Managing Infinite Abstractions in Numeric Pattern Database Heuristics

Markus Fritzsche[1], Daniel Gnad[1,2], Mikhail Gruntov[3], Alexander Shleyfman[4]

[1]Linköping University, [2]Heidelberg University, [3]Technion, [4]Bar-Ilan University

Setting: Simple Numeric Planning (SNP)

- Extension of classical planning with numeric variables.
- Numeric planning task: $\langle V_p \cup V_n, A, s_0, G \rangle$ with propositional (V_p) and numeric (V_n) variables.
- Actions have propositional and numeric preconditions/effects; in SNP effects are simple updates $x := x + c$.
- Finite set of numeric conditions $\bar{\Psi}_n = G_n \cup \{\psi \mid \psi \in \mathrm{pre}_n(a), a \in A\}$ where each ψ has the form $\sum_{x \in V_n} w_x^{\psi} x \geq w_0^{\psi}$.

Pattern Databases (PDBs) — Classical Planning

- Choose a pattern $P \subseteq V$ and projection $\Pi_P : S \to S|_P$ (restrict state to variables in P).
- The PDB stores exact optimal distances in the projected transition system:

$$h_{\mathrm{PDB}}(s) = h^*_P\big(\Pi_P(s)\big)$$

- Properties: precomputed by backward/regression search, admissible, and combined across patterns by **max** or additive pattern collections.

Pattern Databases (PDBs) — Numeric Planning

- Numeric variables have (in general) infinite domains, so a direct projection yields an infinite abstract transition system.
- Practical finite-fragment recipe: explore a finite fragment from the abstract initial state (expanded nodes S_E, fringe S_F) via uninformed search (UCS).
- Estimate for $s \in S_E \cup S_F$:

$$h(s) = \min_{s' \in S_F} \{ \operatorname{cost}^*(s, s') + d(s') \}, \quad d(s') = \begin{cases} 0 & s' \in S_* \\ \text{min action cost} & \text{otherwise} \end{cases}$$

- Use a min-action-cost fallback for states not in the fragment; admissibility follows from the fragment-to-concrete homomorphism.

Motivation

- Pattern Databases (PDBs) are state-of-the-art admissible heuristics in classical planning, but currently they underperform on numeric benchmarks (coverage lags behind LM-cut).
- We hypothesize two root causes: (1) abstractions explore the wrong regions of the numeric state space, and (2) many lookups are uninformed or fail, producing weak heuristic estimates.
- **Contribution: adapt numeric PDB construction and lookup strategies to improve coverage and informativeness on simple numeric planning tasks, shrinking the gap to LM-cut.**

Why Exploring the Right Regions Matters

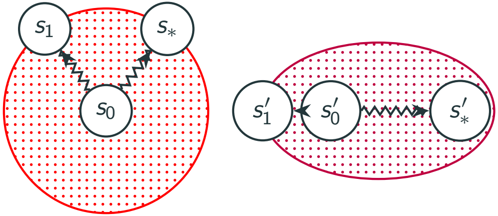

Figure 1: Illustration of heuristic values for s_0 and s'_0 in uninformed (left) vs informed (right) partial expansion.

- Exploring without an uninformed expansion (UCS) causes many expanded states in S_E to be close to fringe non-goal states (right panel), reducing the "padding" around explored regions.
- Consequence: $h(s)$ often reports distance to the fringe rather than true goal distances — compare the goal distances of s_0 and s'_0 between the left and right images.

The XYZ Framework

- **XYZ** naming: choice of heuristics for exploration, fringe, and failed lookup.
- **Guided expansion:** run A*-style exploration with admissible h_{Ex}, stopping when a generation bound B is reached.
- **Fringe refinement:** for expanded states $s \in S_E$ we use the fringe heuristic

$$\tilde{h}(s) = \min_{s' \in S_F} \{ \operatorname{cost}^*(s, s') + h_{\mathrm{Fr}}(s') \} \quad (s \in S_E),$$

  replacing the prior min-action-cost fallback and yielding more informative estimates.
- **Failed lookups:** use h_{Fl} for states whose projection was not reached ($s|_P \notin S_E \cup S_F$).
- **Notation:** an instance is written XYZ where each letter selects the variant for $h_{\mathrm{Ex}}, h_{\mathrm{Fr}}, h_{\mathrm{Fl}}$ (e.g., BBB baseline).

Coverage

	Domain	#	B	L	BBB	LLB	BBL	BLL	LLL
IPC 2023	delivery	20	2	**3**	2	2	2	2	2
	drone	20	3	3	**4**	**4**	**4**	**4**	**4**
	expedition	20	5	**6**	**6**	**6**	**6**	**6**	**6**
	farmland-ipc23	15	4	**15**	8	**15**	**15**	**15**	**15**
	hydropower	20	9	**11**	9	9	8	10	10
	mprime	20	6	**15**	12	12	12	12	12
	rover-ipc23	20	**4**	**4**	**4**	**4**	**4**	**4**	**4**
	sailing-ipc23	20	0	**8**	0	0	0	0	0
	sugar	20	2	**12**	3	3	3	3	3
	zenotravel-ipc23	20	6	**8**	6	6	6	6	6
from literature	counters	20	3	**5**	**5**	**5**	**5**	**5**	**5**
	counters-sym	11	2	**11**	8	8	8	8	9
	depots	20	4	**7**	**7**	**7**	**7**	**7**	**7**
	depots-sym	20	4	**7**	6	6	6	6	6
	farmland	30	12	**30**	**30**	**30**	26	**30**	**30**
	fn-counters-small	8	6	**7**	**7**	**7**	**7**	**7**	**7**
	forestfire	20	10	**11**	10	10	10	10	10
	minecraft-pogo	20	14	5	**18**	17	**18**	17	17
	minecraft-sword	20	**20**	9	**20**	**20**	**20**	**20**	**20**
	petri-net	20	2	8	**9**	8	8	**9**	**9**
	plant-watering	63	**63**	**63**	**63**	**63**	**63**	**63**	**63**
	rover-unit	20	4	**7**	6	6	6	6	6
	sailing	40	10	**40**	15	18	15	17	17
	satellite	20	1	**2**	**2**	1	1	**2**	**2**
	zenotravel	23	6	**13**	10	9	10	10	10
	others	72	**1**	**1**	**1**	**1**	**1**	**1**	**1**
	$\sum$	622	203	**311**	253	269	272	280	281

Search vs Total Time

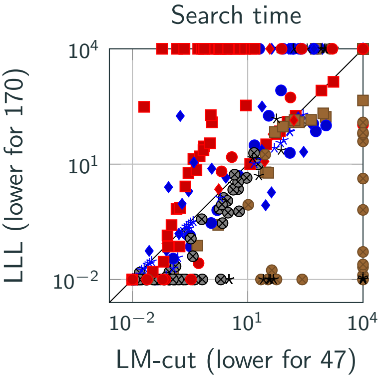

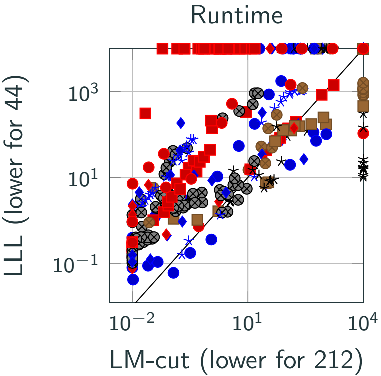

Takeaways

- Combining numeric PDB generation with informed heuristics to explore the right states and increase informedness using a failed lookup heuristic narrows the gap between LM-cut and PDBs.
- However, the gap is not closed yet; further research is required to finally close it.