

# MACQ: A Holistic View of Model Acquisition Techniques

Ethan Callanan\*<sup>1</sup> Rebecca De Venezia\*<sup>1</sup> Victoria Armstrong<sup>1</sup> Alison Paredes<sup>1</sup>  
Tathagata Chakraborti<sup>2</sup> Christian Muise<sup>1</sup>

<sup>1</sup> Queen’s University, <sup>2</sup> IBM Research AI  
Contact: christian.muise@queensu.ca

## Abstract

For over three decades, the planning community has explored countless methods for data-driven model acquisition. These range in *sophistication* (e.g., simple set operations to full-blown reformulations), *methodology* (e.g., logic-based -vs- planing-based), and *assumptions* (e.g., fully -vs- partially observable). With no fewer than 43 publications in the space, it can be overwhelming to understand what approach could or should be applied in a new setting. We present a holistic characterization of the action model acquisition space and further introduce a unifying framework for automated action model acquisition. We have re-implemented some of the landmark approaches in the area, and our characterization of all the techniques offers deep insight into the research opportunities that remain; i.e., those settings where no technique is capable of solving.

**Project:** [macq.planning.domains](#)

## 1 Introduction

Model acquisition has been a hallmark sub-field of automated planning for decades. From the early approaches to extracting simple planning models in an automated way (Shen and Simon 1989) to modern techniques for extracting lifted action theories from state tokens alone (Rodriguez et al. 2021), the planning community has developed dozens of approaches to model acquisition for various settings. Simultaneously, there have been a few surveys that capture the status of this rich sub-field and characterize the approach across many axes (Celorrio et al. 2012; Jilani et al. 2014; Arora et al. 2018). What is missing from the collective focus on the suite of techniques available is a unifying framework that allows researchers to explore both the existing techniques, as well as the gaps in what we are capable of solving. While the explosion in model learning techniques has dated past surveys pretty quickly, our open-source approach and live interface to the repository aims to make MACQ the one-stop-shop for learning planning models.

In this work, we present the first major step towards such a unified framework. Both from a theoretical standpoint – characterizing model acquisition techniques in terms familiar to the planning community – and from the practical

standpoint with growing implementation of the core algorithms. Our contribution goes substantially further than just a survey of the existing techniques. By integrating them under a single theoretical framework and implementation, we effectively open the door to systematically characterizing the entire sub-field of research.

We accomplish this characterization by identifying the key properties that distinguish model acquisition techniques and then applying knowledge compilation techniques to map out what is known from the field. We are thus able to holistically view what is possible theoretically (i.e., papers exist to address the setting), what is possible practically (an implementation exists in the framework), and what remains an open research question.

Finally, having the suite of tools for state trace generation, modification, and analysis all under a single framework allows us to rapidly test different approaches to a specific setting. New use-cases for model acquisition can leverage the growing library of implemented techniques for model acquisition, and this provides tangible benefits to those outside the planning community who wish to apply planning techniques in their domain of expertise.

## 2 MACQ Framework

There are three core components to the MACQ framework: (1) trace generation; (2) observation tokenization; and (3) model extraction. Not all of these are mandatory, but each complements the others to offer a rich array of functionality. We discuss each of them in turn in Section 2, but lead this section with a discussion of key features and assumptions of concern in describing a model learning task.

### 2.1 Feature Analysis

A model learning task involves three key considerations – 1) what are the features of the agent whose model is being learned; 2) what are the features of the model being learned; and 3) what are the features of the data from which the model is being learned. Note the distinction between (2) and (3). This separation of features allows our framework with the flexibility to provide the user with the choice of what kind of model to learn from what kind of data. For example, a user can choose to use a stochastic model extraction technique on noisy data instead of modelling noise directly. More on this “token casting” feature later in Section 2.3.

---

\*These authors contributed equally.

## Agent Features

- **Rationality** The primary consideration here is whether the observed agent is rational or not. By default, we assume that this is unknown or that there is no assumption of rationality in an extraction technique. If there is, we allow for two types of rationality – optimal traces (in the classical sense) or causally relevant traces where there are no redundant actions in a plan i.e. there are no (subsets of) actions that can be removed from the trace and the agent can still reach its goal.

## Model Features

- **Uncertainty** The first feature of interest is whether uncertainty is captured by the model – a model can take three forms: 1) deterministic, 2) non-deterministic, and 3) probabilistic. Note that (3) implies (2) and (1); and (2) implies (1). Thus, a model can have at least one and at most one of these features – this is generally not true for the rest of the features.
- **Parameterization** These features capture whether actions (as well as predicates) in the model are parameterized by the objects they operate on and whether those objects in turn are typed or atomic.

## Data Features

- **Observability** One of the primary features of concern in a model learning task is how much of the environment is observable. The fluents describing the state may be 1) fully or 2) partially observable or 3) not observable at all. Similar to the discussion on model uncertainty, the ability of a model extraction technique to deal with (2) implies (1); but (3) implies neither of (1) or (2).
- **Action Observability** In addition to whether the fluents are observable, an additional consideration is whether the action labels and parameters are known and whether there is a seed model, to begin with.
- **Parameterization** This mirrors the model features by the same names. Additionally, it also includes the possibility to have action costs.
- **Noise** Noise in data may manifest either as noise in fluent observations or noise in observed action labels.
- **Access** to the initial and goal state of a trace.
- **Trace** Finally trace-level features include whether there is access to the cost of the plan and to what extent (full or partial) the trace is observable i.e. whether there are missing actions or not.

## 2.2 Trace Generation

Plan traces may be sourced from a variety of sources, and the first component (trace generation) serves as both a rich set of techniques to *generate* planning traces as well as a suite of tools to parse and process existing data for model extraction. Here, we describe some of the highlights in functionality for the “trace generation” part of the MACQ ecosystem (one of which is currently under development).

**CSV Processing** As a trivial baseline for trace generation, the MACQ library offers functionality to load and package up simple CSV files. Columns with the values 0 or 1 will be retained (presumed to be boolean fluents), and a single specially designated column for the “action” label is required. This does not cover the full set of data format assumptions (e.g., parameterized actions), but is nonetheless a common starting point for many model acquisition tasks.

**Statespace Sampling** Given any valid classical planning problem, MACQ has the ability to generate random statespace trajectories. The PDDL model supplied (either as raw PDDL, pointers to existing files, or a reference to the [api.planning.domains](#) problem ID (Muisse 2016)) is parsed by the TARSKI library (Francés, Ramirez, and Collaborators 2018), and actions are selected (uniformly at random) starting in the initial state. The goal is ignored in this case, and MACQ will generate the predefined number of traces at a given length. As an added feature, we have also incorporated the methodology adopted by the FastDownward system (Helmert 2006) to perform random state sampling to a depth based on heuristic computation.<sup>1</sup>

**Goal-Oriented Sampling** Another approach for generating a single trace is to compute a plan for the given domain/problem pair. However, as a generation technique, it is limited to generating just a single trace. MACQ has expanded on this by allowing for random goal sampling. The approach works as follows:

1. Use the statespace sampling to compute a sequence of actions/states of length  $k$ .
2. Sample subsets of fluents of size  $g$  from the final state in this sequence such that,
  - (a) It reflects the same type of fluents in the original goal.
  - (b) It is not easily achieved from the original initial state.
3. Use those sampled subsets as goals for computing a plan.

There are several design decisions to be made: user-selected values for  $k$  and  $g$  (which may be domain-specific); sampling procedure in step 2 to adhere to 2(a); measuring the quality of a goal candidate in step 2(b); etc. MACQ currently includes a preference to use fluents corresponding to the goal predicates for step 2(a), and uses a computed plan as a proxy for 2(b) – the closer the found plan is to length  $k$ , then the better the goal candidate is. This technique provides a rich mechanism for data generation given a single seed planning problem. Not only is it useful for exploring model extraction techniques, but we predict it may have wider use in the area of planning and learning.

## 2.3 Observation Tokenization

There is a vast array of model acquisition techniques that exist (some surveys on the space are discussed in Section 1). In an effort to provide a common foundation for a library that encompasses this rich variety, we appeal to the notion of *observation tokens* (Geffner and Bonet 2013). To unify various

<sup>1</sup>Details: [https://github.com/aibasel/downward/blob/main/src/search/task\\_utils/sampling.cc](https://github.com/aibasel/downward/blob/main/src/search/task_utils/sampling.cc)

approaches for planning with partial observability, [Geffner and Bonet](#) introduces a notion of an observation token. It succinctly captures what the agent sees and can act upon. At times, this may be an indication of a sensing outcome. At other times, it may capture the partial state information viewable by the agent. We adopt this concept wholeheartedly for use in `MACQ`.

Every extraction method works on a set of observation token lists of a specific token type.

**Observation Token Types** Driven by the variety of extraction methods captured by `MACQ`, we have identified several useful forms of observation tokens. These include:

- `IDENTITY`: Full state / action information is provided.
- `PARTIALSTATE`: Some fluents are masked as unknown.
- `STATE ID`: No action or fluent information is provided – same states correspond to the same token.
- `NOISYSTATE`: Some fluents may be incorrectly assigned.
- `ACTIONONLY`: No state information is provided.

This list, although not exhaustive, provides a sense of the variety embedded within `MACQ`. Each extraction method will designate the token types it is capable of processing, and this allows for the automatic discovery of methods applicable to data of a particular form. This means that `MACQ` has the ability to automatically suggest the extraction methods that can work with a particular source of trace data.

**Tokenization** To explore the effectiveness of extraction techniques, every form of observation token has the functionality to “tokenize” ground-truth data. For example, `PARTIALSTATE` tokens can be created by specifying the likelihood of masking a fluent, along with the set of fluents that are eligible (defaulting to the entire state).

This functionality is essential for the development of new extraction methods, as well as testing pre-existing ones. Combined with the methods for trace generation in `MACQ`, this provides a robust means for data generation in the model acquisition space.

**Token Casting** The extraction techniques are tightly coupled with the type of representations they can handle. An approach for partially observable states will only operate on the appropriate class of `PartiallyObservable` token types for the trace. To extend this, we allow for “token casting”, which will (if possible) transform tokens of one type to another. Taking our example further, if we have fully observable states and we wanted to test a technique implemented in `MACQ` dedicated to partially observable state spaces (because of other functionality it offers), we can “token cast” the trace data to partially observable tokens.

Token casting is not always feasible, but the `MACQ` framework is set up such that finding these casting paths is naturally available. Every contribution to the space of model acquisition is characterized by the limited scope of the paper/work. Through methods like tokenization and token casting, we open the door to applying techniques in a richer variety of settings, not originally envisioned by the authors.

## 2.4 Action Model Extraction

The bulk of the `MACQ` project is dedicated to the extraction of action theories. At the time of writing, we have re-implemented a representative sample of model acquisition techniques spanning several features discussed in Section 2.1 to demonstrate the potential of the `MACQ` system:

- `OBSERVER` ([Wang 1994](#)): One of the first and most simple methods for model acquisition, this technique assumes full observability, noise-free data, and deterministic actions. The extracted theories are in STRIPS form, and the methods are mostly set-based.
- `ARMS` ([Yang, Wu, and Jiang 2005, 2007a](#)): This line of work handles partially observable states (fluents are hidden) as well as traces (entire states may be missing). There is an assumption that the goal and initial states are known, and while the actions and predicates are parameterized, they are not typed. The technique uses MaxSAT to solve a particular encoding for model extraction.
- `SLAF` ([Amir and Chang 2008](#)): This method also focuses on partially observable state traces and appeals to a logical encoding to extract the action theories. It is capable of producing one action theory, or several that fit the data. It operates by iteratively calling a SAT solver to find the entailments that lead to a valid theory (i.e., one that adheres to the observations). The input traces are assumed to be noise-free.
- `AMDN` ([Zhuo, Peng, and Kambhampati 2019](#)): This technique relies on MaxSAT to find the most likely action models given potentially disordered and noisy plan traces (here, noise means actions may be out of order). Further, the states may be partially observable and noisy, and the actions occur in parallel.

A common element of these approaches is the heavy reliance on SAT or MaxSAT technology. Because of this commonality, several elements of functionality have been included in the `MACQ` project for extraction techniques to make use of. These include model building, solving with wrapped binaries, and solution extraction. Specifically, the project relies on the Bauhaus ([Daga and Muise 2021](#)), `python-nnf` ([Verbeek, de Haan, and Muise 2022](#)), and `PySAT` ([Ignatiev, Morgado, and Marques-Silva 2018](#)) libraries, as well as the `kissat` ([Biere et al. 2020](#)) and `RC2` ([Morgado, Dodaro, and Marques-Silva 2014](#)) SAT / MaxSAT solvers. While the current focus is SAT-based, the full scope of model acquisition techniques are planned for eventual implementation.

Finally, custom and flexible representations for learned actions and fluents are shared across the approaches. This allows for a uniform treatment of what is produced – regardless of extraction technology – and further allows us to efficiently generalize the serialization of the action theories. This final step (writing to PDDL) is achieved using the Tarski library ([Francés, Ramirez, and Collaborators 2018](#)).

## 3 `MACQ` in Action

Here, we briefly showcase some of the interface for the toolkit, as well as the summary view of the work.

---

```

1 from macq import generate, extract
2 from macq.trace import PlanningObject, Fluent, TraceList
3 from macq.observation import PartialObservation
4
5 def get_fluent(name: str, objs: list[str]):
6     objects = [PlanningObject(o.split()[0], o.split()[1])
7                 for o in objs]
8     return Fluent(name, objects)
9
10 traces = TraceList()
11 generator = generate.pddl.TraceFromGoal(problem_id.1801)
12
13 generator.change_goal({
14     get_fluent(
15         "communicated_soil_data",
16         ["waypoint waypoint2"]
17     ),
18     get_fluent(
19         "communicated_rock_data",
20         ["waypoint waypoint3"]
21     ),
22     get_fluent(
23         "communicated_image_data",
24         ["objective objective)", "mode high_res"]
25     ),
26 })
27 traces.append(generator.generate_trace())
28
29 generator.change_goal({
30     get_fluent(
31         "communicated_soil_data",
32         ["waypoint waypoint3"]
33     ),
34     get_fluent(
35         "communicated_rock_data",
36         ["waypoint waypoint2"]
37     ),
38     get_fluent(
39         "communicated_image_data",
40         ["objective objective)", "mode high_res"]
41     ),
42 })
43 traces.append(generator.generate_trace())
44
45 observations = traces.tokenize(
46     PartialObservation,
47     percent_missing = 0.60
48 )
49 model = extract.Extract(
50     observations,
51     extract.modes.ARMS,
52     upper_bound = 2,
53     min_support = 2,
54     action_weight = 110,
55     info_weight = 100,
56     threshold = 0.6,
57     info3_default = 30,
58     plan_default = 30,
59 )
60 print(model.details())

```

---

Figure 1: MACQ usage

```

Actions:
(communicate_soil_data waypoint
                                lander
                                rover
                                waypoint):

precond:
    at rover waypoint
add:
    at_rock_sample waypoint
    have_rock_analysis rover waypoint
    communicated_soil_data waypoint
    channel_free lander
    at_soil_sample waypoint
delete:
(communicate_image_data lander
                                waypoint
                                rover
                                objective
                                mode
                                waypoint):

precond:
    calibrated camera rover
    communicated_rock_data waypoint
    at_rock_sample waypoint
    have_soil_analysis rover waypoint
    channel_free lander
    at rover waypoint
add:
    have_image rover objective mode
    calibrated camera rover
    communicated_image_data objective mode
delete:
    calibrated camera rover
(drop store rover):
precond:
    have_image rover objective mode
    have_soil_analysis rover waypoint
    available rover
    calibrated camera rover
    at rover waypoint
add:
    have_image rover objective mode
    calibrated camera rover
delete:

...

```

Figure 2: Sample of the output from the ARMS extraction.

### 3.1 Library Usage

Corresponding to the three main components detailed in Section 2 – trace generation, observation tokenization, and model extraction – the MACQ library offers a range of functionality for each. Usage of MACQ also follows this natural order of first generating or loading traces, then optionally applying tokenization, and then doing the model extraction.

Figure 1 shows the code required to (1) generate traces for a problem found in the online repository at [api.planning.domains](https://api.planning.domains), (2) tokenize by removing 60% of the fluents seen, and (3) apply the ARMS algorithm to extract potential actions. A portion of the output is shown in Figure 2.

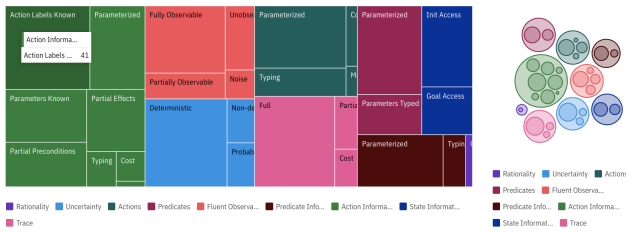


Figure 3: MACQ treemap view.

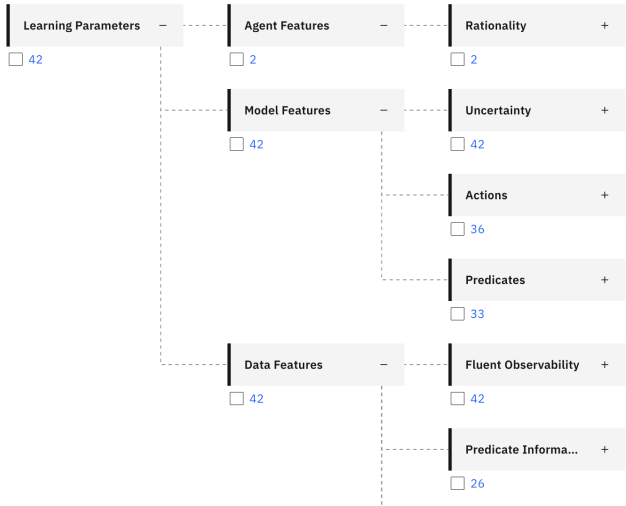


Figure 4: MACQ hierarchical view.

As long as the type of tokens in a trace allows for it<sup>2</sup>, various extraction techniques can be substituted and compared. Similarly, various data sources (from generative to pre-existing) can be used to seed the entire approach.

The MACQ library was built from the ground up to be (1) extensible and generalizable to all of the common model acquisition techniques; (2) serve as a rich resource for practitioners looking to apply model acquisition; and (3) provide a foundation for new research in the area of model acquisition.

### 3.2 Visual Interface

The MACQ library also comes with a visual interface for users to explore the available papers on the topic through various lenses. The primary view provides a taxonomic account of the various topics identified in the field and how papers are classified along those topics – for us, these topics correspond to the features discussed in Section 2.1. This is shown in Figures 3 and 4.

The next view displays the papers in MACQ’s knowledge in the latent space of features. This document embedding is computed according to the approach in (Cohan et al. 2020), inspired by a similar application in (Rush and Strobel 2020). In this view, the user can select subsets of papers in feature space, filter by features by clicking on the tags and

<sup>2</sup>See the Section on “Token Casting” for details on how traces can be transformed to different types.

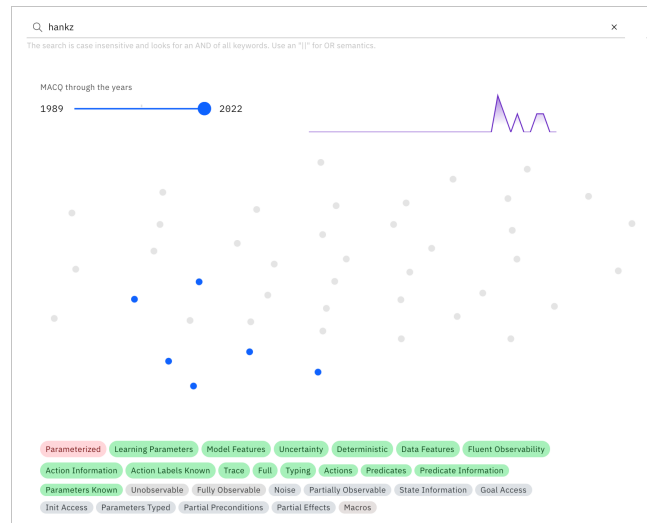


Figure 5: Exploring papers written by Hankz Hankui Zhuo, a prominent author in the field of model acquisition, in feature space. Top right inset illustrates simulation of his over time while feature tags (sized and colored by frequency) summarize salient topics in his papers.

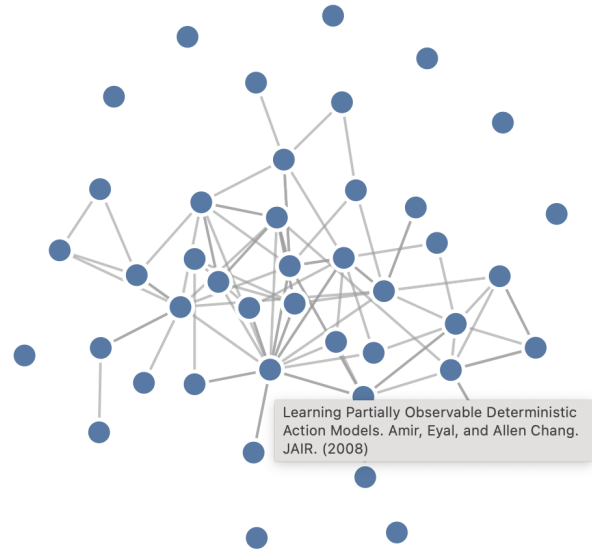


Figure 6: Exploring the influence network in MACQ– this particular cluster being hovered over belongs to (Amir and Chang 2008), one of the seminal papers in the field.

even simulate the evolution of the feature space over time. Figure 5 provides an illustration of the same.

Finally, from the PDF documents of the papers, we also automatically extract a citation network to illustrate the most influential hubs in the world of MACQ. This is shown in Figure 6. As in the case of the similarity view, here too the user can modify the network view using the feature space as well as simulate how the network evolves over time.



## 4 Research Recommendations

As a research field matures, our understanding of the gaps in our knowledge dwindles. Our efforts include not only a taxonomy of existing techniques for model acquisition (and an implementation of some of the most popular ones) but also a mechanism for exploring the research space as well. For every work documented by the MACQ project, we have a feature vector that characterizes the technique. These are detailed above in Section 2.1. Further, we have a growing set of *semantic constraints* over this taxonomy.

For example, “if a technique can operate on partially observable traces, it must be able to work on fully observable traces” and “every technique must have full, partial, or no observability”. Specifying these constraints has one immediate benefit: it allows us to systematically verify the documented features of the existing approaches. This has led to several “bug fixes” of the data collected already. However, the true power lies in the potential for viewing the research field both holistically and logically.

Alongside MACQ, and in collaboration with the visualization project used to exhibit the research area, we have developed a logical theory that corresponds to the *valid space of research* according to the features detailed in Section 2.1 and manually specified constraints over them. All of the model acquisition techniques are validated but further encoded as constraints themselves (their features being converted to a conjunction of Boolean variables or their negation).

### 4.1 Logical Encoding of Research Potential

Given the conjunction of constraints on the features, and the negation of the disjunction on the pre-existing literature (thus ruling out existing feature profiles), we have a logical theory where a satisfying assignment corresponds to a valid selection of features/assumptions about a model acquisition technique, and further is one that has not yet been explored in the known literature. Further preferences on specific features (e.g., wanting to only handle settings without parameters) can be included as unit clauses to further constrain the space of satisfying research configurations.

While this is a seemingly simple concept, there is tremendous potential in taking this viewpoint. We have implemented the above encoding and found that modern SAT solvers and knowledge compilers are readily capable of handling the theory. We further developed a novel research recommendation procedure for the space of model acquisition techniques. Making use of a knowledge compiler and repeated logical conditioning, the procedure is as follows:

1. Encode the constraints and existing techniques into a logical theory  $T$ .
2. Define a full set of soft preferences  $P$  over the features that stipulate “simple” or “nominal” settings (e.g., fully observable over partially observable).
3. Run a full knowledge compilation on  $T$ , to get a d-DNNF representing all possible solutions  $S$ .
4. Iterate over  $p \in P$ , and if  $p$  is consistent with  $S$ , enforce  $p$  by setting  $S = S \wedge p$ .

At the end of this procedure, we will have a single assignment to the full set of features that (1) adheres to a maximal number of preferences; and (2) differs from every other existing approach. Different orders for step 2 will potentially result in different final outcomes, and we will see one such example later on in this paper.

Finally, with a candidate area of unexplored research in hand, we can perform a matching algorithm to find the closest existing approaches to the one being proposed. Our system limits this to 3 and displays the core differences between the existing work and the newly proposed one. An example of this functionality is also provided below.

### 4.2 Unwritten Paper Recommendations

Figure 7 illustrates this process in action on the MACQ visual interface, introduced in Section 3.2. The visualization unfolds in three sections:

- The first part of the exposition describes the features of this newly imagined paper in terms of its features. The tag hierarchy is displayed.
- The hypothetical paper is now visualized in feature space: this view shows where it belongs when all the papers known to MACQ are projected onto a latent space only<sup>3</sup> consisting of the features from Section 2.1.
- Finally, and perhaps most interestingly, the above visual leads into neighbouring (in feature space) papers that the user can tap into as the state of the art closest to this new imagined paper. In addition to the metadata of the neighbouring papers, MACQ surfaces the features of those neighbours that need to change (either relaxed or extended) in order to make a hop from a known relevant paper to this non-existent paper.

Currently, we are working on making this interface more interactive so the user can query the system with a partial selection of papers and features of interest; and iteratively make hops to the next imagined paper. With this feature implemented, we intend to do pilot studies on how quickly we can onboard new students into a field using this exposition and exploration technique over survey data.

## 5 Open Research Questions

Given the holistic analysis of the field that MACQ offers, we have identified several promising areas for further research. Here, we highlight just a few of them.

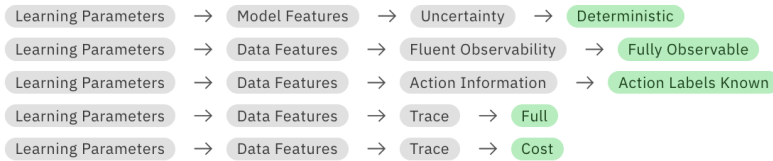
### 5.1 Operationalizing MACQ

Most of the existing approaches to model learning deal with a one-time model learning task while not taking into account the operational considerations of deploying a system with that model. In reality, models are deployed and maintained over time. As such, such models drift (Bryce, Benton, and Boldt 2016) and thus systems require a certain level of

<sup>3</sup>This view is slightly different from the similarity view described Section 3.2. There, a document includes these features but also all the rest of the paper metadata in terms of authors, title, abstract, venue, and so on.

## New Paper Features

This is a paper, described in terms of the tags in this taxonomy, that does not exist yet! It is shown below in green in "tag space".



## Neighboring Papers

To get to this new paper, our AI thinks you should be looking at the following papers known to our system as the state of the art that immediately makes the new work possible. Each paper is tagged with features that need relaxation or extension to get to the new paper.

Efficient, Safe, and Probably Approximately Complete Learning of Action Models by Stern, Roni, and Brendan Juba. IJCAI (2017) [↓](#)

Learning Parameters / Data Features / Trace / Cost / **True**

Constructing Symbolic Representations for High-Level Planning by Konidaris, George, Leslie Kaelbling, and Tomas Lozano-Perez. AAAI (2014) [↓](#)

Learning Parameters / Data Features / State Information / Init Access / **False**  
Learning Parameters / Data Features / Trace / Cost / **True**

Learning First-Order Representations for Planning from Black-Box States: New Results by Rodriguez, Ivan D., Blai Bonet, Javier Romero, and Hector Geffner. arXiv (2021) [↓](#)

Learning Parameters / Model Features / Actions / Parameterized / **False**  
Learning Parameters / Model Features / Predicates / Parameterized / **False**  
Learning Parameters / Data Features / Fluent Observability / Fully Observable / **True**  
Learning Parameters / Data Features / Fluent Observability / Unobservable / **False**  
Learning Parameters / Data Features / Fluent Observability / Noise / **False**  
Learning Parameters / Data Features / Trace / Cost / **True**

Figure 7: Imagining future KEPS papers on the MACQ visual interface.

hand-holding in terms of how to deal with such evolution of models they are deployed with. For learning systems, this is increasingly becoming a trend (Elle O'Brien 2020), with approaches that try to reconcile new models with past decisions (Bansal et al. 2019) thereby offering a certain level of consistency. We envisage similar processes to dovetail with the core MACQ functionality when systems are deployed on top of model learning tools in their portfolio.

## 5.2 XAIP Crossover

Interestingly, most of the models learned from data are underdetermined – i.e. there are many equivalent models that

can “explain” a set of observed circumstances. This also holds for iterative or “online” approaches with a domain writer in the loop. In fact, some approaches e.g. (Yang, Wu, and Jiang 2007b) specifically looked at pattern mining tools to bias the learning approaches towards more likely models. Even so, the decisions made by the model learning algorithm remain rather opaque and it may well be the case that some models rejected by it could have made more sense when presented to the domain writer. We are currently exploring the possible adoption of model-space reasoning techniques from the emerging field of explainable AI planning or XAIP (Chakraborti, Sreedharan, and Kambhampati 2020) to en-

gage in a more transparent model learning interface where the domain writer can be empowered to query and explore the trade-offs made among the equivalence class of models that satisfy a set of observed behaviours.

### 5.3 Will AI write the papers of the future?

In her presidential address (Gil 2022) at AAAI 2020, Yolanda Gil, one of the early pioneers (Carbonell and Gil 1990; Gil 1994) in the field of model learning for automated planning, asked: “Will AI write scientific papers in the future?”. The question was posed to facilitate an exploration of the influence that AI algorithms, from process management to knowledge discovery, increasingly have on our scientific endeavours. As we demonstrated in Section 3.2, the set of exploratory features made available by MACQ also belongs to this emerging theme of collaboration between AI and the scientist: not to synthesize the papers directly, but rather to provide the automated insights necessary for researchers to know where next to look. To the extent that that question applies to the KEPS community, MACQ is most certainly going to (help) write the papers of the future!

### References

- Amir, E.; and Chang, A. 2008. Learning Partially Observable Deterministic Action Models. *Journal of Artificial Intelligence Research*, 33: 349–402.
- Arora, A.; Fiorino, H.; Pellier, D.; Métivier, M.; and Pesty, S. 2018. A review of learning planning action models. *The Knowledge Engineering Review*, 33: e20.
- Bansal, G.; Nushi, B.; Kamar, E.; Weld, D.; Lasecki, W.; and Horvitz, E. 2019. A case for backward compatibility for human-ai teams. *arXiv:1906.01148*.
- Biere, A.; Fazekas, K.; Fleury, M.; and Heisinger, M. 2020. CaDiCaL, Kissat, Paracooba, Plingeling and Treengeling Entering the SAT Competition 2020. In *Proceedings of SAT Competition – Solver and Benchmark Descriptions*.
- Bryce, D.; Benton, J.; and Boldt, M. W. 2016. Maintaining evolving domain models. In *IJCAI*.
- Carbonell, J. G.; and Gil, Y. 1990. Learning by experimentation: The operator refinement method. *Machine learning*, 191–213.
- Celorrío, S. J.; de la Rosa, T.; Fernández, S.; Fernández, F.; and Borrajo, D. 2012. A review of machine learning for automated planning. *The Knowledge Engineering Review*, 27(4): 433–467.
- Chakraborti, T.; Sreedharan, S.; and Kambhampati, S. 2020. The Emerging Landscape of Explainable AI Planning and Decision Making. In *IJCAI*.
- Cohan, A.; Feldman, S.; Beltagy, I.; Downey, D.; and Weld, D. S. 2020. SPECTER: Document-level Representation Learning using Citation-informed Transformers. In *ACL*.
- Daga, K.; and Muise, C. 2021. Bauhaus: a library for building logical theories on the fly with Python. <https://github.com/qumulab/bauhaus>.
- Elle O’Brien. 2020. How machine learning ops works with GitLab and continuous machine learning. <https://about.gitlab.com/blog/2020/12/01/continuous-machine-learning-development-with-gitlab-ci>. GitLab.
- Francés, G.; Ramirez, M.; and Collaborators. 2018. Tarski: An AI Planning Modeling Framework. <https://github.com/aig-upf/tarski>.
- 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.
- Gil, Y. 1994. Learning by experimentation: Incremental refinement of incomplete planning domains. *Machine Learning*, 87–95.
- Gil, Y. 2022. Will AI write scientific papers in the future? *AI Magazine*, 42(4): 3–15.
- Helmert, M. 2006. The fast downward planning system. *Journal of Artificial Intelligence Research*, 26: 191–246.
- Ignatiev, A.; Morgado, A.; and Marques-Silva, J. 2018. PySAT: A Python Toolkit for Prototyping with SAT Oracles. In *SAT*, 428–437.
- Jilani, R.; Crampton, A.; Kitchin, D. E.; and Vallati, M. 2014. Automated Knowledge Engineering Tools in Planning: State-of-the-art and Future Challenges. In *ICAPS Workshop on Knowledge Engineering for Planning and Scheduling (KEPS)*.
- Morgado, A.; Dodaro, C.; and Marques-Silva, J. 2014. Core-Guided MaxSAT with Soft Cardinality Constraints. In O’Sullivan, B., ed., *CP*.
- Muise, C. 2016. Planning.Domains. In *ICAPS System Demonstrations Track*.
- Rodriguez, I. D.; Bonet, B.; Romero, J.; and Geffner, H. 2021. Learning First-Order Representations for Planning from Black-Box States: New Results. *arXiv:2105.10830*.
- Rush, A. M.; and Strobelt, H. 2020. MiniConf – A Virtual Conference Framework. *arXiv:2007.12238*.
- Shen, W.-M.; and Simon, H. A. 1989. Rule Creation and Rule Learning Through Environmental Exploration. In *IJCAI*.
- Verbeek, J.; de Haan, R.; and Muise, C. 2022. NNF: a Python Package for Reasoning with NNF Sentences. <https://github.com/qumulab/python-nnf>.
- Wang, X. 1994. Learning Planning Operators by Observation and Practice. In Hammond, K. J., ed., *AIPS*.
- Yang, Q.; Wu, K.; and Jiang, Y. 2005. Learning Actions Models from Plan Examples with Incomplete Knowledge. In Biundo, S.; Myers, K. L.; and Rajan, K., eds., *ICAPS*.
- Yang, Q.; Wu, K.; and Jiang, Y. 2007a. Learning action models from plan examples using weighted MAX-SAT. *Artificial Intelligence Journal*, 171(2-3): 107–143.
- Yang, Q.; Wu, K.; and Jiang, Y. 2007b. Learning action models from plan examples using weighted MAX-SAT. *Artificial Intelligence*, 171(2-3): 107–143.
- Zhuo, H. H.; Peng, J.; and Kambhampati, S. 2019. Learning Action Models from Disordered and Noisy Plan Traces. *arXiv:1908.09800*.