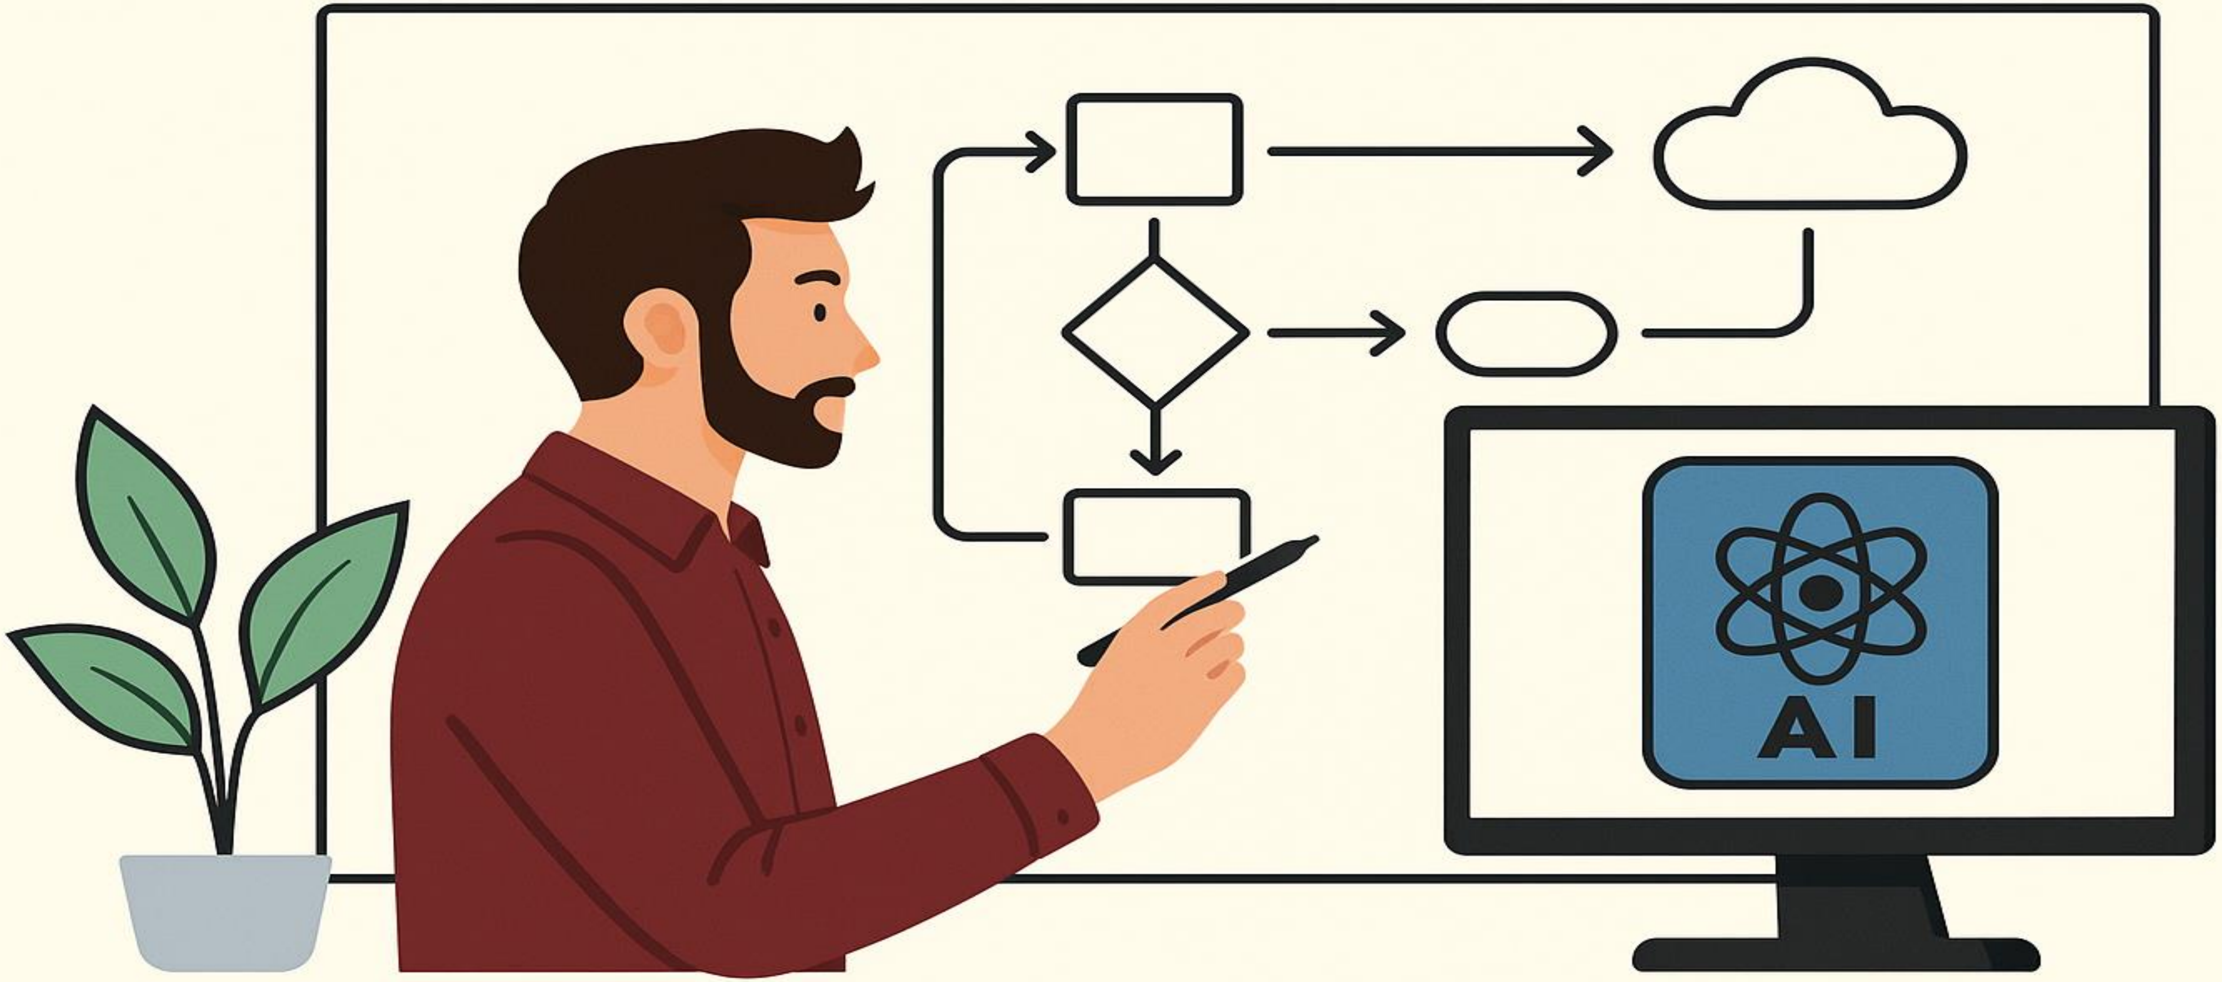

Solving NL/PDDL Planning Problems

Silver et al. FMDM@NeurIPS 2022
Valmeekam et al. NeurIPS 2023
Kambhampati et al. ICML 2024
Bohnet et al. Arxiv 2024
Zhao et al. ICLR 2025

Yao et al. ICLR 2023
Xu et al. Arxiv 2023
Hao et al. EMNLP 2023
Yao et al. NeurIPS 2023
Shinn et al. NeurIPS 2023
Zhou et al. ICML 2024
Sel et al. ICML 2024
Besta et al. AAAI 2024

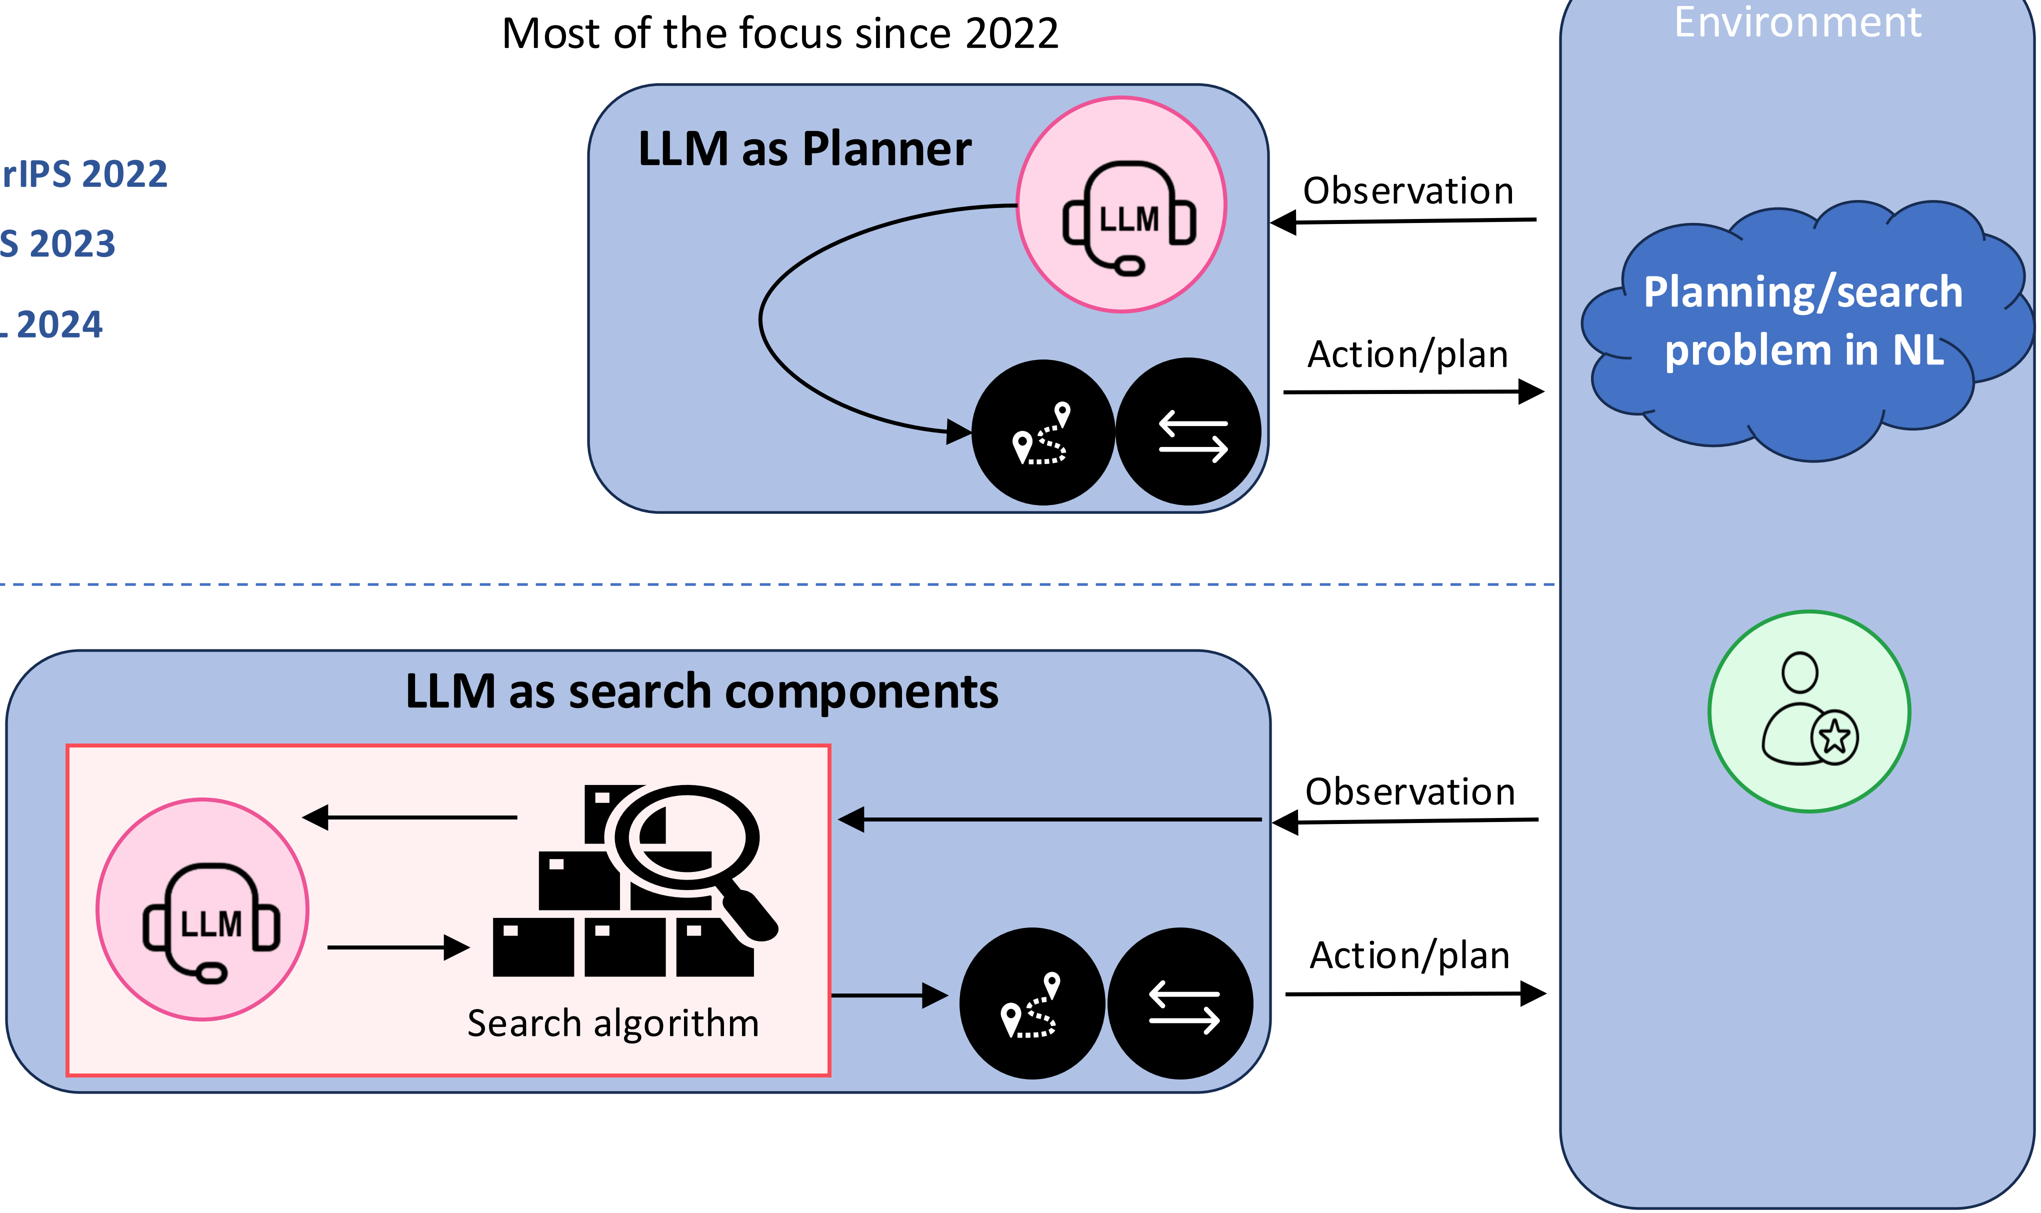

What are the properties of these algorithms?

Sound? No! If validator exists, we can make it sound

Complete? No!

Optimal? No!

Computational complexity?

- Hard to measure. Most expensive operation is LLM call
- Computational effort vs. state space portion explored

		24Game		Crossword		BlocksWorld		PrOntoQA	
Approach	Complexity	States	Calls	States	Calls	States	Calls	States	Calls
IO	O(D)	0.02%	1362	4e-9%	20	0.5%	502	4%	4000
CoT	O(D)	0.02%	1362	4e-9%	20	0.5%	502	4%	4000
ReAct	O(LD)	0.07%	4086	4e-8%	200	7.8%	8032	24.6%	24K
ReWOO	O(LD)	0.07%	4086	4e-8%	200	7.8%	8032	24.6%	24K
RAP	O(TbLD)	3.3%	245K	2e-6%	12K	388%	482K	1229%	1.44M
ToT	O(bmLD)	1.6%	102K	1e-6%	5K	194%	201K	615%	600K
GoT	O(bLD)	0.3%	20K	2e-7%	1K	39%	40K	122%	120K
Reflection	O(LTD)	0.7%	68K	4e-7%	2.4K	77.6%	90K	245%	320K
LATS	O(TbLD)	3.3%	286K	2e-6%	14K	388%	562K	1229%	1.68M

Katz et al, NeurIPS 2024, **Thought of Search: Planning with Language Models Through The Lens of Efficiency**

Why did it work in the first place?

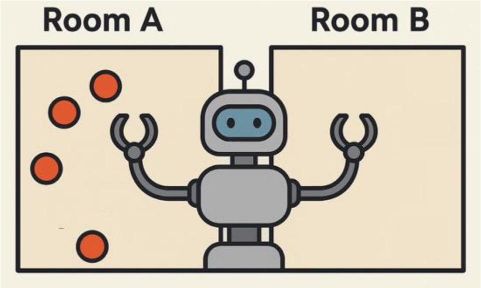

(pick o1 A L)
(pick o2 A R)
(move A B)
(drop o1 B L)
(drop o2 B R)
(move B A)
(pick o3 A L)
(pick o4 A R)
(move A B)
(drop o3 B L)
(drop o4 B R)

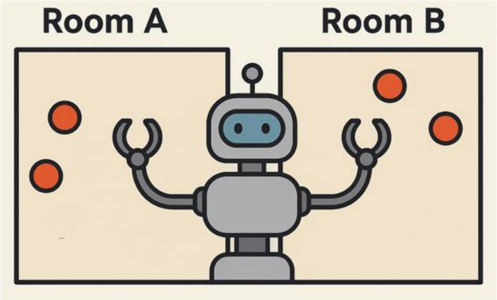

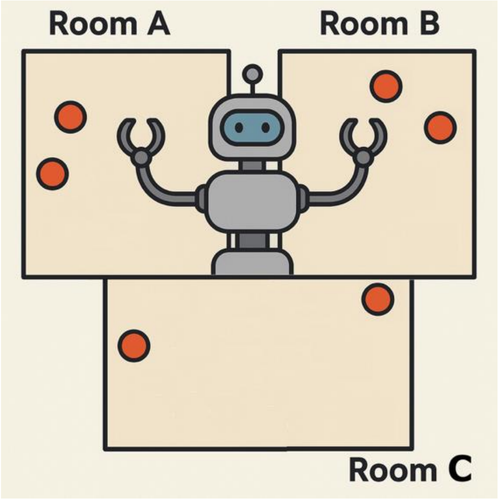

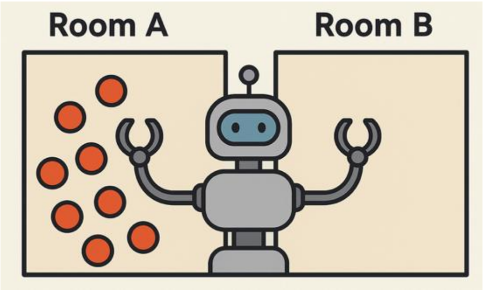

(pick o1 A L)
(pick o2 A R)
(move A B)
(drop o1 B L)
(drop o2 B R)
(move B A)
…
(drop o7 B L)
(drop o8 B R)

Conclusions (see [1]):

- Be aware of instance generator limitations
- Show generalization outside of training set
- Show performance on multiple domains

[1] Katz et al, Arxiv 2025, **Make Planning Research Rigorous Again!**

Why did it work in the first place?

(pickup ?x)
(putdown ?x)
(stack ?x ?y)
(unstack ?x ?y)

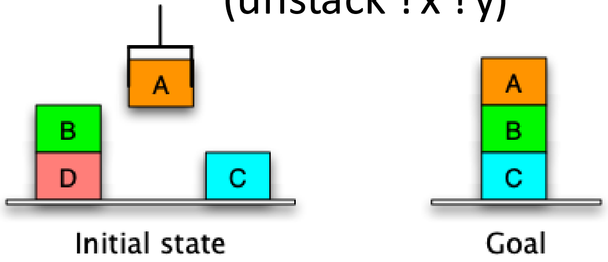

How do plans look like?

- pickup -> stack
- stack -> unstack | pickup
- unstack -> stack | putdown
- putdown -> unstack | pickup

What does it mean to restrict $|\pi| \leq 10$?

- At most 5 blocks are moved (even if the instance has 100s of blocks)
- Total of 1331 possible plan patterns
- When trained on a large collection of instances, most (all?) possible plan patterns appear in the training set

Conclusions (see [1]):

- There are many planning domains out there and BlocksWorld is among the simplest
- Show generalization outside of training set
- Show performance on multiple domains

[1] Katz et al, Arxiv 2025, **Make Planning Research Rigorous Again!**

Why did it work in the first place?

- Tested on data scraped from the internet
    - Data contamination concerns
- No external validator, lenient success criteria

	Model	IO 24Game	IO CD[4]	CoT 24Game	CoT CD[4]	ToT 24Game	ToT CD[4]
acc@5	Qwen	6	2	8	2	83	28
	Llama	7	2	32	7	90	40
	DeepSeek	38	5	48	13	77	20
mean	Qwen	2	1	2	0	47	9
	Llama	1	0	9	1	48	12
	DeepSeek	10	1	18	4	28	4

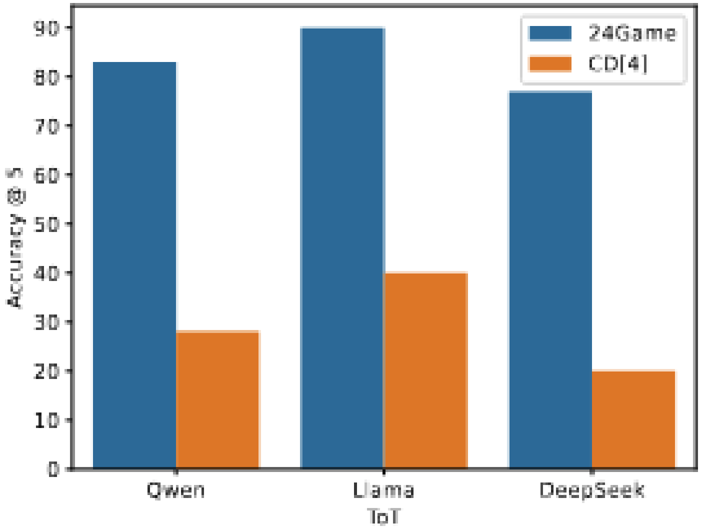

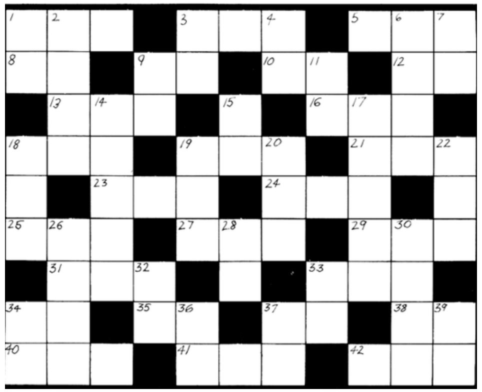

Katz et al, Arxiv 2025, **Seemingly Simple Planning Problems are Computationally Challenging: The Countdown Game**

Conclusions:

- Must have precise definition of solution and sound validators for candidate solutions
- Should have large instance space and dynamic generation procedure to avoid memorization concerns.

PLANNING IN THE MODERN LLM ERA

Solving NL/PDDL Planning Problems

Katz et al, NeurIPS 2024
Cao et al, OWA@NeurIPS 2024
Tuisov et al, Arxiv 2024
Correa et al, NeurIPS 2025

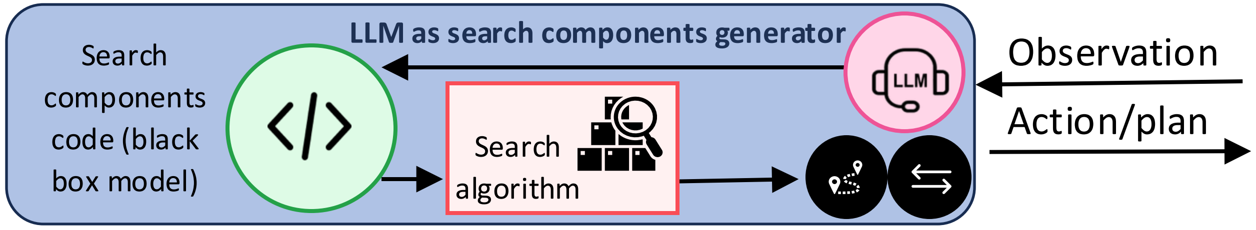

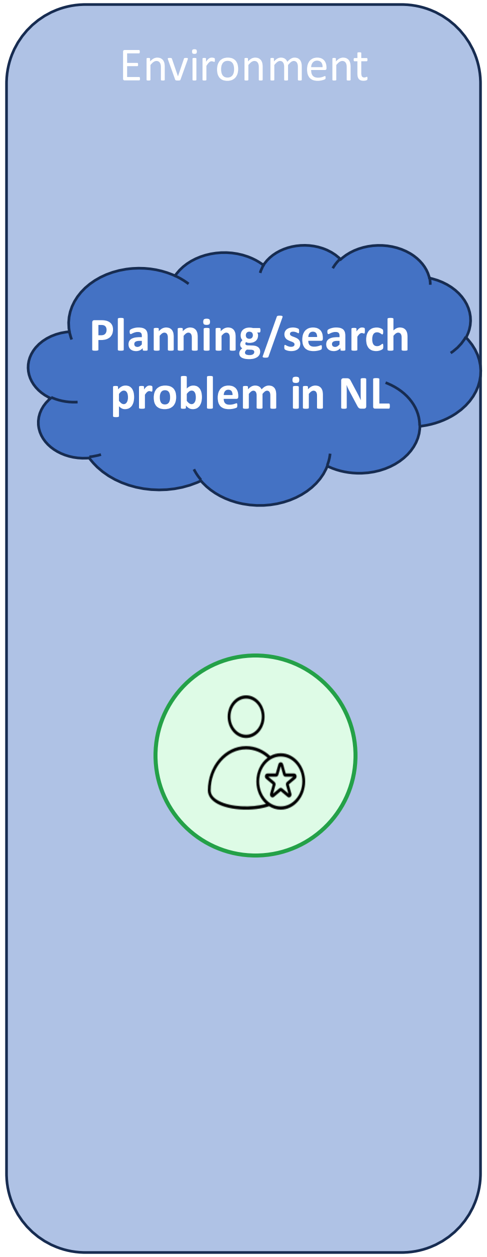

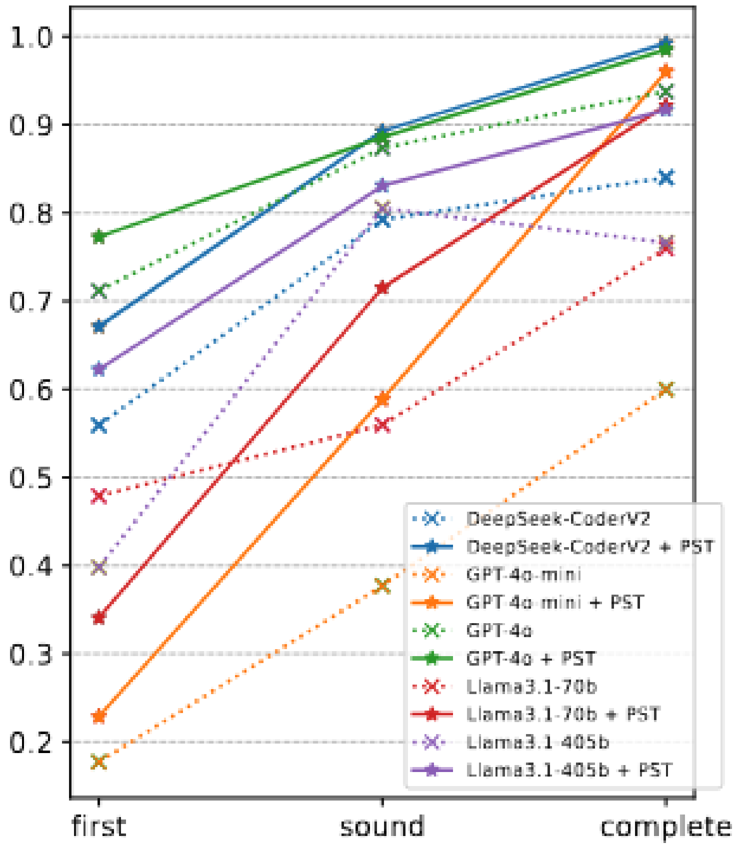

Approach	Complexity	24Game States	24Game Calls	Crossword States	Crossword Calls	BlocksWorld States	BlocksWorld Calls	PrOntoQA States	PrOntoQA Calls
IO	O(D)	0.02%	1362	4e-9%	20	0.5%	502	4%	4000
CoT	O(D)	0.02%	1362	4e-9%	20	0.5%	502	4%	4000
ReAct	O(LD)	0.07%	4086	4e-8%	200	7.8%	8032	24.6%	24K
ReWOO	O(LD)	0.07%	4086	4e-8%	200	7.8%	8032	24.6%	24K
RAP	O(TbLD)	3.3%	245K	2e-6%	12K	388%	482K	1229%	1.44M
ToT	O(bmLD)	1.6%	102K	1e-6%	5K	194%	201K	615%	600K
GoT	O(bLD)	0.3%	20K	2e-7%	1K	39%	40K	122%	120K
Reflection	O(LTD)	0.7%	68K	4e-7%	2.4K	77.6%	90K	245%	320K
LATS	O(TbLD)	3.3%	286K	2e-6%	14K	388%	562K	1229%	1.68M
ToS (ours)	O(1)	27.0%	**2.2**	3e-4%	**3.8**	125%	**3.8**	175%	**2.6**

		24 Game	PrOntoQA	Sokoban	Crossword	BlocksWorld
AutoToS	GPT-4o-mini	8.8	4.8	6.4	9.6	10.0
	GPT-4o	3.4	2.6	2.2	5.8	2.0
	Llama3.1-405b	3.4	2.0	2.6	4.0	3.2
	Llama3.1-70b	7.4	2.0	8.2	6.2	5.8
	DeepSeek-CoderV2	4.4	2.0	2.8	6.6	4.2
ToS	GPT-4	2.2	2.6	NA	3.8	3.8

Solving NL/PDDL Planning Problems

Katz et al, NeurIPS 2024
Cao et al, OWA@NeurIPS 2024
Tuisov et al, Arxiv 2024
Correa et al, NeurIPS 2025

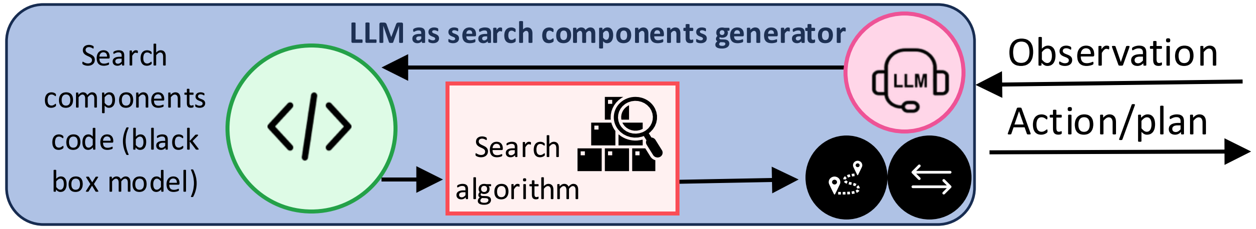

Silver et al, AAAI 2024
Hodel, BSc Thesis 2024
Stein et al, Arxiv 2025

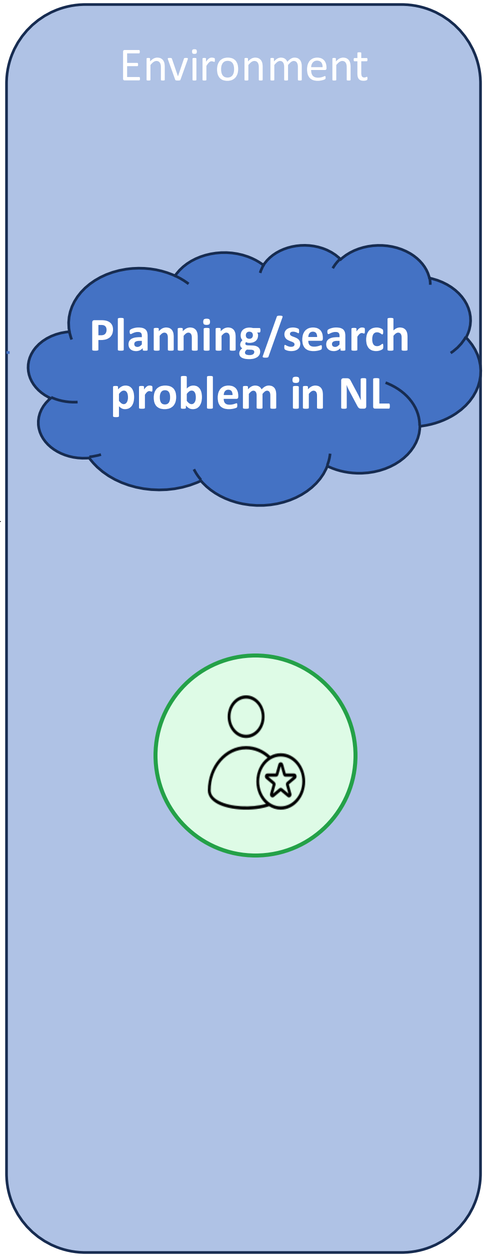

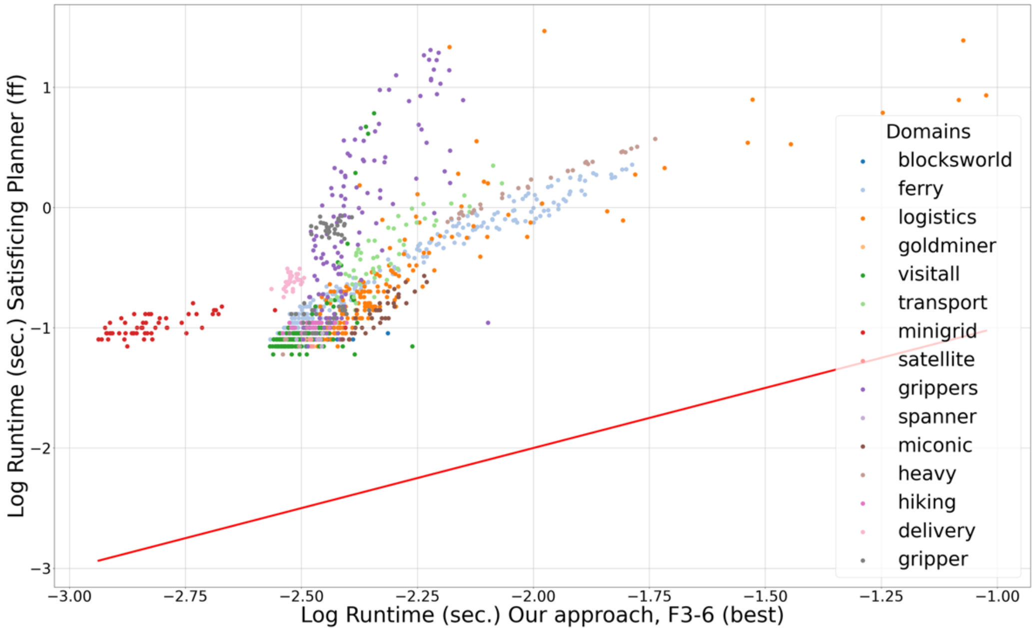

Coverage best run							Cov. symbolic			
Sil	Bas	F5-3	F3-6	-MC	-SD	-CR	lm=	ff=	lm	ff
100	100	100	100	100	100	100	0	100	0	100
100	100	100	100	100	100	100	31	100	43	100
100	100	100	100	100	100	100	15	100	40	100
100	100	100	100	100	100	100	100	100	100	100
100	0	100	100	0	0	100	100	100	100	100
32	12	100	100	0	3	12	56	100	62	100
0	15	100	100	100	100	100	15	15	41	59

Solving NL/PDDL Planning Problems

Katz et al, NeurIPS 2024
Cao et al, OWA@NeurIPS 2024
Tuisov et al, Arxiv 2024
Correa et al, NeurIPS 2025

Silver et al, AAAI 2024
Hodel, BSc Thesis 2024
Stein et al, Arxiv 2025

Guan et al, NeurIPS 2023
Gestrin et al, Arxiv 2024
Oswald et al, ICAPS 2024
Huang et al, AAAI 2025
Tantakoun et al, ACL Findings 2025

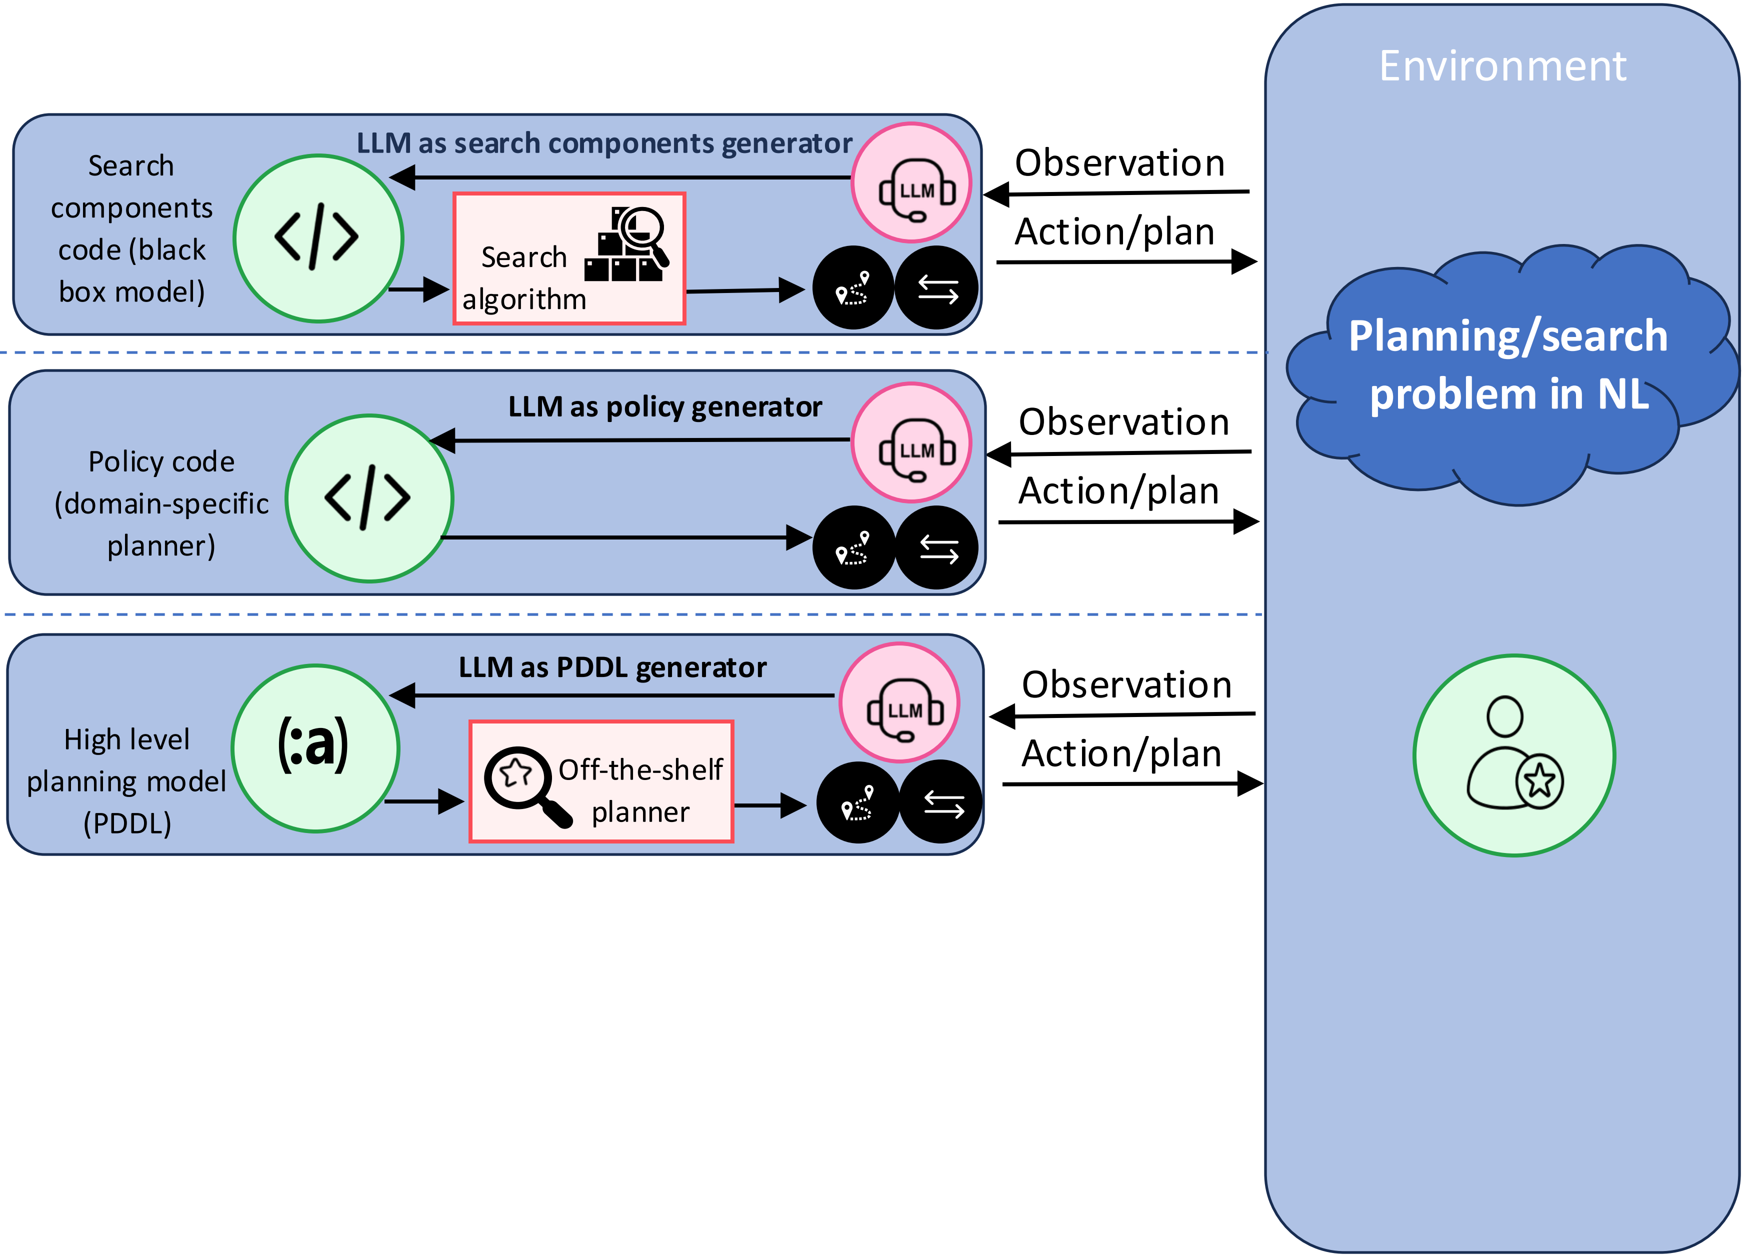

Would You Like to Know More?

ICAPS 2026 Summer School
June 22-25, 2026

Up to the Challenge?

International Planning Competition 2026

IPC 2026

International Planning Competition 2026

- Classical Tracks
- Learning Tracks
- Probabilistic Tracks
- Numeric Tracks
- HTN Tracks

International Planning Competition 2026

The International Planning Competition is organized in the context of the International Conference on Planning and Scheduling (ICAPS). It empirically evaluates state-of-the-art planning systems on a number of benchmark problems. The goals of the IPC are to promote planning research, highlight challenges in the planning community and provide new and interesting problems as benchmarks for future research.