Learning to Dynamically Optimise Algorithms

André Biedenkapp

10.11.2022

Motivation

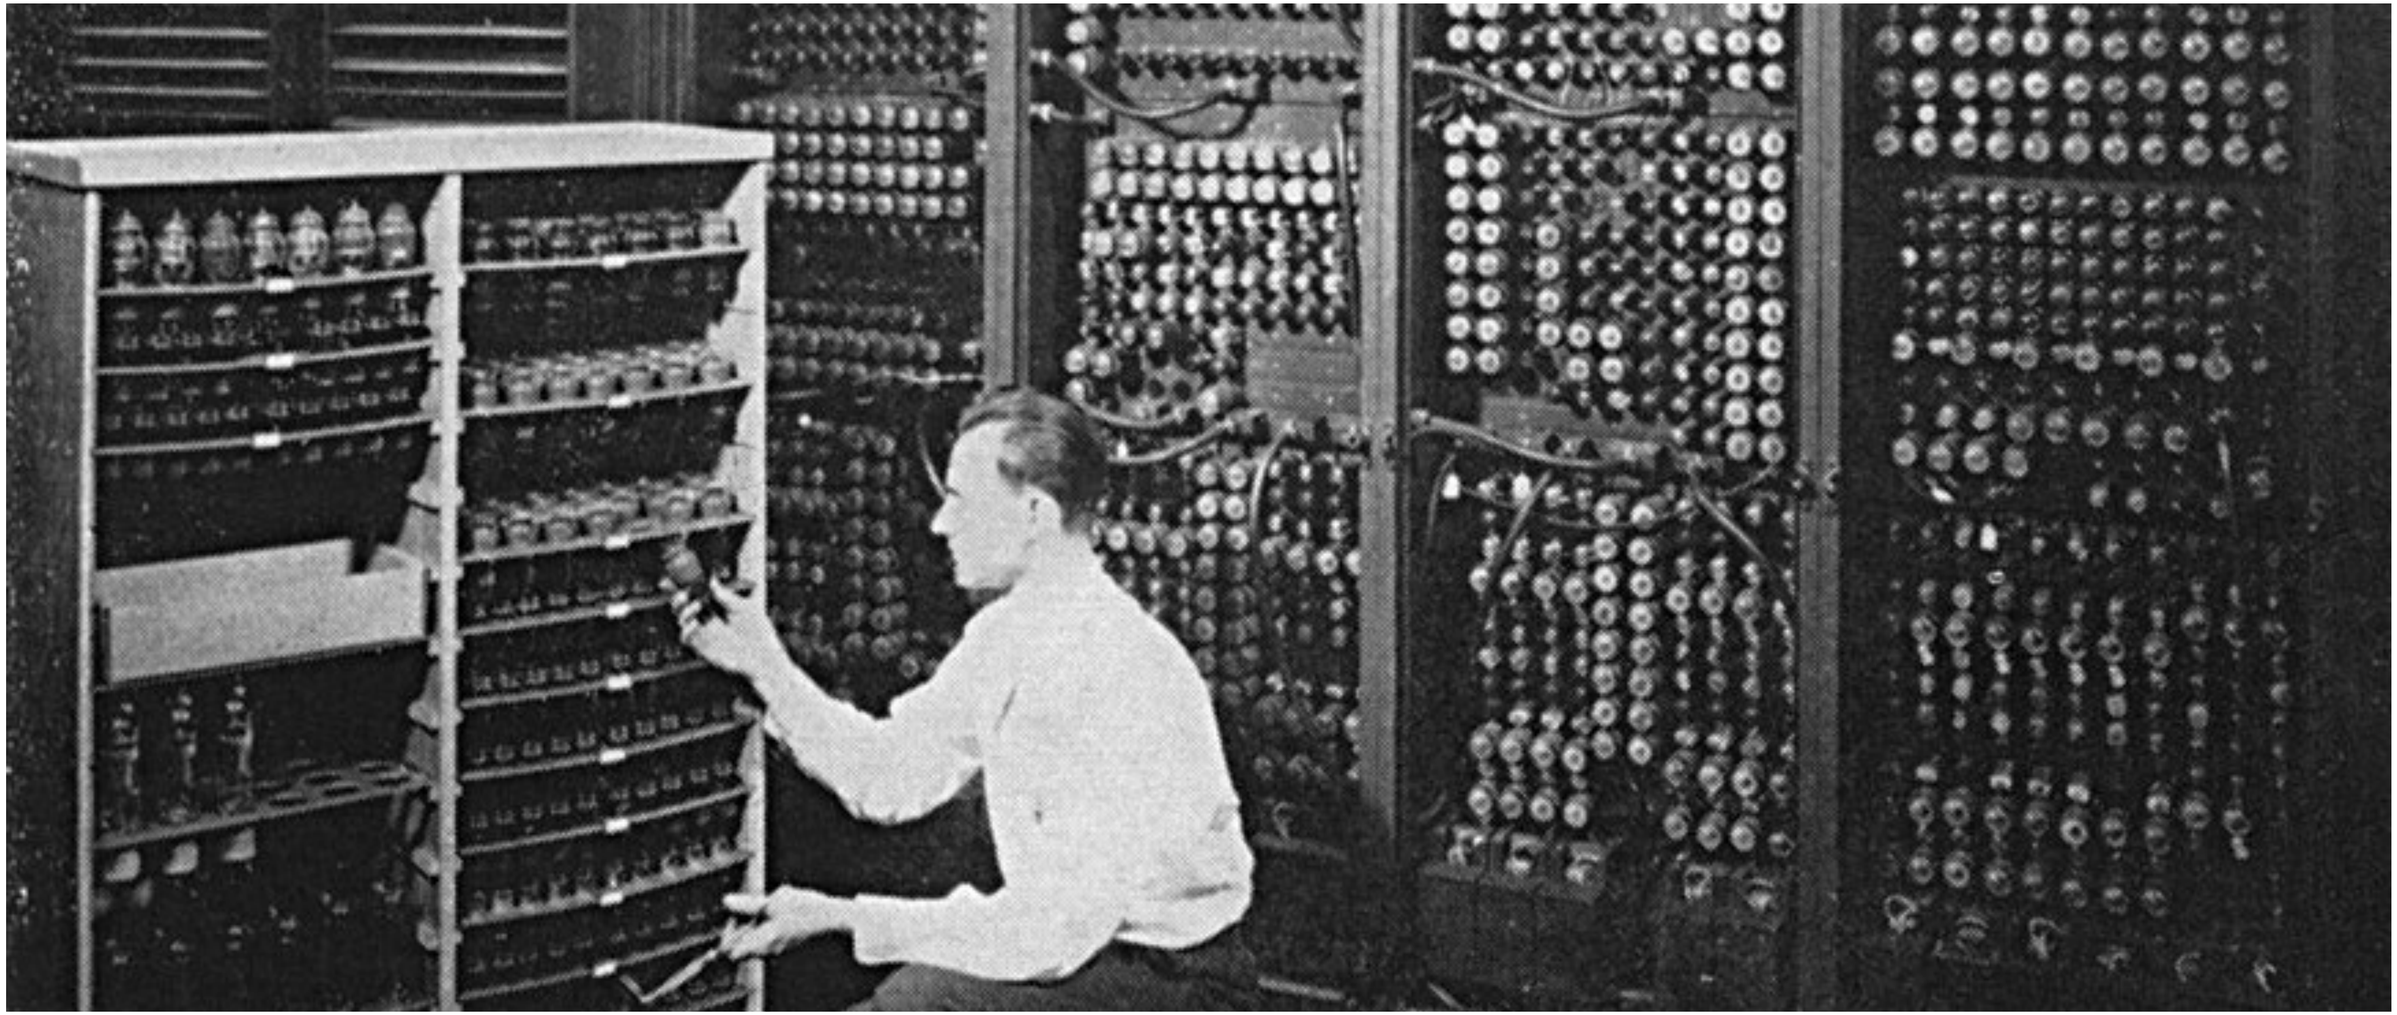

Image Source: *https://commons.wikimedia.org/wiki/File:ENIAC-changing_a_tube.jpg*

Motivation

Image Source: By Cornellanense (talk) 07:14, 29 October 2008 (UTC) - Own work, CC BY-SA 4.0, https://commons.wikimedia.org/w/index.php?curid=5089958

Motivation

Image Source: By A.Savin - Own work, FAL, *https://commons.wikimedia.org/w/index.php?curid=98539643*

Motivation

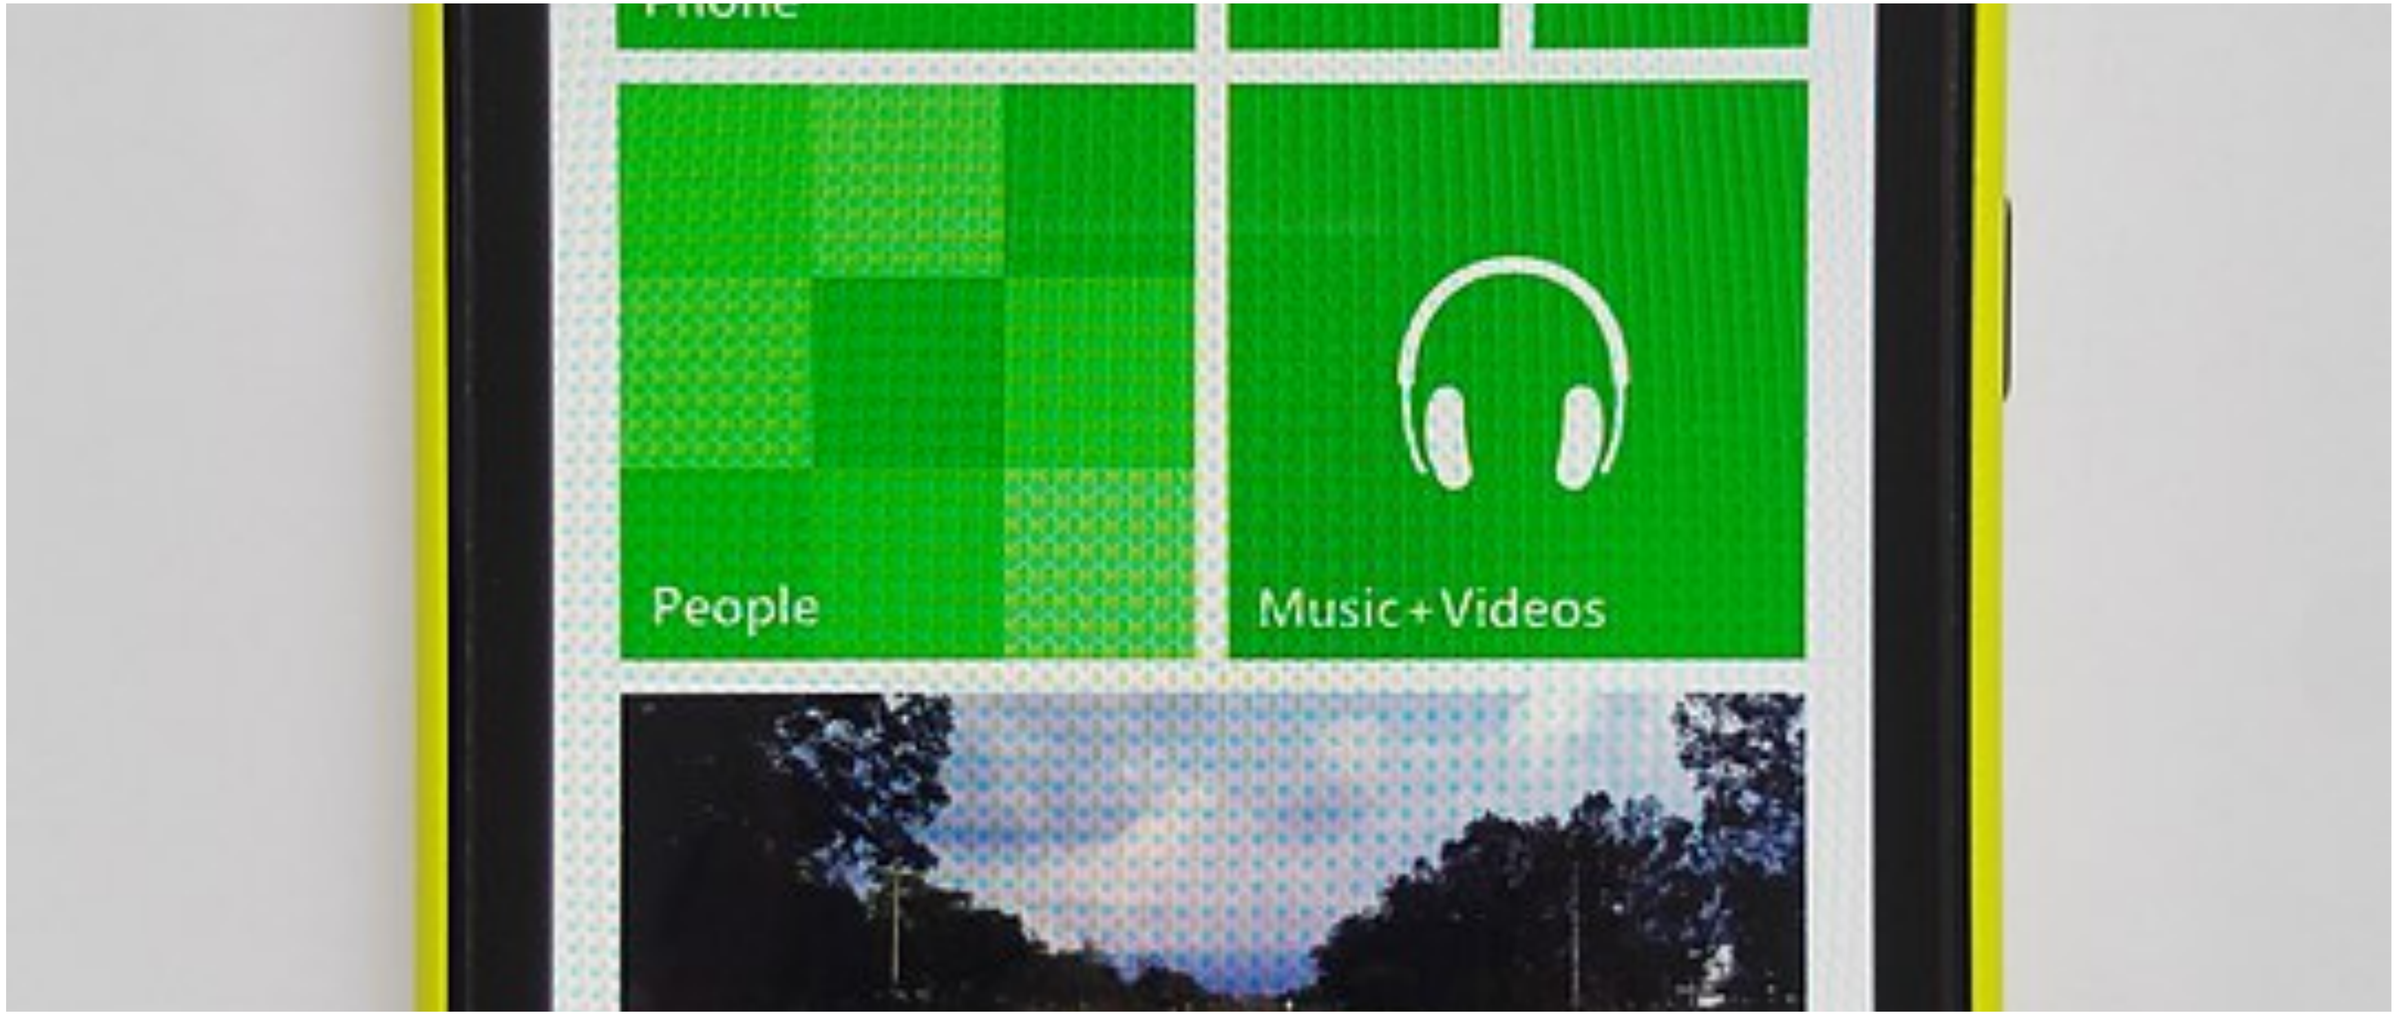

Image Source: By MY PHONE - For the screen: Windows Phone 8 en.svg, CC BY-SA 4.0, https://commons.wikimedia.org/w/index.php?curid=101547605

Motivation

Image Source: By Dllu - Own work, CC BY-SA 4.0, https://commons.wikimedia.org/w/index.php?curid=64517567

Motivation

Image Source: By stockcatalog - https://www.flickr.com/photos/151691693@N02/40970932765, CC BY 2.0, https://commons.wikimedia.org/w/index.php?curid=111183098

Motivation

Problem → Algorithm

Motivation

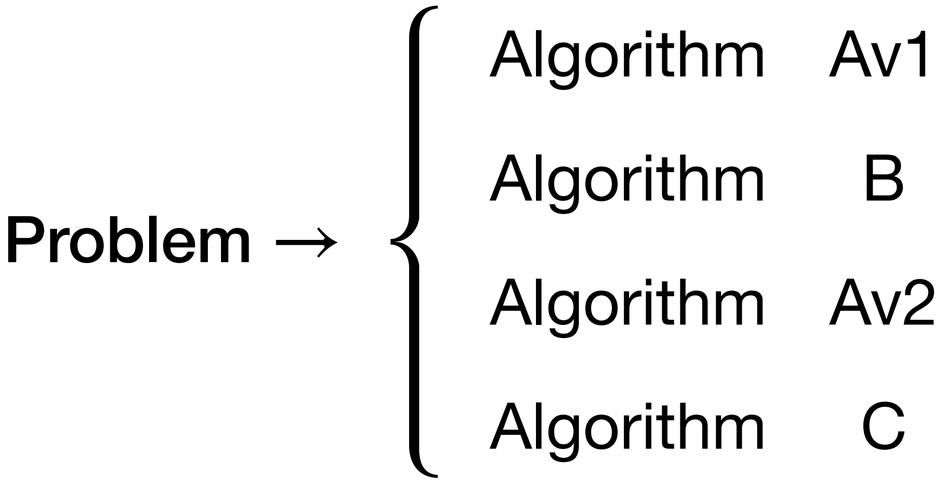
Problem →
Algorithm Av1
Algorithm B
Algorithm Av2
Algorithm C

Prior Frameworks

0 Gold

0 Silver

0 Bronze

Algorithm Parameter Instances Objective

Prior Frameworks

Algorithm

Parameter

Instances

1 Gold

0 Silver

0 Bronze

Objective

Prior Frameworks

Algorithm

Parameter

Instances

1 Gold

1 Silver

0 Bronze

Objective

Prior Frameworks

Algorithm

Parameter

Instances

1 Gold

0 Silver

0 Bronze

Objective

Prior Frameworks

Algorithm

Parameter

Instances

1 Gold

1 Silver

0 Bronze

Objective

Prior Frameworks

0 Gold

1 Silver

0 Bronze

Algorithm Parameter Instances Objective

Prior Frameworks

1 Gold
0 Silver
1 Bronze

Algorithm Parameter Instances Objective

Prior Frameworks

3 Gold

0 Silver

0 Bronze

Algorithm Parameter Instances Objective

A More Realistic Problem

How can we solve such problem instances?

A More Realistic Problem

How can we solve such problem instances?

A More Realistic Problem

How can we solve such problem instances?

A More Realistic Problem

How can we solve such problem instances?

A More Realistic Problem

How can we solve such problem instances?

A More Realistic Problem

How can we solve such problem instances?

How to Handle Non-Stationary Cases?

How to Handle Non-Stationary Cases?

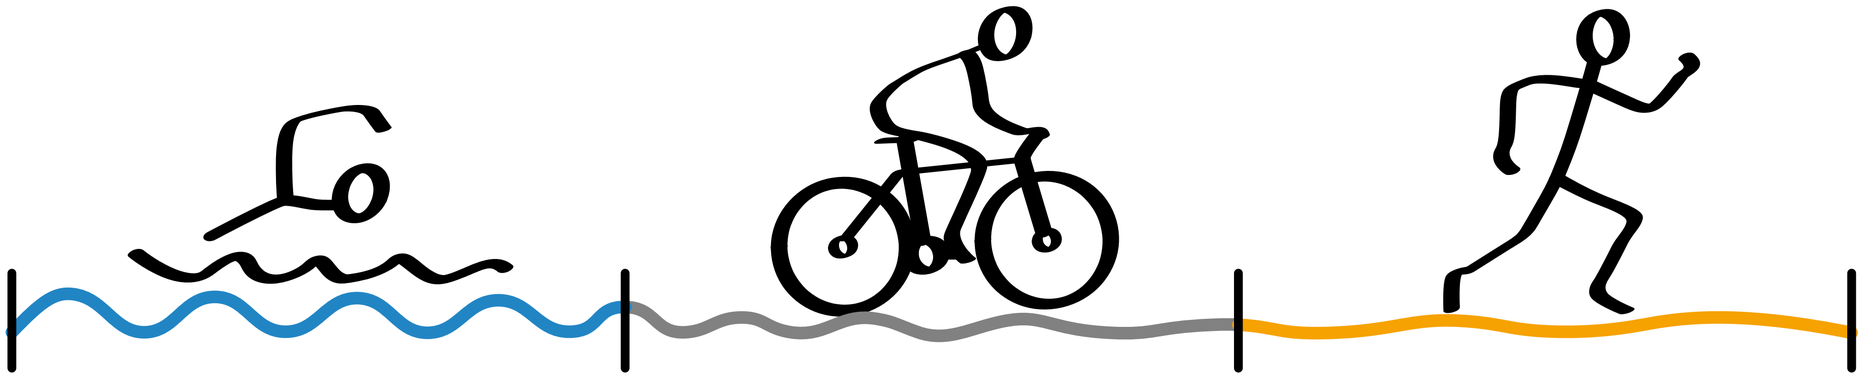

- Manually design heuristics
  - Requires substantial expert knowledge
    - Results in specialised heuristics (for single domains)
  - Very time-consuming

- Use Algorithm Configuration/Selection
  - Limited Approach
    - Does not make use of information available during the algorithms run

How to Handle Non-Stationary Cases?

- Learn new heuristics
  - Data driven approach
  - Loss of interpretability
    - Theoretical guarantees are lost
- Configure an algorithm at each execution step
  - Data driven approach
  - Makes use of existing heuristics
  - Requires access to an algorithms internal statistics

Dynamic Configuration - The Problem

Dynamic Configuration - The Problem

State Space

Dynamic Configuration - The Problem

State Space

Action Space

Dynamic Configuration - The Problem

State Space 

Action Space 

Transition Function **?**

Dynamic Configuration - The Problem

State Space 

Reward Function 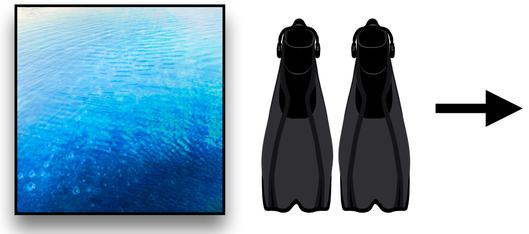

Action Space 

Transition Function **?**

Dynamic Configuration - The Problem

State Space 

Reward Function 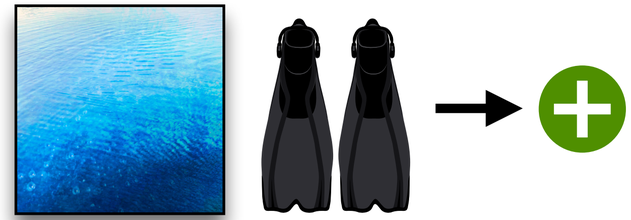

Action Space 

Transition Function **?**

Dynamic Configuration - The Problem

State Space 

Reward Function 

Action Space 

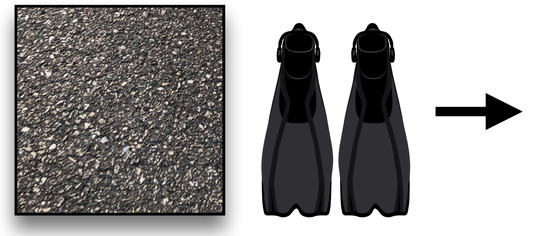

Transition Function **?**

Dynamic Configuration - The Problem

State Space

Action Space

Reward Function

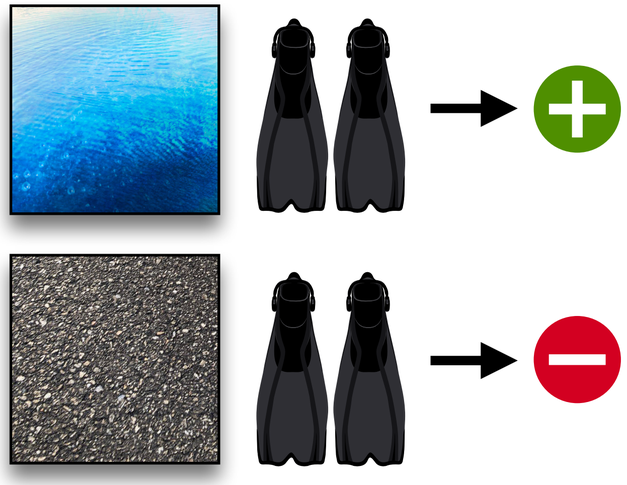

Transition Function **?**

Problem Instances Add Context

Instance B

Instance C

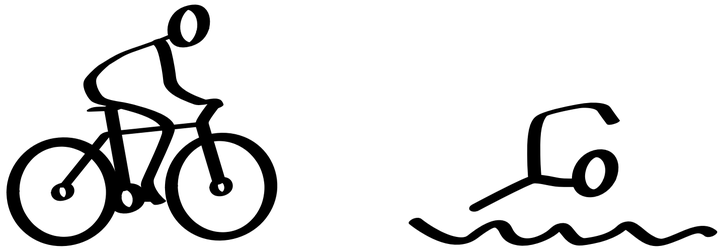

Instance D

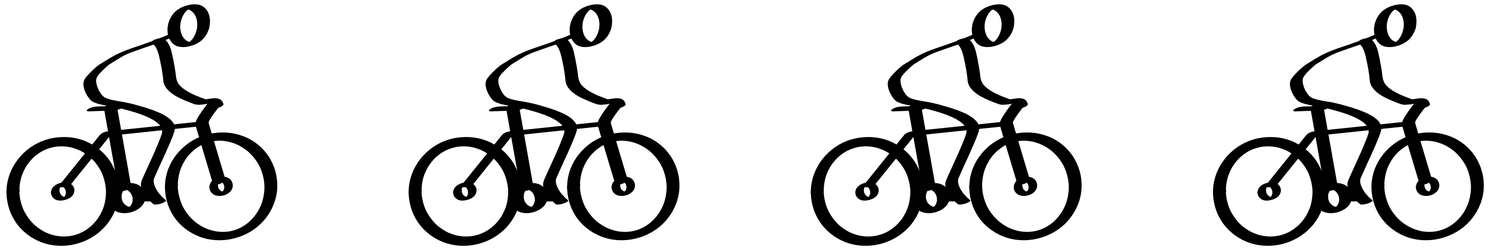

Solving the Problem by RL

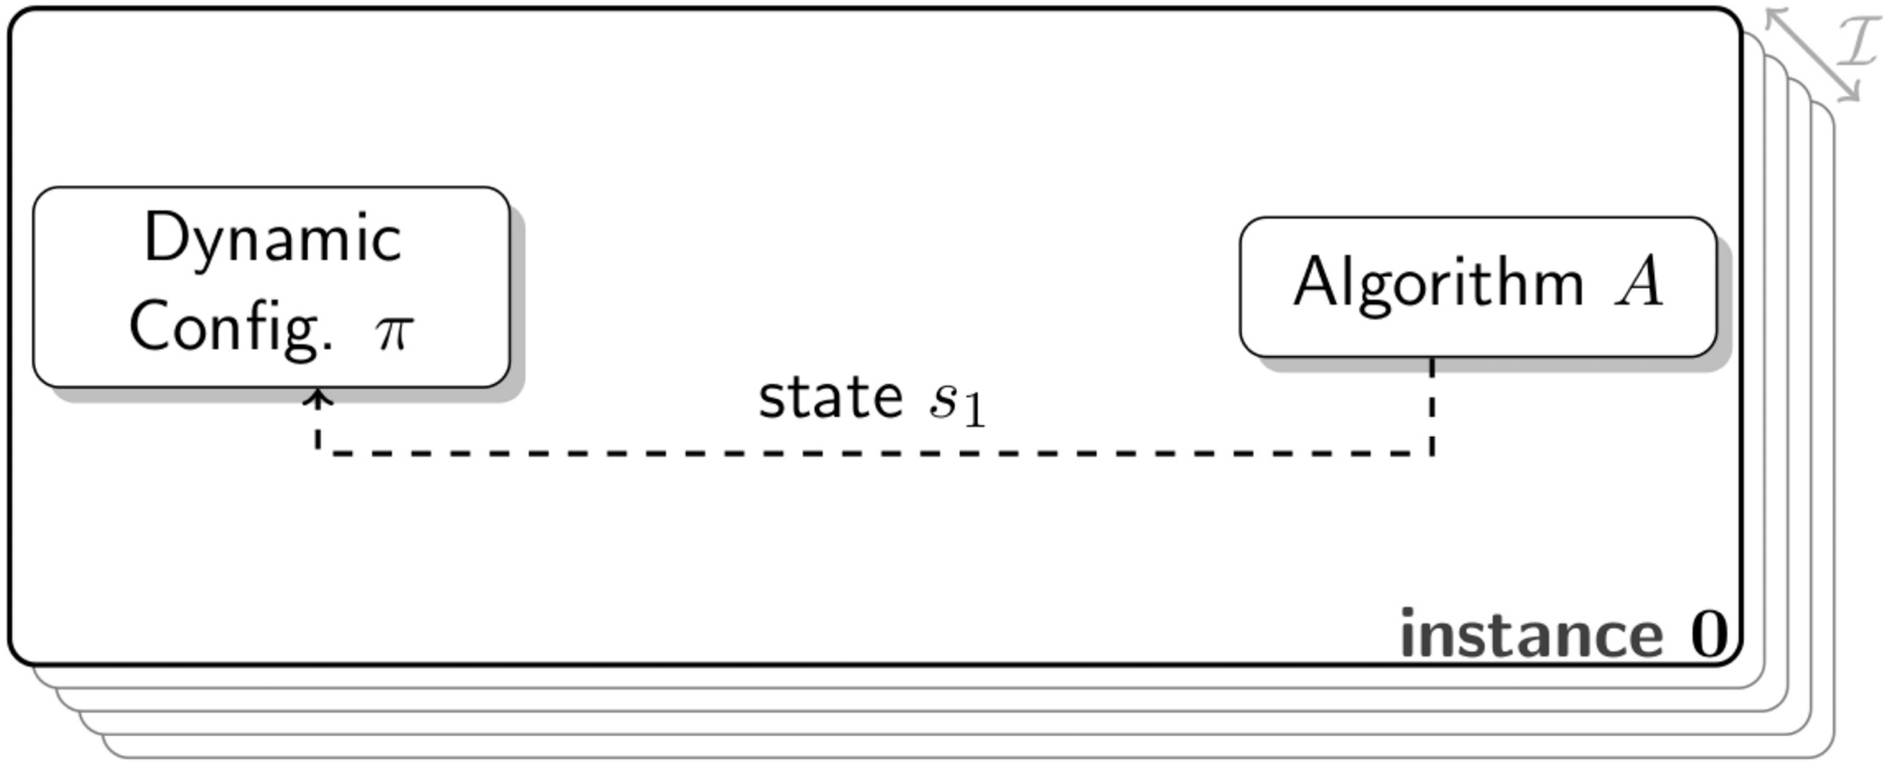

Dynamic Algorithm Configuration: Foundation of a New Meta-Algorithmic Framework
André Biedenkapp, H. Furkan Bozkurt, Theresa Eimer, Frank Hutter, and Marius Lindauer
In Proceedings of the Twenty-fourth European Conference on Artificial Intelligence (ECAI'20) 2020

Solving the Problem by RL

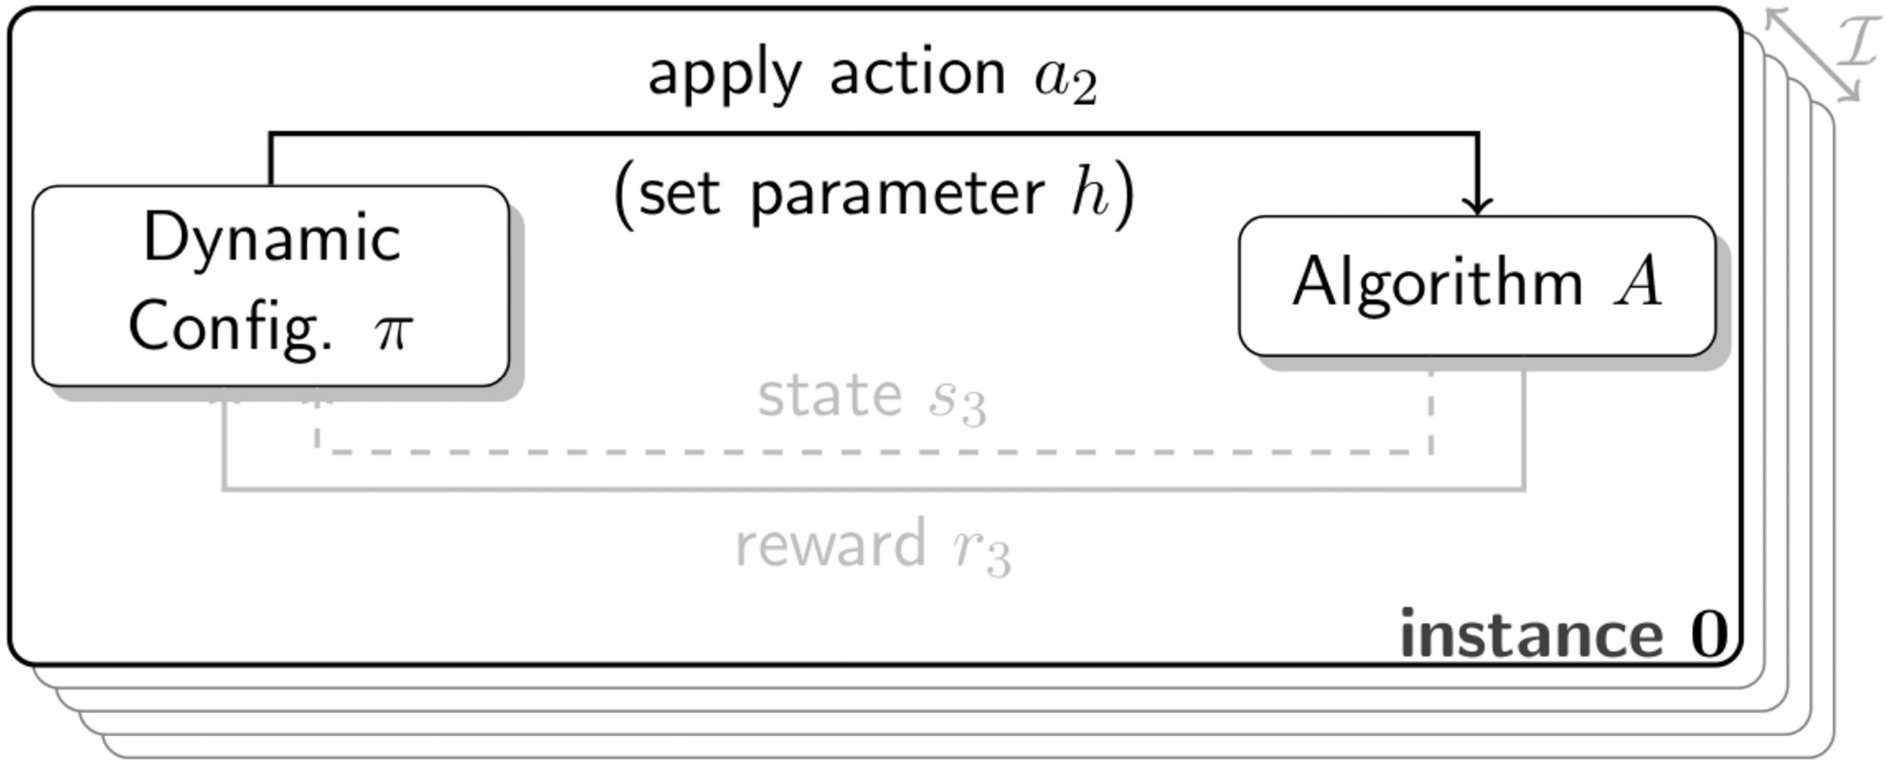

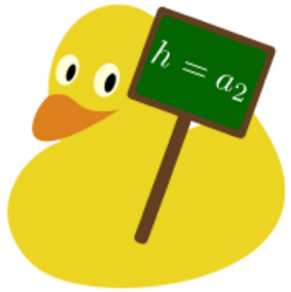

Dynamic Algorithm Configuration: Foundation of a New Meta-Algorithmic Framework
André Biedenkapp, H. Furkan Bozkurt, Theresa Eimer, Frank Hutter, and Marius Lindauer
In Proceedings of the Twenty-fourth European Conference on Artificial Intelligence (ECAI'20) 2020

Solving the Problem by RL

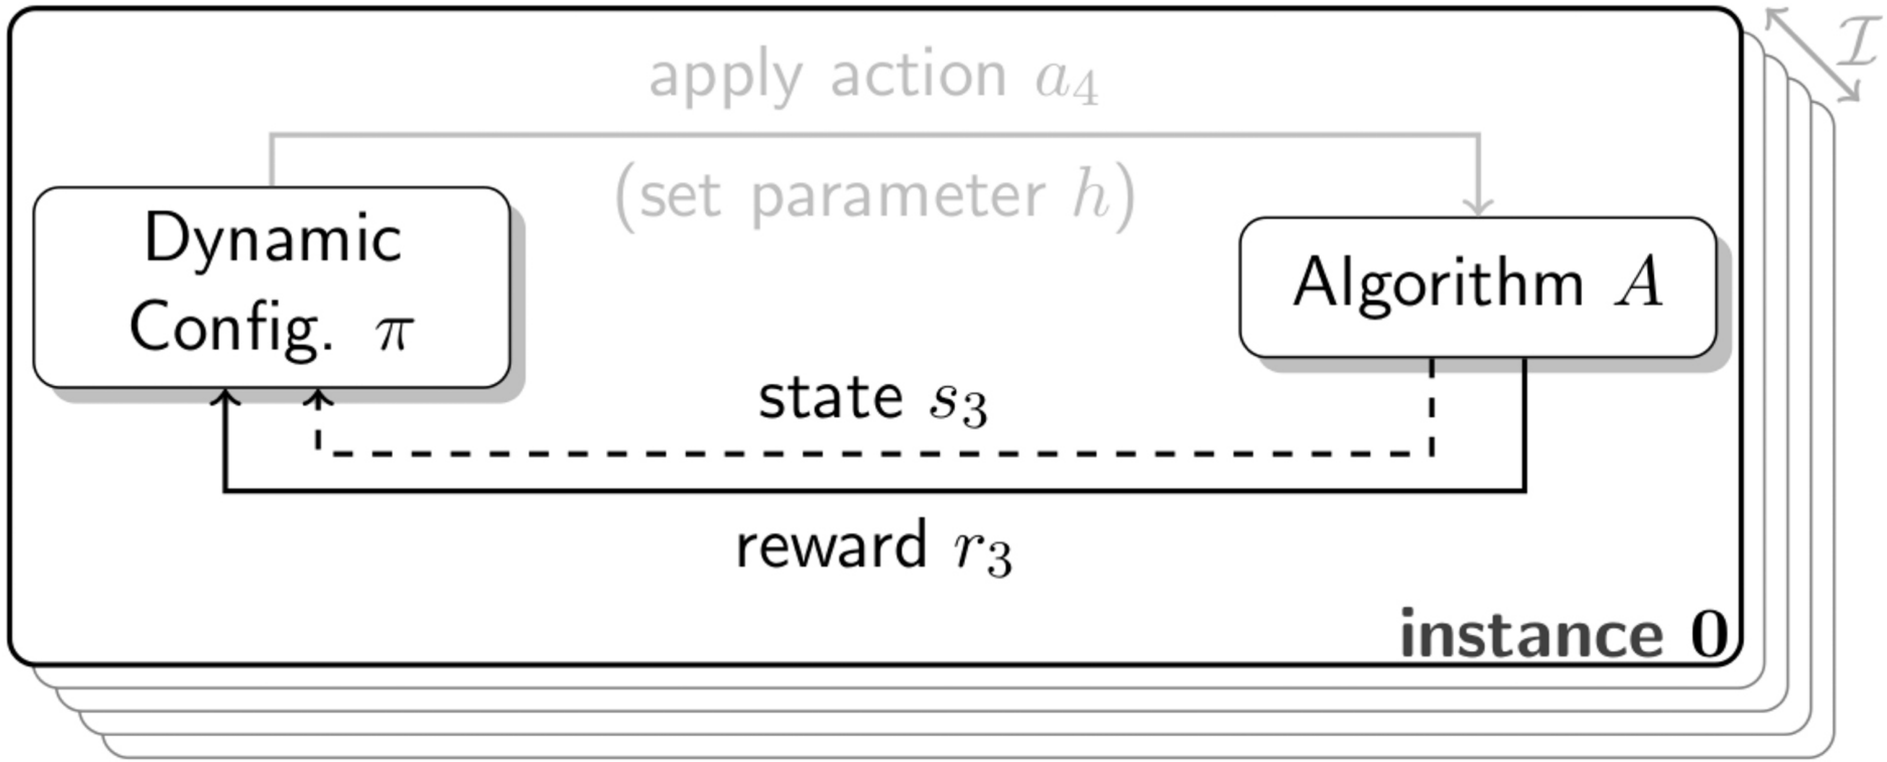

Dynamic Algorithm Configuration: Foundation of a New Meta-Algorithmic Framework
André Biedenkapp, H. Furkan Bozkurt, Theresa Eimer, Frank Hutter, and Marius Lindauer
In Proceedings of the Twenty-fourth European Conference on Artificial Intelligence (ECAI'20) 2020

Solving the Problem by RL

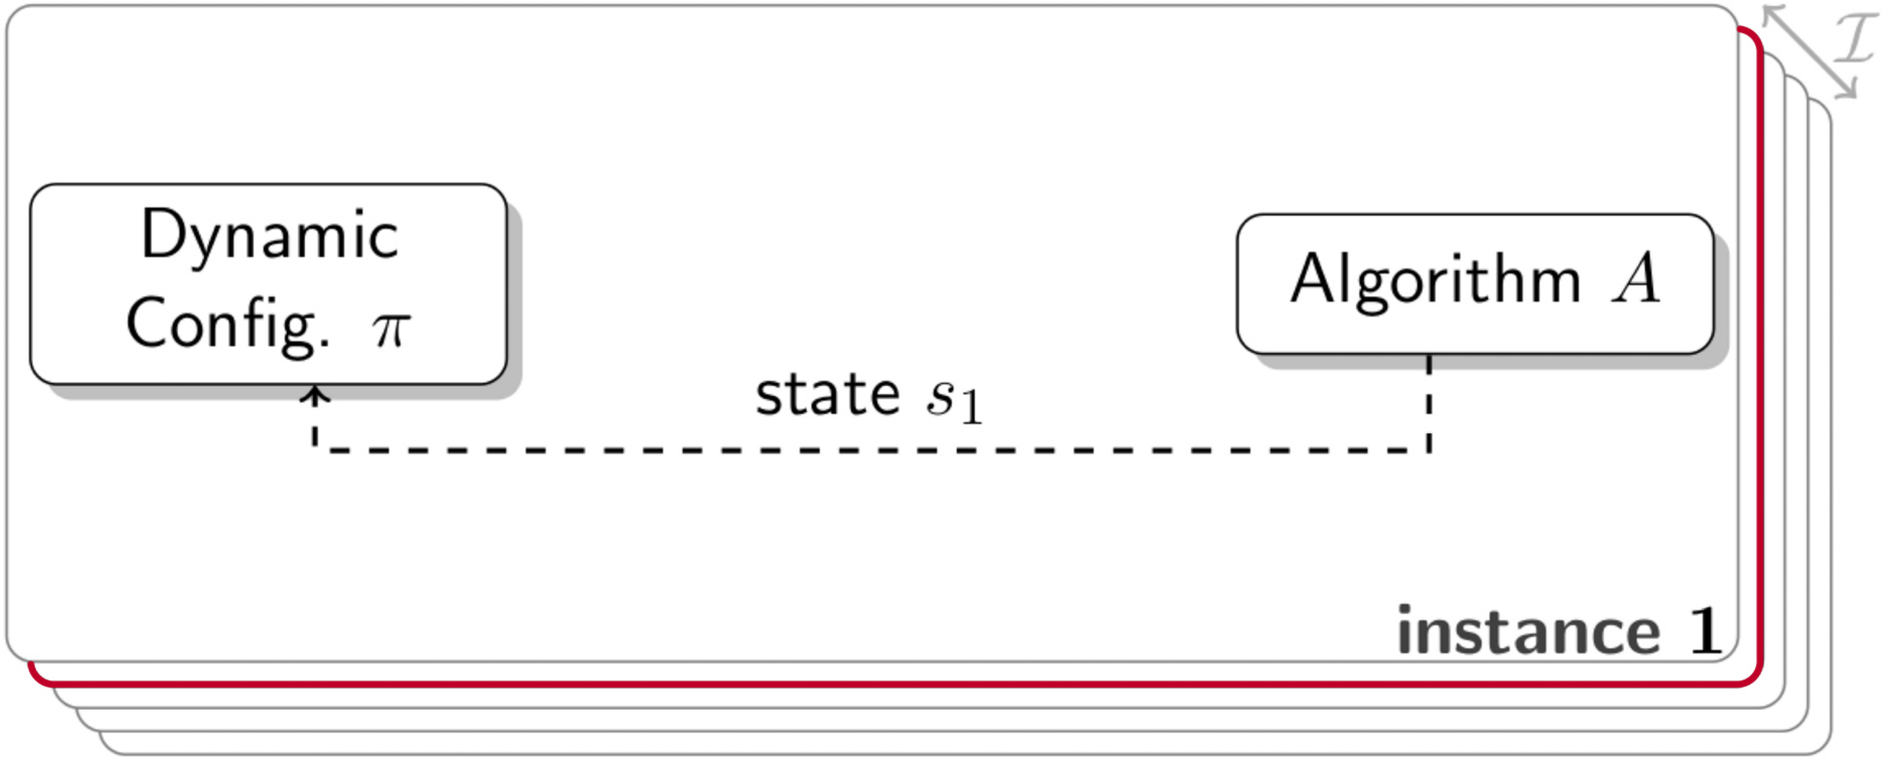

Dynamic Algorithm Configuration: Foundation of a New Meta-Algorithmic Framework
André Biedenkapp, H. Furkan Bozkurt, Theresa Eimer, Frank Hutter, and Marius Lindauer
In Proceedings of the Twenty-fourth European Conference on Artificial Intelligence (ECAI'20) 2020

Solving the Problem by RL

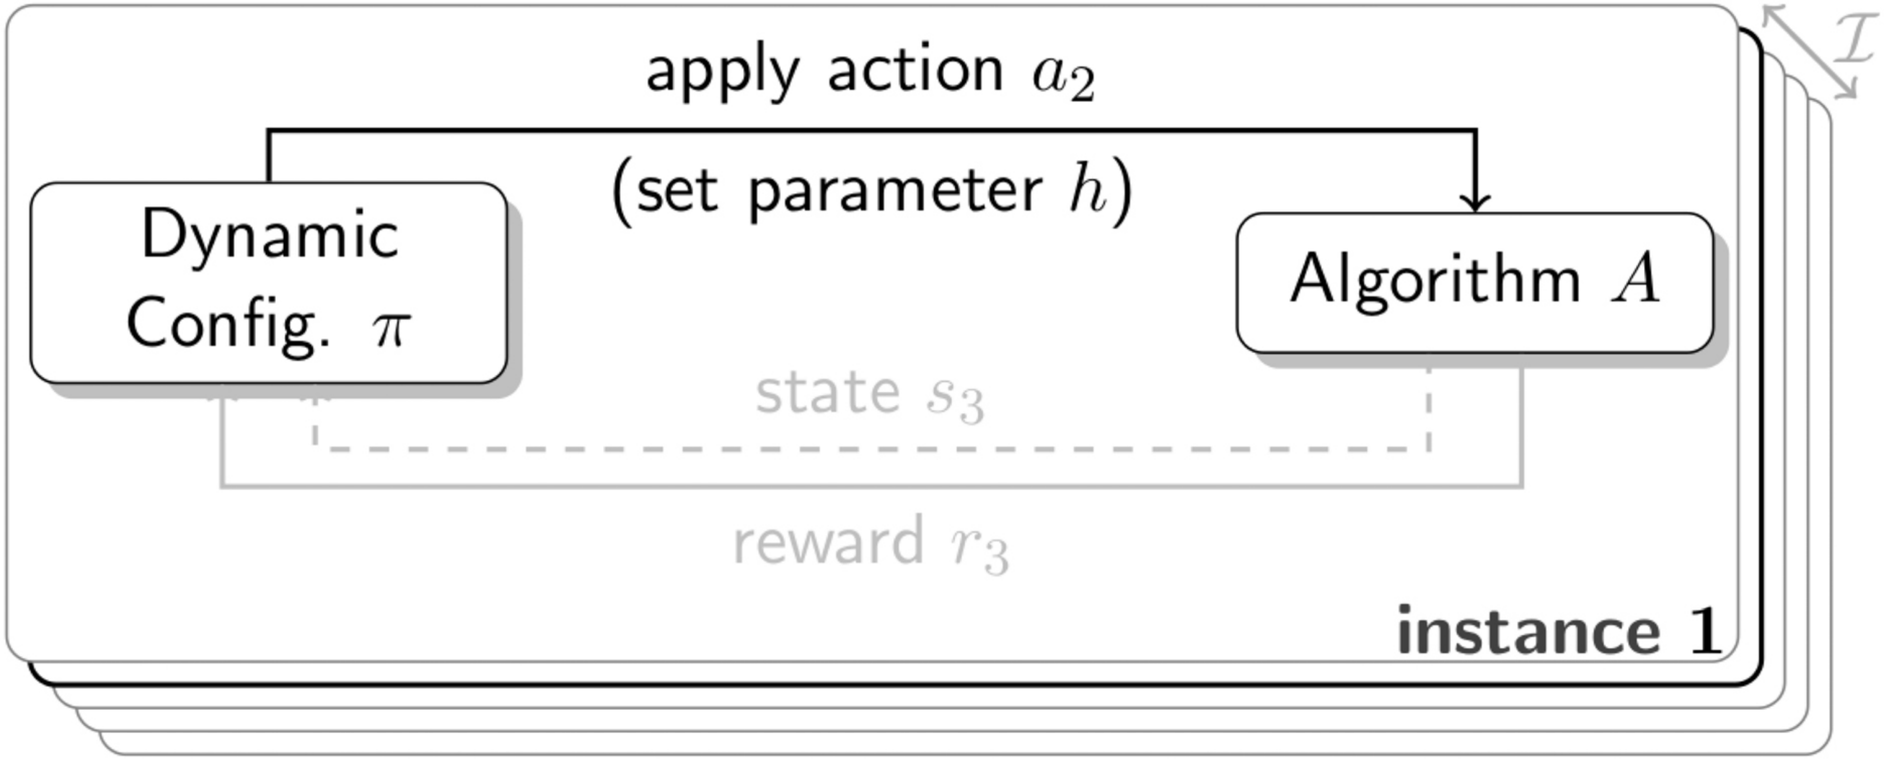

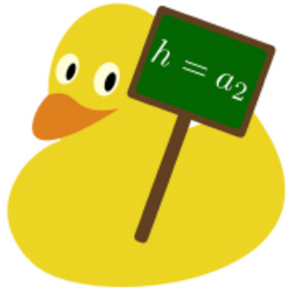

Dynamic Algorithm Configuration: Foundation of a New Meta-Algorithmic Framework
André Biedenkapp, H. Furkan Bozkurt, Theresa Eimer, Frank Hutter, and Marius Lindauer
In Proceedings of the Twenty-fourth European Conference on Artificial Intelligence (ECAI'20) 2020

Solving the Problem by RL

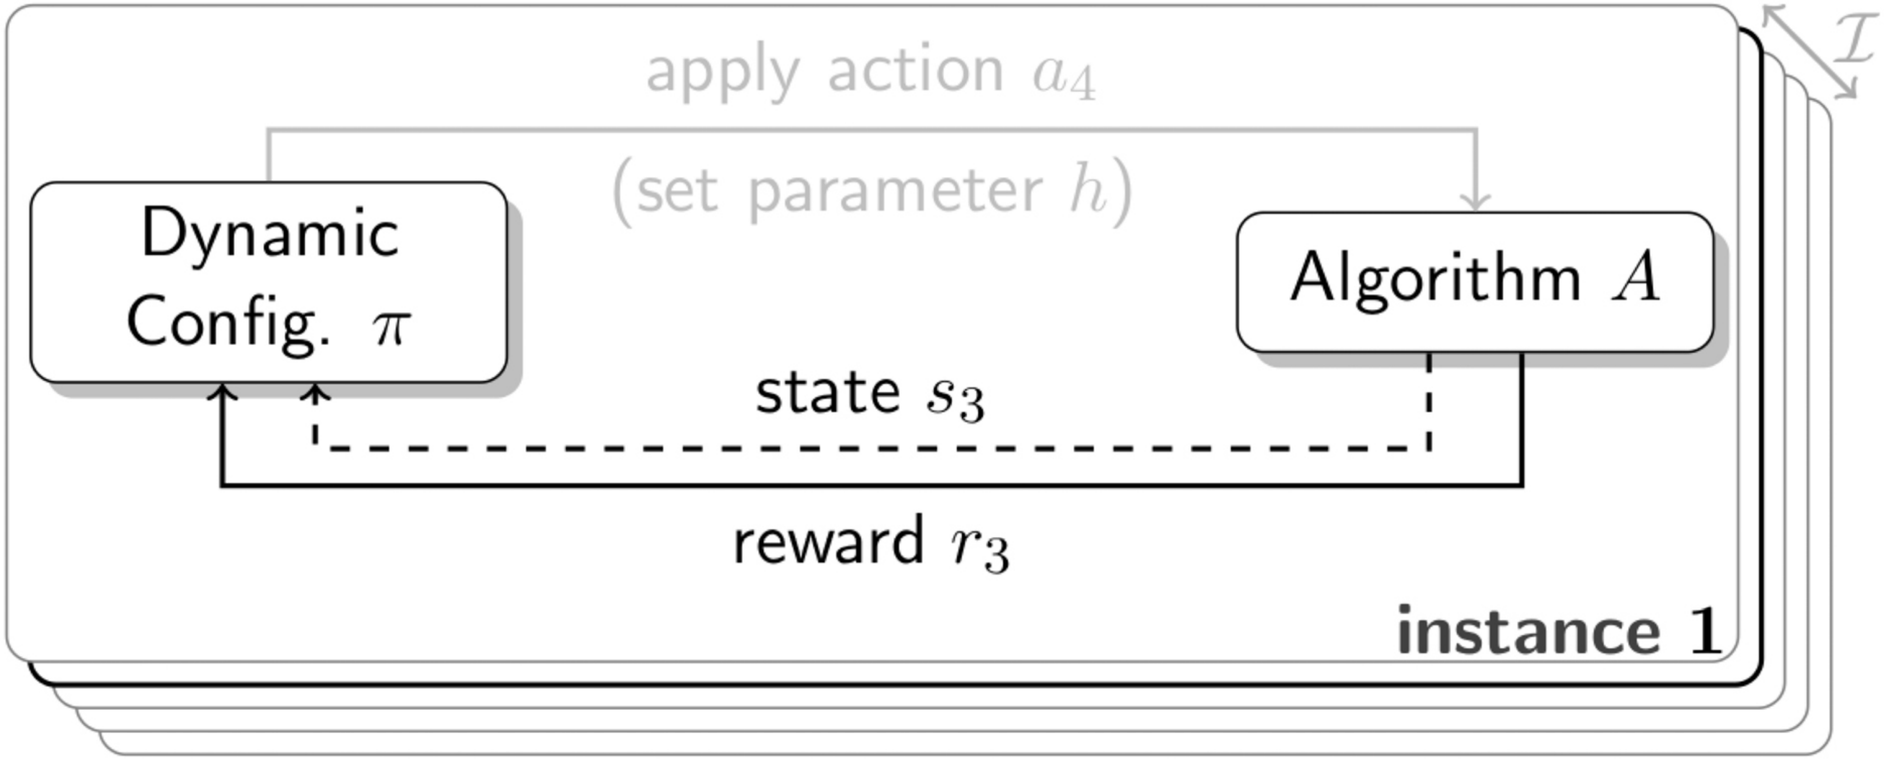

Dynamic Algorithm Configuration: Foundation of a New Meta-Algorithmic Framework
André Biedenkapp, H. Furkan Bozkurt, Theresa Eimer, Frank Hutter, and Marius Lindauer
In Proceedings of the Twenty-fourth European Conference on Artificial Intelligence (ECAI'20) 2020

Solving the Problem by RL

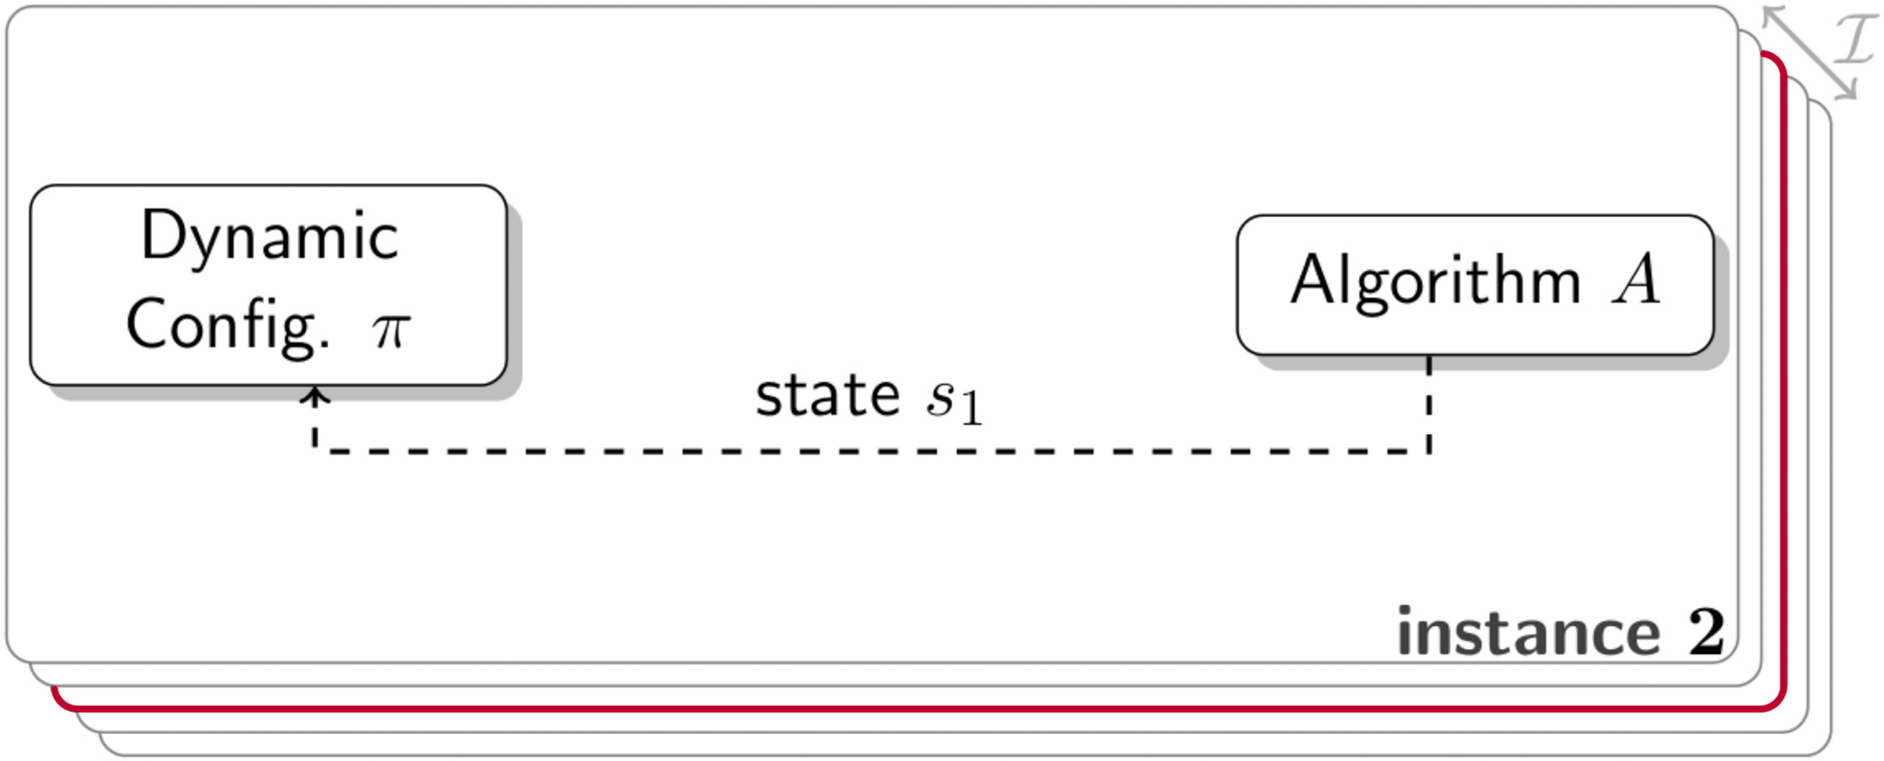

Dynamic Algorithm Configuration: Foundation of a New Meta-Algorithmic Framework
André Biedenkapp, H. Furkan Bozkurt, Theresa Eimer, Frank Hutter, and Marius Lindauer
In Proceedings of the Twenty-fourth European Conference on Artificial Intelligence (ECAI'20) 2020

Solving the Problem by RL

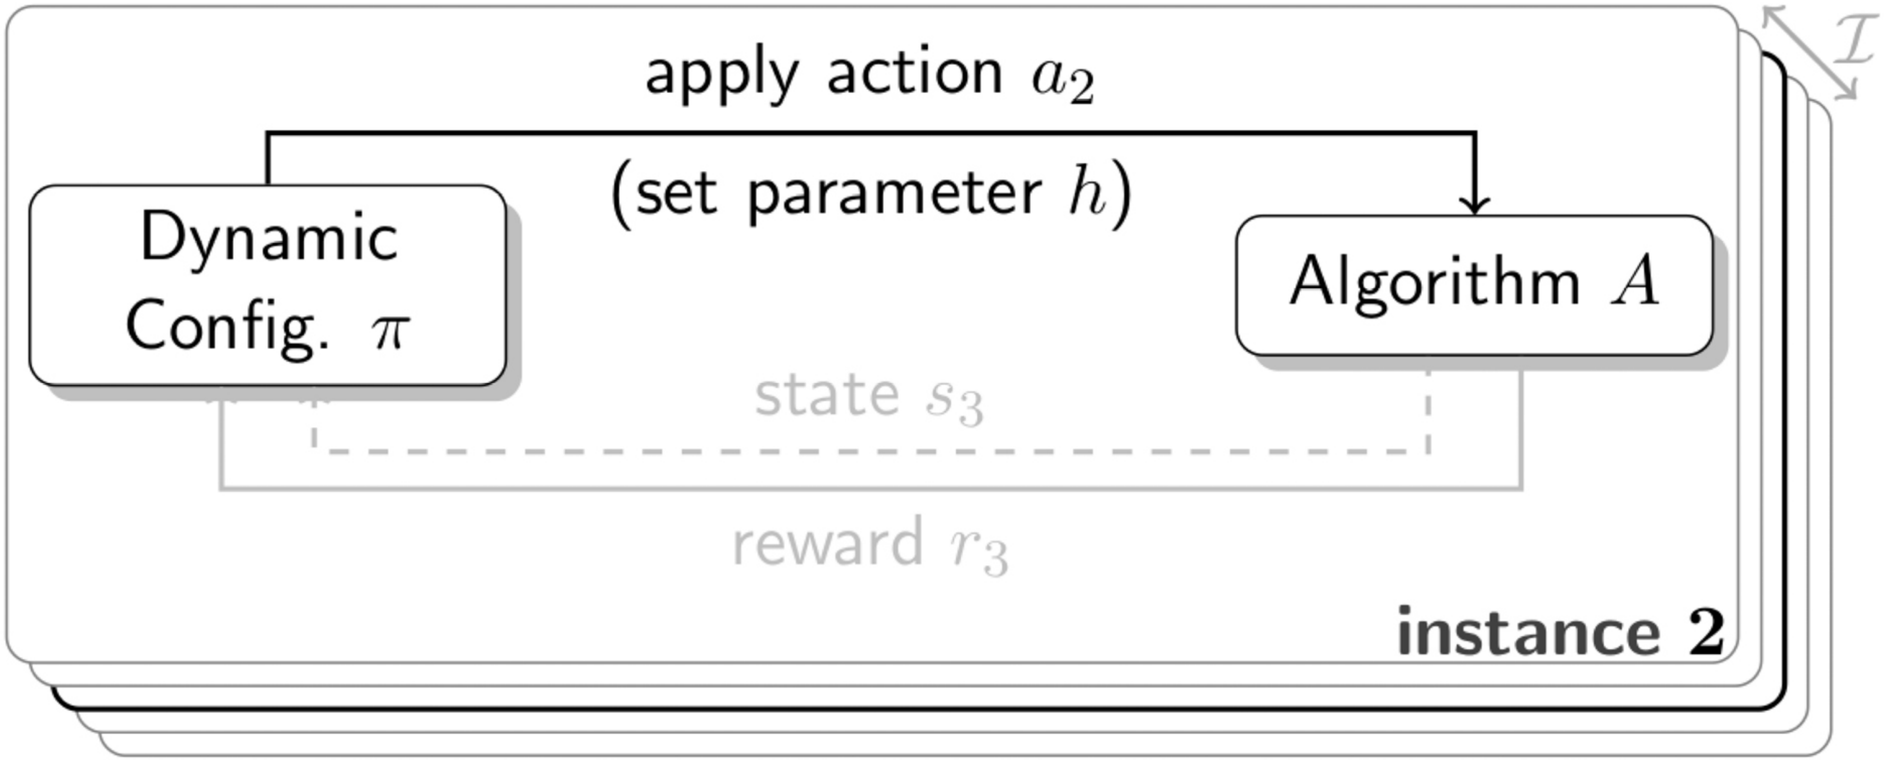

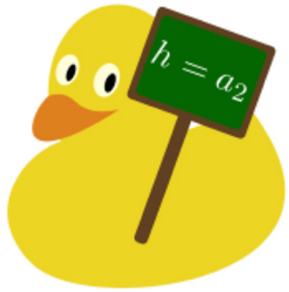

Dynamic Algorithm Configuration: Foundation of a New Meta-Algorithmic Framework
André Biedenkapp, H. Furkan Bozkurt, Theresa Eimer, Frank Hutter, and Marius Lindauer
In Proceedings of the Twenty-fourth European Conference on Artificial Intelligence (ECAI'20) 2020

Solving the Problem by RL

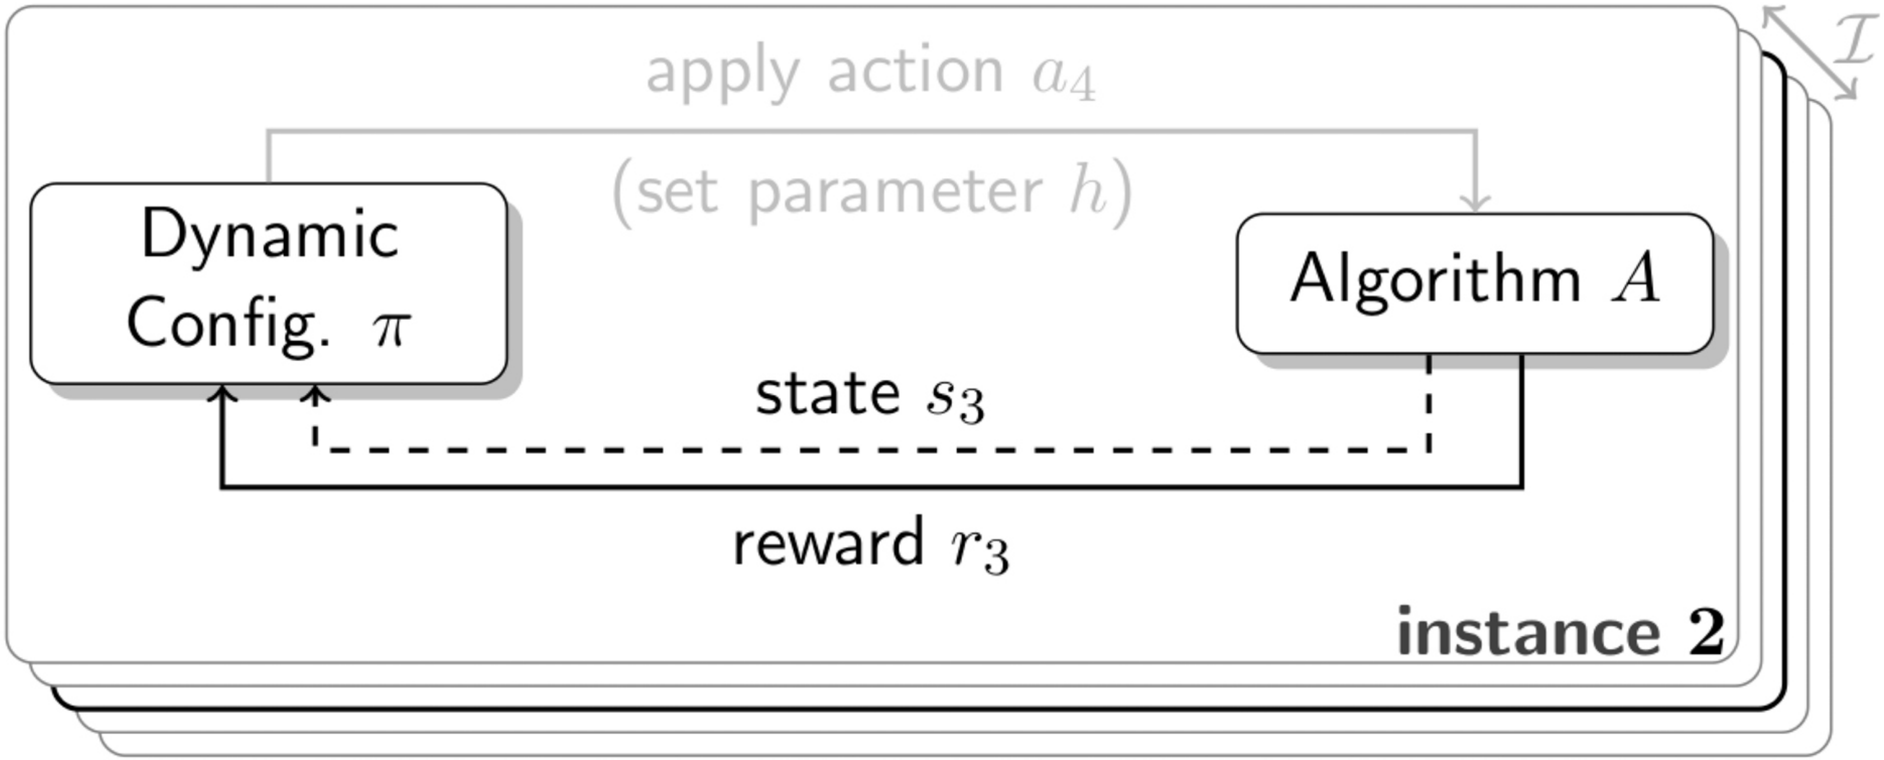

Dynamic Algorithm Configuration: Foundation of a New Meta-Algorithmic Framework
André Biedenkapp, H. Furkan Bozkurt, Theresa Eimer, Frank Hutter, and Marius Lindauer
In Proceedings of the Twenty-fourth European Conference on Artificial Intelligence (ECAI'20) 2020

Looking inside the Algorithm

Looking inside the Algorithm

Desiderata for state features:

Looking inside the Algorithm

Desiderata for state features:

- Cheap to compute

Looking inside the Algorithm

Desiderata for state features:

- Cheap to compute
- Quantify the progress of the algorithm

Looking inside the Algorithm

Desiderata for state features:

- Cheap to compute
- Quantify the progress of the algorithm
- Available at each decision point

Looking inside the Algorithm

Desiderata for state features:

- Cheap to compute
- Quantify the progress of the algorithm
- Available at each decision point

→ Make use of internal statistics

Looking inside the Algorithm

Looking inside the Algorithm

EAs

- population fitness
  [Sharma et al. 2019, Shala et al. 2020]
- stdev population fitness
  [Sharma et al. 2019]
- cumulative evolution path length
  [Shala et al. 2020]
- …

Looking inside the Algorithm

EAs

- population fitness
  [Sharma et al. 2019, Shala et al. 2020]
- stdev population fitness
  [Sharma et al. 2019]
- cumulative evolution path length
  [Shala et al. 2020]
- …

AI Planning

[Speck et al. 2021]

- average heuristic value
- minimal heuristic value
- #possible next planning states
- …

Looking inside the Algorithm

EAs

- population fitness
  [Sharma et al. 2019, Shala et al. 2020]
- stdev population fitness
  [Sharma et al. 2019]
- cumulative evolution path length
  [Shala et al. 2020]
- …

AI Planning

[Speck et al. 2021]

- average heuristic value
- minimal heuristic value
- #possible next planning states
- …

NN Optimization

[Daniel et al. 2016]

- predictive change in function value
- disagreement of function values
- uncertainty
- ...

Looking inside the Algorithm

EAs

- population fitness
  [Sharma et al. 2019, Shala et al. 2020]
- stdev population fitness
  [Sharma et al. 2019]
- cumulative evolution path length
  [Shala et al. 2020]
- …

AI Planning

[Speck et al. 2021]

- average heuristic value
- minimal heuristic value
- #possible next planning states
- …

NN Optimization

[Daniel et al. 2016]

- predictive change in function value
- disagreement of function values
- uncertainty
- ...

Learning Heuristic Selection with Dynamic Algorithm Configuration
David Speck*, André Biedenkapp*, Frank Hutter, Robert Mattmüller, and Marius Lindauer
In Proceedings of the Thirty-First International Conference on Automated Planning and Scheduling (ICAPS'21) 2021

AI Planning

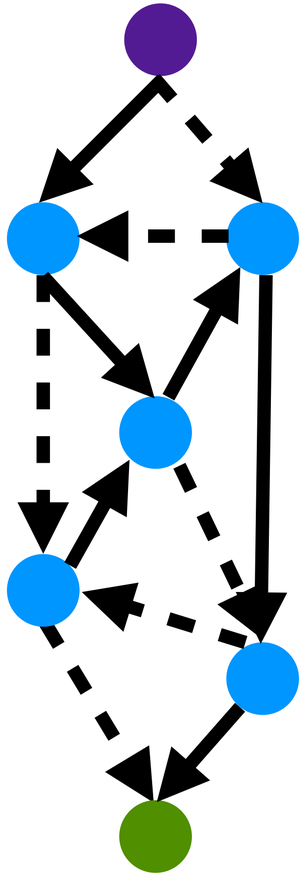

h1: Dashed is faster

h2: Solid is faster

AI Planning

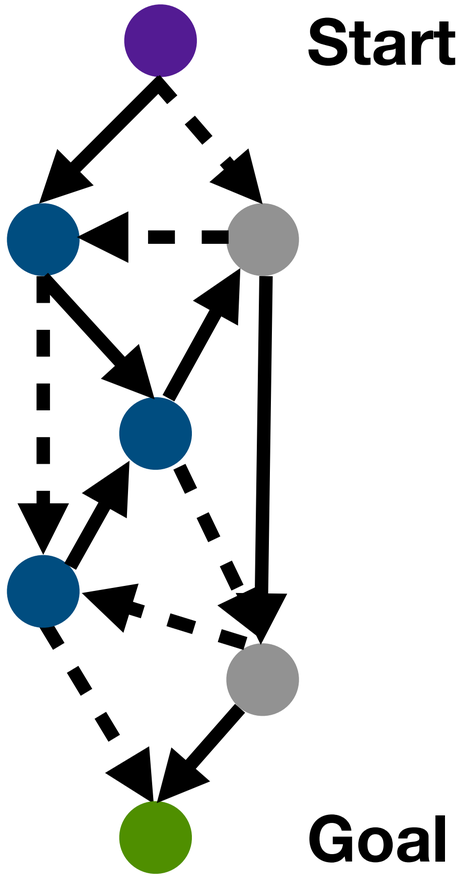

h1: Dashed is faster

h2: Solid is faster

AI Planning

Algorithm	CONTROL POLICY		SINGLE HEURISTIC	
Domain (# Inst.)	**RL**	RND	h_{ff}	h_{cg}

AI Planning

Algorithm	CONTROL POLICY		SINGLE HEURISTIC	
Domain (# Inst.)	**RL**	RND	h_{ff}	h_{cg}
SOKOBAN (100)	87.7	87.1	88.0	**90.0**
VISITALL (100)	56.9	51.0	37.0	**60.0**

AI Planning

Algorithm	CONTROL POLICY		SINGLE HEURISTIC	
Domain (# Inst.)	**RL**	RND	h_{ff}	h_{cg}
SOKOBAN (100)	87.7	87.1	88.0	**90.0**
VISITALL (100)	56.9	51.0	37.0	**60.0**
ROVERS (100)	95.2	**96.0**	84.0	72.0

AI Planning

Algorithm	CONTROL POLICY		SINGLE HEURISTIC	
Domain (# Inst.)	**RL**	RND	h_{ff}	h_{cg}
SOKOBAN (100)	87.7	87.1	88.0	**90.0**
VISITALL (100)	56.9	51.0	37.0	**60.0**
ROVERS (100)	95.2	**96.0**	84.0	72.0
BARMAN (100)	**84.4**	83.8	66.0	17.0
BLOCKSWORLD (100)	**92.9**	83.6	75.0	60.0
CHILDSNACK (100)	**88.0**	86.2	75.0	86.0

AI Planning

Algorithm	CONTROL POLICY		SINGLE HEURISTIC	
Domain (# Inst.)	**RL**	RND	h_{ff}	h_{cg}
SOKOBAN (100)	87.7	87.1	88.0	**90.0**
VISITALL (100)	56.9	51.0	37.0	**60.0**
ROVERS (100)	95.2	**96.0**	84.0	72.0
BARMAN (100)	**84.4**	83.8	66.0	17.0
BLOCKSWORLD (100)	**92.9**	83.6	75.0	60.0
CHILDSNACK (100)	**88.0**	86.2	75.0	86.0
SUM (600)	**505.1**	487.7	425.0	385.0

	Algorithm	
Domain (# Inst.)	**RL**	Static Oracle
SUM (600)	**505.1**	489.0

AI Planning

	Algorithm	
Domain (# Inst.)	**RL**	Static Oracle
SUM (600)	**505.1**	489.0

Theoretical Static Optimum
unbeatable with prior art

Struggled to learn simple static policies

TempoRL

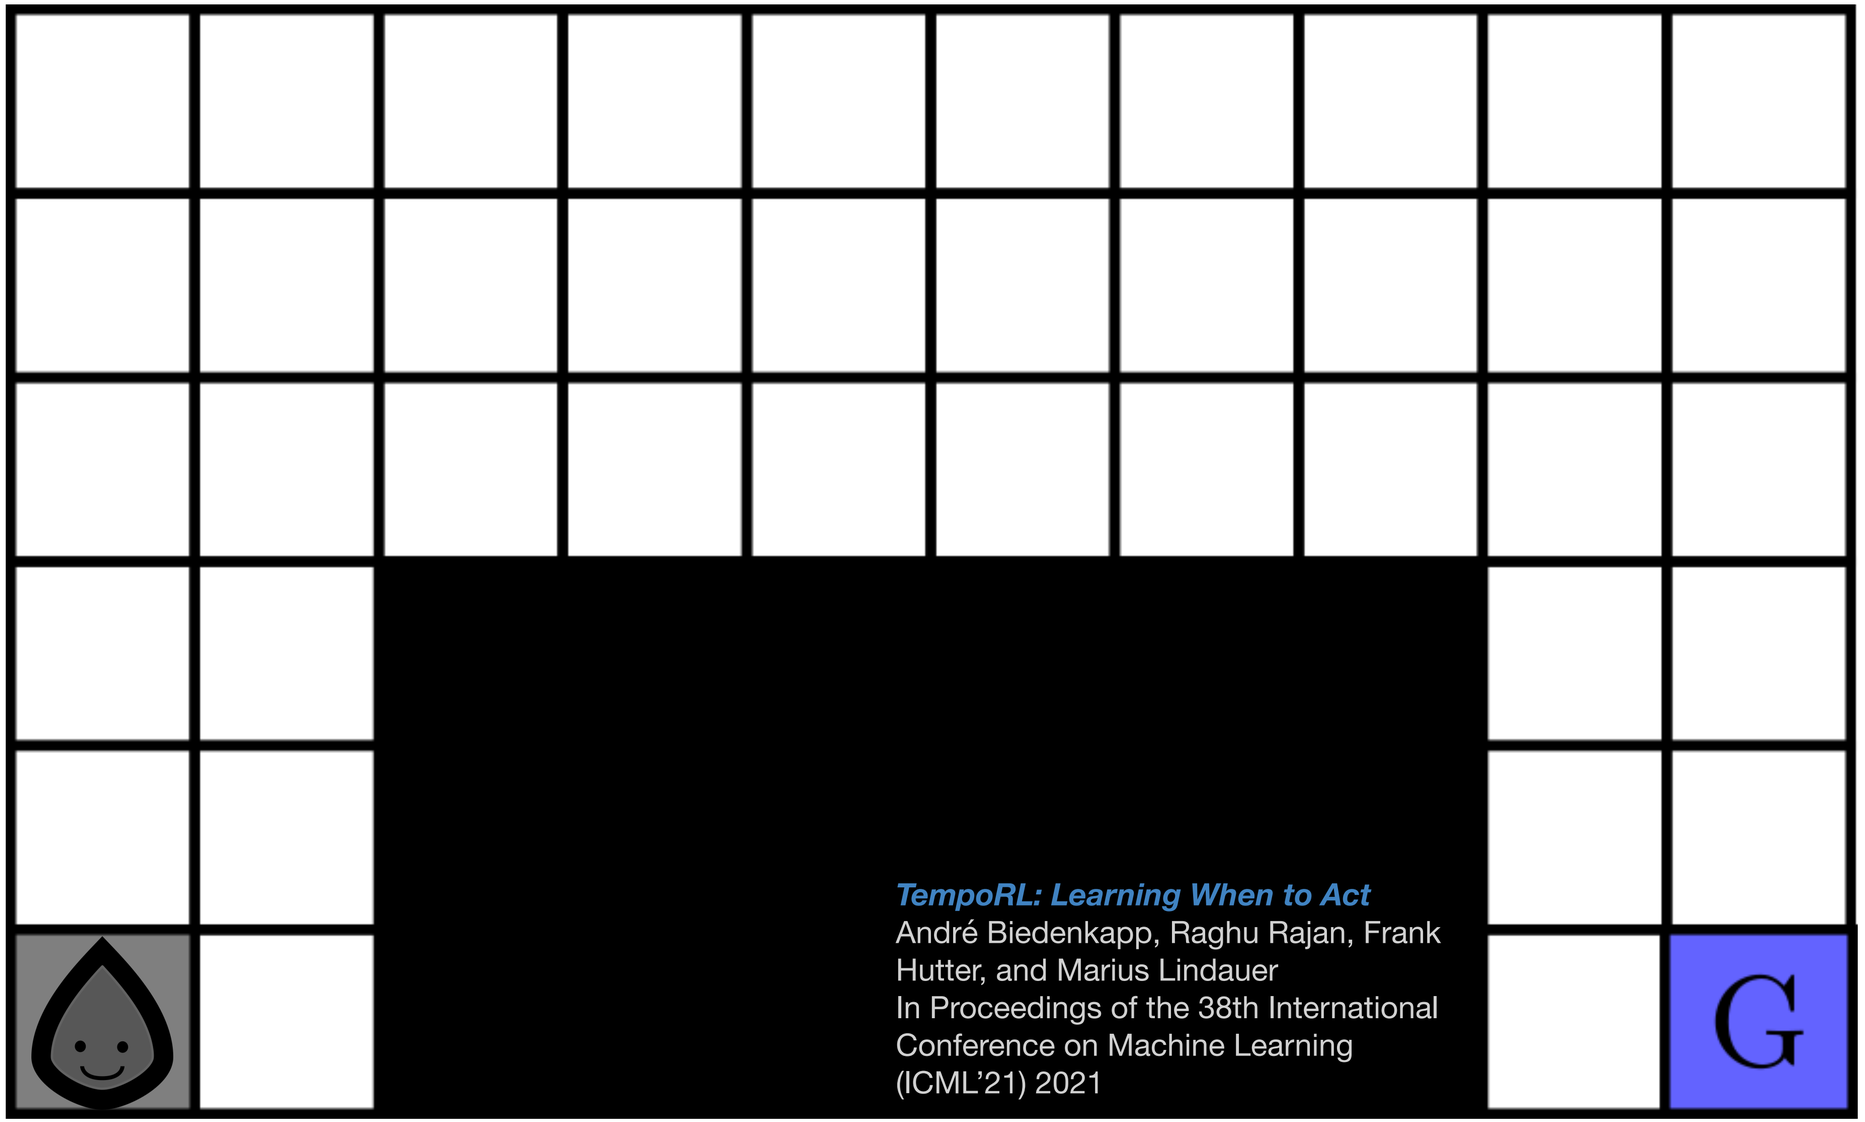

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

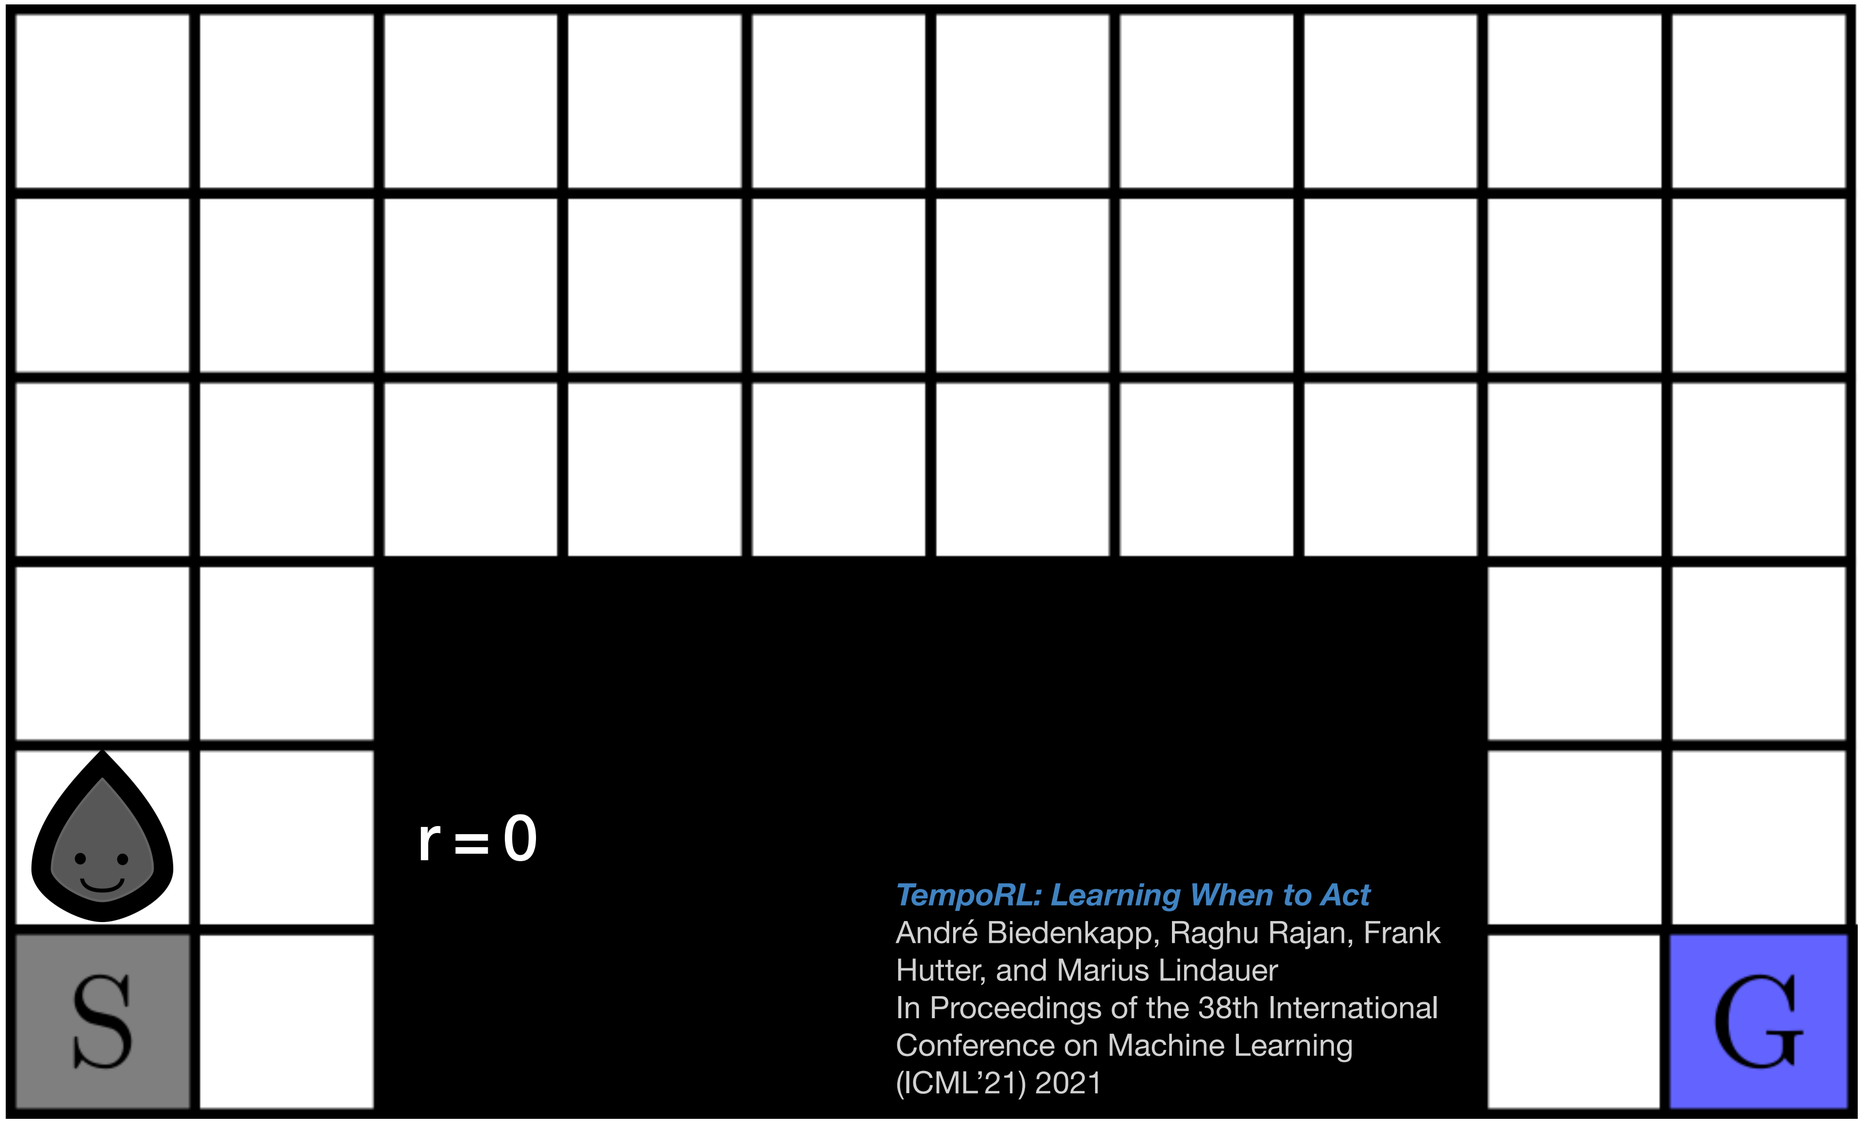
r = 0
S
G
TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

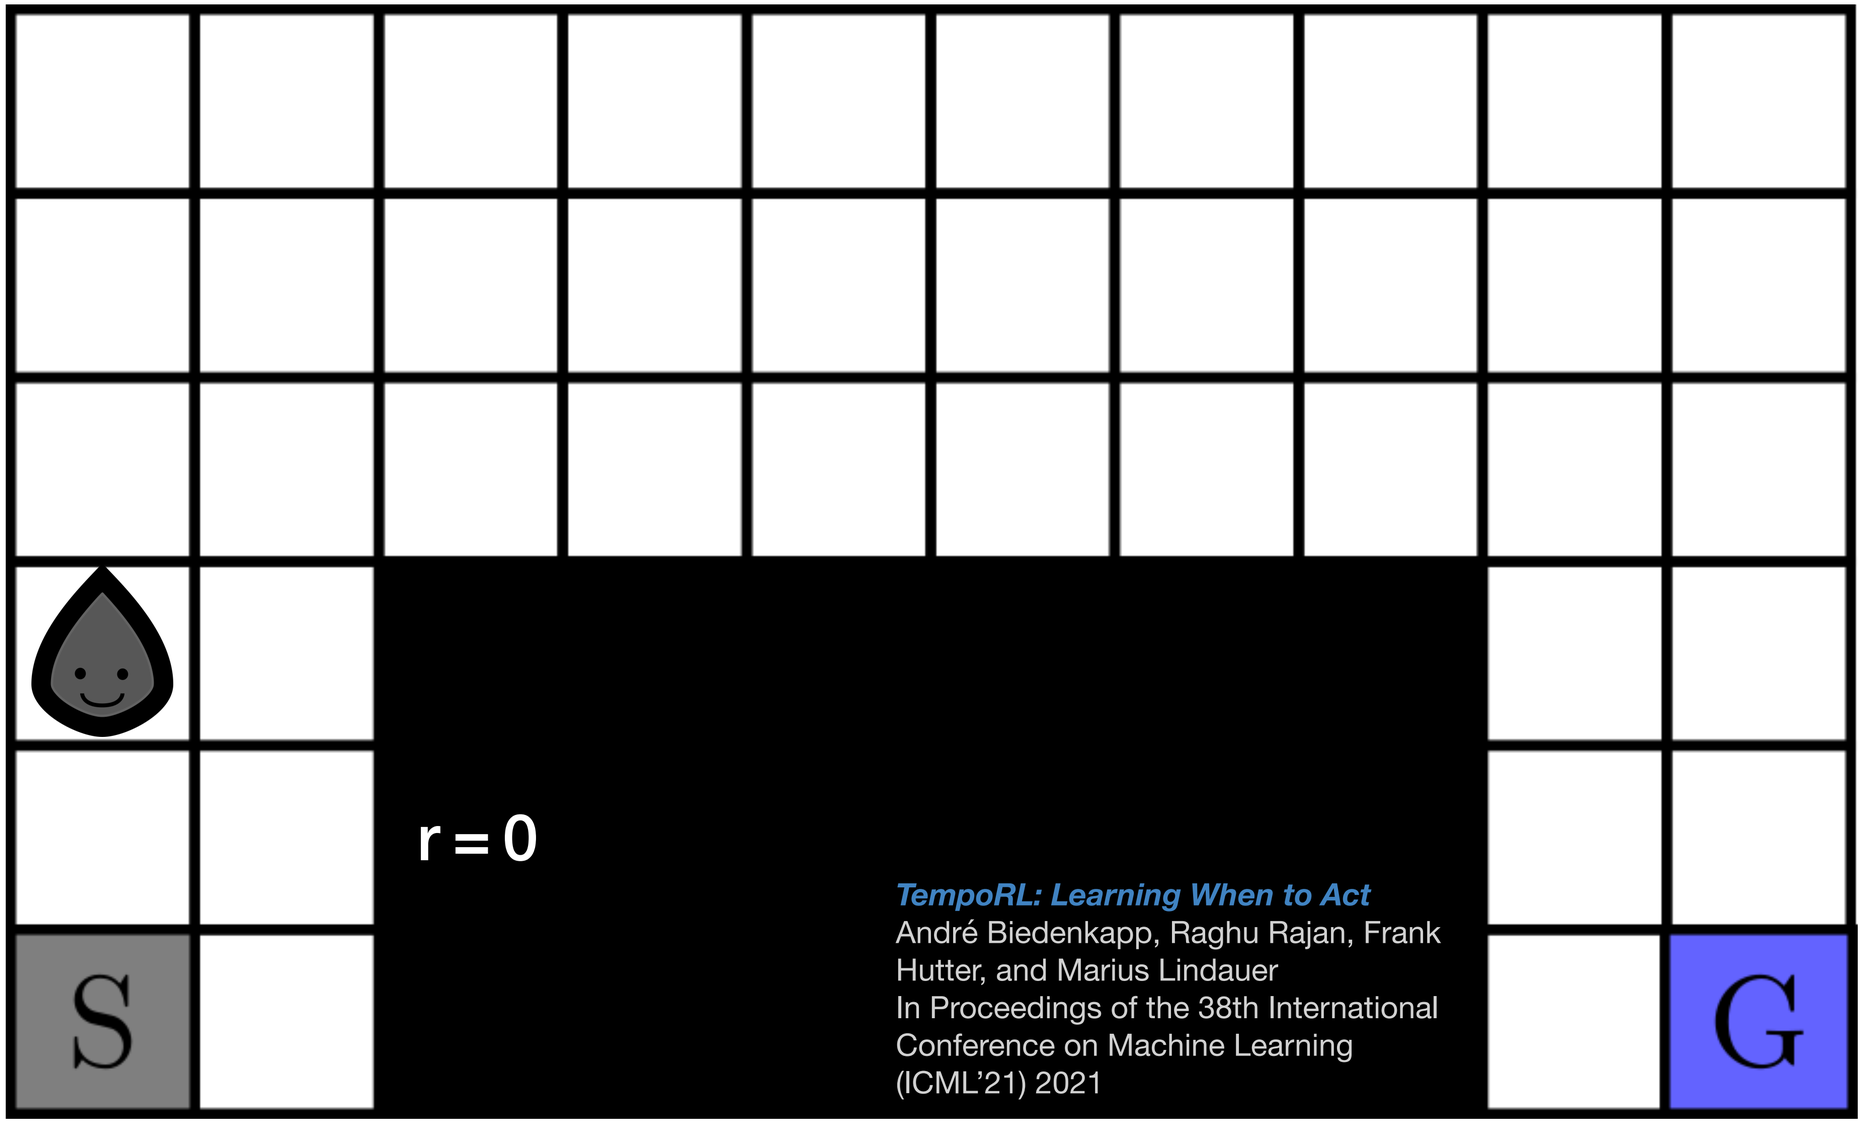
r = 0
S
G
TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

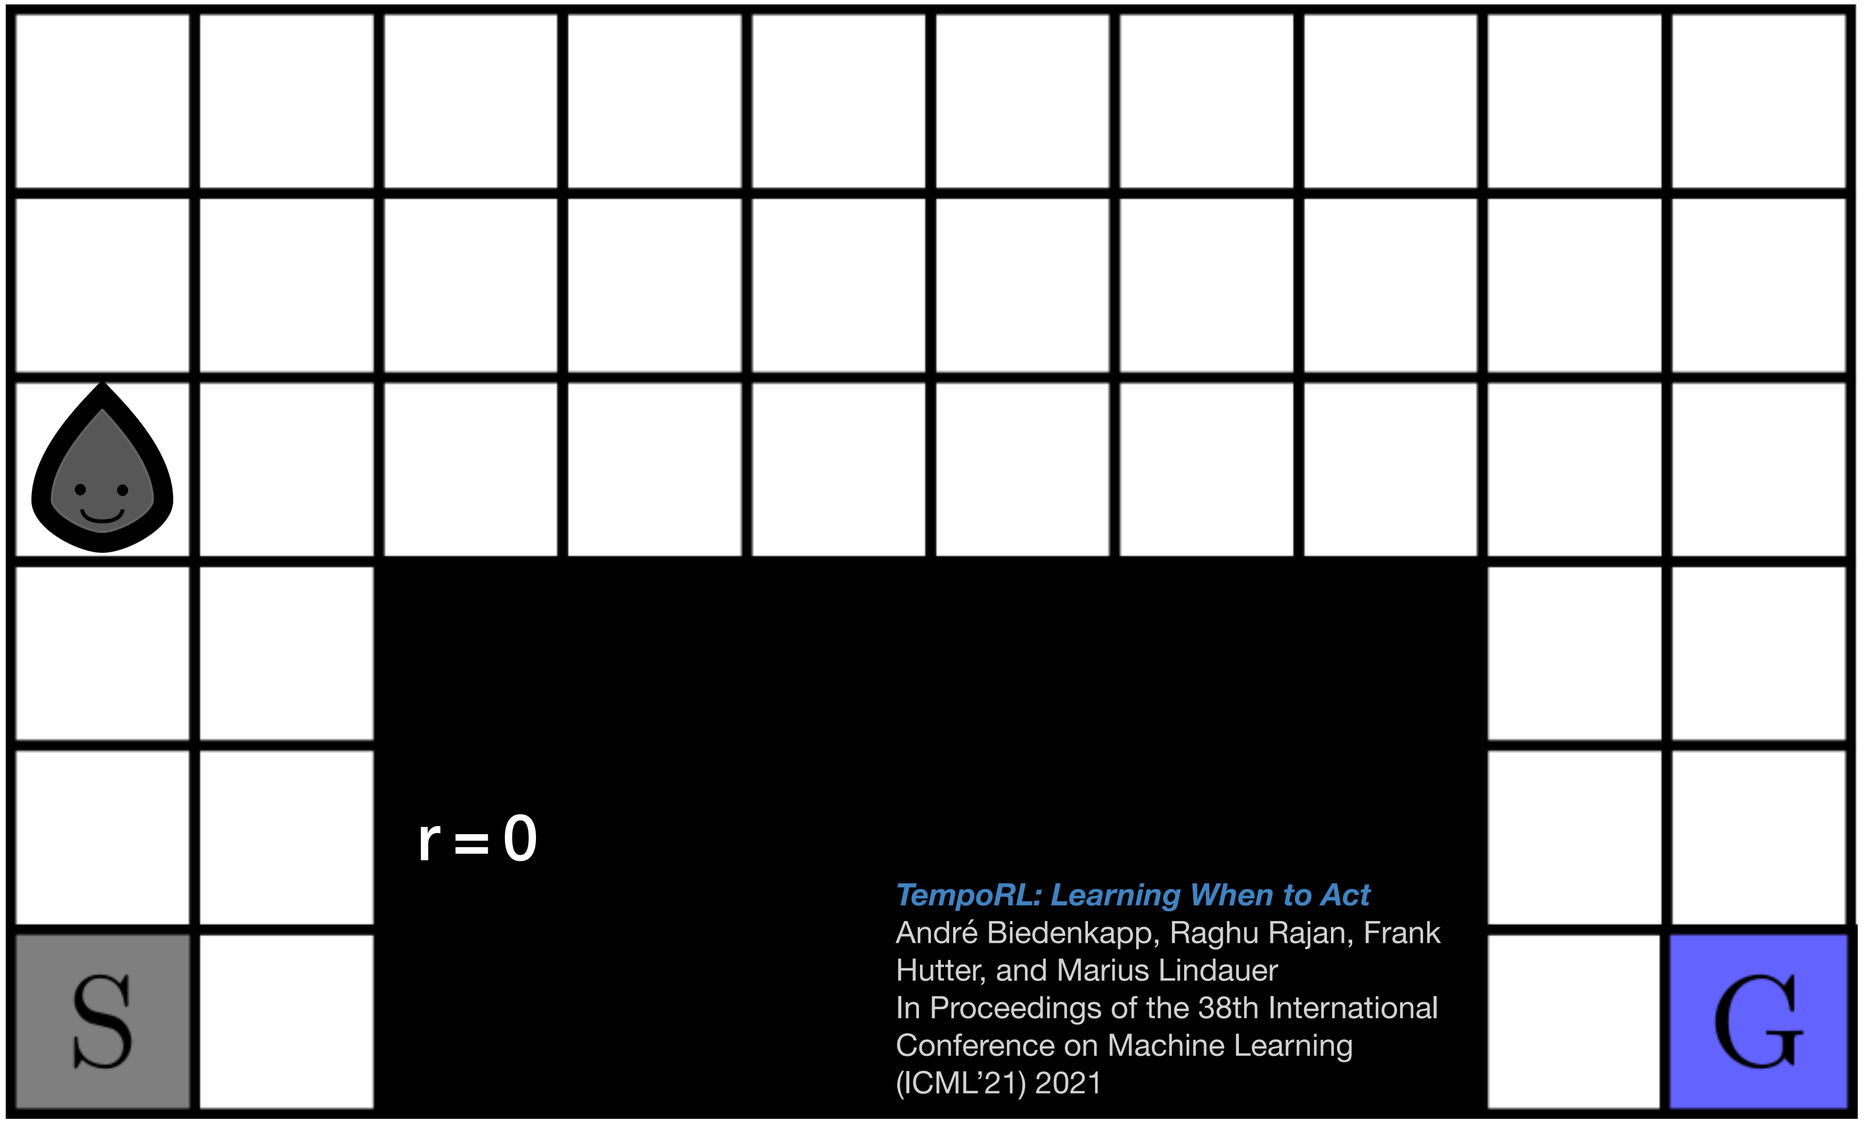

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

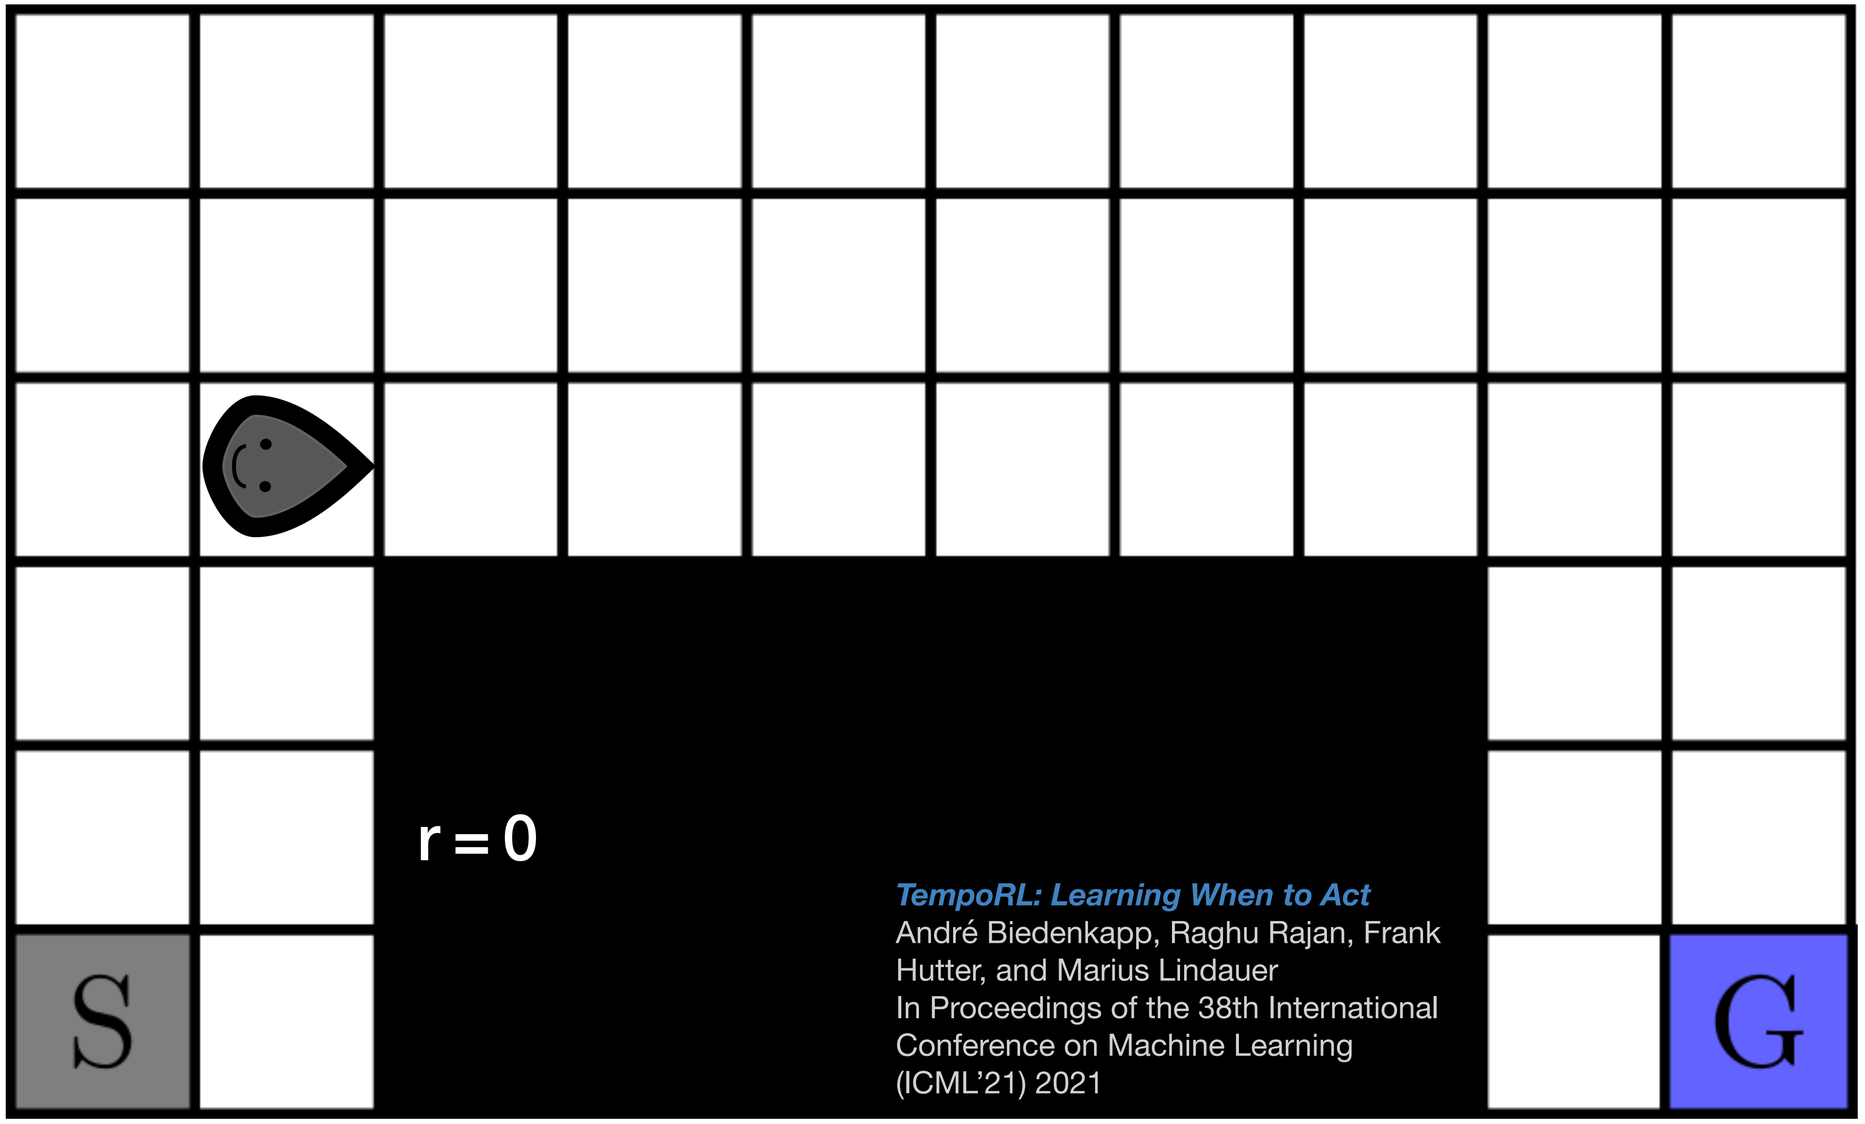

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

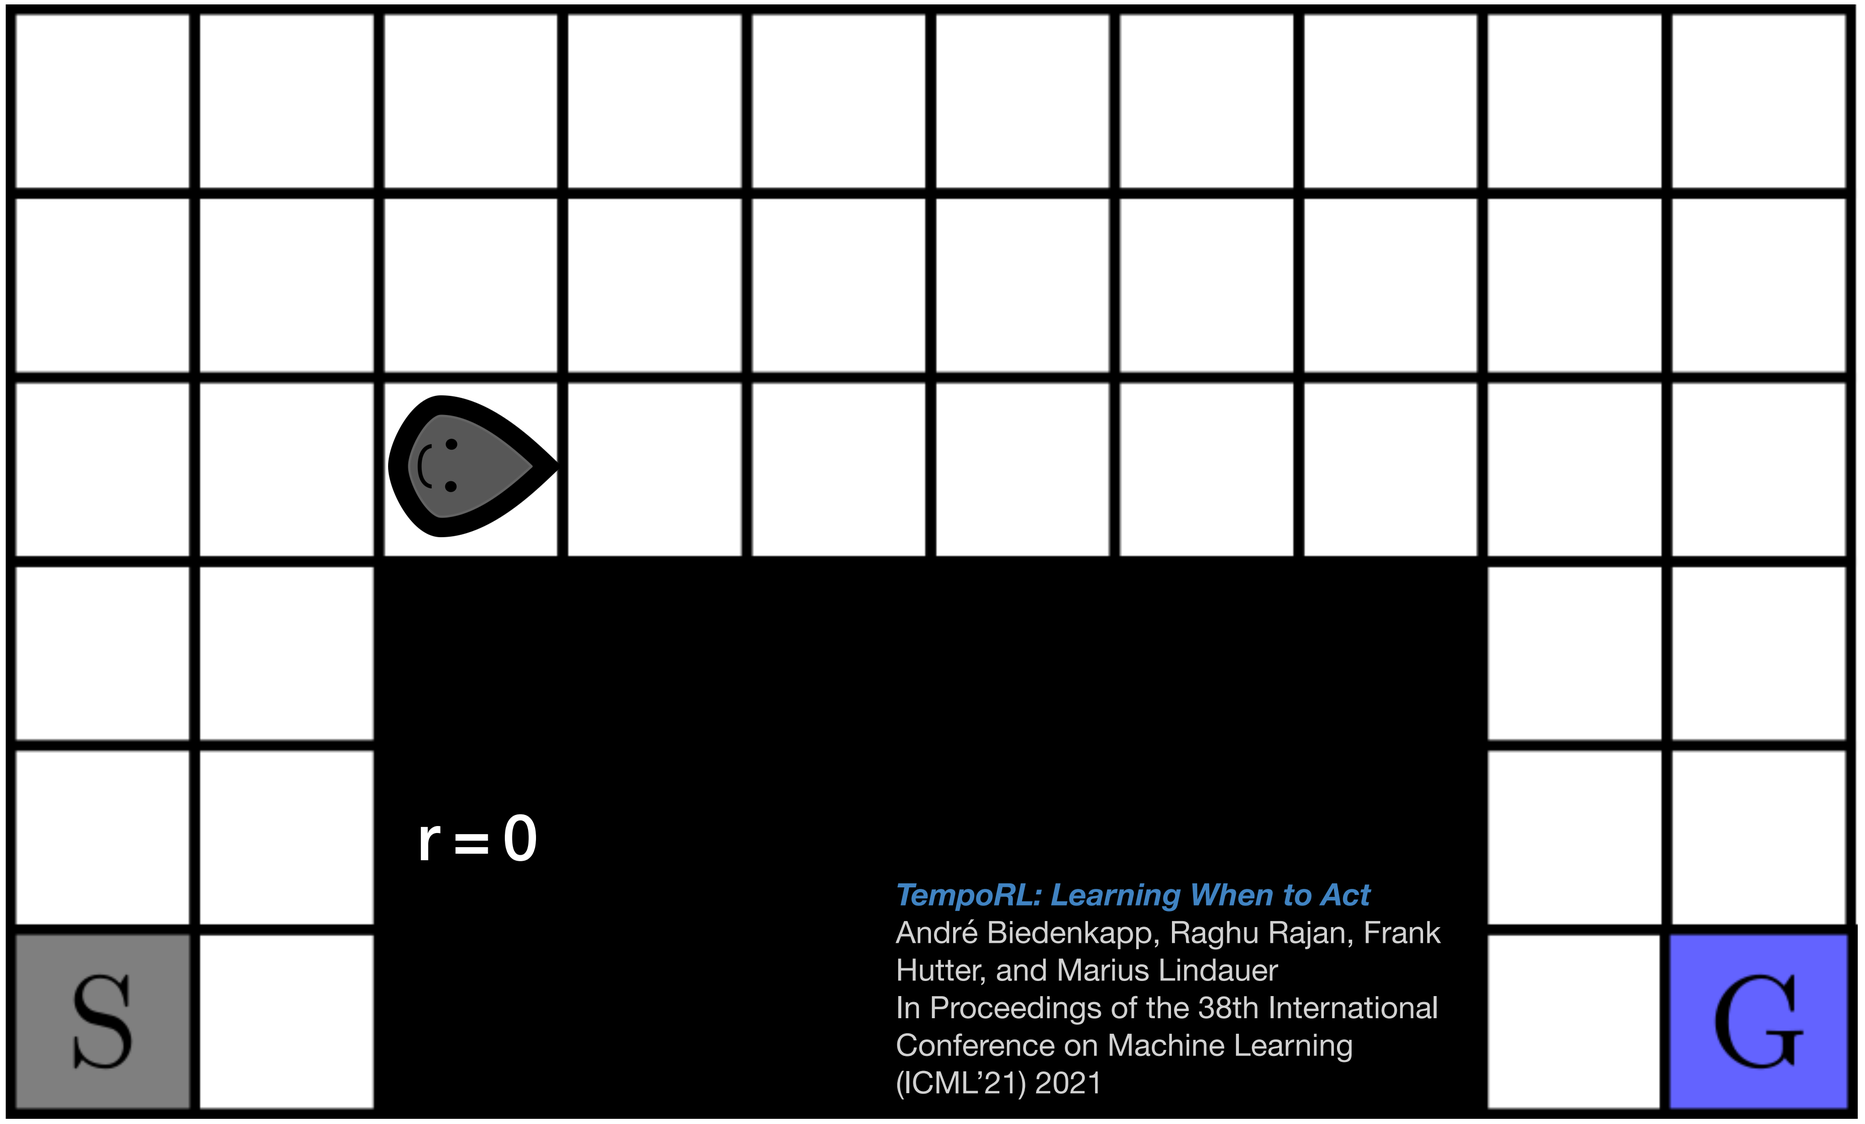

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

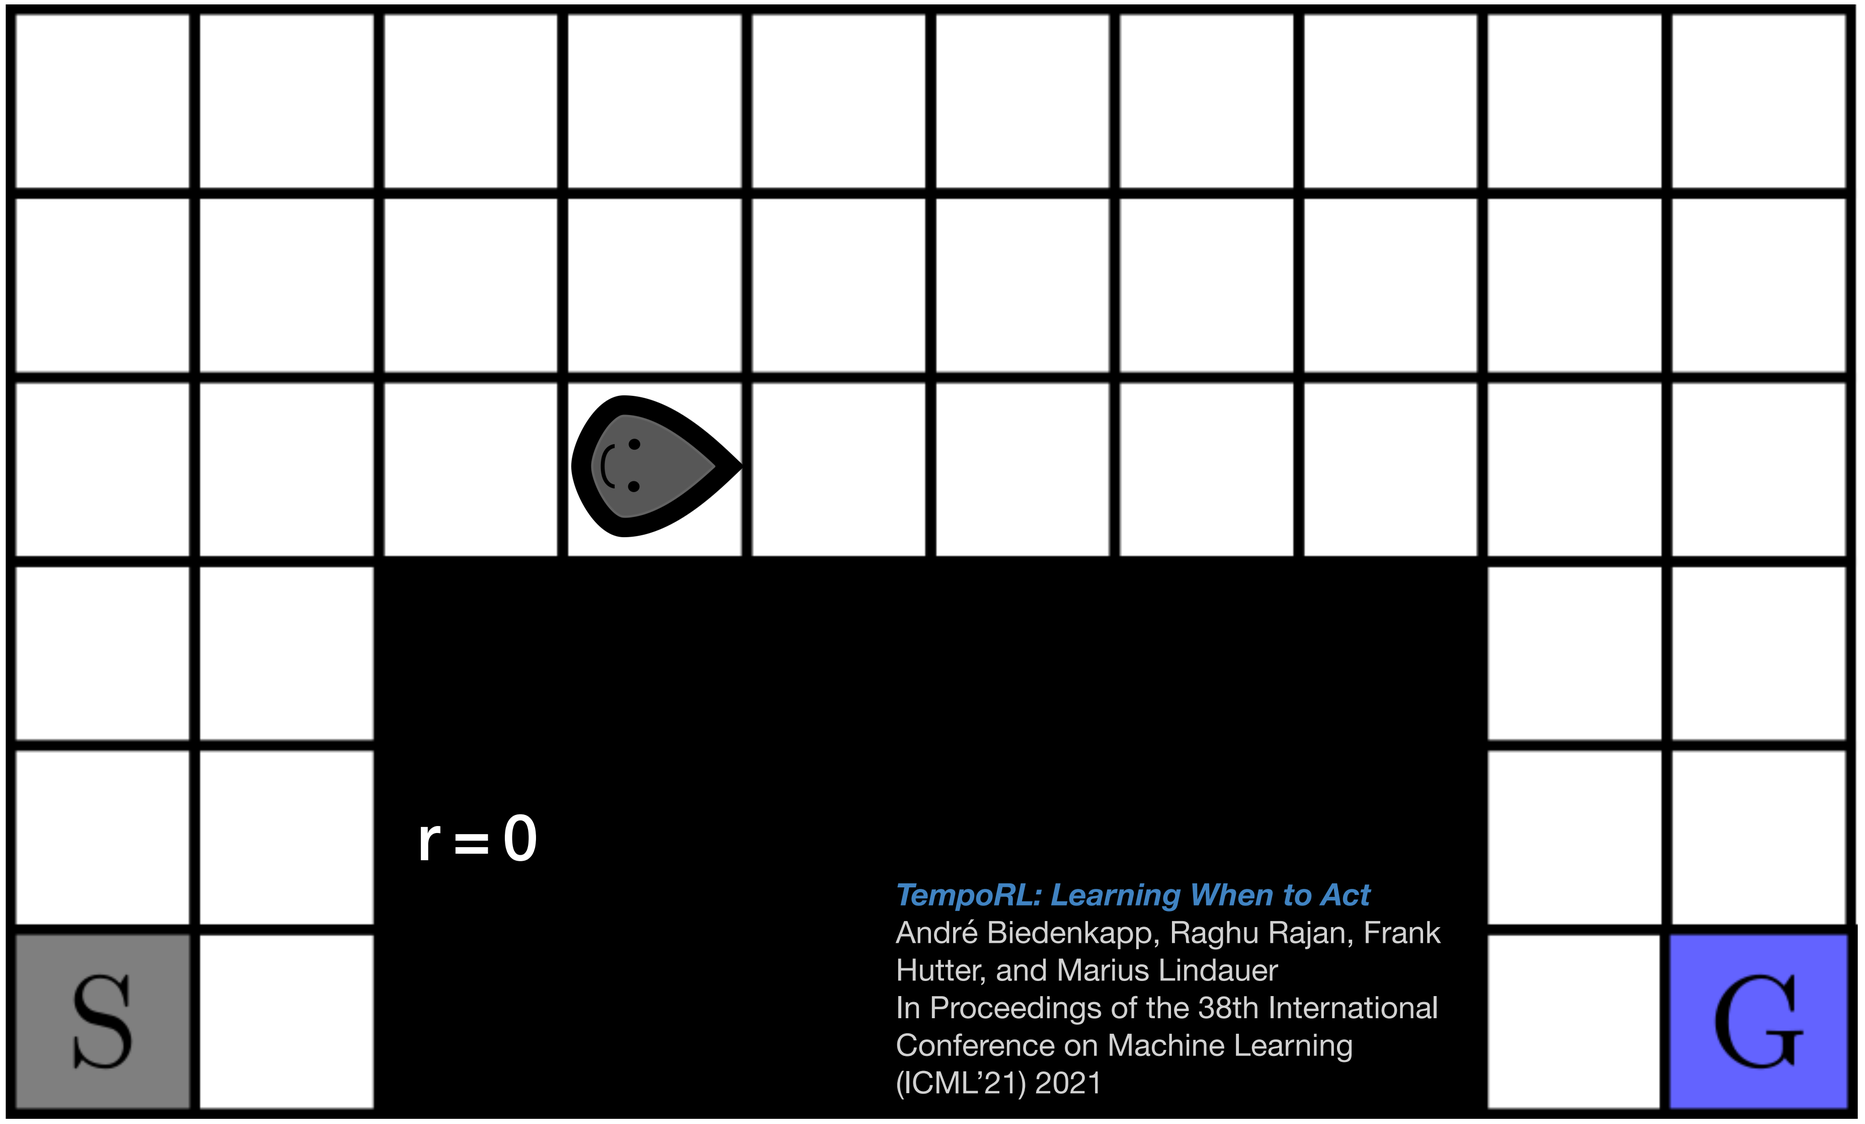

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

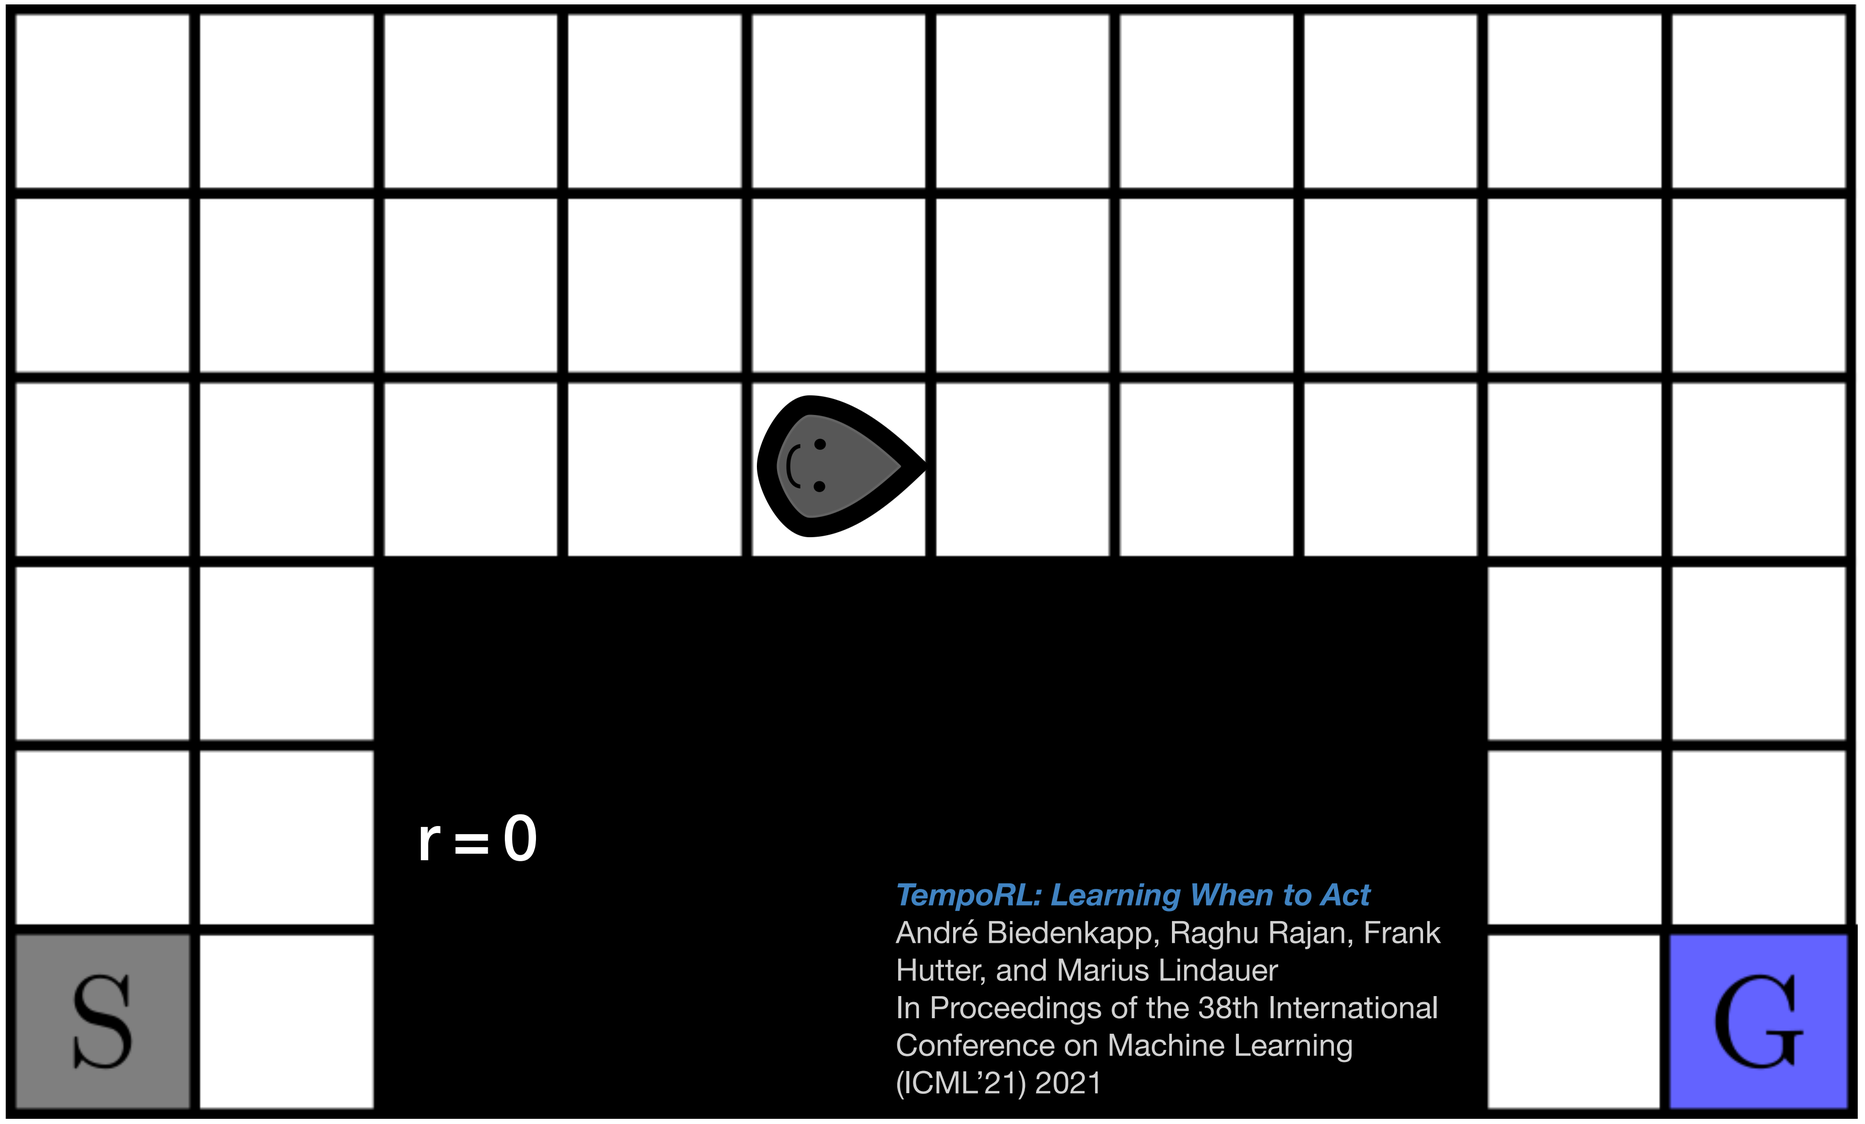
r = 0
S
G
TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

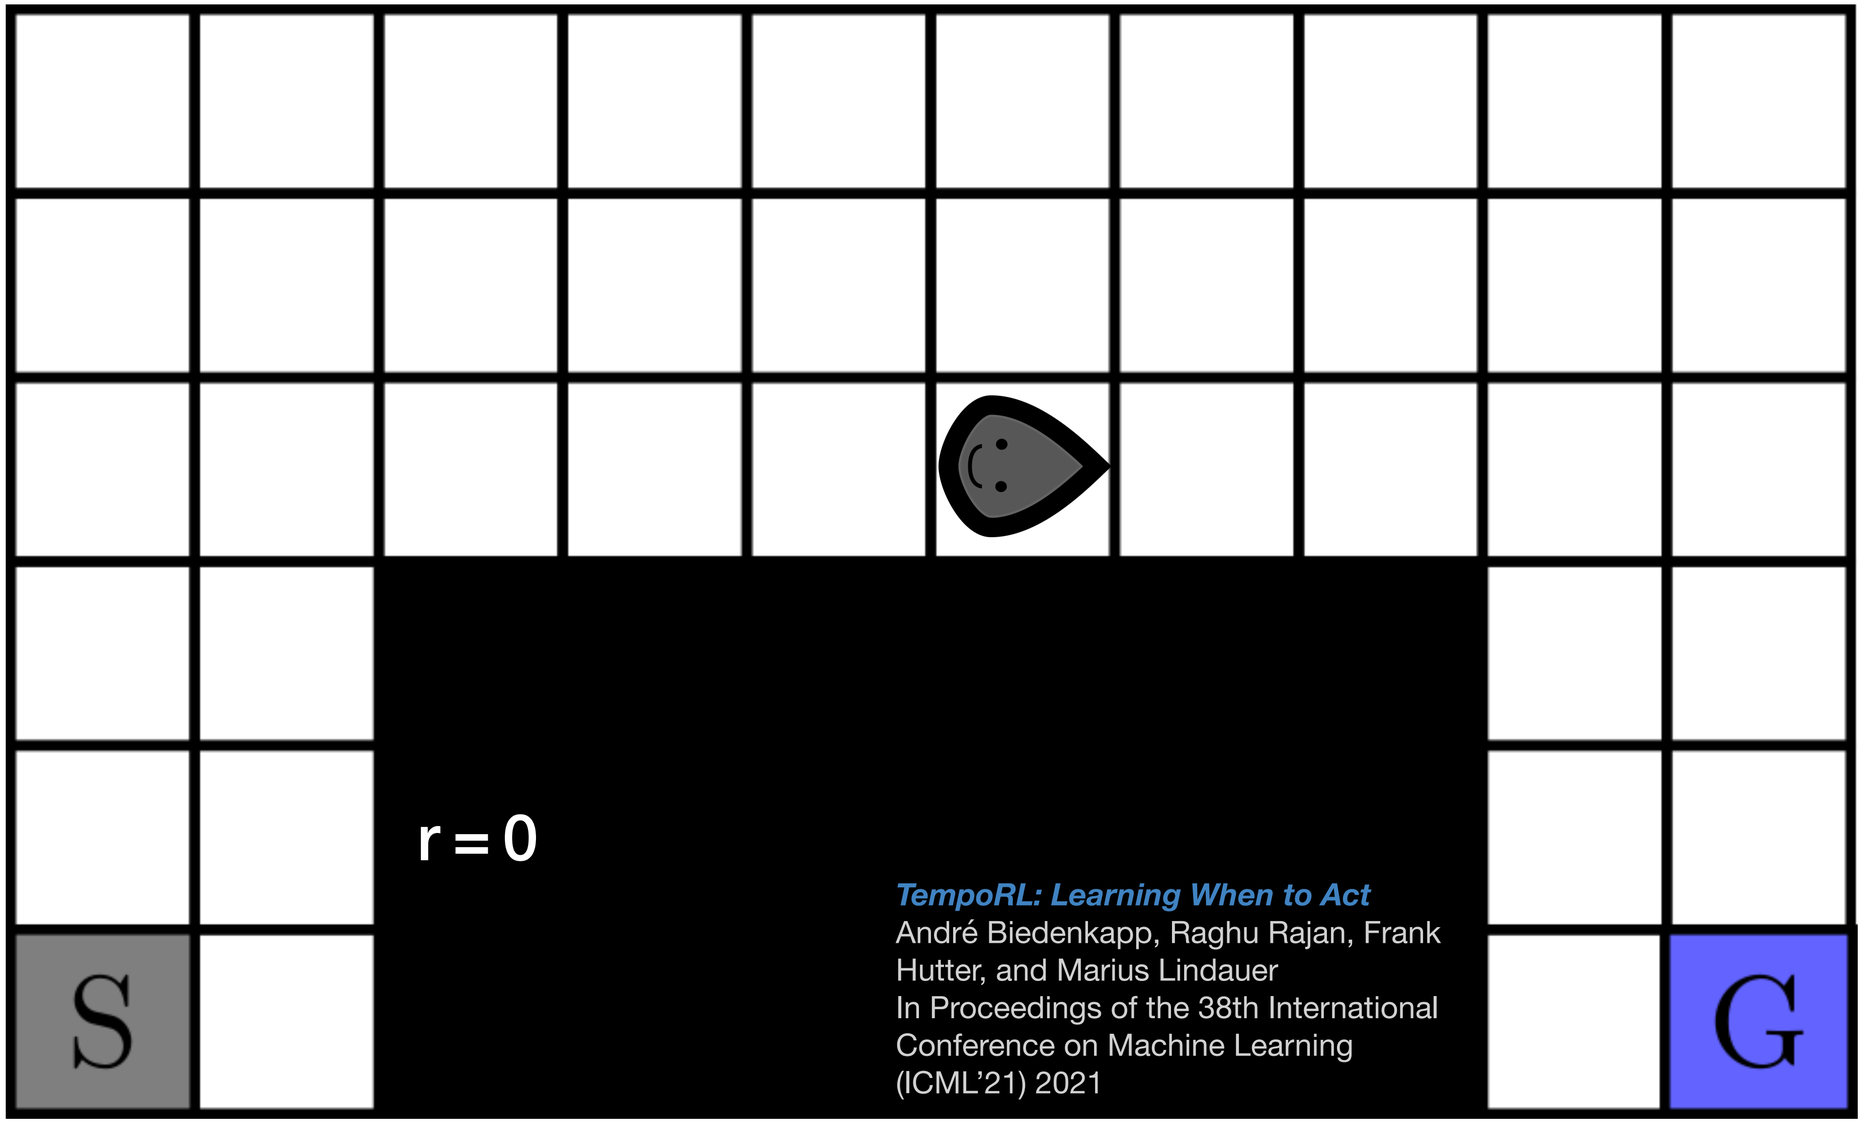

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

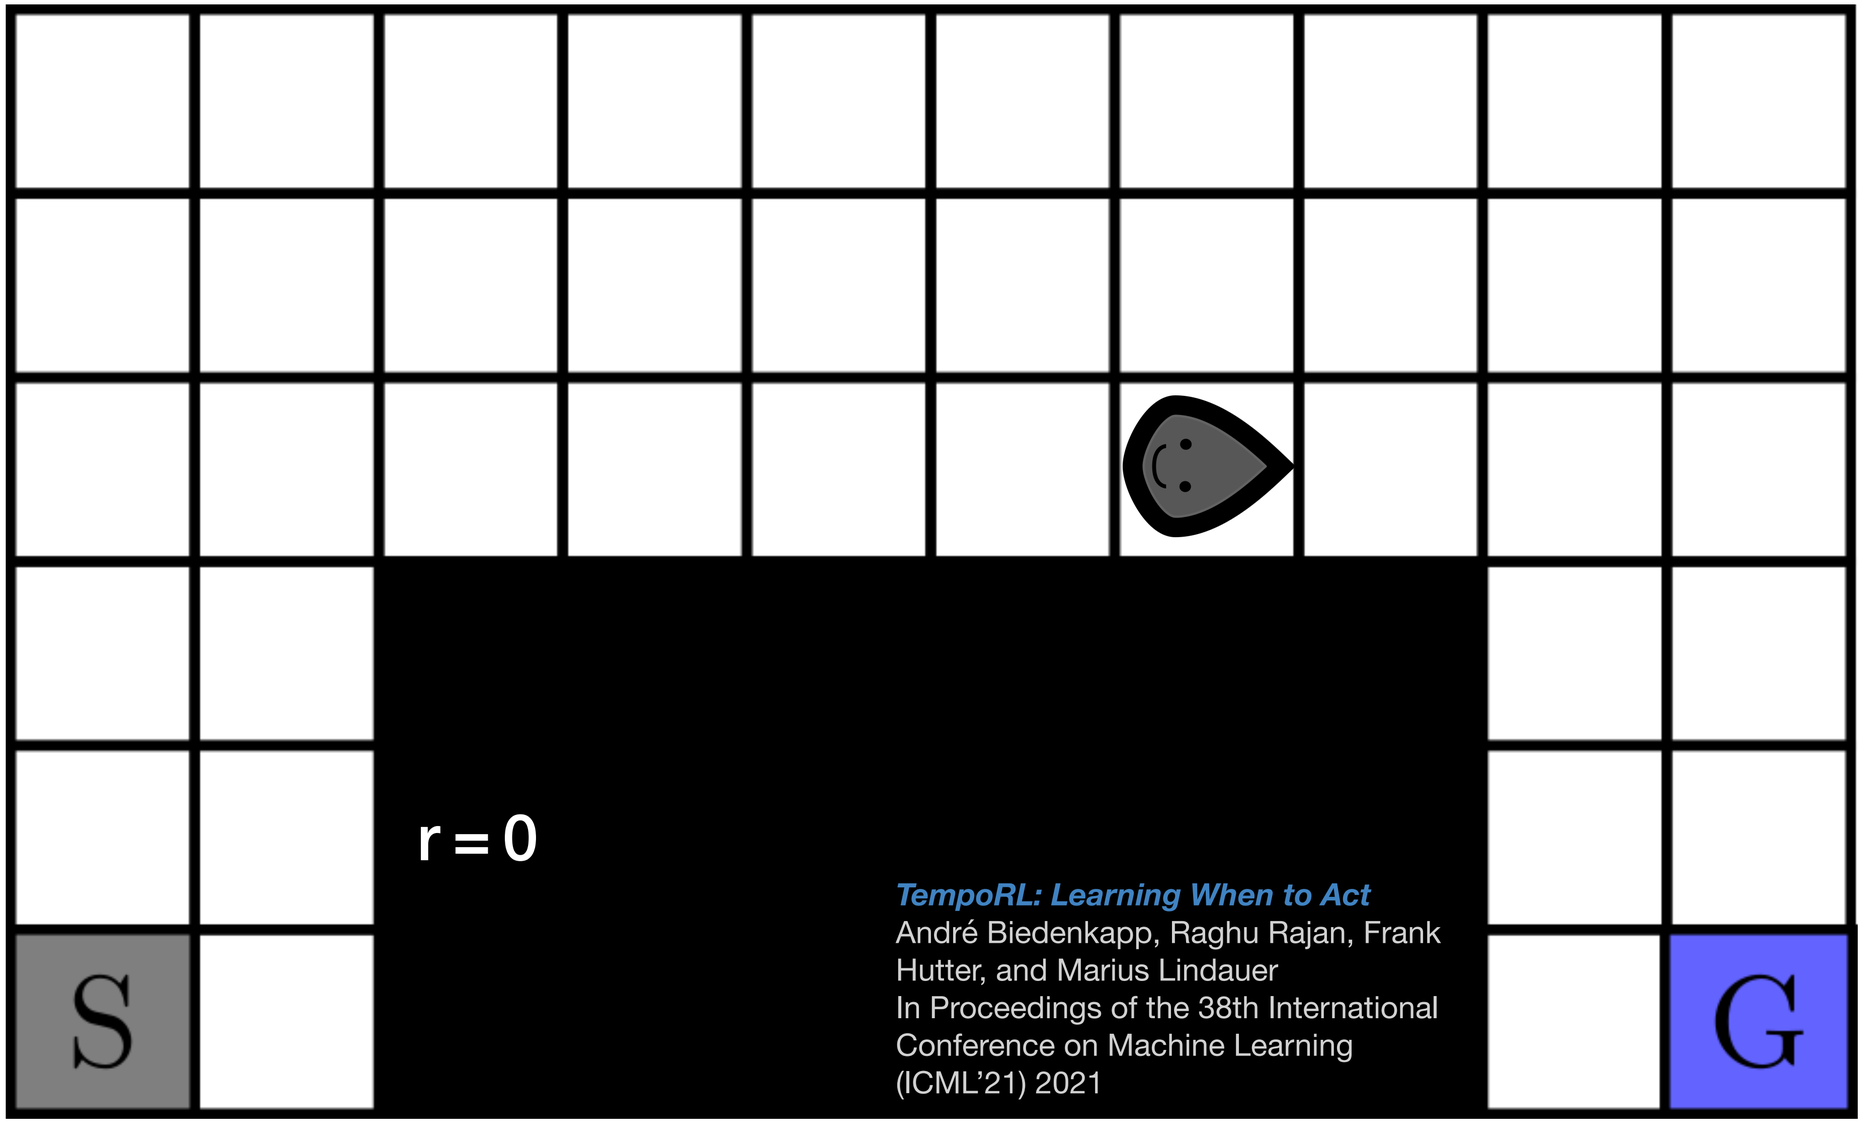

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

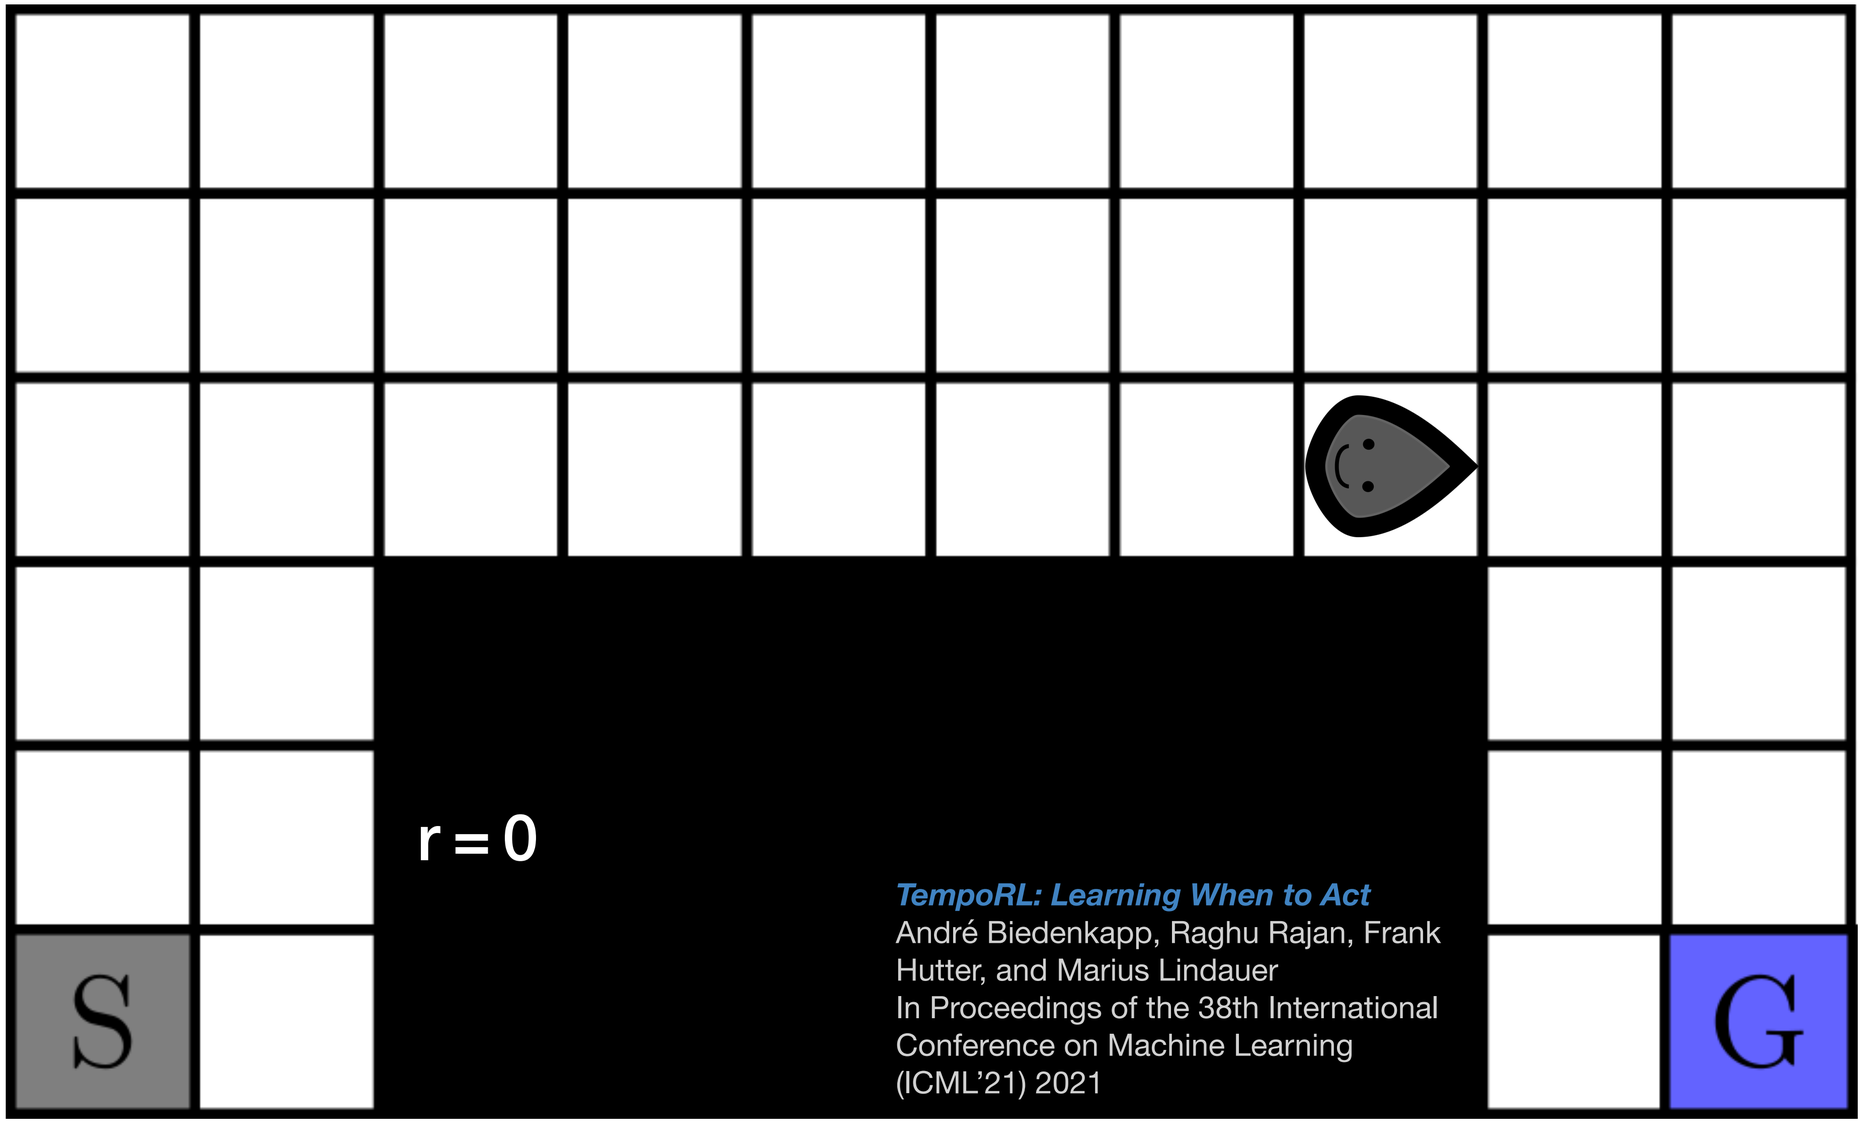

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

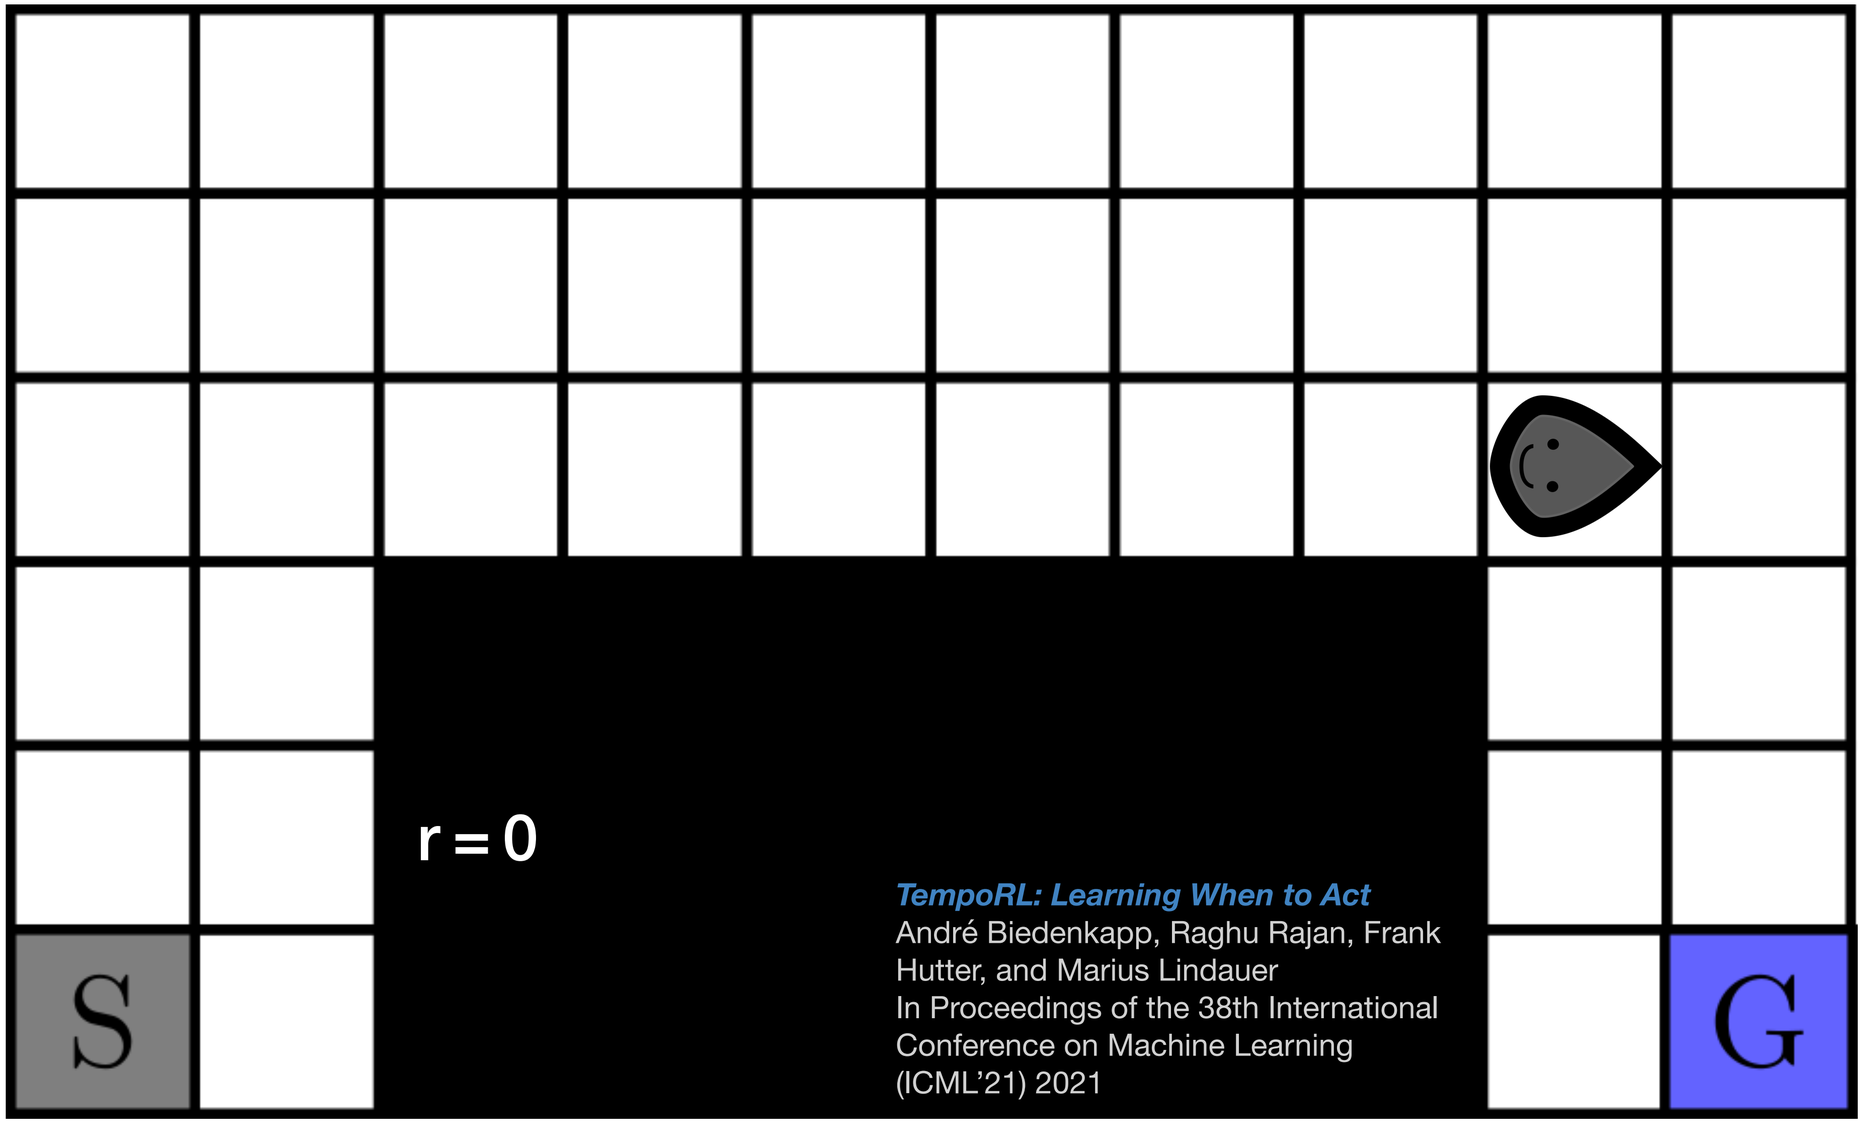

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

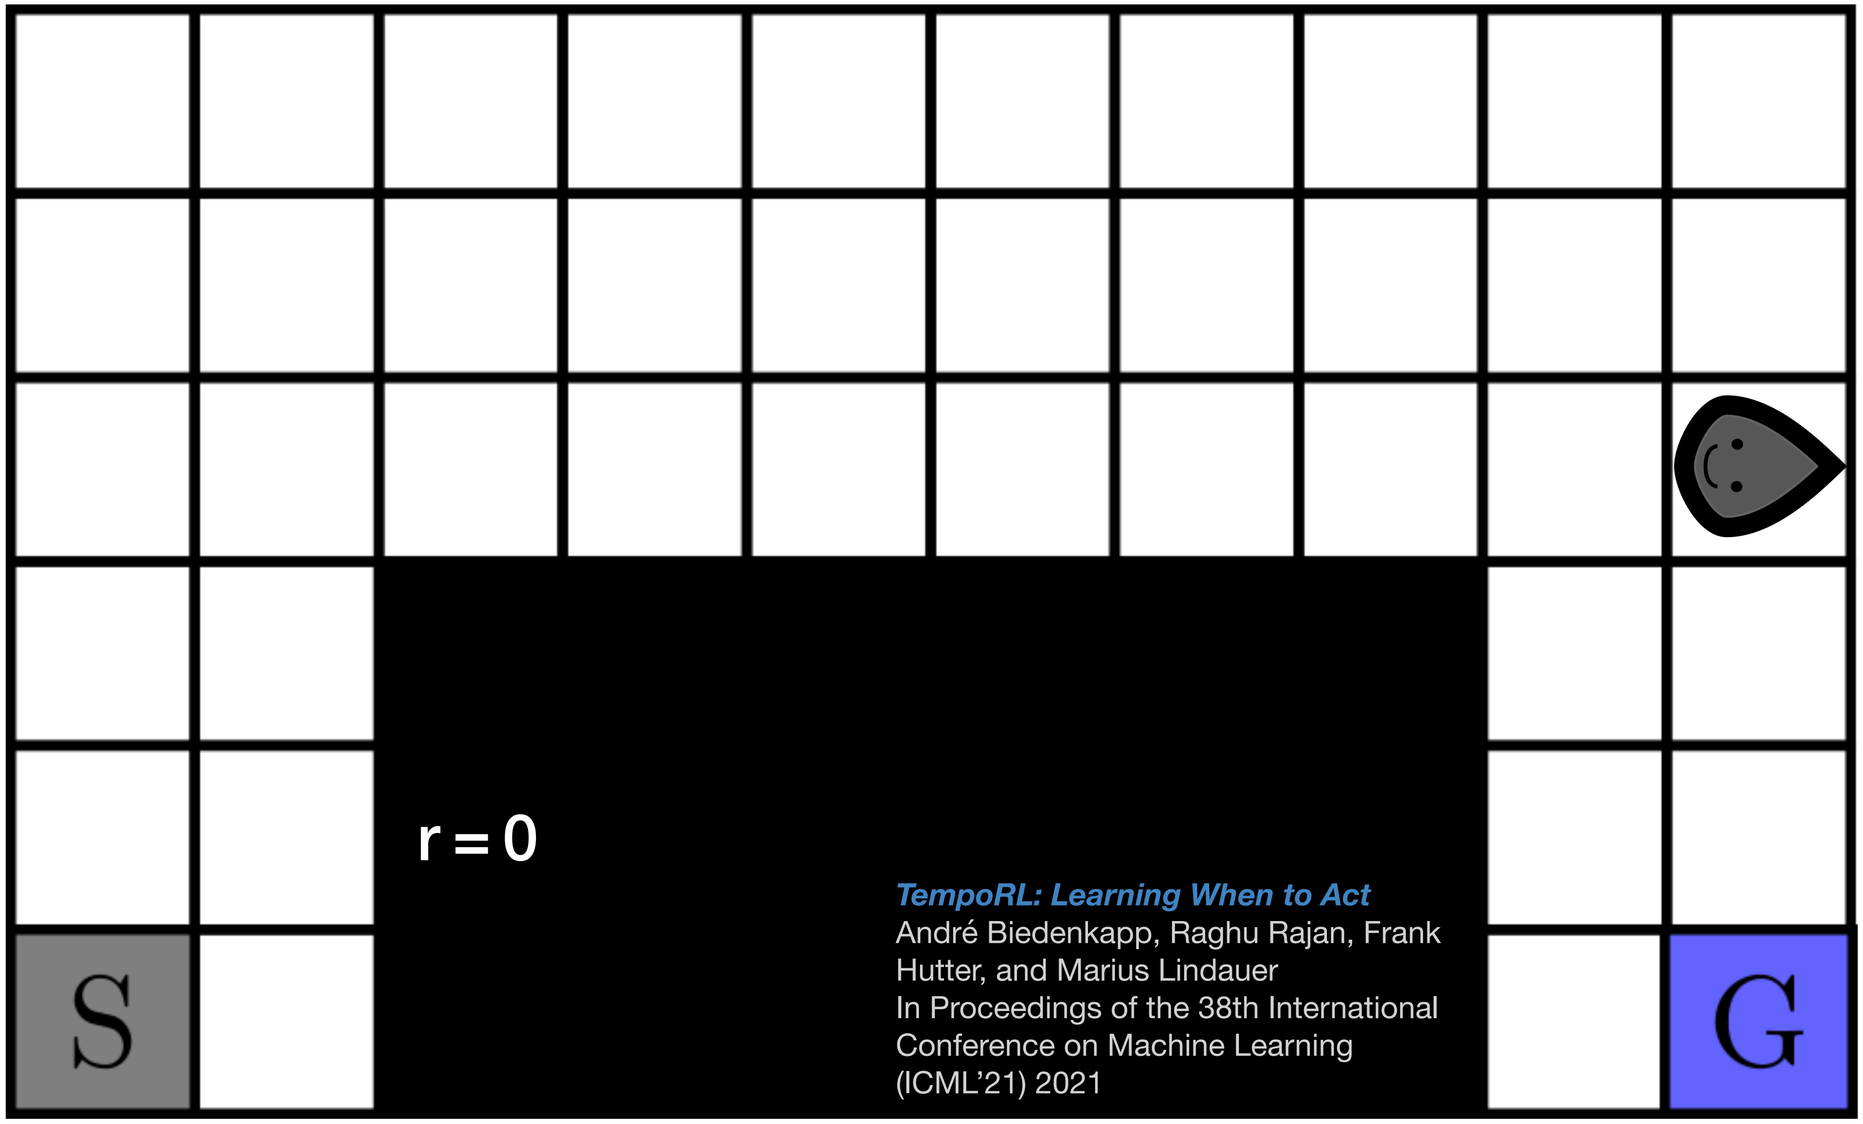

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

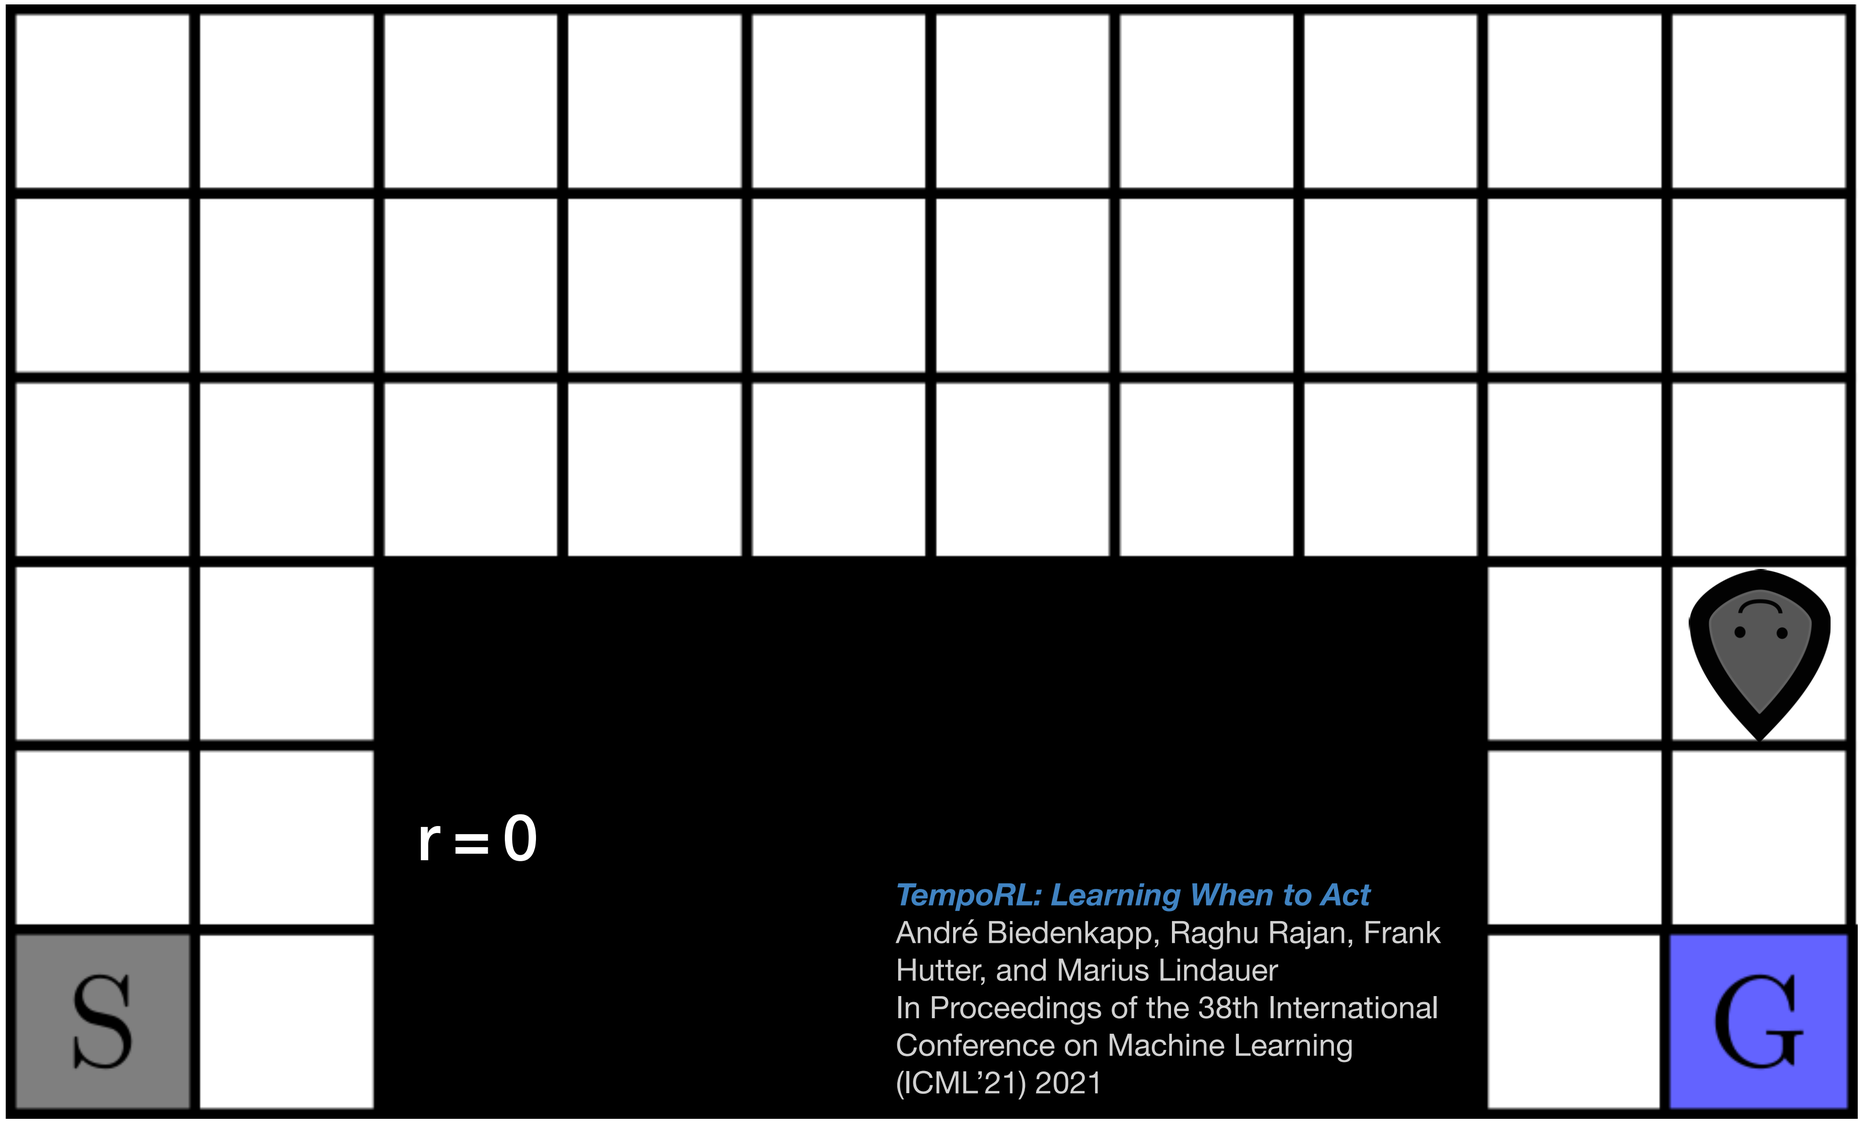

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

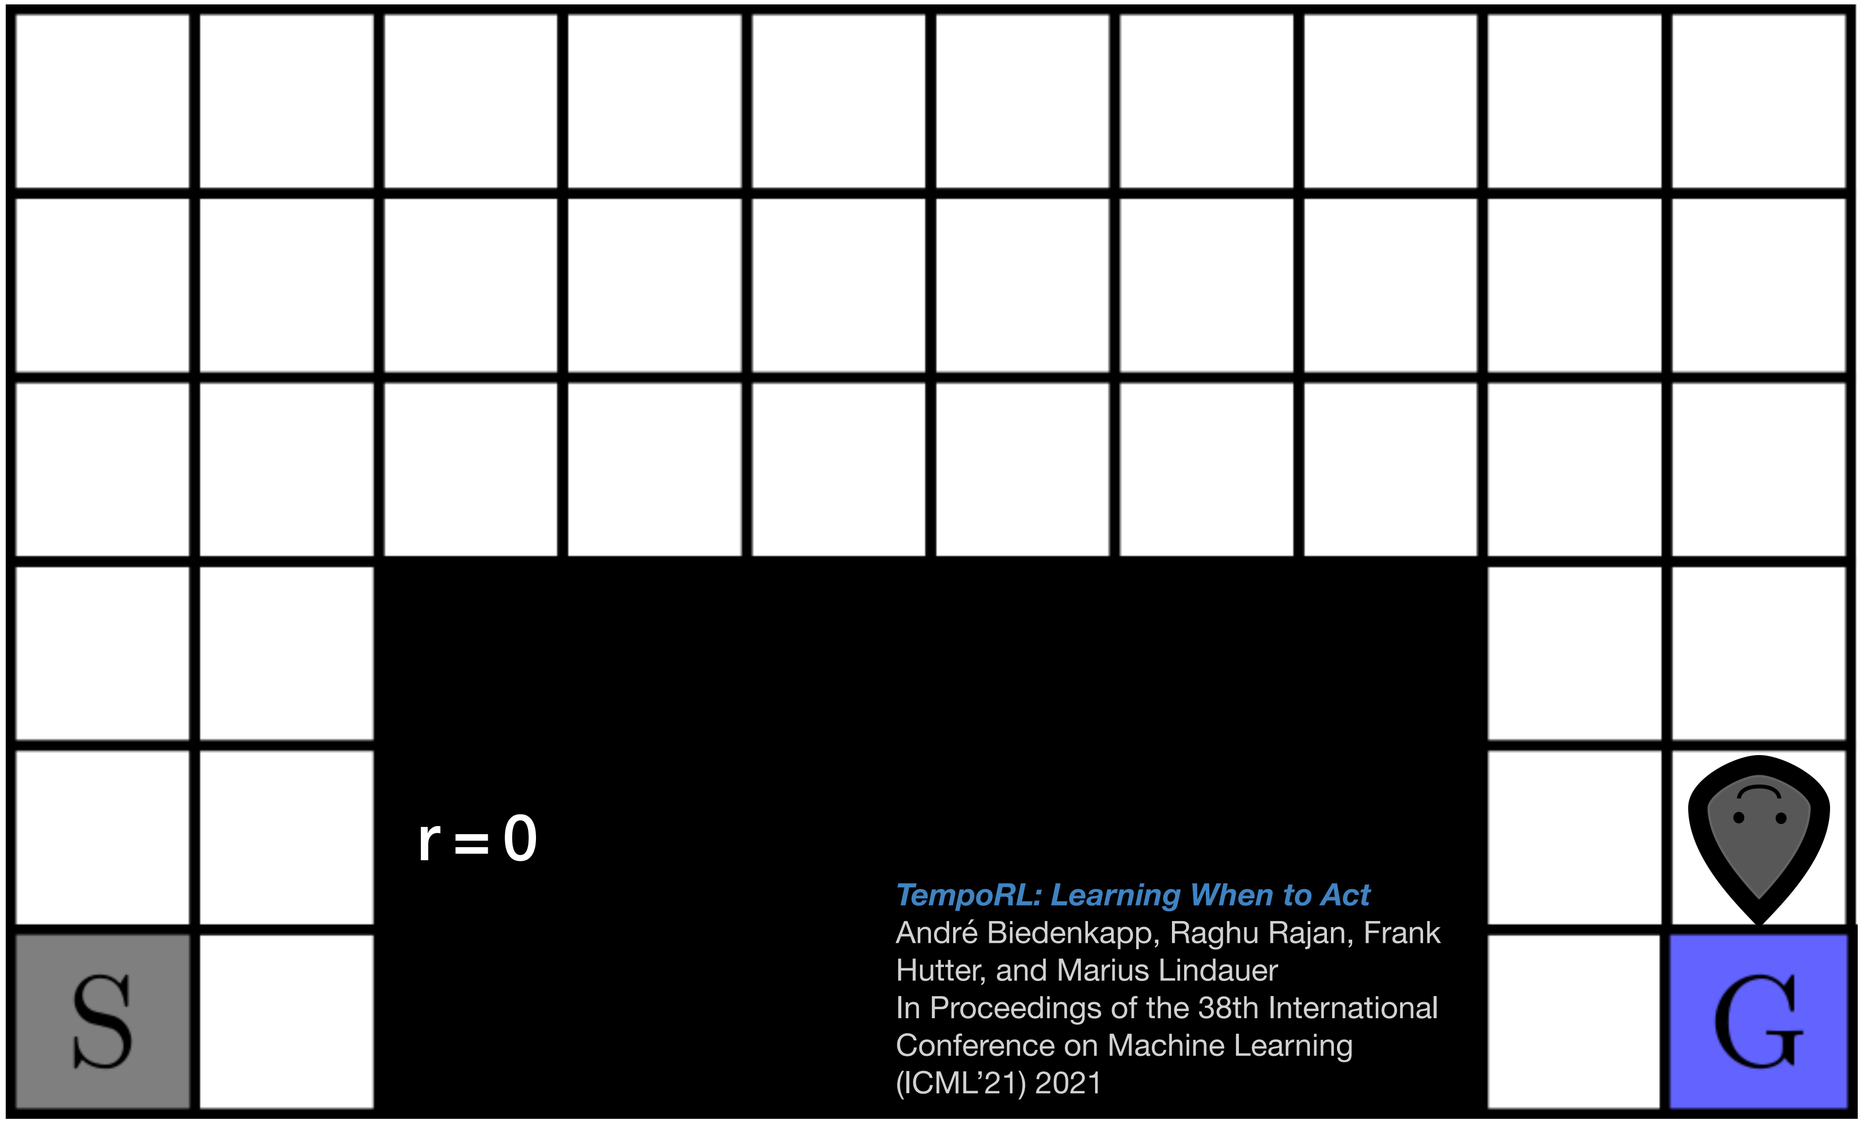

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

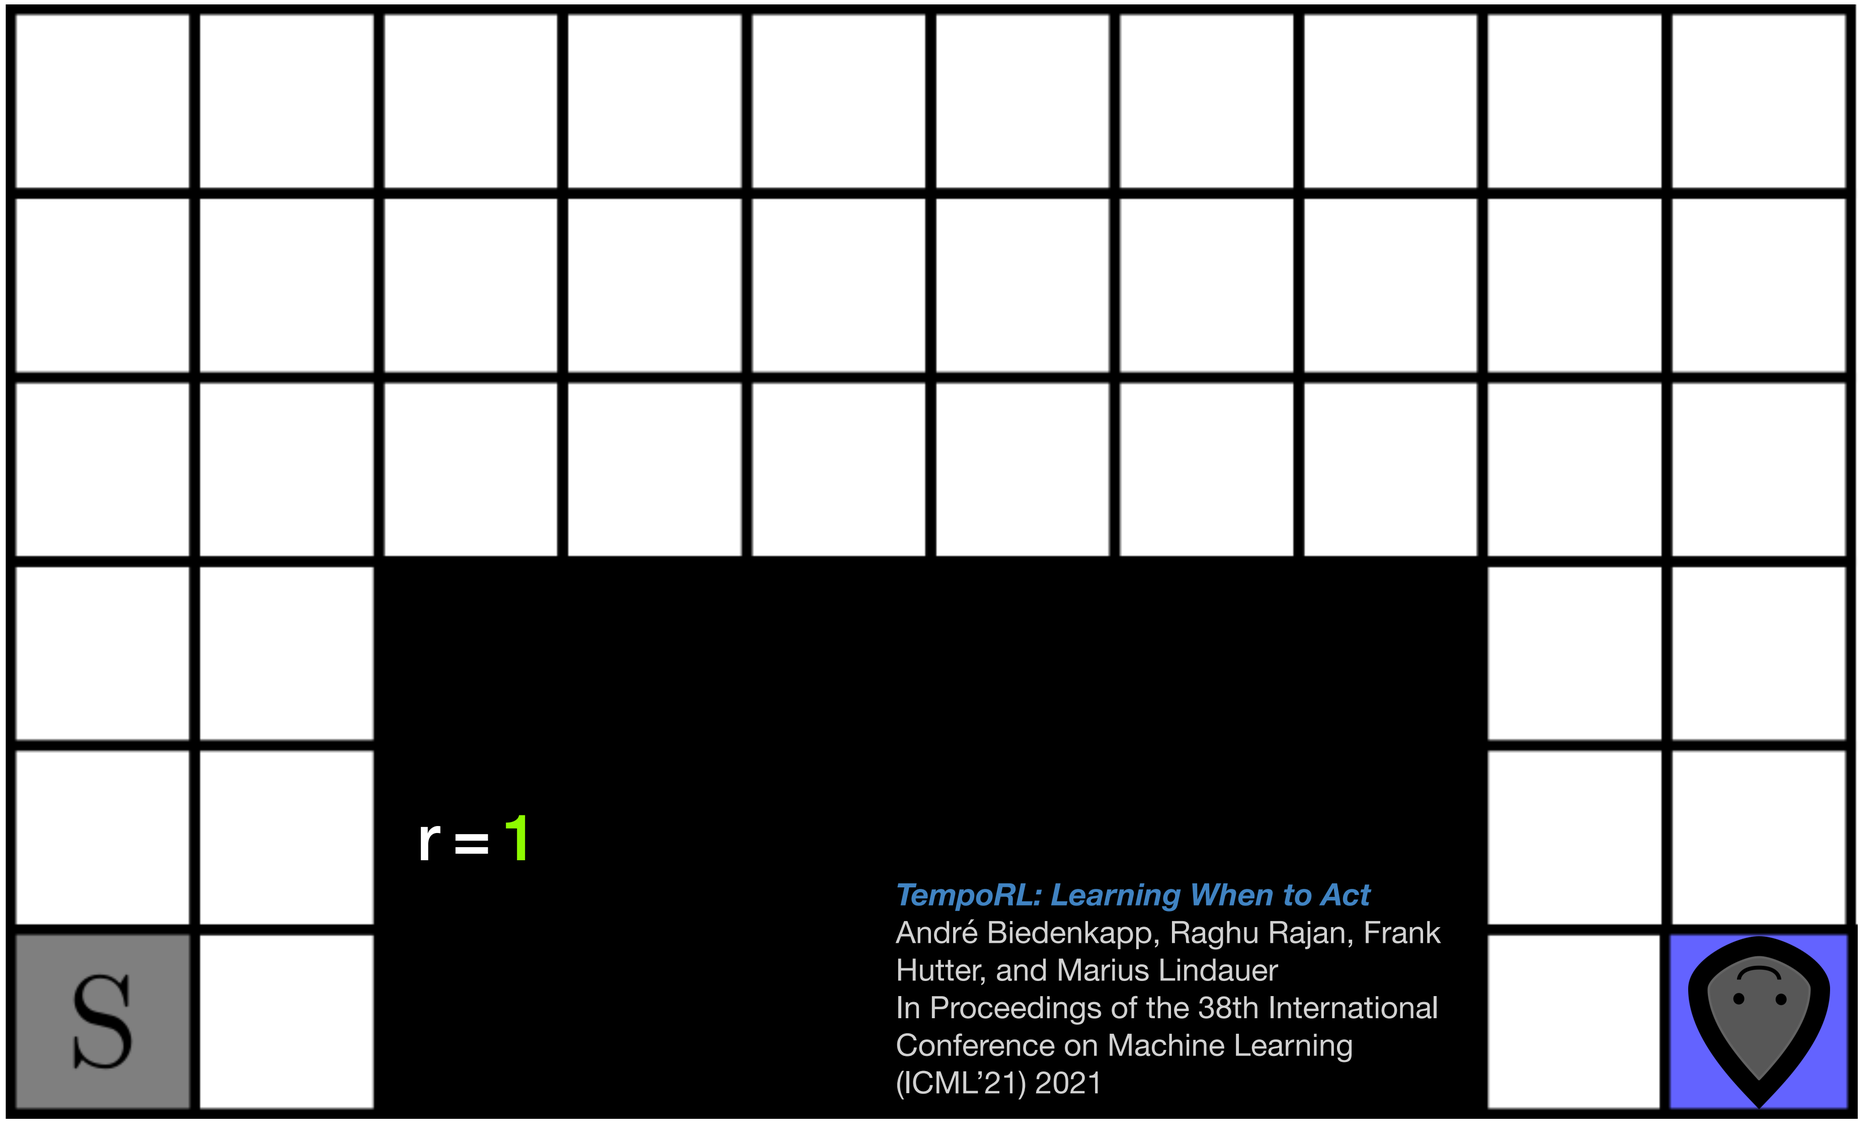

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

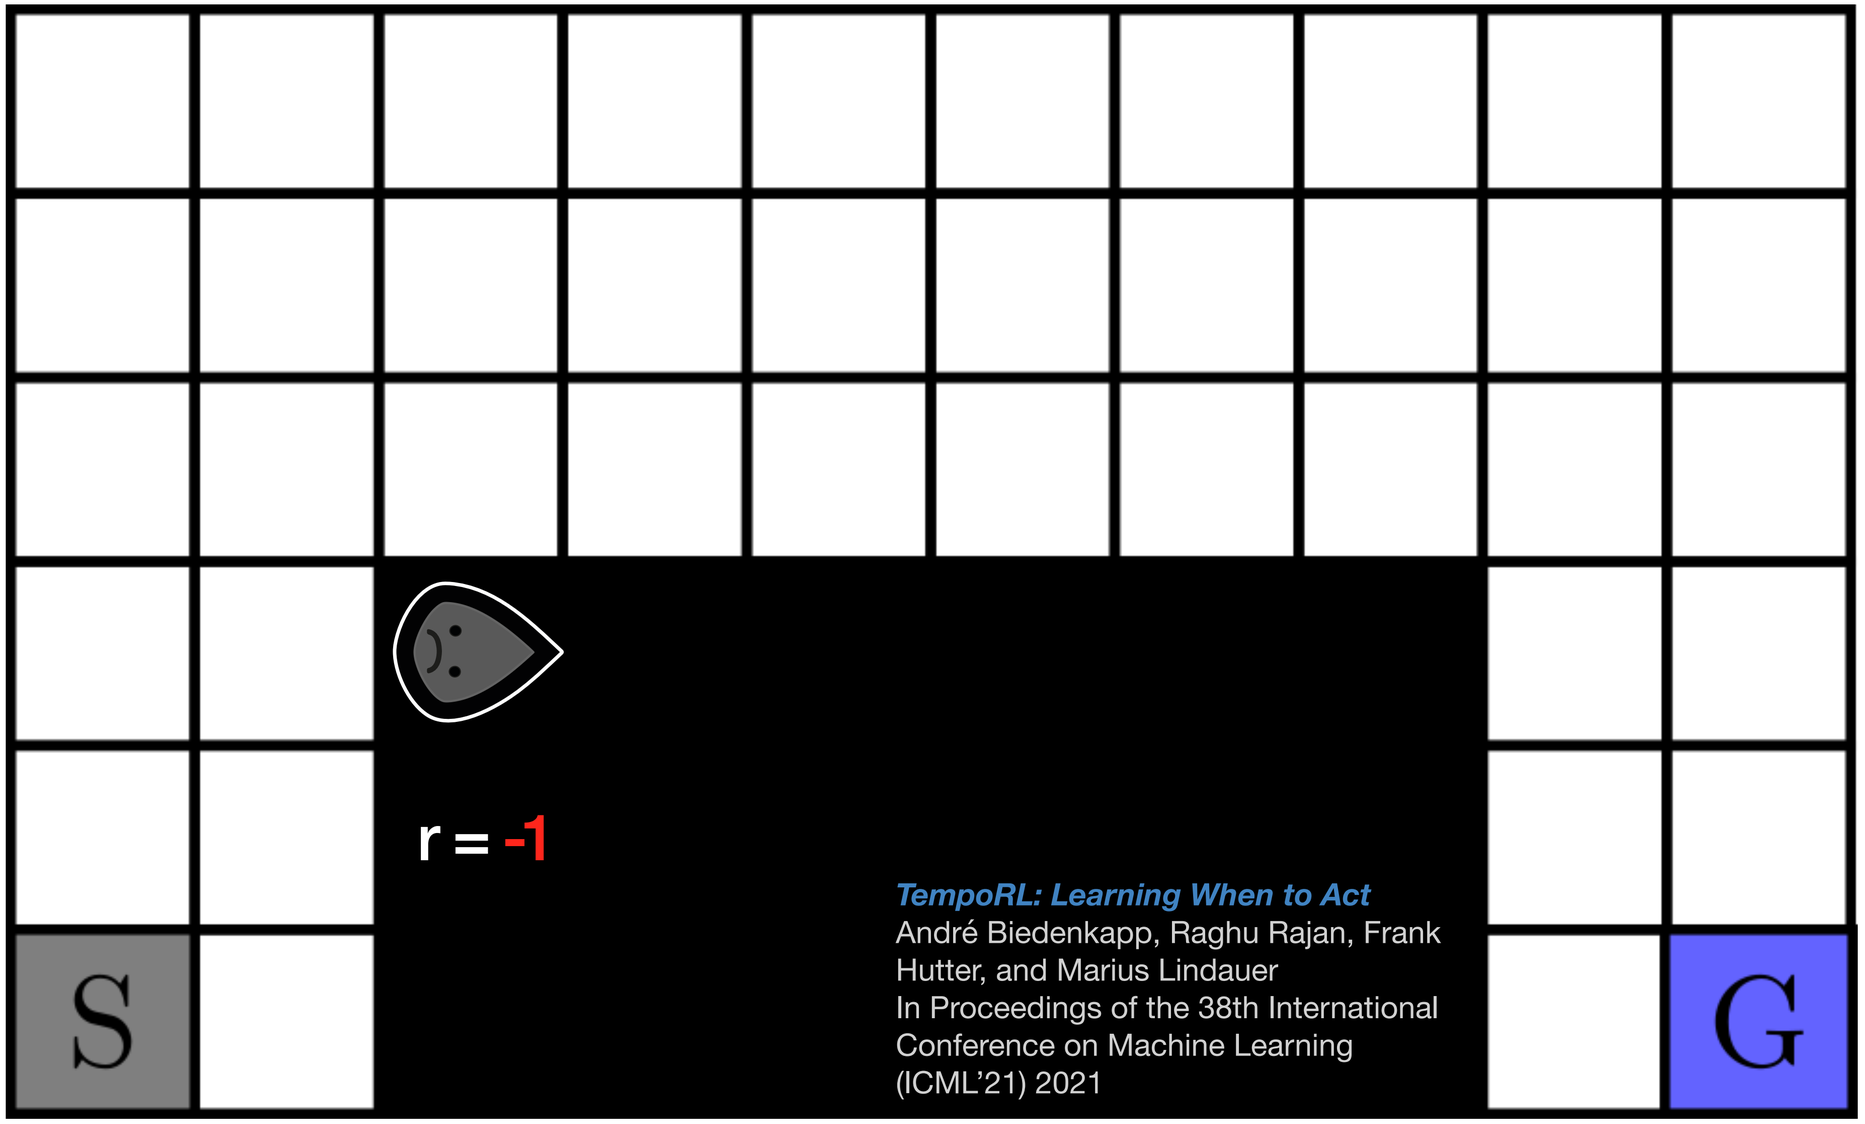

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

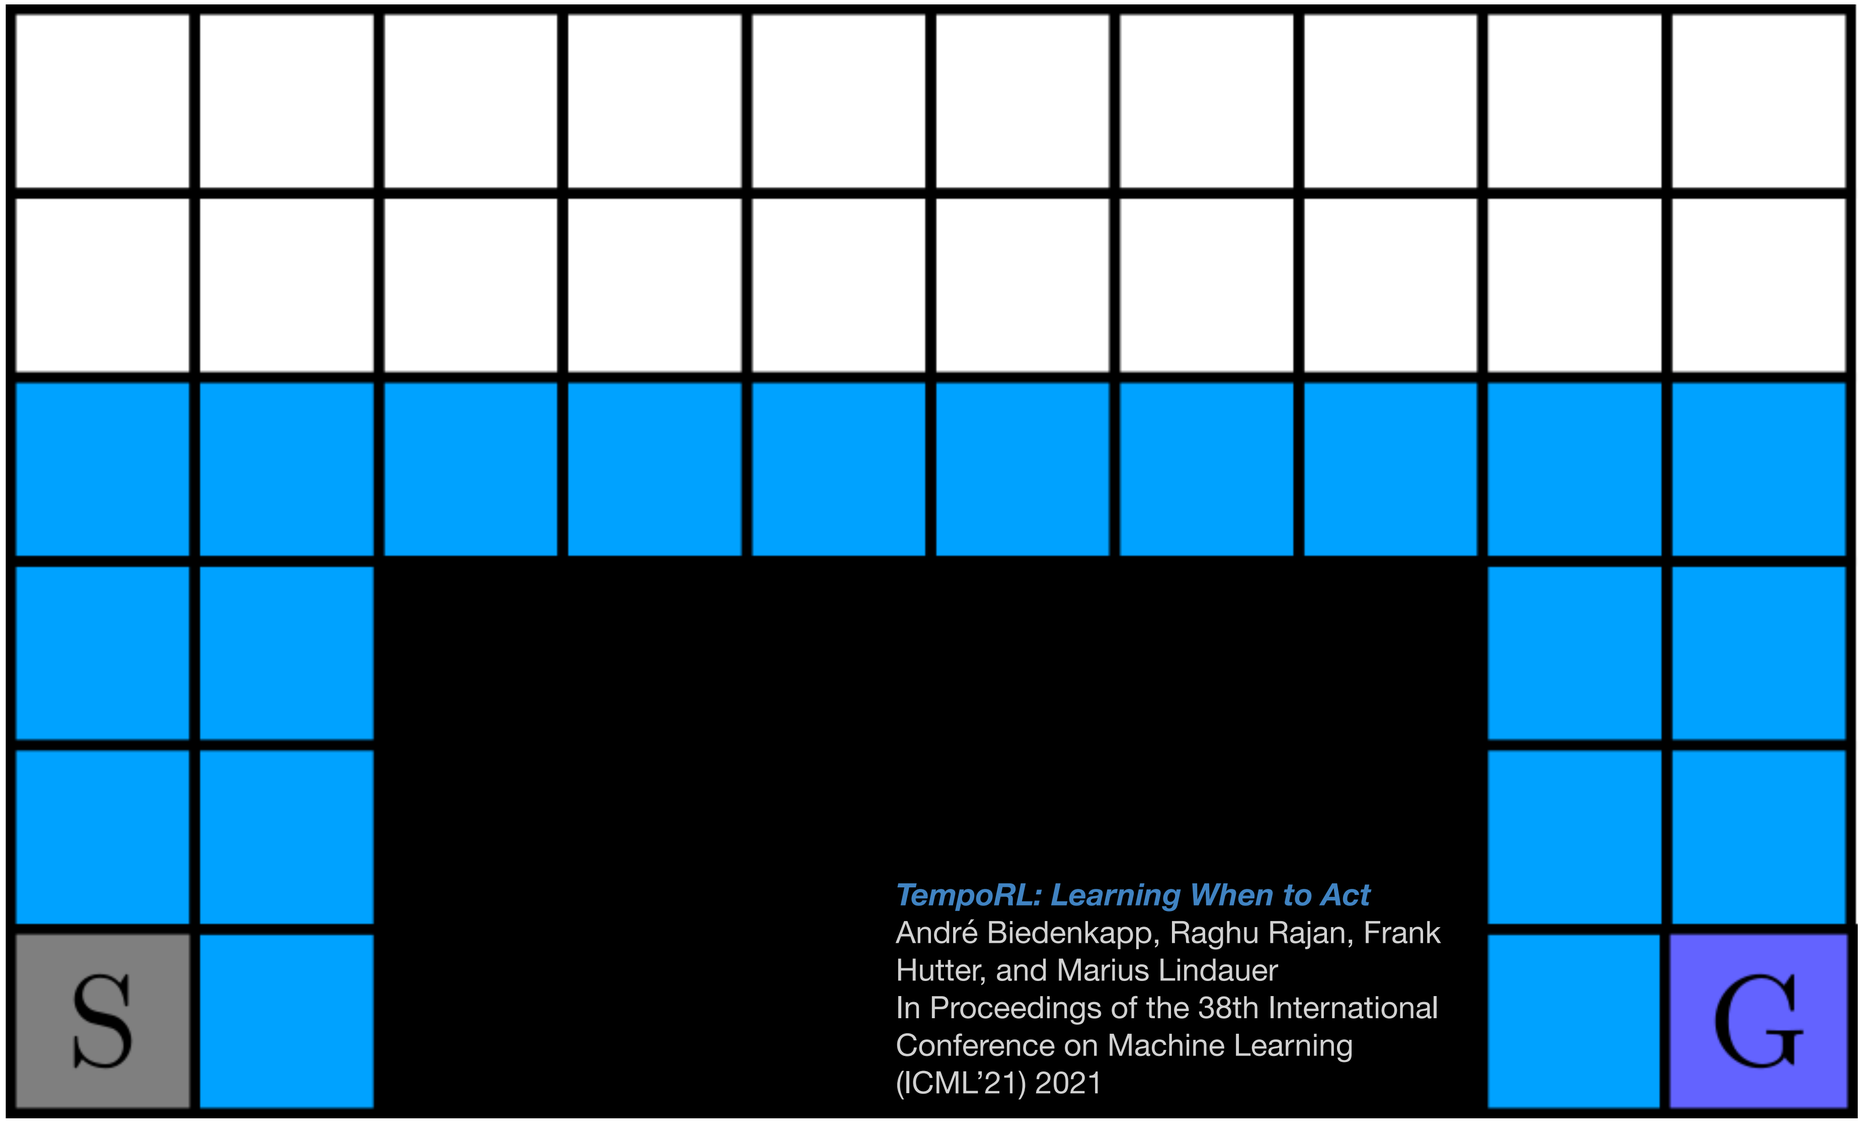

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

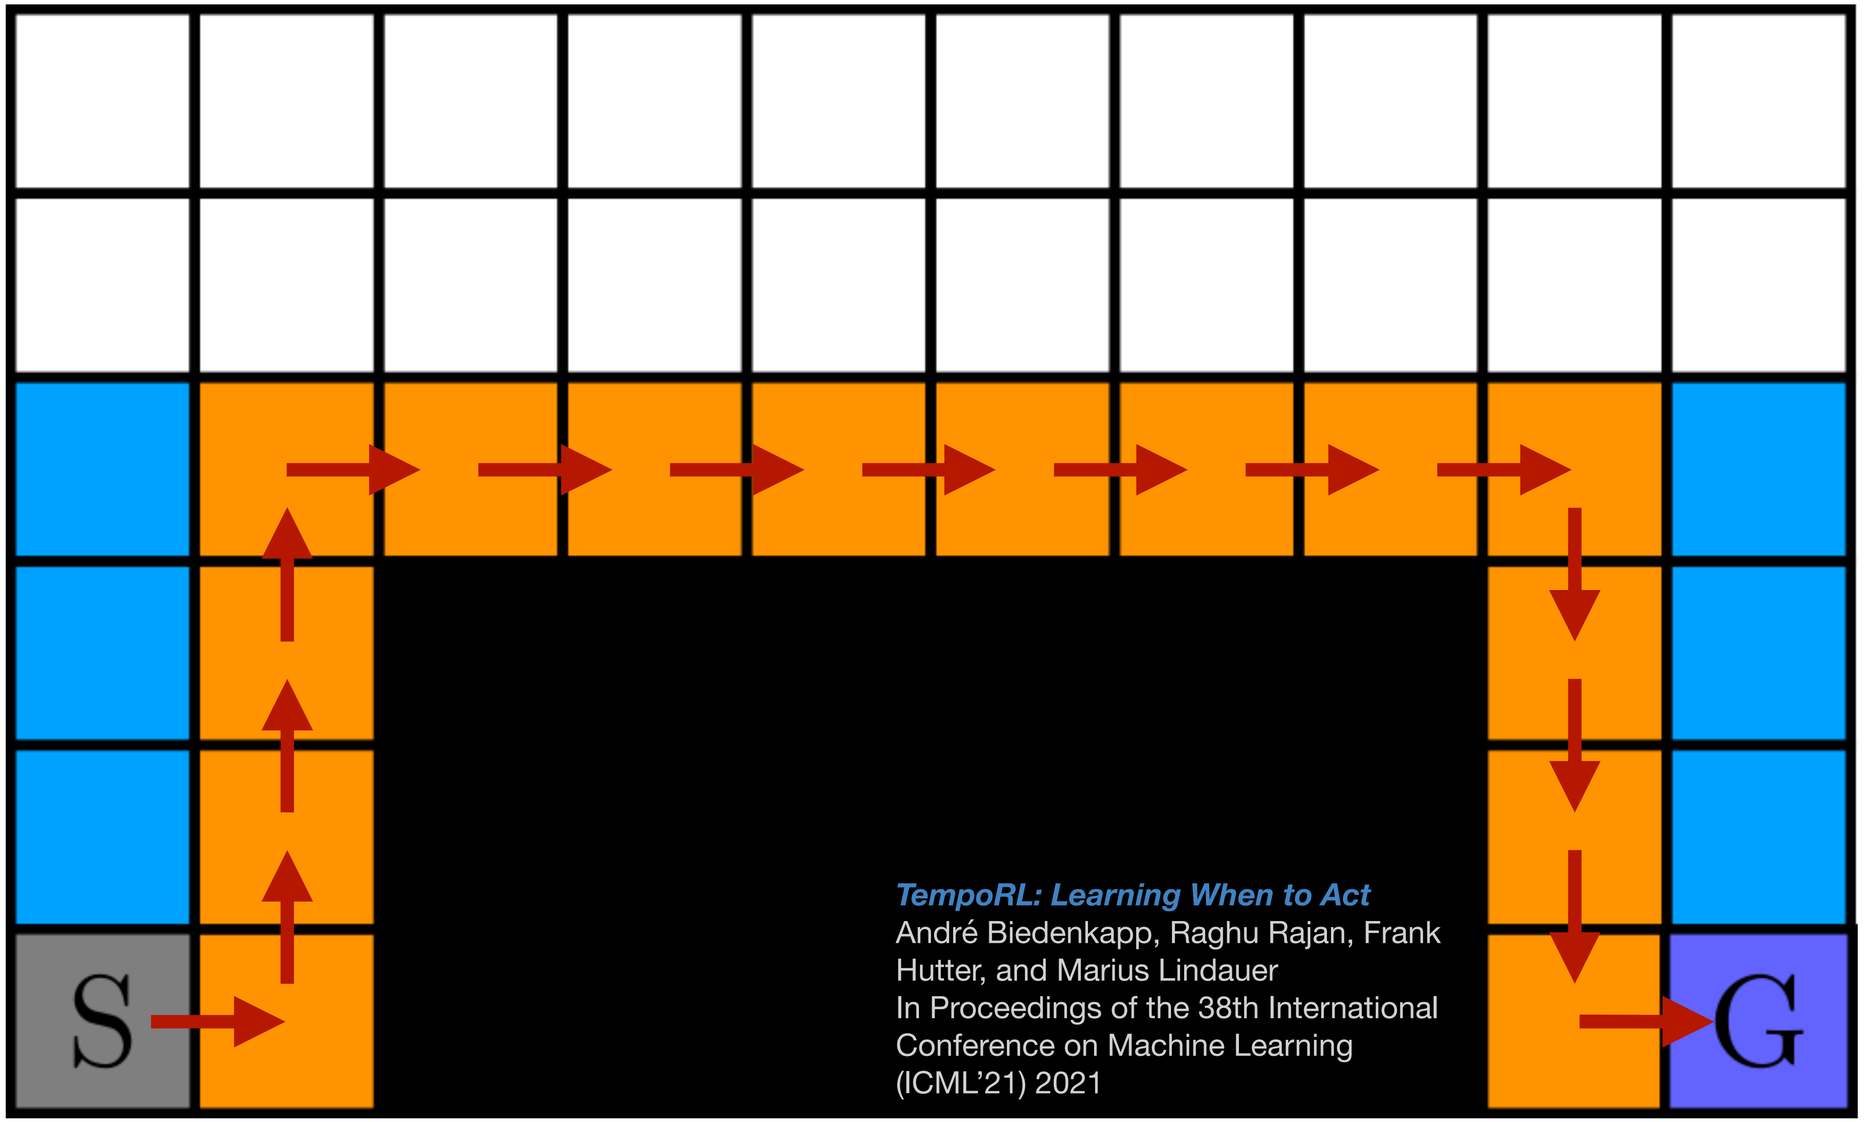

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

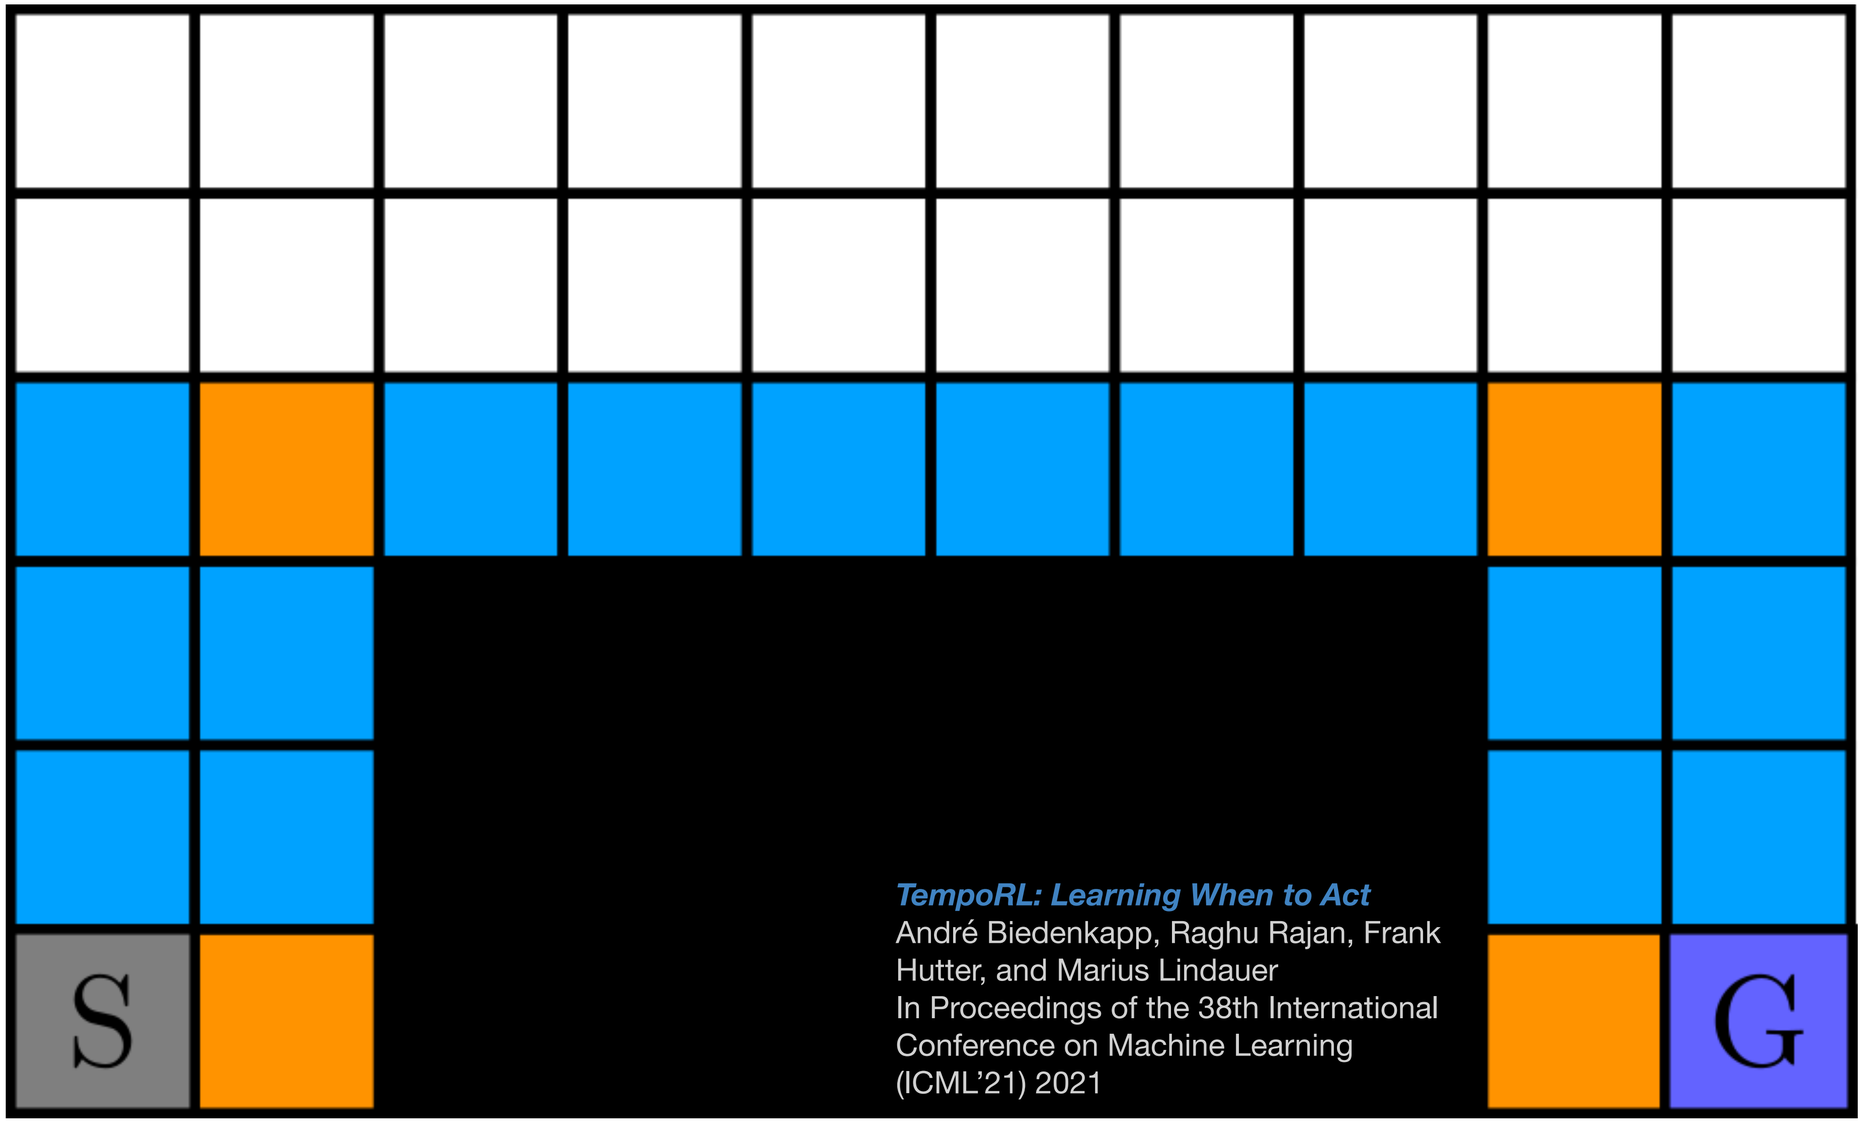

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL

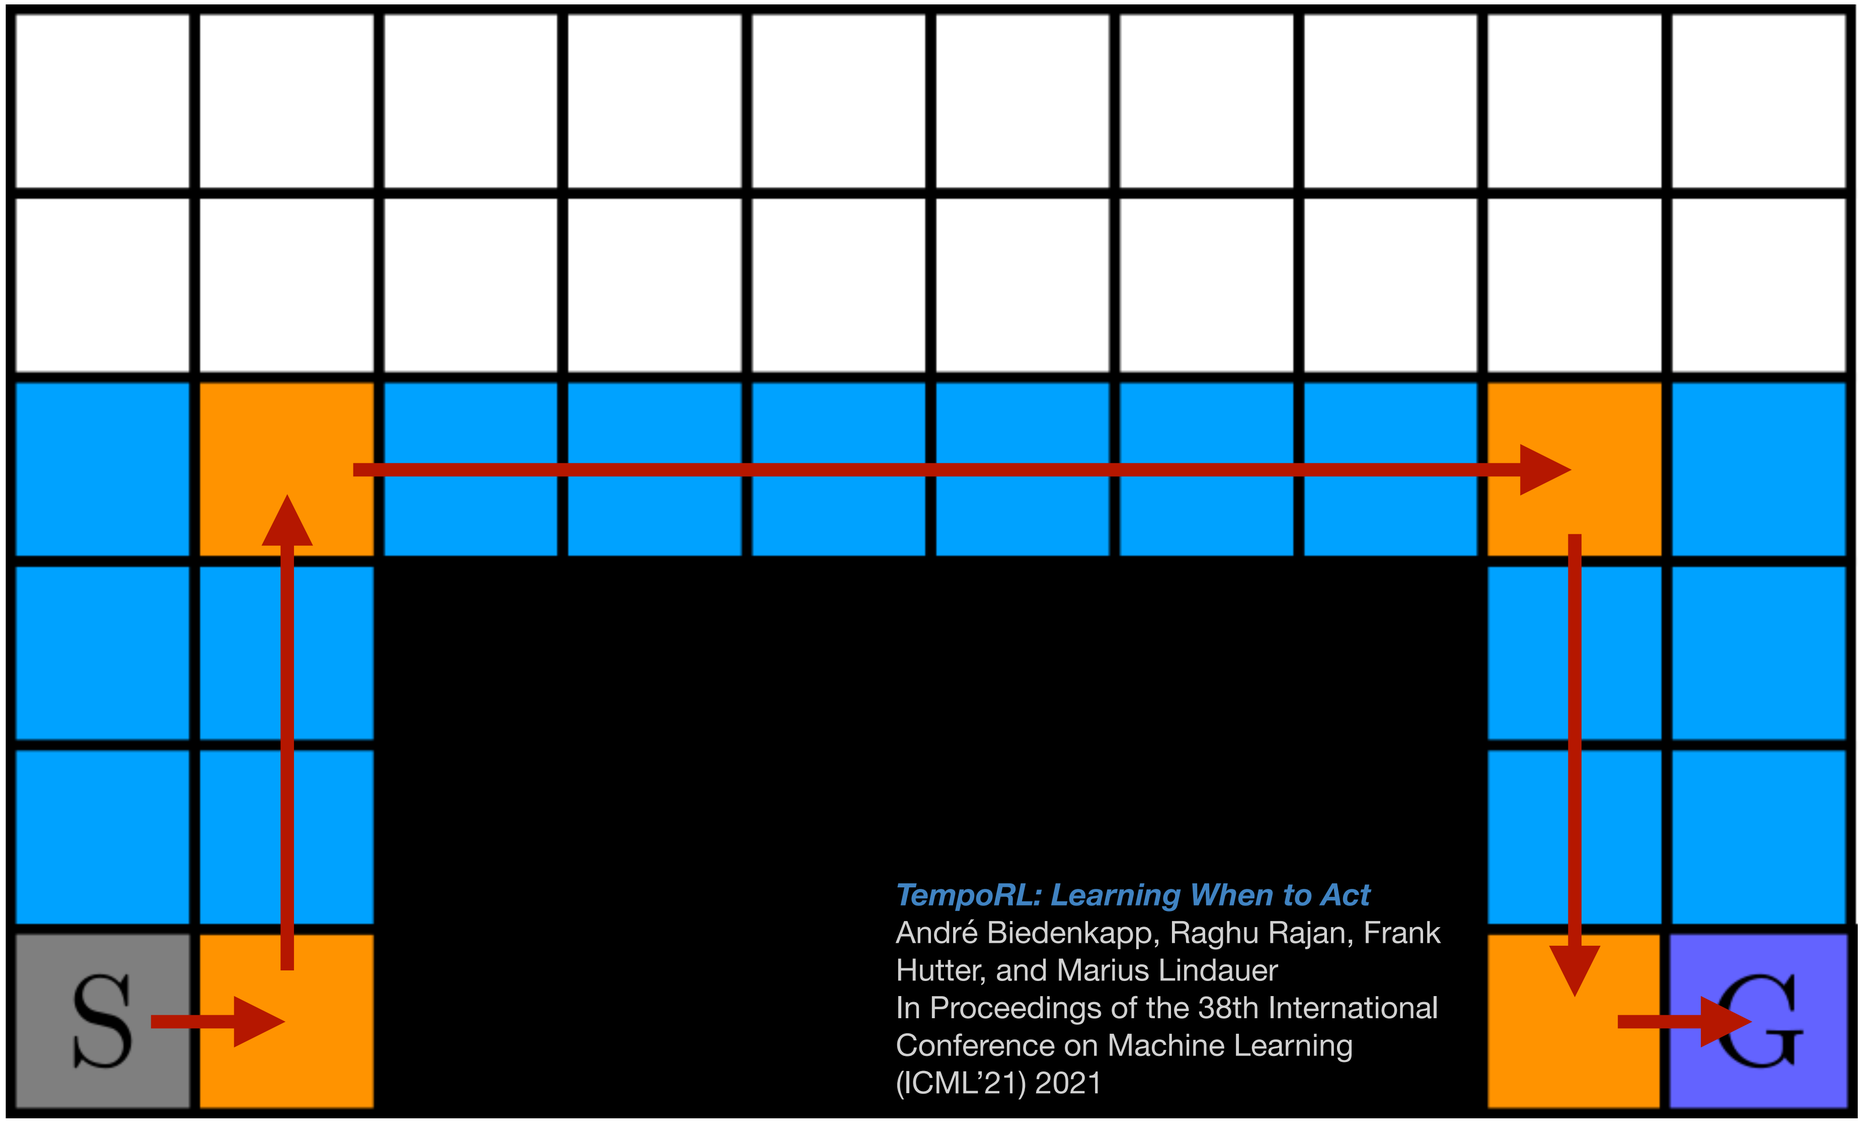

TempoRL: Learning When to Act
André Biedenkapp, Raghu Rajan, Frank Hutter, and Marius Lindauer
In Proceedings of the 38th International Conference on Machine Learning (ICML'21) 2021

TempoRL - Tabular Q-Learning

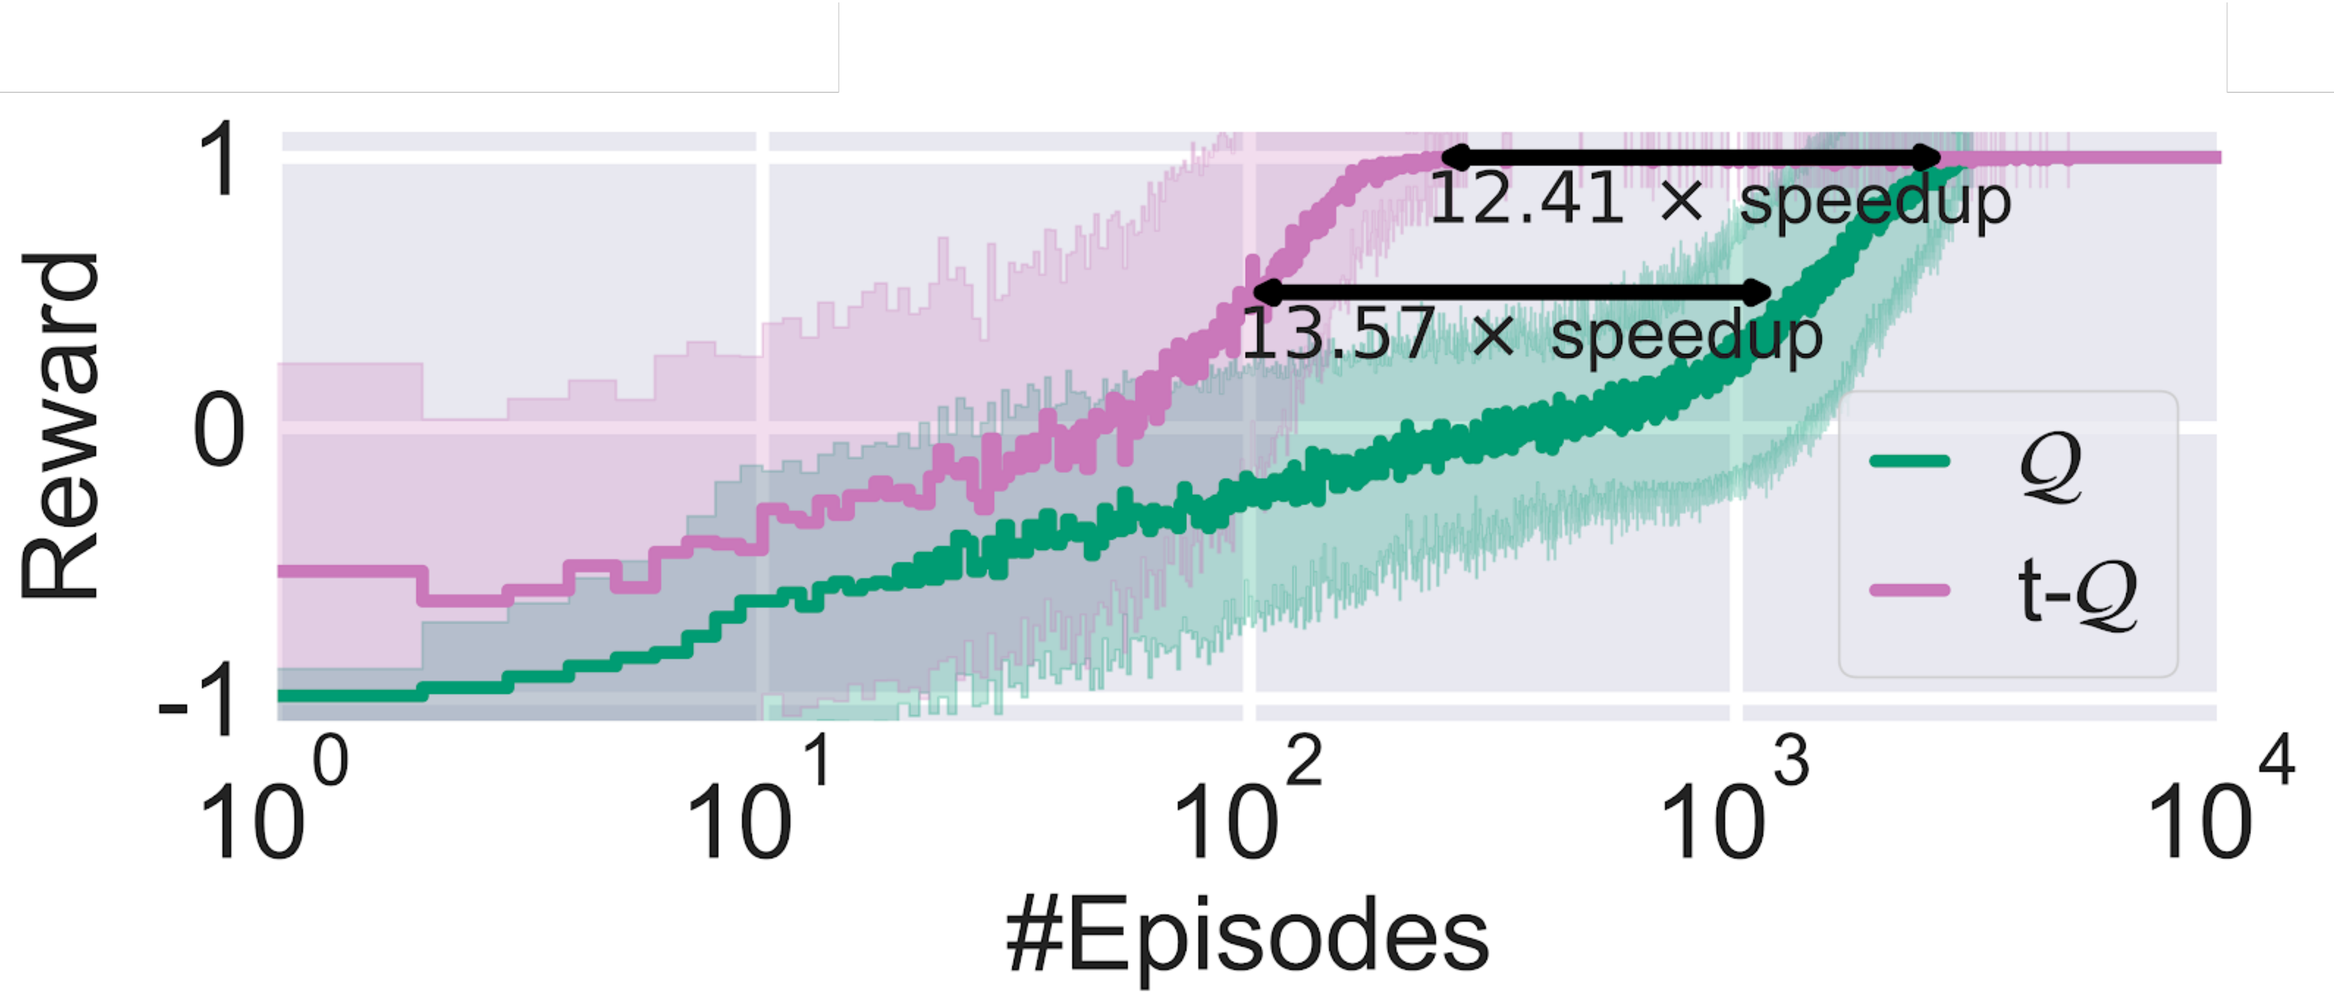

Deep TempoRL - Atari

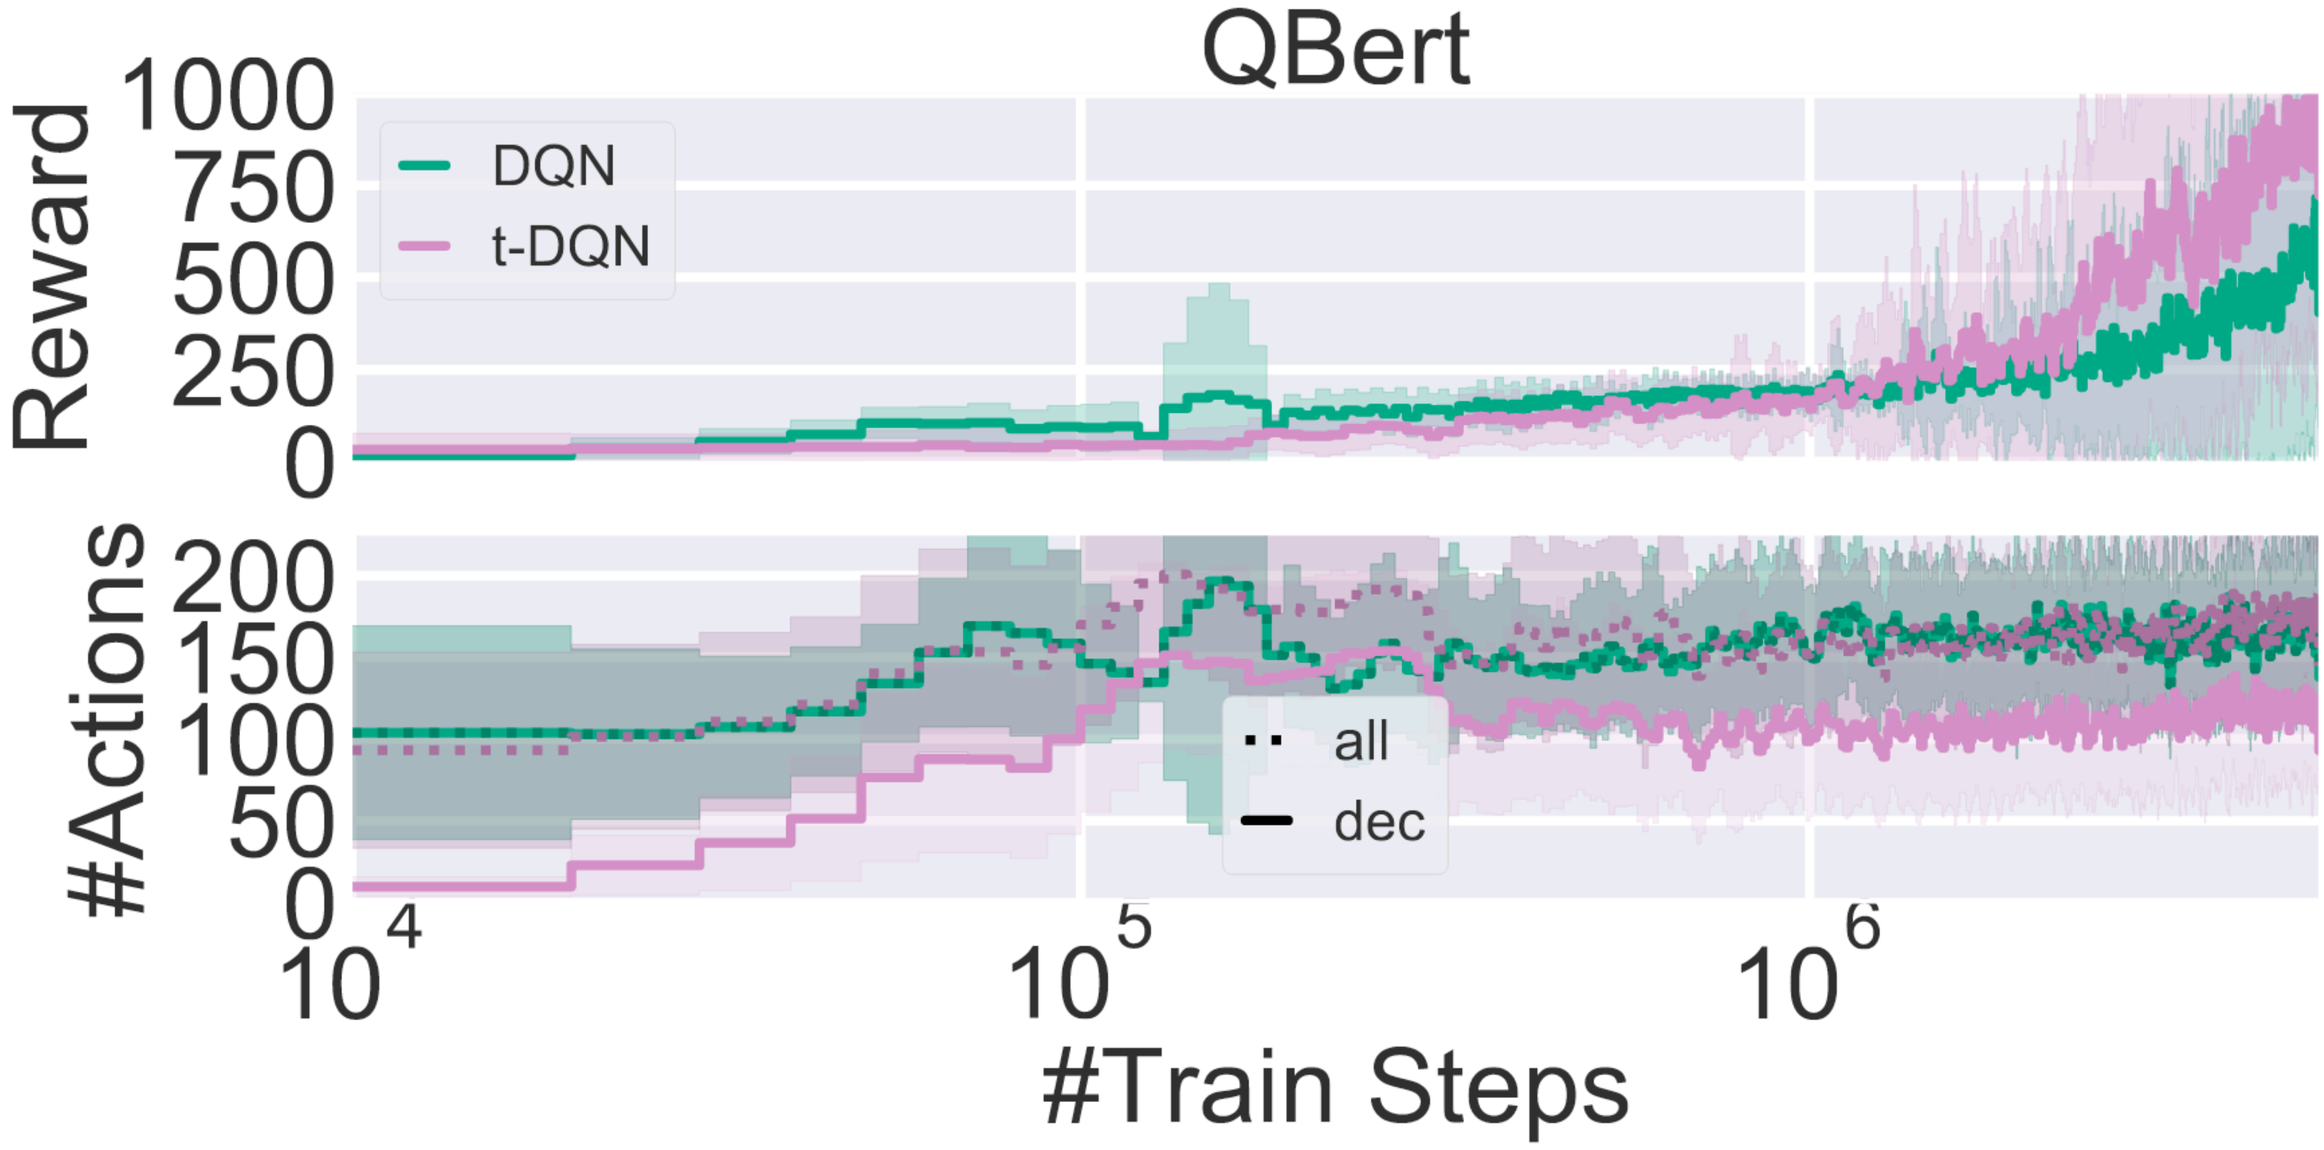

Deep TempoRL - Atari

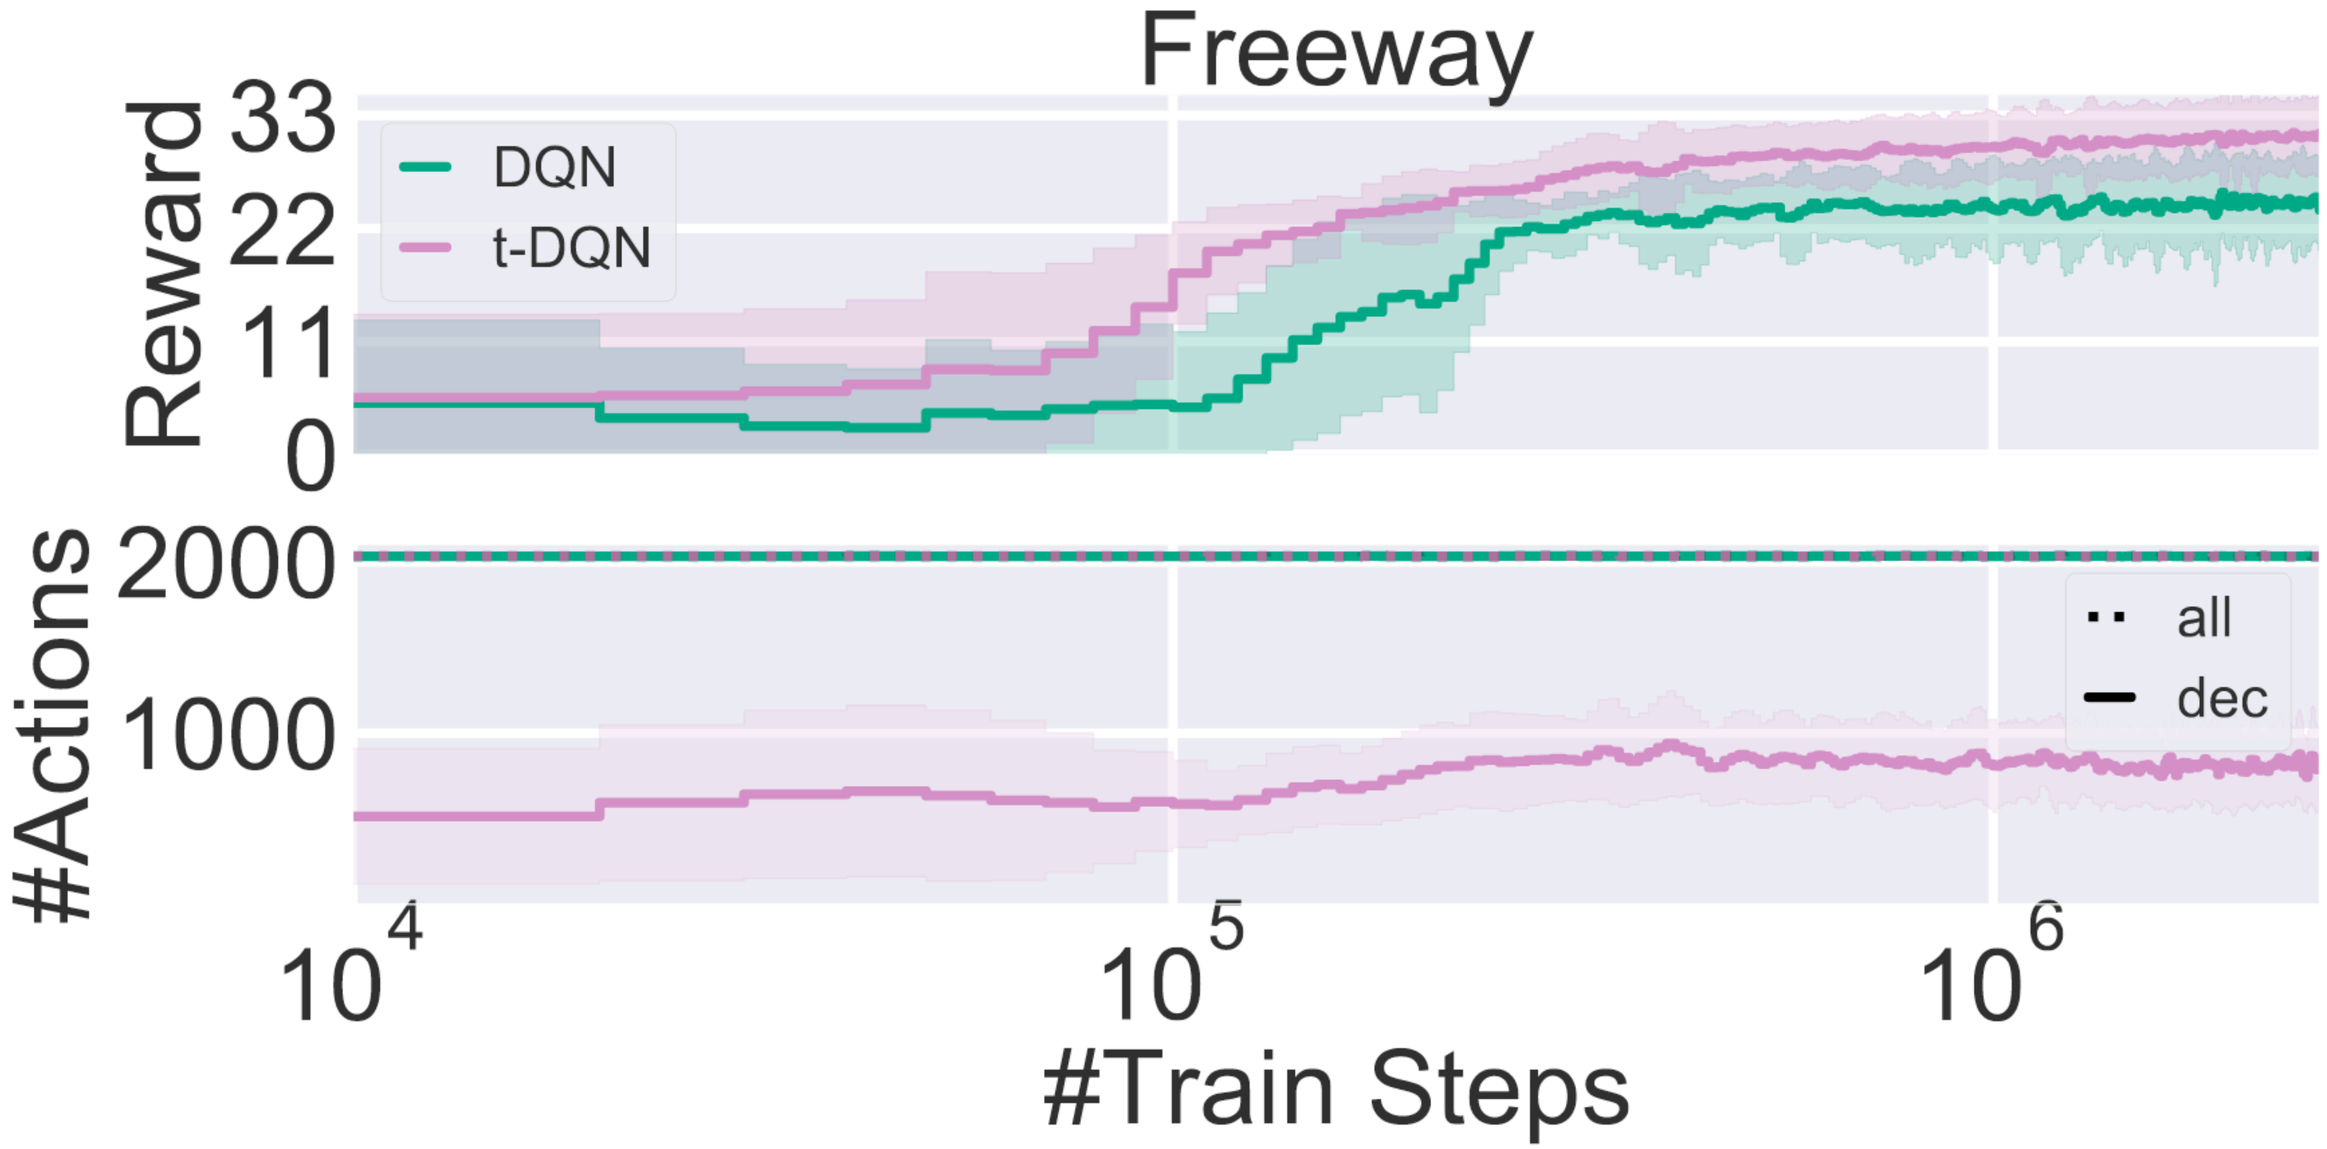

Future Work

Warmstarting and Learning From Handcrafted Policies

Future Work

Warmstarting and Learning From Handcrafted Policies

DAC for AutoRL

Future Work

Warmstarting and Learning From Handcrafted Policies

DAC for AutoRL

Better Methods for Contextual Reinforcement Learning

AutoML & Meta-Algorithmics Venues

AutoML Conf 2023

AutoML & Meta-Algorithmics Venues

AutoML Conf 2023

COSEAL Workshop 2023

AutoML & Meta-Algorithmics Venues

AutoML Conf 2023

COSEAL Workshop 2023

AutoML Fall School 2023

I'm looking forward to your questions