The background features a complex, light gray graphic of interlocking gears and circular patterns. Some gears have numerical markings around their perimeters, such as 40, 150, 160, 170, 180, 190, 200, 210, 220, 230, 240, 250, and 260. The overall aesthetic is technical and mechanical, suggesting a focus on engineering or computer science.

Bridging Logical Reasoning and Machine Learning in Program Synthesis

Yanju Chen

Computer Science Department

University of California, Santa Barbara

12/03/2020

Overview

- Program Synthesis in a Nutshell
- MARS: Encoding Multi-Layer Specifications
- CONCORD: Deduction-Guided Reinforcement Learning
- Related Works & Conclusions

Program Synthesis in a Nutshell

- Problem Formalization
 - Related Works
 - Program Synthesis with Machine Learning (I)
 - A Data Wrangling Example & DSL
 - NEO: A Brief Overview
- Observations & Motivations
 - Q1: Why logical reasoning?
 - Q2: Why machine learning?
 - Q3: Why bridging?

Problem Formalization


multi-modal

sample_ID	site	coll_date	species	TOT	inf_status
382870	site1	27/10/2007	SpeciesB	1	positive
382872	site2	27/10/2007	SpeciesB	1	negative
487405	site3	28/10/2007	SpeciesA	1	positive
487405	site3	28/10/2007	SpeciesA	1	positive

site	coll_date	SP_A_pos	SP_A_neg	SP_B_pos	SP_B_neg
site1	27/10/2007	0	0	1	0
site2	27/10/2007	0	0	0	1
site3	28/10/2007	2	0	0	0



examples



occurs(unite) \wedge
 occurs(group_by) \wedge
 hasChild(group_by, unite) \wedge
 ... logical constraints



I need to reformat the data so that there is just one row per site visit (i.e. in a given site name and date combo) with columns for total found by species and the fish status (i.e. speciesA_pos, SpeciesA_neg, Sp_B_pos.. etc).

natural languages



...

multi-paradigm



neural




deductive



...



synthesizer



specifications ϕ



program P

Find a program P that satisfies all the specifications ϕ .

Program Synthesis in a Nutshell

Related Works

*The table only lists some of the recent related works.

IO: Input-Output Example | NL: Natural Language | 😊: Yes | 😞: Not explicitly claimed

synthesizer	domain evaluated	specification	logical reasoning	machine learning	bridging level	multi-modal
DEEPCODER (Balog et al. 2017)	list	IO	😞	😊	NA	😞
SEQ2SQL (Zhong et al. 2017)	SQL	I + NL	😞	😊	NA	😊
DIALSQL (Gur et al. 2018)	SQL	NL	😞	😊	NA	😊
EXEC (Chen et al. 2018)	Karel	IO	😞	😊	NA	😞
NEO (Feng et al. 2018)	table + list	IO	😊	😊	★	😞
SKETCHADAPT (Nye et al. 2018)	list + string + Algolisp	IO / IO + NL	😊	😊	★★	😊
SQLIZER (Yaghmazadeh et al. 2018)	SQL	NL	😊	😊	★	😞
AutoPandas (Bavishi et al. 2019)	table	IO	😞	😊	NA	😞
MARS (Chen et al. 2019)	table	IO + NL	😊	😊	★★★	😊
METAL (Si et al. 2019)	circuit	logical formula	😊	😊	★★★	😞
CONCORD (Chen et al. 2020)	list	IO	😊	😊	★★★	😞
PROBE (Barke et al. 2020)	string + circuit + bitvector	IO	😊	😊	★★★	😞
REGEL (Chen et al. 2020)	regex	IO + NL	😊	😊	★★	😊
VISER (Wang et al. 2020)	visualization	IO + visual sketch	😊	😊	★	😊

Program Synthesis in a Nutshell

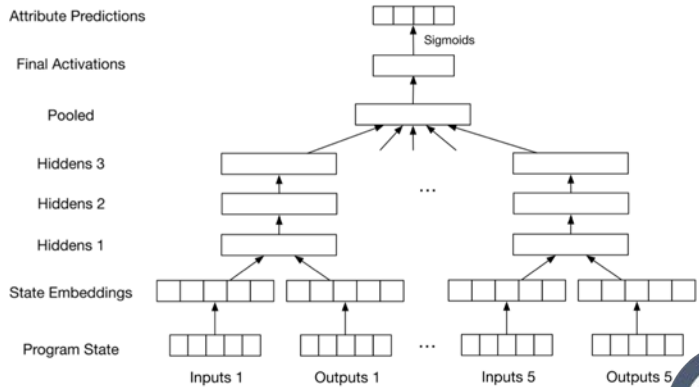
Program Synthesis with Machine Learning (I)



Predictions of Language Constructs / Partial Programs

```
a ← [int]
b ← FILTER (<0) a
c ← MAP (*4) b
d ← SORT c
e ← REVERSE d
```

Proposed Program



Neural Encoder

Input:
[-17, -3, 4, 11, 0, -5, -9, 13, 6, 6, -8, 11]
Output:
[-12, -20, -32, -36, -68]

Examples

representation learning^[1]

multi-modal encoding^[2]

Natural Language



User

- [1] DEEPCODER (Balog et al. 2017); EXEC (Chen et al. 2018); AutoPandas (Bavishi et al. 2019); METAL (Si et al. 2019); CONCORD (Chen et al. 2020);
- [2] SEQ2SQL (Zhong et al. 2017); DIALSQL (Gur et al. 2018); SQLIZER (Yaghmazadeh et al. 2018); MARS (Chen et al. 2019); REGEL (Chen et al. 2020); VISER (Wang et al. 2020);

A Running Example from StackOverflow^[1]

[Title] r script to count columns within dataset

[Example]

sample_ID	site	coll_date	species	TOT	inf_status
382870	site1	27/10/2007	SpeciesB	1	positive
382872	site2	27/10/2007	SpeciesB	1	negative
487405	site3	28/10/2007	SpeciesA	1	positive
487405	site3	28/10/2007	SpeciesA	1	positive



site	cat	sts
site1	SpeciesB_positive	1
site2	SpeciesB_negative	1
site3	SpeciesA_positive	2

[Description]

I need to reformat the data so that there is just one row per site visit (i.e. in a given site name and date combo) with columns for total found by species and the fish status (i.e. speciesA_pos, SpeciesA_neg, Sp_B_pos.. etc).

figured I need to sum within site. My thoughts were to use split/apply/aggregate/for loops etc but tried various combinations and not getting anywhere. apologies I'm not familiar with R. any comments appreciated!

[1] Example adapted from <https://stackoverflow.com/questions/39369502/r-script-to-reshape-and-count-columns-within-dataset>
Bridging Logical Reasoning and Machine Learning in Program Synthesis

A Running DSL for Data Wrangling^[1]

$t \rightarrow x_i$ (input table)
 | `select(t, \vec{c}_{arg})` (column projection)
 | `unite($t, c_{tgt}, \vec{c}_{arg}$)` (column merging)
 | `separate($t, \vec{c}_{tgt}, c_{arg}$)` (column splitting)
 | `mutate($t, c_{tgt}, op, \vec{c}_{arg}$)` (column arithmetic)
 | `group_by(t, \vec{c}_{arg})` (row grouping)
 | `summarise($t, c_{tgt}, a, \vec{c}_{arg}$)` (row aggregation)
 | `filter(t, f, \vec{c}_{arg})` (row filtering)
 $op \rightarrow + \mid - \mid \times \mid \div$
 $a \rightarrow \min \mid \max \mid \text{sum} \mid \text{count} \mid \text{avg}$

x_i : the i -th input table

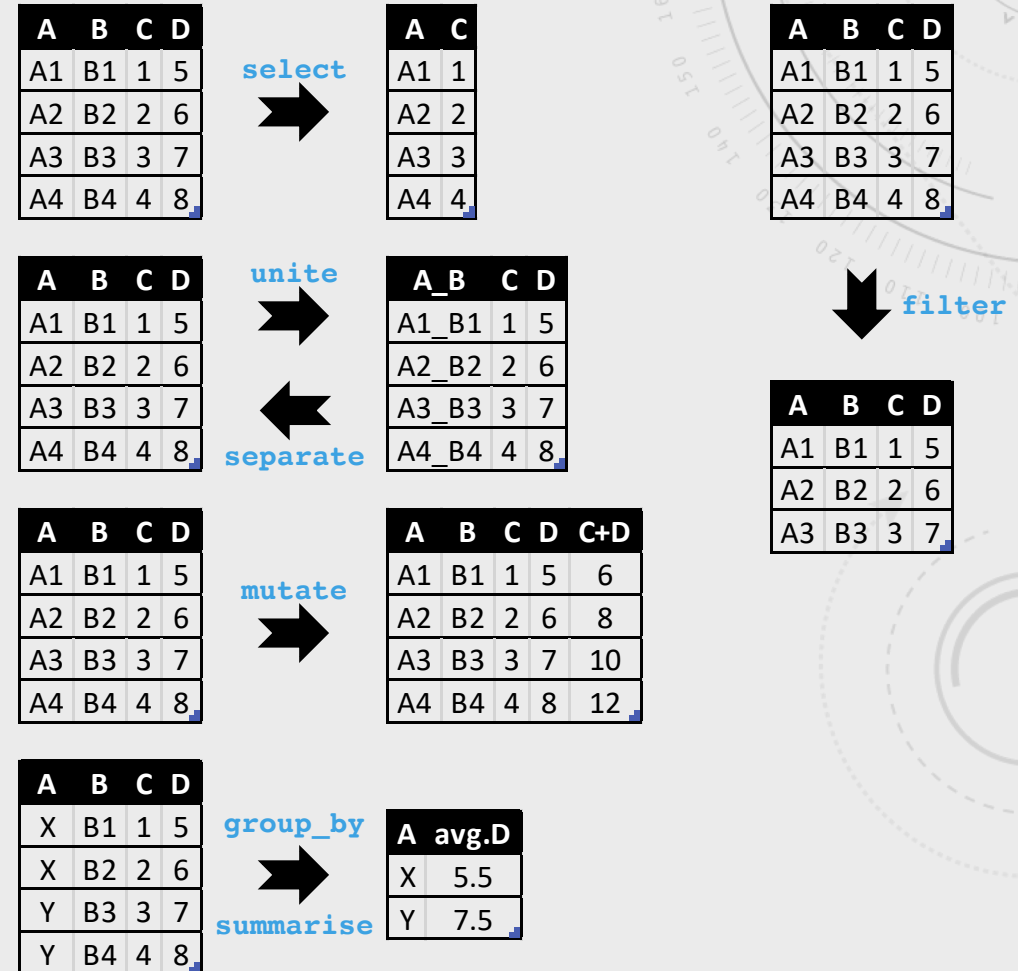
t : table

c, \vec{c} : column(s) of table

op : arithmetic operation

a : aggregation function

f : higher-order boolean function



[1] DSL adapted from Wang, C. et al. Visualization by Example. POPL'20
 Bridging Logical Reasoning and Machine Learning in Program Synthesis

A Running Example from StackOverflow

[Example]

sample_ID	site	coll_date	species	TOT	inf_status
382870	site1	27/10/2007	SpeciesB	1	positive
382872	site2	27/10/2007	SpeciesB	1	negative
487405	site3	28/10/2007	SpeciesA	1	positive
487405	site3	28/10/2007	SpeciesA	1	positive

unite



sample_ID	site	coll_date	cat	TOT
382870	site1	27/10/2007	SpeciesB_positive	1
382872	site2	27/10/2007	SpeciesB_negative	1
487405	site3	28/10/2007	SpeciesA_positive	1
487405	site3	28/10/2007	SpeciesA_positive	1

group_by

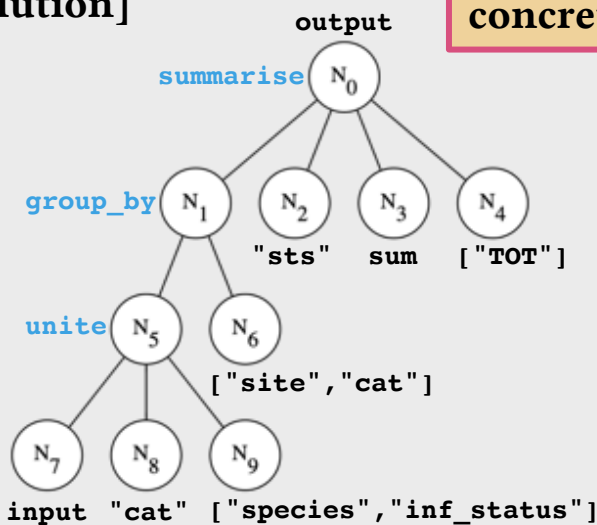


summarise

site	cat	sts
site1	SpeciesB_positive	1
site2	SpeciesB_negative	1
site3	SpeciesA_positive	2

[Solution]


concrete program



```
T0 = unite( input, "cat", ["species", "inf_status"] )
T1 = group_by( T0, ["site", "cat"] )
output = summarise( T1, "sts", sum, ["TOT"] )
```

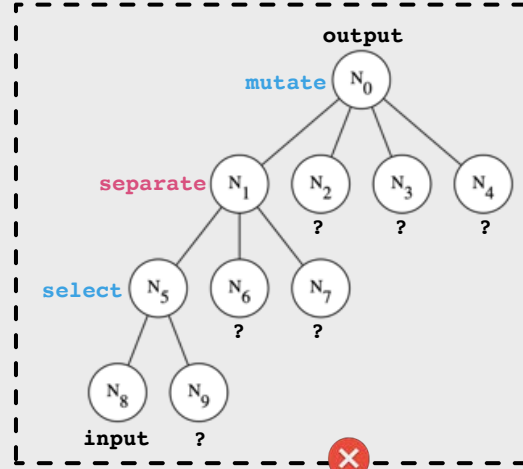
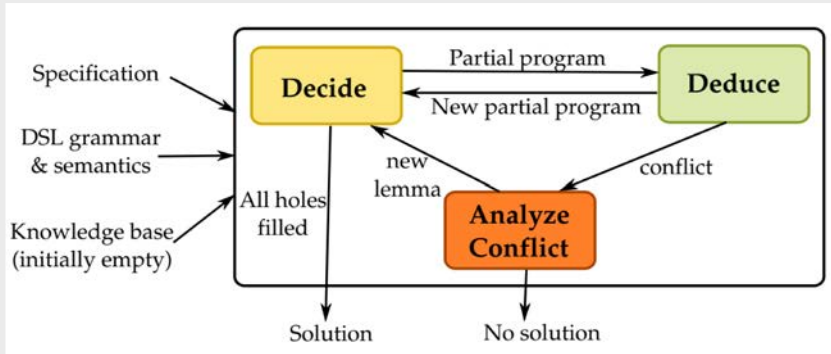
partial program / sketch

Labels of some AST nodes are yet to be determined.



```
T0 = unite( input, ?, ? )
T1 = group_by( T0, ? )
output = summarise( T1, ?, ?, ? )
```

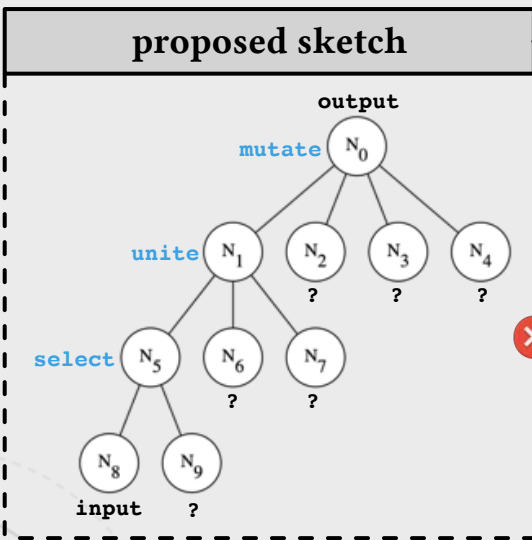
NEO^[1]: A Brief Overview



Equivalent Modulo Conflict (EMC)^[1]

select	$out.row == in.row \wedge out.col \leq in.col - 1$
unite	$out.row == in.row \wedge out.col == in.col - 1$
separate	$out.row == in.row \wedge out.col == in.col + 1$
mutate	$out.row == in.row \wedge out.col == in.col + 1$
group_by	$out.row == in.row \wedge out.col == in.col$
summarise	$out.row \leq in.row \wedge out.col \leq in.col + 1$
filter	$out.row \leq in.row - 1 \wedge out.col == in.col$


Component-Based Specifications^[2] for Data Wrangling DSL



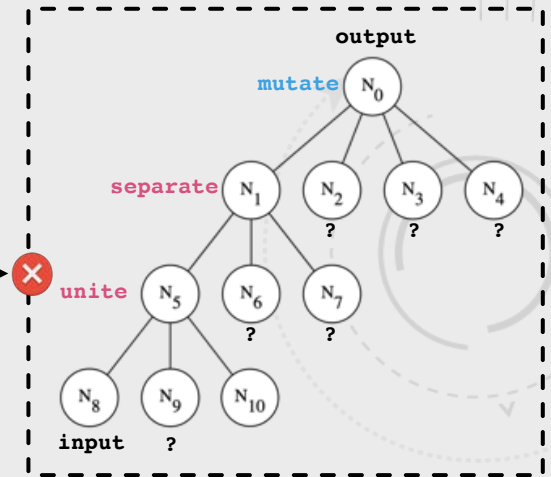
generated constraints

$input.row == 4 \wedge input.col == 6 \wedge$
 $N8.row == input.row \wedge N8.col == input.col \wedge$
 $N5.row == N8.row \wedge N5.col \leq N8.col - 1 \wedge$
 $N1.row == N5.row \wedge N1.col == N5.col - 1 \wedge$
 $N0.row == N1.row \wedge N0.col == N1.col + 1 \wedge$
 $output.row == N0.row \wedge output.col == N0.col \wedge$
 $output.row == 3 \wedge output.col == 3$

(input example)
 (input alignment)
 (select semantics)
 (unite semantics)
 (mutate semantics)
 (output alignment)
 (output example)



SMT-Based Deduction^[1] & Analyze Conflicts^[2]



[1] Feng, Y. et al. Program Synthesis using Conflict-Driven Learning. PLDI'18

[2] Feng, Y. et al. Component-based Synthesis of Table Consolidation and Transformation Tasks from Examples. PLDI'17

Observations & Motivations

- Q1: Why logical reasoning?
 - Example: EXEC^[1]
 - Concrete interpretation is less efficient, especially for complex problems
 - Logical reasoning results generalize better in pruning search space
- Q2: Why machine learning?
 - Example: AutoPandas^[2]
 - Machine learning backend provides better estimations prioritizing search order
- Q3: Why bridging?
 - Example: NEO^[3]
 - Programs are precise, but specifications can be vague
 - Statistical components can't reflect deduction feedbacks on the fly

We need both, and better!

[1] Chen, X. et al. Execution-Guided Neural Program Synthesis. ICLR'18

[2] Bavishi, R. et al. AutoPandas: Neural-backed Generators for Program Synthesis. OOPSLA'19

[3] Feng, Y. et al. Program Synthesis using Conflict-Driven Learning. PLDI'18

Intertitles

Bridging the Logical and Statistical Lands

- Observations & Motivations
 - Existing tools do have logical and statistical components combined
 - Example: NEO^[1] / TRINITY^[2]
 - But they are no more than "wired" together: still talk in different languages, act independently
- Two Bridging Directions
 - MARS^[3]: Encode multi-layer specifications (via machine learning) into logical components
 - Talk in logical language!
 - Encode specifications as soft/hard constraints in maximum satisfiability modulo theory (Max-SMT)
 - CONCORD^[4]: Guide the statistical components using deductions
 - Talk in statistical language!
 - Generate training samples for machine learning models by explaining deduction results

[1] Feng, Y. et al. Program Synthesis using Conflict-Driven Learning. PLDI'18

[2] Martins, R. et al. Trinity: An Extensible Synthesis Framework for Data Science. VLDB'19

[3] Chen, Y. et al. Maximal Multi-layer Specification Synthesis. FSE'19

[4] Chen, Y. et al. Program Synthesis Using Deduction-Guided Reinforcement Learning. CAV'20

MARS^[1]: Encoding Multi-Layer Specifications

- Motivations
- Formalization
- Framework Overview
- Multi-Layer Specification Encoding
 - Encoding Examples as Hard Constraints
 - Encoding Natural Language Specifications
- Evaluations
 - Evaluation Setup
 - Evaluation Results & Analysis
- Discussions

Maximal Multi-Layer Specification Synthesis

- Motivations
 - Examples can be imprecise
 - Multi-modal specifications contain more useful information

[Title] r script to count columns within dataset

[Example]

sample_ID	site	coll_date	species	TOT	inf_status
382870	site1	27/10/2007	SpeciesB	1	positive
382872	site2	27/10/2007	SpeciesB	1	negative
487405	site3	28/10/2007	SpeciesA	1	positive
487405	site3	28/10/2007	SpeciesA	1	positive



site	cat	sts
site1	SpeciesB_positive	1
site2	SpeciesB_negative	1
site3	SpeciesA_positive	2

[Description]

I need to reformat the data so that there is just one row per site visit (i.e. in a given site name and date combo) with columns for total found by species and the fish status (i.e. speciesA_pos, SpeciesA_neg, Sp_B_pos.. etc).

figured I need to sum within site. My thoughts were to use split/apply/aggregate/for loops etc but tried various combinations and not getting anywhere. apologies I'm not familiar with R. any comments appreciated!

Maximal Multi-Layer Specification Synthesis

- Motivations
 - Examples can be imprecise
 - Multi-modal specifications contain more useful information

[Title] r script to count columns within dataset

[Example]

sample_ID	site	coll_date	species	TOT	inf_status
382870	site1	27/10/2007	SpeciesB	1	positive
382872	site2	27/10/2007	SpeciesB	1	negative
487405	site3	28/10/2007	SpeciesA	1	positive
487405	site3	28/10/2007	SpeciesA	1	positive



site	cat	sts
site1	SpeciesB_positive	1
site2	SpeciesB_negative	1
site3	SpeciesA_positive	2

[Description]

I need to reformat the data so that there is just one row per site visit (i.e. in a given site name and date combo) with columns for total found by species and the fish status (i.e. speciesA_pos, SpeciesA_neg, Sp_B_pos.. etc).

figured I need to sum within site. My thoughts were to use split/apply/aggregate/for loops etc but tried various combinations and not getting anywhere. apologies I'm not familiar with R. any comments appreciated!

Maximal Multi-Layer Specification Synthesis

- Motivations
 - Examples can be imprecise
 - Multi-modal specifications contain more useful information

[Title] r script to count columns within dataset

[Example]

sample_ID	site	coll_date	species	TOT	inf_status
382870	site1	27/10/2007	SpeciesB	1	positive
382872	site2	27/10/2007	SpeciesB	1	negative
487405	site3	28/10/2007	SpeciesA	1	positive
487405	site3	28/10/2007	SpeciesA	1	positive



site	cat	sts
site1	SpeciesB_positive	1
site2	SpeciesB_negative	1
site3	SpeciesA_positive	2

[Description]

I need to reformat the data so that there is just one row per site visit (i.e. in a given site name and date combo) with columns for total found by species and the fish status (i.e. speciesA_pos, SpeciesA_neg, Sp_B_pos.. etc).

figured I need to sum within site. My thoughts were to use split/apply/aggregate/for loops etc but tried various combinations and not getting anywhere. apologies I'm not familiar with R. any comments appreciated!

Maximal Multi-Layer Specification Synthesis

- Motivations

- Example: `group_by` imprecise
- Multi-modal specifications contain more useful information

[Title] r script to `count` columns within dataset

[Example]

sample_ID	site	coll_date	species	TOT	inf_status
382870	site1	27/10/2007	SpeciesB	1	positive
382872	site2	27/10/2007	SpeciesB	1	negative
487405	site3	28/10/2007	SpeciesA	1	positive
487405	site3	28/10/2007	SpeciesA	1	positive

`unite`

site	cat	sts
site1	SpeciesB_positive	1
site2	SpeciesB_negative	1
site3	SpeciesA_positive	2

[Description]

I need to reformat the data so that there is just one row per site visit (i.e. in a given site name and date combo) with columns for `total found` by species and the fish status (i.e. `speciesA_pos`, `SpeciesA_neg`, `Sp_B_pos..` etc).

figured I need to `sum` within site. My thoughts were to use `split/apply/aggregate` for loops etc but tried various combinations and not getting anywhere. apologies I'm not familiar with R. any comments appreciated!

`summarise`

`sum`

Formalization

- Maximal Multi-Layer Specification Synthesis

hard constraints/specifications: examples

DSL construct

Given specification $(\mathcal{E}, \Psi, \Sigma)$ where $\mathcal{E} = (T_{in}, T_{out})$, $\Psi = \cup(\chi_i, \omega_i)$, and Σ represents all symbols in the DSL, the *Maximal Multi-Layer Specification Synthesis* problem is to infer a program \mathcal{P} such that:

- \mathcal{P} is a well-typed expression over symbols in Σ ,
- $\mathcal{P}(T_{in}) = T_{out}$, and
- $\sum \omega_i$ is maximized.

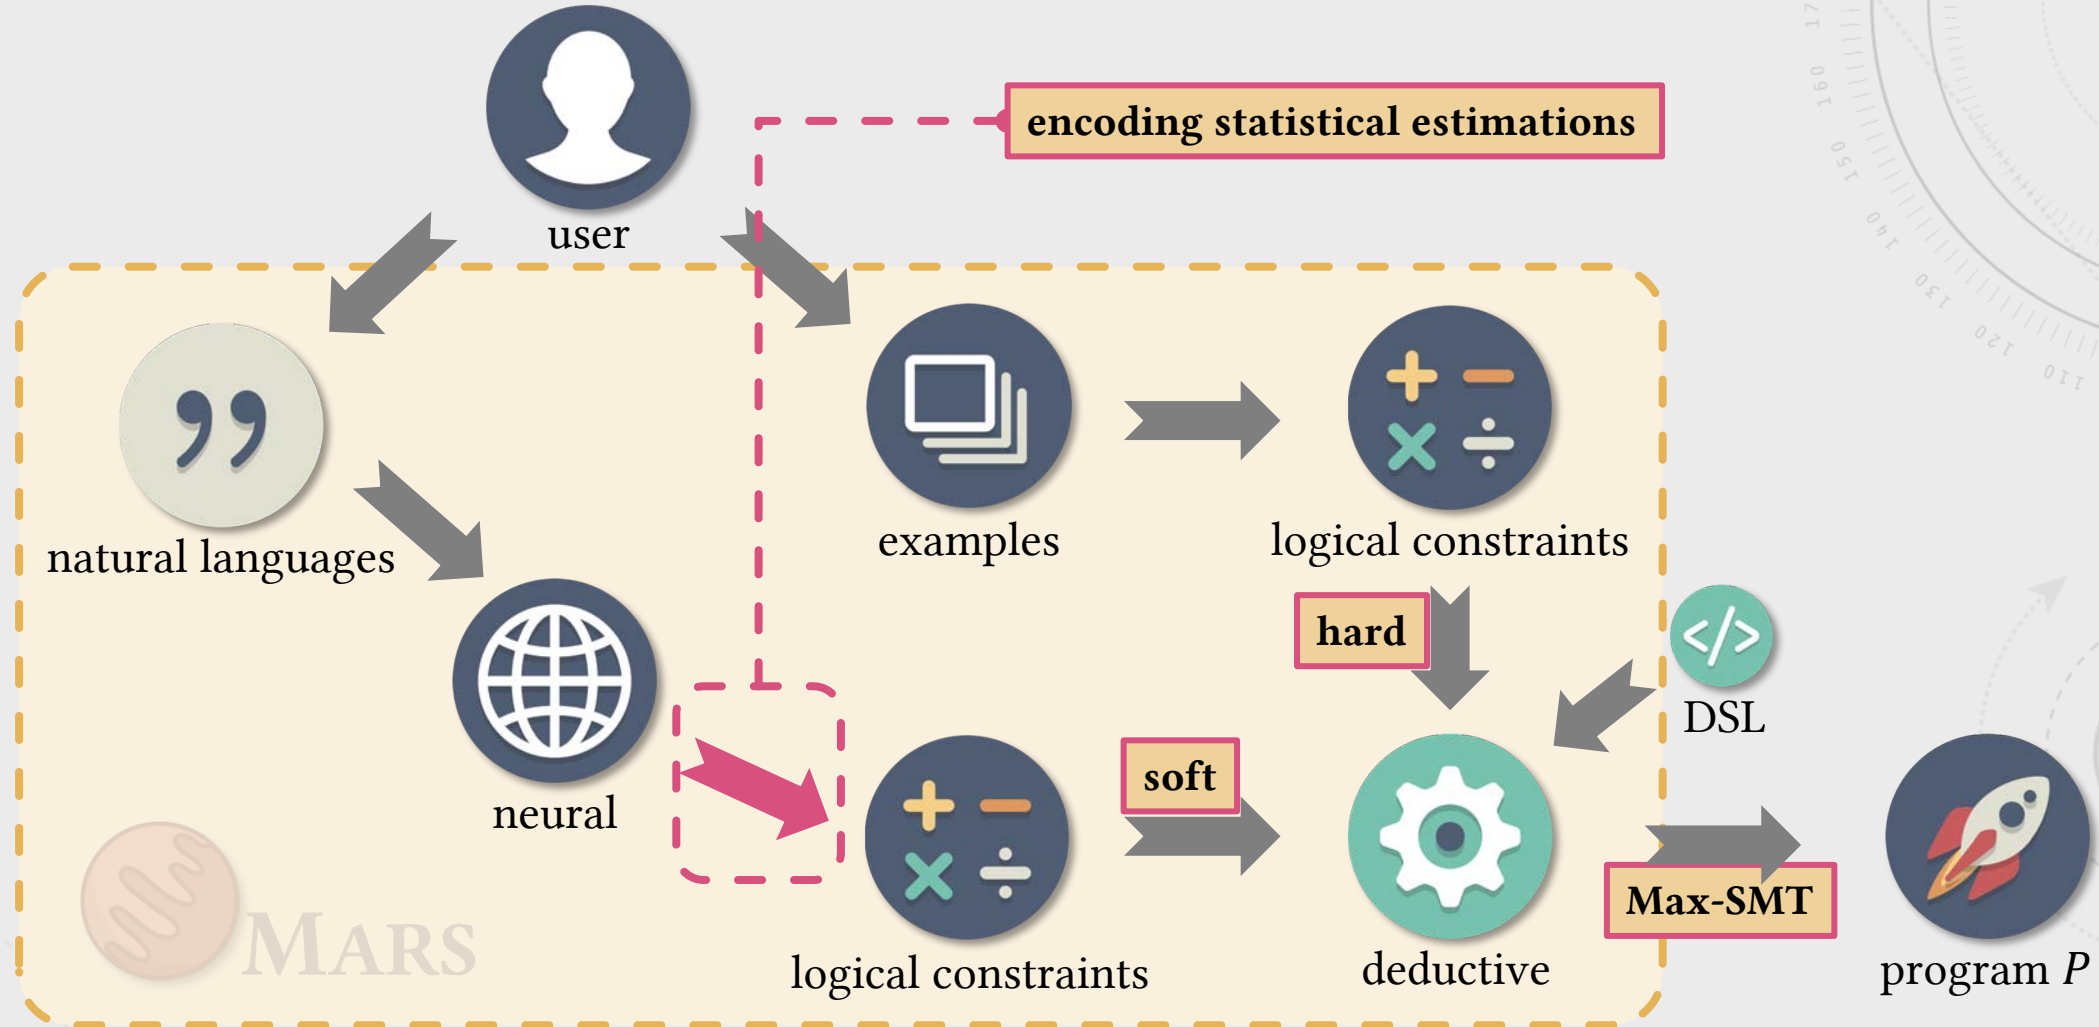
soft constraints/specifications: natural languages

preference/confidence

- We model the problem using *maximum satisfiability modulo theory (Max-SMT)* and solve it with an off-the-shelf SMT solver.

**Hard constraints should be satisfied;
Soft constraints should be maximized.**

MARS: Encoding Multi-Layer Specifications Framework Overview



Encoding Examples as Hard Constraints

symbolic program

site	cat	sts
site1	SpeciesB_positive	1
site2	SpeciesB_negative	1
site3	SpeciesA_positive	2

output.row == 3
output.col == 3

output.row == N0.row
output.col == N0.col

satisfiable?

YES

summarise

group_by

unite

input

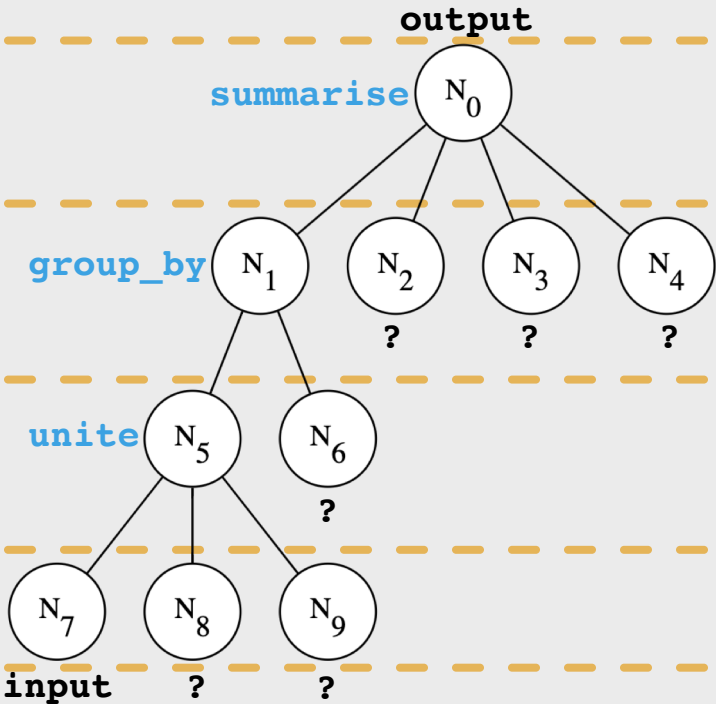
N0.row <= 4
N0.col <= 6

N1.row == 4
N1.col == 5

N5.row == 4
N5.col == 5

N7.row == 4
N7.col == 6

input.row == 4
input.col == 6



```
N5 = unite( ?, ?, ? )
N1 = group_by( N5, ? )
N0 = summarise( N1, ?, ?, ? )
```

sample_ID	site	coll_date	species	TOT	inf_status
382870	site1	27/10/2007	SpeciesB	1	positive
382872	site2	27/10/2007	SpeciesB	1	negative
487405	site3	28/10/2007	SpeciesA	1	positive
487405	site3	28/10/2007	SpeciesA	1	positive

Encoding Examples as Hard Constraints

symbolic program

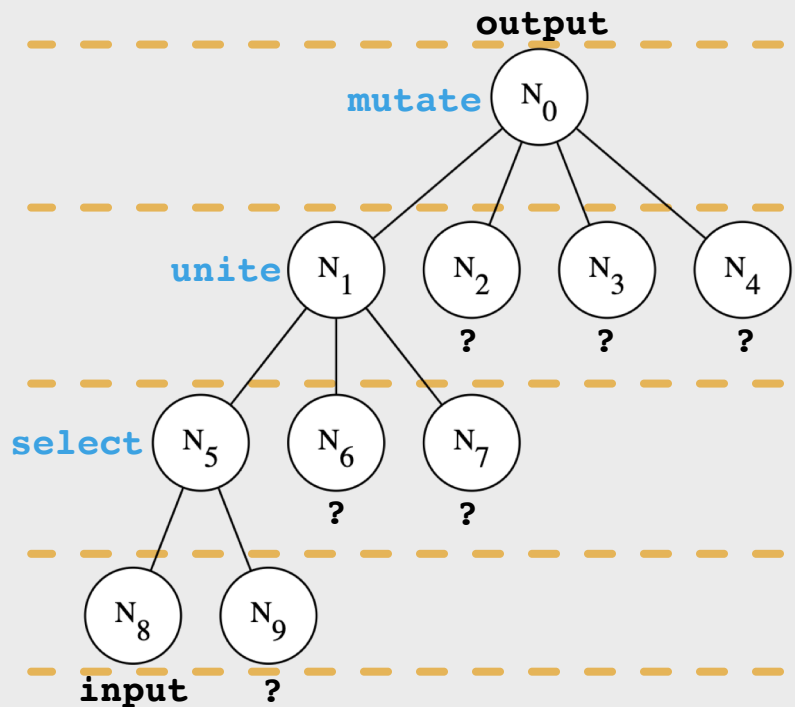
site	cat	sts
site1	SpeciesB_positive	1
site2	SpeciesB_negative	1
site3	SpeciesA_positive	2

output.row == 3
output.col == 3

output.row == N0.row
output.col == N0.col

satisfiable?

NO



mutate

unite

select

input

N0.row == 4
N0.col <= 6

N1.row == 4
N1.col <= 4

N5.row == 4
N5.col <= 5

N8.row == 4
N8.col == 6

sample_ID	site	coll_date	species	TOT	inf_status
382870	site1	27/10/2007	SpeciesB	1	positive
382872	site2	27/10/2007	SpeciesB	1	negative
487405	site3	28/10/2007	SpeciesA	1	positive
487405	site3	28/10/2007	SpeciesA	1	positive

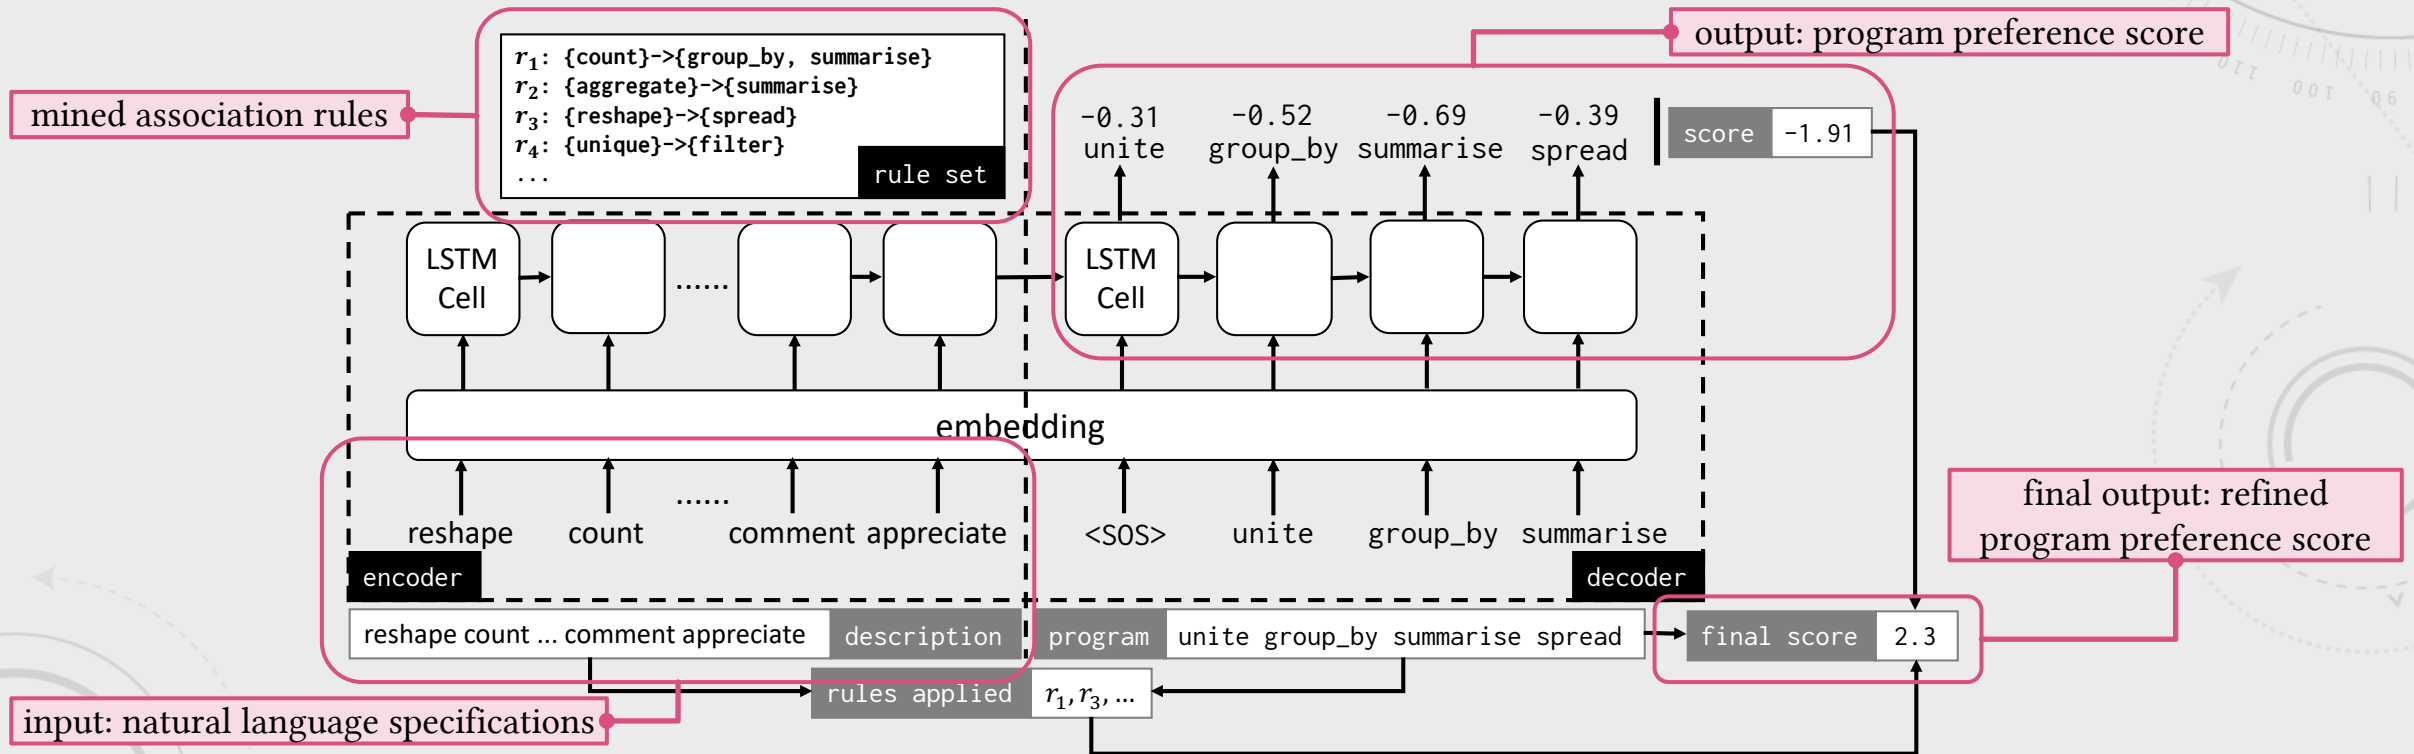
input.row == 4
input.col == 6

```
N5 = select( ?, ? )
N1 = unite( N5, ?, ? )
N0 = mutate( N1, ?, ?, ? )
```

Encoding Natural Language Specifications

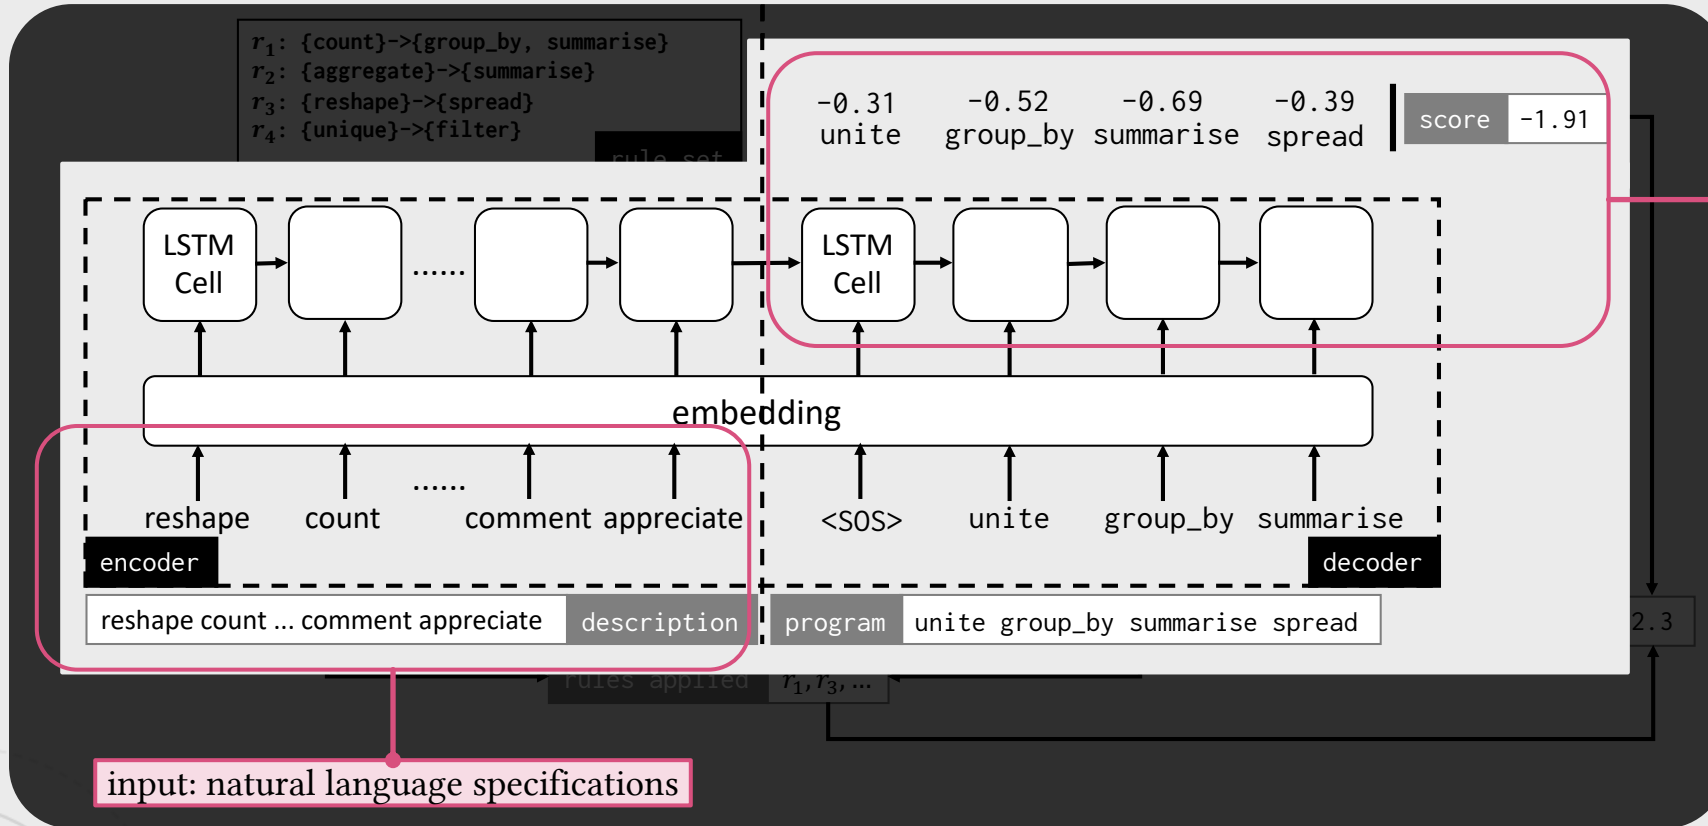
- The Hybrid Neural Architecture

- seq2seq model (supervised): capture common natural language semantics
- association rule module (unsupervised): capture frequent patterns and refine the preference



Encoding Natural Language Specifications

- The seq2seq model



output: program preference score

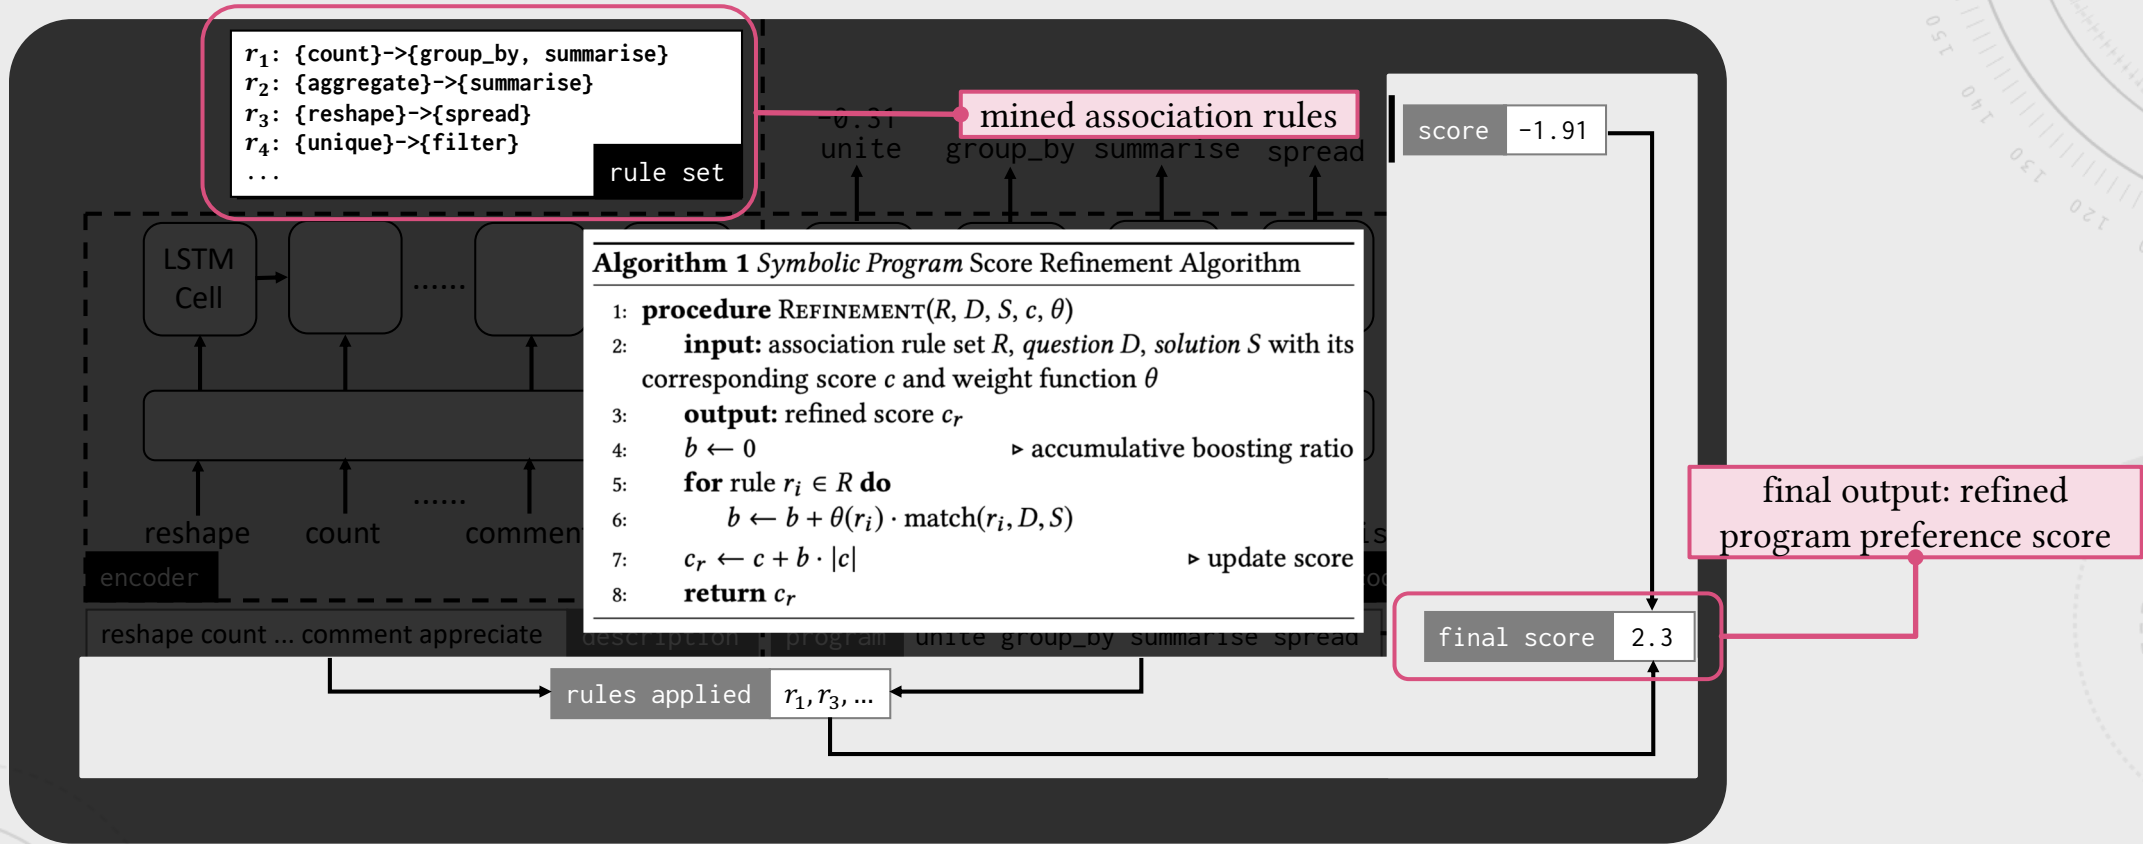
$$\ell(P|D) = \sum_i \log \pi(P_i|D)$$

* π is the seq2seq model

seq2seq model captures global user intents.

Encoding Natural Language Specifications


- The association rule module



Association rules captures local user intents.

Encoding Natural Language Specifications

- Encoding refined preference scores



2.3 unite group_by summarise

occurs predicates
 $\text{occurs}(\text{summarise}, 2.3) \wedge$
 $\text{occurs}(\text{group_by}, 2.3) \wedge$
 $\text{occurs}(\text{unite}, 2.3)$
weight=2.3

hasChild predicates
 $\text{hasChild}(\text{summarise}, \text{group_by}, 2.3) \wedge$
 $\text{hasChild}(\text{group_by}, \text{unite}, 2.3)$
weight=2.3

$(N_0 == \text{idx}(\text{summarise}) \vee N_1 == \text{idx}(\text{summarise}) \vee N_5 == \text{idx}(\text{summarise})) \wedge$
 $(N_0 == \text{idx}(\text{group_by}) \vee N_1 == \text{idx}(\text{group_by}) \vee N_5 == \text{idx}(\text{group_by})) \wedge$
 $(N_0 == \text{idx}(\text{unite}) \vee N_1 == \text{idx}(\text{unite}) \vee N_5 == \text{idx}(\text{unite}))$

$(N_0 == \text{idx}(\text{summarise}) \Rightarrow N_1 == \text{idx}(\text{group_by})) \wedge$
 $(N_0 == \text{idx}(\text{summarise}) \Rightarrow N_5 == \text{idx}(\text{unite})) \wedge$
 $(N_1 == \text{idx}(\text{group_by}) \Rightarrow N_5 == \text{idx}(\text{unite}))$

encoding **occurs**(p_i, ω_i)
 $\bigwedge_{p_i \in \Lambda} \bigvee_{N_i \in \mathcal{N}} N_i == \text{idx}(p_i)$

encoding **hasChild**(p_i, p_j, ω_i)
 $p_i, p_j \in \Lambda, N_i \in \mathcal{N} \Rightarrow \bigwedge_{N_j \in \text{Ch}(N_i)} N_j == \text{idx}(p_j)$

Evaluation Setup

- Research Questions
 - Q1: Do our multi-layer specification and neural architecture suggest candidates that are close to the user intent?
 - Q2: What is the impact of the neural architecture in MARS on the performance of a state-of-the-art synthesizer for data wrangling tasks?
 - Q3: How is the performance of MARS affected by the quality of the corpus?
- Experiment Setup
 - Benchmarks: 80 Real-World Challenging Data Wrangling Tasks
 - Dataset: 20,640 StackOverflow Pages of Data Wrangling Tasks
 - 16,459 question-solution pairs for `seq2seq` model
 - 37,748 transactions for association rule mining (*Apriori* algorithm); we obtain 187 *valid*^[1] rules
 - Comparison to MORPHEUS^[2]

[1] A rule is *valid* if its confidence ≥ 0.9 or support ≥ 0.003 , and satisfies all the criteria defined in Chen, Y. et al.. Maximal Multi-layer Specification Synthesis. FSE'19

[2] Feng, Y. et al.. Component-based Synthesis of Table Consolidation and Transformation Tasks from Examples. PLDI'17

Evaluation Results & Analysis

- Timeout: 5 mins
- Ablation Variants
 - *ngram*: built-in statistical model in MORPHEUS
 - *seq2seq*: MARS with seq2seq model
 - *hybrid*: MARS with seq2seq model and preference score refinement (association rules)

Table 1: Statistics for different model rankings.

model	<i>n-gram</i>	<i>seq2seq</i>	<i>hybrid</i>
average [*]	42	25	18
std. ¹	70	39	26

¹ standard deviation.

^{*} computed based on the rankings of the correct solutions.

Table 2: Counts of top-1s and top-3s in different models.

model	<i>n-gram</i>	<i>seq2seq</i>	<i>hybrid</i>
Top-1 total [*]	0	8	11
Top-3 total [*]	2	18	29

^{*} computed based on the rankings of the correct solutions.

Table 3: Statistics of running time.

model	avg. speedup ¹	#timeouts [*]
<i>ngram</i>	1x	11
<i>seq2seq</i>	6x	8
<i>hybrid</i>	15x	2

¹ average speedup on challenging solved benchmarks.

^{*} number of timeouts on all benchmarks.

Discussions


- Limitations
 - Insufficient Text
 - Description of the question is barely useful
 - Contextual Text
 - Some questions require understanding of *pragmatic* contexts, not only semantic
 - Misleading Text
 - User specifies functionality not supported by the DSL
- Threats to Validity
 - Quality of the Corpus
 - Benchmark Selection

*“... I can solve my problem using dplyr’s **mutate** but it’s a time-intensive, roundabout way to achieve my goal. ...”*

*“... I want to use **mutate** to make variable *d* which is mean of *a,b* and *c*. ...”*

Intertitles

Bridging the Logical and Statistical Lands

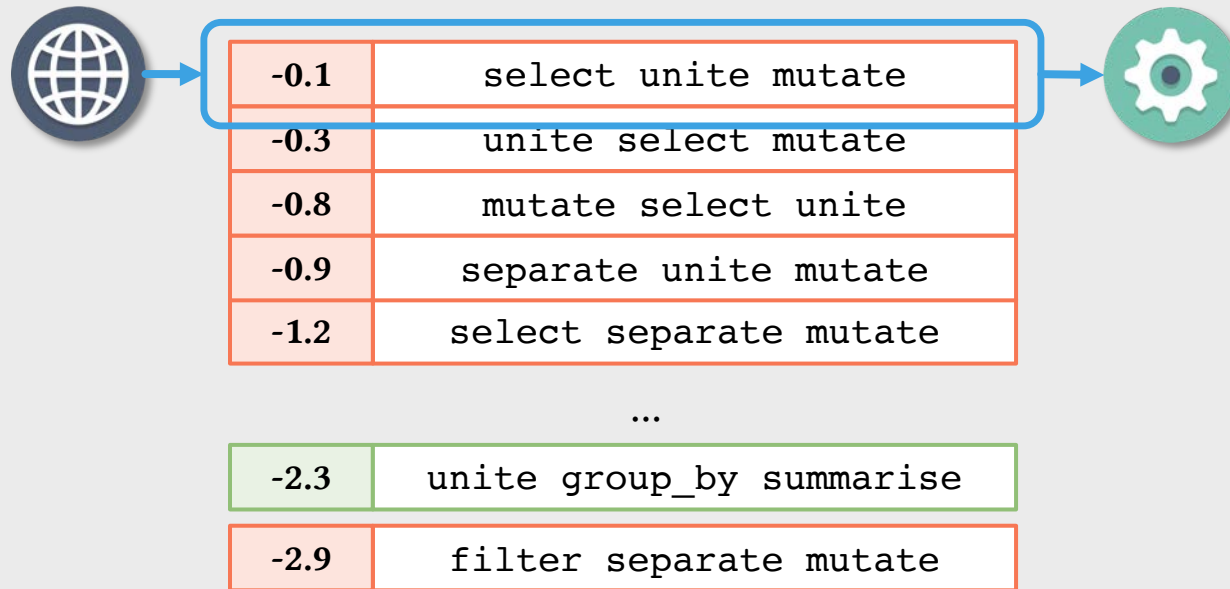


CONCORD^[1]: Deduction-Guided Reinforcement Learning

- Motivations
- Pure Deductive & Statistical Approaches
- Framework Overview
- Formalization
- A Running Example
- Deduction-Guided Reinforcement Learning
 - Deduction Engine
 - Off-Policy Sampling
 - Importance Weighting
- Evaluations
 - Evaluation Setup
 - Evaluation Results & Analysis

Deduction-Guided Reinforcement Learning

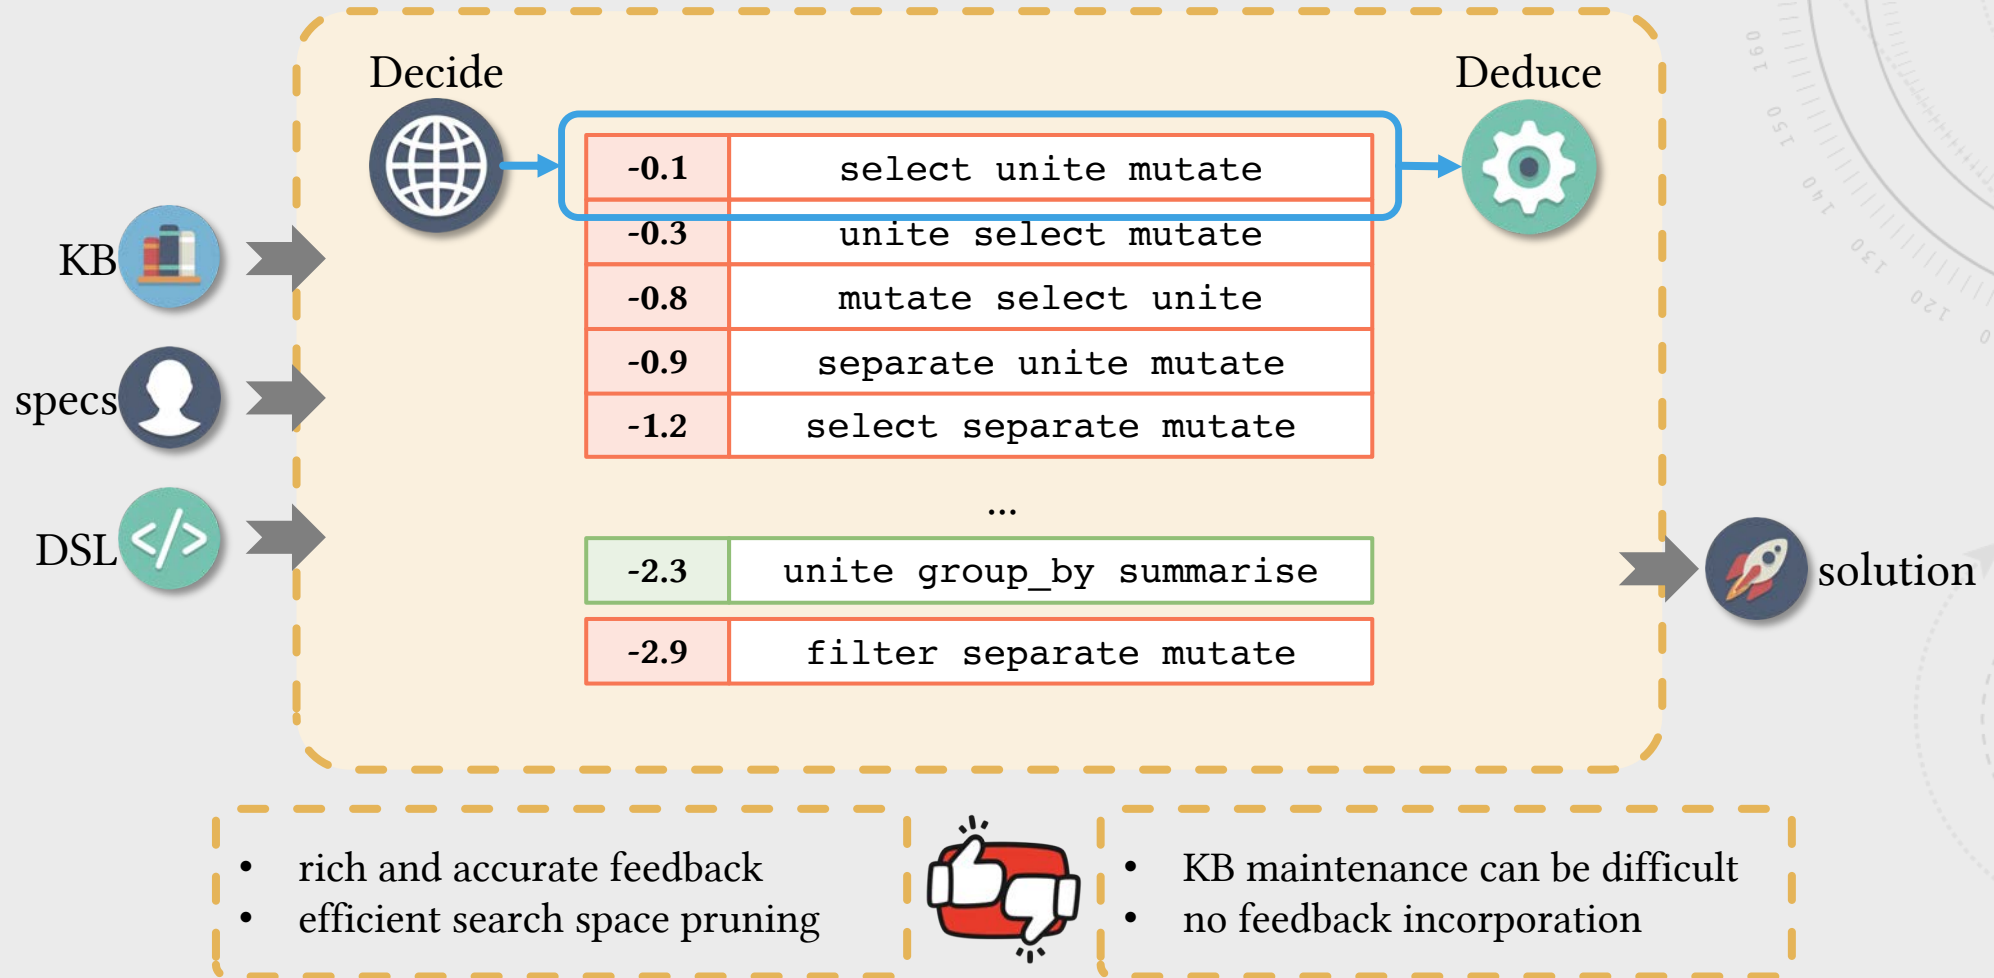
- Motivations
 - Feedback of deduction cannot be seamlessly used by statistical model



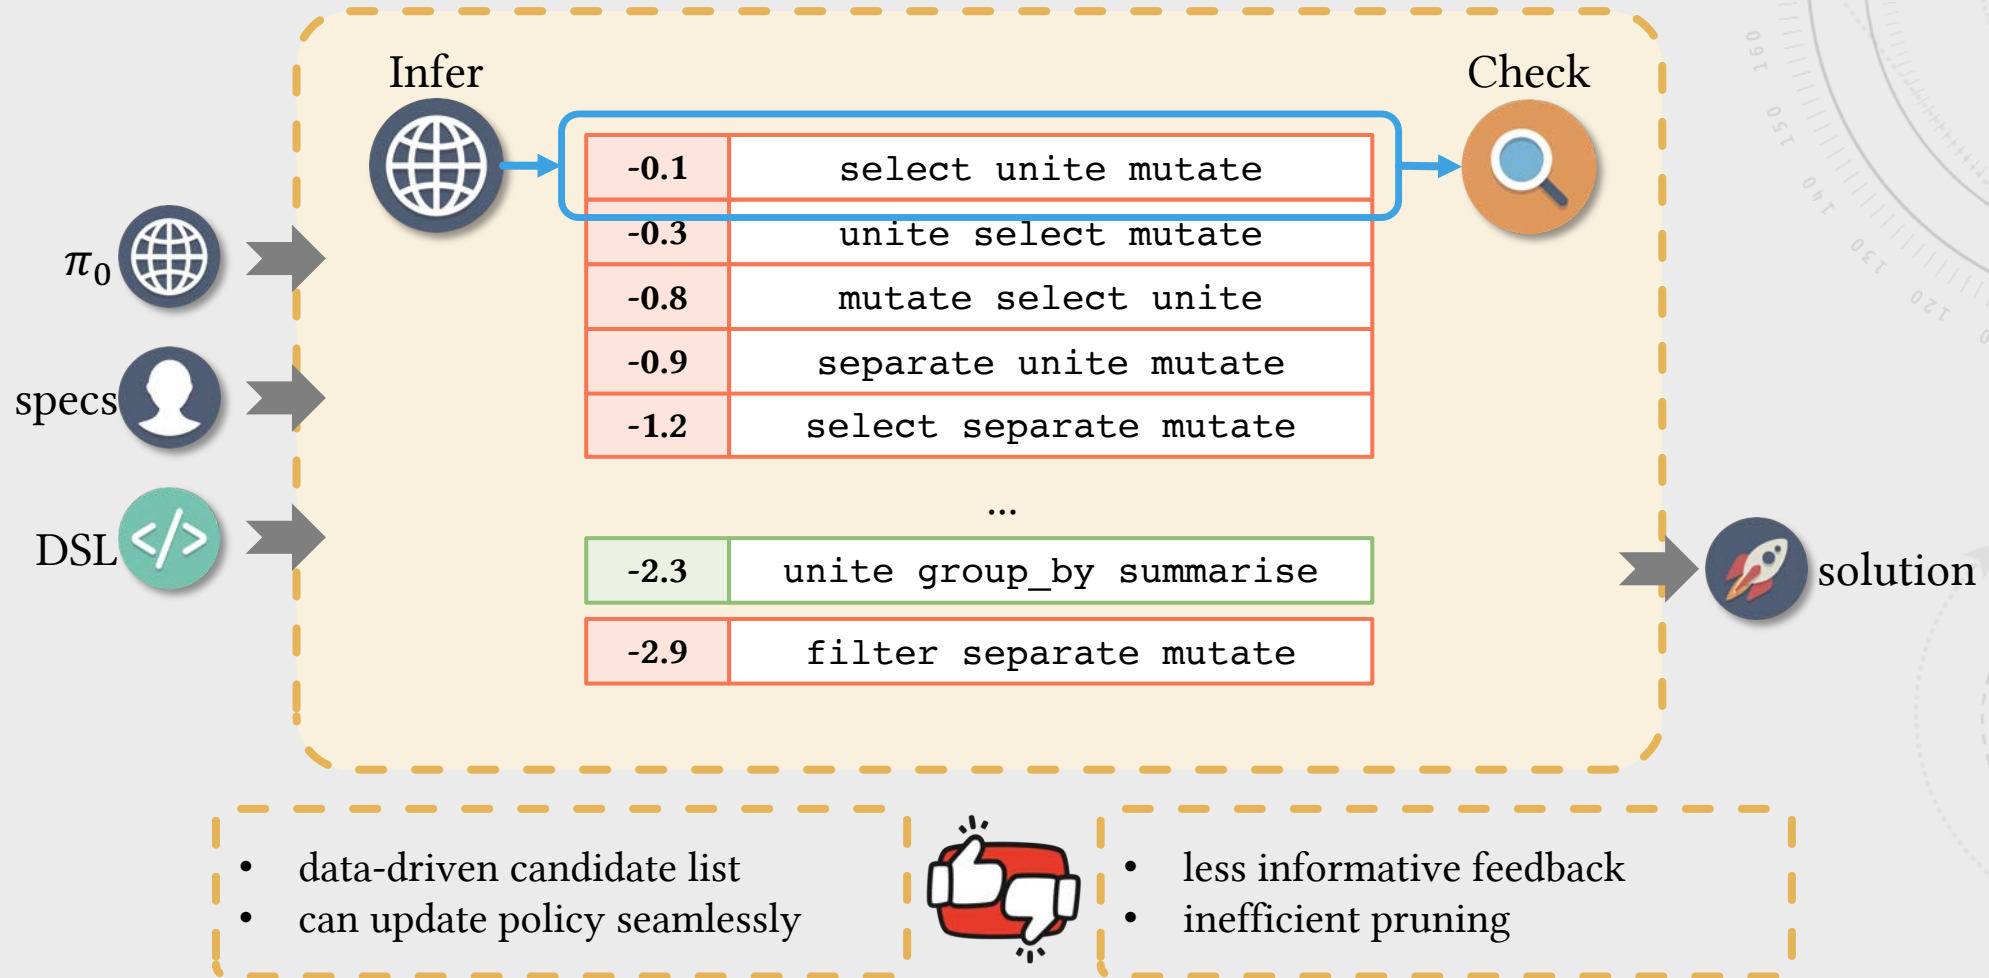
- Statistical estimation is not synchronized with deductive knowledge
- Maintenance of deductive knowledge creates overhead

CONCORD: Deduction-Guided Reinforcement Learning

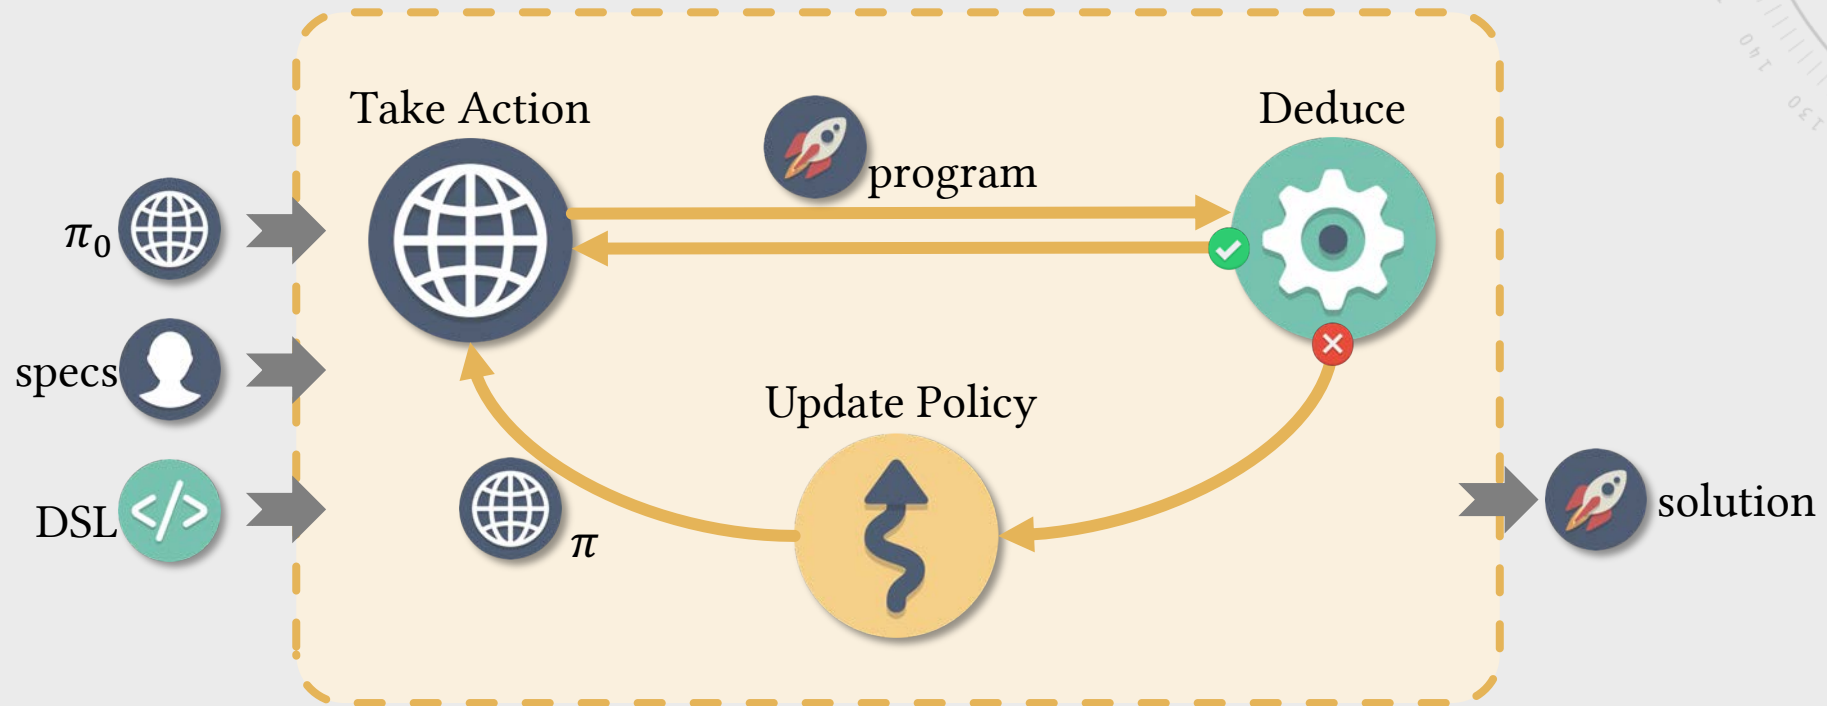
Deductive Approach



CONCORD: Deduction-Guided Reinforcement Learning Statistical Approach

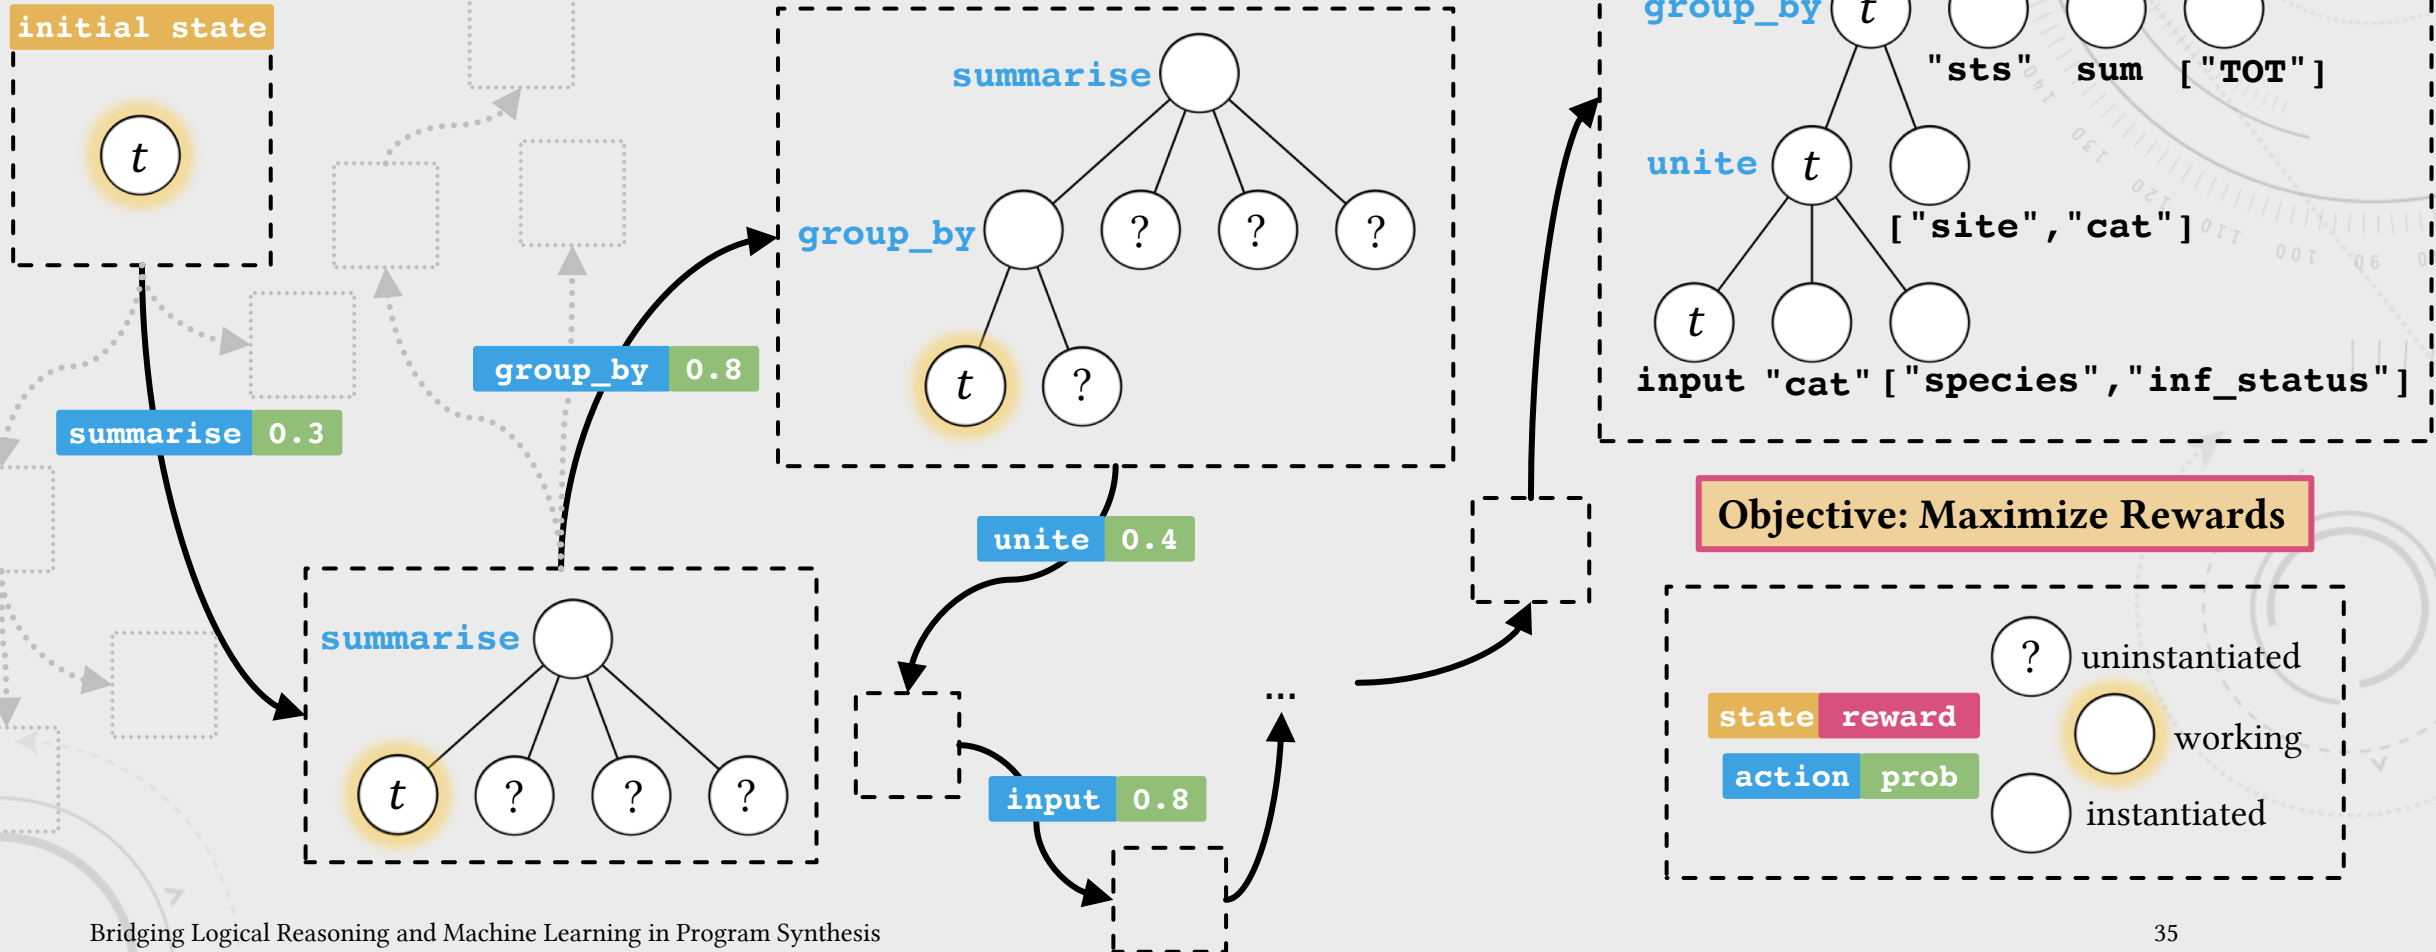


CONCORD: Deduction-Guided Reinforcement Learning Framework Overview

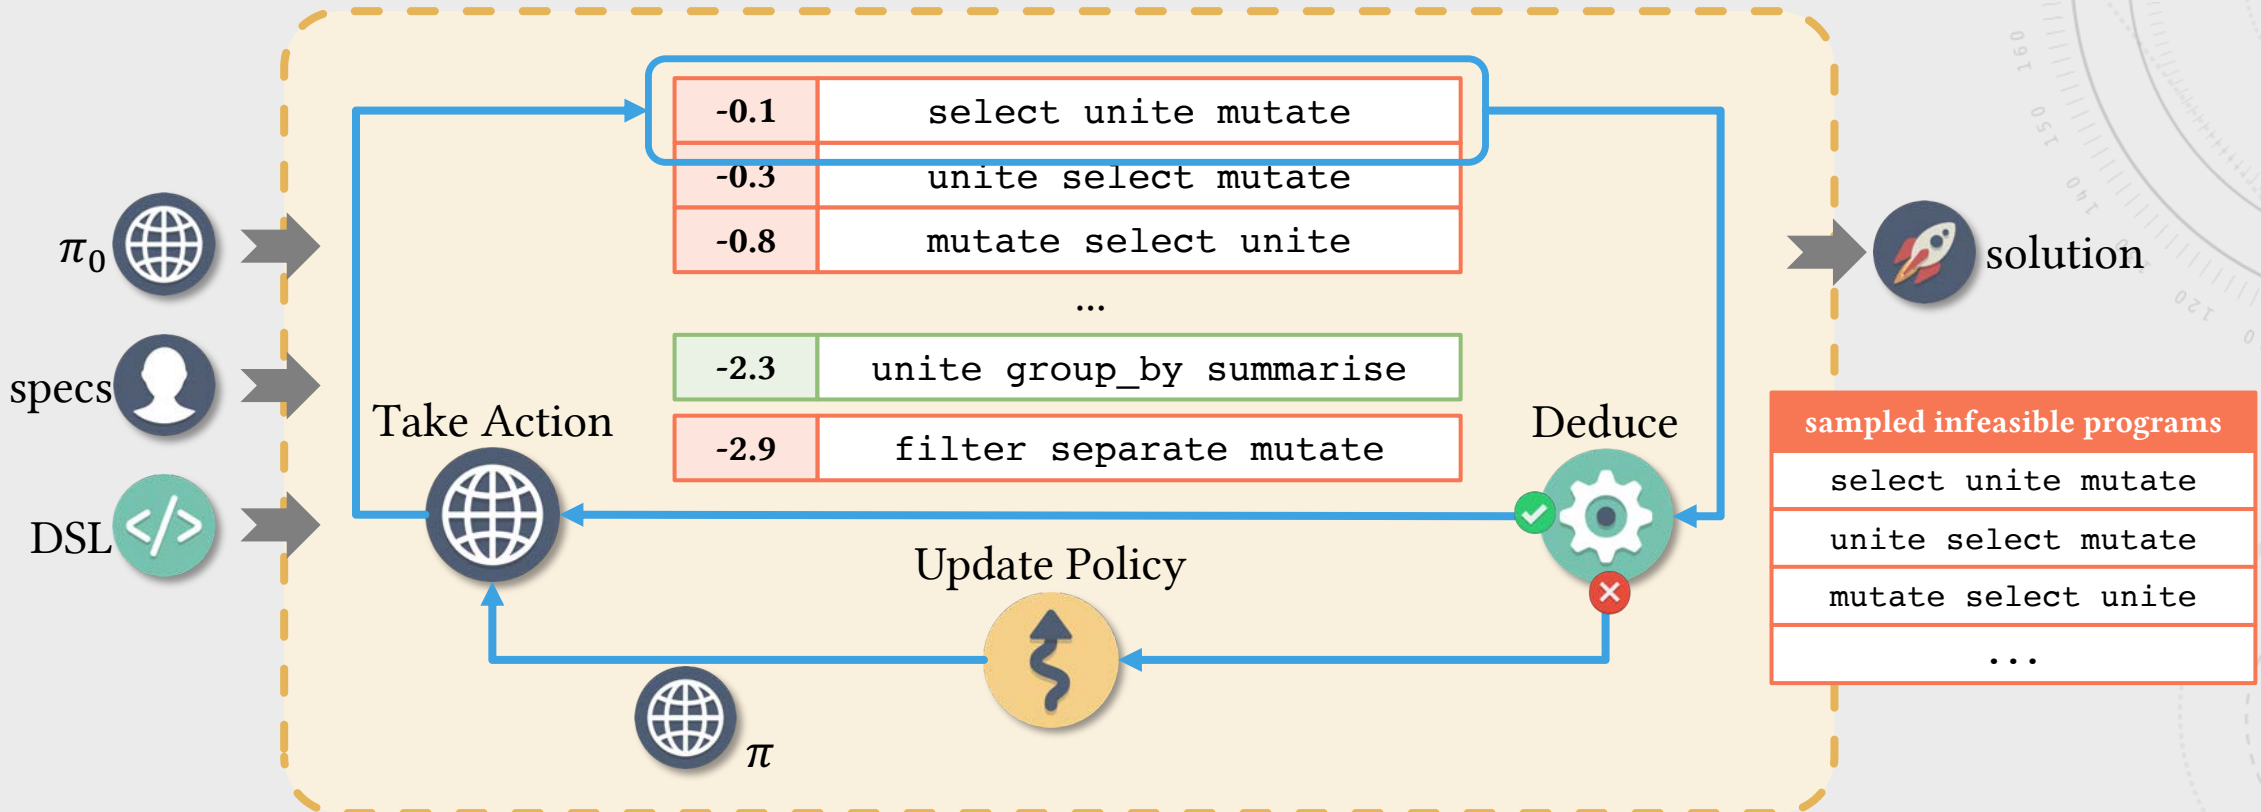


Formalization

- Program Synthesis as Markov Decision Process

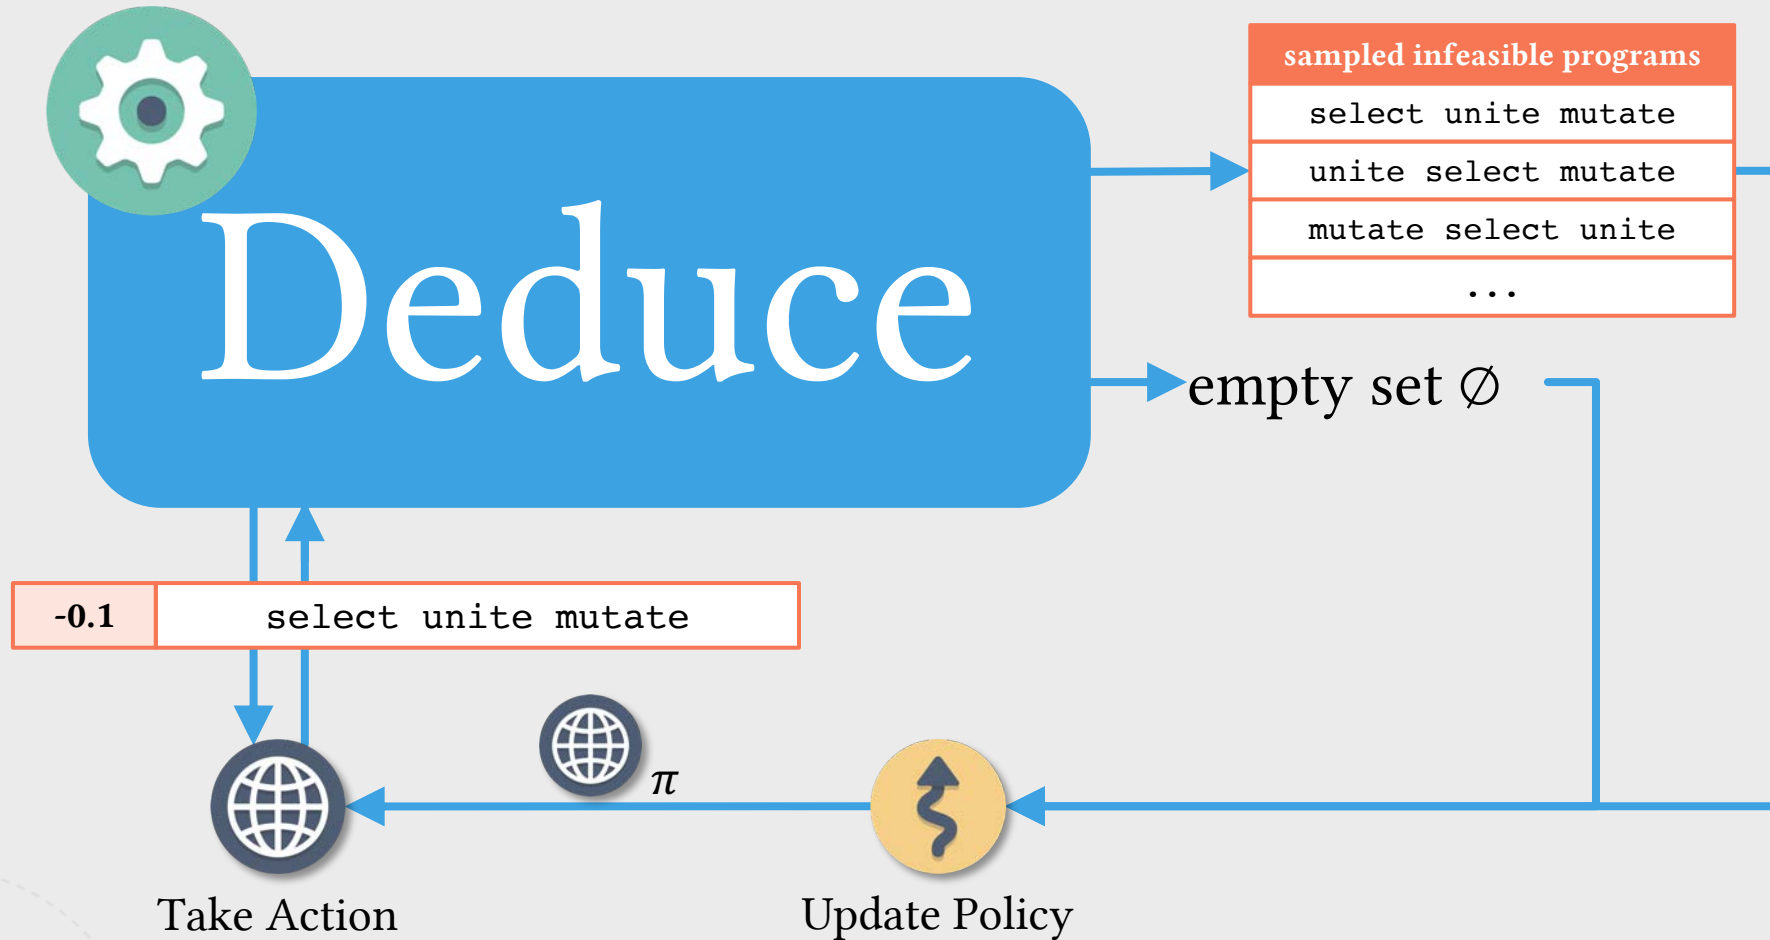


Running Example

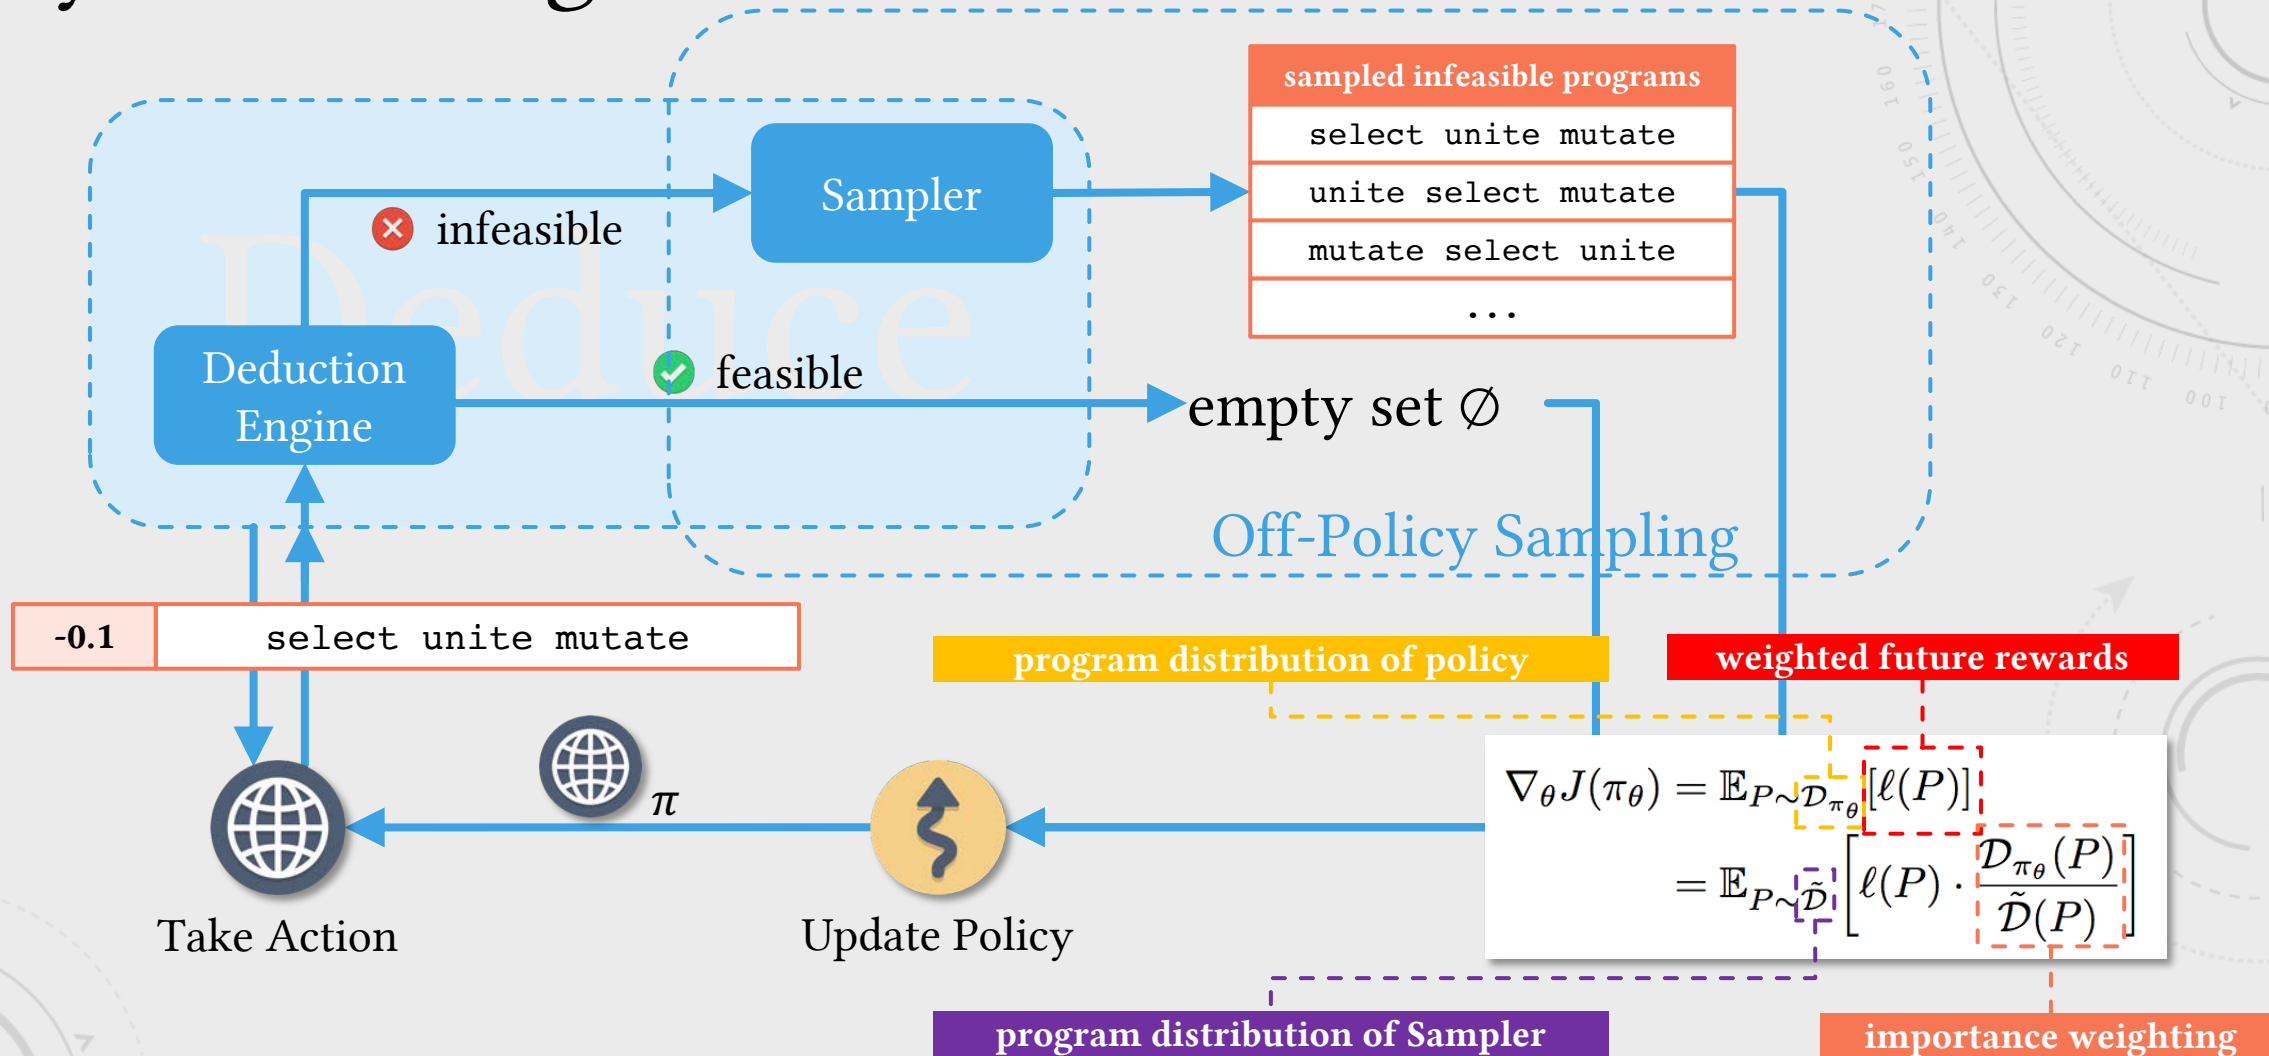


- feedback from deduction flows seamlessly to the policy update
- not only prune the search space, but also promote good candidates

CONCORD: Deduction-Guided Reinforcement Learning Synthesis Algorithm

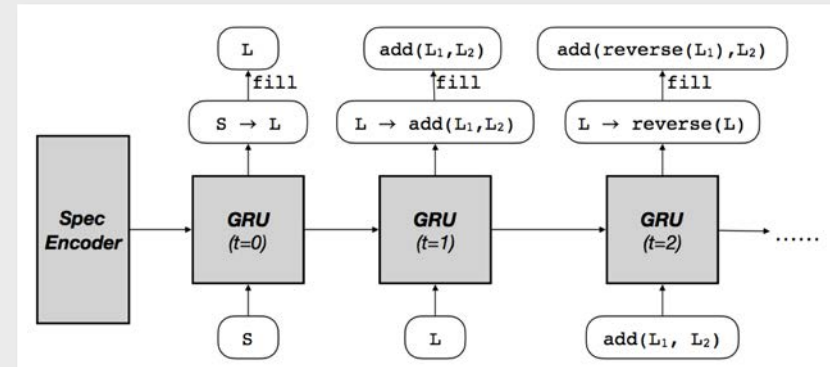


CONCORD: Deduction-Guided Reinforcement Learning Synthesis Algorithm



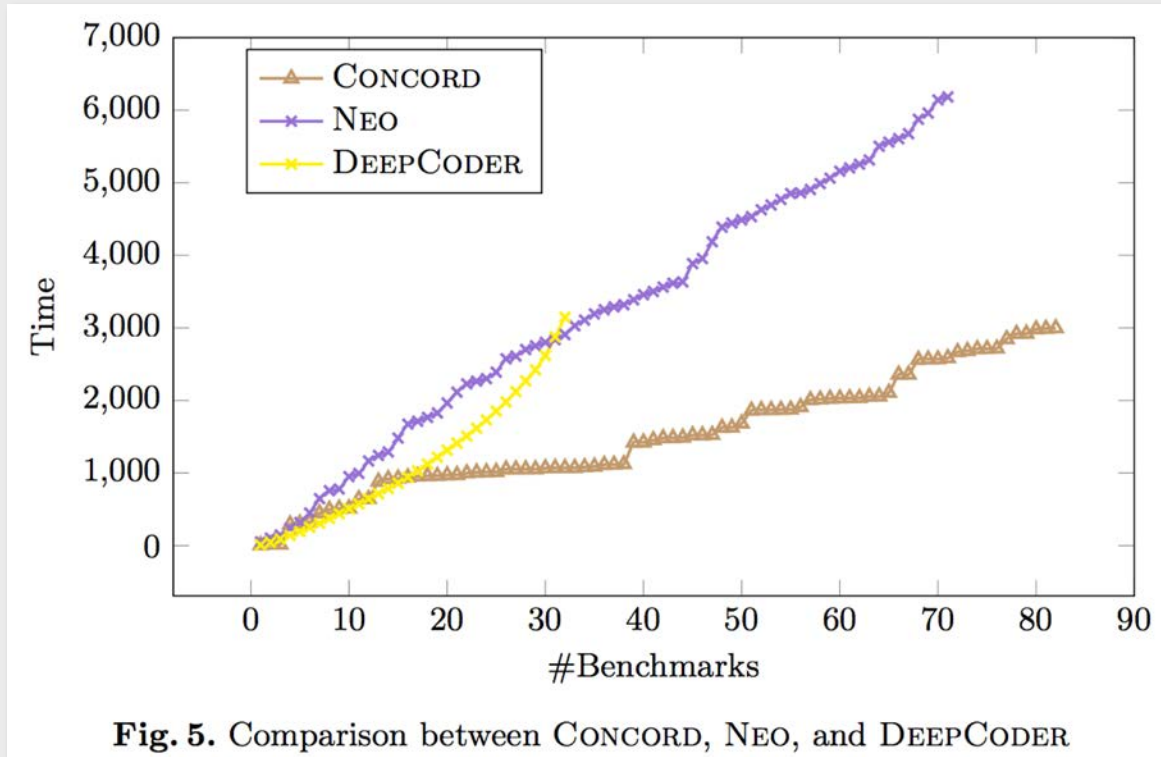
Evaluation Setup

- Research Questions:
 - Q1: How does Concord compare against existing synthesis tools?
 - Q2: How effective is the off-policy RL algorithm compared to standard policy gradient?
- Experiment Setup
 - Deduction Engine: NEO's (Feng et al. 2018) conflict-driven deduction engine
 - Policy: Gated Recurrent Unit (GRU)
 - Benchmarks: DEEPCODER benchmarks used in NEO
 - 100 challenging list processing problems
 - Comparison between:
 - NEO (Feng et al. 2018)
 - DEEPCODER (Balog et al. 2017)



The architecture of the policy network used

Evaluation Results & Analysis



tool	solved	time
CONCORD	82%	36s
NEO	71%	99s
DEEPCODER	32%	205s

tool	solved	speedup over NEO
CONCORD	82%	8.71x


- Concord tightly couples statistical and deductive reasoning based on reinforcement learning.
- The off-policy reinforcement learning technique is effective.

Related Works & Conclusions

- Program Synthesis with Machine Learning (II)
- Related Works
 - METAL
 - PROBE
 - ABL
- Challenges, Conclusions & Future Works

Related Works & Conclusions

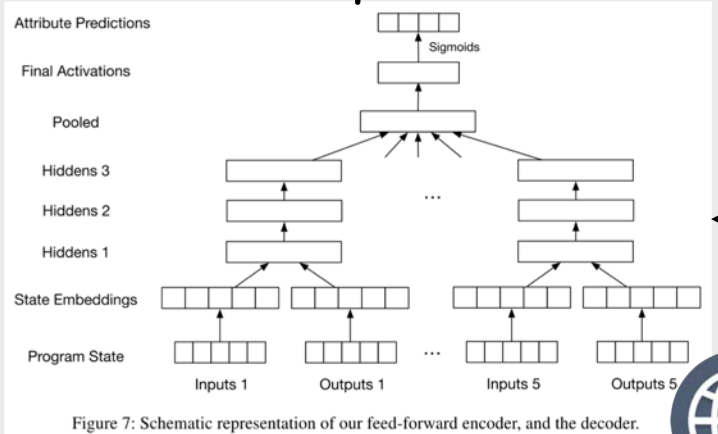
Program Synthesis with Machine Learning (II)



```

a ← [int]
b ← FILTER (<0) a
c ← MAP (*4) b
d ← SORT c
e ← REVERSE d
    
```

Proposed Program



Neural Encoder

Input:
[-17, -3, 4, 11, 0, -5, -9, 13, 6, 6, -8, 11]

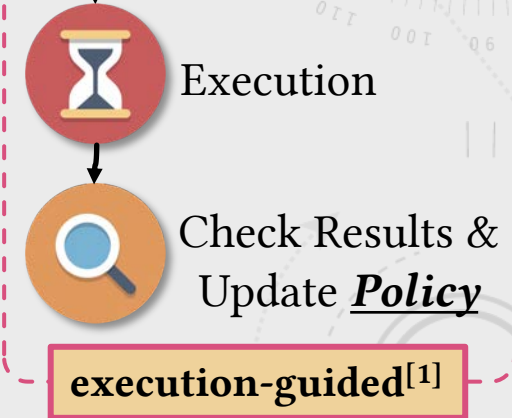
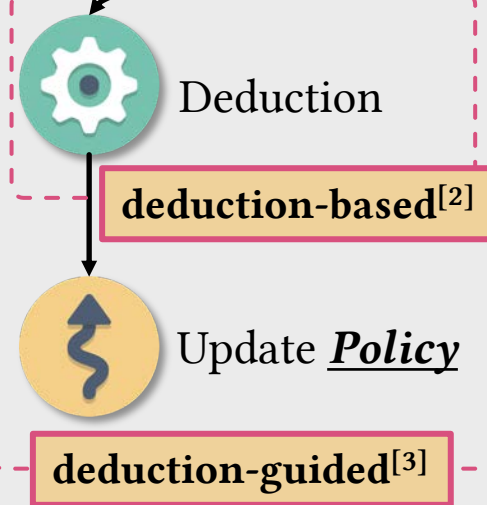
Output:
[-12, -20, -32, -36, -68]



representation learning

Natural Language

multi-modal encoding



[1] SEQ2SQL (Zhong et al. 2017); EXEC (Chen et al. 2018); AutoPandas (Bavishi et al. 2019);

[2] NEO (Feng et al. 2018); SQLIZER (Yaghmazadeh et al. 2018); MARS (Chen et al. 2019); REGEL (Chen et al. 2020); VISER (Wang et al. 2020)

[3] METAL (Si et al. 2019); SKETCHADAPT (Nye et al. 2018); PROBE (Barke et al. 2020); CONCORD (Chen et al. 2020);

[4] DIALSQL (Gur et al. 2018);

Related Works

- METAL^[1]
 - Circuit Synthesis
 - Invoke a SAT solver to generate a counter-example which adds to the test cases
- PROBE^[2]
 - String Transformation & Bitvector & Circuit Synthesis
 - Just-in-Time Learning: updates a PCFG during synthesis by learning from partial solutions
- ABL^[3]
 - Handwritten Equation Decipherment
 - Improve machine learning models using abductive learning

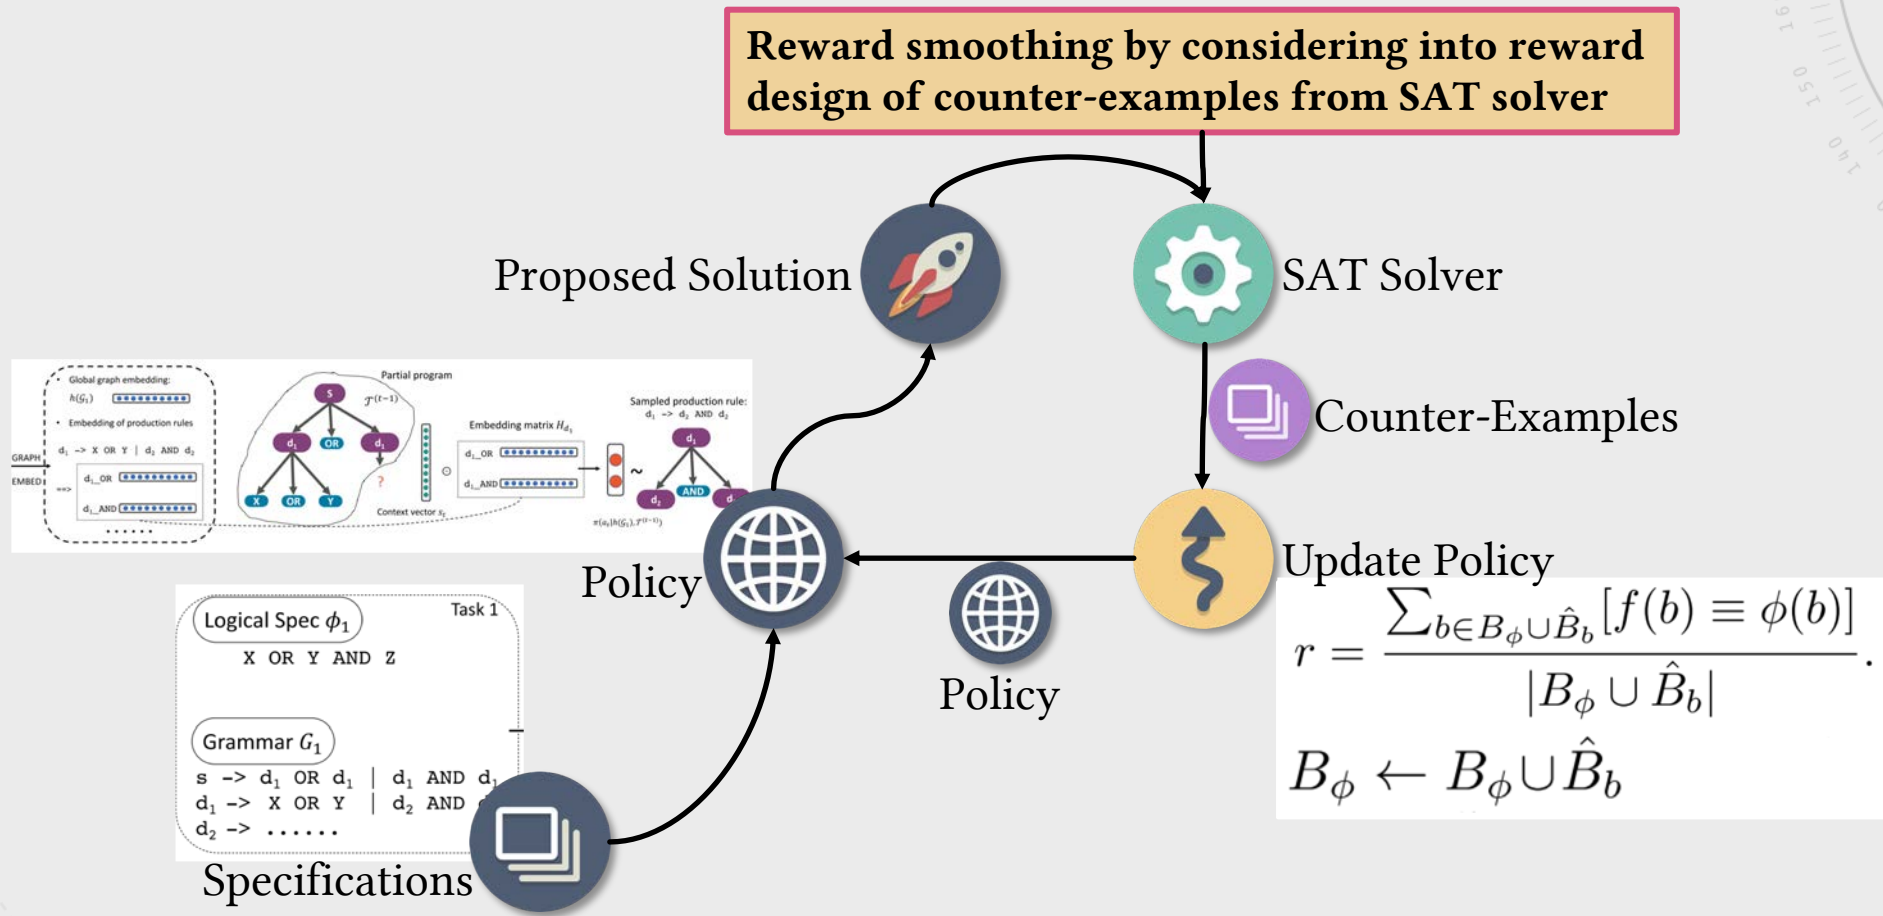
[1] Si, X. et al. Learning a Meta-Solver for Syntax-Guided Program Synthesis. ICLR'19

[2] Barke, S. et al. Just-in-Time Learning for Bottom-up Enumerative Synthesis. OOPSLA'20

[3] Dai, W.-Z. et al. Bridging Machine Learning and Logical Reasoning by Abductive Learning. NeurIPS'19

METAL^[1] (The Reinforcement Learning Part)

Reward smoothing by considering into reward design of counter-examples from SAT solver



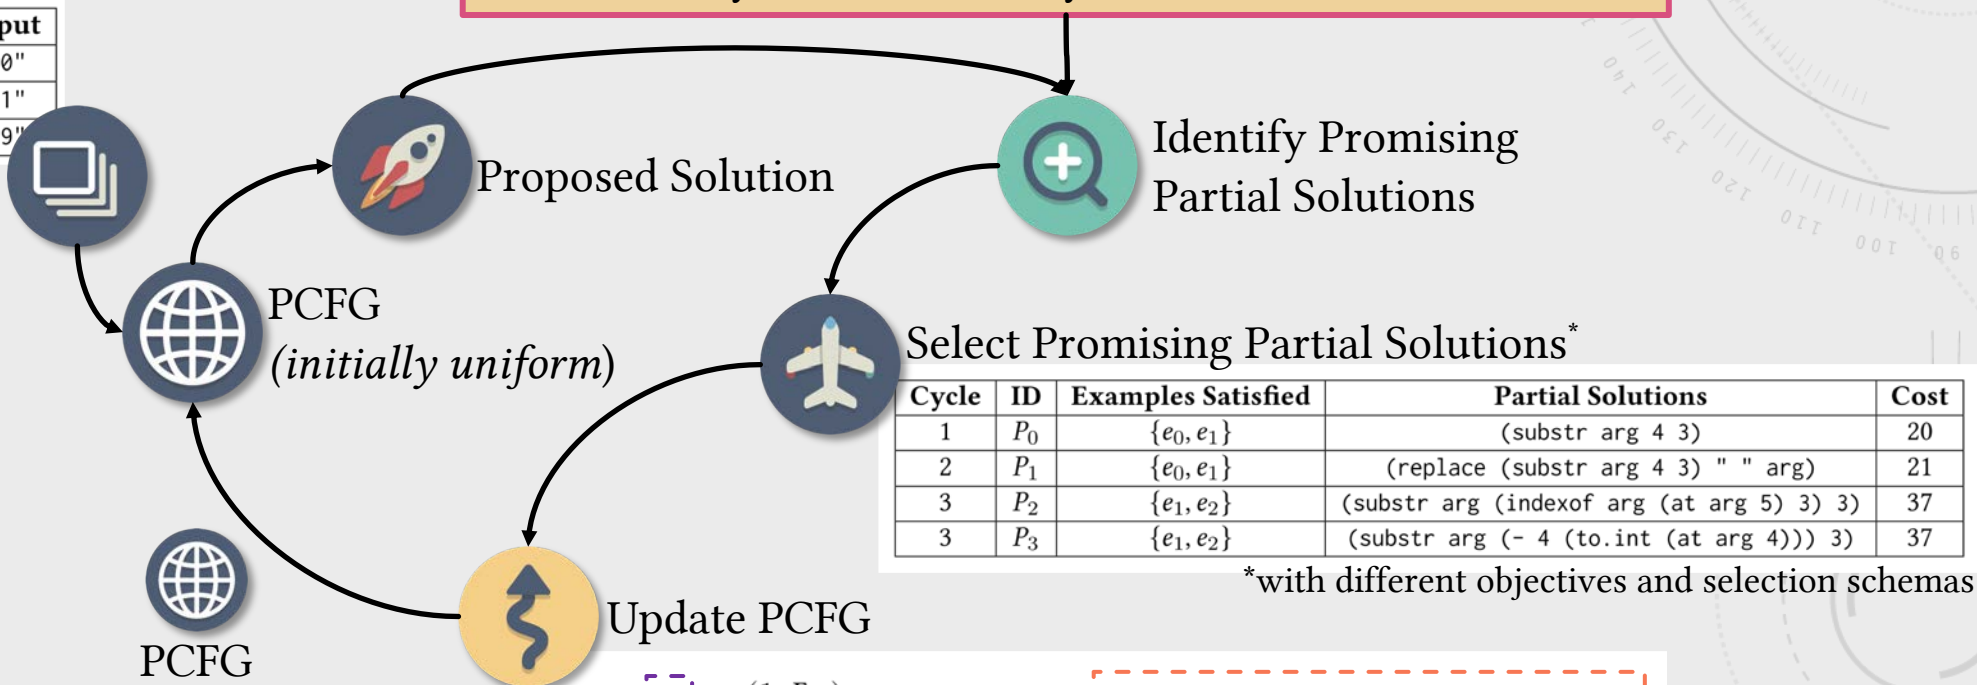
[1] Si, X. et al. Learning a Meta-Solver for Syntax-Guided Program Synthesis. ICLR'19
 Bridging Logical Reasoning and Machine Learning in Program Synthesis

PROBE^[1] (The Just-in-Time Learning Part)

Programs that satisfy a subset of the semantic specification often share syntactic similarity with the full solution.

ID	Input	Output
e_0	"+95 310-537-401"	"310"
e_1	"+72 001-050-856"	"001"
e_2	"+106 769-858-438"	"769"

Examples



$$p(R) = \frac{p_{ui}(R)^{(1-FIT)}}{Z}$$

uniform distribution

where

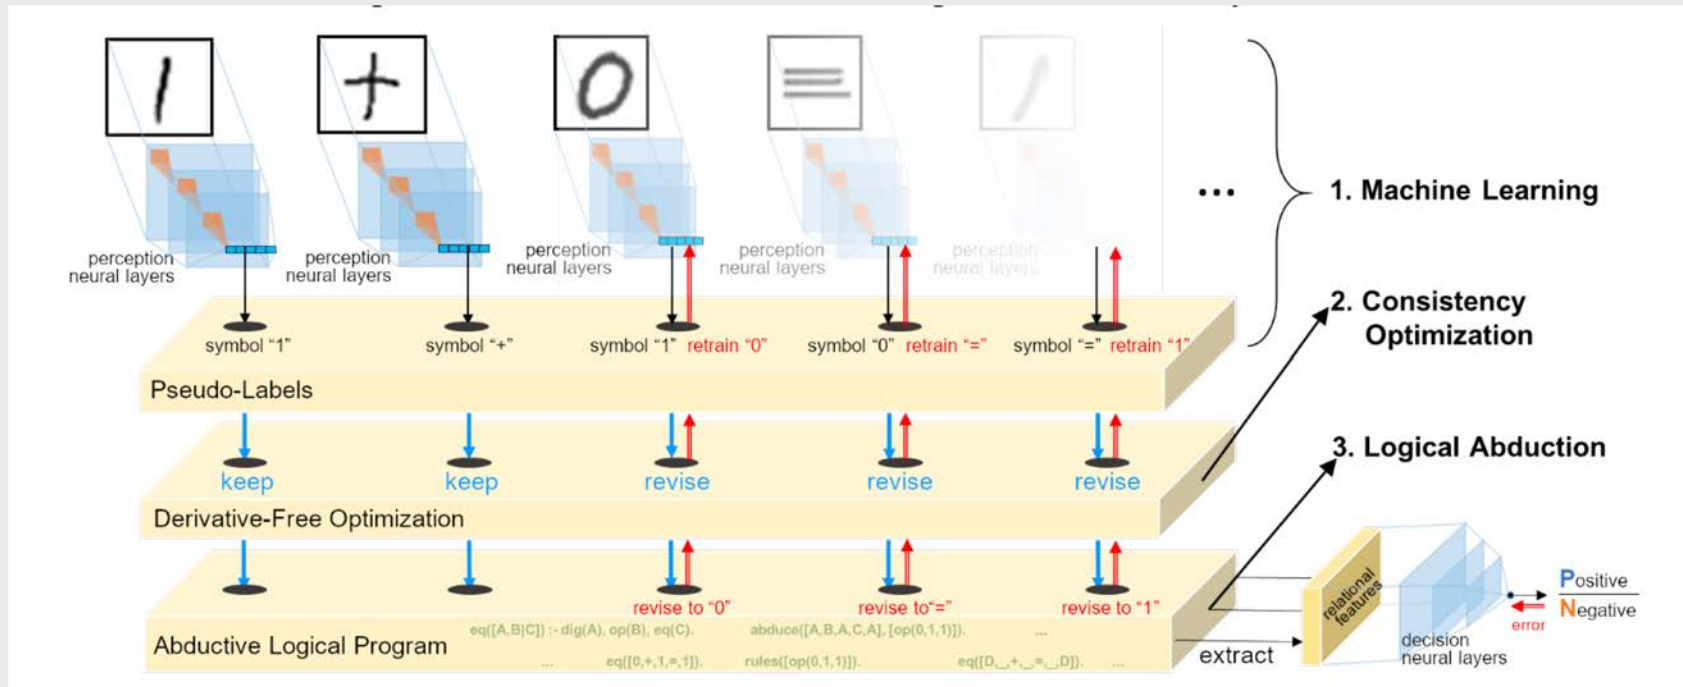

$$FIT = \max_{\{P \in PSol | R \in tr(P)\}} \frac{|\mathcal{E} \cap E[P]|}{|\mathcal{E}|}$$

highest proportion of IO satisfied

[1] Barke, S. et al. Just-in-Time Learning for Bottom-up Enumerative Synthesis. OOPSLA'20
 Bridging Logical Reasoning and Machine Learning in Program Synthesis

Related Works & Conclusions

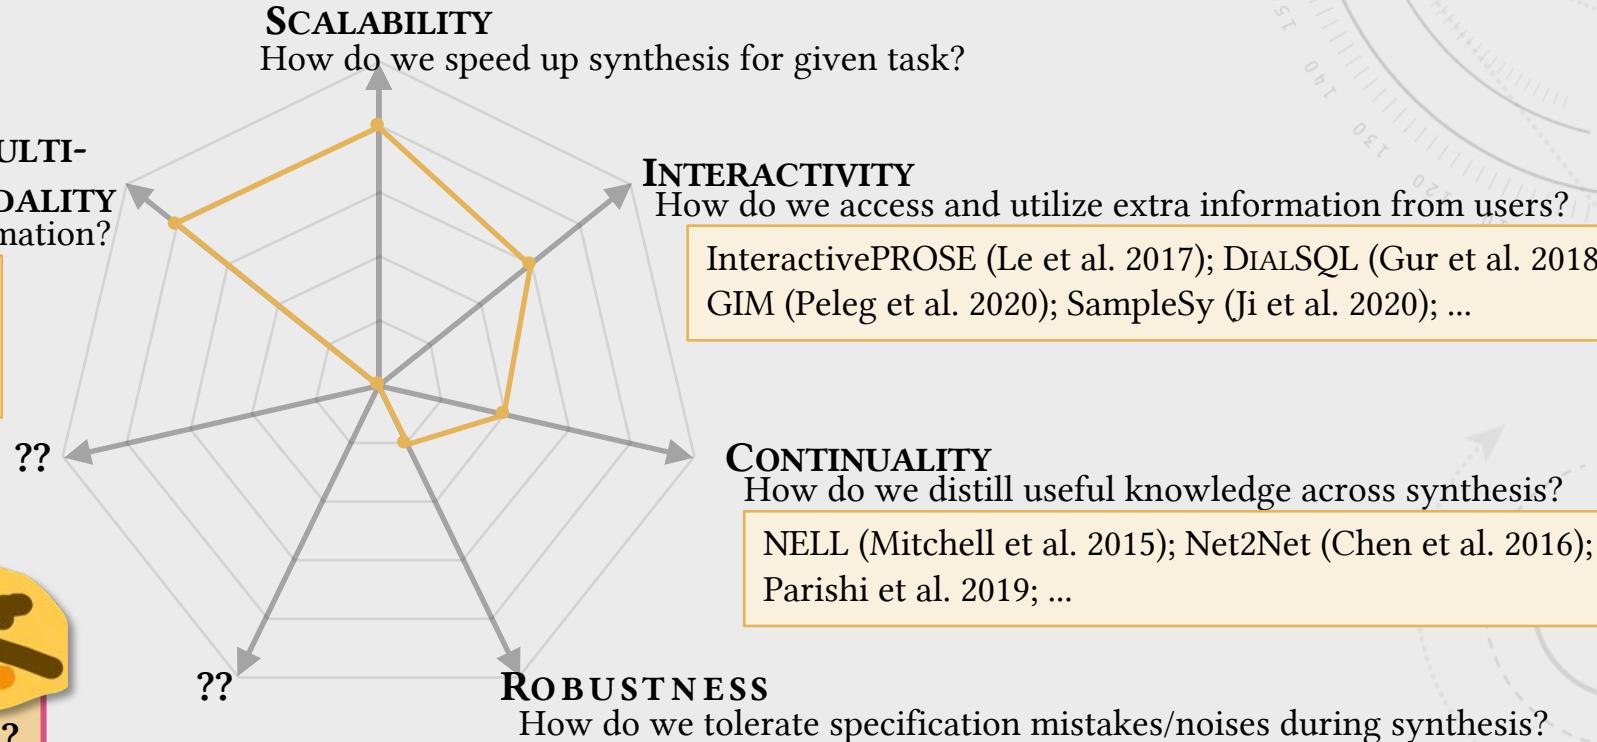
ABL^[1]: A Brief Overview



[1] Dai, W.-Z. et al. Bridging Machine Learning and Logical Reasoning by Abductive Learning. NeurIPS'19
Bridging Logical Reasoning and Machine Learning in Program Synthesis

Challenges, Conclusions & Future Works

DEEPCODER (Balog et al. 2017); EXEC (Chen et al. 2018); NEO (Feng et al. 2018); SQLIZER (Yaghmazadeh et al. 2018); AutoPandas (Bavishi et al. 2019); METAL (Si et al. 2019); SKETCHADAPT (Nye et al. 2018); PROBE (Barke et al. 2020); CONCORD (Chen et al. 2020); ...



SEQ2SQL (Zhong et al. 2017); MARS (Chen et al. 2019); REGEL (Chen et al. 2020); VISER (Wang et al. 2020); ...

InteractivePROSE (Le et al. 2017); DIALSQL (Gur et al. 2018); GIM (Peleg et al. 2020); SampleSy (Ji et al. 2020); ...

NELL (Mitchell et al. 2015); Net2Net (Chen et al. 2016); Parish et al. 2019; ...

FLASHFILL (Gulwani 2011); RULESYNTH (Singh 2017); BESTER (Peleg et al. 2020); ...

... and some more interesting dimensions?



References I

- Gulwani, S. Automating String Processing in Spreadsheets using Input-Output Examples. In POPL'11
- Berant, J., Chou, A., Frostig, R., & Liang, P. Semantic Parsing on Freebase from Question-Answer Pairs. In EMNLP'13
- Mitchell, T., Cohen, W., Hruschka, E., Talukdar, P., Betteridge, J., Carlson, A., ... Welling, J. Never-Ending Learning. In AAAI'15
- Chen, T., Goodfellow, I. J., & Shlens, J. Net2Net: Accelerating Learning via Knowledge Transfer. In ICLR'16
- Balog, M., Gaunt, A. L., Brockschmidt, M., Nowozin, S., & Tarlow, D. DeepCoder: Learning to Write Programs. In ICLR'17
- Feng, Y., Martins, R., Van Geffen, J., Dillig, I., & Chaudhuri, S. Component-based Synthesis of Table Consolidation and Transformation Tasks from Examples. In PLDI'17
- Le, V., Perelman, D., Polozov, O., Raza, M., Udupa, A., & Gulwani, S. Interactive Program Synthesis. CoRR, abs/1703.0
- Singh, R., Meduri, V. V., Elmagarmid, A., Madden, S., Papotti, P., Quiané-Ruiz, J.-A., ... Tang, N. Synthesizing Entity Matching Rules by Examples. In VLDB'17
- Yaghmazadeh, N., Wang, Y., Dillig, I., & Dillig, T. SQLizer: Query Synthesis from Natural Language. In OOPSLA'17
- Zhong, V., Xiong, C., & Socher, R. Seq2SQL: Generating Structured Queries from Natural Language using Reinforcement Learning. CoRR, abs/1709.0
- Dong, L., & Lapata, M. Coarse-to-Fine Decoding for Neural Semantic Parsing. In ACL'18
- Feng, Y., Martins, R., Bastani, O., & Dillig, I. Program Synthesis using Conflict-Driven Learning. In PLDI'18
- Gur, I., Yavuz, S., Su, Y., & Yan, X. DialSQL: Dialogue Based Structured Query Generation. In ACL'18
- Peleg, H., Shoham, S., & Yahav, E. Programming Not Only by Example. In ICSE'18
- Bavishi, R., Lemieux, C., Fox, R., Sen, K., & Stoica, I. AutoPandas: Neural-backed Generators for Program Synthesis. In OOPSLA'19

References II

- Chen, X., Liu, C., & Song, D. Execution-Guided Neural Program Synthesis. In ICLR'19
- **Chen, Y., Martins, R., & Feng, Y. Maximal Multi-layer Specification Synthesis. In FSE'19**
- Dai, W.-Z., Xu, Q., Yu, Y., & Zhou, Z.-H. Bridging Machine Learning and Logical Reasoning by Abductive Learning. In NeurIPS'19
- **Martins, R., Chen, J., Chen, Y., Feng, Y., & Dillig, I. Trinity: An Extensible Synthesis Framework for Data Science. In VLDB'19**
- Nye, M., Hewitt, L., Tenenbaum, J., & Solar-Lezama, A. Learning to Infer Program Sketches. In ICML'19
- Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. Continual lifelong learning with neural networks: A review. *Neural Networks*, 113, 54–71.
- Si, X., Yang, Y., Dai, H., Naik, M., & Song, L. Learning a Meta-Solver for Syntax-Guided Program Synthesis. In ICLR'19
- Barke, S., Peleg, H., & Polikarpova, N. Just-in-Time Learning for Bottom-up Enumerative Synthesis. In OOPSLA'20
- Chen, Q., Wang, X., Ye, X., Durrett, G., & Dillig, I. Multi-Modal Synthesis of Regular Expressions. In PLDI'20
- **Chen, Y., Wang, C., Bastani, O., Dillig, I., & Feng, Y. Program Synthesis Using Deduction-Guided Reinforcement Learning. In CAV'20**
- Ji, R., Liang, J., Xiong, Y., Zhang, L., & Hu, Z. Question Selection for Interactive Program Synthesis. In PLDI'20
- **Mariano, B., Chen, Y., Feng, Y., Lahiri, S., & Dillig, I. Demystifying Loops in Smart Contracts. In ASE'20**
- Peleg, H., & Polikarpova, N. Perfect is the Enemy of Good: Best-Effort Program Synthesis. In ECOOP'20
- Wang, C., Feng, Y., Bodik, R., Cheung, A., & Dillig, I. Visualization by Example. In POPL'20