Return to Tradition: Learning Reliable Heuristics with Classical Machine Learning

Dillon Z. Chen[1,2]**, Felipe Trevizan**[2]**, Sylvie Thiébaux**[1,2]

[1]LAAS-CNRS, Université de Toulouse
[2]Australian National University
dillon.chen@laas.fr, felipe.trevizan@anu.edu.au, sylvie.thiebaux@anu.edu.au

Abstract

Current approaches for learning for planning have yet to achieve competitive performance against classical planners in several domains, and have poor overall performance. In this work, we construct novel graph representations of lifted planning tasks and use the WL algorithm to generate features from them. These features are used with classical machine learning methods which have up to 2 orders of magnitude fewer parameters and train up to 3 orders of magnitude faster than the state-of-the-art deep learning for planning models. Our novel approach, WL-GOOSE, reliably learns heuristics from scratch and outperforms the h^{FF} heuristic in a fair competition setting. It also outperforms or ties with LAMA on 4 out of 10 domains on coverage and 7 out of 10 domains on plan quality. WL-GOOSE is the first learning for planning model which achieves these feats. Furthermore, we study the connections between our novel WL feature generation method, previous theoretically flavoured learning architectures, and Description Logic Features for planning.

1 Introduction

Learning for planning has regained traction in recent years due to advancements in deep learning (DL) and neural network architectures. The focus of learning for planning is to learn domain knowledge in an automated, domain-independent fashion in order to improve the computation and/or quality of plans. Recent examples of learning for planning methods using DL include learning policies (Toyer et al. 2018; Groshev et al. 2018; Garg, Bajpai, and Mausam 2019; Rivlin, Hazan, and Karpas 2020; Silver et al. 2024), heuristics (Shen, Trevizan, and Thiébaux 2020; Karia and Srivastava 2021) and heuristic proxies (Shen et al. 2019; Ferber et al. 2022; Chrestien et al. 2023), with more recent architectures motivated by theory (Ståhlberg, Bonet, and Geffner 2022, 2023; Mao et al. 2023; Chen, Thiébaux, and Trevizan 2024; Horcik and Šír 2024). However, learning for planning is not a new field and works capable of learning similar domain knowledge using classical statistical machine learning (SML) methods predate DL. For instance, learning heuristic proxies using support vector machines (SVMs) (Garrett, Kaelbling, and Lozano-Pérez 2016), policies using reinforcement learning (Buffet and Aberdeen

2009) and decision lists (Yoon, Fern, and Givan 2002). We refer to (Jiménez et al. 2012) for a more comprehensive overview of classical SML methods.

Unfortunately, all deep learning for planning architectures have yet to achieve competitive performance against classical planners and suffer from a variety of issues including (1) a need to tune a large number of hyperparameters, (2) lack of interpretability and (3) being both data and computationally intensive. In this paper, we introduce WL-GOOSE, a novel approach for learning for planning that takes advantage of the efficiency of SML-based methods for overcoming all these issues. WL-GOOSE uses a new graph representation for lifted planning tasks. However, differently from several DL-based methods, we do not use GNNs to learn domain knowledge and use *graph kernels* instead. More precisely, we use a modified version of the Weisfeiler-Leman algorithm for generating features for graphs (Shervashidze et al. 2011) which can be used to train SML models. Another benefit of WL-GOOSE is its support for various learning targets, such as heuristic values or policies, without the need for backpropagation to generate features as in DL-based approaches. We also provide a comprehensive theoretical comparison between our approach, GNNs for learning planning domain knowledge, and Description Logic Features for planning (Martín and Geffner 2000).

To demonstrate the potential of WL-GOOSE, we applied it to learn domain-specific heuristics using two classical SML methods: SVMs and Gaussian Processes (GPs). We evaluated the learned heuristics against the state-of-the-art learning for planning models on the 2023 International Planning Competition Learning Track benchmarks (Seipp and Segovia-Aguas 2023). The learned heuristics generalise better than previous DL-based methods while also being more computationally efficient: our models took less than 15 seconds to train which is up to 3 orders of magnitude times faster than GNNs which train on GPUs. Furthermore, some of our models were trained in a deterministic fashion with minimal parameter tuning, unlike DL-based approaches which require stochastic gradient descent and tuning of various hyperparameters, and have up to 2 orders of magnitude times fewer learned parameters. When used with greedy best-first search, our learned heuristic models achieved higher total coverage than h^{FF} (Hoffmann and Nebel 2001) and vastly outperforms all previous learning for

planning models. Moreover, our learned SVM and GP models outperformed or tied with LAMA (Richter and Westphal 2010) on 4 out of 10 domains with regards to coverage, and 7 out of 10 domains for plan quality. These results make our learned heuristics using WL-GOOSE the first ones to surpass the performance of h^{FF} and the best performing learned heuristics against LAMA.

2 Background and Notation

Planning A classical planning task (Geffner and Bonet 2013) is a state transition model given by a tuple $\Pi = \langle S, A, s_0, G \rangle$ where S is a set of states, A a set of actions, $s_0 \in S$ an initial state and $G \subseteq S$ a set of goal states. An action $a \in A$ is a function $a : S \to S \cup \bot$ where $a(s) = \bot$ indicates that action a is not applicable at state s, and otherwise, $a(s)$ is the successor state when a is applied to s. An action has an associated cost $c(a) \in \mathbb{R}_{\geq 0}$. A solution or *plan* for this model is a sequence of actions $\pi = a_1, \ldots, a_n$ where $s_i = a_i(s_{i-1}) \neq \bot$ for $i = 1, \ldots, n$ and $s_n \in G$. In other words, a plan is a sequence of applicable actions which progresses the initial state to a goal state when executed. The cost of a plan π is the sum of its action costs: $c(\pi) = \sum_{i=1}^{n} c(a_i)$. A planning task is *solvable* if there exists at least one plan. A plan is *optimal* if there does not exist any other plan with strictly lower cost.

We represent planning tasks in a compact form which does not require enumerating all states and actions. A *lifted planning task* (Lauer et al. 2021) is a tuple $\Pi = \langle \mathcal{P}, \mathcal{O}, \mathcal{A}, s_0, G \rangle$ where $\mathcal{P}$ is a set of first-order predicates, $\mathcal{O}$ a set of objects, $\mathcal{A}$ a set of action schemas, s_0 the initial state, and G the goal condition. A predicate $P \in \mathcal{P}$ has a set of parameters $x_1, \ldots, x_{n_P}$ where $n_P \in \mathbb{N}$ depends on P, and it is possible for a predicate to have no parameters. A ground proposition is a predicate which is instantiated by assigning all of the x_i with objects from $\mathcal{O}$ or other defined variables. An action schema $a \in \mathcal{A}$ is a tuple $\langle \Delta(a), \text{pre}(a), \text{add}(a), \text{del}(a) \rangle$ where $\Delta(a)$ is a set of parameter variables, and the preconditions $\text{pre}(a)$, add effects $\text{add}(a)$, and delete effects $\text{del}(a)$ are sets of predicates from $\mathcal{P}$ instantiated with elements from $\Delta(a) \cup \mathcal{O}$. Each action schema has an associated cost $c(a) \in \mathbb{R}_{\geq 0}$. An action is an action schema where each variable is instantiated with an object. A domain $\mathcal{D}$ is a set of lifted planning tasks which share the same sets of predicates $\mathcal{P}$ and action schemas $\mathcal{A}$.

In a lifted planning task, states are represented as sets of ground propositions. The following are sets of ground propositions: states, goal condition, and the preconditions, add effects, and delete effects of all actions. An action a is applicable in a state s if $\text{pre}(a) \subseteq s$, in which case we define $a(s) = (s \setminus \text{del}(a)) \cup \text{add}(a)$. Otherwise $a(s) = \bot$. The cost of an action is given by the cost of its corresponding action schema. A state s is a goal state if $G \subseteq s$.

A heuristic is a function $h : S \to \mathbb{R} \cup \{\infty\}$ which maps a state into a number representing an estimate of the cost of the optimal plan to the goal, or ∞ representing the state is unsolvable. A heuristic can be defined on problems by evaluating their initial state: $h(\Pi) = h(s_0)$. The optimal heuristic h^* returns for each state s the cost of the optimal plan to the goal if the problem is solvable from s, and ∞ otherwise.

Algorithm 1: WL algorithm

1   $c^0(v) \leftarrow c(v), \forall v \in V$
2   **for** $j = 1, \ldots, L$ **do for** $v \in V$ **do**
3    $c^j(v) \leftarrow \text{hash}\left(c^{j-1}(v), \{\!\{ c^{j-1}(u) \mid u \in \mathcal{N}(v) \}\!\} \right)$
4   **return** $\bigcup_{j=0,\ldots,L} \{\!\{ c^j(v) \mid v \in V \}\!\}$

Figure 1: Two non-isomorphic graphs G_1 (6-cycle) and G_2 (two disjoint 3-cycles) which the WL algorithm returns the same outputs, thus failing to distinguishing them.

The Weisfeiler-Leman algorithms We write $\langle V, E, c, l \rangle$ for a graph with coloured nodes and edges, where V is a set of nodes, $E \subseteq \binom{V}{2}$ is a set of undirected edges, $c : V \to \Sigma_V$ maps nodes to a set of colours Σ_V, and $l : E \to \Sigma_E$ maps edges to a set of colours Σ_E. The edge neighbourhood of a node u under edge colour ι is $\mathcal{N}_\iota(u) = \{e = \langle u, v \rangle = \langle v, u \rangle \in E \mid l(e) = \iota\}$. The neighbourhood of a node u in a graph is $\mathcal{N}(u) = \bigcup_{\iota \in \Sigma_E} \mathcal{N}_\iota(u)$.

We only focus on the WL algorithm which is a special case of the class of k-Weisfeiler-Leman (k-WL) algorithms (Leman and Weisfeiler 1968). The k-WL algorithms were originally constructed to provide tests for whether pairs of graphs are isomorphic or not. The $k+1$-WL algorithm subsumes the k-WL algorithm as it can distinguish a greater class of non-isomorphic graphs, and furthermore is in correspondence with k-variable counting logics (Cai, Fürer, and Immerman 1992). However, the complexity of the k-WL algorithms is exponential in k.

The WL algorithm takes as input graphs without edge colours, i.e. $\forall e \in E, l(e) = 0$, and outputs a canonical form in terms of a multiset of colours, a set which is allowed to have duplicate elements. It has also been used to construct a kernel for graphs (Shervashidze et al. 2011) which converts the multiset of colours in the WL algorithm into a feature vector and then uses the simple dot product kernel. We denote a multiset of elements by $\{\!\{ \ldots \}\!\}$.

The WL algorithm is shown in Alg. 1 which takes as input a graph G with coloured nodes only and a predefined number of WL iterations L. The algorithm begins by initialising the current colours of each node with the initial node colours. If no node colours are given in the graph, we can set them to 0. Line 3 updates the colour of each node v by iteratively collecting the current colours of its neighbors in a multiset and then hashing this multiset and v's current colour into a colour using an injective $\text{hash}(\cdot, \cdot)$ function. In practice, hash is built lazily by using a map data structure and multisets are represented as sorted strings. Line 4 returns a multiset of the node colours seen over all iterations.

If the WL algorithm outputs two different multisets for two graphs G_1 and G_2, then the graphs are non-isomorphic. However, if the algorithm outputs the same multisets for two graphs we cannot say for sure whether they are isomorphic or not. The canonical example illustrating this case is
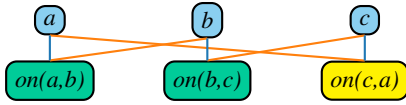
Figure 2: ILG subgraph of facts and goal condition corresponding to the *on* predicate of a Blocksworld instance. The current state says that a stacked on b, which is stacked on c, and the goal condition is for c to be stacked on a.

Figure 3: Expressivity hierarchy of WL, GNN and DL generated features for planning.

in Fig. 1 where the two graphs are not isomorphic but the WL algorithm returns the same output for both graphs since it views all nodes as the same because they have degree 2.

3 WL Features for Planning

In this section we describe how to generate features for planning states in order to learn heuristics. The process involves three main steps: (1) converting planning states into graphs with coloured nodes and edges, (2) running a variant of the WL algorithm on the graphs in order to generate features, and then (3) training a classical machine learning model for predicting heuristics using the obtained features. We start by defining the Instance Learning Graph (ILG), a novel representation for lifted planning tasks.

Definition 3.1. The *instance learning graph (ILG)* of a lifted planning problem $\Pi = \langle \mathcal{P}, \mathcal{O}, \mathcal{A}, s_0, G \rangle$ is the graph $G = \langle V, E, c, l \rangle$ with

- $V = \mathcal{O} \cup s_0 \cup G$
- $E = \bigcup_{p = P(o_1, \ldots, o_{n_P}) \in s_0 \cup G} \{ \langle p, o_1 \rangle, \ldots, \langle p, o_{n_P} \rangle \}$
- $c : V \to (\{\mathtt{ap}, \mathtt{ug}, \mathtt{ag}\} \times \mathcal{P}) \cup \{\mathtt{ob}\}$ defined by

$$
u \mapsto \begin{cases}
\mathtt{ob}, & \text{if } u \in \mathcal{O}; \\
(\mathtt{ag}, P), & \text{if } u = P(o_1, \ldots, o_{n_P}) \in s_0 \cap G; \\
(\mathtt{ap}, P), & \text{if } u = P(o_1, \ldots, o_{n_P}) \in s_0 \setminus G; \\
(\mathtt{ug}, P), & \text{if } u = P(o_1, \ldots, o_{n_P}) \in G \setminus s_0;
\end{cases}
$$

- $l : E \to \mathbb{N}$ with $\langle p, o_i \rangle \mapsto i$.

Fig. 2 provides an example of an ILG. An ILG consists of a node for each object and the union of propositions that are true in the state s_0 and the goal condition G. A proposition is connected to the n object nodes which instantiates the proposition. The labels of the n edges correspond to the position of the object in the predicate argument. The colours of the nodes indicate whether the node corresponds to an object ($\mathtt{ob}$), or determines whether it is a proposition belonging to s_0 ($\mathtt{ap}$) or G ($\mathtt{ug}$) only or both ($\mathtt{ag}$), as well as its corresponding predicate. Hence $\mathtt{ug}$ stands for unachieved goal, $\mathtt{ag}$ for achieved goal, and $\mathtt{ap}$ for achieved (non-goal) proposition. Note that ILGs are agnostic to the transition system of the planning task as they ignore action schemas and actions.

Since ILGs have coloured edges, we need to extend the WL algorithm to account for edge colours to generate features for ILGs. Our modified WL algorithm is obtained by replacing Line 3 in Alg. 1 with the update function

$$
c^j(v) \leftarrow \text{hash}\Big(c^{j-1}(v), \bigcup_{\iota \in \Sigma_E} \{\!\{ (c^{j-1}(u), \iota) \mid u \in \mathcal{N}_\iota(v) \}\!\} \Big),
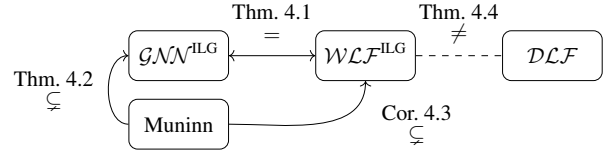$$

where the union of multisets is itself a multiset. Note that edge colours do not update during this modified WL algorithm. It is possible to run a variant of the WL algorithm which modifies edge colours but this comes at an additional computational cost given that usually $|E| \gg |V|$.

Now that ILGs can be represented as multisets of colours, we can generate features by representing these multisets as histograms (Shervashidze et al. 2011). The feature vector of a graph is a vector v with a size equal to the number of observed colours during training, where $v[\kappa]$ counts how many times the WL algorithm has encountered colour κ throughout its iterations. Formally, let $G_1 = \langle V_1, E_1, c_1, l_1 \rangle, \ldots, G_n = \langle V_n, E_n, c_n, l_n \rangle$ be the set of training graphs. Then the colours the WL algorithm encounters in the training graphs are given by

$$
\mathcal{C} = \{ c_i^j(v) \mid i \in \{1, \ldots, n\}; j \in \{0, \ldots, h\}; v \in V_i \}
$$

where $c_i^j(v)$ is the colour of node v in graph G_i during the j-th iteration of WL for $j > 0$ and $c_i^0(v) = c_i(v)$. Given a planning task Π and the set of colours $\mathcal{C}$ observed during training, Π's feature vector representation $\vec{v} \in \mathbb{R}^{|\mathcal{C}|}$ is $\vec{v} = [\text{count}_{\mathcal{C}}(\Pi, \kappa_1), \ldots, \text{count}_{\mathcal{C}}(\Pi, \kappa_{|\mathcal{C}|})]$ where $\text{count}_{\mathcal{C}}(\Pi, \kappa)$ is the number of times the colour $\kappa \in \mathcal{C}$ is present in the output of the WL algorithm on the ILG representation of Π. There is no guarantee that $\mathcal{C}$ contains all possible observable colours for a given planning domain and colours not in $\mathcal{C}$ observed after training are ignored.

4 Theoretical Results

In this section, we investigate the relationship between WL features and features generated using message passing graph neural network (GNN), description logic features (DLF) for planning (Martín and Geffner 2000) and the features used by Muninn (Ståhlberg, Bonet, and Geffner 2022, 2023), a theoretically motivated deep learning model. Fig. 3 summarises our theoretical results.

We begin with some notation. Let $\mathcal{D}$ represent the set of all problems in a given domain. We define $\mathcal{WLF}_{\Theta}^{\text{ILG}} : \mathcal{D} \to \mathbb{R}^d$ as the WL feature generation function described in Sec. 3 which runs the WL algorithm on the ILG representation of planning tasks. We denote Θ the set of parameters of the function which includes the number of WL iterations and the set of colours $\mathcal{C}$ with size d observed during training. We similarly denote parametrised GNNs acting on ILG representations of planning tasks by $\mathcal{GNN}_{\Theta}^{\text{ILG}} : \mathcal{D} \to \mathbb{R}^d$. Parameters for GNNs include number of message passing layers, the message passing update function with fixed weights, and the aggregation function.

We denote DLF generators (Martín and Geffner 2000) by $\mathcal{DLF}_{\boldsymbol{\Theta}} : \mathcal{D} \to \mathbb{R}^d$ where the parameters for $\mathcal{DLF}$ include the maximum complexity length of its features. DLFs have been used in several areas of learning for planning including learning descending dead-end avoiding heuristics (Francès et al. 2019), unsolvability heuristics (Ståhlberg, Francès, and Seipp 2021) and policy sketches (Bonet, Francès, and Geffner 2019; Drexler, Seipp, and Geffner 2022). Lastly, we denote the architecture from Ståhlberg, Bonet, and Geffner (2022) for generating features by Muninn$_{\boldsymbol{\Theta}} : \mathcal{D} \to \mathbb{R}^d$. We omit their final MLP layer which transforms the vector feature into a heuristic estimate. Furthermore in our theorems, we ignore their use of random node initialisation (RNI) (Abboud et al. 2021). The original intent of RNI is to provide a universal approximation theorem for GNNs but the practical use of the theorem is limited by the assumption of exponential width layers and absence of generalisation results. Parameters for Muninn include hyperparameters for their GNN architecture and learned weights for their update functions.

In all of the aforementioned models, the parameters $\boldsymbol{\Theta}$ consist of a combination of model hyperparameters and trained parameters based on a training set $\mathcal{T}_{\mathcal{D}} \subseteq \mathcal{D}$. The expressivity and distinguishing power of a feature generator for planning determines if it can theoretically learn h^* for larger subsets of planning tasks. We begin with an application of a well-known result connecting the expressivity of the WL algorithm and GNNs for distinguishing graphs (Xu et al. 2019) by extending it to edge-labelled graphs.

Theorem 4.1 ($\mathcal{WLF}^{\text{ILG}}$ and $\mathcal{GNN}^{\text{ILG}}$ have the same power at distinguishing planning tasks.)**. *Let Π_1 and Π_2 be any two planning tasks from a given domain. If for a set of parameters $\boldsymbol{\Theta}$ we have that $\mathcal{GNN}^{\text{ILG}}_{\boldsymbol{\Theta}}(\Pi_1) \neq \mathcal{GNN}^{\text{ILG}}_{\boldsymbol{\Theta}}(\Pi_2)$, then there exists a corresponding set of parameters $\boldsymbol{\Phi}$ such that $\mathcal{WLF}^{\text{ILG}}_{\boldsymbol{\Phi}}(\Pi_1) \neq \mathcal{WLF}^{\text{ILG}}_{\boldsymbol{\Phi}}(\Pi_2)$. Conversely for all $\boldsymbol{\Phi}$ such that $\mathcal{WLF}^{\text{ILG}}_{\boldsymbol{\Phi}}(\Pi_1) \neq \mathcal{WLF}^{\text{ILG}}_{\boldsymbol{\Phi}}(\Pi_2)$, there exists $\boldsymbol{\Theta}$ such that $\mathcal{GNN}^{\text{ILG}}_{\boldsymbol{\Theta}}(\Pi_1) \neq \mathcal{GNN}^{\text{ILG}}_{\boldsymbol{\Theta}}(\Pi_2)$.*

Proof. [$\subseteq$] The forward statement follows from (Xu et al. 2019, Lemma 3) which states that GNNs are at most as expressive as the WL algorithm for distinguishing non-isomorphic graphs. We can modify the lemma for the edge labelled WL algorithm and GNNs which account for edge features. Then the result follows after performing the transformation of planning tasks into the ILG representation.

[$\supseteq$] The converse statement follows from (Xu et al. 2019, Corollary 6) and modifying Eq. (4.1) of their GIN architecture by introducing an MLP for each of the finite number of edge labels in the ILG graph and summing their outputs at each GIN layer. The MLPs have disjoint range in order for injectivity to be preserved as to achieve the same distinguishing power of the edge labelled WL algorithm. This can be easily enforced by increasing the hidden dimension size and having each MLP to map to orthogonal dimensions. □

We proceed to show that GNNs acting on ILGs is similar to Muninn's GNN architecture (Ståhlberg, Bonet, and Geffner 2022). The idea of the proof is that encoding different predicates into the ILG representation is equivalent to having different weights for message passing to and from
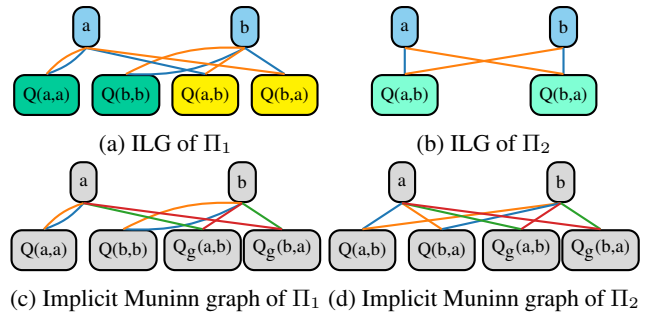

(a) ILG of Π_1 (b) ILG of Π_2

(c) Implicit Muninn graph of Π_1 (d) Implicit Muninn graph of Π_2

Figure 4: ILG and Muninn graph representations of tasks in Thm. 4.2 [$\supsetneq$].

different predicates in Muninn. However, we also show that our model has strictly higher expressivity for distinguish planning tasks due to explicitly encoding achieved goals.

Theorem 4.2 ($\mathcal{GNN}^{\text{ILG}}$ is strictly more expressive than Muninn at distinguishing planning tasks.)**. *Let Π_1 and Π_2 be any two planning tasks from a given domain. For all $\boldsymbol{\Theta}$, if Muninn$_{\boldsymbol{\Theta}}(\Pi_1) \neq$ Muninn$_{\boldsymbol{\Theta}}(\Pi_2)$, then there exists a corresponding set of parameters $\boldsymbol{\Phi}$ such that $\mathcal{GNN}^{\text{ILG}}_{\boldsymbol{\Phi}}(\Pi_1) \neq \mathcal{GNN}^{\text{ILG}}_{\boldsymbol{\Phi}}(\Pi_2)$. Furthermore, there exists a pair of planning tasks Π_1 and Π_2 such that there exists $\boldsymbol{\Phi}$ with $\mathcal{GNN}^{\text{ILG}}_{\boldsymbol{\Phi}}(\Pi_1) \neq \mathcal{GNN}^{\text{ILG}}_{\boldsymbol{\Phi}}(\Pi_2)$ but for all $\boldsymbol{\Theta}$, Muninn$_{\boldsymbol{\Theta}}(\Pi_1) = $ Muninn$_{\boldsymbol{\Theta}}(\Pi_2)$.*

Proof. [$\supseteq$] In order to show the inclusion, we show that a Muninn instance operating on a state can be expressed as a GNN operating on the ILG representation of the state. More explicitly, we show that the implicit graph representation of planning states by Muninn is the same graph as ILG. The message passing steps and initial node features are different but the semantic meaning of executing both algorithms are the same. The node features in the implicit graphs of Muninn are all the same when ignoring random node initialisation. Muninn differentiates object nodes and fact nodes by using different message passing functions depending on whether a node is an object or a fact, and depending on which predicate the fact belongs to. In the language of ILG, Muninn's message passing step on fact nodes $p = P(o_1, \ldots, o_{n_P})$ is

$$h_p^{L+1} = \mathbf{MLP}_P(h_{o_1}^L, \ldots, h_{o_{n_P}}^L) \tag{1}$$

where h_p^{L+1} denotes the latent embedding of the node p in the $L+1$-th layer, $h_{o_i}^L$ denotes the latent embedding of the object node o_i in the l-th layer, and $\mathbf{MLP}_P$ is a multilayer perceptron, with a different one for each predicate. The message passing step of Muninn on object nodes o_i is

$$h_o^{L+1} = \mathbf{MLP}_U(h_o^L, \{\!\{h_p^L \mid o \in p\}\!\}) \tag{2}$$

where $o \in p$ denotes that o is an argument of the predicate associated with p. We note that having a different $\mathbf{MLP}$ in the message passing step for different nodes is equivalent to having a larger but identical $\mathbf{MLP}$ in the message passing step for all nodes. This is because the model can learn to partition latent node features depending on their semantic

meaning and thus be able to use a single **MLP** function to act as multiple different functions for different node feature partitions. Thus, Eq. (1) and (2) can be imitated by a GNN operating on ILG since ILG features differentiate nodes depending on whether they correspond to an object, or a fact associated with a predicate. Different edge labels in the ILG allow it to distinguish the relationship between facts and objects depending on their position in the predicate argument.

[$\supsetneq$] The main idea here is that Muninn does not keep track of achieved goals and sometimes cannot even see that the goal has been achieved. Firstly, let MuG denote the underlying edge-labelled graph representation of planning tasks in Muninn, such that $\mathcal{GNN}^{\text{MuG}}$ = Muninn. To see how $\mathcal{GNN}^{\text{ILG}}$ is strictly more expressive than Muninn, we consider the following pair of planning tasks. Let $\Pi_1 = \langle \mathcal{P}, \mathcal{O}, \mathcal{A}, s_0^{(1)}, G \rangle$ and $\Pi_2 = \langle \mathcal{P}, \mathcal{O}, \mathcal{A}, s_0^{(2)}, G \rangle$ with $\mathcal{P} = \{Q\}$, $\mathcal{O} = \{a, b\}$, $\mathcal{A} = \emptyset$, $G = s_0^{(2)} = \{Q(a,b), Q(b,a)\}$ and $s_0^{(1)} = \{Q(a,a), Q(b,b)\}$. Fig. 4 illustrates the ILG and MuG representation of Π_1 and Π_2. It is clear that the ILG representation of Π_1 and Π_2 are different and hence $\mathcal{GNN}^{\text{ILG}}$ differentiates between Π_1 and Π_2. On the other hand without RNI, any edge-labelled variant of the WL algorithm views the MuG representation of Π_1 and Π_2 illustrated in Fig. 4(c) and (d) as the same. Thus, $\mathcal{GNN}^{\text{MuG}}$ = Muninn views the graphs as the same. $\square$

Corollary 4.3 ($\mathcal{WLF}^{\text{ILG}}$ is strictly more expressive than Muninn at distinguishing planning tasks.)**.** *Let Π_1 and Π_2 be any two planning tasks from a given domain. For all Θ, if $\text{Muninn}_\Theta(\Pi_1) \neq \text{Muninn}_\Theta(\Pi_2)$, then there exists a corresponding set of parameters Φ such that $\mathcal{WLF}_\Phi^{\text{ILG}}(\Pi_1) \neq \mathcal{WLF}_\Phi^{\text{ILG}}(\Pi_2)$. Furthermore, there exists a pair of planning tasks Π_1 and Π_2 such that there exists Φ with $\mathcal{WLF}_\Phi^{\text{ILG}}(\Pi_1) \neq \mathcal{WLF}_\Phi^{\text{ILG}}(\Pi_2)$ but for all Θ, $\text{Muninn}_\Theta(\Pi_1) = \text{Muninn}_\Theta(\Pi_2)$.*

Our next theorem shows that $\mathcal{WLF}^{\text{ILG}}$ and $\mathcal{DLF}$ features are incomparable, in the sense that there are pairs of planning tasks that look equivalent to one model but not the other. We use a similar counterexample to that used for Muninn but with an extra predicate which $\mathcal{WLF}^{\text{ILG}}$ does not distinguish but $\mathcal{DLF}$ can. Conversely we use the fact that DLFs are limited by the need to convert planning predicates into binary predicates to construct a counterexample pair of planning tasks with ternary predicates which $\mathcal{DLF}$ views as the same while $\mathcal{WLF}^{\text{ILG}}$ does not.

Theorem 4.4 ($\mathcal{WLF}^{\text{ILG}}$ and $\mathcal{DLF}$ are incomparable at distinguishing planning tasks.)**.** *There exists a pair of planning tasks Π_1 and Π_2 such that there exists Φ with $\mathcal{WLF}_\Phi^{\text{ILG}}(\Pi_1) \neq \mathcal{WLF}_\Phi^{\text{ILG}}(\Pi_2)$ but for all Θ, $\mathcal{DLF}_\Theta(\Pi_1) = \mathcal{DLF}_\Theta(\Pi_2)$. Furthermore, there exists a pair of planning tasks Π_1 and Π_2 such that there exists Φ with $\mathcal{DLF}_\Phi(\Pi_1) \neq \mathcal{DLF}_\Phi(\Pi_2)$ but for all Θ, $\mathcal{WLF}_\Theta^{\text{ILG}}(\Pi_1) = \mathcal{WLF}_\Theta^{\text{ILG}}(\Pi_2)$.*

Proof. [$\exists <$] We begin by describing a pair of planning tasks Π_1 and Π_2 such that $\mathcal{WLF}_\Theta^{\text{ILG}}(\Pi_1) = \mathcal{WLF}_\Theta^{\text{ILG}}(\Pi_2)$ for any set of parameters Θ but are distinguished by DLF.


(a) ILG of Π_1 (b) ILG of Π_1

Figure 5: ILG representations of tasks in Thm. 4.4 [$\exists <$].

Let $\Pi_1 = \langle \mathcal{P}, \mathcal{O}, \mathcal{A}, s_0^1, G \rangle$ and $\Pi_2 = \langle \mathcal{P}, \mathcal{O}, \mathcal{A}, s_0^2, G \rangle$ with $\mathcal{P} = \{Q, W\}$, $\mathcal{O} = \{a, b\}$, $\mathcal{A}$ contains the single action schema $o = \langle \{x, y\}, \{Q(x,y)\}, \{W(x,y)\}, \emptyset \rangle$, $G = \{W(a,b), W(b,a)\}$, $s_0^1 = \{Q(a,a), Q(b,b)\}$ and $s_0^2 = \{Q(a,b), Q(b,a)\}$.

We have that $h^*(\Pi_1) = \infty$ as the problem Π_1 is unsolvable, while $h^*(\Pi_2) = 2$ as an optimal plan contains actions $o(a,b)$ and $o(b,a)$. DL features are able to distinguish the two planning tasks by considering the role-value map $(Q = W)(s)$ defined by $\{x \mid \forall y : Q(x,y) \in s \iff W(x,y) \in s\}$, and the corresponding numerical feature $|Q = W|(s) = |(Q = W)(s)|$. We have that $|Q = W|(s_0^1) = 0$ and $|Q = W|(s_0^1) = 2$, meaning that $\mathcal{DLF}$ can distinguish between Π_1 and Π_2.

On the other hand, the ILG representations of Π_1 and Π_2 are indistinguishable to our definition of the edge-labelled WL algorithm. Fig. 5 illustrates this example and we note that it is similar to the implicit Muninn graph representations of the pair of planning tasks from Thm. 4.2.

[$\exists >$] We identify a pair of problems with ternary predicates which compile to the same problem with only binary predicates for which DL features are defined. For problems with at most binary predicates, DL introduces base roles on each predicate $P(x,y) \in \mathcal{P}$ by $P^s = \{(a,b) \mid P(a,b) \in s\}$ where s is a planning state. Then given an n-ary predicate $R(x_1, \ldots, x_n)$, DL introduces $n(n-1)/2$ roles defined by $R_{i,j}^s = \{(a,b) \mid \exists o_1, \ldots, o_{i-1}, o_{i+1}, \ldots, o_{j-1}, o_{j+1}, \ldots, o_n \in \mathcal{O}, R(o_1, \ldots, o_{i-1}, a, o_{i+1}, \ldots, o_{j-1}, b, o_{j+1}, \ldots, o_n) \in s\}$ for $1 \leq i < j \leq n$. Now consider the problems $\Pi_1 = \langle \mathcal{P}, \mathcal{O}, \mathcal{A}, s_0^1, G \rangle$ and $\Pi_2 = \langle \mathcal{P}, \mathcal{O}, \mathcal{A}, s_0^2, G \rangle$ now with $\mathcal{P} = \{P\}$, $\mathcal{O} = \{a, b, c, d\}$, $\mathcal{A} = \emptyset$, $G = \{P(a,b,c)\}$, and

$$s_0^1 = \{P(a,b,a), P(c,b,c), P(a,d,c), P(c,d,a)\}$$
$$s_0^2 = \{P(a,b,c), P(c,b,a), P(a,d,a), P(c,d,c)\}.$$

We have that $h^*(\Pi_1) = \infty$ since there are no actions and the initial state is not the goal condition, while $h^*(\Pi_2) = 0$ since $G \subseteq s_0^2$. The ILG for the two tasks are distinguished by the WL algorithm as the ILG of Π_1 has no achieved goal colour while Π_2 does. However, DL features view the two states s_0^1 and s_0^2 as the same due after the compilation from ternary to binary predicates:

$P_{1,2}(a,b)$	$P_{1,2}(a,d)$	$P_{1,2}(c,b)$	$P_{1,2}(c,d)$
$P_{1,3}(a,a)$	$P_{1,3}(a,c)$	$P_{1,3}(c,a)$	$P_{1,3}(c,c)$
$P_{2,3}(b,a)$	$P_{2,3}(b,c)$	$P_{2,3}(d,a)$	$P_{2,3}(d,c).$

Thus any DL features will be the same for both s_0^1 and s_0^2 and thus cannot distinguish Π_1 and Π_2. $\square$

Our final theorem combines previous results and states that there exist domains for which all feature generators defined thus far are not powerful enough to perfect learn h^*, with proof in the appendix. Although this is not a surprising result, we hope to bring intuition on what is further required for constructing more expressive planning features.

Corollary 4.5 (All feature generation models thus far cannot generate features that allow us to learn h^* for all domains.). *Let $\mathcal{F} \in \{\mathcal{WLF}^{ILG}, \mathcal{GNN}^{ILG}, Muninn, \mathcal{DLF}\}$. There exists a domain $\mathcal{D}$ with a pair of planning tasks Π_1, Π_2 such that for all parameters Θ for $\mathcal{F}$, we have that $\mathcal{F}_{\Theta}(\Pi_1) = \mathcal{F}_{\Theta}(\Pi_2)$ and $h^*(\Pi_1) \neq h^*(\Pi_2)$.*

In this section, we concluded that our $\mathcal{WLF}^{ILG}$ features are one of the most expressive features thus far in the literature for representing planning tasks, the other being $\mathcal{DLF}$ features. We have done so by drawing an expressivity hierarchy between our $\mathcal{WLF}^{ILG}$ features and previous work on GNN architectures (Ståhlberg, Bonet, and Geffner 2022). We further constructed explicit counterexamples illustrating the difference between $\mathcal{WLF}^{ILG}$ and $\mathcal{DLF}$ features, highlighting their respective advantages and limitations.

5 Experiments

In this section, we empirically evaluate WL-GOOSE[1] for learning domain-specific heuristics using WL features against the state-of-the-art. We consider the domains and training and test sets from the learning track of the 2023 International Planning Competition (IPC) (Seipp and Segovia-Aguas 2023). The domains are Blocksworld, Childsnack, Ferry, Floortile, Miconic, Rovers, Satellite, Sokoban, Spanner, and Transport. Actions in all domains have unit cost. Each domain contains instances categorised into easy, medium and hard difficulties depending on the number of objects in the instance. For each domain, the training set consists of at most 99 easy instances and the test set consists of exactly 30 instances from each of the three easy, medium and hard difficulties that are not in the training set.

The hyperparameters considered for WL-GOOSE are the number of iterations L for generating features using the WL algorithm and the choice of a machine learning model used and its corresponding hyperparameters. In all our experiments with WL-GOOSE, we use $L = 4$ and, since our learning target is h^*, we consider the following regression models: support vector regression with the dot product kernel (SVR) and the radial basis kernel (SVR_{∞}), and Gaussian process regression (Rasmussen and Williams 2006) with the dot product kernel (GPR). We choose SVR over ridge regression for our kernelised linear model due to its sparsity and hence faster evaluation time with use of the ϵ-insensitive loss function (Vapnik 2000). The choice of Gaussian process allows us to explore a Bayesian treatment for learning h^*, providing us with confidence bounds on learned heuristics.

Furthermore, we experiment with the *2-LWL algorithm* (Morris, Kersting, and Mutzel 2017) with $L = 4$ for generating features alongside SVR with the dot product kernel ($SVR_{2\text{-LWL}}$). The 2-LWL algorithm is a computationally

feasible approximation of the 2-WL algorithm (Morris, Kersting, and Mutzel 2017), which in turn is a generalisation of the WL algorithm where colours are assigned to pairs of vertices. While the features computed by the 2-WL algorithm subsume those of the WL algorithm, it is requires quadratically more time than the WL algorithm.

For any configuration of WL-GOOSE, we use optimal plans returned by scorpion (Seipp, Keller, and Helmert 2020) on the training set with a 30-minute timeout on each instance for training. States and the corresponding cost to the goal from each optimal plan are used as training data. As baselines for heuristics, we use the domain-independent heuristic h^{FF} and GNNs. For the GNNs, we use GOOSE (Chen, Thiébaux, and Trevizan 2024) operating on the ILG representations of planning tasks with the max aggregator and Muninn adapted to learn heuristics for use in GBFS only (Ståhlberg 2024). Every GOOSE model and SVR model is trained and evaluated 5 times with mean scores reported. GPR's optimisation is deterministic and thus is only trained and evaluated once. All GNNs use 4 message passing layers and a hidden dimension of 64. All heuristics are evaluated using GBFS. We include LAMA (Richter and Westphal 2010) using its first plan output as a strong satisficing planner baseline that uses multi-queue heuristic search and other optimisation techniques. All methods use a timeout of 1800 seconds per evaluation problem. Non-GNN models were run on a cluster with single Intel Xeon 3.2 GHz CPU cores and a memory limit of 8GB. GOOSE used an NVIDIA RTX A6000 GPU, and Muninn an NVIDIA A10. Other competition planners were not considered because they do not learn a heuristic.

Tab. 1 summarises our results with the coverage per domain for each planner and their total IPC score. We discuss our results in detail and conclude this section by describing how to analyse the learned features of our models using an example. More results can be found in the appendix.

How well do heuristics learned from WL features perform? Considering total coverage and total IPC score (Tab. 1), we notice that SVR and GPR outperform all the other planners with the exception of LAMA-first, i.e., all learning-based approaches as well as h^{FF}. Domain-wise, both SVR and GPR outperform Muninn and GOOSE on 9 domains. Both SVR and GPR outperform or tie with LAMA on 4 domains, namely Blocksworld, Ferry, Miconic and Spanner. GPR is able to return better plans than LAMA on 5 domains (Blocksworld, Childsnack, Ferry, Miconic, Sokoban), while the reverse is true only on 3 domains (Rovers, Satellite, Spanner). For Spanner, this is because LAMA's heuristics are not informative for this domain, which leads to it performing like blind search and hence returning better plans on problems it can solve.

SVR and GPR also outperform or tie with h^{FF} on 6 and 7 domains, respectively. We compare GPR and h^{FF} in more detail in Fig. 6 by showing plan cost and nodes expanded per problem. We observe that the better performing planner on a domain generally has better plan quality and fewer node expansions. An exception is Sokoban where GPR expands more nodes but solves more problems due to its faster

---

[1] Source code available at (Chen, Trevizan, and Thiébaux 2024)

	classical		GNN		WLF			
Domain	LAMA-F	h^{FF}	Muninn	GOOSE‡	SVR‡	SVR‡$_\infty$	SVR‡$_{\text{2-LWL}}$	GPR
blocksworld	61	28	53	63	72	19	22	**<u>75</u>**
childsnack	<u>35</u>	26	12	23	25	13	10	29
ferry	68	68	38	70	**76**	32	60	<u>76</u>
floortile	11	**<u>12</u>**	1	0	2	0	0	2
miconic	90	**90**	90	89	**90**	30	67	<u>90</u>
rovers	<u>67</u>	34	24	26	**38**	28	34	37
satellite	<u>89</u>	**65**	16	31	46	29	19	53
sokoban	<u>40</u>	36	31	33	**38**	30	31	**38**
spanner	30	30	**76**	46	73	30	52	73
transport	<u>66</u>	**41**	24	32	31	27	34	29
all	<u>557</u>	430	365	413	491	238	328	**502**

Table 1: Coverage of planners. The bottom-most row provides their overall IPC 2023 learning track score. Our new models are the WLF models. Models marked ‡ are run 5 times with mean scores presented. LAMA-first is the only planner not performing single-queue GBFS. The top three single-queue heuristic search planners in each row are indicated by the cell colouring intensity, with the best one in bold. The best planner overall in each row is underlined.


Figure 6: Returned plan cost and number of expanded nodes of h^{FF} and GPR. Problems that were not solved by one planner has their respective metric set to the axis limit. Points on the top left triangle favour GPR while points on the bottom right triangle favour h^{FF}.

heuristic evaluations. Overall, the domains in which GPR performs worse are domains that require traversing a map which WL features cannot express with limited iterations.

Are our methods more computationally efficient to train? To answer this question, we compare the training time of GNNs using ILGs, SVR and GPR. Their mean and standard deviation in seconds are 77.2 ± 33.7 (GOOSE), 0.2 ± 0.1 (SVR) and 3.8 ± 4.6 (GPR). Comparing against the more efficient GNN model per domain, we have that SVR is between 187x (Satellite) to 922x (Childsnack) more efficient and GPR is between 8x (Floortile) and 615x (Childsnack) more efficient. Note that the GNNs have access to GPUs and would take even more time to train on a CPU. Lastly, the GNNs have between 54529 and 74561 number of parameters in the range, while the WLF models have between 108 and 23202 parameters.

	h error				Expanded			
Domain	E.	M.	H.	all	E.	M.	H.	all
blocksworld	**+0.9**	**+0.9**	**+0.9**	**+1.0**	*+0.3*	*+0.2*	*+0.3*	**+0.6**
childsnack	**+0.7**	**+0.9**	-	**+0.9**	**+0.6**	*+0.5*	-	*+0.2*
ferry	**+0.9**	**+1.0**	**+1.0**	**+1.0**	**+0.9**	**+0.9**	**+0.8**	**+0.9**
floortile	-	-	-	-	-	-	-	-
miconic	**+0.6**	**+0.7**	**+1.0**	**+1.0**	**+0.6**	**+0.8**	**+1.0**	**+1.0**
rovers	**+0.9**	**+0.9**	-	**+1.0**	*+0.3*	*+0.2*	-	**+0.5**
satellite	**+0.7**	**+1.0**	-	**+1.0**	*+0.1*	*+0.1*	-	*+0.2*
sokoban	*+0.3*	**+0.9**	-	**+1.0**	*+0.3*	**+0.8**	-	**+0.8**
spanner	**+0.3**	**+0.5**	**+1.0**	**+0.9**	**+0.4**	**+0.5**	**+1.0**	**+0.9**
transport	**+0.8**	-	-	**+0.8**	**+0.4**	-	-	**+0.4**

Table 2: Pearson's correlation coefficient ρ rounded do 1 d.p. between standard deviation obtained by GPR against heuristic estimate error and node expansions of initial states from solved problems. Statistically significant coefficients ($p < 0.05$) are highlighted in bold font and italics otherwise. Strongly correlated values ($|\rho| \geq 0.5$), medium correlated values ($0.3 \leq |\rho| < 0.5$), and low correlation ($|\rho| < 0.3$) values are denoted by cell colouring intensity. Entries for which we solved fewer then 10 problems are omitted.

Does kernelising help? As commonly done in classical machine learning, we combine our WL features with non-linear kernels to obtain new non-linear features that can increase the expressivity of the regression models. Unfortunately, as shown in Tab. 1, this generally results in a decrease in the performance of the learned heuristic: the SVR$_\infty$ model has significantly worse coverage than SVR despite theoretically having more expressive implicit features. The drop in performance can be explained by overfitting to the more expressive features which do not bring any obvious semantic information for planning tasks.

Do higher order WL features help? The motivation for using higher-order WL features is similar to using higher-order kernels: to introduce more expressive features that may be correlated with the optimal heuristic. In Tab. 1, we see that the performance of 2-LWL is generally worse on all domains except for Transport. This again can be attributed to poorer generalisation. Furthermore, computing the 2-LWL features are slower to generate than WL features as they take time cubic in the size of the ILGs in the worst case, and in the case of Floortile runs out of memory when generating features. We also note that attempting to generate 3-LWL features causes out of memory problems during training as the size of features generated is extremely large, on the order of 10^7 and above.

Are Bayesian variance estimates meaningful? One advantage of Bayesian models is that by assuming a prior distribution on the weights of our models, we are able to derive uncertainty bounds on the outputs of the learned posterior model. In Tab. 2, we analyse the Pearson's correlation coefficient between the standard deviation obtained by GPR and (1) the error between output mean and h^*, and (2) the number of expanded nodes using the learned heuristic with greedy best first search. We see that there is a sta-
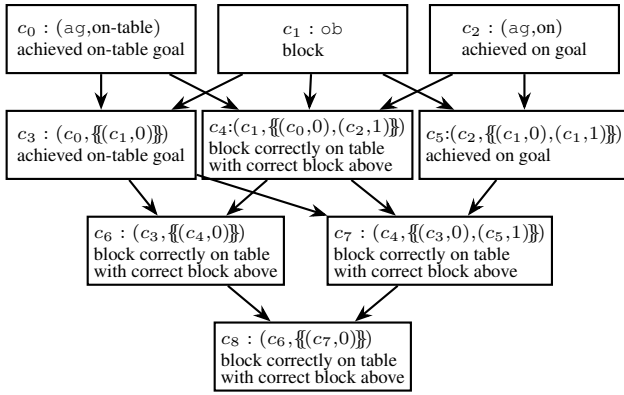
Figure 7: The dependency subgraph of generated WL features on Blocksworld. The first row of each node indicates the feature colour, followed by the initial colour the feature corresponds to or the input to the hash function which generated the colour. The second row describes the semantic meaning of the feature. Edges describe the dependency of the feature on previous features based on the hash function.

tistically significant strong correlation between the heuristic estimate error and the GPR variance outputs. This is reasonable given that the derivation of the Bayesian model computes the uncertainty on its output prediction. The story is different for the number of expansions during search where for easy problems there is no significant correlation depending on the domain. Interestingly, the correlation is more significant and stronger on harder problems for more domains. Thus, the Bayesian model is able to determine the difficulty of solving a problem within a domain by looking at the predicted standard deviation for $h(s_0)$ but the quality of this prediction will depend on the domain.

Understanding Learned Models

Another advantage of WL-GOOSE is that its set of features is explainable, and it is possible to see which features are chosen when using a linear inference model. The models can be understood by analysing the features with the highest corresponding linear weights, and by observing the distribution of such weights. The semantic meaning of the features can be understood by examining the *generation* of WL colours. This can be achieved by representing the observed WL colours as a directed acyclic graph (DAG) where each WL colour is a node and there is a directed edge from κ to κ' if $\kappa' = \text{hash}(x, M)$ and $x = \kappa$ or $\exists \iota, (\kappa, \iota) \in M$. We provide an example of how to interpret the learned models by briefly studying the learned GPR model on Blocksworld. In this domain, a total of 10444 features were generated from the training data and Fig. 7 illustrates the DAG representation of feature c_8's generation. Consider feature c_4 in Fig. 7, it computes the number of blocks that are correctly on the table and also have the correct block above it. We have that $c_4 = \text{hash}(c_1, \{\!\{(c_0, 0), (c_2, 1)\}\!\})$, meaning that the colour c_4 is generated from an object node ($c_1 = \text{ob}$) which is part of an achieved on-table goal ($c_0 = (\text{ag}, \text{on-table})$) and achieved on goal ($c_2 = (\text{ag}, \text{on})$). The corresponding edge

label of the node colours indicate the position of the block object in the proposition indexed from 0. Thus, blocks b with colour c_4 are in the first and only argument of on-table and the second argument of on. This means that the colour c_4 is assigned to blocks correctly on the table and correctly underneath another block.

Moreover, we observed that certain subsets of features were evaluated to the same value on all training states. As a result, the same learned weight value was assigned to each feature in these subsets. This can be seen in Fig. 7 where features c_4, c_7, c_6 and c_8 are semantically equivalent. The sum of their weight values is -1.76, the largest in value from subsets of features. Thus, the learned weight rewards states satisfying this condition as blocks correctly on the table do not have to be moved.

Note that it is possible for features to evaluate to the same values on the training set but have different semantic meanings because the training set is finite. For example, in Blocksworld, a training set may satisfy that a block is correctly on the table if and only if it has the correct block above it. In this case, the count of colours c_0 and c_4 would be the same on all states despite not being semantically equivalent.

6   Conclusion

We introduced WL-GOOSE, a novel approach that makes use of the efficiency of classical machine learning for learning to plan. We developed the Instance Learning Graph (ILG), a novel representation of lifted planning tasks and provided a method to generate features for ILGs based on the WL algorithm, agnostic to the downstream model. Similar to Description Logic Features for planning, our generated features are agnostic to the learning target and can be used without the need for backpropagation. Furthermore, some of our models can be trained in a deterministic fashion with minimal parameter tuning in contrast to DL-based approaches. To validate the benefits of WL-GOOSE, we used two classical SML models, support vector regression (SVR) and Gaussian process regression (GPR), to learn domain-specific heuristics and compared them to the state of the art.

The experimental results showed that WL-GOOSE can efficiently and reliably learn domain-specific heuristics from scratch. Compared to GNNs applied to ILGs, our learned heuristics are up to 3 orders of magnitude times faster to train and have up to 2 orders of magnitude fewer parameters. Our results also showed that both SVR and GPR are the first learned heuristics capable of outperforming h^{FF} in terms of total coverage. Moreover, our learned heuristics outperform or tie with LAMA on 4 domains. To our knowledge, this is the best performance of learned heuristics against LAMA. We also showed the theoretical connections between our novel feature generation method with Description Logic Features and GNNs. Our future work agenda includes exploring how to best use the uncertainty bounds provided by GPR to improve search, making use of generated WL features for learning different forms of domain knowledge such as policies, landmarks and sketches (Bonet and Geffner 2021), and combining stronger satisfying search algorithms to further improve the performance of WL-GOOSE.

Acknowledgements

References

Abboud, R.; Ceylan, İ. İ.; Grohe, M.; and Lukasiewicz, T. 2021. The Surprising Power of Graph Neural Networks with Random Node Initialization. In *IJCAI*.

Bonet, B.; Francès, G.; and Geffner, H. 2019. Learning Features and Abstract Actions for Computing Generalized Plans. In *AAAI*.

Bonet, B.; and Geffner, H. 2021. General Policies, Representations, and Planning Width. In *AAAI*.

Buffet, O.; and Aberdeen, D. 2009. The factored policy-gradient planner. *Artificial Intelligence*.

Cai, J.-Y.; Fürer, M.; and Immerman, N. 1992. An optimal lower bound on the number of variables for graph identification. *Combinatorica*.

Chen, D. Z.; Thiébaux, S.; and Trevizan, F. 2024. Learning Domain-Independent Heuristics for Grounded and Lifted Planning. In *AAAI*.

Chen, D. Z.; Trevizan, F.; and Thiébaux, S. 2024. Code for Return to Tradition: Learning Reliable Heuristics with Classical Machine Learning. https://doi.org/10.5281/zenodo.10757383.

Chrestien, L.; Edelkamp, S.; Komenda, A.; and Pevný, T. 2023. Optimize Planning Heuristics to Rank, not to Estimate Cost-to-Goal. In *NeurIPS*.

Drexler, D.; Seipp, J.; and Geffner, H. 2022. Learning sketches for decomposing planning problems into subproblems of bounded width. In *ICAPS*.

Ferber, P.; Geißer, F.; Trevizan, F.; Helmert, M.; and Hoffmann, J. 2022. Neural Network Heuristic Functions for Classical Planning: Bootstrapping and Comparison to Other Methods. In *ICAPS*.

Francès, G.; Corrêa, A. B.; Geissmann, C.; and Pommerening, F. 2019. Generalized Potential Heuristics for Classical Planning. In *IJCAI*.

Garg, S.; Bajpai, A.; and Mausam. 2019. Size Independent Neural Transfer for RDDL Planning. In *ICAPS*.

Garrett, C. R.; Kaelbling, L. P.; and Lozano-Pérez, T. 2016. Learning to rank for synthesizing planning heuristics. In *IJCAI*.

Geffner, H.; and Bonet, B. 2013. *A Concise Introduction to Models and Methods for Automated Planning*. Synthesis Lectures on Artificial Intelligence and Machine Learning. Morgan & Claypool Publishers.

Groshev, E.; Goldstein, M.; Tamar, A.; Srivastava, S.; and Abbeel, P. 2018. Learning Generalized Reactive Policies Using Deep Neural Networks. In *ICAPS*.

Hoffmann, J.; and Nebel, B. 2001. The FF Planning System: Fast Plan Generation Through Heuristic Search. *JAIR*.

Horcik, R.; and Šír, G. 2024. Expressiveness of Graph Neural Networks in Planning Domains. In *ICAPS*.

Jiménez, S.; de la Rosa, T.; Fernández, S.; Fernández, F.; and Borrajo, D. 2012. A review of machine learning for automated planning. *The Knowledge Engineering Review*.

Karia, R.; and Srivastava, S. 2021. Learning Generalized Relational Heuristic Networks for Model-Agnostic Planning. In *AAAI*.

Lauer, P.; Torralba, A.; Fiser, D.; Höller, D.; Wichlacz, J.; and Hoffmann, J. 2021. Polynomial-Time in PDDL Input Size: Making the Delete Relaxation Feasible for Lifted Planning. In *IJCAI*.

Leman, A.; and Weisfeiler, B. 1968. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Technicheskaya Informatsiya*.

Mao, J.; Lozano-Pérez, T.; Tenenbaum, J. B.; and Kaelbling, L. P. 2023. What Planning Problems Can A Relational Neural Network Solve? In *NeurIPS*.

Martín, M.; and Geffner, H. 2000. Learning Generalized Policies in Planning Using Concept Languages. In *KR*.

Morris, C.; Kersting, K.; and Mutzel, P. 2017. Glocalized Weisfeiler-Lehman Graph Kernels: Global-Local Feature Maps of Graphs. In *IEEE Int. Conf. on Data Mining*.

Rasmussen, C. E.; and Williams, C. K. I. 2006. *Gaussian Processes for Machine Learning*. MIT press Cambridge, MA.

Richter, S.; and Westphal, M. 2010. The LAMA Planner: Guiding Cost-Based Anytime Planning with Landmarks. *JAIR*.

Rivlin, O.; Hazan, T.; and Karpas, E. 2020. Generalized planning with deep reinforcement learning. In *Workshop on Bridging the Gap Between AI Planning and Reinforcement Learning*.

Seipp, J.; Keller, T.; and Helmert, M. 2020. Saturated Cost Partitioning for Optimal Classical Planning. *JAIR*.

Seipp, J.; and Segovia-Aguas, J. 2023. International Planning Competition 2023 - Learning Track. https://ipc2023-learning.github.io/.

Shen, W.; Trevizan, F.; and Thiébaux, S. 2020. Learning Domain-Independent Planning Heuristics with Hypergraph Networks. In *ICAPS*.

Shen, W.; Trevizan, F.; Toyer, S.; Thiébaux, S.; and Xie, L. 2019. Guiding Search with Generalized Policies for Probabilistic Planning. In *SoCS*.

Shervashidze, N.; Schweitzer, P.; Van Leeuwen, E. J.; Mehlhorn, K.; and Borgwardt, K. M. 2011. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research*.

Silver, T.; Dan, S.; Srinivas, K.; Tenenbaum, J. B.; Kaelbling, L.; and Katz, M. 2024. Generalized Planning in PDDL Domains with Pretrained Large Language Models. In *AAAI*.

Ståhlberg, S.; Bonet, B.; and Geffner, H. 2022. Learning General Optimal Policies with Graph Neural Networks: Expressive Power, Transparency, and Limits. In *ICAPS*.

Ståhlberg, S.; Bonet, B.; and Geffner, H. 2023. Learning General Policies with Policy Gradient Methods. In *KR*.

Ståhlberg, S.; Francès, G.; and Seipp, J. 2021. Learning Generalized Unsolvability Heuristics for Classical Planning. In *IJCAI*.

Ståhlberg, S. 2024. Revised Code and Models for the IPC 2023 Learning Track method "Muninn". https://doi.org/10.5281/zenodo.10688802.

Toyer, S.; Trevizan, F.; Thiébaux, S.; and Xie, L. 2018. Action Schema Networks: Generalised Policies with Deep Learning. In *AAAI*.

Vapnik, V. 2000. *The Nature of Statistical Learning Theory*. Statistics for Engineering and Information Science. Springer.

Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How powerful are graph neural networks? In *ICLR*.

Yoon, S.; Fern, A.; and Givan, R. 2002. Inductive policy selection for first-order MDPs. In *UAI*.