

Plausibility-Based Heuristics for Latent Space Classical Planning

Yuta Takata¹ Alex Fukunaga²

¹ Fabalis ² The University of Tokyo

Abstract

Recent work on LatPlan has shown that it is possible to learn models for domain-independent classical planners from unlabeled image data. Although PDDL models acquired by LatPlan can be solved using standard PDDL planners, the resulting latent-space plan may be invalid with respect to the underlying, ground-truth domain (e.g., the latent-space plan may include hallucinatory/invalid states). We propose Plausibility-Based Heuristics, which are domain-independent plausibility metrics which can be computed for each state evaluated during search and uses as a heuristic function for best-first search. We show that PBH significantly increases the number of valid found plans on image-based tile puzzle and Towers of Hanoi domains.

1 Introduction

Automated acquisition of domain models for planners from subsymbolic data (e.g., images taken by cameras) poses an important challenge for autonomous systems that must operate in new, unknown environments. Recent work has shown that it is possible to learn domain models for domain-independent classical planners from unlabeled, unannotated pairs of images representing the state of the world before and after actions are executed (Asai and Fukunaga 2018; Asai et al. 2022). LatPlan uses deep learning to learn a *latent* propositional representation of the domain, which can be output as a STRIPS PDDL model. These latent propositions as well as the action model are learned from scratch, without domain-specific prior knowledge, i.e., LatPlan grounds the symbols (Steels 2008) and learns an action model with respect to the latent propositions. Problem instances for these learned models can be solved using standard search-based planners using standard domain-independent heuristics.

However, an inherent issue with learned models is that propositions and actions in the learned, latent PDDL model do not necessarily correspond to those of the underlying problem. Thus, a significant problem is that although standard search algorithms such as Fast Downward can find solutions which are correct with respect to the latent PDDL model, these solutions may not be valid with respect to the underlying problem. For example, Figure 1 shows an invalid plan found by LatPlan for the image-based MNIST-8 puzzle.



Figure 1: Visualizations of 2 latent plans for MNIST-8 puzzle (start: left \rightarrow goal:right). The top plan is valid and optimal. The bottom plan is invalid because the 7th image has two “0”s and an unclear digit which could be interpreted as a 2, 3, or 6.

Previous work focused on improving the architecture and learning objectives so that models learned by LatPlan accurately correspond to the underlying domains (Asai and Kajino 2019; Asai et al. 2022). Despite these efforts, the current implementation of LatPlan has limitations. Although LatPlan can successfully solve image-based 15-puzzles, it has not been able to find valid solutions for a 4-disk Towers of Hanoi problem, even though it is much “simpler” than the 15-puzzle with respect to the size of the search space of the underlying problem (Asai et al. 2022, Appendix M).

In this paper, we investigate a complementary approach. Instead of the learning component, we focus on the search component of LatPlan. Previous work relied on using unmodified, standard planner (Fast Downward) to solve the PDDL instances generated by LatPlan. We improve the ability to find valid plans by making LatPlan-specific (but domain-independent) extensions to the search component.

We propose Plausibility-Based Heuristics (PBH), which uses the learned LatPlan model during the search to evaluate the *plausibility* of a state, i.e., the likelihood that the state is valid with respect to the underlying domain. The plausibility measure can be used as the heuristic evaluation function h for best-first search (e.g., A^* , GBFS). Instead of guiding search according to standard distance-based heuristics, we prioritize search according to plausibility, as finding a valid plan takes precedence over finding short but invalid plans. This is somewhat similar to the idea of using novelty and width-based exploration in black-box domains where standard heuristics do not lead to a solution (Lipovet-

zky, Ramírez, and Geffner 2015).

As LatPlan is a domain-independent, image-based domain learner/planner, an important goal of this work is to identify useful, *domain-independent* plausibility metrics which are useful on some subset of LatPlan domains. We propose image-based plausibility measures which can be computed *during search* from the propositional latent state.

The remainder of the paper is structured as follows. Section 2 reviews LatPlan. Section 3 discusses the validity of latent space plans, and show that the ability to find valid plans is highly sensitive to LatPlan hyperparameter settings. Section 4 proposes metrics for the plausibility of states in latent search spaces, and their use as heuristic functions for best-first search. Section 5 experimentally evaluates PBH and shows that it significantly improves the number of valid solutions found by the search algorithm. Section 6 concludes with a discussion of our results.

2 Preliminaries and Background

A STRIPS planning task (Fikes and Nilsson 1971) is defined by a 4-tuple $T = \langle P, A, I, G \rangle$. P is a set of propositions. A is a set of actions. $I \subset P$ represents the initial state, and $G \subset P$ is the goal condition. A state is represented by a subset of P , and applying an action to a state adds some propositions and removes some propositions in the state. Each action $a \in A$ is composed of three subsets of P , $\langle pre(a), add(a), del(a) \rangle$, which are the *preconditions*, *add effects*, and *delete effects*. An action a is applicable to a state S iff $pre(a) \subseteq S$. By applying a to S , propositions in S change to $S(a) = ((S \setminus del(a)) \cup add(a))$. For a sequence of actions $\pi = (a_0, \dots, a_n)$, we use $S(\pi)$ to denote $((((S \setminus del(a_0)) \cup add(a_0)) \setminus del(a_1)) \cup \dots) \cup add(a_n)$. A solution to T is a sequence of actions that transforms I to a state S satisfying $G \subseteq S$. A feasible solution, i.e., a plan, is a sequence of actions $\pi = (a_0, \dots, a_n)$ that satisfies (i) $\forall i, pre(a_i) \subseteq I((a_0, \dots, a_{i-1}))$ and (ii) $G \subseteq I(\pi)$.

Latent-space planning using LatPlan Given only an unlabeled set of image pairs showing a subset of transitions allowed in the environment (training inputs), LatPlan learns a complete propositional PDDL action model of the environment. Later, when a pair of images representing the initial and the goal states (planning inputs) is given, LatPlan finds a plan to the goal state in a symbolic latent space and returns a visualized plan execution.

Assume that the environment poses a STRIPS planning problem. Assume that images (e.g., images taken using a camera) are represented as a vector of integers. Let \mathcal{O} be the space of observations, where an *observation* $o \in \mathcal{O}$ is an image representing a state of the environment.

The input to LatPlan is \mathcal{X} , a set of pairs of images. Each image pair $(\mathbf{x}_{i,before}, \mathbf{x}_{i,after}) \in \mathcal{X}$ represents a transition from an observation $\mathbf{x}_{i,before}$ to another observation $\mathbf{x}_{i,after}$ caused by an unknown action. Based only on \mathcal{X} , LatPlan learns two functions: (1) $Encode(o) : \mathcal{O} \rightarrow (0, 1)^N$, which maps an image o to a propositional vector, and (2) $Decode(o) : (0, 1)^N \rightarrow \mathcal{O}$, which maps a propositional vector to an image (Fig. 2). N is a hyperparameter for the number of latent propositions. LatPlan learns a Discrete Vari-




Figure 2: State Autoencoder

ational Autoencoder (DVAE) which perform both *Encode* and *Decode*. Unlike a standard autoencoder (Hinton and Salakhutdinov 2006) where unit activations are continuous, the DVAE uses a binary (propositional) activation (Jang, Gu, and Poole 2017; Maddison, Mnih, and Teh 2017). The propositional output/input of *Encode/Decode* is called the *latent representation*.

Using *Encode* and *Decode*, and \mathcal{X} , LatPlan generates a *latent domain* $D^l = \langle P^l, A^l \rangle$, where P^l is a set of (latent) propositions and A^l is a set of actions in latent space. Each action $a^l \in A^l$ is composed of three subsets of P^l , $\langle pre(a^l), add(a^l), del(a^l) \rangle$, and action a^l is applicable to a state S^l iff $pre(a^l) \subseteq S^l$. By applying a^l to S^l , the propositions in S^l change to $S^l(a^l) = ((S^l \setminus del(a^l)) \cup add(a^l))$.

Let o_I be an image of a start state in the environment (e.g., jumbled 8-puzzle configuration). Let o_G be an image representing a desired state of the environment (e.g., solved 8-puzzle state). A latent STRIPS planning task $T^l = \langle P^l, A^l, I^l, G^l \rangle$, where $I^l = Encode(o_I)$ and $G^l = Encode(o_G)$. A solution $\pi^l = (a_0^l, \dots, a_n^l)$ to T^l is a sequence of actions that transforms I^l to a state S^l that satisfies $G^l \subseteq S^l$. As T^l is just a standard STRIPS task, T^l can be solved using any planner which takes PDDL input.

The latent space plan π^l is a sequence of actions in the latent representation, and can not be directly interpreted and executed by the agent. Therefore, LatPlan uses its learned action model and *Decode* to map the latent states visited by π^l to a sequence of images $(\mathbf{y}_0, \dots, \mathbf{y}_n)$, called the *visualization* of π^l , which can be interpreted and executed by the agent, e.g., a human or robot can move the sliding tile puzzle tiles according to the sequence of images.

3 Validity of Latent Space Plans

An inherent issue with domain learning systems is that the learned domain may differ from the underlying, canonical domain. The latent domain D^l learned by LatPlan can contain many states/actions which do not correspond to any valid canonical state/actions. For example, although the MNIST-8 puzzle corresponds to a canonical 8-puzzle with a branching factor b of at most 4, the search space of the learned MNIST-8 latent space domain used in the Section 5 experiments has a maximum (over all expanded nodes) b of 124.9 (mean of 20 instances). Thus, the search algorithm may incorrectly find latent search space plans which can not be executed.

We define the following notion of validity of a plan generated by LatPlan, relative to some target STRIPS task T which the latent domain learned by LatPlan is intended to correspond to.

Definition 1. We say that π^l is *valid* with respect to a canonical STRIPS task $T = \langle P, A, I, G \rangle$ if: for each image \mathbf{y}_i in

the visualization of π^l except the goal, \mathbf{y}_i corresponds to a state $S_i \subset P$, \mathbf{y}_{i+1} corresponds to a valid state $S_{i+1} \subset P$, and $\exists a \in A$ such that $S_i(a) = S_{i+1}$.

Similarly, length optimality of π^l wrto T^l does not necessarily imply optimality wrto the underlying canonical task T , so we define the following condition which is satisfied when π^l is optimal wrto T .

Definition 2. We say that π^l is *c-optimal* relative to T if π^l is valid and $|\pi^l|$ is equal to that of the minimal-length plan for T .

In Figure 1, the top plan is valid and *c-optimal*. The bottom plan is invalid, because the 7th image from the left has two ‘‘O’’s and an unclear digit.

In the experiments below, the LatPlan input \mathcal{X} is generated by a domain-specific generator which has full knowledge of a canonical STRIPS domain (e.g., 8-puzzle), and outputs a set of valid sample image pairs \mathcal{X} for that domain. The start/goal images o_I and o_G are also generated using this generator. Therefore, for each LatPlan domain, plan validity for a latent space plan π^l is checked using a domain-specific validator.

Since our experiments use validation code to test latent plan validity, one might wonder whether we could simply use this validator during search to prune invalid states. However, unlike domain-independent validators for PDDL domains such as VAL (Howey, Long, and Fox 2004), the LatPlan image-based plan validators are carefully tuned, hand-coded domain-specific tools (see Asai et al. (2022, Appendix L.1)), which use extensive domain knowledge e.g., for MNIST-8, the knowledge that a valid state image is composed of 8 unique tiles, each of which corresponds to a portion of the unscrambled image. If we have such knowledge, then instead of a weak learning model such as LatPlan, a model learning method which can exploit such prior knowledge would probably be a better alternative. Therefore, such domain-dependent image validators are unlikely to be available in domains for which LatPlan is a viable approach.

3.1 Plan Validity and LatPlan Model Learning Hyperparameters

LatPlan has 4 key hyperparameters related to model learning: (1) The number of latent propositional variables $N = |P^l|$ (2) ϵ , the parameter for the Bernoulli(ϵ) distribution used as the prior for the latent random variables in the Binary-Concrete VAE (Asai et al. 2022, Sec 5.4). (3,4) the regularization parameters β_1 and β_3 for the KL-divergence used in the Evidence Lower Bound (ELBO) loss function. The best values for these hyperparameters are domain-dependent. Asai et al (2022) showed that the accuracy of the learned models depended significantly on hyperparameter settings.

We investigate how the ability to find valid solutions to latent space planning problems is affected by the hyperparameters used during the latent model learning. We ran LatPlan on the MNIST-8 domain, with 60 different hyperparameter value combinations (all combinations of $(N, \epsilon, \beta_1, \beta_3)$ for $N \in \{50, 100, 300\}$, $p \in \{0.1, 0.5\}$, $\beta_1 \in \{1, 10\}$, $\beta_3 \in$

$\{1, 10, 100, 1000, 10000\}$) to learn 60 different latent domains, all using the AMA4+ action model (Asai et al. 2022). We then solved the 20 instances for each of these latent domains with Fast Downward using A^* with a blind heuristic.

The search results (for each domain, the number of instances for which solutions were found/valid/*c-optimal*) are shown in Table 1. With the 3 best hyperparameters ($N = 300, \epsilon = 0.1, \beta_1 = 10, \beta_3 = 1$), ($N = 300, \epsilon = 0.1, \beta_1 = 10, \beta_3 = 10$) and ($N = 300, \epsilon = 0.5, \beta_1 = 1, \beta_3 = 1$), optimal solutions are found for all 20 instances. However, for the other 57 hyperparameter combinations, there are failures to find any plans (regardless of validity), failures to find valid plans, and/or failures to find optimal plans. Thus, hyperparameter settings for model learning have a significant impact on the ability of LatPlan to find valid plans.

Note that although previous work by showed that LatPlan can find valid plans on all of the domains tested, the results shown in this previous work (e.g., (Asai et al. 2022, Table 11.1)) are for the hyperparameter configurations that yielded the highest number of valid plans *for each domain*, i.e., the hyperparameters were tuned for each domain (Asai et al. 2022, p.1642).

In general, we do not know *a priori* the hyperparameter values necessary to facilitate finding valid plans. Furthermore, hyperparameter tuning for this purpose is not possible in general, as that would require the ability to evaluate whether the plans found by the planner are valid. While we use a domain-specific ground-truth based validator in our experiments, a ground-truth domain model is not accessible in many situations where learning systems such as LatPlan are necessary. Therefore, developing a method to compensate for suboptimal model learning hyperparameter settings by improving the search component offers a significant opportunity to improve the robustness of LatPlan.

4 Plausibility-Based Heuristics (PBH)

We propose an approach to estimating the plausibility of a state based on invariance checking. We assume that valid latent states satisfy some invariant property of the domain. Consider the image-based sliding tile-puzzle or blocks world domains. If we compare, for example, start state image with the image of a valid state, although corresponding pixels differ because parts of the scene are moved by actions, we would expect that some overall statistical properties of the pixels in the images are invariant – the locations of the pixels are changing (tile t moves from location i to j), but tiles are not being deformed/created/destroyed by the actions.

Among the domains to which LatPlan has been applied so far, all of the sliding-tiles variants, blocks world and Towers of Hanoi satisfies this assumption. In contrast, the LightsOut domain does not, as each action significantly changes ratio of white/black pixels. The color version of Sokoban used in (Asai et al. 2022) slightly violates this assumption, as empty goal location cells have color c_g , differently from the other unoccupied floor cells (color c_f), so when the agent pushes a box (color c_b) from location b to a goal cell g , the distribution of colors before the move ($color(b) = c_b, color(g) = c_g$) and after the move ($color(b) = c_f, color(g) = c_b$) differ.

		0.1			0.1			0.5			0.5			0.5					
		50			100			300			50			100			300		
β_3	β_1	found	valid	c-optimal	found	valid	c-optimal	found	valid	c-optimal	found	valid	c-optimal	found	valid	c-optimal	found	valid	c-optimal
1	1	20	9	9	20	19	19	20	20	19	20	13	13	20	19	19	20	20	20
	10	20	4	4	3	3	3	20	20	20	20	7	7	2	0	0	0	0	0
10	1	20	8	8	17	17	17	20	20	19	20	4	4	18	16	16	18	18	13
	10	20	0	0	20	11	11	20	20	20	20	6	6	1	0	0	0	0	0
100	1	20	15	15	20	9	8	11	11	0	20	9	9	20	19	19	0	0	0
	10	20	4	4	20	5	5	0	0	0	20	10	10	0	0	0	0	0	0
1000	1	20	2	2	0	0	0	0	0	0	20	11	11	0	0	0	0	0	0
	10	20	4	4	0	0	0	5	5	0	20	11	11	0	0	0	0	0	0
10000	1	12	0	0	6	0	0	0	0	0	5	0	0	7	2	2	0	0	0
	10	12	1	1	4	0	0	0	0	0	12	0	0	7	0	0	0	0	0

Table 1: Hyperparameter sweep study of LatPlan (AMA4+ action model) on MNIST 8-puzzle. A latent domain was learned for each all combinations of $(N, \epsilon, \beta_1, \beta_3)$ for $N \in \{50, 100, 300\}$, $\epsilon \in \{0.1, 0.5\}$, $\beta_1 \in \{1, 10\}$, $\beta_3 \in \{1, 10, 100, 1000, 10000\}$, where $N = |P^l|$, the number of latent propositional variables, ϵ is the parameter for the Bernoulli(ϵ) distribution used as the prior for the latent random variables in the Binary-Concrete VAE, and β_1 and β_3 are the regularization parameters for the KL-divergence used in the ELBO loss function. For each domain, 20 instances were solved using Fast Downward (A^* with blind heuristic).

Our invariant-based approach to computing the plausibility of a latent space state s^l is as follows: Let r^l be a *reference state*, a latent space state which is known/assumed to correspond to a valid state in the canonical domain.

We apply a similarity metric to (s^l, r^l) which seeks to capture an invariance property, such that a latent state which is more similar to r^l is more likely to correspond to a valid canonical state than a state which is less similar to r^l .

Regarding the choice of the reference state r^l , any image which is believed to be valid image (e.g., any image in $\mathcal{X} \cup o_I \cup o_G$) is a candidate to be used as r^l , but the effectiveness of the similarity metric may depend on the specific choice of r^l . In the experiments below, we use the reference state $r^l = \text{Encode}(o_G)$, which is the propositional latent space state corresponding to the goal image o_G . Investigating other choices for r^l is future work.

Since latent propositions P^L do not correspond straightforwardly to propositions P of the canonical problem T , it is not obvious how to apply the idea of invariant-based plausibility to latent state vectors. Preliminary attempts to designing similarity metrics for latent state vectors have not yet been successful. Thus, instead of directly using the latent proposition vector s^l to determine the plausibility of s^l , we can instead use the image corresponding to s^l .

To obtain the image corresponding to latent state s^l , we use the *Decode* function learned by LatPlan to convert s^l to an image $M_{s^l} = \text{Decode}(s^l)$, where M_{s^l} is an array of pixels (integers in $[0, A]$). We then apply simple image analysis techniques to M_{s^l} in order to determine the plausibility of s^l . Let $H_{M_{s^l}}$ denote the *histogram* of an image M_{s^l} . For simplicity, we write H_{s^l} to mean $H_{M_{s^l}}$. H_{s^l} is an array of bins, where bin $H[b]$ is the frequency of pixels in $[[A/B] \times b, [A/B] \times (b + 1)]$ where B is the # of bins.

We use the standard χ^2 difference and *KL*-divergence to compare s^l and the reference state r :

$$p_{\chi^2}(s^l) = - \sum_b \frac{(H_r[b] - H_{s^l}[b])^2}{H_r[b]} \quad (\chi\text{-squared difference})$$

$$p_{KL}(s^l) = - \sum_b H_r[b] \log \left(\frac{H_r[b]}{H_{s^l}[b]} \right) \quad (KL\text{-divergence})$$

The plausibility metrics proposed above can be used to guide a best-first algorithm such as A^* (Hart, Nilsson, and Raphael 1968) or Greedy Best-First Search (GBFS). A^* uses a node evaluation function $f(n) = g(n) + h(n)$, where $g(n)$ is the cost to reach node n from the initial node and $h(n)$ is the estimated cost from n to a goal. GBFS uses the evaluation function $f(n) = h(n)$. The heuristic functions h_{χ^2} and h_{KL} directly use the corresponding plausibility metrics. We scale h_{χ^2} and h_{KL} to an integer (as we implement these in FastDownward, which uses integer-valued heuristics).

$$h_{\chi^2}(s^l) = \lfloor -p_{\chi^2}(s^l) \rfloor \quad (1)$$

$$h_{KL}(s^l) = \lfloor -p_{KL}(s^l) \rfloor \quad (2)$$

Our current implementation of PBH extends the state evaluation procedure of the Fast Downward planner so that when evaluating s^l , it sends a request to an external server process which uses LatPlan to convert s^l to an image M_{s^l} and then uses OpenCV to compute an image-based plausibility metric based on M_{s^l} . This evaluation accounts for the vast majority of the runtime consumed by the search algorithm. Due to the slow state evaluation rate, solving a latent PDDL instance using the current unoptimized implementation of PBH is much slower than searching the same PDDL instance using unmodified Fast Downward. Since the focus of this paper is on improving the ability to find valid plans, optimization of the evaluation procedure, as well as search algorithms which do not evaluate plausibility for every state are directions for future work.

5 Experiments

We evaluate the search performance of PBH. We implemented PBH by extending the Fast Downward search code. We compared: (1) Baseline latent space search (A^* , GBFS). using standard heuristics (the Blind ($h = 1$) heuristic, the admissible LMCut (Helmert and Domshlak 2009) and Merge-and-Shrink (Helmert et al. 2014) heuristics, as well as the non-admissible FF heuristic (Hoffmann and Nebel 2001)) and (2) A^* and GBFS using plausibility-based heuristics h_{χ^2} and h_{KL} (bin size $B = 10$).

		Hanoi(4,4)				MNIST-8			
Search Alg.	Heuristic	found	valid	c -optimal	length	found	valid	c -optimal	length
Baselines									
A^*	blind	20	1	1	6.45	20	9	9	7
A^*	lmc	20	1	1	6.45	20	6	6	7
A^*	M&S	20	2	2	6.45	20	7	7	7
A^*	FF	20	0	0	7.85	20	5	5	8.1
GBFS	blind	20	1	1	6.45	20	9	9	7
GBFS	lmc	20	0	0	10.9	20	0	0	45
GBFS	M&S	20	0	0	6.8	20	4	4	8.4
GBFS	FF	20	0	0	28.9	20	0	0	57.15
PBH: Plausibility-Based Heuristics									
A^*	h_{χ^2}	20	15	8	7.35	20	20	20	7
GBFS	h_{χ^2}	20	20	8	7.7	20	20	6	21.5
A^*	h_{KL}	20	19	8	7.65	20	20	20	7
GBFS	h_{KL}	20	20	4	8.75	20	20	9	30.2

Table 2: Search Results: Number of found/valid/ c -optimal instances (out of 20) and average plan length (optimal length = 7 for both domains).

Comparison of ability to find valid plans We compare the search algorithms on the following domains (all from Asai et al. (2022), all using the AMA4+ action model): (1) **MNIST-8**: 3x3 sliding tile puzzle where the tiles have handwritten digits 0-8 from the MNIST dataset, (2) **Hanoi(4,4)**: Towers of Hanoi (4 towers, 4 disks), (3) **Mandril-15**: 4x4 sliding tile puzzle where the goal is to unscramble the well-known Mandril face image. The learning hyperparameters were ($N = 50, \epsilon = 0.1, \beta_1 = 1, \beta_3 = 1$) for MNIST-8, ($N = 50, \epsilon = 0.1, \beta_1 = 1, \beta_3 = 100$) for Hanoi(4,4), ($N = 100, \epsilon = 0.1, \beta_1 = 1, \beta_3 = 1$) for Mandril-15.

First, we compared the search performance on the MNIST-8 and Hanoi(4,4) domains (20 instances each) Table 2 shows the results (“found”: latent space plan found, “valid”: plan is valid, “ c -optimal”: plan is c -optimal). Although baseline A^* and GBFS find latent space plans, many of them are not valid. Note that the average plan lengths include invalid plans which can be either longer or shorter than valid plans, so the average lengths differ from the optimal valid length (7) even for A^* using admissible heuristics (blind, lmc, M&S).

PBH (both h_{χ^2} and h_{KL}) found significantly more valid plans than the baselines. In particular, PBH successfully, reliably finds valid solutions for Hanoi(4,4), on which previous work had failed to find any valid solutions (Asai et al. 2022).

On Mandril-15, we ran a subset of the algorithm configurations, as it requires significantly more runtime than Hanoi(4,4) and MNIST-8. As explained in Section 4 the runtime of our current unoptimized implementation of PBH is very slow (state evaluation rate is less than 100 nodes per second for Mandril-15, and the average runtime is 8.3 hours). The results were as follows: (a) A^* with blind heuristic: 20 found, 15 valid, 15 c -optimal, (b) A^* with landmark cut: 20 found, 14 valid, 14 c -optimal, (c) A^* with merge-and-shrink: 20 found, 15 valid, 15 c -optimal, (d) GBFS with

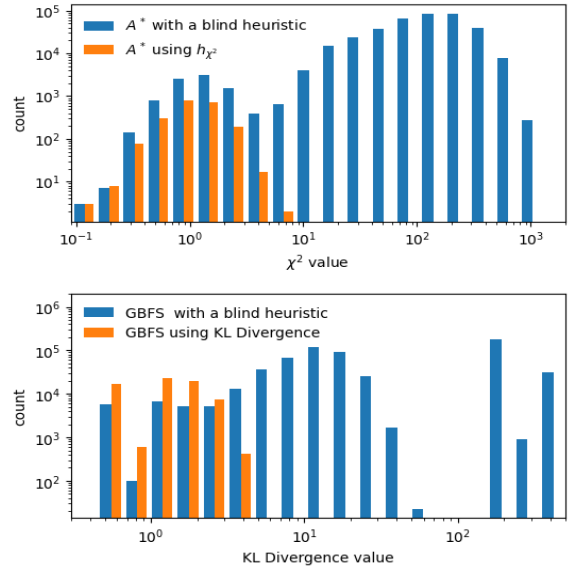


Figure 3: Top: h_{χ^2} values of states expanded during search on MNIST-8 puzzle (20 instances). Bottom: h_{KL} values of states expanded during search on Hanoi(4,4) (20 instances).

blind heuristic: 20 found, 15 valid, 15 c -optimal, (e) A^* with h_{KL} : 20 found, 20 valid, 20 c -optimal. Thus, PBH significantly improves the ability to find valid solutions on Mandril-15.

Search Behavior In order to compare the search behavior of PGH vs. baseline search, we compare the plausibility scores of all states expanded the search on all 20 instances of MNIST-8 and Hanoi(4,4) by a baseline search ($A^*/$ blind for MNIST, GBFS/blind for Hanoi vs. PBH (A^*/h_{χ^2} for MNIST, GBFS/ h_{KL} for Hanoi). For both baseline and PBH, we stored the plausibility values of the expanded states (the baseline computes but ignores them during search). Figure 3 shows that the plausibility values of the states expanded by the baseline are much more widely distributed than those of the states expanded by h_{χ^2} and h_{KL} , confirming that search with PBH focuses on highly plausible states.

6 Conclusion

In a system such as LatPlan which is intended to operate autonomously in unknown environments without a teacher who can identify whether a plan is valid or invalid, flawed learned models which lead to invalid plans poses a significant challenge to the utility and robustness of the system. This work showed that even if a learned, domain model is flawed, the search algorithm can compensate by avoiding suspicious (low plausibility) states, resulting in significantly improved ability to find valid plans. Although PBH was evaluated on LatPlan, we believe the overall approach may be useful in similar planning systems where learned problem representations have spurious states.

Our approach is domain-independent. The plausibility metrics do not use any domain-specific notions, e.g., there are no “tile/disc detectors”. Instead, we rely on properties

common to some subset of LatPlan domains, i.e., invariance properties of the histogram of the image representation of valid states, and our plausibility metrics successfully improve search results in multiple domains. Developing plausibility metrics for other types of domains, e.g., LightsOut, is a direction for future work.

This paper evaluated using plausibility metrics as heuristic evaluation functions to guide search instead of standard distance-based heuristics. In some cases, this may result in degraded search performance or solution quality compared to distance-based heuristics. However, there are many situations where an invalid plan may be no better than finding a poor quality solution or failing to find a plan altogether, so in this paper we focus on prioritizing validity. Future work will investigate combining plausibility with distance based heuristics by using plausibility as a tie-breaking criterion (Asai and Fukunaga 2017) or by considering both plausibility and distance in the heuristic evaluation function in order to improve solution quality and search performance.

This work opens avenues for future work on search strategies as complementary approaches to improvements to the learning component in domain-learning systems. This paper explored the use of plausibility metrics to evaluate the plausibility of a single state during search. Extending the notion of plausibility to paths, e.g., prioritizing paths according to plausibility, may further improve search behavior. In this paper, we focused on improving the ability to find valid plans by straightforwardly applying plausibility as a heuristic evaluation function in standard A^* and GBFS search. However, PBH incurs significant runtime costs while computing the plausibility of a state. Development of more sophisticated search strategies which seek to reduce the number of plausibility computations while still avoiding invalid plans is a direction for future work.

Acknowledgments

Thanks to Masataro Asai for helpful discussions and assistance with the LatPlan source code and sample model files. Thanks to the anonymous reviewers for helpful comments.

References

- Asai, M.; and Fukunaga, A. 2017. Exploration among and within Plateaus in Greedy Best-First Search. In Barbulescu, L.; Frank, J.; Mausam; and Smith, S. F., eds., *Proceedings of the Twenty-Seventh International Conference on Automated Planning and Scheduling, ICAPS 2017, Pittsburgh, Pennsylvania, USA, June 18-23, 2017*, 11–19. AAAI Press.
- Asai, M.; and Fukunaga, A. 2018. Classical Planning in Deep Latent Space: Bridging the Subsymbolic-Symbolic Boundary. In *AAAI Conference on Artificial Intelligence*, 6094–6101. AAAI Press.
- Asai, M.; and Kajino, H. 2019. Towards Stable Symbol Grounding with Zero-Suppressed State AutoEncoder. In Benton, J.; Lipovetzky, N.; Onaindia, E.; Smith, D. E.; and Srivastava, S., eds., *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling, ICAPS 2018, Berkeley, CA, USA, July 11-15, 2019*, 592–600. AAAI Press.
- Asai, M.; Kajino, H.; Fukunaga, A.; and Muise, C. 2022. Classical Planning in Deep Latent Space. *J. Artif. Intell. Res.*, 74: 1599–1686.
- Fikes, R.; and Nilsson, N. J. 1971. STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence*, 2(3): 189–208.
- Hart, P. E.; Nilsson, N. J.; and Raphael, B. 1968. A Formal Basis for the Heuristic Determination of Minimum Cost Paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2): 100–107.
- Helmert, M.; and Domshlak, C. 2009. Landmarks, Critical Paths and Abstractions: What’s the Difference Anyway? In *International Conference on Automated Planning and Scheduling (ICAPS)*, 162–169. AAAI.
- Helmert, M.; Haslum, P.; Hoffmann, J.; and Nissim, R. 2014. Merge-and-Shrink Abstraction: A Method for Generating Lower Bounds in Factored State Spaces. *Journal of the ACM*, 61(3): 16:1–16:63.
- Hinton, G. E.; and Salakhutdinov, R. R. 2006. Reducing the Dimensionality of Data with Neural Networks. *Science*, 313(5786): 504–507.
- Hoffmann, J.; and Nebel, B. 2001. The FF Planning System: Fast Plan Generation Through Heuristic Search. *Journal of Artificial Intelligence Research*, 14: 253–302.
- Howey, R.; Long, D.; and Fox, M. 2004. VAL: Automatic Plan Validation, Continuous Effects and Mixed Initiative Planning Using PDDL. In *16th IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2004), 15-17 November 2004, Boca Raton, FL, USA*, 294–301. IEEE Computer Society.
- Jang, E.; Gu, S.; and Poole, B. 2017. Categorical Reparameterization with Gumbel-Softmax. In *International Conference on Learning Representations (ICLR)*. OpenReview.net.
- Lipovetzky, N.; Ramírez, M.; and Geffner, H. 2015. Classical Planning with Simulators: Results on the Atari Video Games. In Yang, Q.; and Wooldridge, M. J., eds., *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, 1610–1616. AAAI Press.
- Maddison, C. J.; Mnih, A.; and Teh, Y. W. 2017. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. In *International Conference on Learning Representations (ICLR)*. OpenReview.net.
- Steels, L. 2008. The Symbol Grounding Problem has been Solved. So What’s Next? In de Vega, M.; Glenberg, A.; and Graesser, A., eds., *Symbols and Embodiment*. Oxford University Press.