# Envisioning a Domain Learning Track for the IPC

#### Gregor Behnke and Pascal Bercher

University of Amsterdam and The Australian National University

04/06/2024

ICAPS 2024





## Expanding the IPC

- IPC started out quite monothematic
- There is only STRIPS optimal and satisficing planning

## Expanding the IPC

- IPC started out quite monothematic
- There is only STRIPS optimal and satisficing planning

• But it grew to probabilistic, unsolvable, parallel, HTN, learning, ...

## Expanding the IPC

- IPC started out quite monothematic
- There is only STRIPS optimal and satisficing planning

• But it grew to probabilistic, unsolvable, parallel, HTN, learning, ...

• We should continue this by starting tracks on currently important research questions

Planning<sup>-1</sup>

#### Planning

<u>Given</u> Domain  $\mathcal{D} = (P, A)$  and Problem  $\Pi = (s_0, s_g)$ <u>Find</u> Plan  $\pi = (a_1, \dots, a_n)$ 

#### Planning

<u>Given</u> Domain  $\mathcal{D} = (P, A)$  and Problem  $\Pi = (s_0, s_g)$ <u>Find</u> Plan  $\pi = (a_1, \dots, a_n)$ 

#### **Planning**<sup>-1</sup>

<u>Given</u> Plan  $\pi = (a_1, \dots, a_n)$  and Problem  $\Pi = (s_0, s_g)$ <u>Find</u> Domain  $\mathcal{D} = (P, A)$ 

#### Planning

<u>Given</u> Domain  $\mathcal{D} = (P, A)$  and Problem  $\Pi = (s_0, s_g)$ <u>Find</u> Plan  $\pi = (a_1, \dots, a_n)$ 

#### Planning<sup>-1</sup> aka Model or Domain Learning <u>Given</u> Plan $\pi = (a_1, ..., a_n)$ and Problem $\Pi = (s_0, s_g)$ <u>Find</u> Domain $\mathcal{D} = (P, A)$

#### Planning

<u>Given</u> Domain  $\mathcal{D} = (P, A)$  and Problem  $\Pi = (s_0, s_g)$ <u>Find</u> Plan  $\pi = (a_1, \dots, a_n)$ 

#### Planning<sup>-1</sup> aka Model or Domain Learning <u>Given</u> Plan $\pi = (a_1, ..., a_n)$ and Problem $\Pi = (s_0, s_g)$ <u>Find</u> Domain $\mathcal{D} = (P, A)$

It is not as simple as that:

#### Planning

<u>Given</u> Domain  $\mathcal{D} = (P, A)$  and Problem  $\Pi = (s_0, s_g)$ <u>Find</u> Plan  $\pi = (a_1, \dots, a_n)$ 

#### Planning<sup>-1</sup> aka Model or Domain Learning <u>Given</u> Plan $\pi = (a_1, ..., a_n)$ and Problem $\Pi = (s_0, s_g)$ <u>Find</u> Domain $\mathcal{D} = (P, A)$

It is not as simple as that: Details!

#### 1 Observability

• Full, partial, none

- Full, partial, none
- Noise?

#### 1 Observability

- Full, partial, none
- Noise?

#### 2 What is observed?

- Actions
- States

#### 1 Observability

- Full, partial, none
- Noise?
- 2 What is observed?
  - Actions
  - States

#### 3 Which plans?

- Goal achieving
- Optimal
- Locally Minimal
- Invalid (Black List)

- Full, partial, none
- Noise?
- 2 What is observed?
  - Actions
  - States
- 3 Which plans?
  - Goal achieving
  - Optimal
  - Locally Minimal
  - Invalid (Black List)
- 4 Structure
  - Lifted vs Grounded

- Full, partial, none
- Noise?
- 2 What is observed?
  - Actions
  - States
- 3 Which plans?
  - Goal achieving
  - Optimal
  - Locally Minimal
  - Invalid (Black List)
- 4 Structure
  - Lifted vs Grounded
- 5 Given Model

- Full, partial, none
- Noise?
- 2 What is observed?
  - Actions
  - States
- 3 Which plans?
  - Goal achieving
  - Optimal
  - Locally Minimal
  - Invalid (Black List)
- 4 Structure
  - Lifted vs Grounded
- 5 Given Model
  - Partial Preconditions and Effects

- Full, partial, none
- Noise?
- 2 What is observed?
  - Actions
  - States
- 3 Which plans?
  - Goal achieving
  - Optimal
  - Locally Minimal
  - Invalid (Black List)
- 4 Structure
  - Lifted vs Grounded
- 5 Given Model
  - Partial Preconditions and Effects
  - Can the given information be wrong?

As simple as possible!

#### As simple as possible! and Where there are learners!

1 Full observable, Fully Lifted

#### As simple as possible! and Where there are learners!

- 1 Full observable, Fully Lifted
- 2 Full observable plans, partial observable state, Fully Lifted

#### As simple as possible! and Where there are learners!

- 1 Full observable, Fully Lifted
- 2 Full observable plans, partial observable state, Fully Lifted
- 3 Partial observable plans, partial observable state, FullyLifted

#### As simple as possible! and Where there are learners!

- 1 Full observable, Fully Lifted
- 2 Full observable plans, partial observable state, Fully Lifted
- 3 Partial observable plans, partial observable state, FullyLifted

 $\times$ 

- 1 No given model
- Partially correct given model
- 3 Partially possibly incorrect given model

#### As simple as possible! and Where there are learners!

- 1 Full observable, Fully Lifted
- 2 Full observable plans, partial observable state, Fully Lifted
- 3 Partial observable plans, partial observable state, FullyLifted

 $\times$ 

- 1 No given model
- Partially correct given model
- 3 Partially possibly incorrect given model

 $\times$ 

- 1 Valid Plans
- 2 Invalid Plans

• STRIPS in PDDL

- STRIPS in PDDL
- STRIPS with negative preconditions in PDDL

- STRIPS in PDDL
- STRIPS with negative preconditions in PDDL
- ADL in PDDL

- STRIPS in PDDL
- STRIPS with negative preconditions in PDDL
- ADL in PDDL
- RDDL

- STRIPS in PDDL
- STRIPS with negative preconditions in PDDL
- ADL in PDDL
- RDDL
- HDDL

- STRIPS in PDDL
- STRIPS with negative preconditions in PDDL
- ADL in PDDL
- RDDL
- HDDL
- ...

- STRIPS in PDDL
- STRIPS with negative preconditions in PDDL
- ADL in PDDL
- RDDL
- HDDL
- ...

For now, probably the first two only

When is a learned domain correct and/or good?

When is a learned domain correct and/or good?

Classical answer:

There is a ground truth model  $\mathcal{D}^*$  that created all given plans. Then just take  $|\mathcal{D}_{\textit{learned}}\Delta\mathcal{D}^*|$ 

When is a learned domain correct and/or good?

Classical answer:

There is a ground truth model  $\mathcal{D}^*$  that created all given plans. Then just take  $|\mathcal{D}_{\textit{learned}}\Delta\mathcal{D}^*|$ 

When is a learned domain correct and/or good?

Classical answer:

There is a ground truth model  $\mathcal{D}^*$  that created all given plans. Then just take  $|\mathcal{D}_{\textit{learned}}\Delta\mathcal{D}^*|$ 

See any problem?

• Additional preconditions? (if I am at A, I cannot be at B)

When is a learned domain correct and/or good?

Classical answer:

There is a ground truth model  $\mathcal{D}^*$  that created all given plans. Then just take  $|\mathcal{D}_{\textit{learned}}\Delta\mathcal{D}^*|$ 

- Additional preconditions? (if I am at A, I cannot be at B)
- Other equivalent preconditions? (package in T, or package is not at any location)

When is a learned domain correct and/or good?

Classical answer:

There is a ground truth model  $\mathcal{D}^*$  that created all given plans. Then just take  $|\mathcal{D}_{\textit{learned}}\Delta\mathcal{D}^*|$ 

- Additional preconditions? (if I am at A, I cannot be at B)
- Other equivalent preconditions? (package in T, or package is not at any location)
- Additional deleting effects? (things I know must be false)

When is a learned domain correct and/or good?

Classical answer:

There is a ground truth model  $\mathcal{D}^*$  that created all given plans. Then just take  $|\mathcal{D}_{\textit{learned}}\Delta\mathcal{D}^*|$ 

- Additional preconditions? (if I am at A, I cannot be at B)
- Other equivalent preconditions? (package in T, or package is not at any location)
- Additional deleting effects? (things I know must be false)
- More adding effects (only pathological ones that need to be true)

When is a learned domain correct and/or good?

Classical answer:

There is a ground truth model  $\mathcal{D}^*$  that created all given plans. Then just take  $|\mathcal{D}_{\textit{learned}}\Delta\mathcal{D}^*|$ 

- Additional preconditions? (if I am at A, I cannot be at B)
- Other equivalent preconditions? (package in T, or package is not at any location)
- Additional deleting effects? (things I know must be false)
- More adding effects (only pathological ones that need to be true)
- Can the learner even identify  $\mathcal{D}^*$ ?

When is a learned domain correct and/or good?

Classical answer:

There is a ground truth model  $\mathcal{D}^*$  that created all given plans. Then just take  $|\mathcal{D}_{\textit{learned}}\Delta\mathcal{D}^*|$ 

See any problem?

- Additional preconditions? (if I am at A, I cannot be at B)
- Other equivalent preconditions? (package in T, or package is not at any location)
- Additional deleting effects? (things I know must be false)
- More adding effects (only pathological ones that need to be true)
- Can the learner even identify  $\mathcal{D}^*$ ?

If the domain is partially observable and we don't see one predicate at all, there might be multiple isomorphic ways to encode it.

#### Syntax vs Semantic

• Correct and planned domain should have the same plans!

- Correct and planned domain should have the same plans!
- Compute k shortest plans for both domains and compare (Top-k planning)

- Correct and planned domain should have the same plans!
- Compute k shortest plans for both domains and compare (Top-k planning)
- Two types of errors

- Correct and planned domain should have the same plans!
- Compute k shortest plans for both domains and compare (Top-k planning)
- Two types of errors
  - Precision: Learned model admits incorrect plans

- Correct and planned domain should have the same plans!
- Compute k shortest plans for both domains and compare (Top-k planning)
- Two types of errors
  - Precision: Learned model admits incorrect plans
  - Recall: Learned model does not admit correct plans.

- Correct and planned domain should have the same plans!
- Compute k shortest plans for both domains and compare (Top-k planning)
- Two types of errors
  - Precision: Learned model admits incorrect plans
  - Recall: Learned model does not admit correct plans.
- Functional Equivalence (Shrinah, Long, Eder'23)

- Correct and planned domain should have the same plans!
- Compute k shortest plans for both domains and compare (Top-k planning)
- Two types of errors
  - Precision: Learned model admits incorrect plans
  - Recall: Learned model does not admit correct plans.
- Functional Equivalence (Shrinah, Long, Eder'23)
- Comparing Domain Models (Chrpa, Dodaro, Maratea, Mochi, Vallati'23)

Can the learners even be expected to find  $\mathcal{D}^*$ ?

• Given plans might not be enough to objectively identify the ground truth domain.

- Given plans might not be enough to objectively identify the ground truth domain.
- What are we evaluating then?

- Given plans might not be enough to objectively identify the ground truth domain.
- What are we evaluating then?
  Bias towards models written by the modeller ...

Can the learners even be expected to find  $\mathcal{D}^*$ ?

- Given plans might not be enough to objectively identify the ground truth domain.
- What are we evaluating then?
  Bias towards models written by the modeller ...

• Admit every model that is a provably correct inference from the given plans.

- Given plans might not be enough to objectively identify the ground truth domain.
- What are we evaluating then?
  Bias towards models written by the modeller ...

- Admit every model that is a provably correct inference from the given plans.
- Check the learned domain only w.r.t. the given plans!

- Given plans might not be enough to objectively identify the ground truth domain.
- What are we evaluating then?
  Bias towards models written by the modeller ...

- Admit every model that is a provably correct inference from the given plans.
- Check the learned domain only w.r.t. the given plans! **Test:** are all given plans valid for the learned model?

- Given plans might not be enough to objectively identify the ground truth domain.
- What are we evaluating then?
  Bias towards models written by the modeller ...

- Admit every model that is a provably correct inference from the given plans.
- Check the learned domain only w.r.t. the given plans! **Test:** are all given plans valid for the learned model?
- Doable for

- Given plans might not be enough to objectively identify the ground truth domain.
- What are we evaluating then?
  Bias towards models written by the modeller ...

- Admit every model that is a provably correct inference from the given plans.
- Check the learned domain only w.r.t. the given plans! **Test:** are all given plans valid for the learned model?
- Doable for
  - learning from a given incorrect model and minimize changes

- Given plans might not be enough to objectively identify the ground truth domain.
- What are we evaluating then?
  Bias towards models written by the modeller ...

- Admit every model that is a provably correct inference from the given plans.
- Check the learned domain only w.r.t. the given plans! **Test:** are all given plans valid for the learned model?
- Doable for
  - learning from a given incorrect model and minimize changes
  - learning from a plans with guarantees (e.g. optimal, non-redundant)

### What Happens Next?

- We will keep discussing with the community.
- So: contact us!

### What Happens Next?

- We will keep discussing with the community.
- So: contact us!

• We target ICAPS'25 if we can get a consensus on the setting until end of the year.