Self-directed Learning of Action Models using Exploratory Planning

Dustin Dannenhauer Matthew Molineaux

DUSTIN.DANNENHAUER@PARALLAXRESEARCH.ORG MATTHEW.MOLINEAUX@PARALLAXRESEARCH.ORG

Parallax Advanced Research, Beavercreek, OH USA

Michael W. Floyd

Knexus Research Corporation, National Harbor, MD USA

Noah Reifsnyder

Lehigh University, Bethlehem, PA USA

David W. Aha

Naval Research Laboratory, Washington, D.C. USA

MICHAEL.FLOYD@KNEXUSRESEARCH.COM

NDR217@LEHIGH.EDU

DAVID.AHA@NRL.NAVY.MIL

Abstract

Complex, real-world domains may not be fully modeled for an agent, especially if the agent has never operated in the domain before. The agent's ability to effectively plan and act in such a domain is influenced by its knowledge of when it can perform specific actions and the effects of those actions. We describe a novel exploratory planning agent that is capable of learning action preconditions and effects without expert traces or a given goal. The agent's architecture allows it to perform both exploratory planning and goal planning should be controlled, as well as how the agent's behavior should be explained to any teammates it may have. The contributions of this work include a new representation for contexts called Lifted Linked Clauses, a novel exploration action selection approach using these clauses, an exploration planner that uses lifted linked clauses as goals in order to reach new states, and an empirical evaluation in a scenario from an explorationfocused video game demonstrating that lifted linked clauses improve exploration and action model learning against non-planning baseline agents.

1. Introduction

Sufficiently complex domains, like those that may be encountered in the real world, are rarely fully defined and modeled in advance. Even minor changes in the domain, like weather conditions, can significantly impact how an agent is able to act. For example, the action preconditions for steering a vehicle in snow differ from the preconditions when the weather is clear. It is unrealistic to assume that a knowledge engineer would be able to provide the agent with information related to all possible domain variants. Instead, it would be advantageous for the agent to be able to dynamically learn how novel domains impact its ability to act and update its action model accordingly.

We present a self-directed exploratory planning approach to learning action models. Our planning agent is capable of operating in new domains, even if those domains differ noticeably from previously encountered domains. We do not assume the agent will have access to an expert-provided domain model or plan traces. Instead, we consider agents that perform exploration to obtain a collection of action interactions and uses those interactions for learning relational action models (i.e., preconditions and effects). We define interactions as triples containing the prior state, action, and post state.

A core concept for the exploration planning and exploration action selection is the notion of a Lifted Linked Clause (LLC), which we discuss more formally in Section 3. Briefly, a LLC is a representation of a context in a general manner. For example, a LLC context might be that an agent is in front of a closed door. Regardless of the exact state of the environment or exact type of door, the context of facing a doorway is relevant to which actions the agent can successfully carry out. We define LLCs as lifted subsets of the predicates in the state space where predicates share at least one variable across their arguments. LLCs allow us to represent situations in a general way, such as the agent is in a corner, or the agent is in the same location as an object, without having to refer to an exact state. LLCs often correspond to one or more preconditions for an action in the agent's transition model. By tracking which LLCs are active in the current state, and which actions have been taken when an LLC is active, the agent can prioritize reaching new states with active LLCs the agent has not seen before. This enables the agent to reach new states needed to learn a more complete action model, as our results show in Section 5. It also increases the rate at which the agent explores the domain.

The vision for the agents described here is not expected to be purely exploratory, but instead we expect these agents will operate as members of teams where other teammates may request the agents to pursue provided goals. At any time, the agents may be performing a mixture of exploration (i.e., attempting actions in different contexts), learning (i.e., learning action models), and goal pursuit (i.e., exploiting a learned model to achieve teammate-provided goals).

This work has five primary contributions:

- A formalization of the Lifted Linked Clause (LLC).
- A novel exploratory action selection approach that prioritizes taking actions that have been attempted the least in the currently active LLCs.
- A novel exploratory planning agent architecture that attempts to reach new states, using LLCs as goals.
- A formalism describing the algorithms behind the exploratory action selection and exploratory planning using LLCs.
- An evaluation in a scenario inspired from the Dungeon Crawl Stone Soup video game, including a random baseline.

The rest of the paper is organized as follows. In Section 2 we discuss related work, followed by a formalism of LLCs in Section 3. Section 4 presents the architecture of our agent including the

algorithms for key components, including the process of inductive learning of action models. We describe our evaluation in Section 5, results in Section 6 and conclude in Section 7.

2. Related Work

Since we are concerned with both exploration and learning relational action models for planning, we draw on prior work from both reinforcement learning (RL) and inductive learning. A difference of this work from RL is that our agent is goal-driven. Given that our agent does not know when actions can be applied (preconditions) and the changes made to the state (effects), we draw on ideas from RL to determine what actions to take in order to obtain enough examples for learning action models.

Hester and Stone (2017) discuss multiple exploration reward schemes and reward metrics in intrinsically motivated RL systems. They present two exploration-based intrinsic reward schemes that seek to take the most novel action. Briefly, action novelty takes into account the difference of the current state compared to the most recent state in which an action was executed as well as how uncertain the agents predictions for the state reached after executing the action.

Sequeira, Melo, and Paiva (2011) discuss emotions as a structure for intrinsic motivations which are balanced against extrinsic motivations. Numeric intrinisic motivations are based on appraisal dimensions including novelty, motivation, control, and valence. For example, novelty is proportional to the number of times an action has been used in a given situation before. The use of contexts in our agent is an attempt to reach a similar notion of novelty: taking actions that have not been tried in similar situations. Our agent prioritizes first to take actions that have been taken the least among all the currently active contexts (LLCs, see Section 3) and when all actions have been tried, the agent then prioritizes planning to reach states containing new contexts.

Inductive learning approaches that use heuristics have been effective (Vere, 1980; Hayes-Roth & McDermott, 1978; Watanabe & Rendell, 1990) but may not apply to all learning problems. FOIL (Quinlan, 1990) is a greedy algorithm that, like the heuristic approaches just mentioned, requires both positive and negative examples for training. Wang (1995) developed a system called OB-SERVER that uses expert traces of actions and a simulator for running practice problems. OB-SERVER does not need approximate action models; it learns action models from scratch that are as effective as human-expert coded action models assuming expert traces are available. Our learning problem is similar to those of OBSERVER and FOIL, hence our use of an inductive learning approach. There has been recent work on learning action models using genetic algorithms (Kučera & Barták, 2018) with example plans. The main difference in this work is that we focus on obtaining good examples for learning unknown action models, rather than on the learning algorithm itself. A benefit of the approach we describe is relaxing the need for expert traces of high quality plans.

To perform learning we use the Inductive Learning of Answer Set Programs (ILASP) system (Law et al., 2015). ILASP is a powerful inductive learning system that combines brave and cautious induction to learn answer set programs.

To deal with incompleteness of action models, Weber and Bryce (2011) give a planning algorithm that reasons about the action model's incompleteness in order to avoid failure during planning.

Incorporating such a technique into our agent for the goal-driven planner would allow the agent to perform planning prior to the learning of a complete action model (which may not ever happen).

Human-agent teaming, and specifically the need for explanation in such teams, has been identified as a potential challenge problem for Goal Reasoning (Molineaux et al., 2018). Their work presents a model for how explanations can be exchanged among teammates and used in decision making, and presents several common instantiations of the model. Our work falls under the **Single Supervisor** instantiation, where an agent has a supervisor that may make requests of it, and it may need to explain its reasoning or behavior.

Existing applications of explanation to Goal Reasoning have focused primarily on *internal explanations* (Aamodt, 1993) (i.e., for the benefit of the agent itself) rather than *external explanations* (i.e., for the benefit of an external user). DiscoverHistory (Molineaux et al., 2012) generates explanations for why observed domain transitions occurred. More specifically, it generates explanations that contain the external actions (i.e., of other agents) and exogenous event that are most likely to have resulted in the changes in the domain. In a human-robot teaming domain, such explanations have been used to allow a Goal Reasoning agent to determine the plans and goals of other agents, and reason over its own goals (Gillespie et al., 2015). However, even though these explanations are used by an agent that is a member of a team, they are only used internally and never provided to any teammates.

3. Lifted Linked Clause (LLC) Formalism

The approaches presented here use *LLCs* to guide the exploration process. A LLC c is a first-order relational conjunct $c_1 \wedge c_2 \wedge \ldots \wedge c_n$ that refers to one or more existentially quantified variables and is satisfied by some states $S' \subset S$. We describe a LLC c having n terms as being of size n (e.g. Figure 3 shows LLCs of size n = 3, 4, 5). LLC terms may include the following: finite relations over objects, finite equalities and inequalities between integers and functions over objects, and the infinite "successor" relation that describes consecutive integers. An upper bound¹ C on the number of possible LLCs for a given size n is given by

$$C = \binom{M! \times 2 \times |P|}{n}$$

where M is the maximum number of arguments for a single predicate $p \in P$. The 2 is because each predicate can be negated. We define the size n of a LLC as the number of relations (conditions) in the clause. In a LLC with multiple predicates, each predicate must refer to at least one existentially quantified variable that is also present in another term; this requirement means that no context of size n is semantically equivalent to the conjunction of two or more smaller contexts of size < n. For example, in the LLC of Figure 3a, T1 is such an existentially quantified variable. We describe the set of all legal LLCs for a domain Σ up to a size n as $clauses(\Sigma, n)$.

The constraint that each predicate must refer to at least one existentially quantified variable is hypothesized to reduce the space of potential LLCs and follows from an assumption that useful

^{1.} This upper bound assumes all predicates $p \in P$ have a number of arguments equal to the maximum number of arguments M, however usually some p have less than M arguments.

contexts will be defined by predicates that are relationally connected in some way, rather than any arbitrary collection of predicates. We leave for future work additional experiments to evaluate different methods for constraining the LLC space.

We say that an LLC is **active** if it unifies with the current state s (more than one unification is possible; as long as there is at least one unification, the LLC is considered active). The agent maintains $\mathcal{L}toA$, a count of the number of times each LLC action-pair $\langle c, a \rangle$ has occurred in the Interaction History, where \mathcal{L} refers to the set of LLCs up to size n, A refers to the set of actions, and $c \in \mathcal{L}$. This is given by the following function:

$$examples(c,a) = \left| \{ \langle s_i, a_i, s_{i+1} \rangle \in I \mid c \models s_i \land a_i = a \} \right|$$

Given a planning domain, it is possible to generate all LLCs up to a certain size (in our experiments we choose an LLC size of 2). Although the number of LLCs grows exponentially with n, even small LLCs of size 2 can be useful, as shown in our evaluation in Section 5.

4. Agent Architecture for Self-Directed Exploration

We now describe the agent architecture for our self-directed agent. We adopt the planning formalism from Ghallab et al. (2004) with a state transition system $\Sigma = (S, A, E, \gamma)$ where S is the set of all states, A is the set of actions, E is the set of events, and γ is the state-transition function. The architecture for our exploratory planning agent is shown in Figure 1.

This agent assumes a relational representation of states, observations, and actions are provided to it, but not the transition function. The agent attempts to learn the transition function over time.



Figure 1: Architecture for our exploratory planning agent

We separate the functionality of this agent into three submodules. The **exploration planner** is responsible for taking actions to obtain new information with which to update the action model,

Alg	orithm 1 Controller and Action Selection Pr	ocess
1:	procedure Controller()	
2:	Global: $s, \mathcal{L}toA$	
3:	while running do	
4:	$s \leftarrow observe()$	▷ Obtain current state
5:	$a \leftarrow ActionSelection()$	▷ Choose next action
6:		
7:	for c in \mathcal{L} do	
8:	if $active(c, s)$ then	
9:	$\mathcal{L}toA[c][a]$.increment()	\triangleright Record that this action was taken in this LLC c
10:	execute(a)	\triangleright Execute action a
11:	procedure ActionSelection()	
12:	Global: $s, \pi, \mathcal{L}toA$	
13:	if $\pi.length() > 0$ then	
14:	return $\pi.pop()$	▷ Retrieve next action of plan
15:	$lta \leftarrow leastTakenActions(s, \mathcal{L}toA)$	▷ Retrieve set of actions taken the least
16:		in the currently active contexts
17:	if $lta.length() > 0$ then	
18:	return random.choice(lta)	▷ Return a least taken action
19:	else	
20:	$\pi \leftarrow explorationPlanner()$	▷ Call planning
21:	if $\pi.length() > 0$ then	
22:	return $\pi.pop()$	▷ Retrieve next action of plan
23:	else	
24:	return randomAction()	▷ Take a random action

the **transition model learner**, responsible for updating the agent's model of the world, and the **controller**, responsible for determining when to perform learning (updating the model), planning, or if necessary, take a random action. Each submodule is designed in a domain-independent fashion, with data represented in a relational fashion. We now describe the current design of each submodule.

4.1 Controller

At each time step, the controller receives a state s_t from the domain, then generates a new action a_t to act in the domain. The controller and action selection procedures are given in Algorithm 1. Lines 3 to 10 describe a common runtime loop where the agent observes the current state (Line 4) and ultimately executes an action (Line 10), and continues indefinitely (Line 3). The exploration approaches (action selection in Algorithm 2 and planning in Algorithm 3) must know which actions have been executed in which contexts. This is recorded just before an action is executed by the controller in lines 7-9. The controller calls the action selection function in Line 5 which first checks if there are remaining actions in a plan to be executed, and if so, executes the next action (Lines 13-

14). If there are no remaining actions of the plan then the agent may either not be able to generate a plan or finds itself in a new state worth exploring. Line 15 retrieves all actions that have not been taken in currently active contexts (see Algorithm 2 for the definition of the leastTakenActions() procedure). If the current state is fully explored (in this case the condition on line 17 will be false) there will be no available actions to take, so the agent proceeds to attempt to generate a new plan (line 20) and execute the first action (lines 21-22). Finally, in some situations, the agent must be forced to take a purely random action which occurs on line 24.

The behavior that results for an agent in a new environment with no action model is that it takes random actions until it has learned a partial model. With each successful action taken, the agent finds itself in a different state than previously, and may have actions to take via exploration (lines 15-18). Eventually the agent's model will have learned enough to be able to generate a plan to reach a state with a context that has never been active before. While the plan may be incorrect, in the worst case it will generate new examples for learning, and in the best case it will enable the agent to reach a new state (and new context). From on observer's perspective, by attempting a plan, the agent goes from flailing randomly to directed movement towards a new state.

4.2 Exploration-based Action Selection using LLCs (part of Controller)

The exploration-based action selection is shown in Algorithm 2. When planning is not possible (such as when the agent may have no model of its action's preconditions or effects) or the agent has found itself in a new situation, this algorithm chooses actions that have not yet been taken in these contexts. Actions are prioritized by the number of contexts they have not been taken in. The algorithm begins by constructing a dictionary tracking the number of active contexts for each action taken (line 3). Lines 4 through 10 iterate over contexts and actions while recording the number of contexts the action has not been taken in. Line 11 creates the variable *leastActions* which keeps track of the actions taken least in currently active contexts. Line 12 introduces a variable tracking the maximum number of contexts an action has not been tried in. This is used to ensure that if there is a tie among actions for number of unknown contexts, then all actions that tie will be returned. Lines 14 through 21 iterate over the actions building the list of actions that have been tried the least. Finally, the actions are return on line 22.

Note that this algorithm is not guaranteed to return a non-empty set of actions. This happens when the agent has fully explored (i.e. executed every action in all of the currently active contexts) the current state. When this occurs the controller performs planning to reach a new state with contexts that have not been fully explored.

4.3 Exploration Planner

Given an initial state s, an interaction history $(\langle s_0, a_0, s_1 \rangle, \langle s_1, a_1, s_2 \rangle, ...)$, a domain model of the transition function λ , and a goal LLC g, the exploration planner finds a plan π to reach any state that satisfies the goal.

The role of the Exploration Planner is to obtain novel *interactions*. An interaction is defined as a triple $\langle s_i, a_i, s_{i+1} \rangle$ where $s_i \in S$ is the previous state before taking action $a_i \in A$ and $s_{i+1} \in S$ is the state after a_i is executed. As the agent acts in the domain all interactions are recorded and

Algorithm 2 Exploration-based Action Selection using LLCs

```
1: procedure leastTakenActions()
        Global: \pi, s, \mathcal{L}, \mathcal{L}toA
 2:
        actionTakenCounts \leftarrow Dict()
 3:
        for c in \mathcal{L} do
 4:
 5:
            if active(c, s) then
                for a in A do
 6:
                    if a \notin actionTakenCounts then
 7:
                        actionTakenCounts[a] \leftarrow 0
                                                                               ▷ Start tracking this action
 8:
                    if \mathcal{L}toA[c][a] == 0 then
                                                                 \triangleright True if a never taken when c is active
 9:
                        actionTakenCounts[a].increment()
10:
        leastActions \leftarrow []
11:
        max \leftarrow 0
                                                           ▷ Score enables tracking number of contexts
12:
13:
                                                                that are active and compatible per action
14:
        for a in A do
            if a \in actionTakenCounts and actionTakenCounts[a] \neq 0 then
15:
                if actionTakenCounts[a] > max then
16:
                    max = actionTakenCounts[a]
17:
                    leastActions \leftarrow [a]
                                                              ▷ Reset leastActions and insert this action
18:
19:
20:
                if actionTakenCounts[a] == max then
                    leastActions.append(a)
21:
22:
        return leastActions
```

added to the Interaction History I. To determine the next action, the Exploration Planner chooses the action a that minimizes examples(c, a) for all active legal contexts of a size less than or equal to n:

```
\mathop{\arg\min}_{a\in A}\min_{c\in contexts(\Sigma,n)}examples(c,a)
```

We collect complete states from the domain after each action is executed. The Interaction History provides training examples that are later used by the Transition Model Learner to learn a Domain Model. Thus, by guiding the agent's actions, the Exploration Planner can impact which interactions are stored in the Interaction History and used for learning.

4.4 Transition Model Learner

We are interested in agents that can learn their action model online while operating in a new domain. For our Transition Model Learner component, we assume that action predicates and argument types are known to the agent a priori (i.e., provided as expert domain knowledge) as well as a model of the domain that includes objects, types of these objects, and predicates for these objects. Currently we assume the world is fully observable and deterministic; we will relax these constraints in future work.

As we discussed previously, every action the agent takes is recorded into a triple $\langle s_i, a_i, s_{i+1} \rangle$ and stored in the Interaction History, which can be used later for learning. We have taken an inductive learning approach, inspired by prior work on learning action models. The primary difference in our work is that the agent is not provided with expert traces of action sequences and instead directs its own behavior to acquire examples of executing actions.

For the inductive learning algorithm, we use the Inductive Learning of Answer Set Programs (ILASP) system (Law et al., 2015) to learn rules that act as preconditions of individual actions. We assume that if an action did not change the state of the world, then it failed and use that assumption to label interactions as *positive interactions* that succeeded ($s_i \neq s_{i+1}$ after performing a_i) or *negative interactions* that failed ($s_i = s_{i+1}$ after performing a_i). Additionally, ILASP is given inductive biases describing the types of predicates that should be considered in the body (in the form of **#modeb** declarations), which are translated from the predicates and arguments types from our domain model of our scenarios used in the evaluation.

An example for a precondition rule generated by the ILASP system is shown in Figure 2. This rule encodes the learned preconditions for the *southeast* move action for a grid-world domain. The various values used in the rule are:

- X: The X coordinate of the tile the agent is currently located in
- *Y*: The Y coordinate of the tile the agent is currently located in
- X2: The X coordinate of destination tile
- *Y2*: The X coordinate of destination tile

In our current setup, we support learning the eight cardinal directions and the only change in the domain that occurs from executing actions is the agent's location or opening / closing a door.

Before the agent performs planning, the agent makes an external call to the ILASP system to perform learning on all actions that have newly recorded interactions. External calls to ILASP may take up to a few hours to complete, so limiting the number of calls to the ILASP planner should be prioritized. For future work there are a number of ways to reduce ILASP run time including selecting only a subset of interactions for ILASP to learn over. Currently, every interaction, including hundreds of failed interactions because of invalid arguments (i.e. the agent attempts to move to a non-adjacent tile) are given to ILASP. Additionally, ILASP is under active development and future versions may significantly speed up run time.

The Transition Model Learner uses an inductive learning procedure to learn the preconditions of each of the agent's actions (i.e., update the Domain Model with the actions' preconditions). An

southeast(X, Y) :- agentat(X2, Y2), north(Y2, Y), west(X2, X), not wall(X, Y), not cdoor(X, Y).

Figure 2: Learned rule from ILASP for the southeast move action in a tile grid world environment.

Algorithm 3 Exploration-based Planning using LLCs as Goals 1: procedure LLC-PLANNER 2: **Global:** $s, A, I, \mathcal{L}toA$ $A \leftarrow learnActionModels(I)$ 3: ▷ Perform learning to ensure up-to-date action models $G_{done} \leftarrow []$ ▷ Store goals already attempted 4: 5: do $G_{remaining}, G_{done} \leftarrow contextsWithLeastActions(\mathcal{L}toA, G_{done})$ \triangleright Retrieve LLCs 6: 7: for g in $G_{remaining}$ do $\pi \leftarrow ASP$ -Planner(A, s, g)▷ Call planner 8: 9: if π .length() > 0 then return π ▷ Return first plan found 10:

while $G_{remaining}.length() > 0$

11:

12:

return []

example of the learned preconditions for the *southeast* move action from a tile grid world exploration domain is shown in Figure 2. The learning component is only concerned with learning action preconditions. Once preconditions are learned, the effects of the actions can be computed by finding the difference between s_{i+1} and s_i and unifying the differences so that the argument variables in the effects clause are consistent with the variables in the preconditions clause.

The intuition behind using contexts for exploration is that the same action needs to be executed in multiple contexts in order to learn the complete precondition rule, and also effects, for that action. Consider the preconditions for *southeast* shown in Figure 2. In order for the agent to learn the condition *not wall(V0)*, it needs to execute the action in a context where there is a wall southeast of the agent, thereby preventing the action from succeeding. This produces a negative interaction. To learn positive interactions, the agent must execute the action in contexts where there is not a wall in a cell to the southeast. These positive interactions enable learning the other preconditions besides *not wall(V0)*.

{agent-at(T1), at(T1,X1,Y1), at(T2,X2,Y2)} (a) LLC of size 3

{agent-at(T1), at(T1,X1,Y1), at(T2,X2,Y2), not wall(T2), X2 = X1-1} (c) LLC of size 5

Figure 3: Example Lifted Linked Clauses of varying sizes

SELF-DIRECTED LEARNING OF ACTION MODELS



(b) PDDL Representation

Figure 4: Scenario 1 Visualization and PDDL Representation

5. Evaluation

We perform an evaluation of our self-directed planning agent in an ablation study in a scenario inspired from the exploration-based genre of rogue-like video games. In these games humans must figure out when (preconditions) and how (effects) actions change the game state to explore an unknown environment that is partially observable and open. Figure 4 depicts both a graphic visualization and a relational representation of this scenario. We use the Planning Domain Definition Language (PDDL) to represent the state and actions (McDermott et al., 1998). The PDDL representation of each scenario shown is exactly what is given to our agents as the starting state. Agents are given the actions they can take and typed arguments for these actions without any preconditions or effects.

The scenario here is static and fully observable. This allows us to label each interaction $\langle s_i, a_i, s_{i+1} \rangle$ as positive if $s_i \neq s_{i+1}$ and negative if $s_i = s_{i+1}$. We leave evaluations in dynamic domains for future work. Briefly, a challenge in a dynamic environment (which is the case for most rogue-like video games) is to identify whether a change in the environment can be attributed to the agent or an external actor. Only in the cases when a change in state occurs and is likely to be caused by the agent should the interaction be considered positive.

Each agent has 24 possible actions consisting of eight movement actions, eight open door actions and eight close door actions. The eight variations of each action type (movement, open door, close door) is because each action is for one of the eight cardinal directions: N, S, E, W, NE, NW, SE, SW.

5.1 Exploration Progress Metric

We measure how much of the scenario each agent has explored by counting the number of unique tiles they visited. Visiting more unique tiles means that an agent has explored more of the domain and is a partial indicator of the likelihood of that agent's model being accurate, since without visiting certain tiles, some actions can never be accurately learned (i.e. visiting adjacent tiles to the closed door tile are needed to learn the *open-door* action). Scenario 1 has 33 unique tiles the agent can occupy. These include the closed door tile since if the door is open, the agent can then move into that tile.

5.2 Learning Accuracy Metric

Since the agents are learning action preconditions from scratch, we develop a metric that compares the learned action models from perfect action models in such a way that the rules do not need to be equivalent regarding the logical representation. Instead, we test each learned precondition for an action against perfect preconditions on a number of test states and compare whether the output of the learned rule (i.e. whether the action can be taken in the test state) is the same as the perfect rule. In this case the output is either that the preconditions hold in which case the action can be taken or the preconditions do not hold and the action will fail to change the state. The perfect action models were hand authored: movement actions contain preconditions to ensure there is a destination cell in that direction and that a wall is not in the destination, closing and opening door actions contain the preconditions there is a destination cell and that there is an open or closed door in the destination, respectively.

For this evaluation we used 16 expert authored test states. Eight of these states involved a closed door and eight involved an open door. In all test states some movement actions are possible while others are not. The closed door test states are shown in Figure 5. W represents a wall, A represents the agent's location, and CD represents a tile with a closed door. Empty tiles are spaces that the agent can move into. An agent may move into a door tile only if the door is open.

For example, when applying perfect action preconditions on this state the agent is able to successfully execute the following actions: *move south, move east, and close door southeast*. All other actions when applied to this state will fail. Then we apply the learned preconditions for each action. We use the results (i.e. whether or not the perfect/learned rule can be executed in that state) from the test states to calculate precision, recall, and F1 scores between 0 and 100% for each actions' learned preconditions. The benefit of this metric is that we can measure learning performance irrespective of the rule's representation.

SELF-DIRECTED LEARNING OF ACTION MODELS

										_										
w	W	w	w	W	W	w	w	w	w		w	w	w	w	w	w	w	w	w	w
w	А			w	w		А		w		w			Α	w	w				w
w		CD		w	w		CD		w		w		CD		w	w		CD	Α	w
w				w	w				w		w				w	w				w
w	w	w	w	W	w	w	w	w	w	1	w	w	w	w	w	w	w	w	w	w
										-										
w	w	w	w	w	w	w	w	w	w		w	w	w	w	w	w	w	w	w	w
w				w	w				w		w				w	w				w
w	А	CD		w	w		CD		w		w		CD		w	w		CD		w
w				W	w	Α			w		w		А		w	w			Α	w
w	w	w	w	w	w	w	w	w	W		w	W	w	w	w	W	w	w	w	w

Figure 5: Closed Door Test States Used for Learning Accuracy Metric

5.3 Hypotheses

Our evaluation aims to test whether these agents can learn action preconditions and effects from scratch when situated in a previously unseen domain and if there is a benefit to our exploration planner using LLCs. We hypothesize that:

- H1 All agents will be able to learn some action preconditions from collected interactions.
- **H2** The LLC one-step action selection agent will learn better preconditions for more actions than a baseline agent taking only random actions.
- **H3** The LLC one-step action selection agent will explore more of its environment than the random baseline agent.
- **H4** The LLC exploratory planning agent will learn better preconditions for more actions than both the one-step LLC action selection agent and the random baseline agent.
- **H5** The LLC exploratory planning agent will explore more of its environment than both the one-step LLC action selection agent and the random baseline agent.

H1 relates to our claim that the approaches described here can be used to learn action models by collecting interactions, even if the agent is a random baseline. The rational behind H2 and H3 is that because the one-step LLC agent maintains information about which actions have been taken in which contexts (LLCs) and prioritizes taking actions that have never been taken in particular contexts, it will obtain the needed interactions for learning an action quicker than the random baseline. We expect H4 and H5 to be true because the agent will eventually choose to plan to reach new states, and this planning will be more efficient than only considering one action ahead.



Figure 6: Exploration Progress Per Approach for Scenario 1

6. Results

Our results show the performance of three agents (random, one-step LLC, and planning LLC) in a scenario in Figure 4. The performance of the agents displayed in Figure 6 are averaged over three runs per agent for 4,000 actions. In each run the agent starts with no preconditions or effects for any of its actions. That is, each run is independent and no knowledge is maintained between runs.

6.1 Scenario 1 Results

The graph in Figure 6 shows the exploration progress of the agent per action executed. Solid lines are composed of the average data point over three runs, where each data point represents the cumulative total number of unique tiles visited by that agent since the beginning of the run. The shaded area shows the maximum and minimum values between all the runs. The larger the shaded area, the more variance there is among the runs.

The random agent explores the slowest initially (within the first 500 actions) and manages to perform as well as the explore local agent (using one-step LLC action selection) until about 2500 actions. Then the explore local agent explores slightly more than the random agent. This shows that our hypothesis **H3** is true in Scenario 1 although only minimally.

The planning agent in this figure has a slight delay in number of tiles executed but then quickly explores the environment much faster than the other two agents. Even though there is fairly high variance between the runs, even the worst case is still almost entirely better than the other two agents, and this variance reduces by the time the agent has executed

	F	Randor	n]	Explor	e	Planning			
	Р	R	F1	Р	R	F1	Р	R	F1	
move_w	88	100	93	88	100	93	100	100	100	
move_e	94	100	97	94	100	97	100	100	100	
move_n	0	0	0	67	29	41	100	100	100	
move_s	0	0	0	62	67	65	100	100	100	
move_nw	33	23	27	67	46	54	65	52	58	
move_ne	0	0	0	0	0	0	98	77	86	
move_sw	0	0	0	0	0	0	56	50	53	
move_se	31	33	32	67	46	54	60	48	53	
close_door_w	0	0	0	0	0	0	33	33	33	
close_door_e	0	0	0	0	0	0	67	67	67	
close_door_n	0	0	0	0	0	0	0	0	0	
close_door_s	0	0	0	0	0	0	67	67	67	
close_door_nw	0	0	0	0	0	0	0	0	0	
close_door_ne	0	0	0	0	0	0	0	0	0	
close_door_sw	0	0	0	0	0	0	0	0	0	
close_door_se	0	0	0	0	0	0	0	0	0	
open_door_w	0	0	0	0	0	0	0	0	0	
open_door_e	0	0	0	0	0	0	33	33	33	
open_door_n	0	0	0	0	0	0	0	0	0	
open_door_s	0	0	0	0	0	0	96	100	98	
open_door_nw	0	0	0	0	0	0	0	0	0	
open_door_ne	0	0	0	0	0	0	0	0	0	
open_door_sw	0	0	0	0	0	0	0	0	0	
open_door_se	0	0	0	0	0	0	0	0	0	

Table 1: Precision, Recall and F1 Score per action model learned by each approach in Scenario 1

approximately 3000 actions (shown by the shaded area reducing in size). This result shows that our hypothesis **H5** is significantly true in Scenario 1. Finally, there are 33 possible tiles the agents' could visit in Scenario 1, which shows that the planning agent almost explores every location using the same number of actions it takes for the explore local and random agents to visit 7 to 5 tiles respectively.

Turning to Table 1 we see that the random agent was able to learn partial models for four actions, the explore local agent learned partial models for 6 actions, and the planning agent learned models for 13 actions, four of which were fully correct models. The explore local agent was able to learn 2 more models than the random agent, and learned more correct models in all cases where random learned a partial model of the same action, which verifies hypothesis **H2** in Scenario 1. Hypothesis **H4** is almost entirely true for Scenario 1, except when regarding the move southeast action, where the explore agent was able to

achieve an F1 score above the planning agent by 1%. This is due to the explore agent learning a rule with higher precision than the rule learned by the planning agent. Since the planning agent learned significantly more action models than explore, we conclude that this shows hypothesis **H4** to be true in general.

7. Conclusions

We introduce a novel formalism of a situational context which we call a Lifted Linked Clause. LLCs can be used as goals by an agent to learn a more accurate action model by attempting to generate a plan to reach a state with a LLC that has never been seen before. Our evaluation shows that an agent can more efficiently explore a space using LLCs as goals than without such a planning mechanism (i.e. against a non-planning explorer or a random agent). We show that both exploration of the state space and learning accuracy are improved by using such an approach. These agents are able to learn in a single run as soon as they gather a small number (two or more) examples per action for learning. A primary motivation for learning such relational models is that it would enable increased explainability for autonomous planning and acting agents.

Acknowledgements

We thank Michael T. Cox for his comments and support in drafting this article. Additionally, we thank DARPA for supporting this research. The views, opinions and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government.

References

- Aamodt, A. (1993). Explanation-driven case-based reasoning. Proceedings of the First European Workshop on Case-Based Reasoning (pp. 274–288). Kaiserslautern, Germany: Springer.
- Ghallab, M., Nau, D., & Traverso, P. (2004). *Automated planning: Theory and practice*. Elsevier.
- Gillespie, K., Molineaux, M., Floyd, M. W., Vattam, S. S., & Aha, D. W. (2015). Goal reasoning for an autonomous squad member. *Proceedings of the Workshop on Goal Reasoning (held at the 3rd Conference on Advances in Cognitive Systems)* (pp. 52–67).
- Hayes-Roth, F., & McDermott, J. (1978). An interference matching technique for inducing abstractions. *Communications of the ACM*, 21, 401–411.
- Hester, T., & Stone, P. (2017). Intrinsically motivated model learning for developing curious robots. *Artificial Intelligence*, 247, 170–186.

- Kučera, J., & Barták, R. (2018). Louga: Learning planning operators using genetic algorithms. *Pacific Rim Knowledge Acquisition Workshop* (pp. 124–138). Springer.
- Law, M., Russo, A., & Broda, K. (2015). The ILASP system for learning answer set programs. https://www.doc.ic.ac.uk/~ml1909/ILASP.
- McDermott, D., Ghallab, M., Howe, A., Knoblock, C., Ram, A., Veloso, M., Weld, D., & Wilkins, D. (1998). Pddl-the planning domain definition language.
- Molineaux, M., Floyd, M. W., Dannenhauer, D., & Aha, D. W. (2018). Human-agent teaming as a common problem for goal reasoning. *Proceedings of the AAAI Spring Symposium on Integrating Representation, Reasoning, Learning, and Execution for Goal Directed Autonomy*. AAAI Press.
- Molineaux, M., Kuter, U., & Klenk, M. (2012). DiscoverHistory: Understanding the past in planning and execution. Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (pp. 989–996). Valencia, Spain: IFAAMAS.
- Quinlan, J. R. (1990). Learning logical definitions from relations. *Machine learning*, 5, 239–266.
- Sequeira, P., Melo, F. S., & Paiva, A. (2011). Emotion-based intrinsic motivation for reinforcement learning agents. *Proceedings of the International Conference on Affective Computing and Intelligent Interaction* (pp. 326– 336). Springer. From https://pdfs.semanticscholar.org/28e6/ 3b995213efe0eac8169eae3ae8ffe6eb04e3.pdf.
- Vere, S. A. (1980). Multilevel counterfactuals for generalizations of relational concepts and productions. *Artificial intelligence*, *14*, 139–164.
- Wang, X. (1995). Learning by observation and practice: An incremental approach for planning operator acquisition. *Proceedings of the 12th International Conference on Machine Learning* (pp. 549–557). From http://pdf.aminer.org/000/335/ 904/learning{_}by{_}observation{_}and{_}practice{_ }an{_}incremental{_}approach{_}for{_}planning.pdf.
- Watanabe, L., & Rendell, L. A. (1990). Effective generalization of relational descriptions. *Proceedings of the 8th National Conference on Artificial Intelligence* (pp. 875–881).
- Weber, C., & Bryce, D. (2011). Planning and Acting in Incomplete Domains. *Proceedings* of the International Conference on Automated Planning and Scheduling (pp. 274–281).