Northeastern University
Khoury College of Computer Sciences

Using Geometric Structure In Robot Learning and Planning

Linfeng Zhao
Khoury College of Computer Sciences, Northeastern University
2024/01 Tsinghua University

Useful Geometric Structure

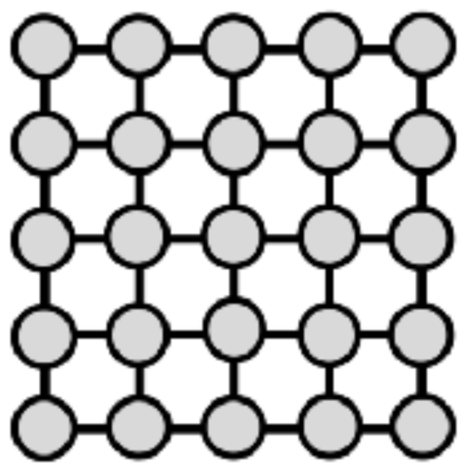

2D Grid

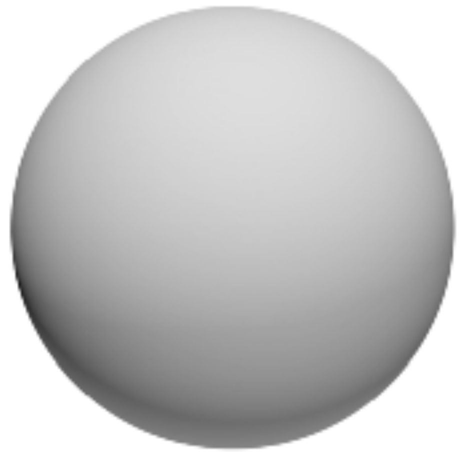

Sphere

Mathematical Space

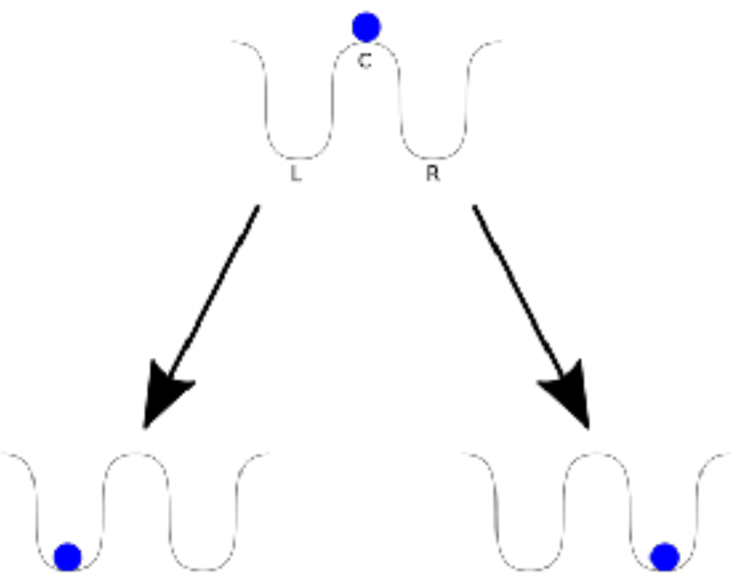

Physical System

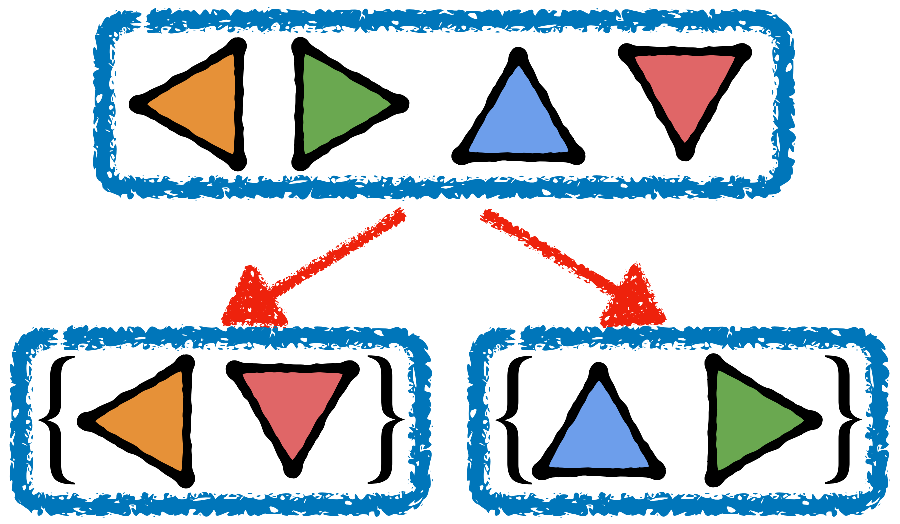

Object Interaction

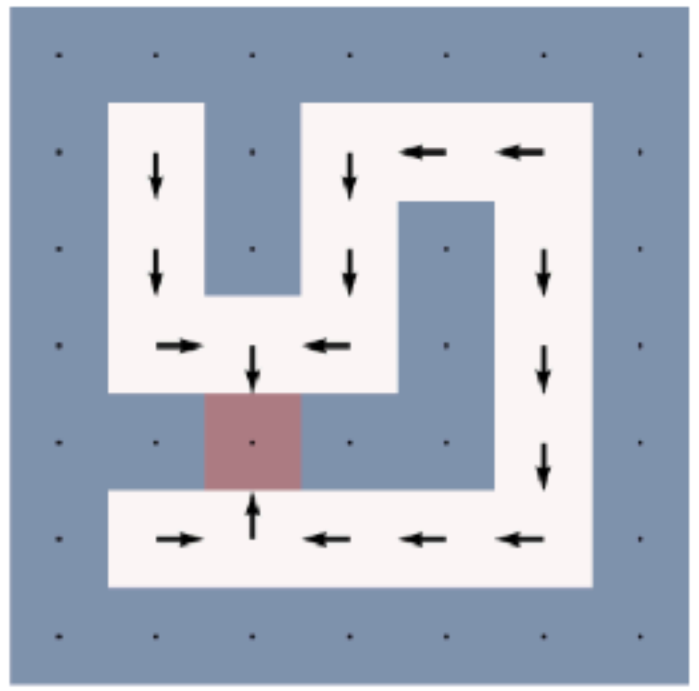

Learning / Planning

- Robots operate in physical spaces
- They have geometric structure that sometimes are neglected
- Resulting less efficiency, generalizability, or even failure

Task and Challenges

Target: Long-horizon Mobile Manipulation — Requires combined *learning* and *planning*

Challenges:

- *Efficiency, Generalization, Scalability in Learning*
- *Learning Desirable Representation for Planning*

Geometric Structures Could Help

- Use e.g,. *symmetry* and *compositionality* of the tasks
- May *reduce number of free parameters* and *solution space*
- Result in better *efficiency, generalization, scalability, …*

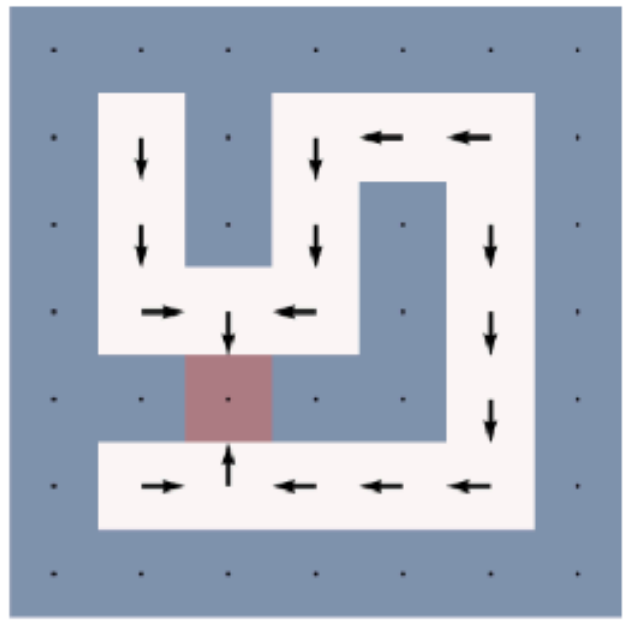

2D Discrete Map Rotation/

2D Discrete Symmetry D_4

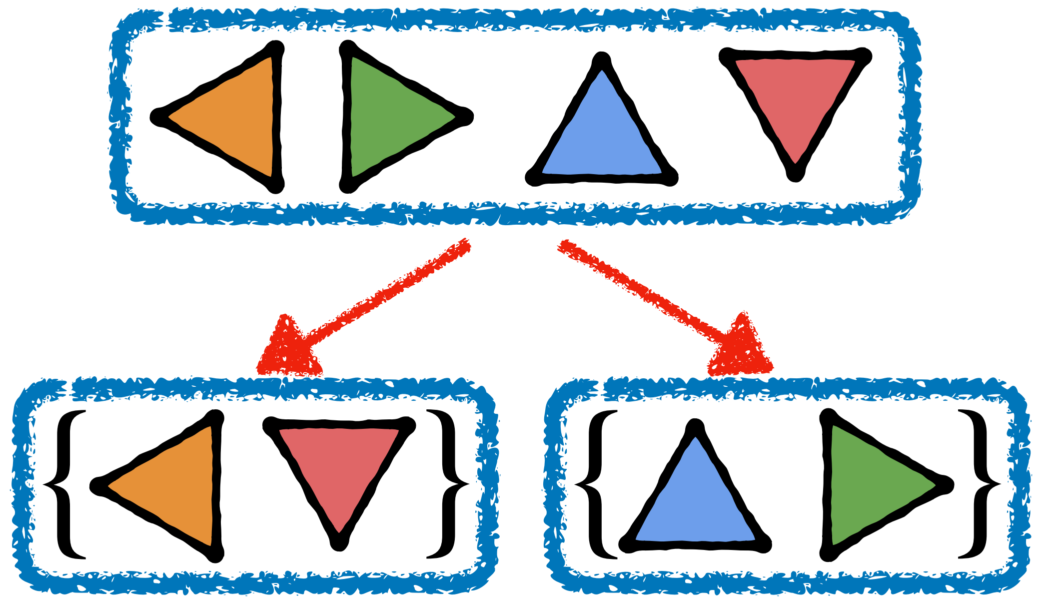

Object Interchangeability

Permutation Symmetry S_N

Negative Examples

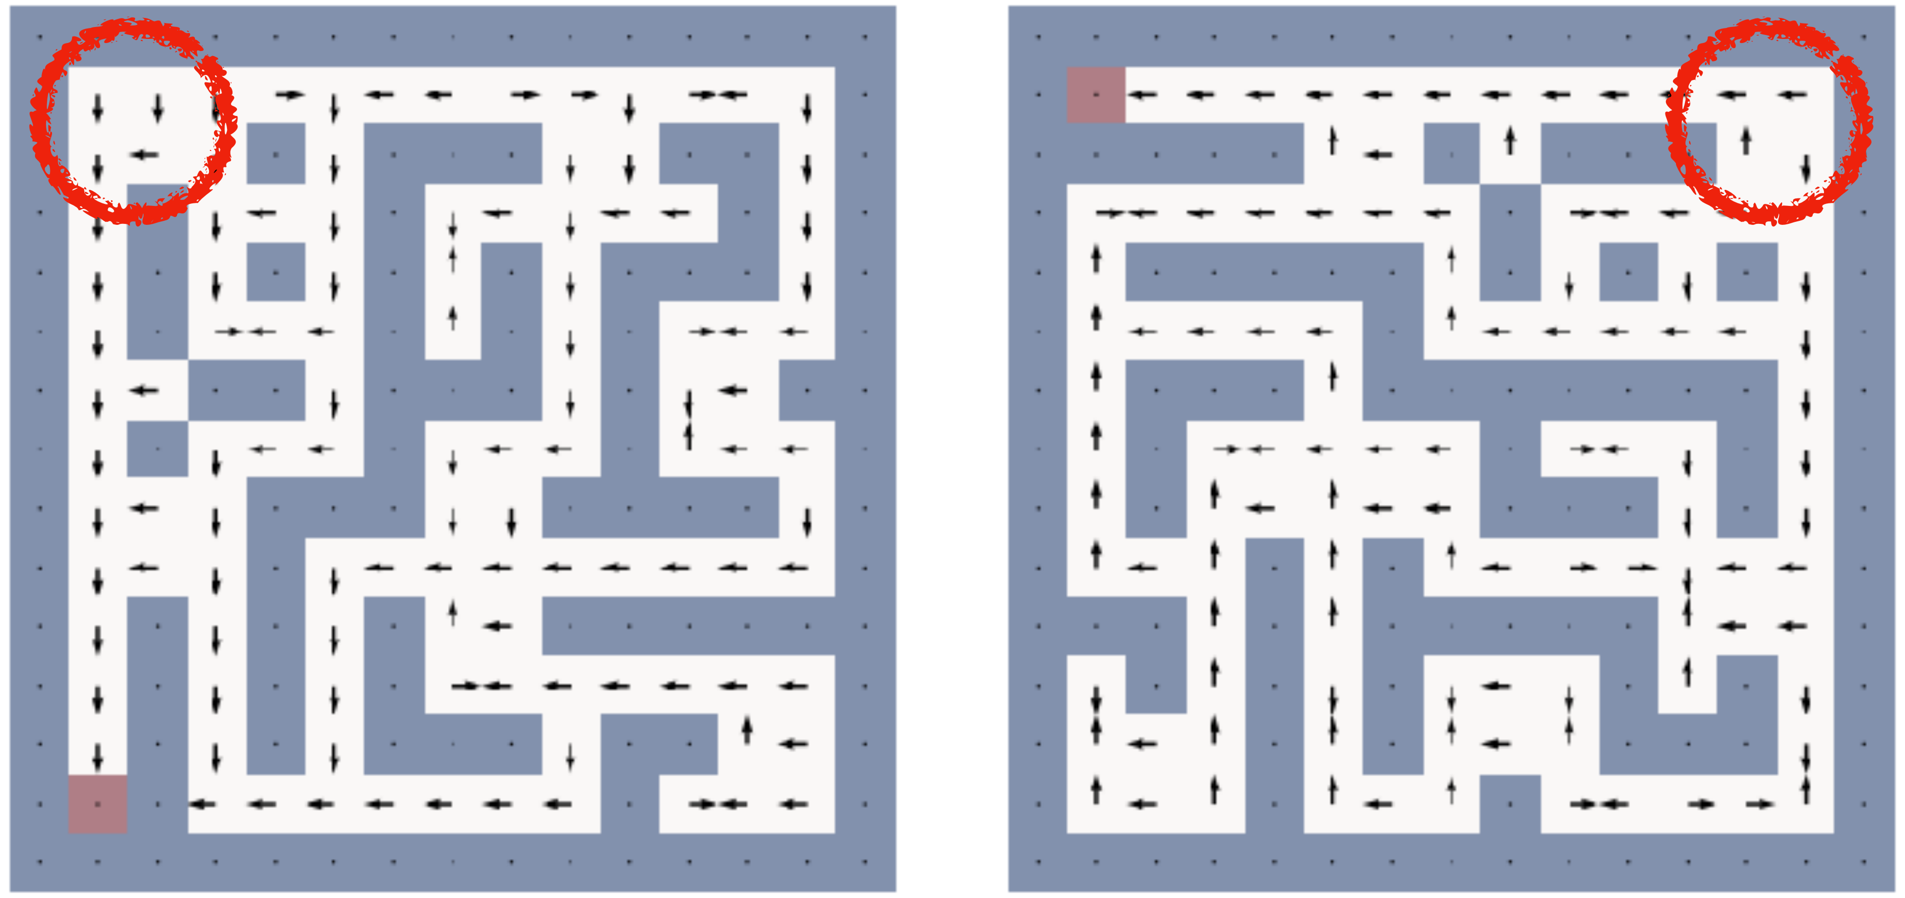

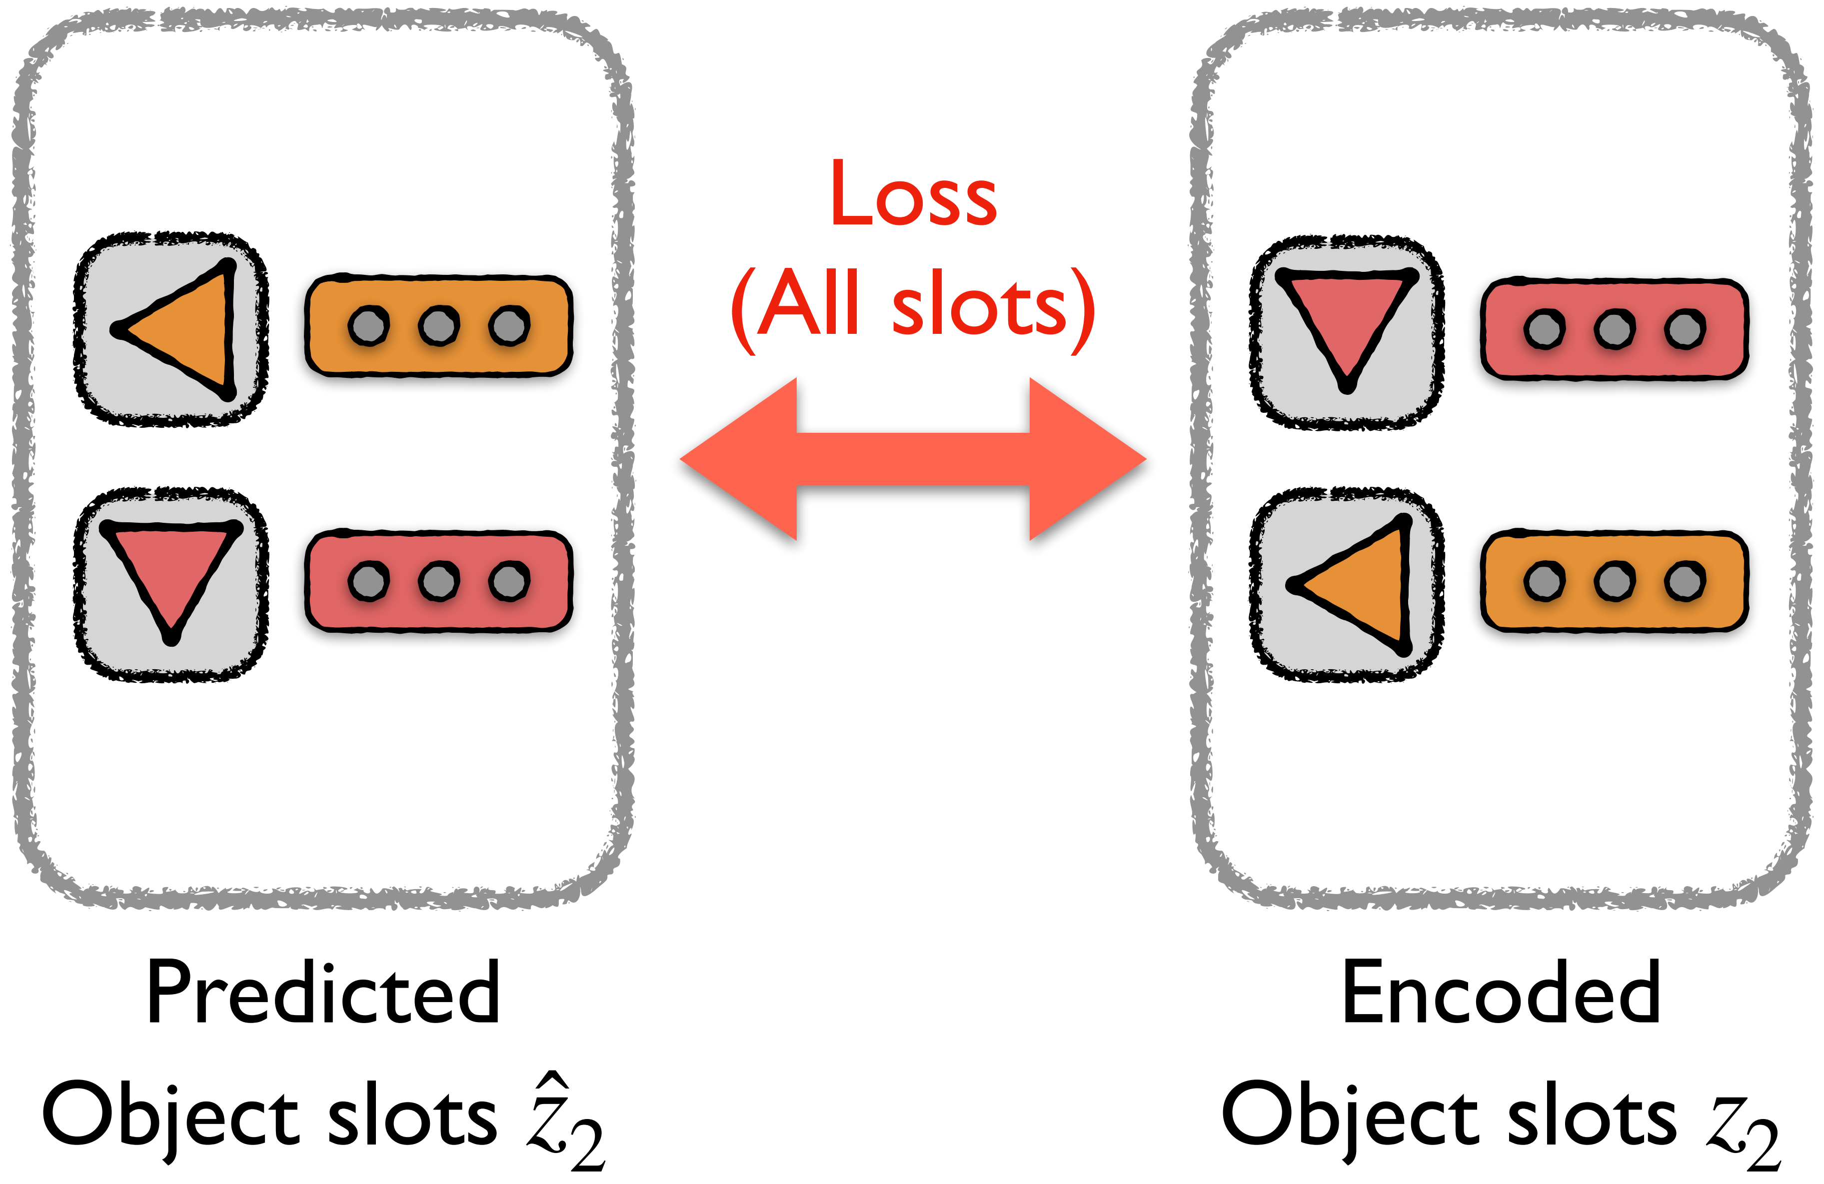

Ignoring rotation symmetry in planning results in inconsistency

Ignoring object interchangeability in world modeling results in misalignment

Symmetry and Equivariance

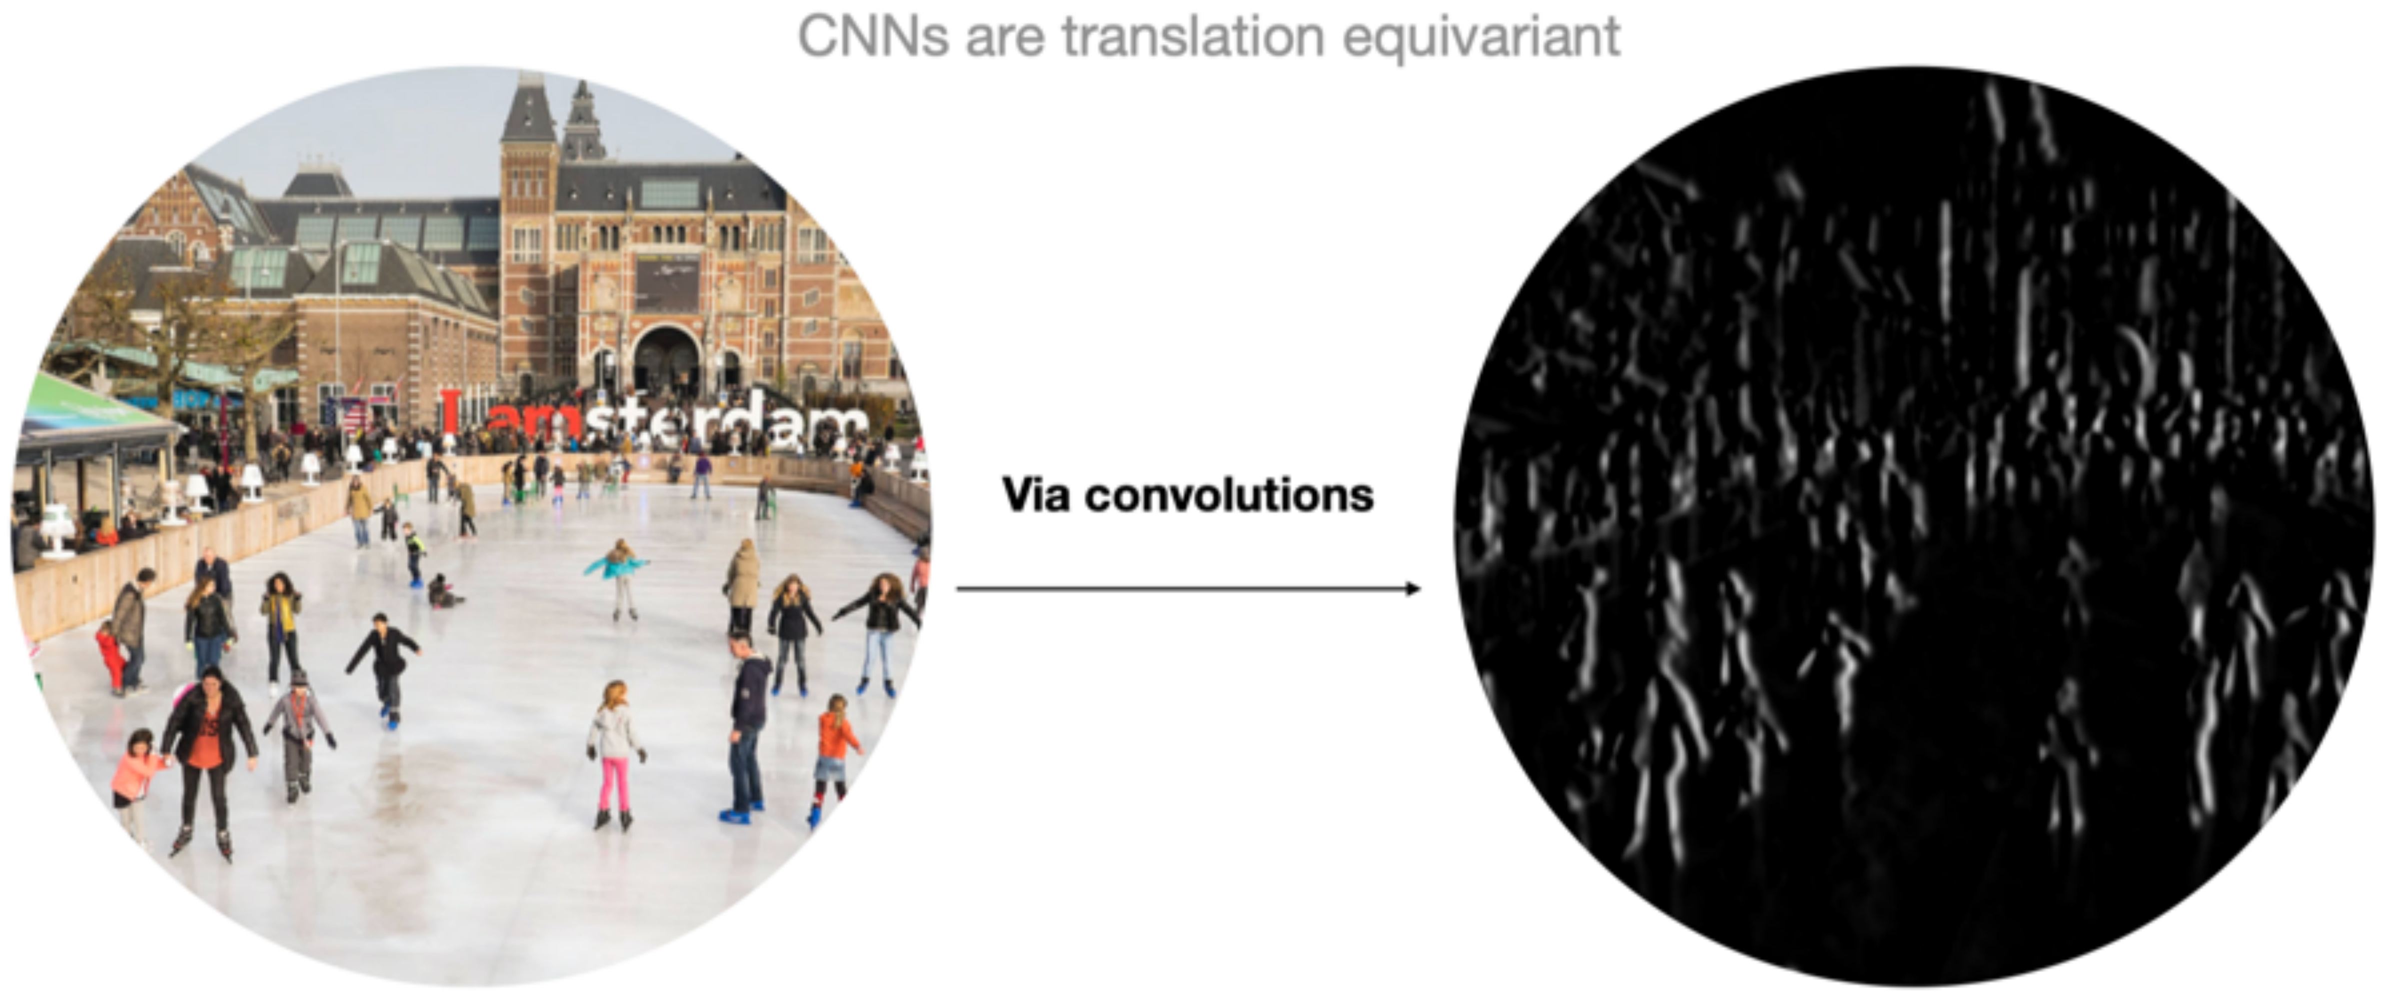

- Object segmentation task has symmetry: moving/rotating objects
- 2D Convolution Networks are translation equivariant by design

[Credit: UvA Group Equivariant NNs lecture; https://github.com/QUVA-Lab/escnn; Geometric Deep Learning, Bronstein et al. 2021]

Additional Symmetry: Rotation

Normal CNN

Rotation-equivariant CNN

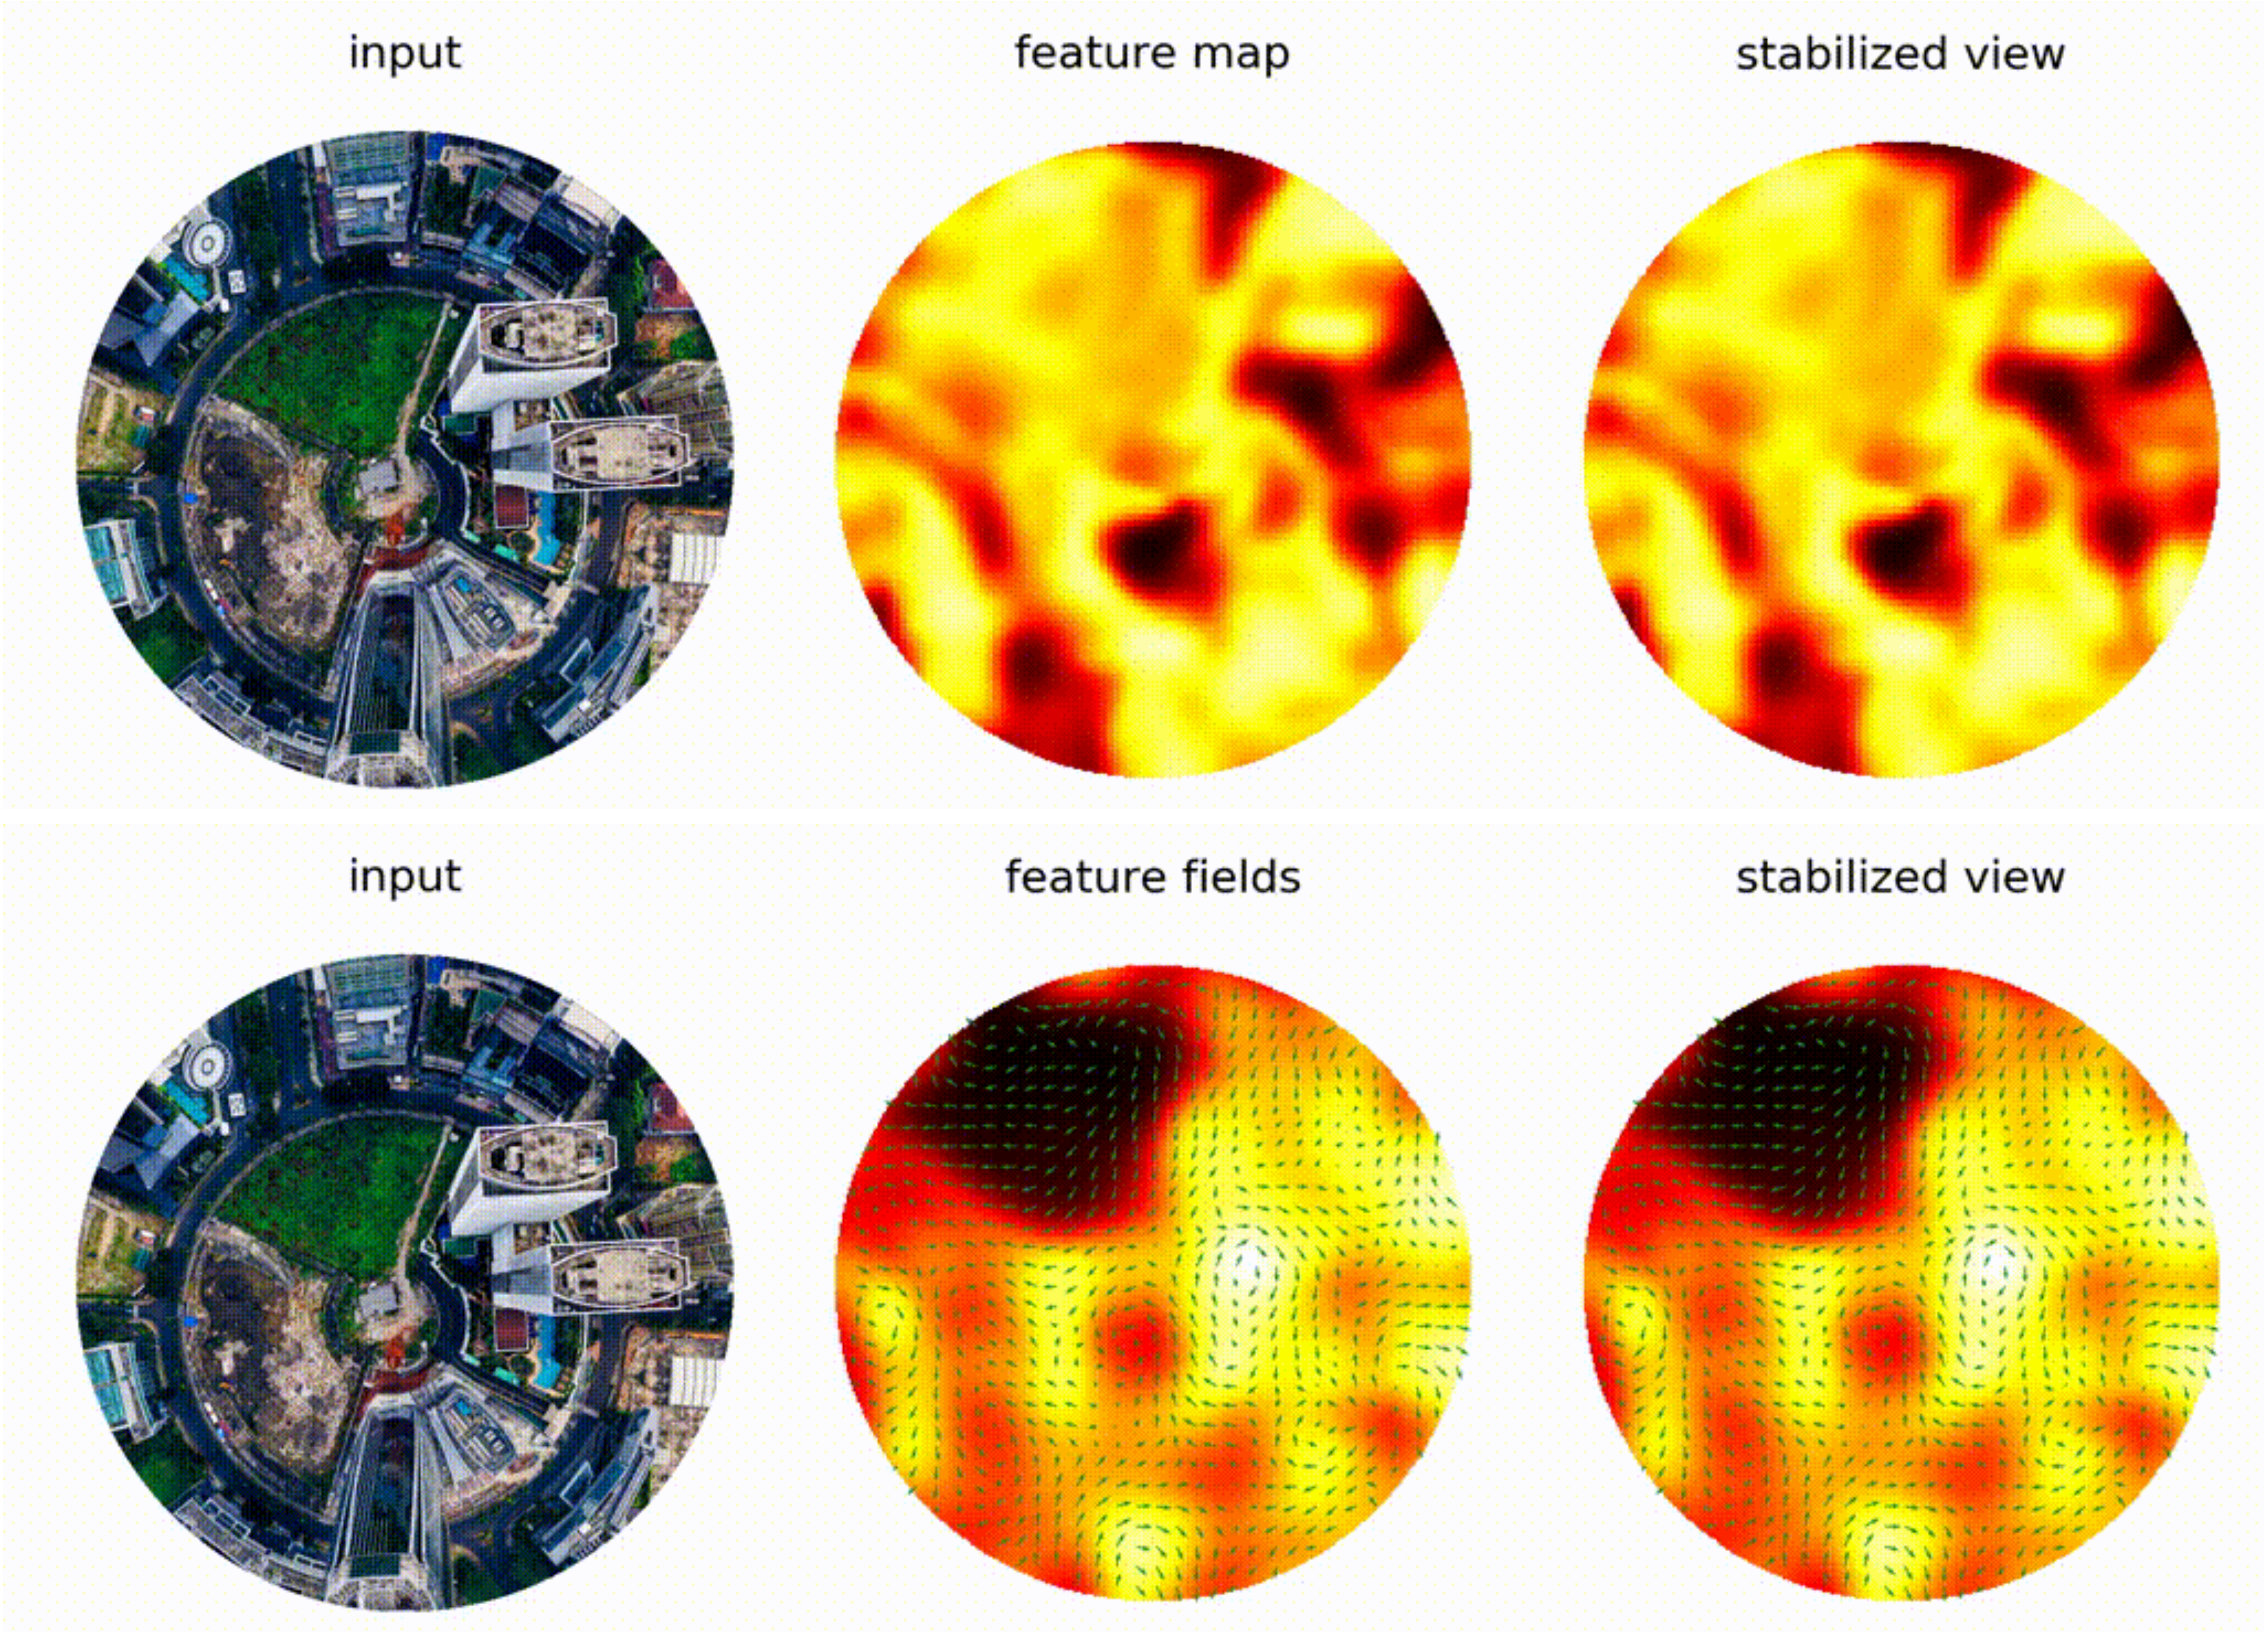

[Credit: https://github.com/QUVA-Lab/escnn; Geometric Deep Learning, Bronstein et al. 2021]

Brief History of Equivariant NNs

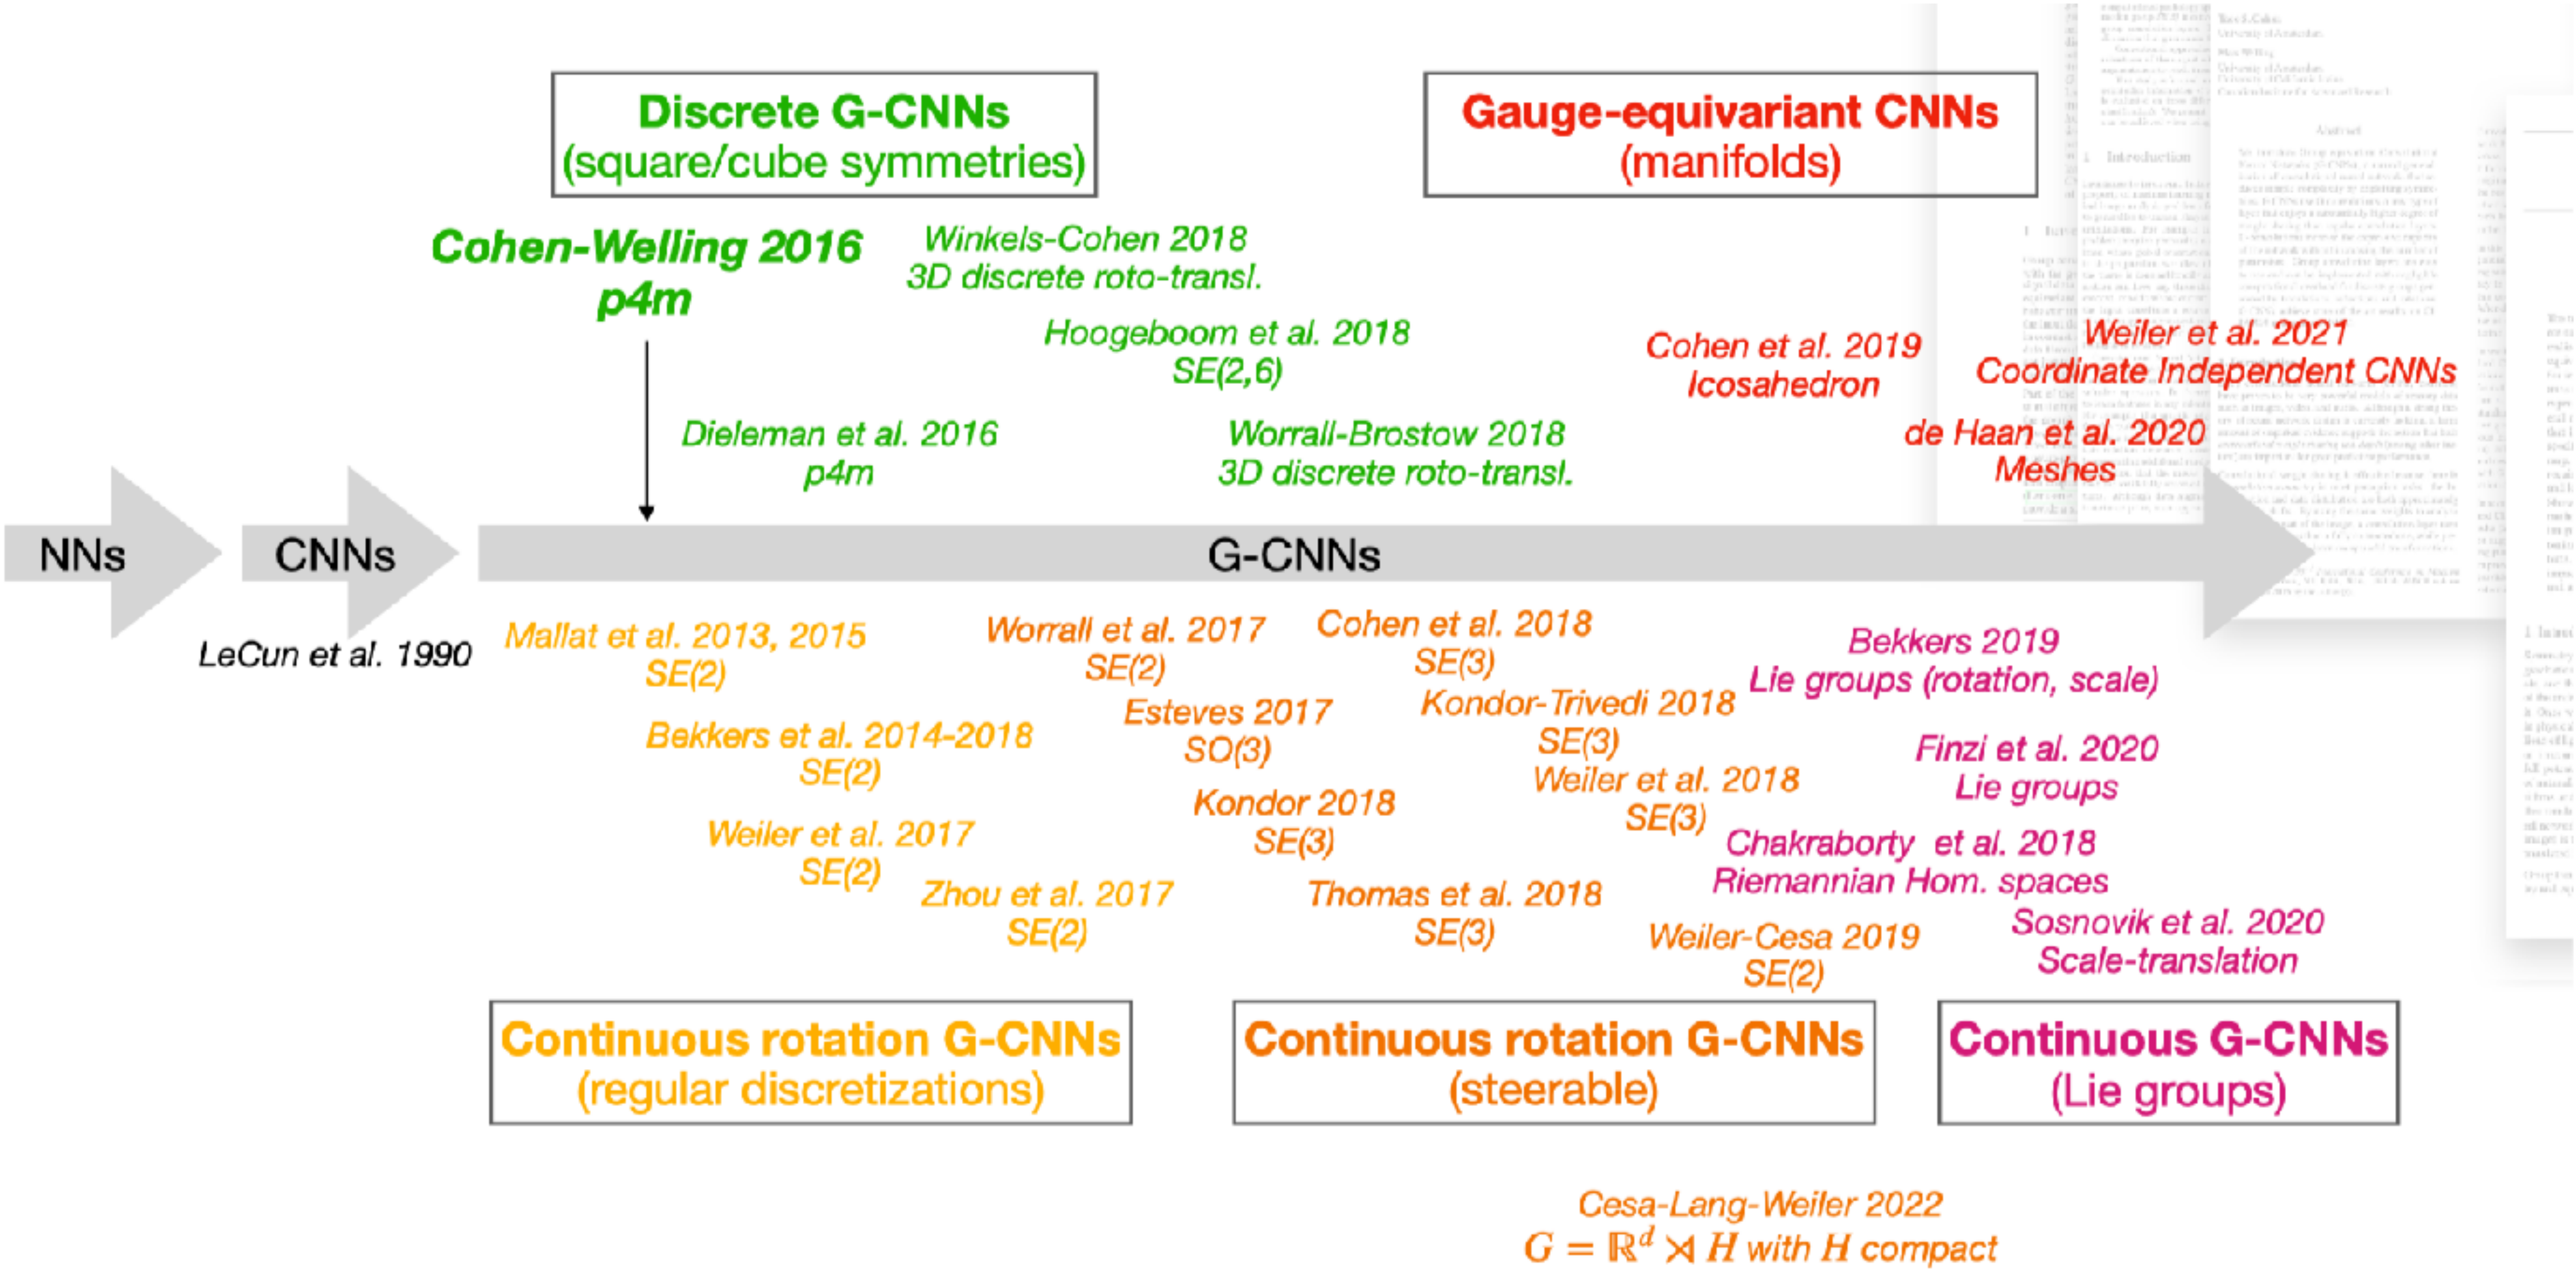

[Credit: UvA Group Equivariant NNs lecture]

This only covers Lie groups, but not e.g,. GNN and permutation groups.

This talk

Motivating what component needs geometric structure: Learning + Planning

Identifying Geometric Structure in *Tasks*

Encoding Geometric Structure into *Algorithms*

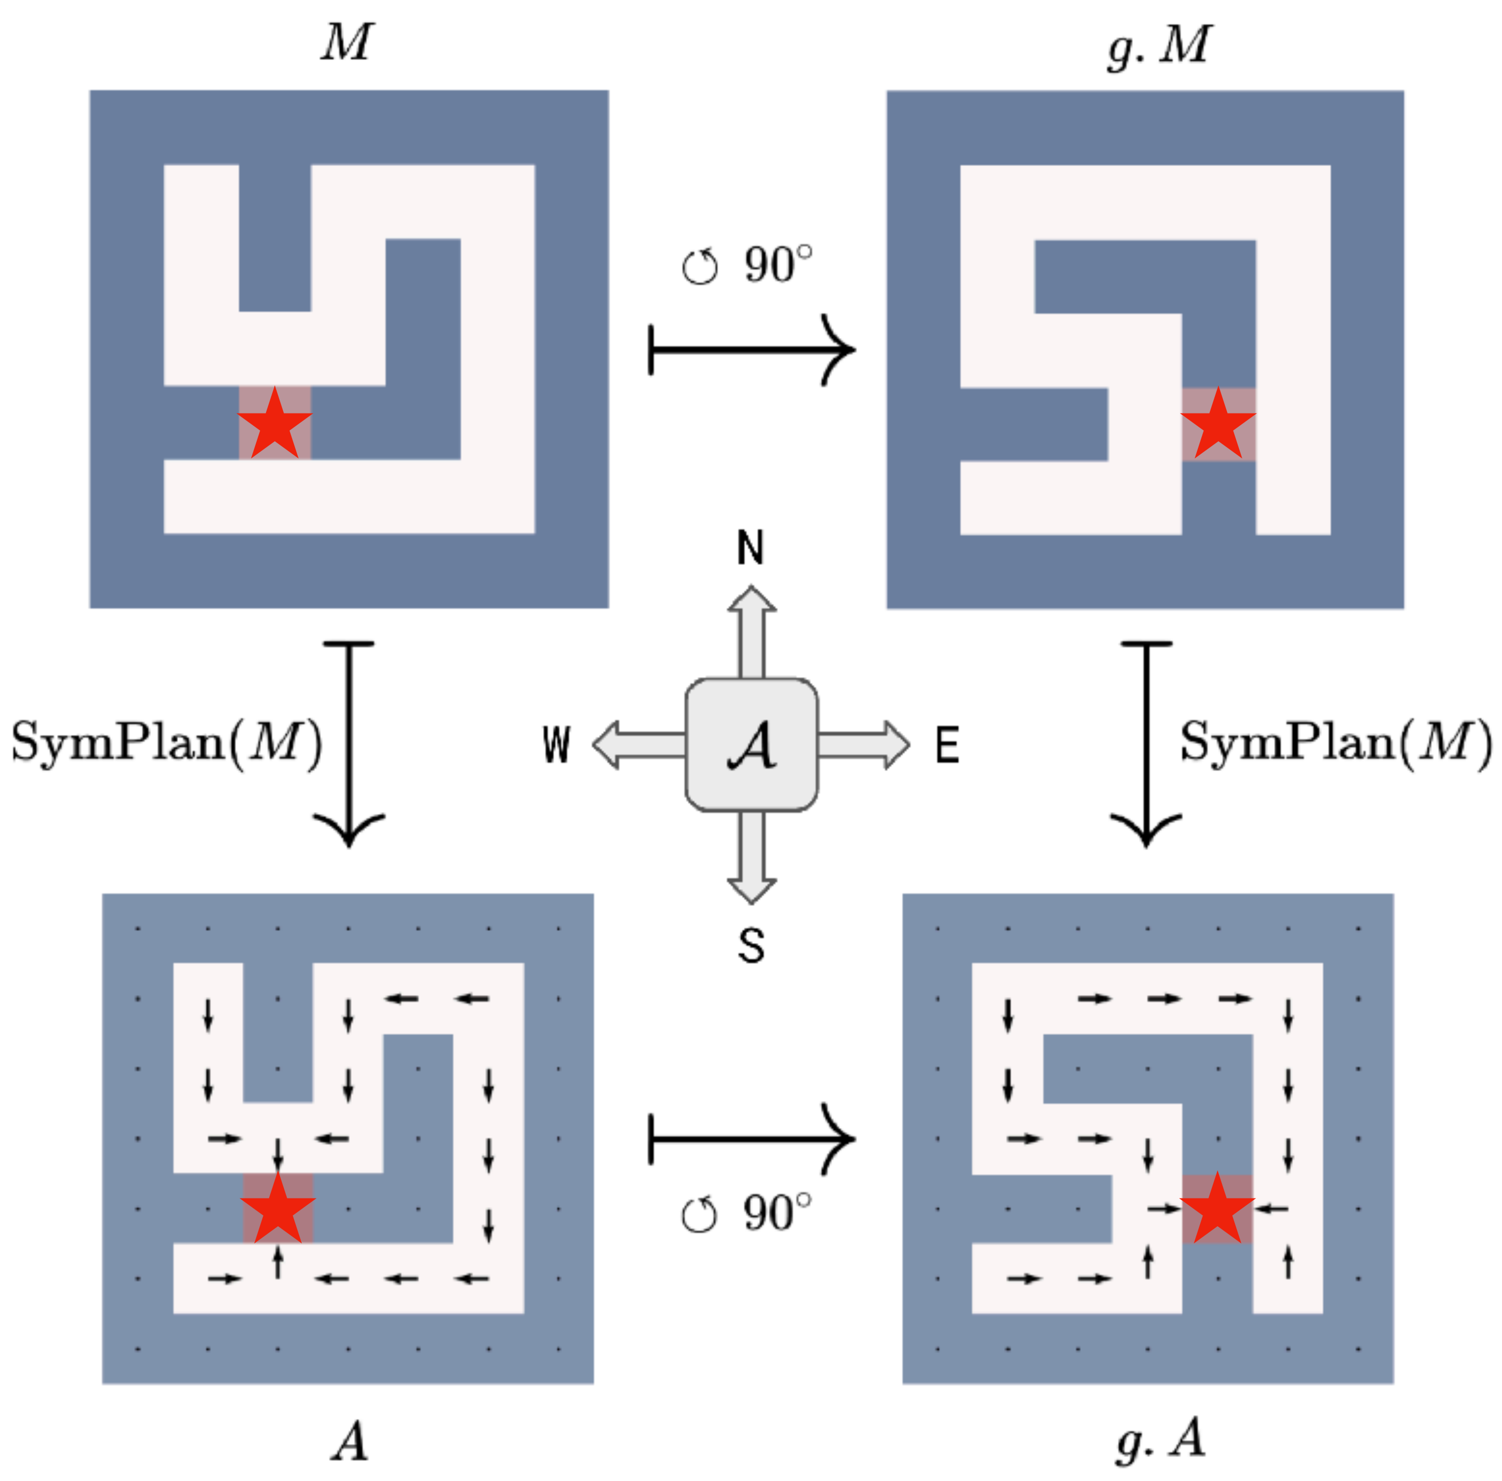

Topics

- *Learning for* End-to-end Planning
  — How to learn to plan and improve stability
- *Symmetric (and Compositional)* Structure in Robot Planning
  — How does symmetry improve (path) planning and world modeling
- *Symmetric* Structure in Robot Learning
  — How could symmetry improve learning robot skills

Grey items are secondary

Topics

- *Learning for* End-to-end Planning
  — How to learn to plan and improve stability
- *Symmetric* *(and Compositional)* Structure in Robot Planning
  — How does symmetry improve (path) planning and world modeling
- *Symmetric* Structure in Robot Learning
  — How could symmetry improve learning robot skills

Grey items are secondary

What is planning?

- A (learned) state space $\mathcal{S}$
- A (learned) transition model $f(s, a) = s'$ that gives next state
- Goal of planning
  - 1, Maximize reward/utility function, or minimize cost
  - 2, Reach goal region

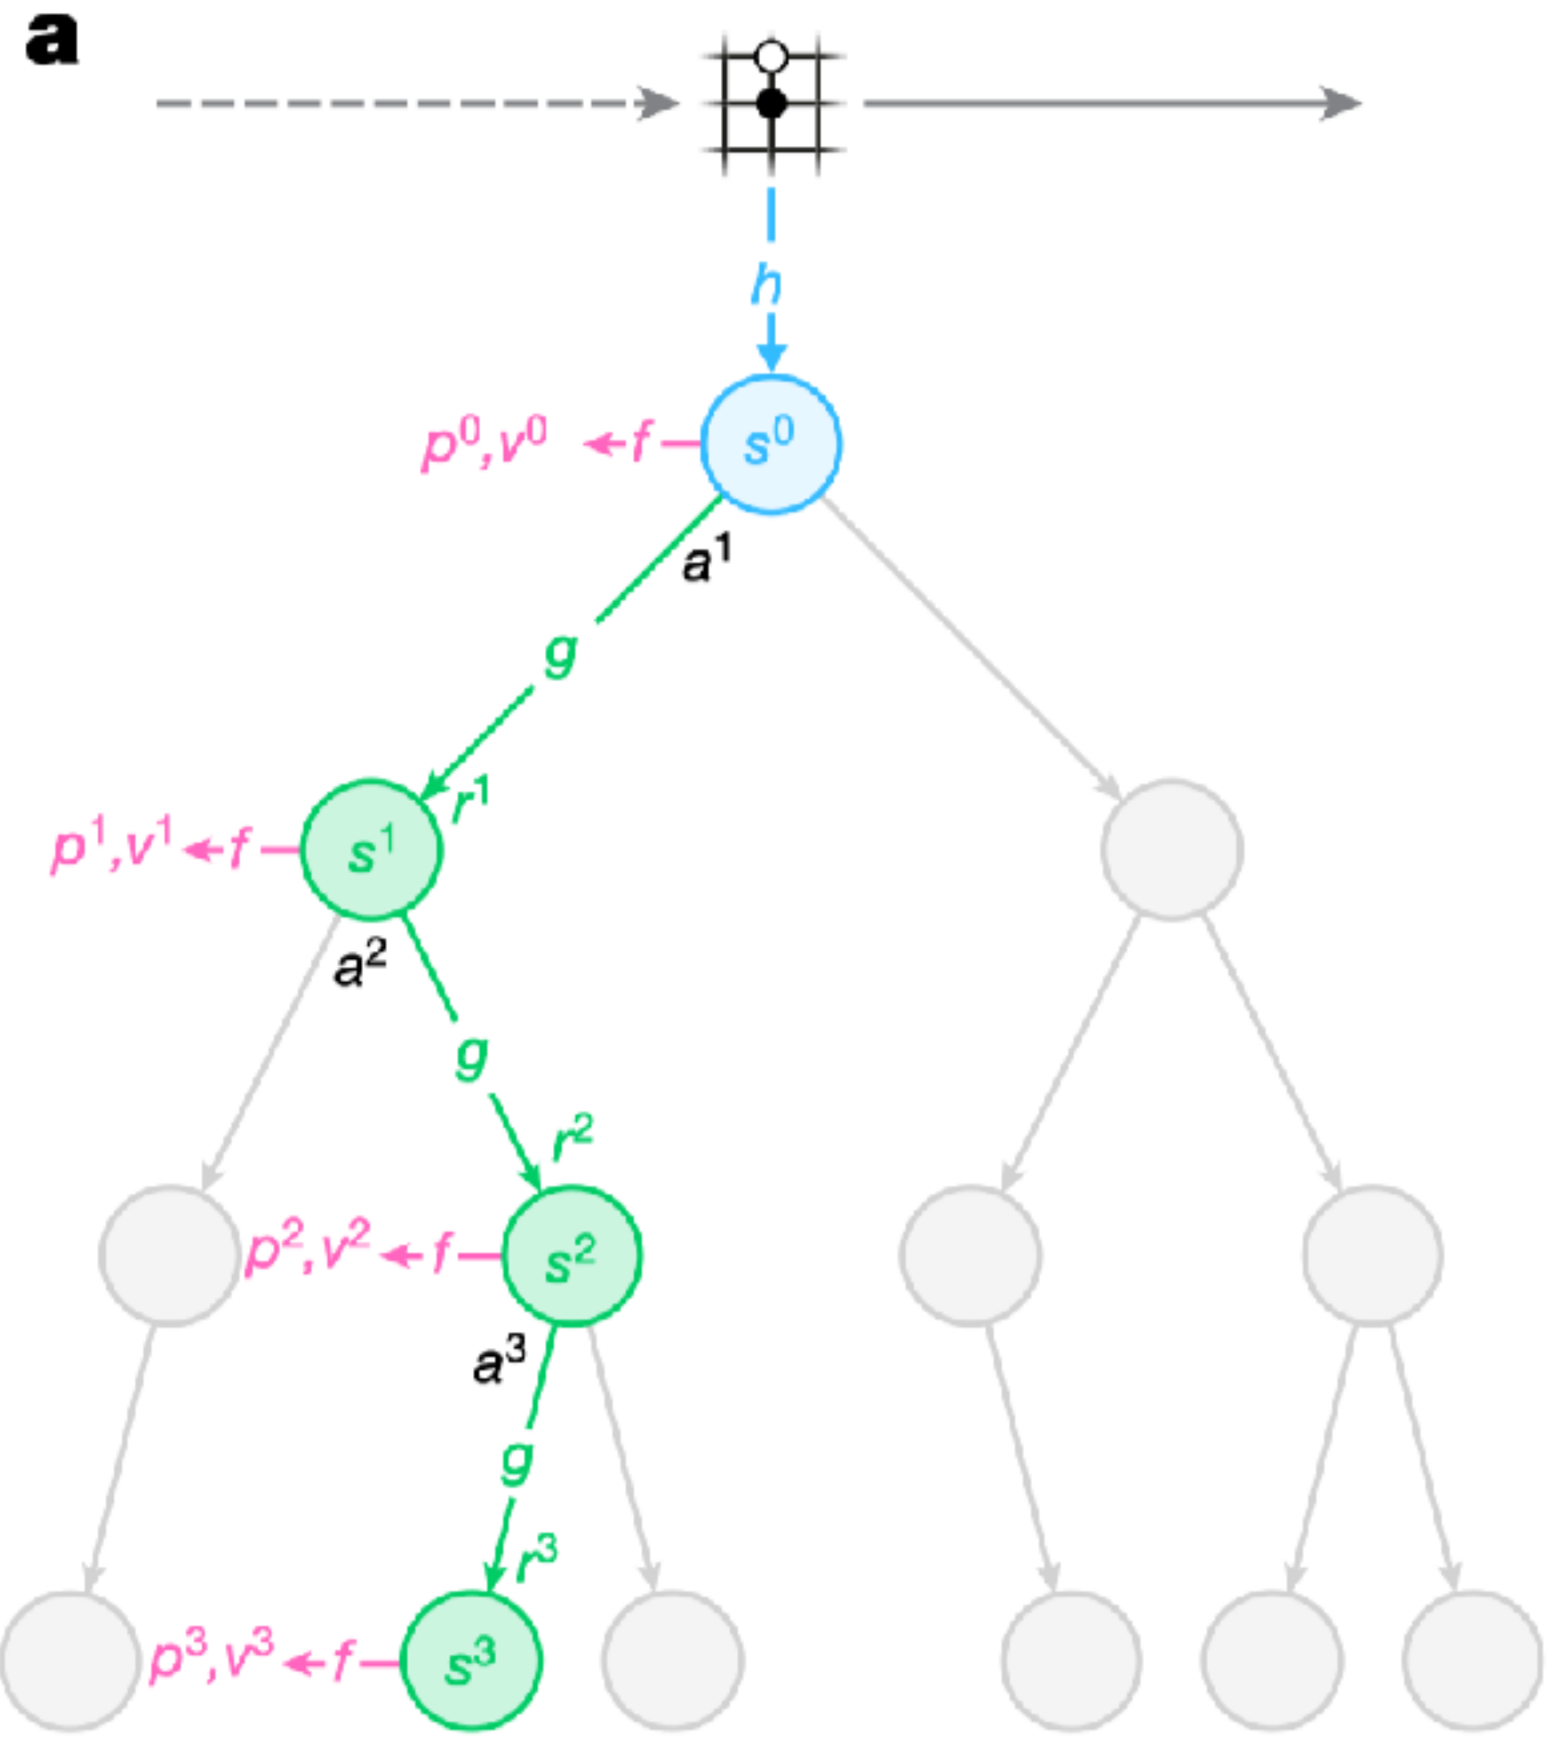

[Credit: MuZero, DeepMind]

Challenges in Learning + Planning

In complex tasks

- E.g., mobile manipulation, visual navigation
- Require e.g., long-horizon goal reaching

It is hard to simultaneously

- Define compact state space
- Obtain accurate predictive model
- Reliably reach far goals or sparse rewards

[Credit: Robohub Visual Nav]

The Need of “Learning to Plan”

One core challenge: How to *learn good representation* for *planning*

- Balance what is contained in states

Prior work: *VIN [Tamar et al. NIPS 2016], MuZero [Schrittwieser et al. Nature 2020], Value Equivalence Principle [Grimm et al. NeurIPS 2020]*

One Idea: *Learning from data using end-to-end architecture*

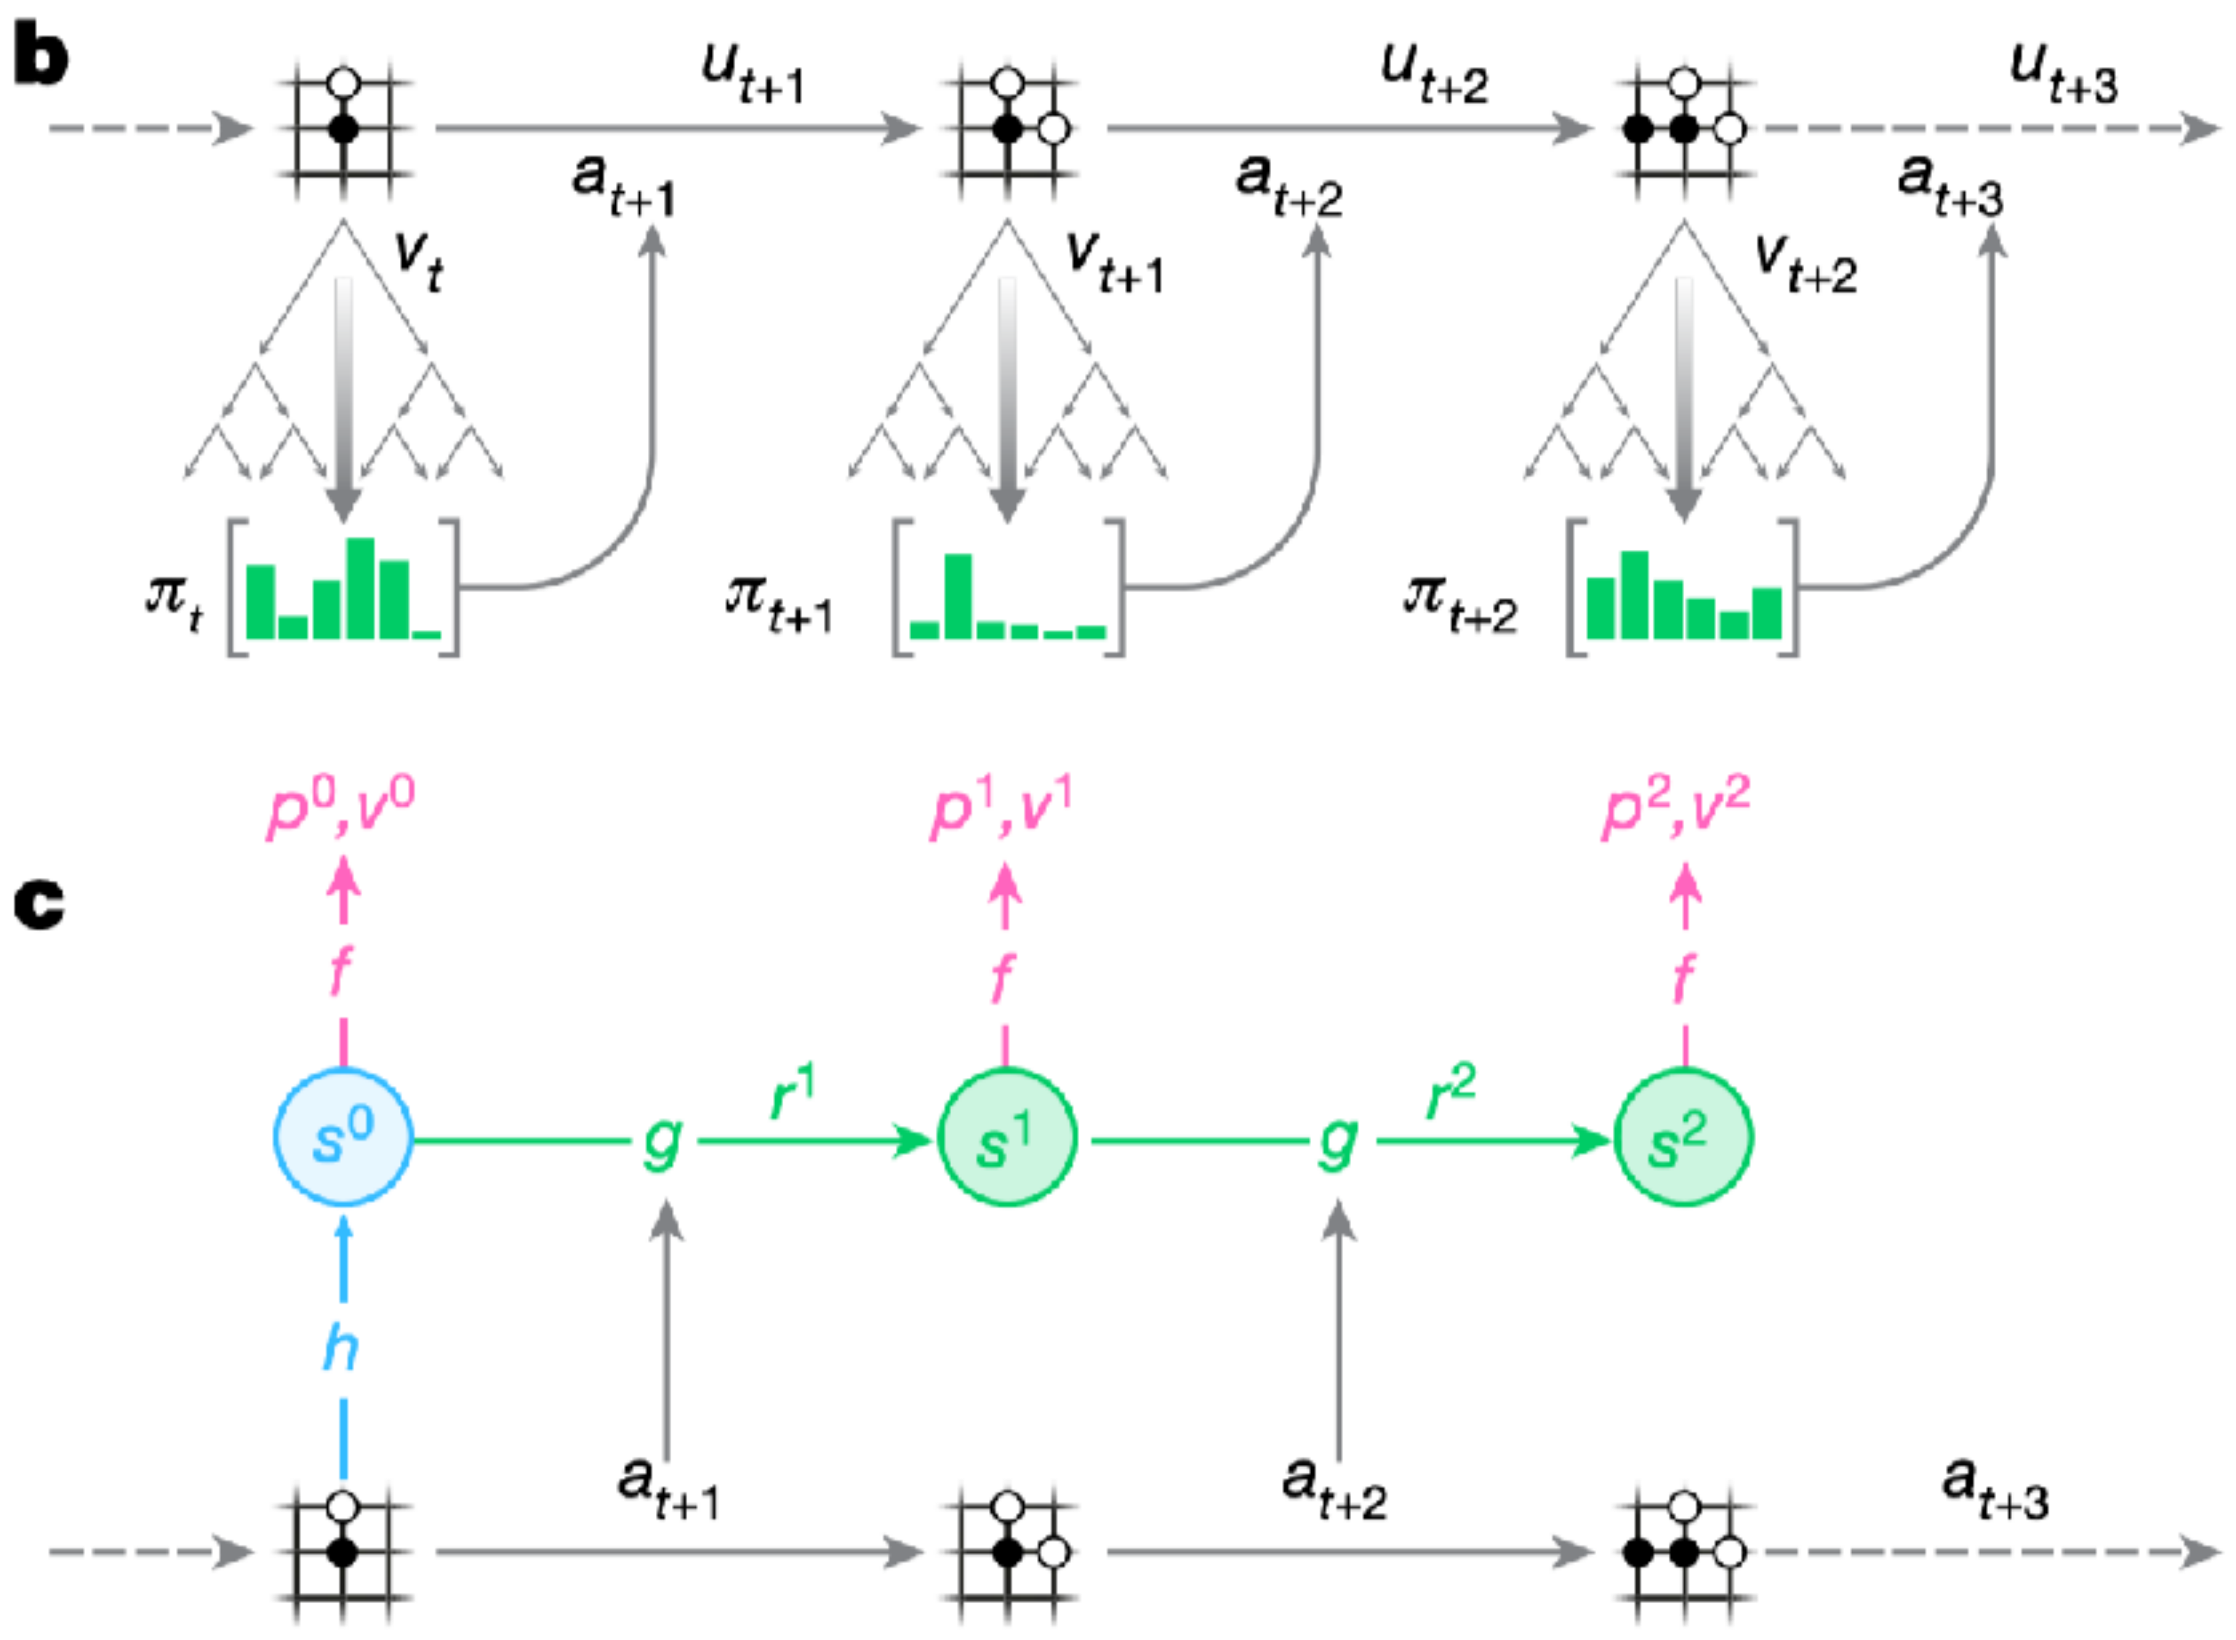

[Credit: MuZero, DeepMind]

Improving “Differentiable Planning”

Scaling up and Stabilizing Differentiable Planning with Implicit Differentiation

ICLR 2023

Linfeng Zhao, Huazhe Xu, Lawson Wong

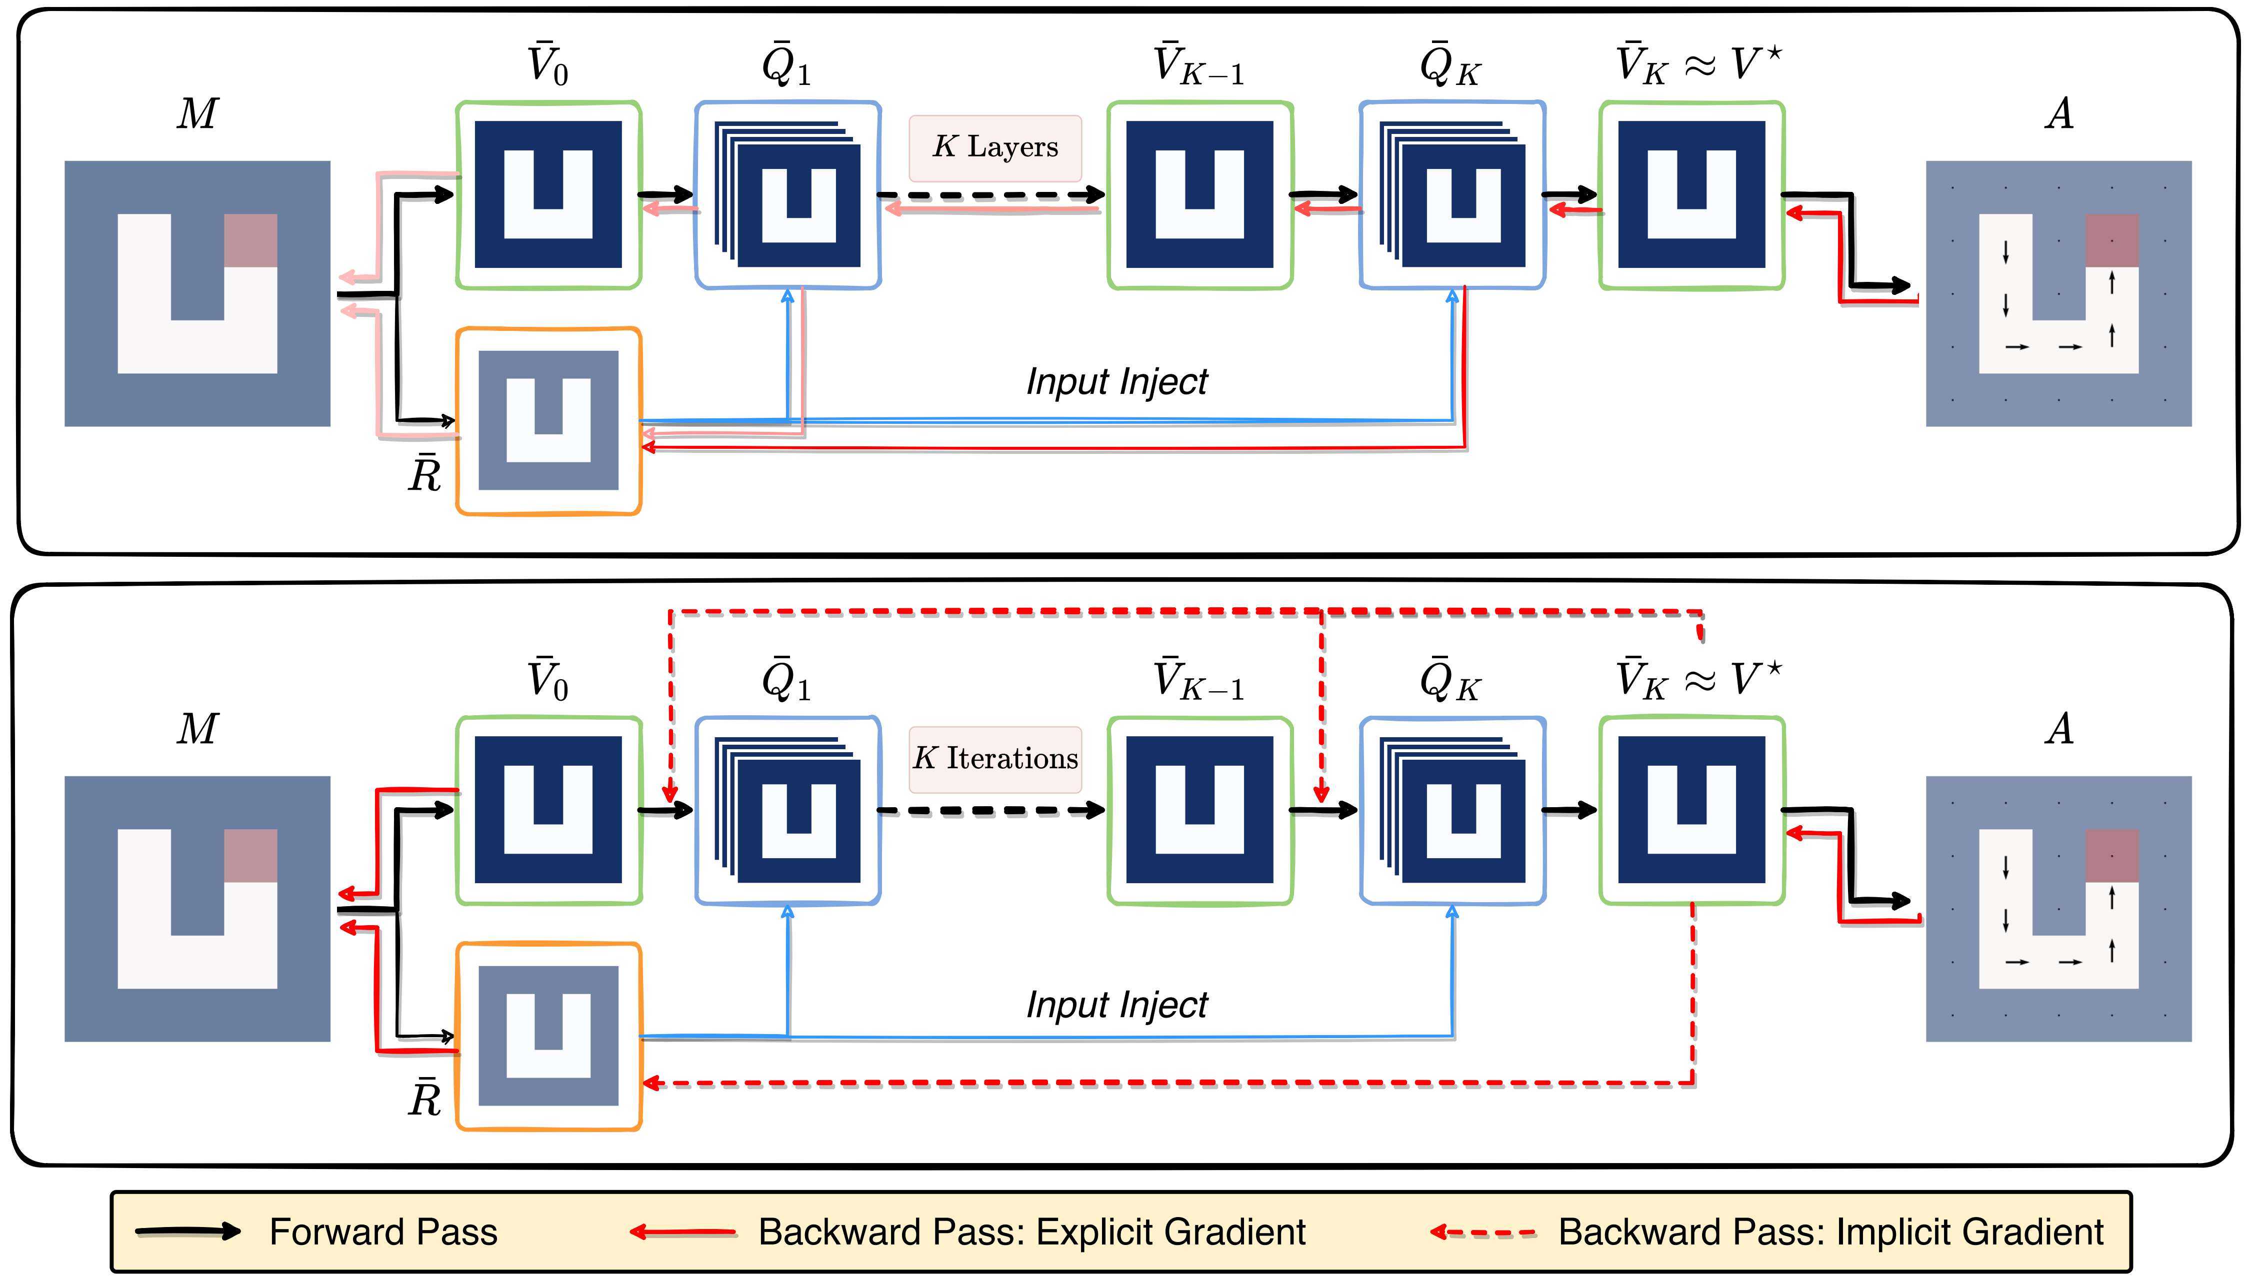

Path Planning

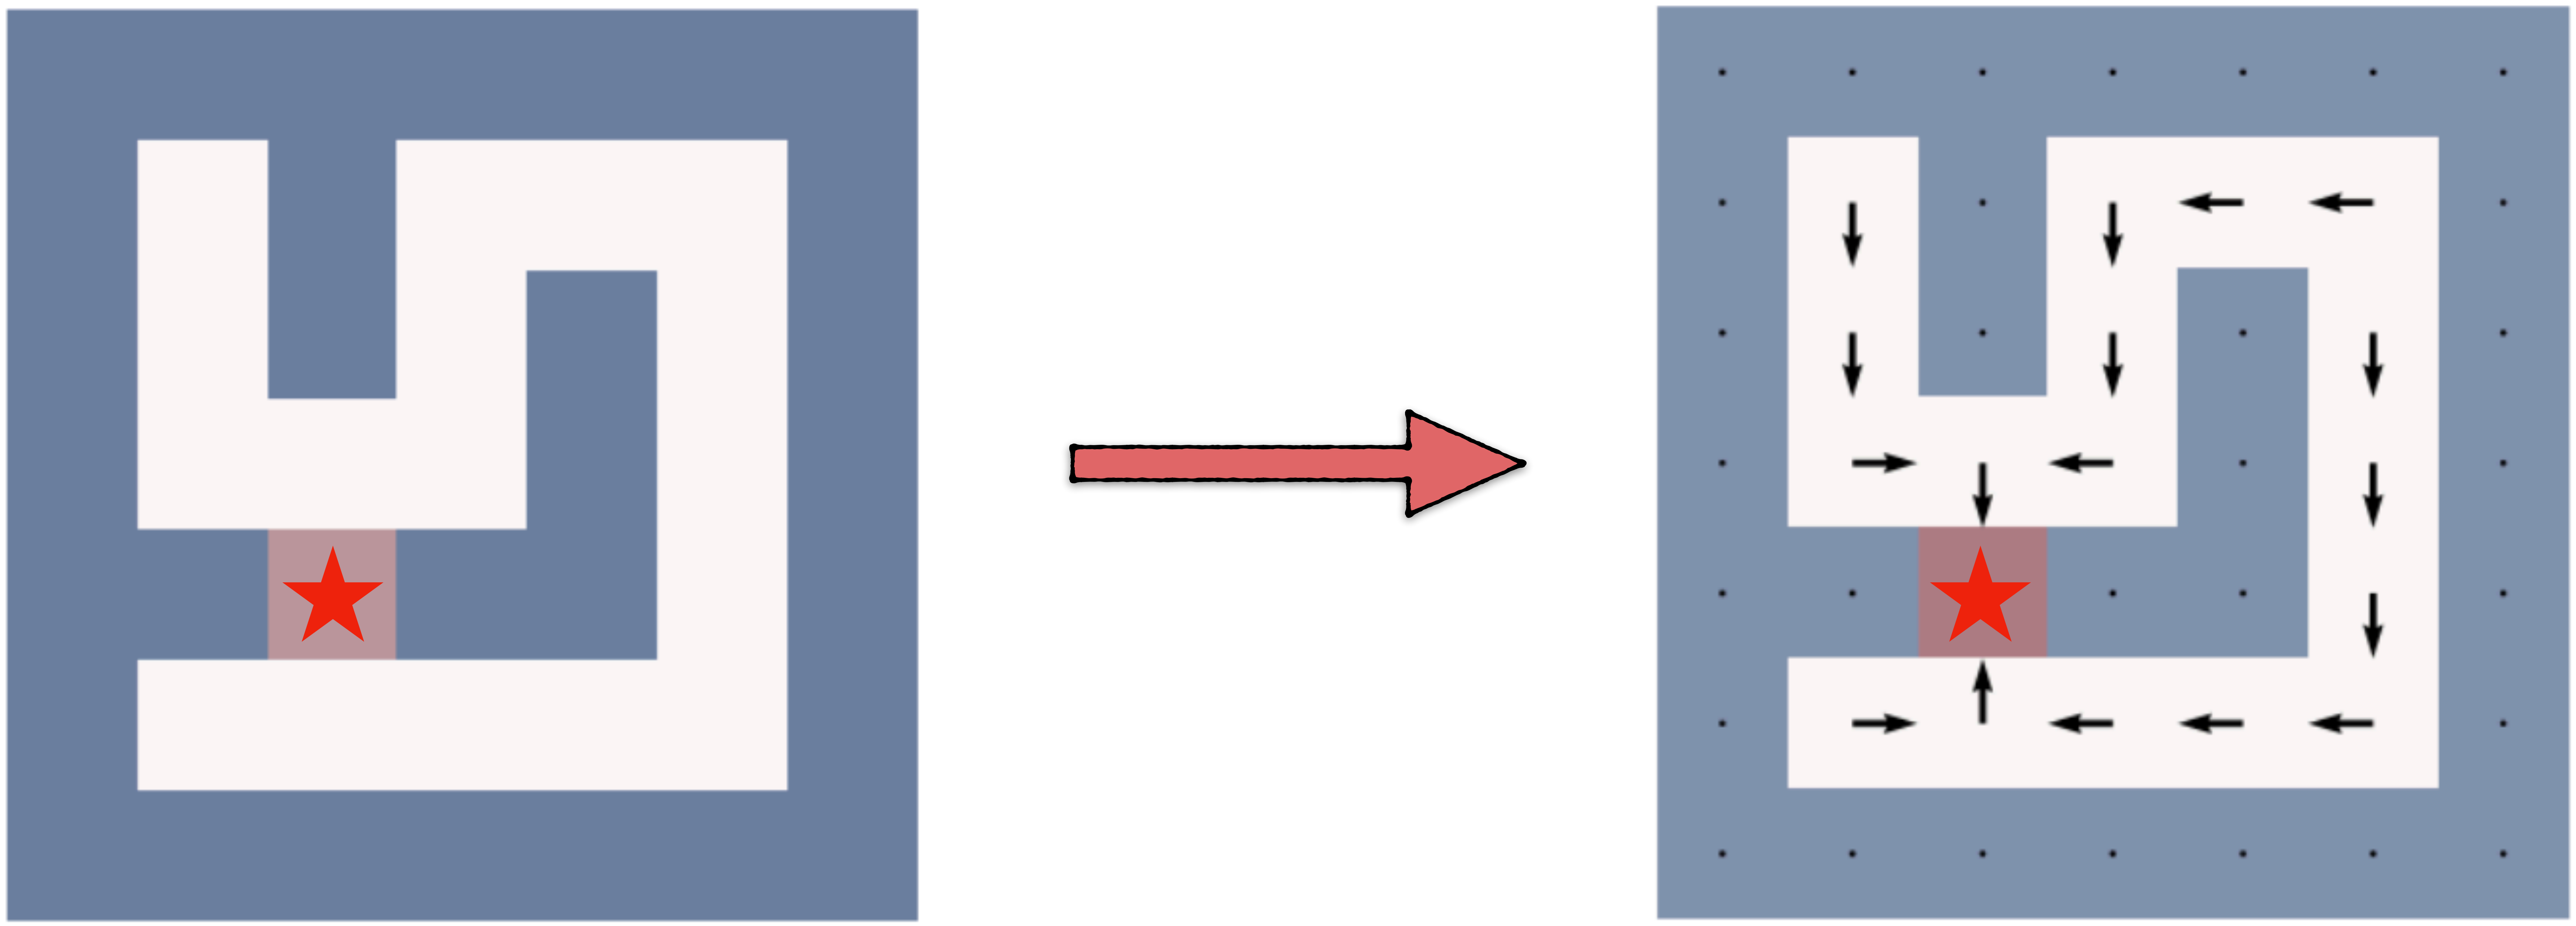

Find shortest path / optimal actions to the goal location (red)

Tasks

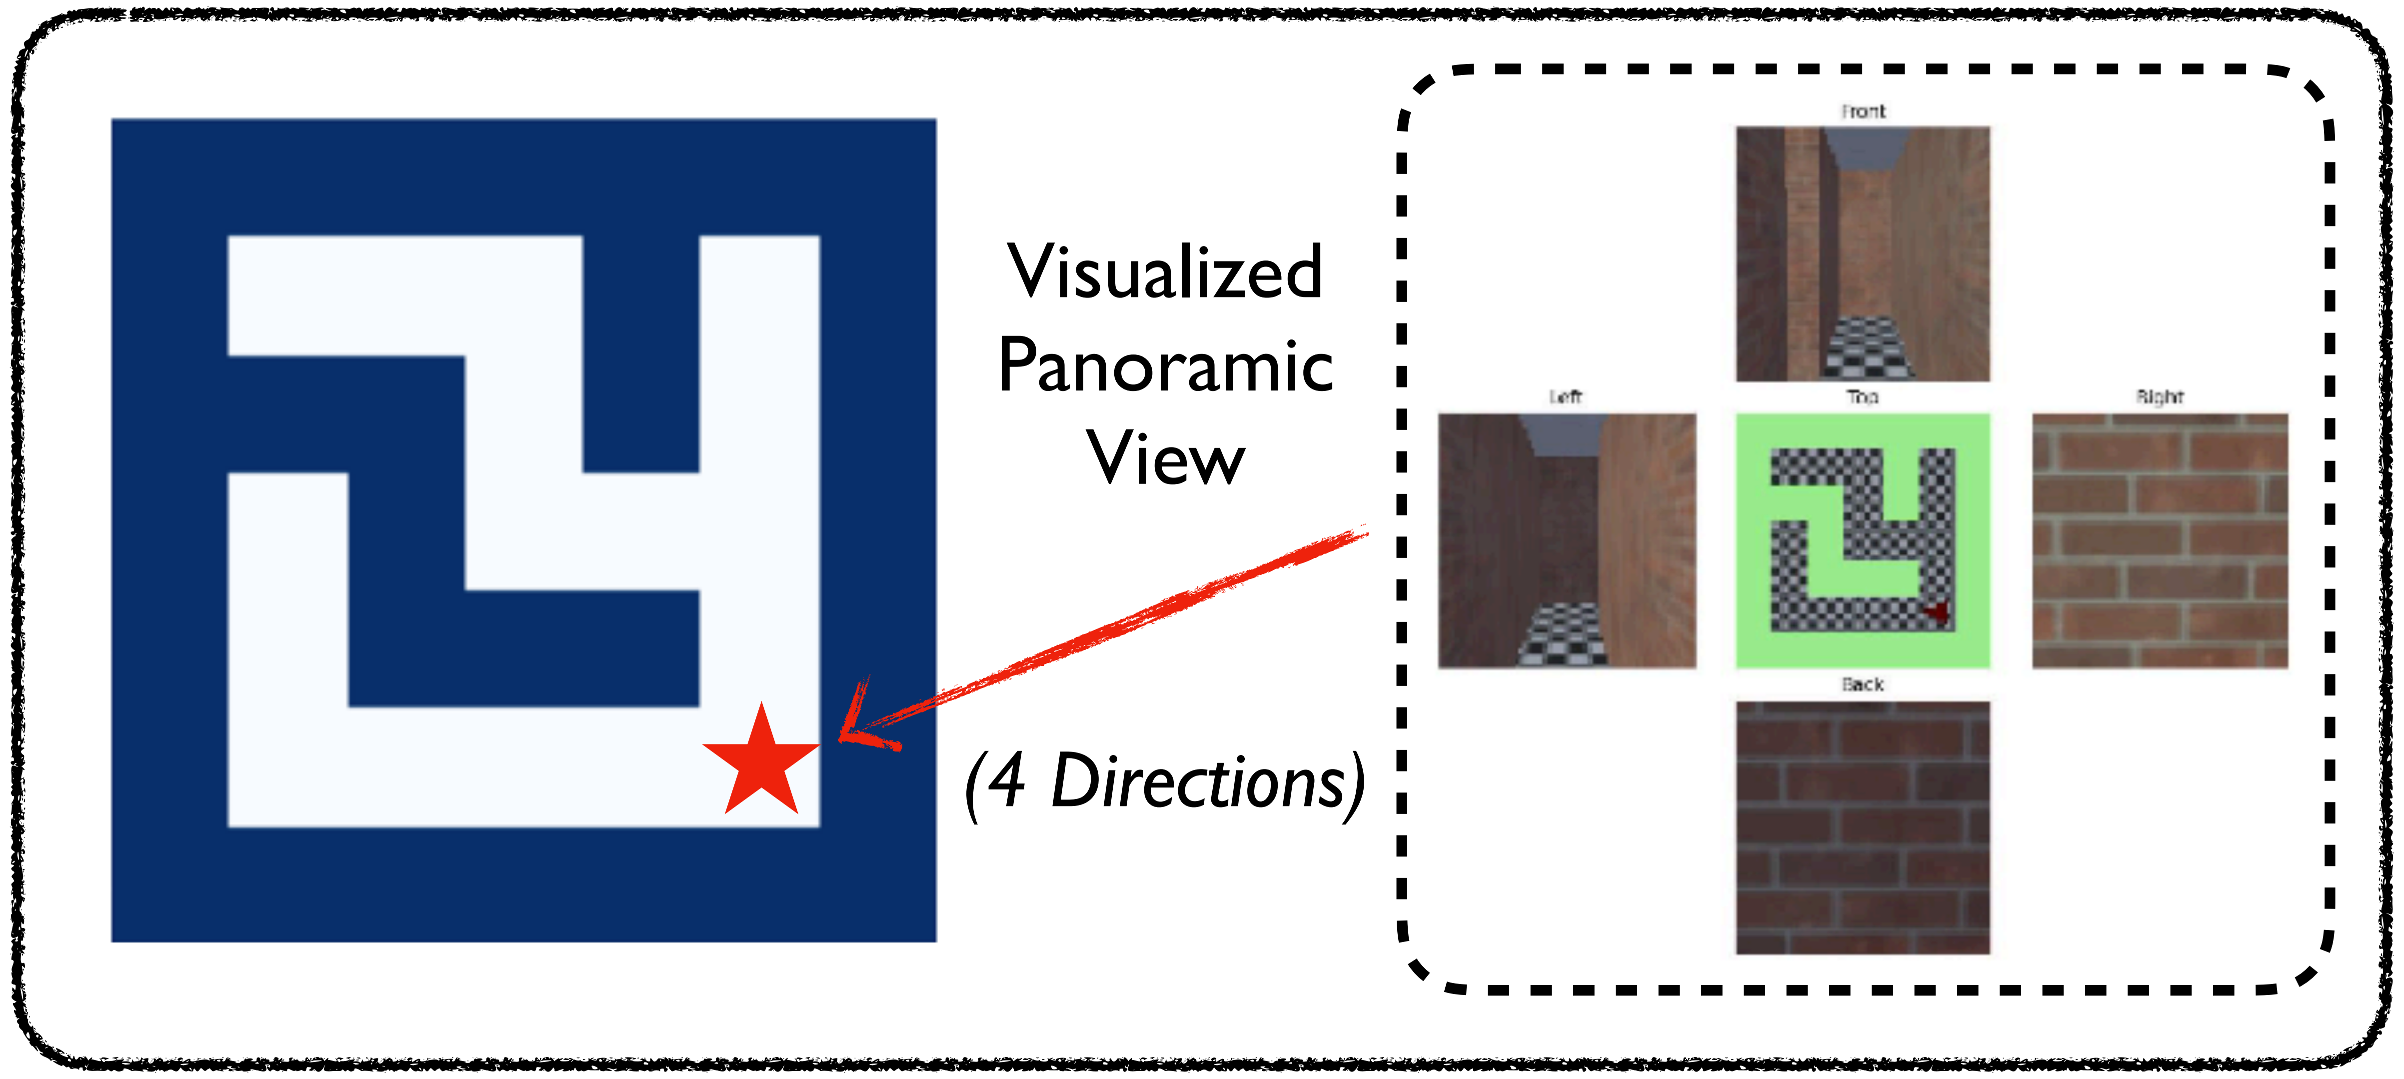

2D and Visual
Maze Navigation

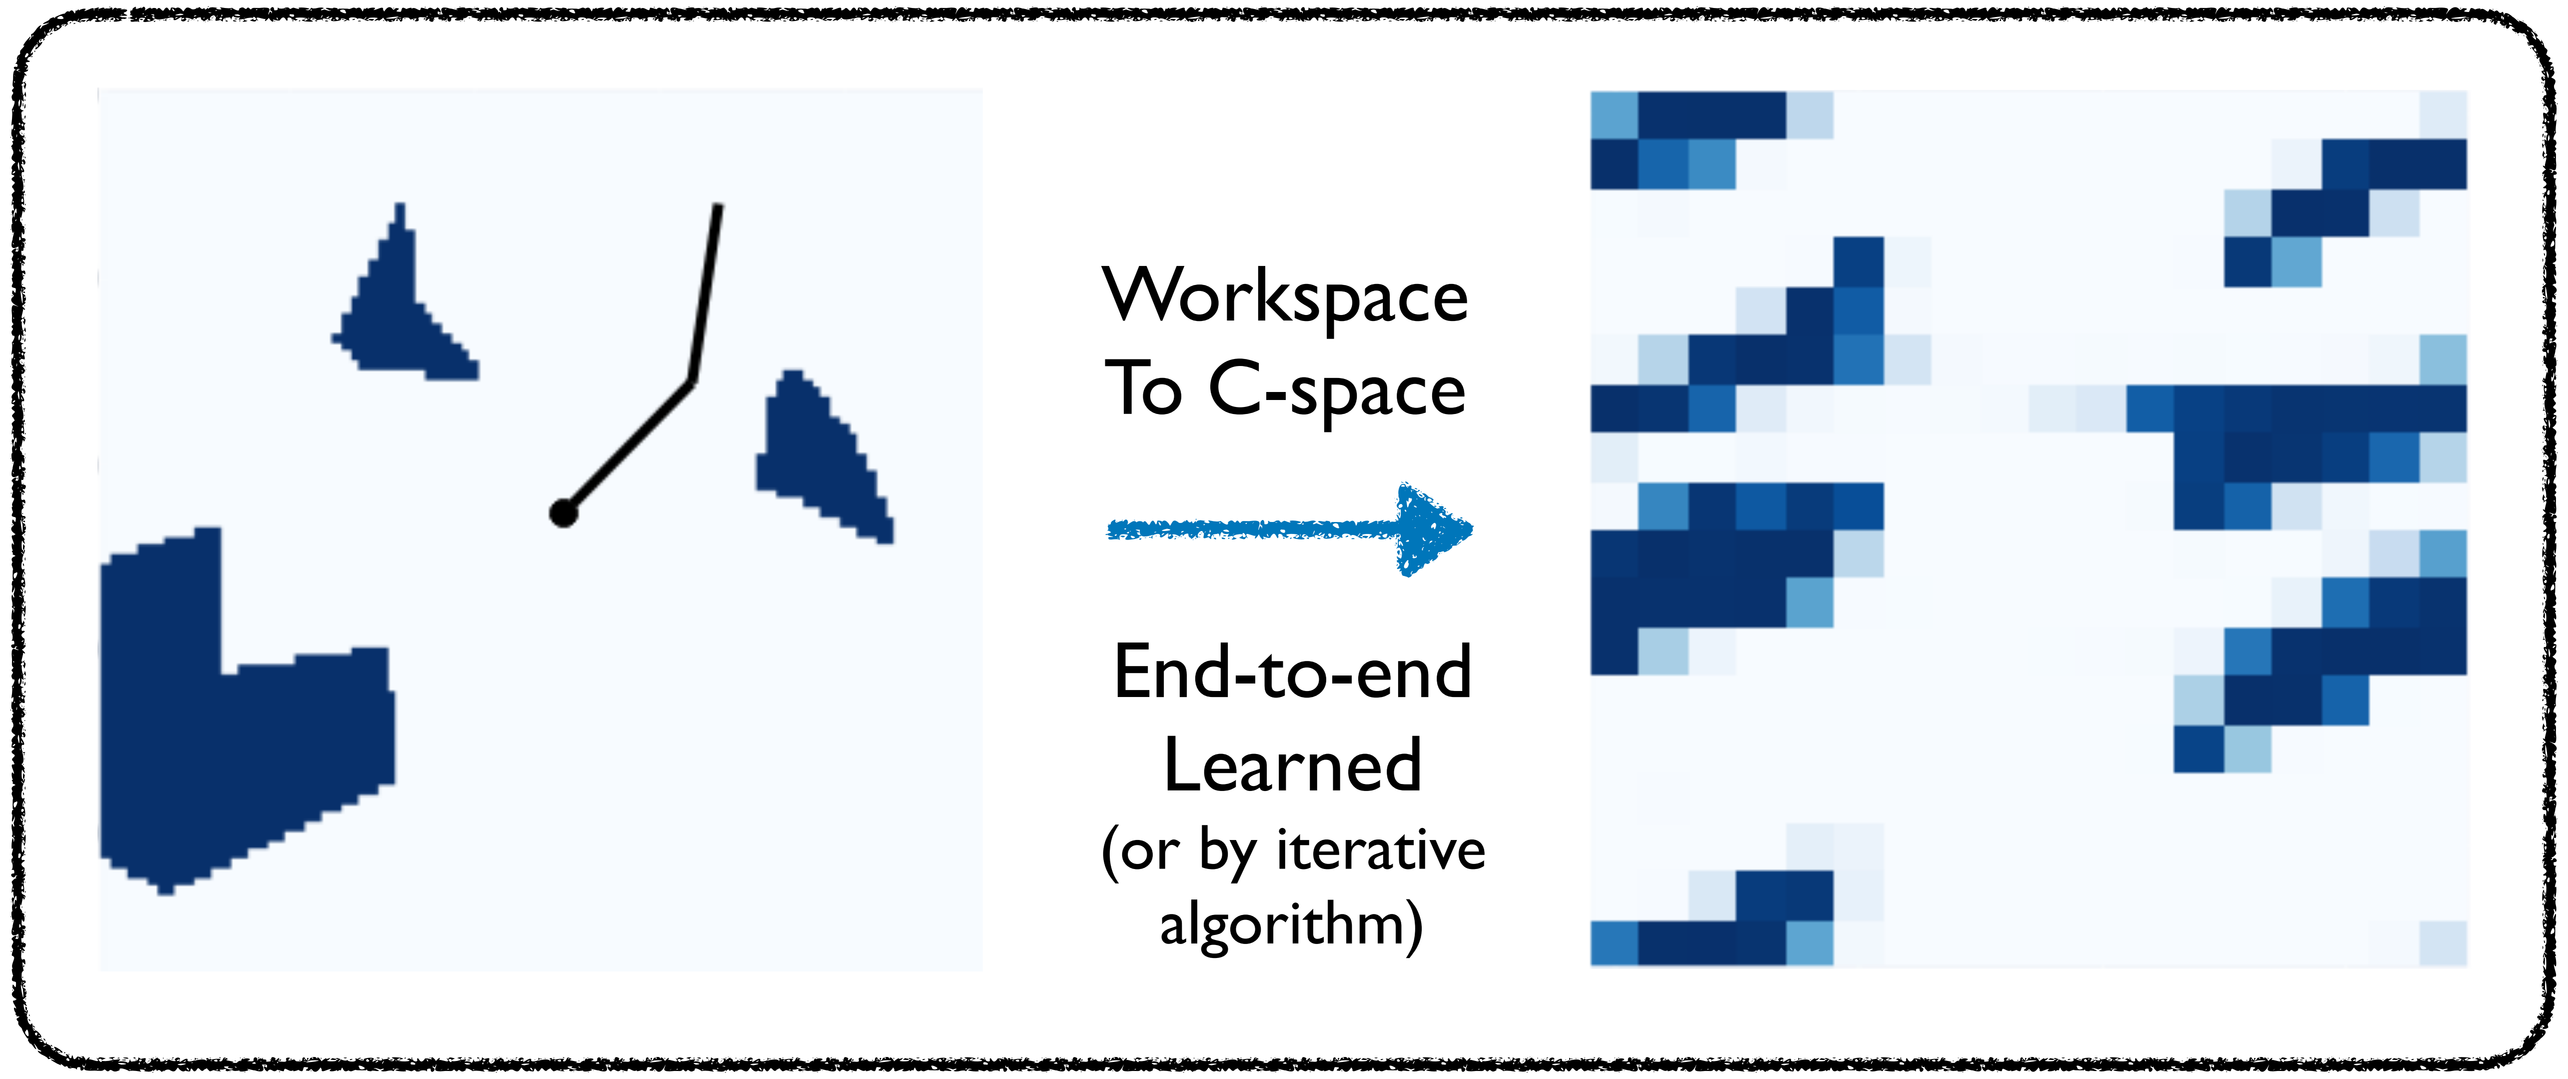

2-DOF Manipulation
In Workspace and C-space

Background: Value Iteration Networks

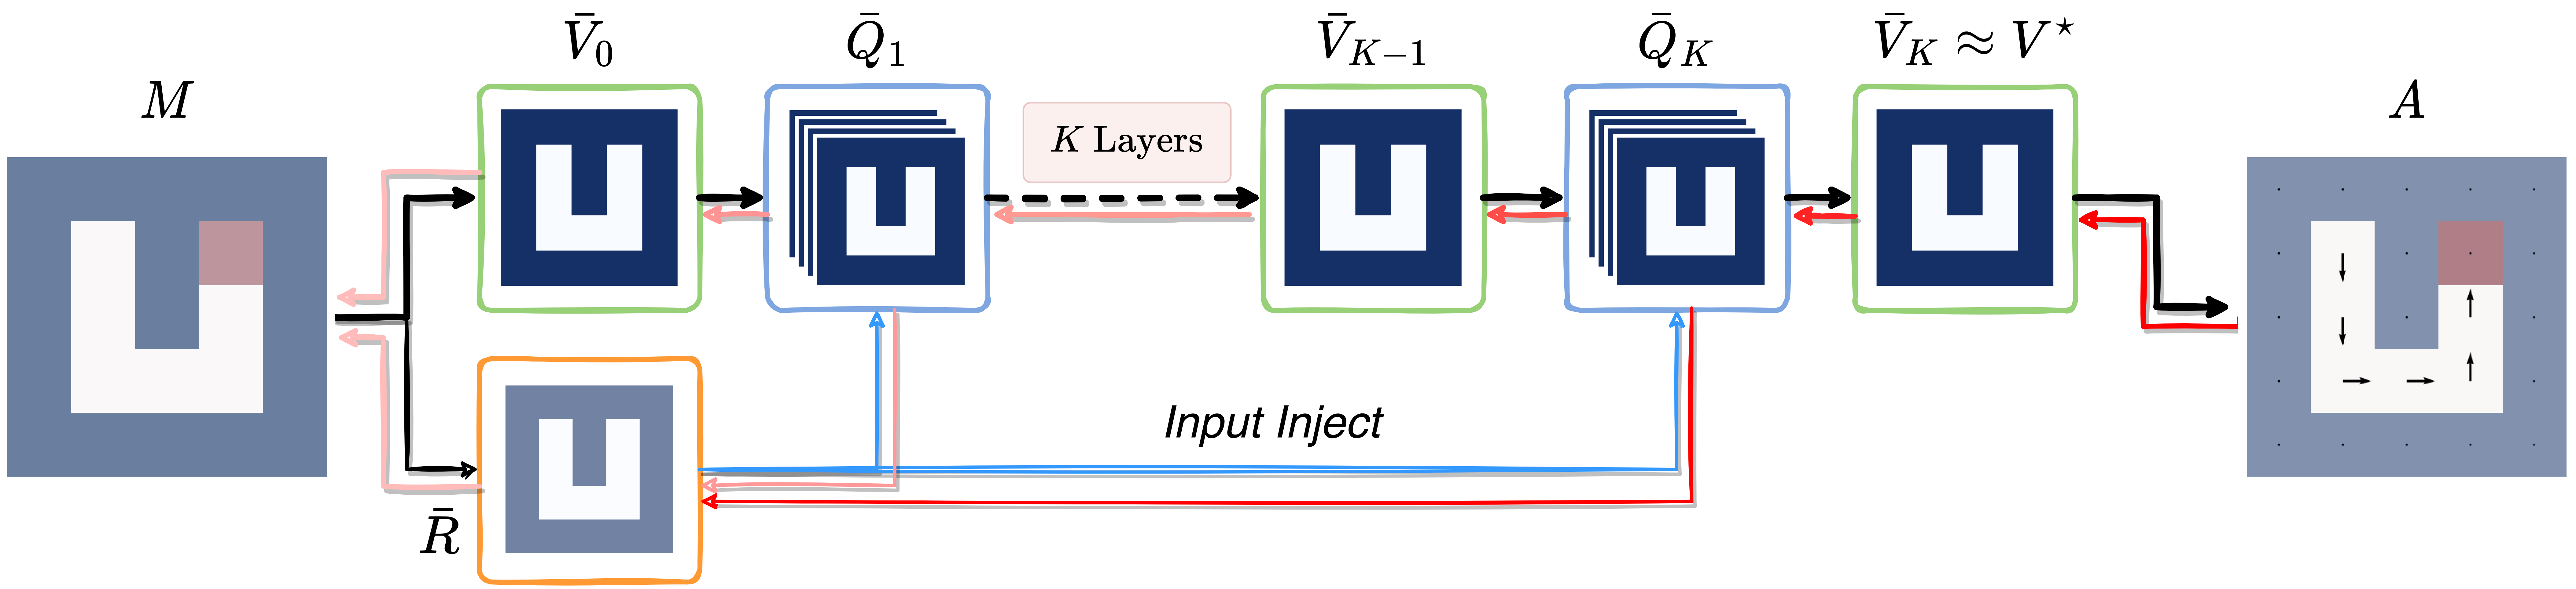

- *Value Iteration Networks* implement Value Iteration by CNNs
- It iteratively applies Bellman operator and differentiates through multiple layers

Tamar et al. Value Iteration Networks. NIPS 2016.

Algorithmic Differentiation in VIN

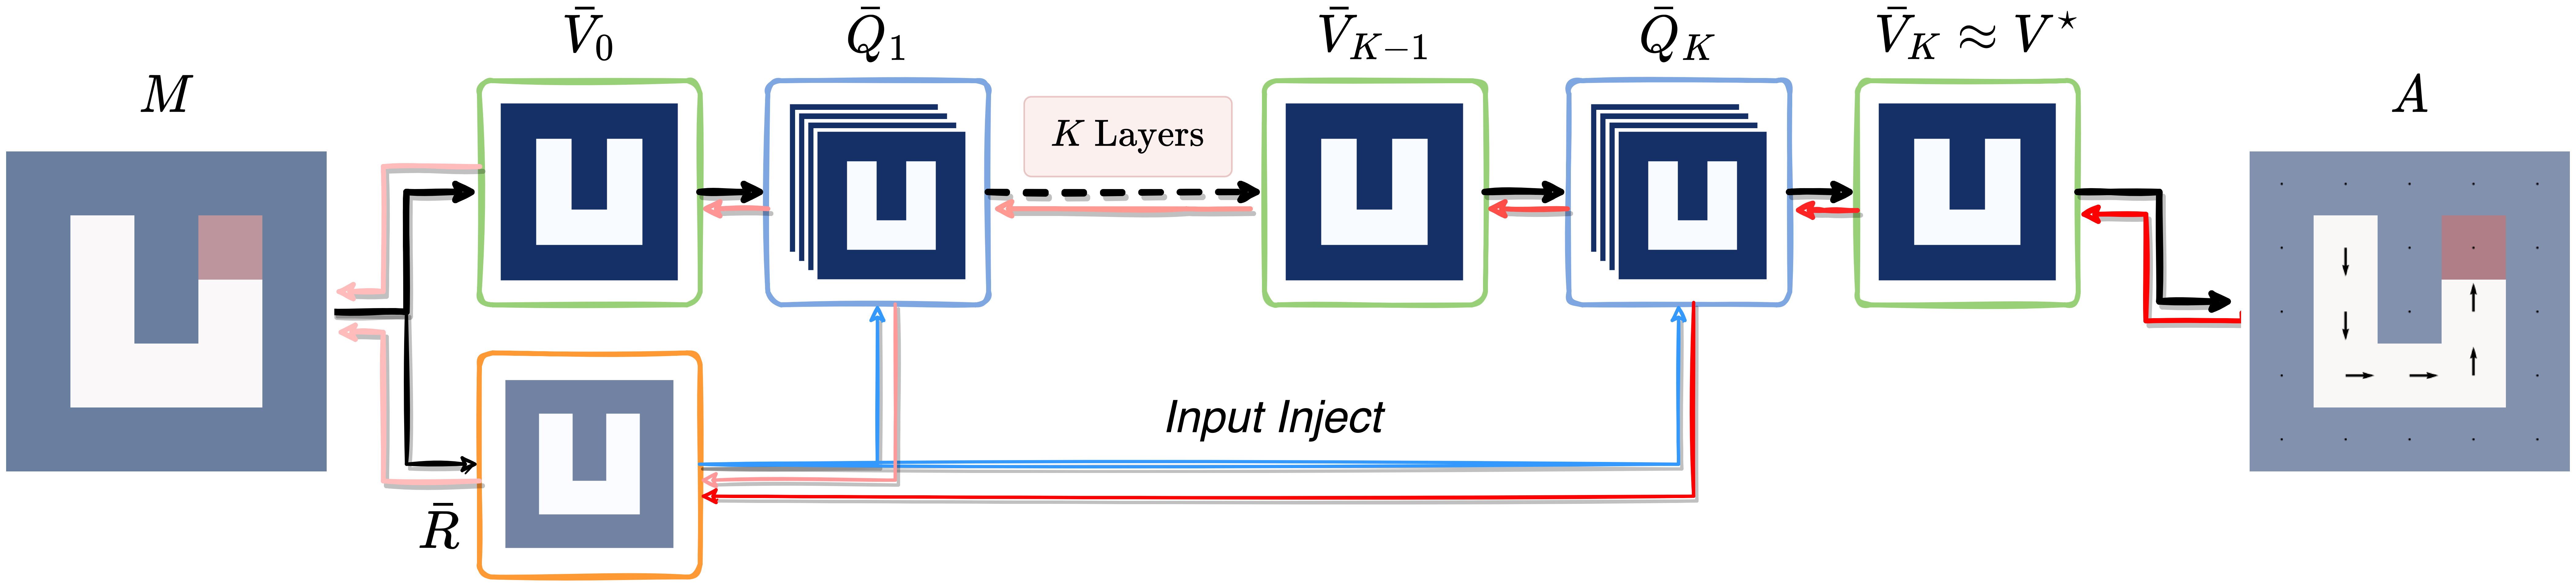

Key issue:
Forward and Backward passes are coupled together

When the planning horizon is long, backpropagation is not scalable, stable, or efficient

Implicit Differentiation => remove math

- Bellman equation: $$\boldsymbol{v}^\star = f(\boldsymbol{v}^\star, \boldsymbol{x})$$
- Differentiating both sides: $$\frac{\partial \boldsymbol{v}^\star(\cdot)}{\partial(\cdot)} = \frac{\partial f(\boldsymbol{v}^\star(\cdot), \boldsymbol{x})}{\partial(\cdot)} = \frac{\partial f(\boldsymbol{v}^\star, \boldsymbol{x})}{\partial \boldsymbol{v}^\star} \frac{\boldsymbol{v}^\star(\cdot)}{\partial(\cdot)} + \frac{\partial f(\boldsymbol{v}^\star, \boldsymbol{x})}{\partial(\cdot)}$$
$$\frac{\partial \ell}{\partial(\cdot)} = \frac{\partial \ell}{\partial \boldsymbol{v}^\star} \frac{\partial \boldsymbol{v}^\star(\cdot)}{\partial(\cdot)} = \frac{\partial \ell}{\partial \boldsymbol{v}^\star} \left(I - \frac{\partial f(\boldsymbol{v}^\star, \boldsymbol{x})}{\partial \boldsymbol{v}^\star} \right)^{-1} \frac{\partial f(\boldsymbol{v}^\star, \boldsymbol{x})}{\partial(\cdot)}$$
- Solving backward fixed-point: $$\boldsymbol{w}^\top \triangleq \frac{\partial \ell}{\partial \boldsymbol{v}^\star} \left(I - \frac{\partial f(\boldsymbol{v}^\star, \boldsymbol{x})}{\partial \boldsymbol{v}^\star} \right)^{-1}; \quad \boldsymbol{w}^\top = \boldsymbol{w}^\top \frac{\partial f(\boldsymbol{v}^\star, \boldsymbol{x})}{\partial \boldsymbol{v}^\star} + \frac{\partial \ell}{\partial \boldsymbol{v}^\star}.$$

Bai et al. Deep Equilibrium Models. 2019.
Nikishin et al. Control-Oriented Model-Based Reinforcement Learning with Implicit Differentiation. 2021.
Gehring et al. Understanding End-to-End Model-Based Reinforcement Learning Methods as Implicit Parameterization. 2021.

Method: Implicit Differentiable Planners

Algorithmic Differentiable Planner: (ADP)
VIN

Implicit Differentiable Planner: (IDP)
***ID**-VIN*

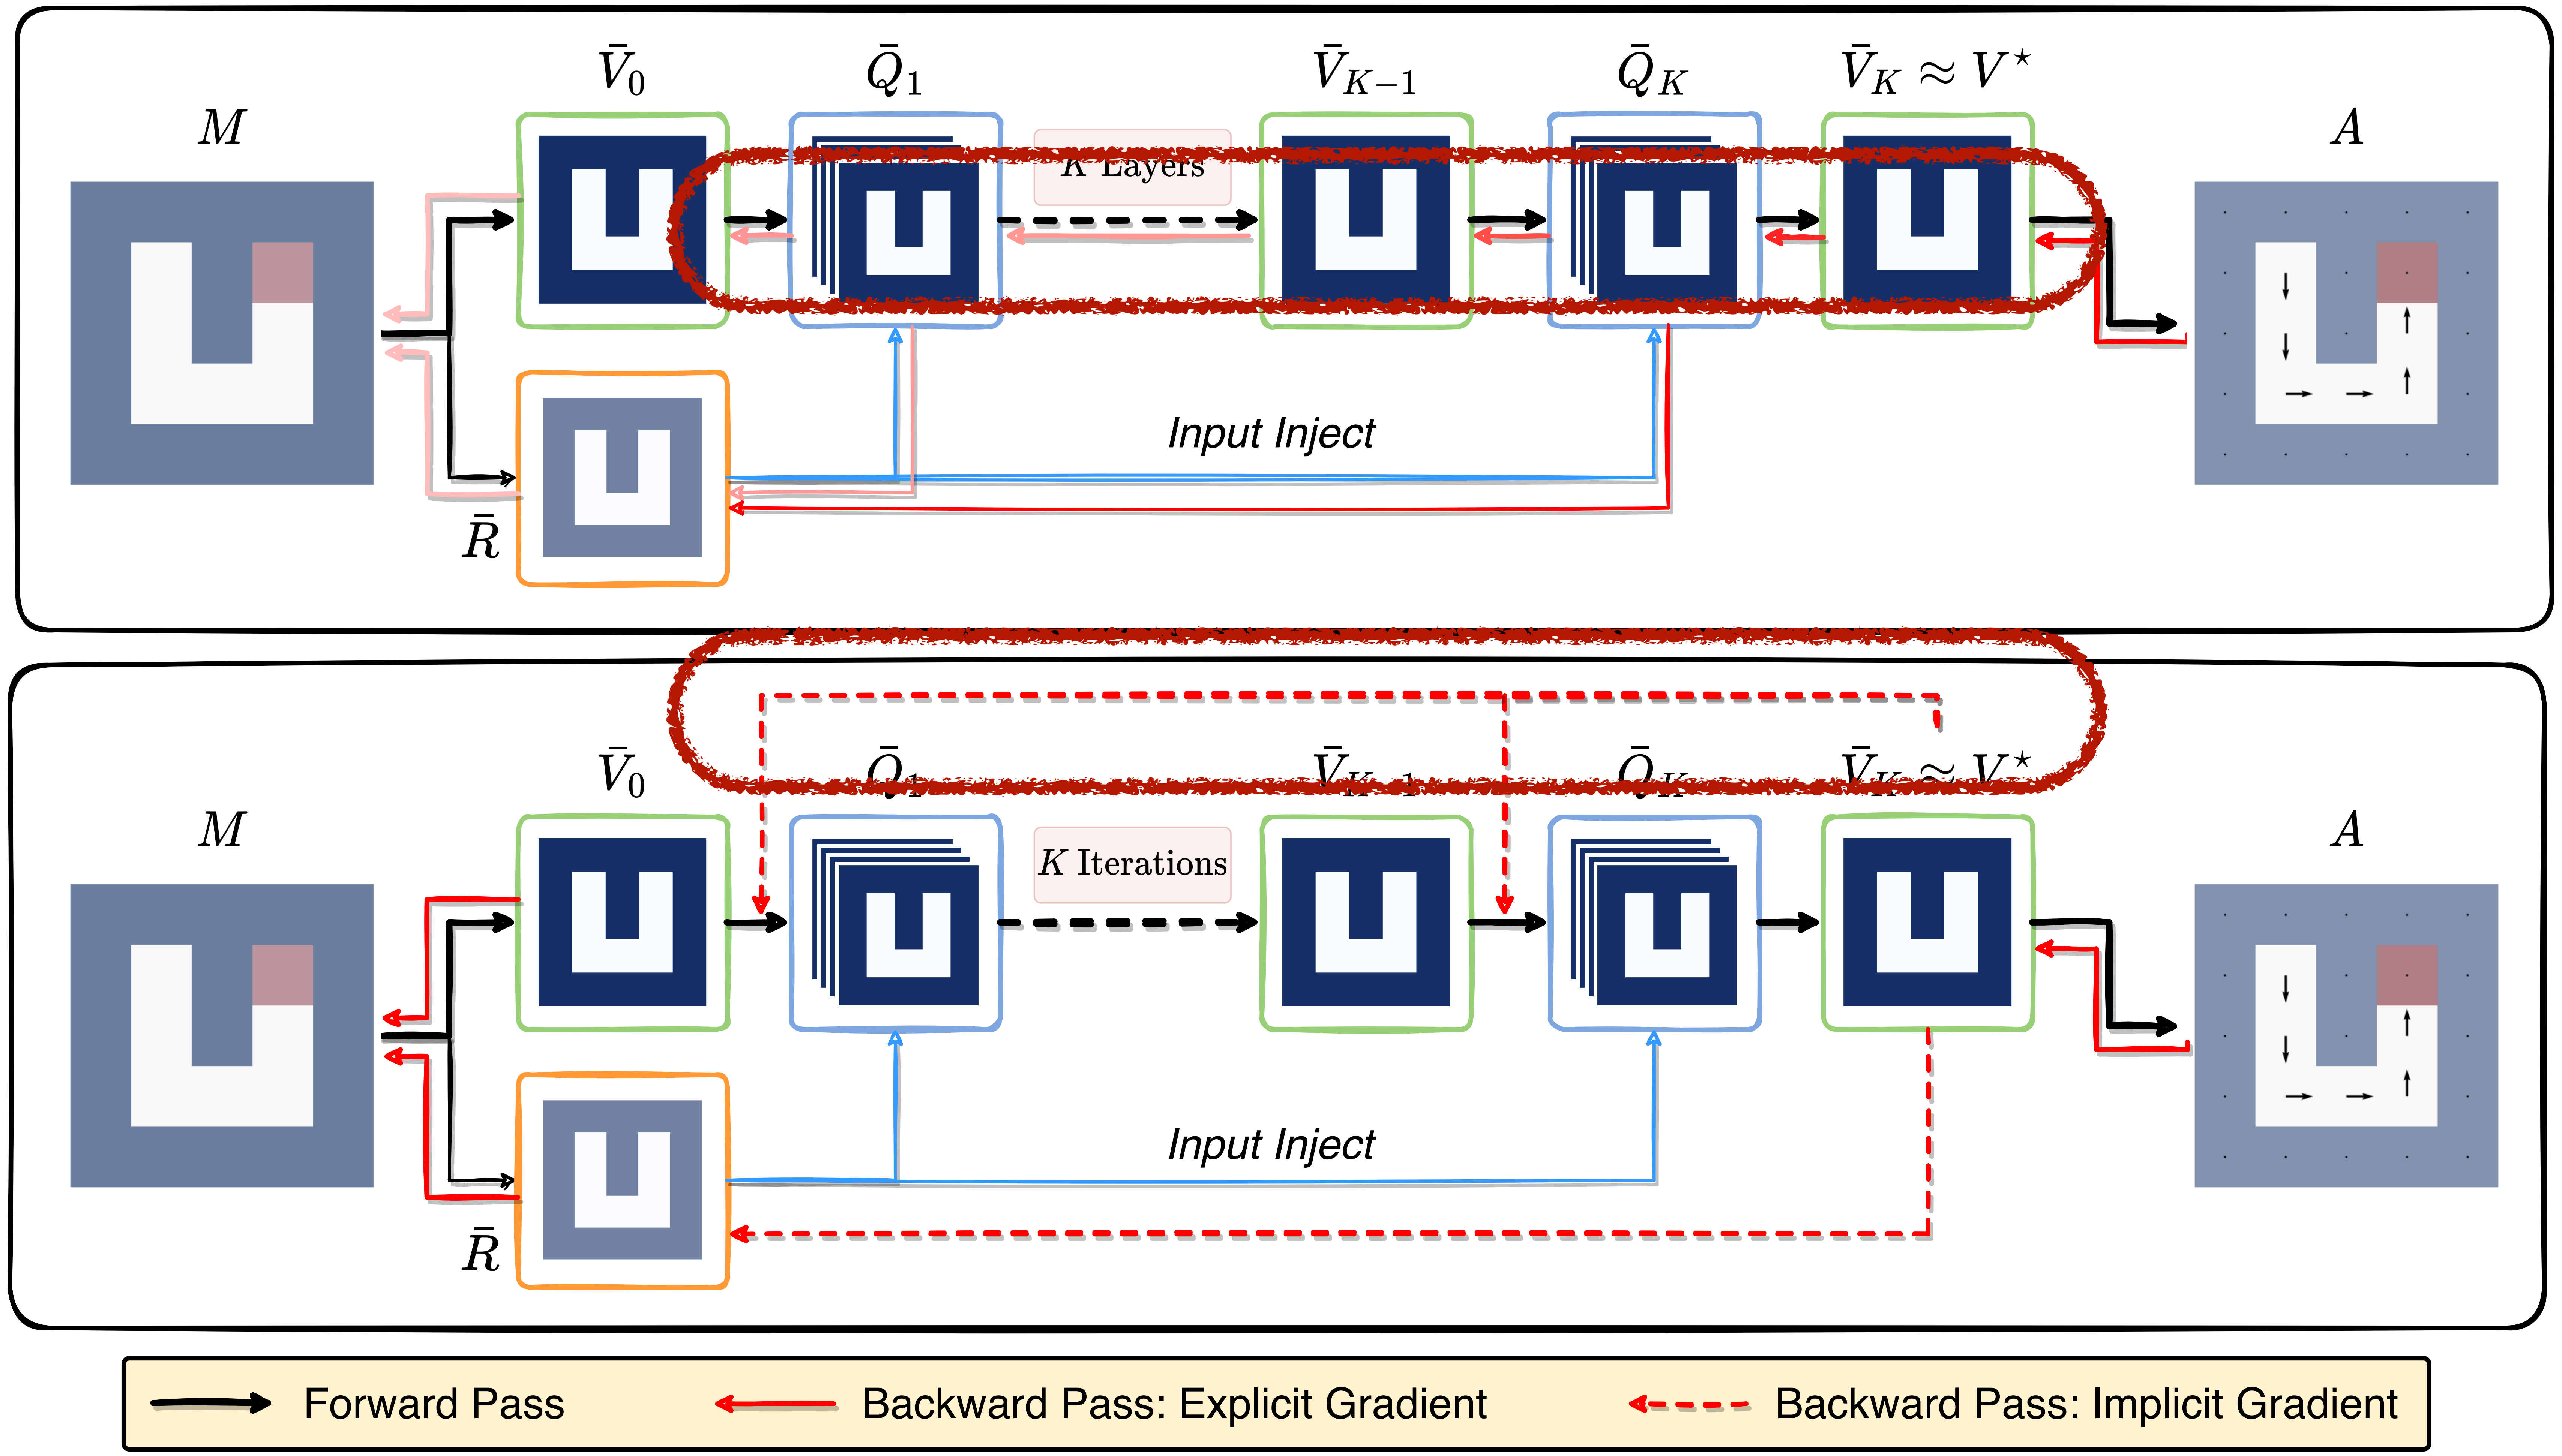

Results: Runtime on 2D Nav

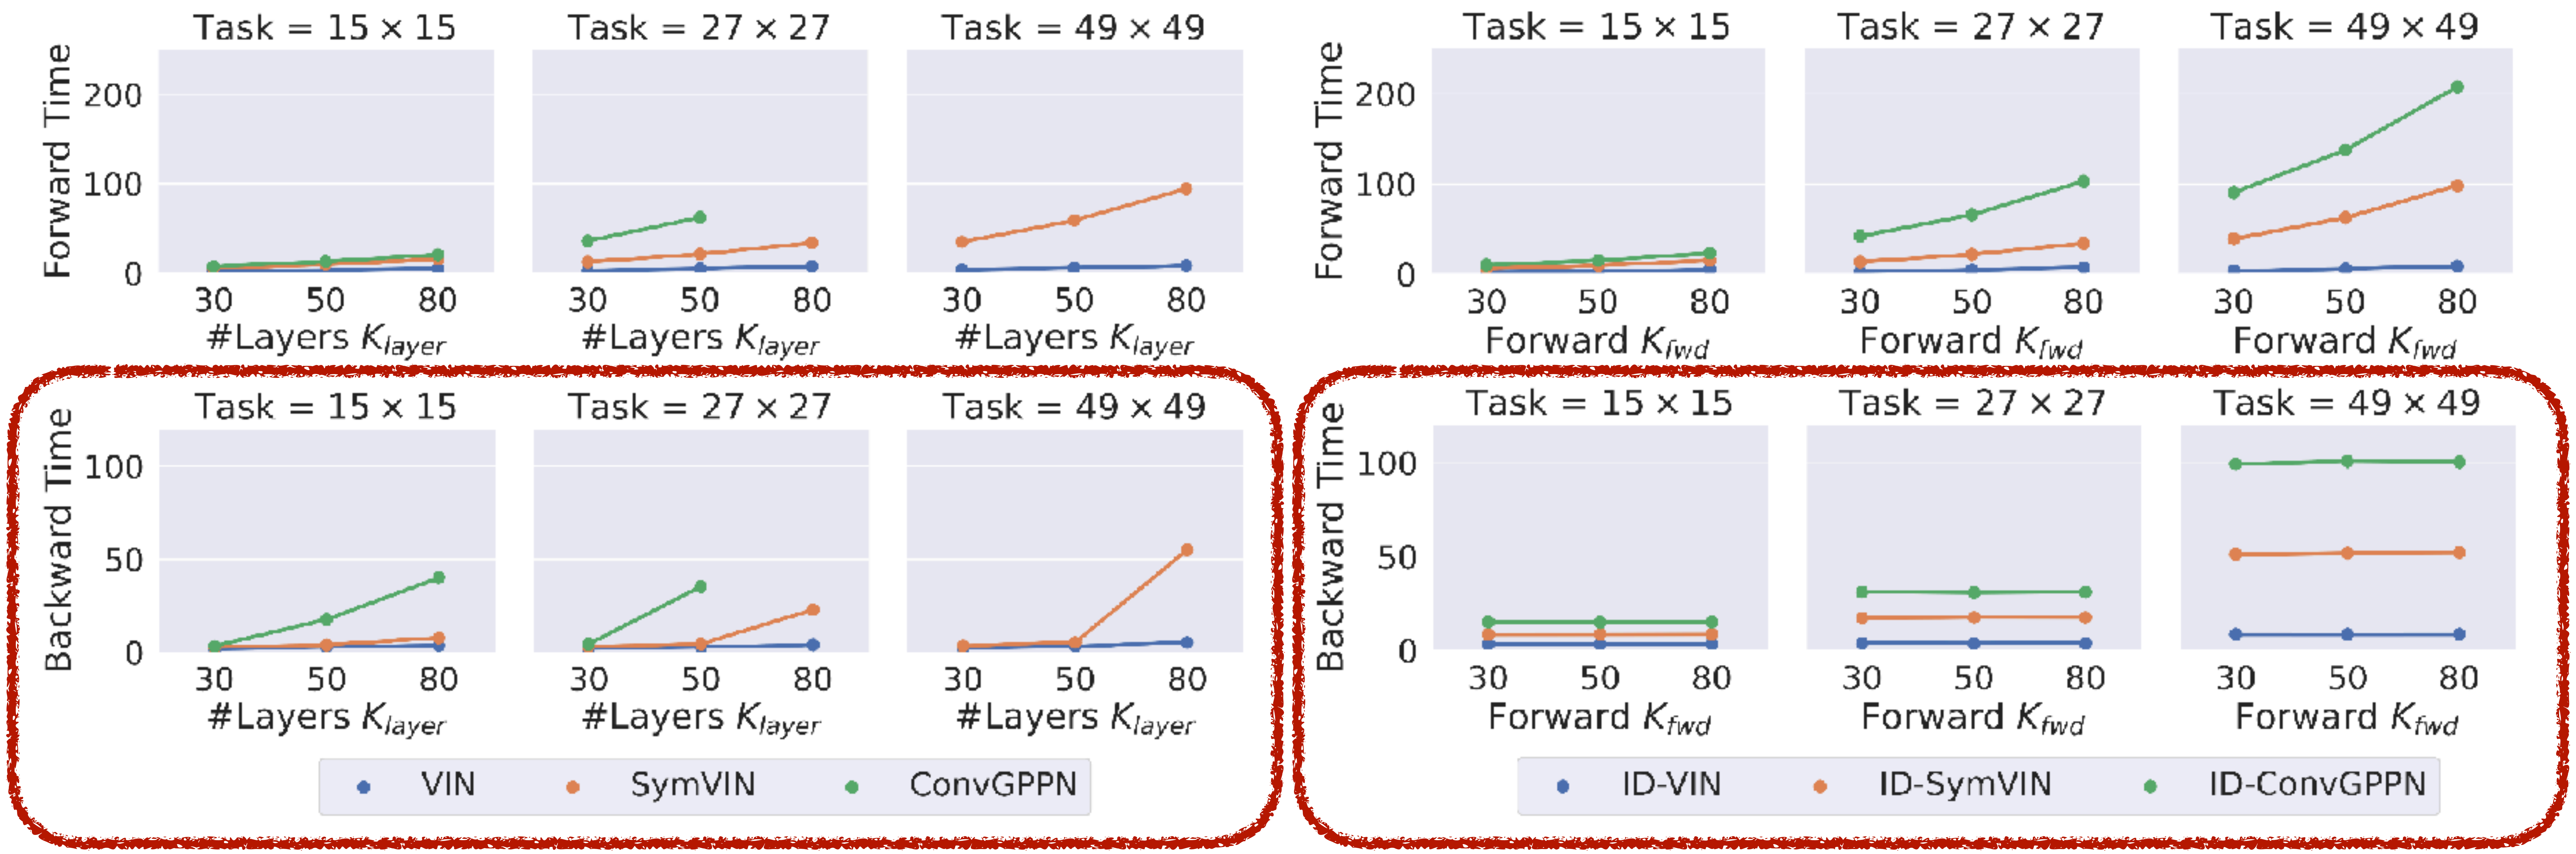

Algorithmic Differentiable Planners

Implicit Differentiable Planners

Results: Success Rate

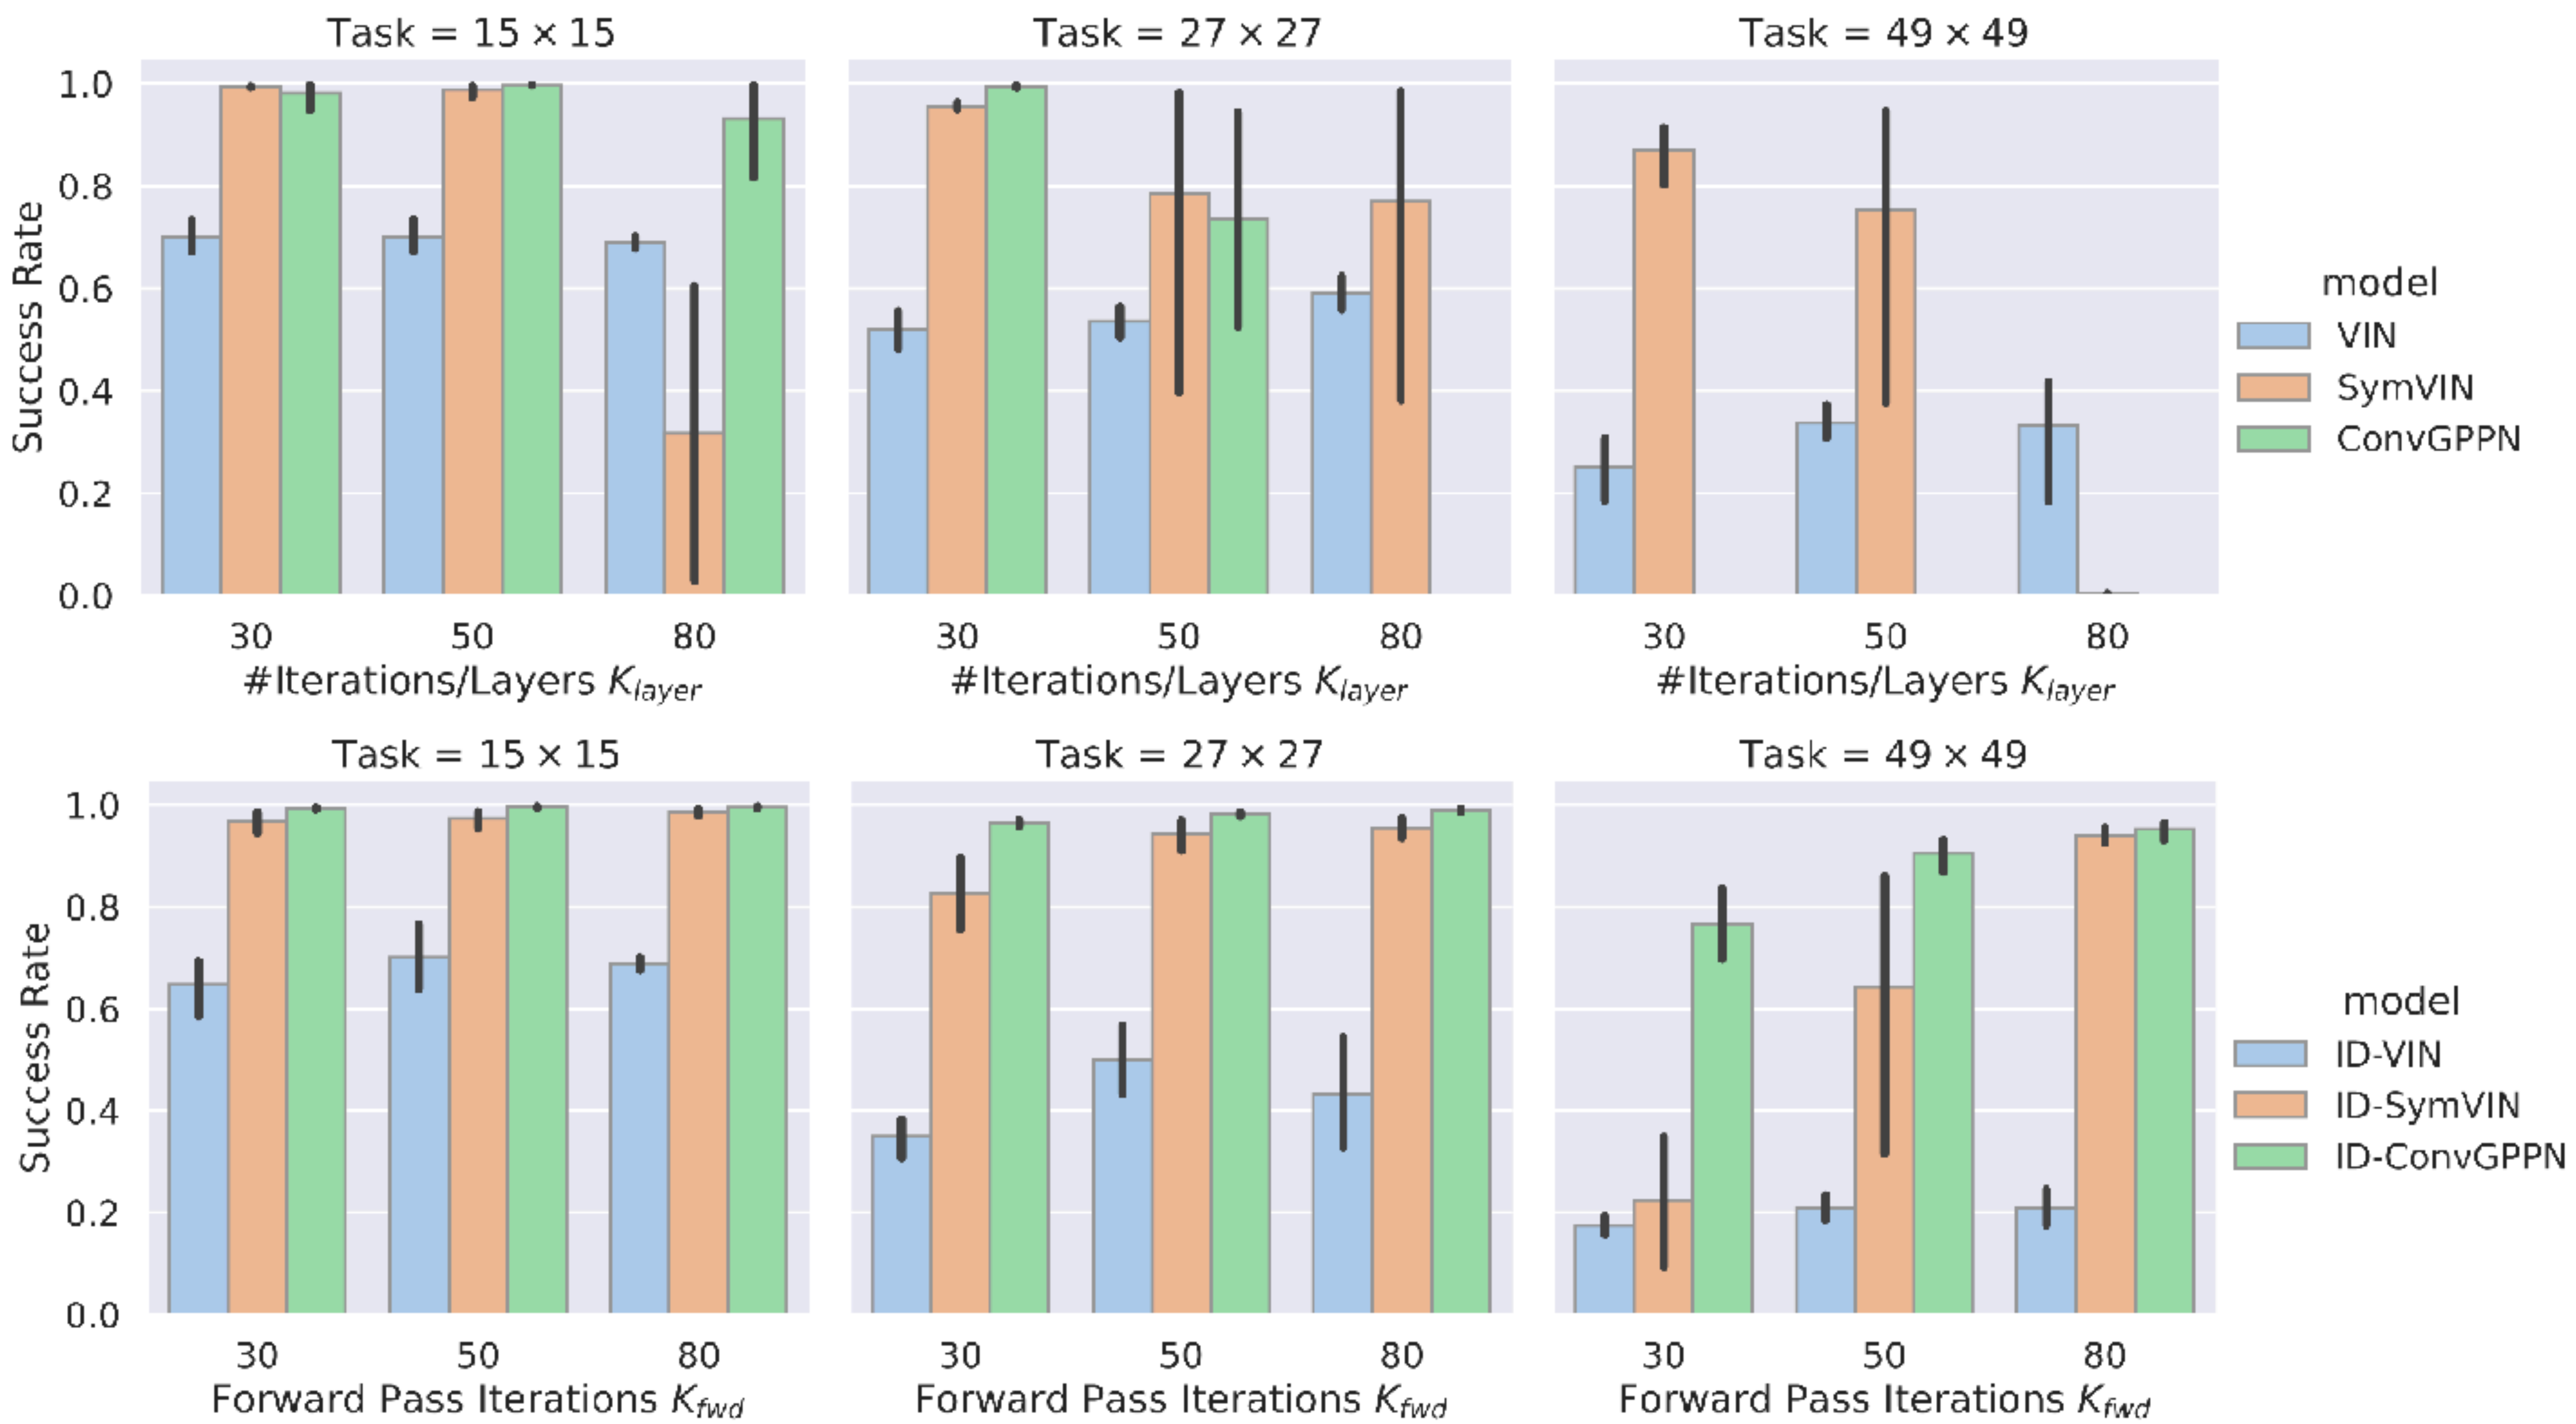

Topics

- *Learning for* End-to-end Planning
  — How to learn to plan and improve stability
- *Symmetric (and Compositional)* Structure in Robot Planning
  — How does symmetry improve (path) planning and world modeling
- *Symmetric* Structure in Robot Learning
  — How could symmetry improve learning robot skills

Grey items are secondary

Symmetry in Path Planning

Integrating Symmetry Into Differentiable Planning With Steerable Convolutions

ICLR 2023

Linfeng Zhao, Xupeng Zhu, Lingzhi Kong, Robin Walters, Lawson Wong

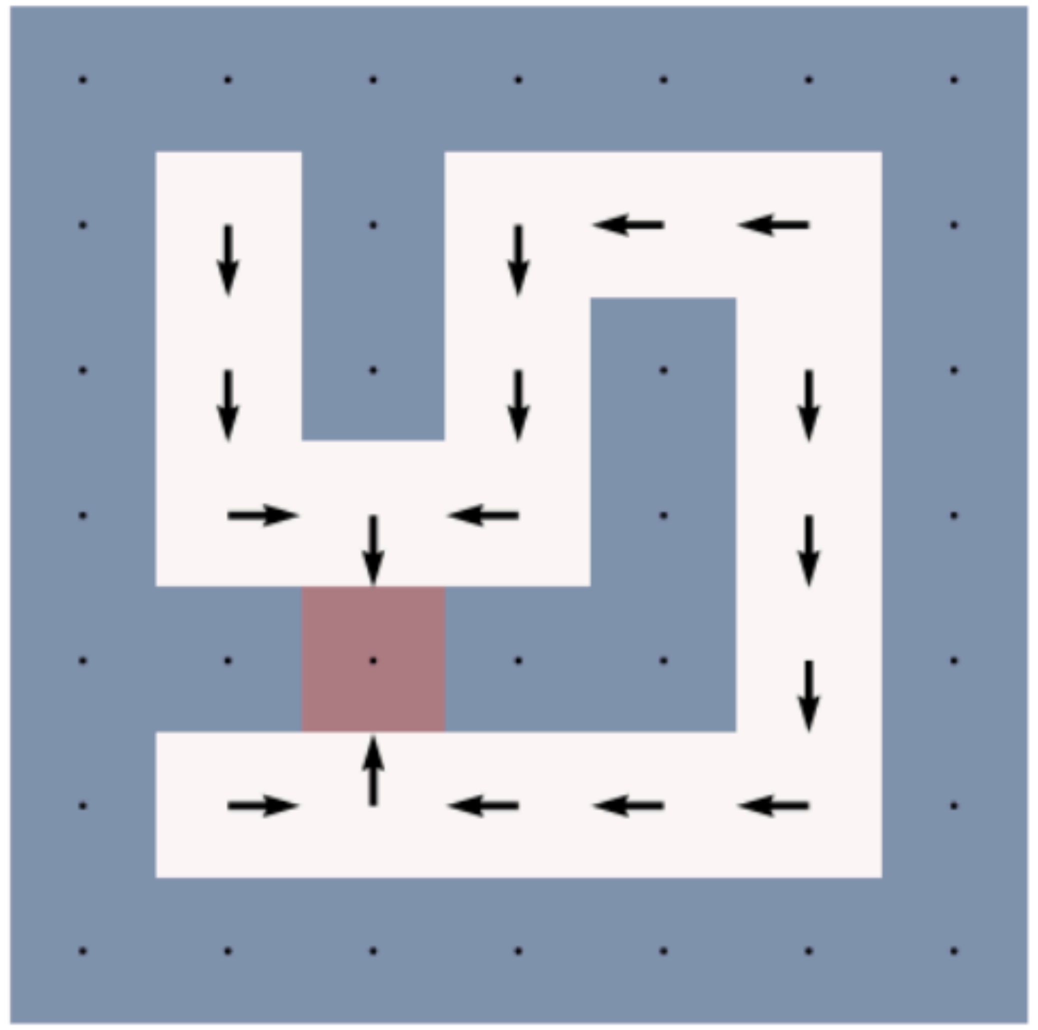

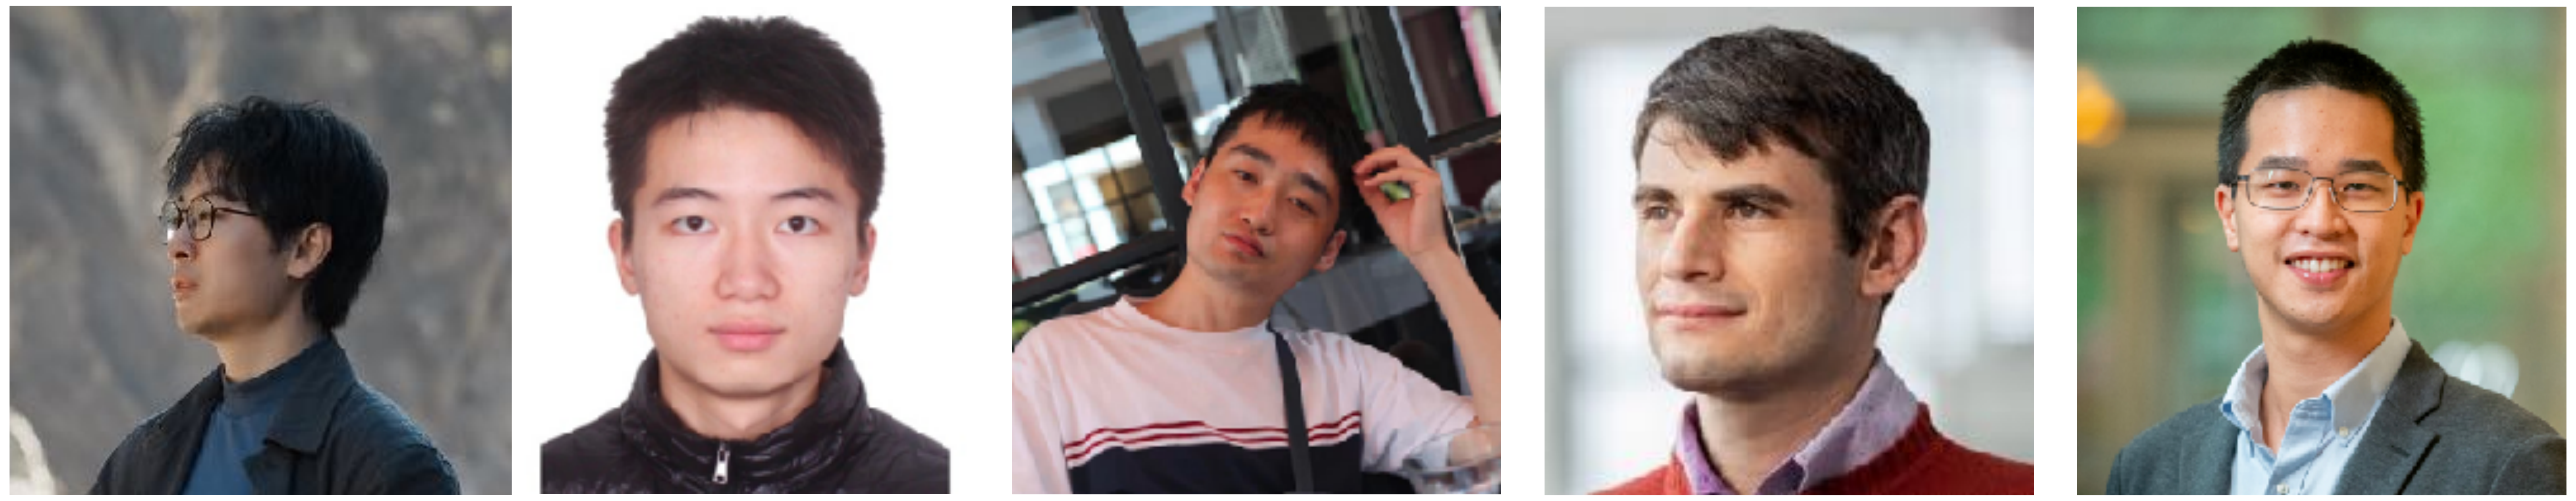

2D Discrete Map Rotation/Reflection

2D Discrete Symmetry D_4

Symmetry in Path Planning

What does the symmetry look like?

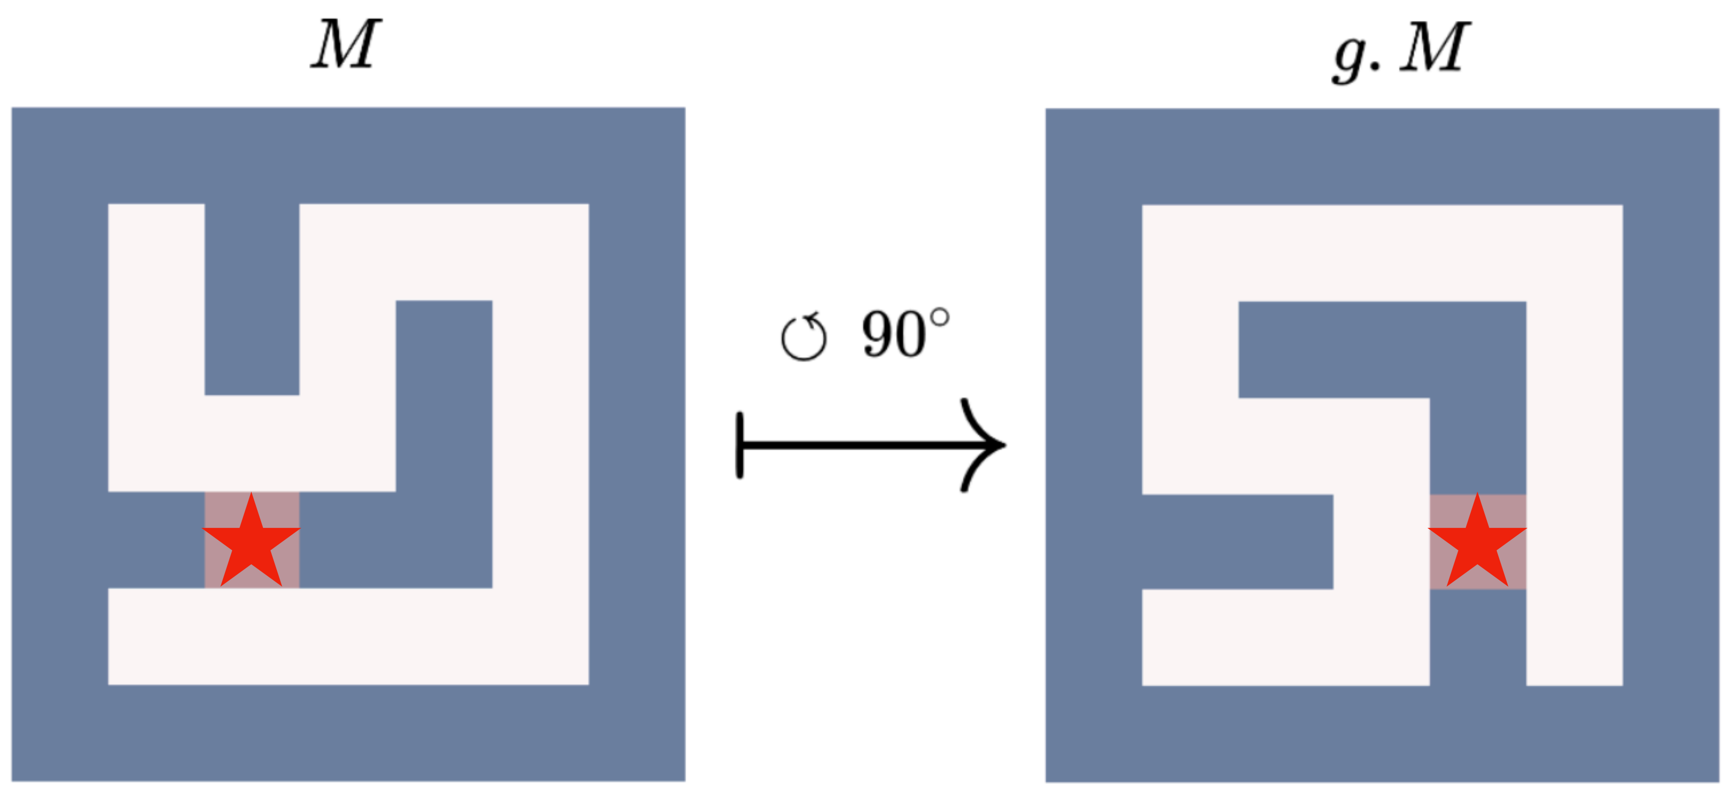

Symmetry in Path Planning

What does the symmetry look like?

They can be described by

Equivariance

$$\circlearrowleft 90^\circ \circ \big(\text{Plan}(M)\big) = \text{Plan}(\circlearrowleft 90^\circ \circ M)$$

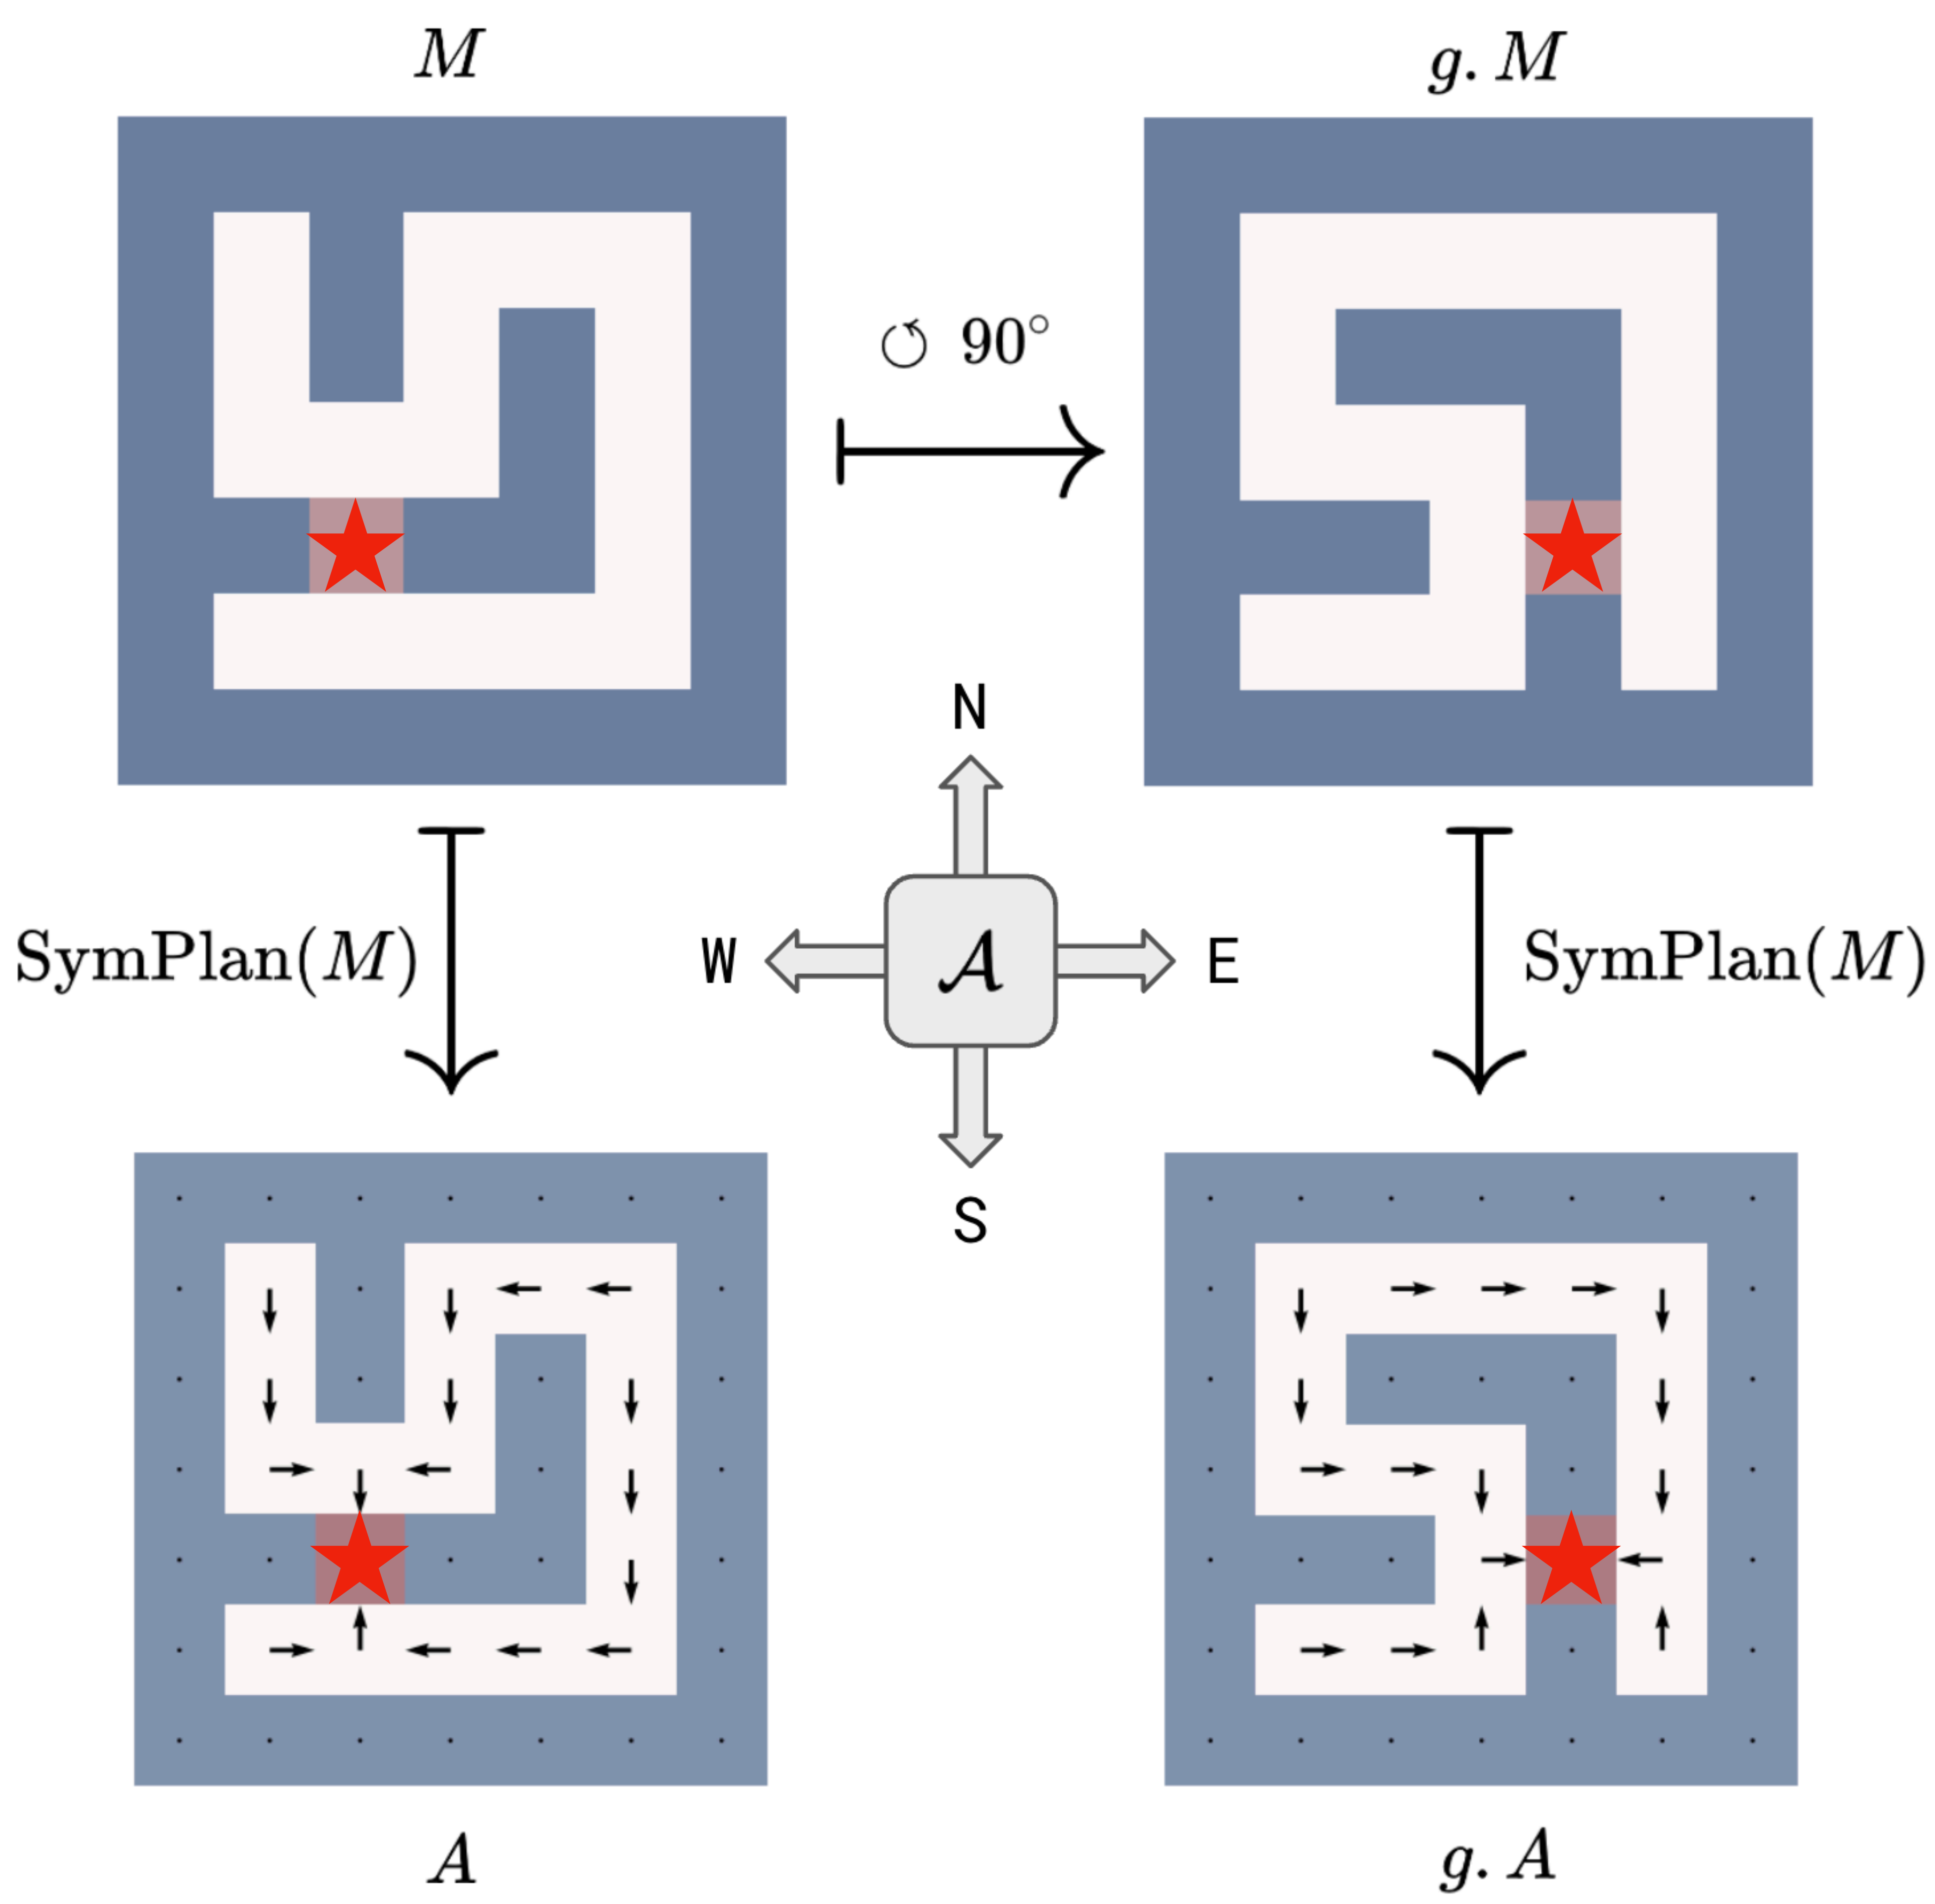

Symmetry: All Rotations and Reflections

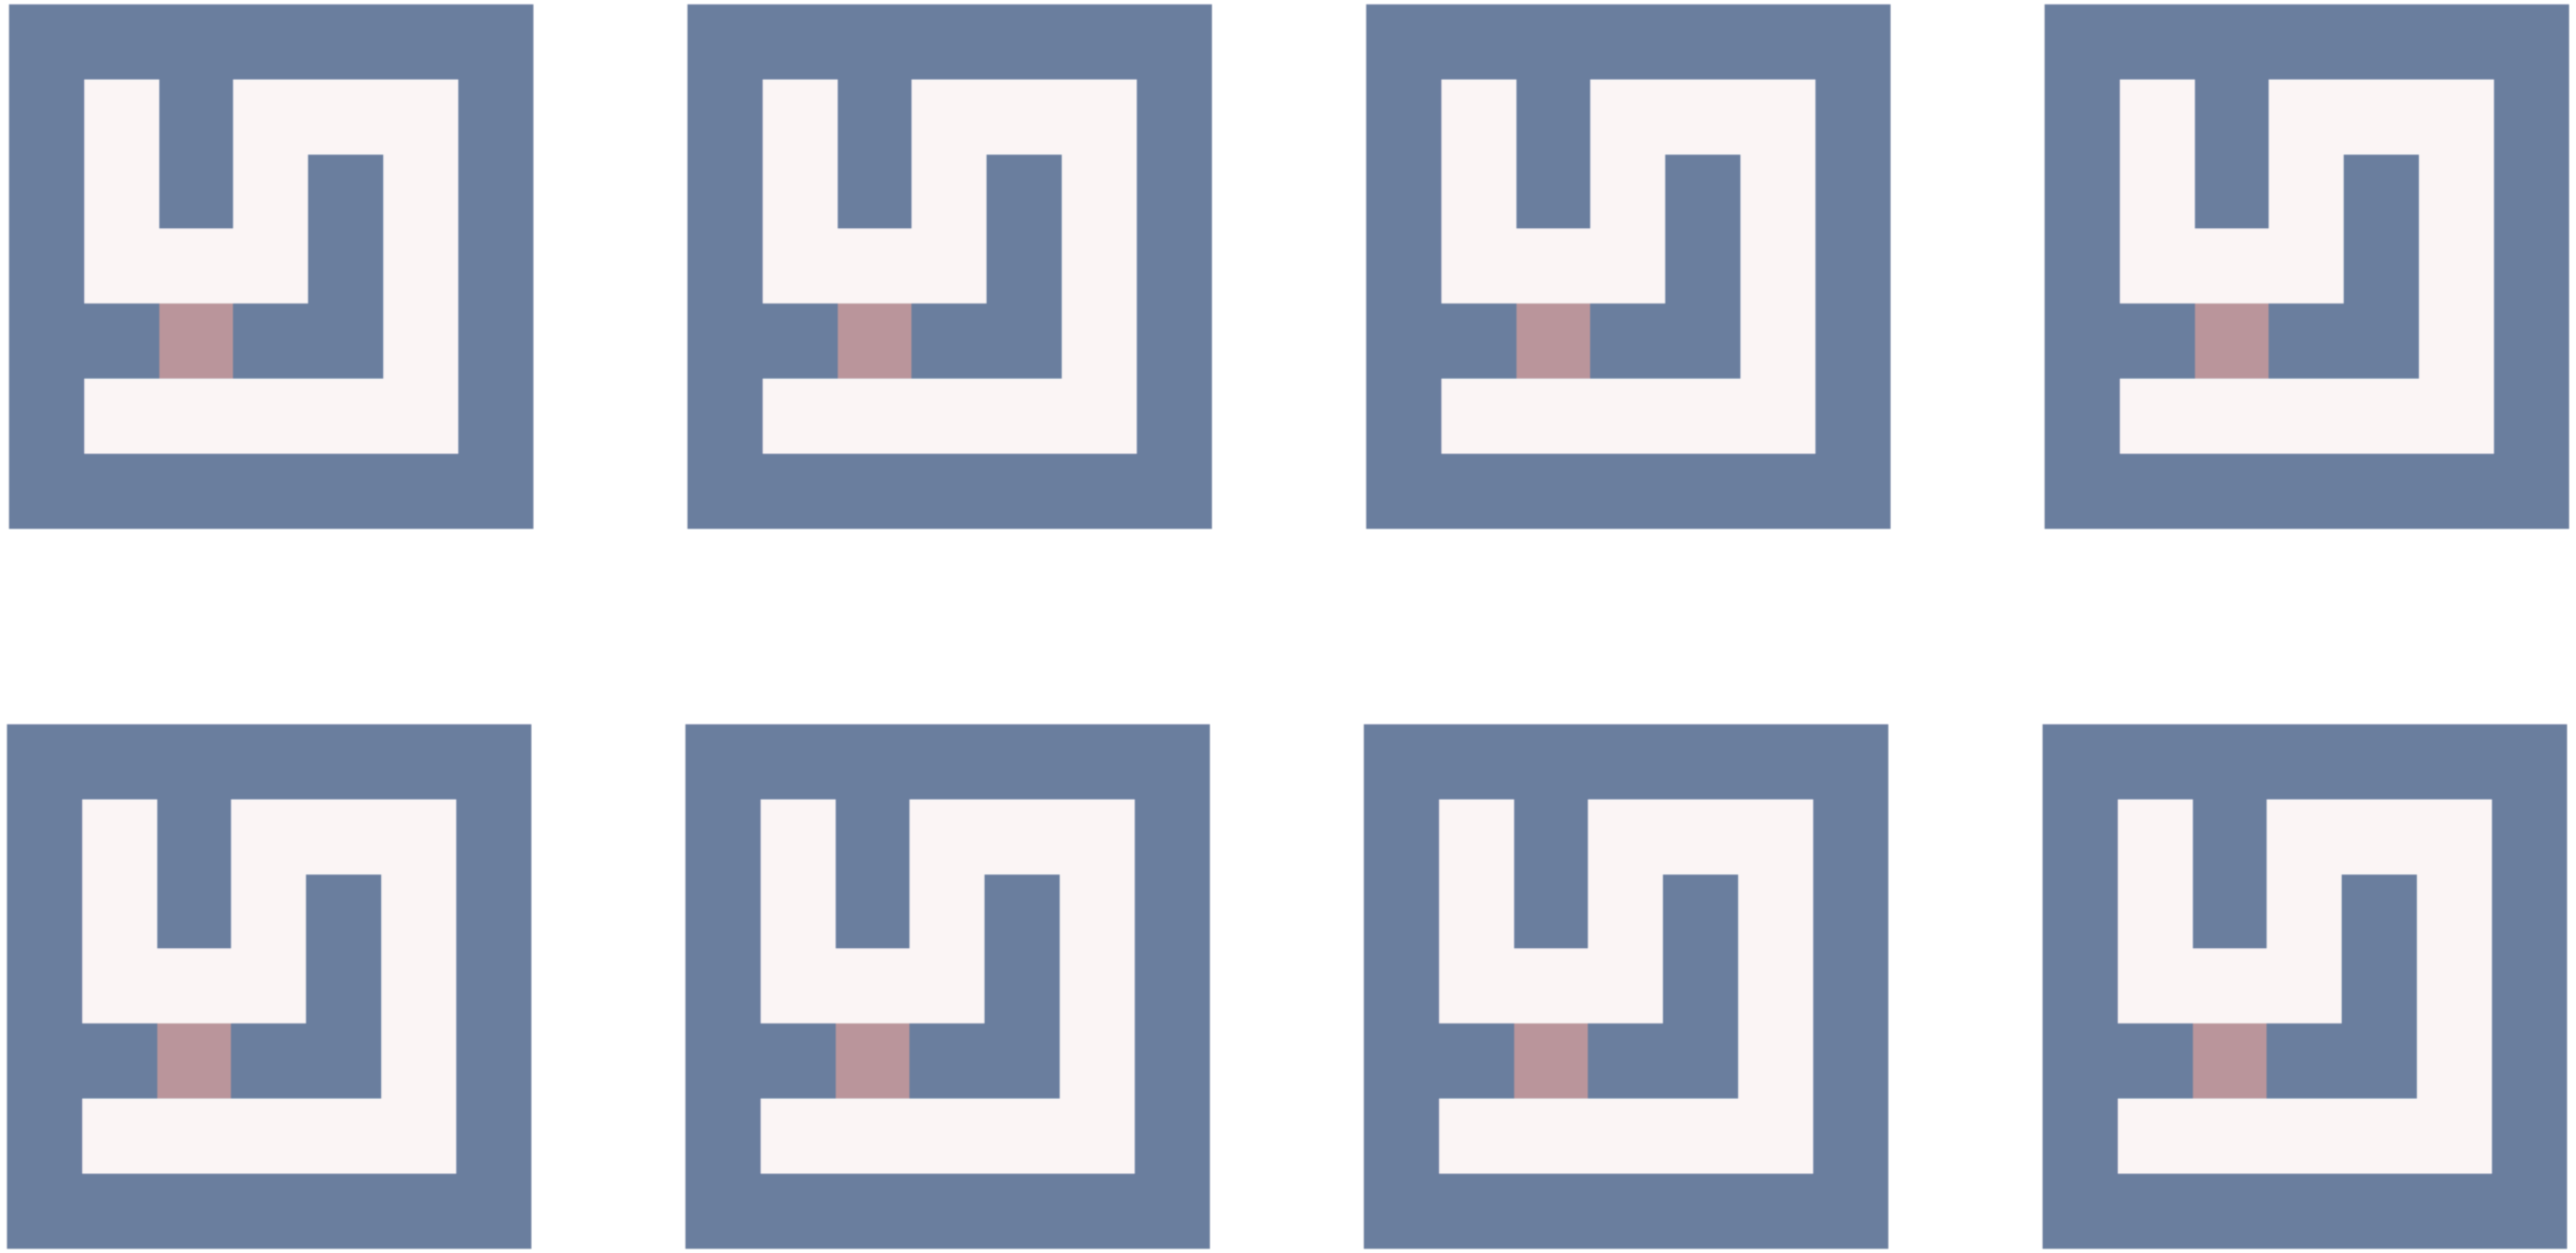

Symmetry: Rotations

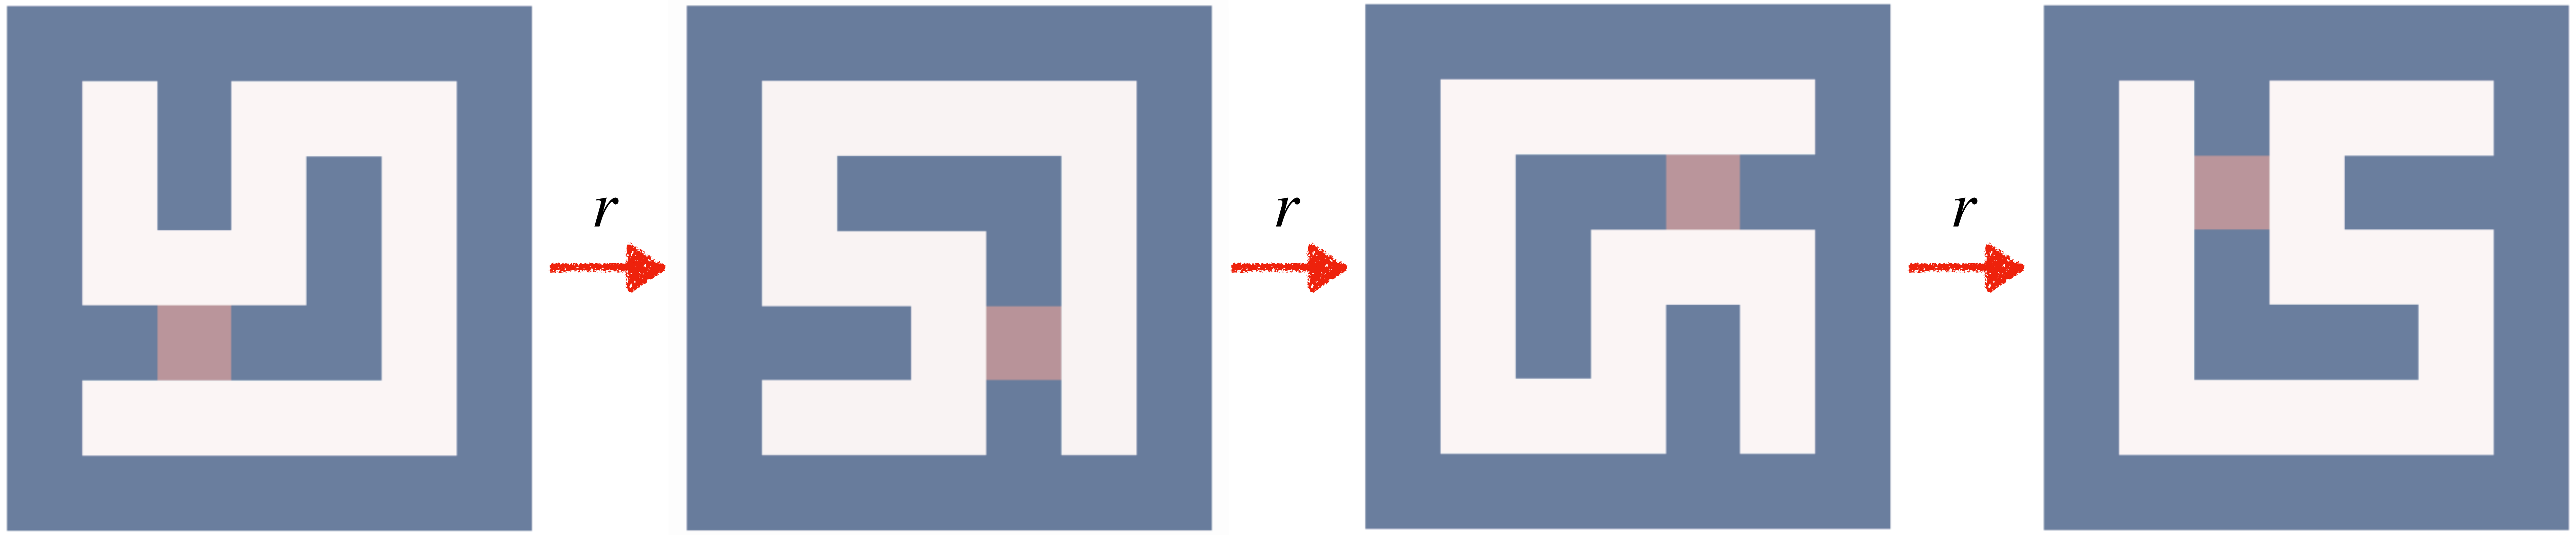

Symmetry: Rotations and Reflections

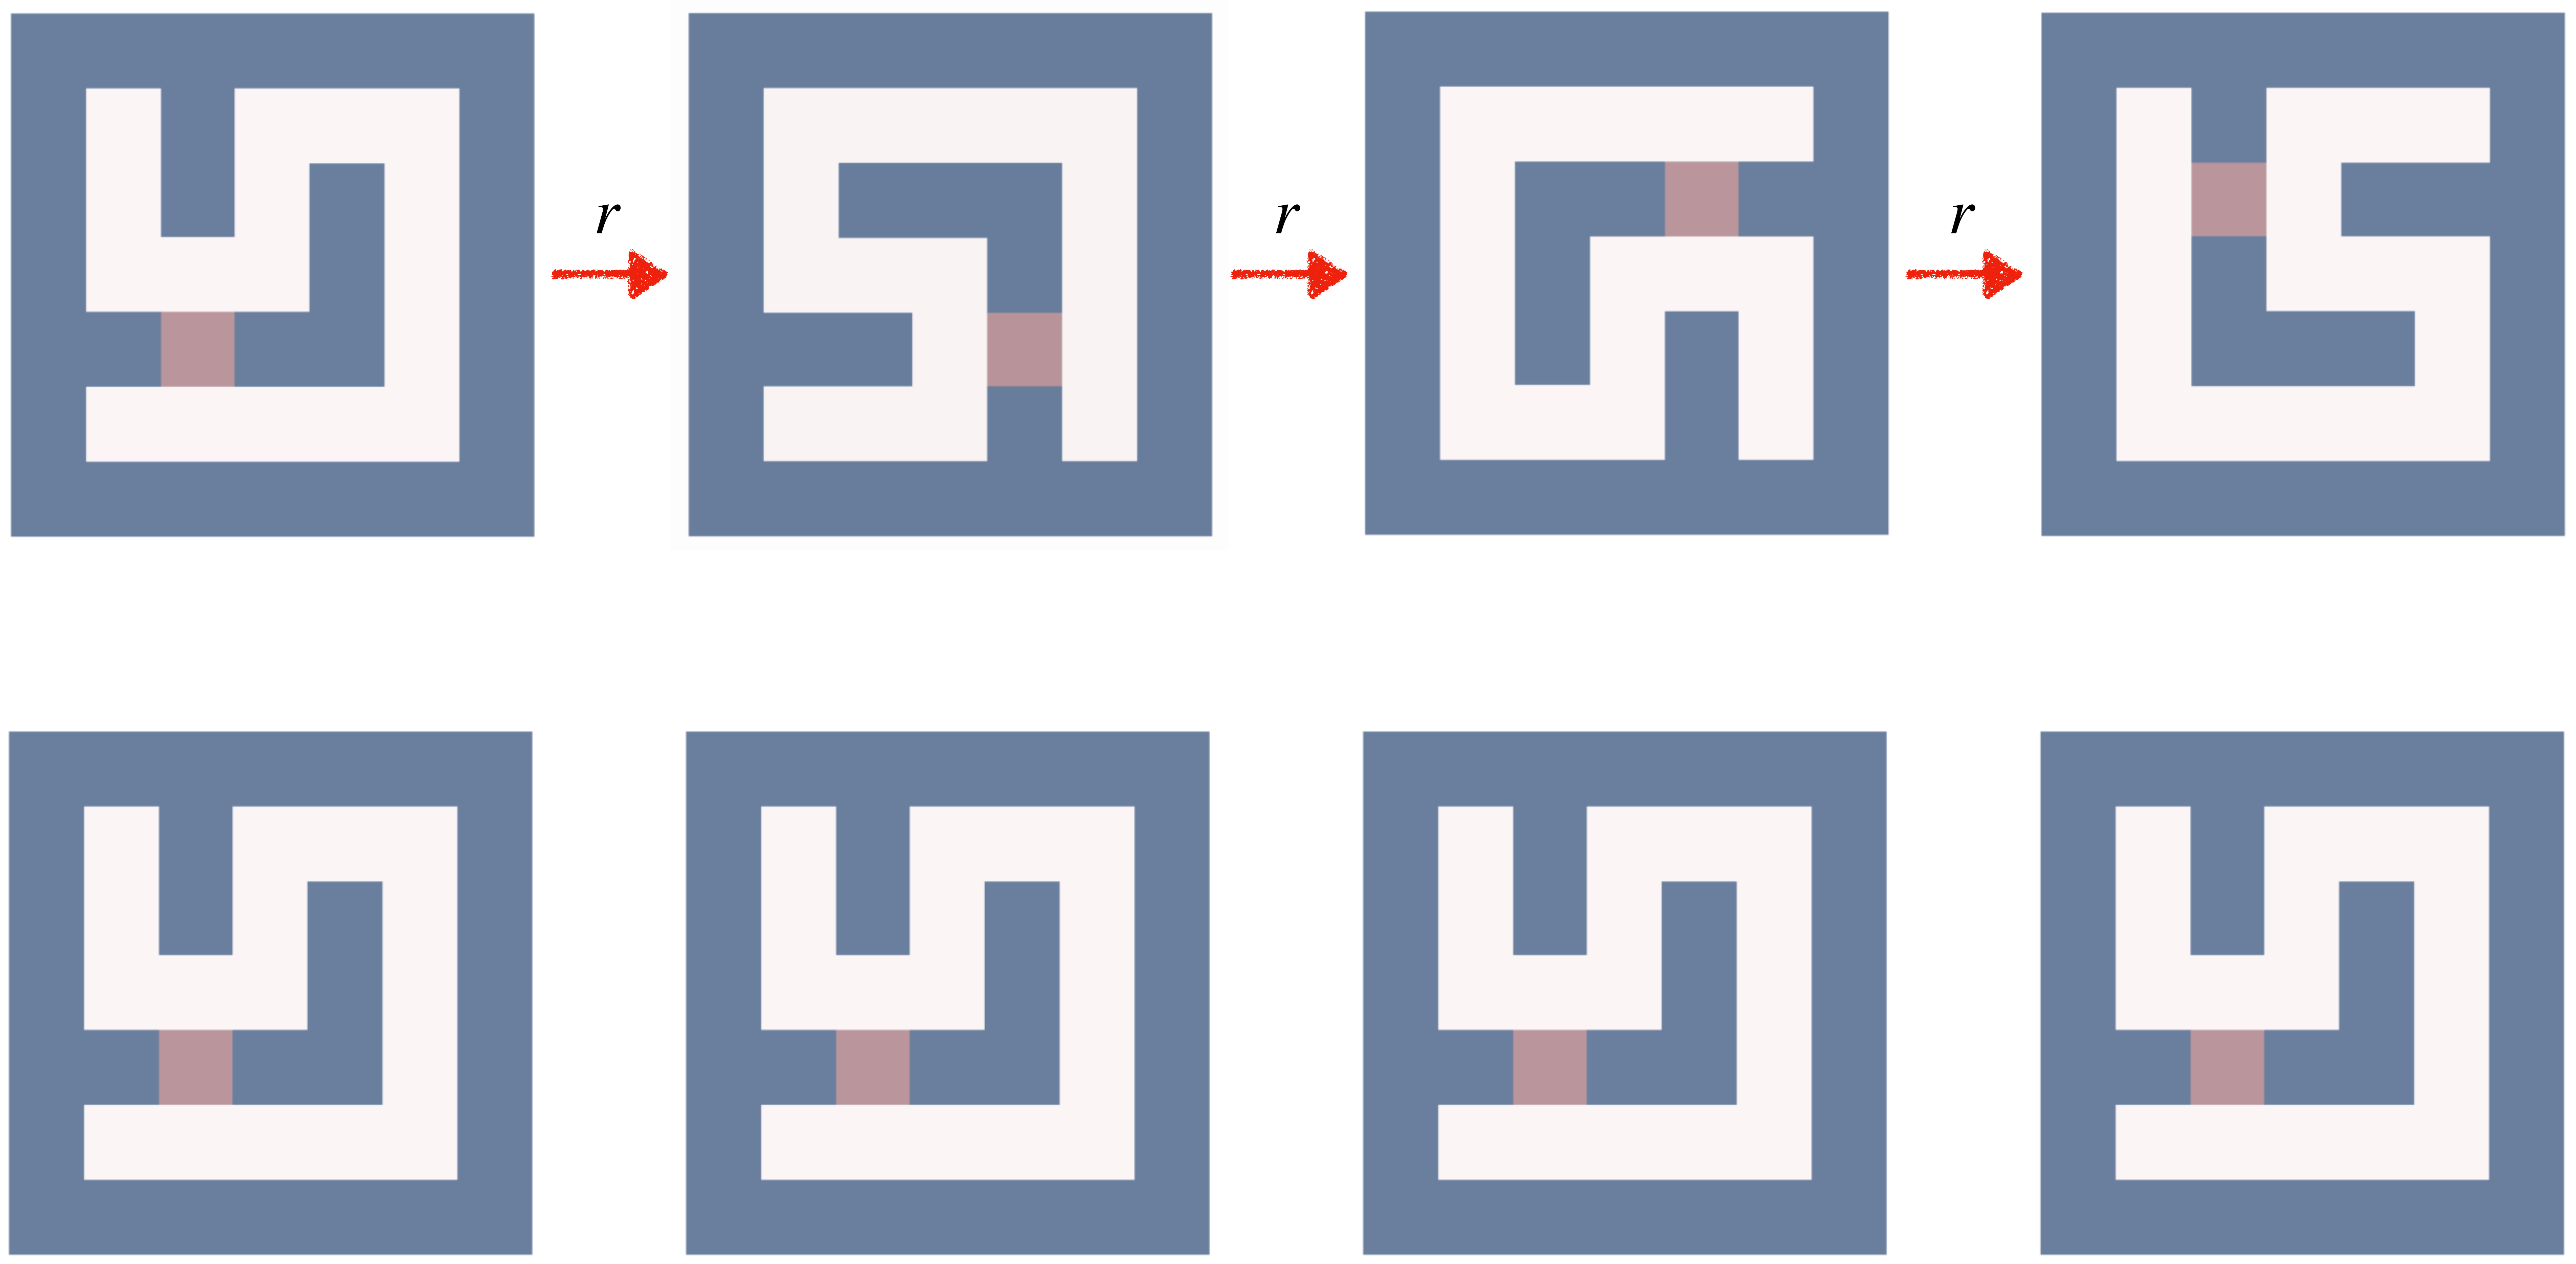

Symmetry: All 8 Transformations in D_4

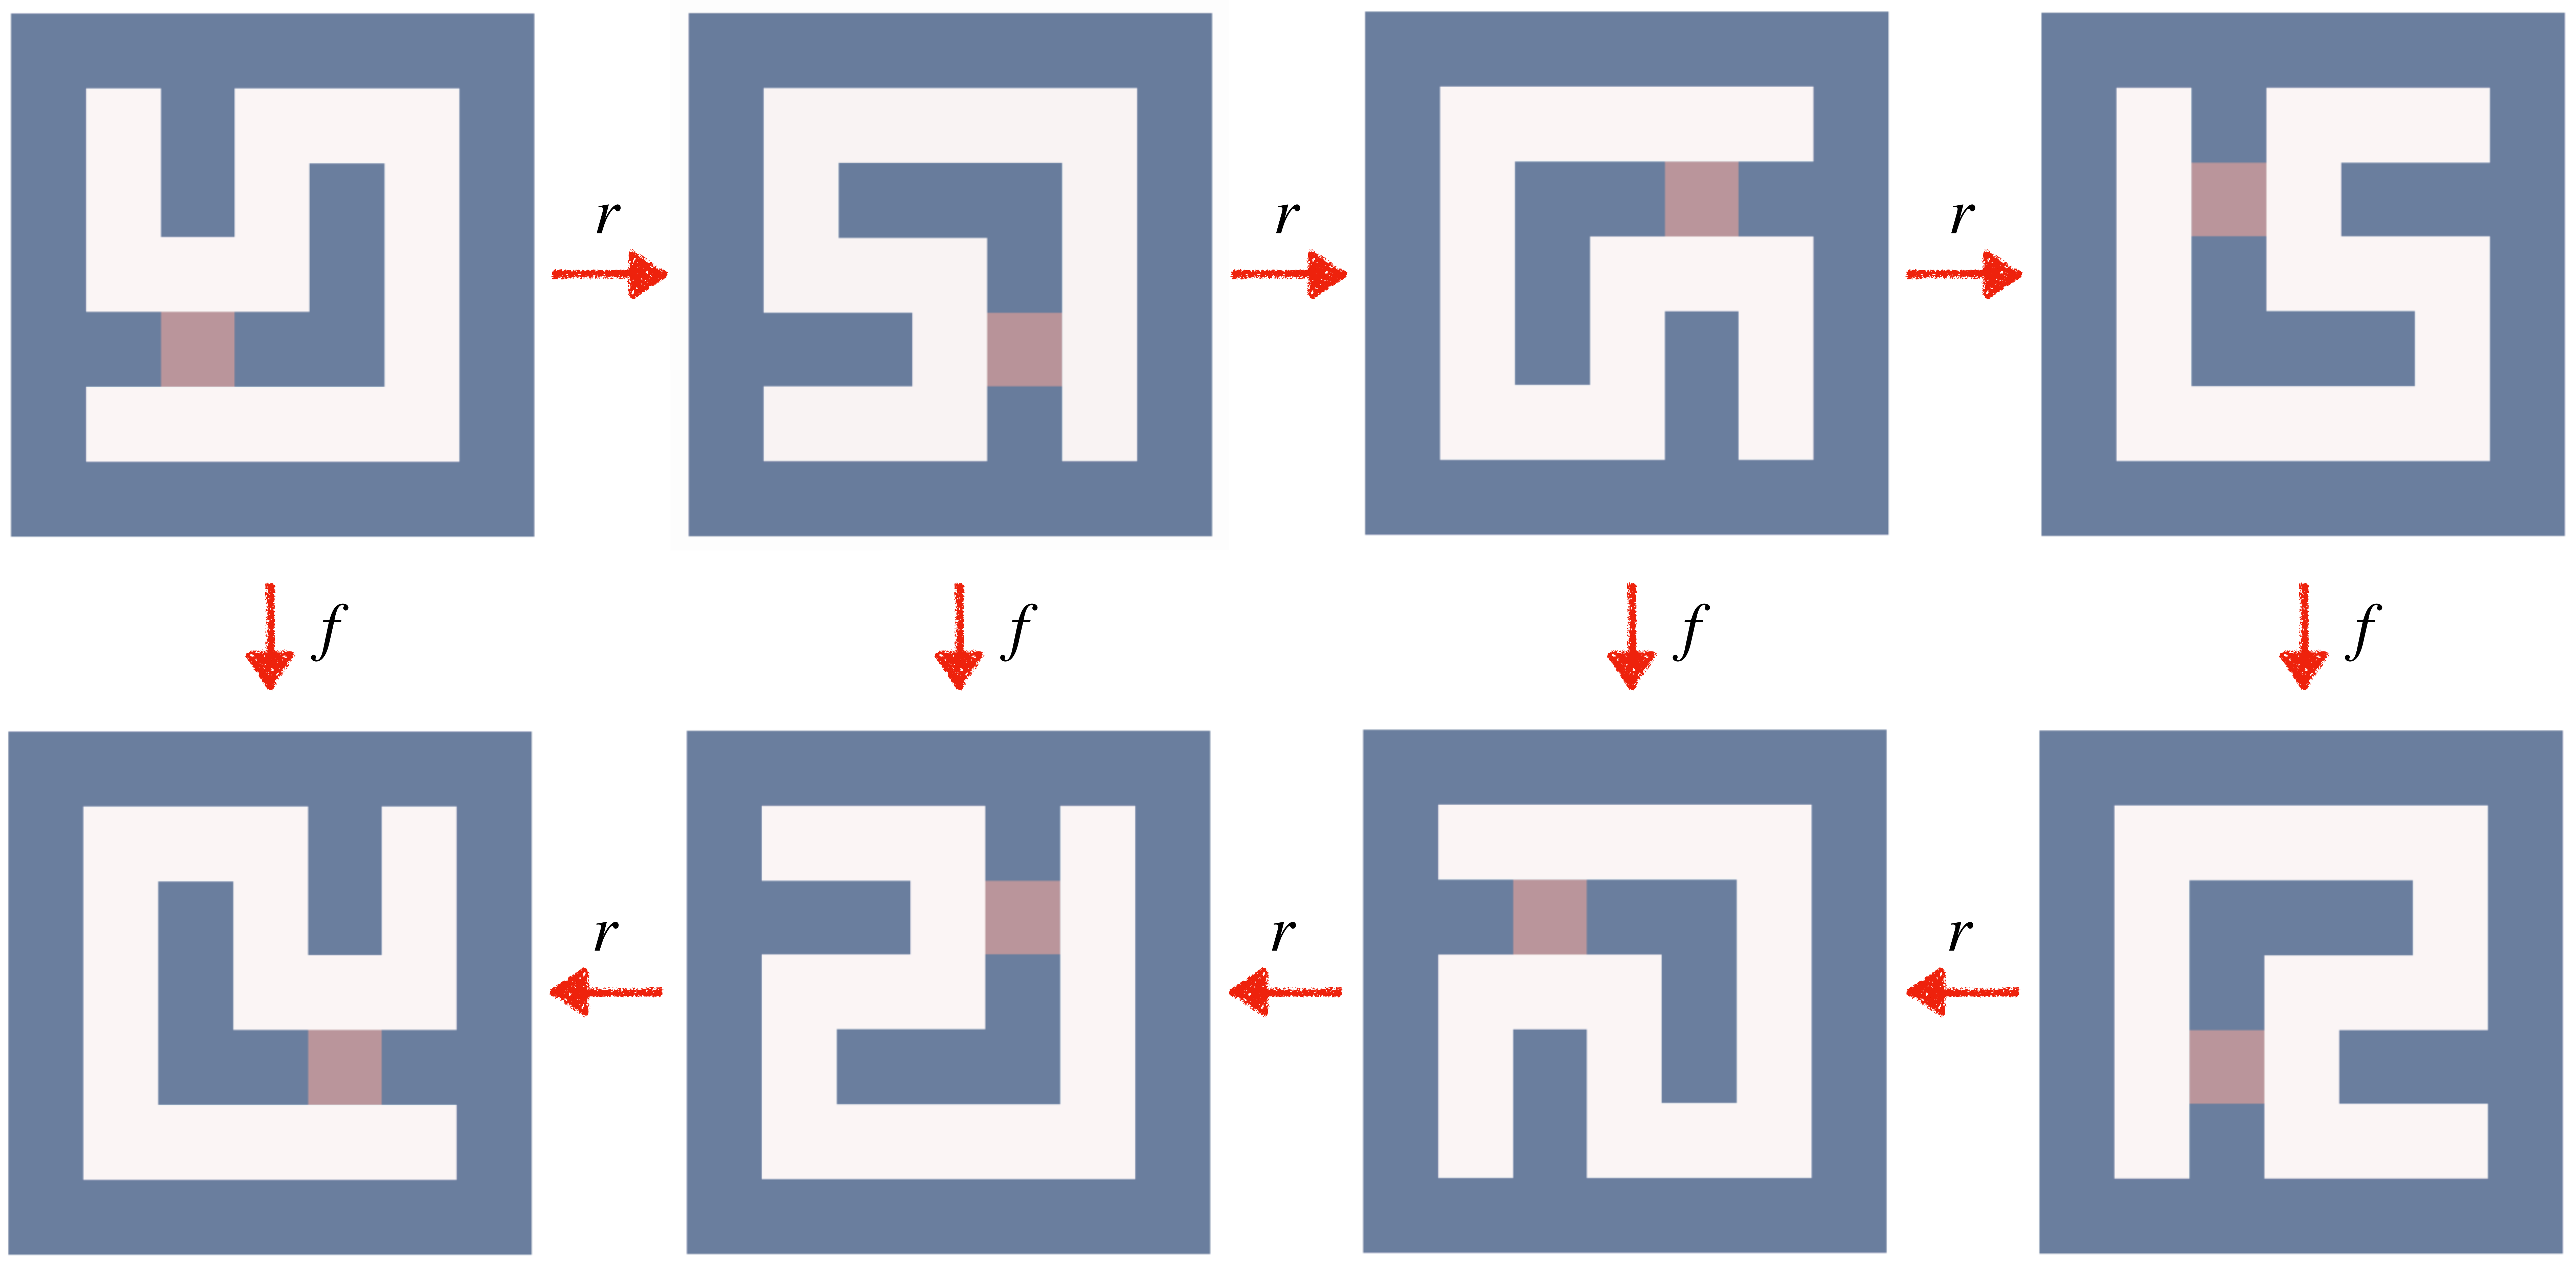

Value Iteration with Symmetry

Every update is equivariant
— Local Equivariance

Entire planning is equivariant
— Global Equivariance

$$\circlearrowleft 90^\circ \circ \textbf{VI}(M) \equiv \circlearrowleft 90^\circ \circ \mathcal{T}^\infty[V_0] = \mathcal{T}^\infty[\circlearrowleft 90^\circ \circ V_0] \equiv \textbf{VI}(\circlearrowleft 90^\circ \circ M)$$

- Use steerable convolution, equivariant to rotation and reflection:

$$\bar{Q}^{(k)} = \bar{R}^a + \textbf{Conv2D}(\bar{V}^{(k-1)}; W_{\bar{a}}^V) \xrightarrow{\text{Replace}} \bar{Q}_{\bar{a}}^{(k)} = \bar{R}_{\bar{a}} + \textbf{SteerableConv}(\bar{V}; W^V)$$

Main Pipeline: Symmetric Value Iteration Network

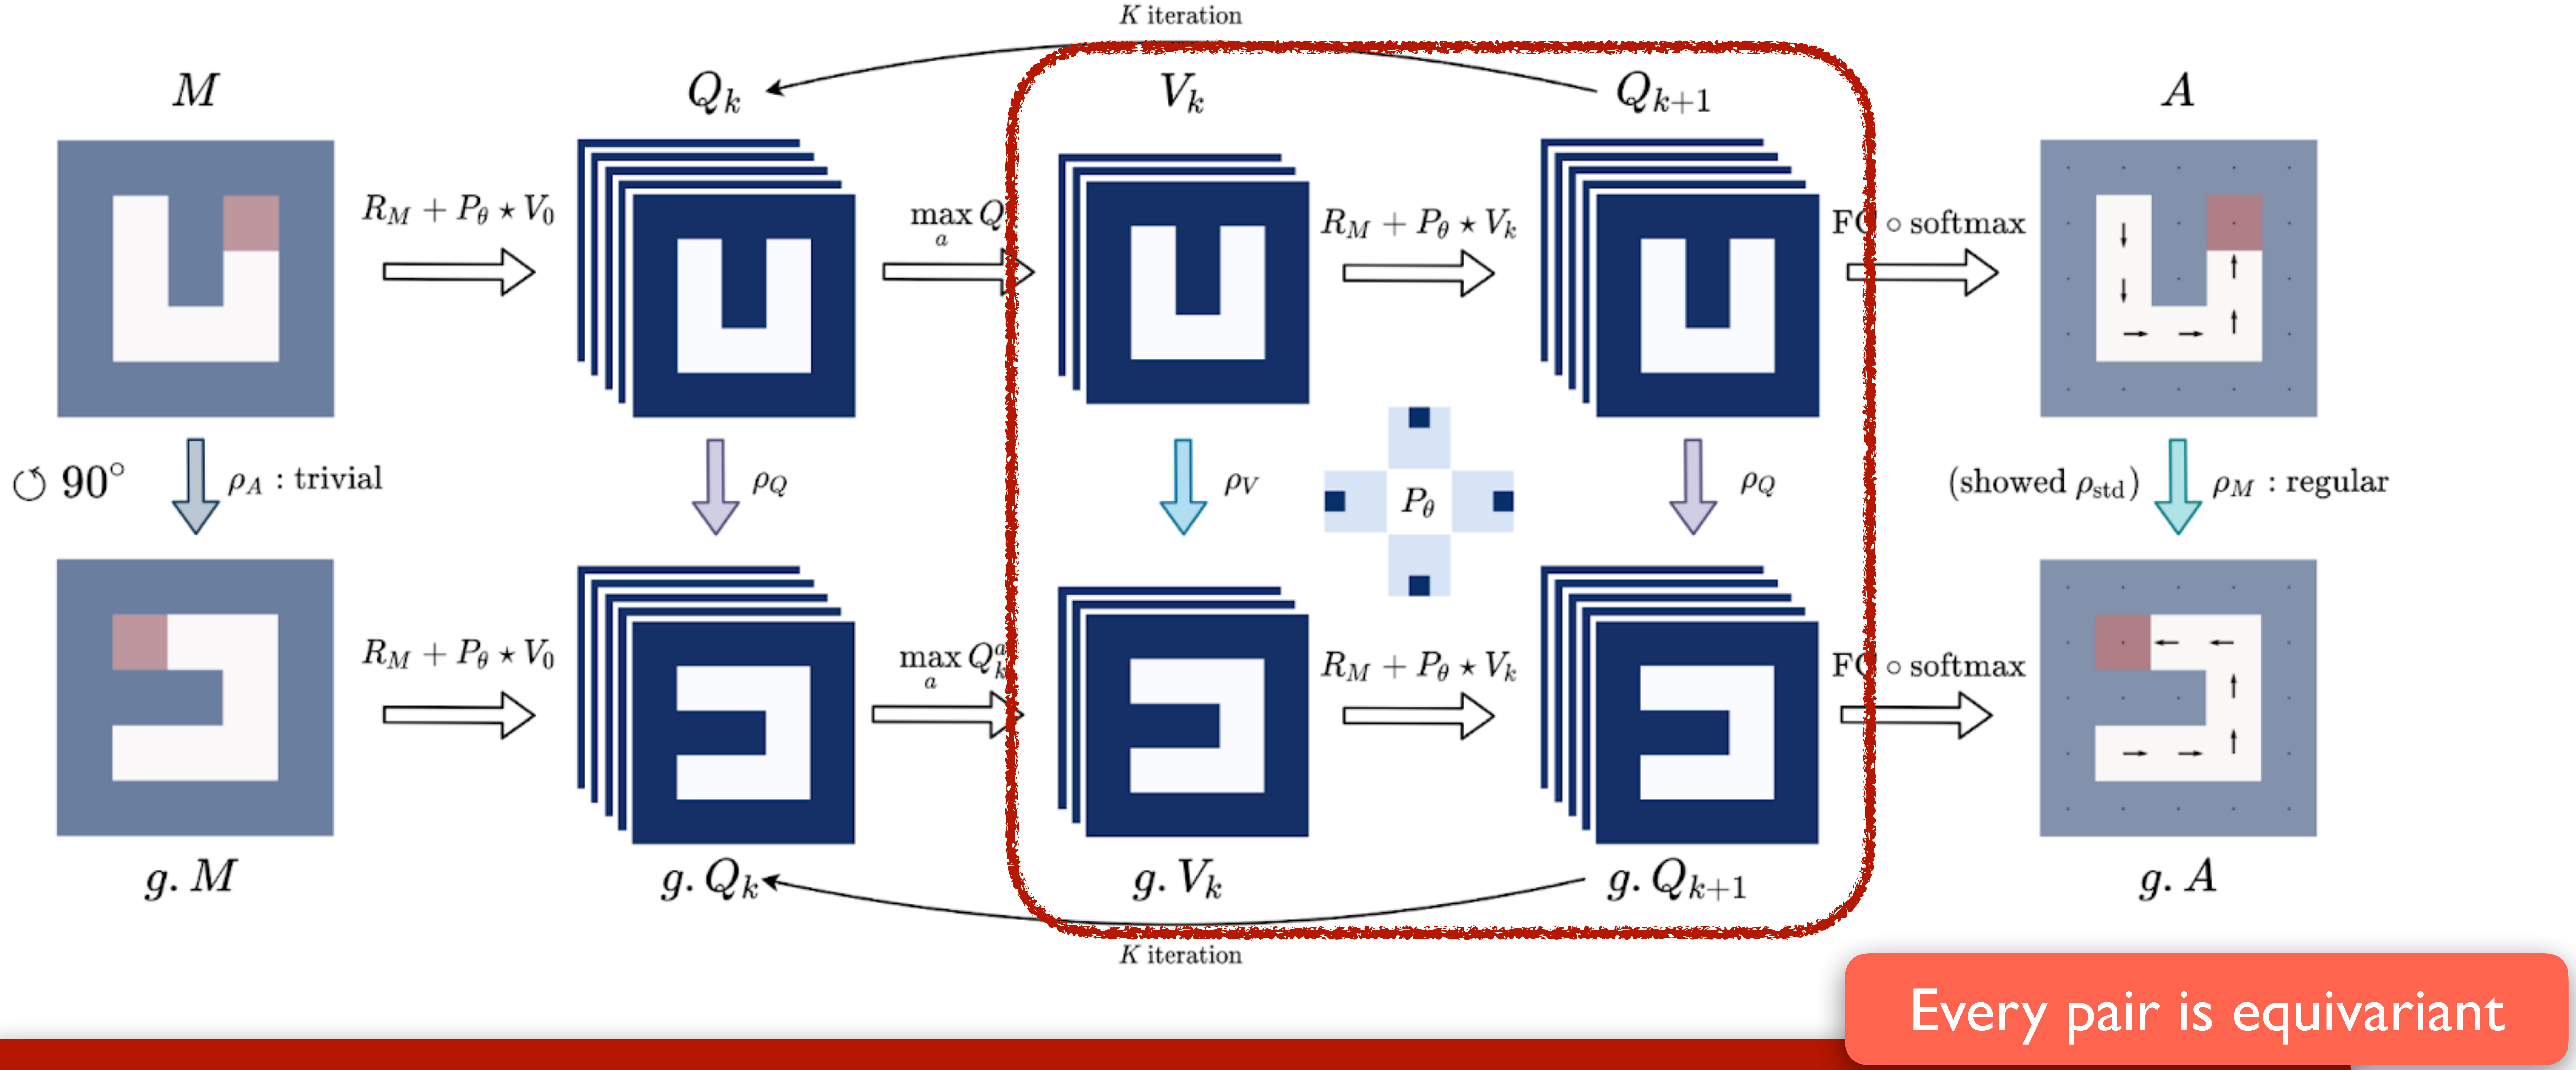

We use steerable convolutions to integrate symmetry in VINs.

Key Insights

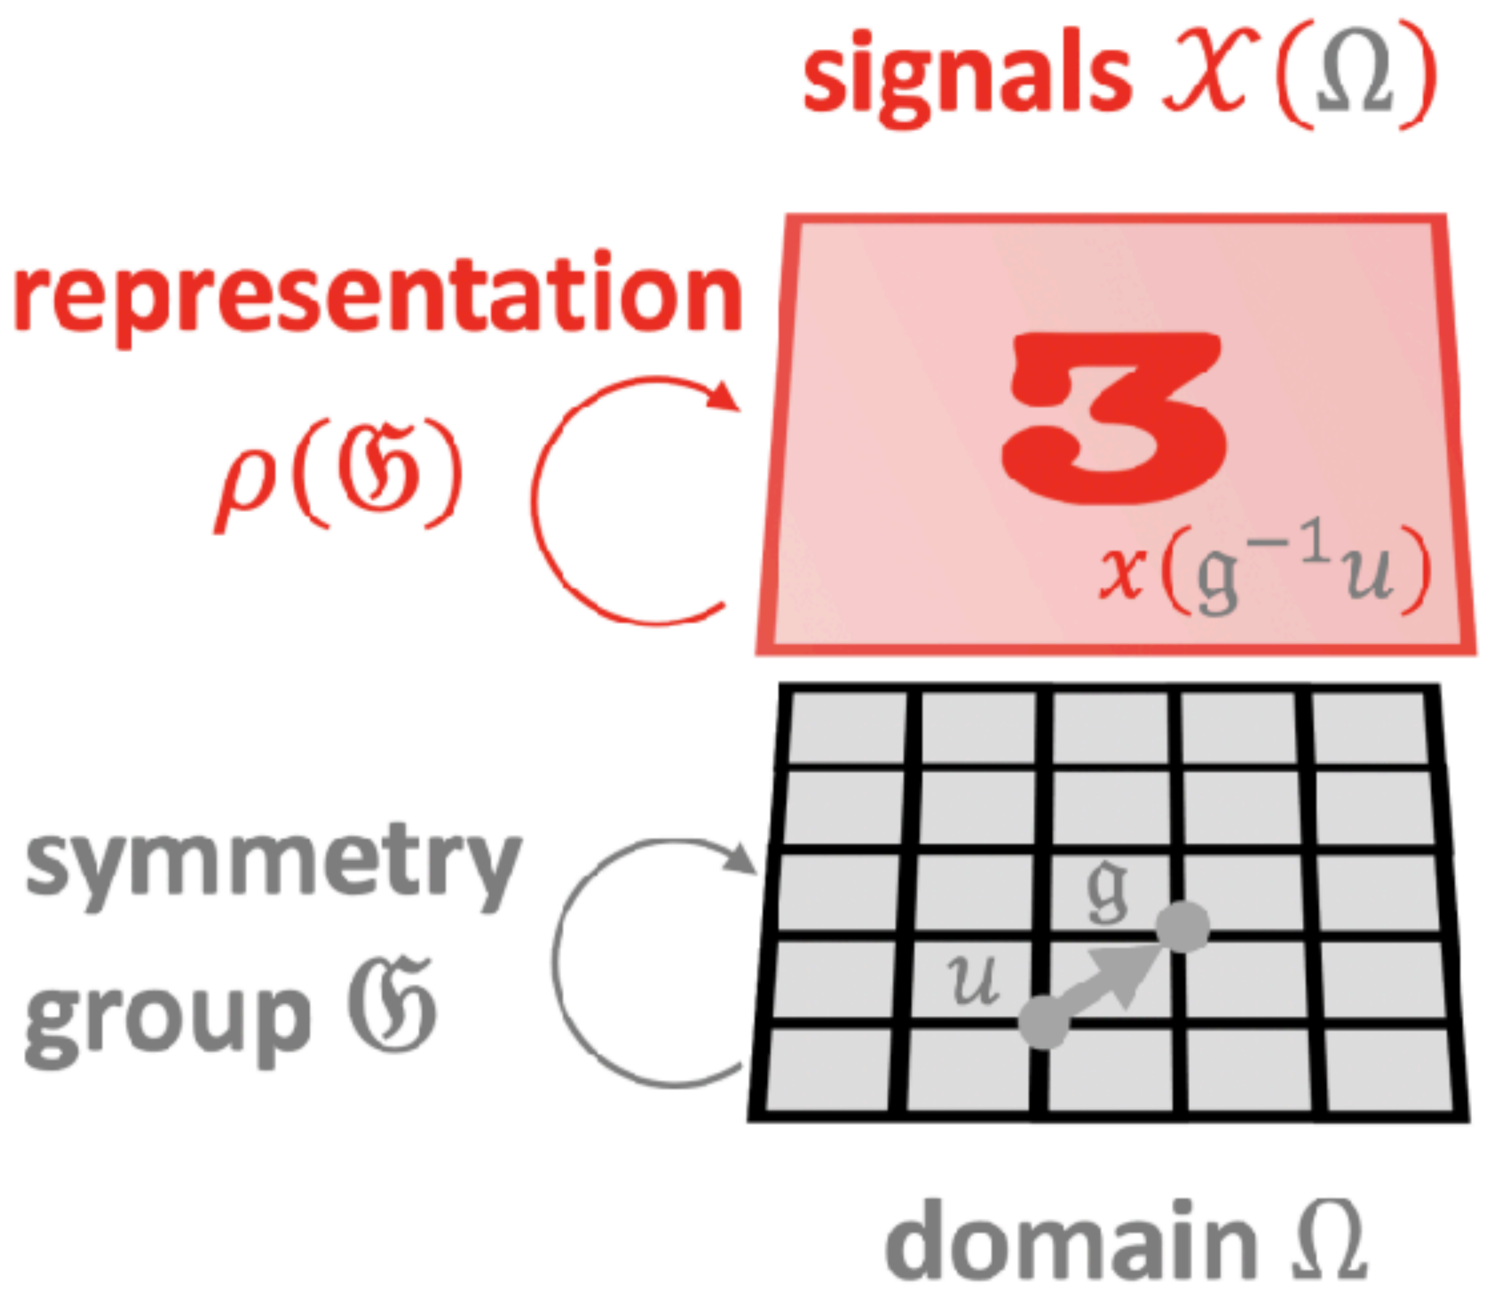

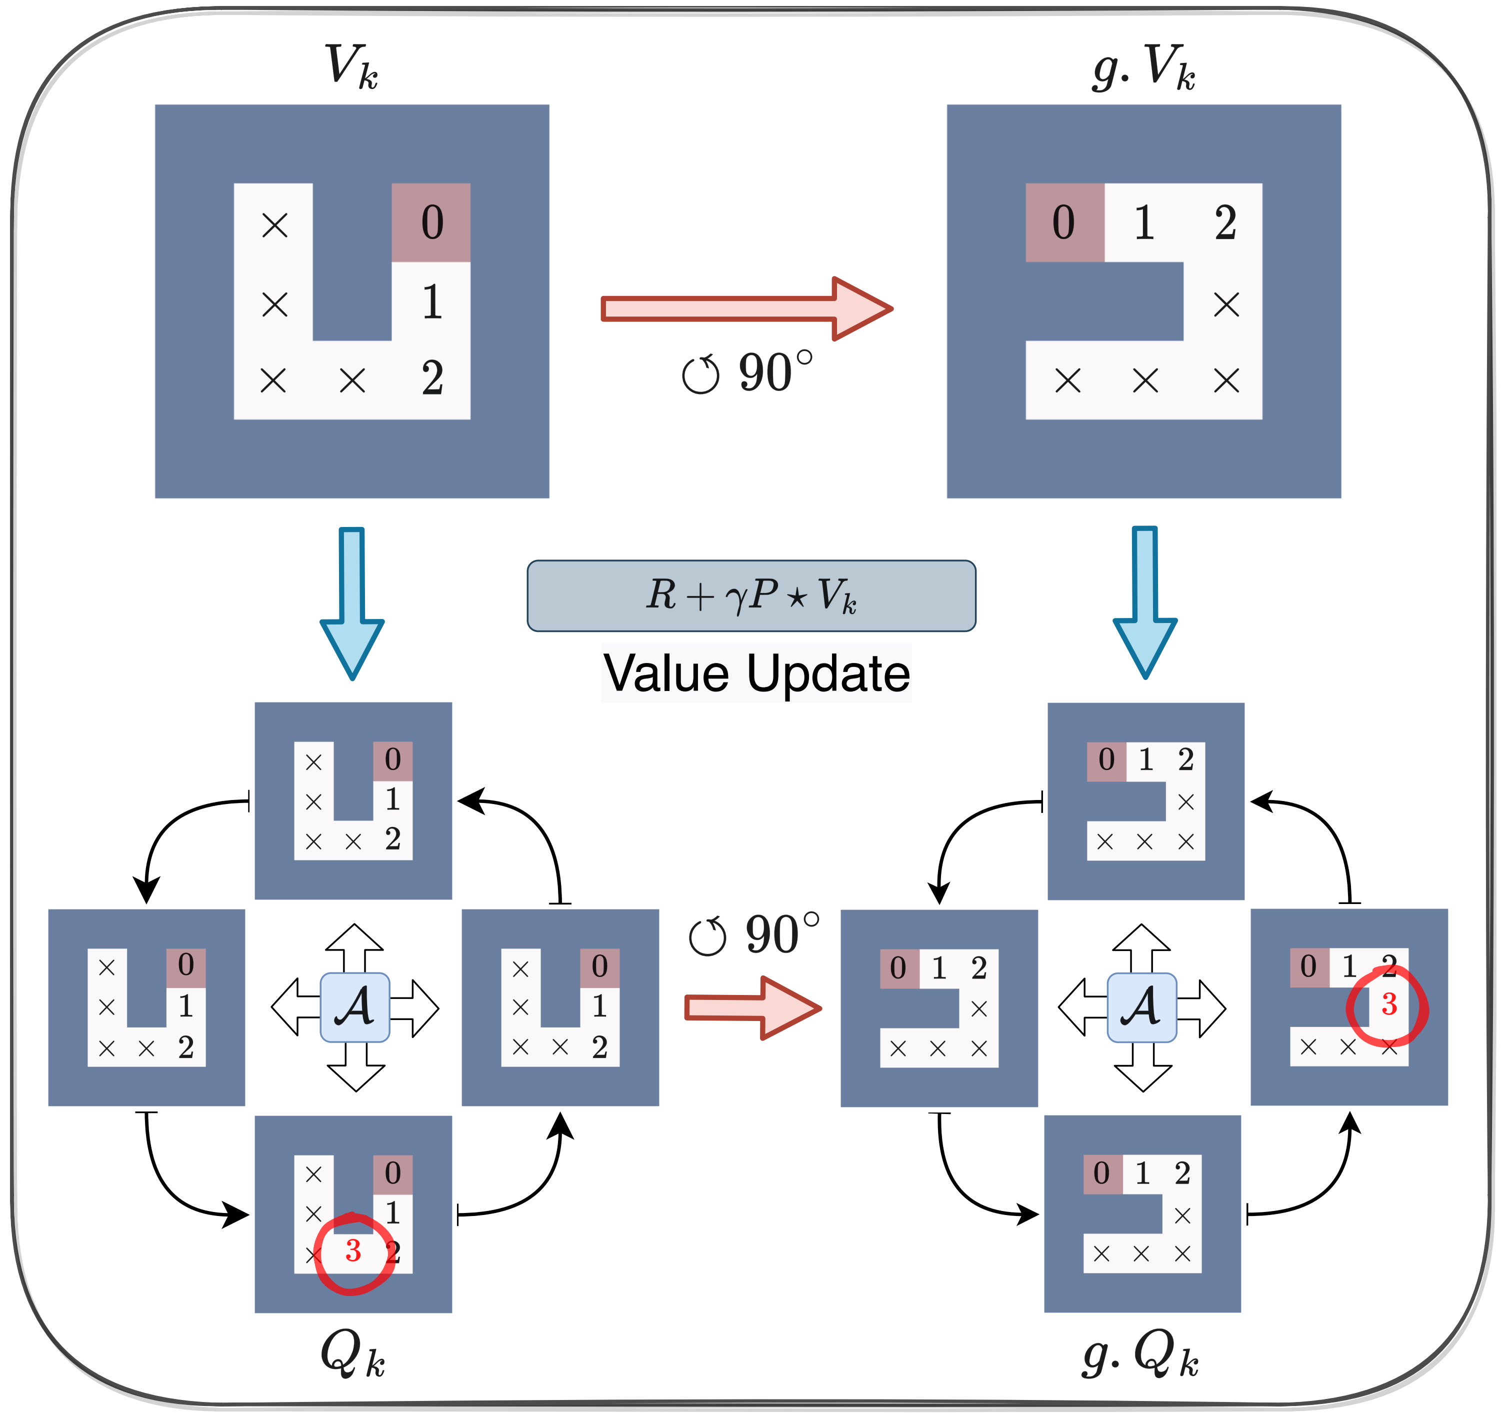

1: Represent (value) functions as "fields"

2: Value iteration as convolution (network)

Bronstein et al. (2021): Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges. arXiv.

Theoretical results

Theorem 1 (informal): **Value iteration** for **path planning*** is equivariant to translation, rotation, and reflection

Theorem 2 (informal): **Value iteration** for **path planning*** is a form of **steerable convolution network****

*: Path planning on 2D grid, an example of homogeneous spaces
**: Steerable CNN over grids, equivariant under induced representations

Cohen et al. (2017): Steerable CNNs, ICLR 2017

Experiment: Setup

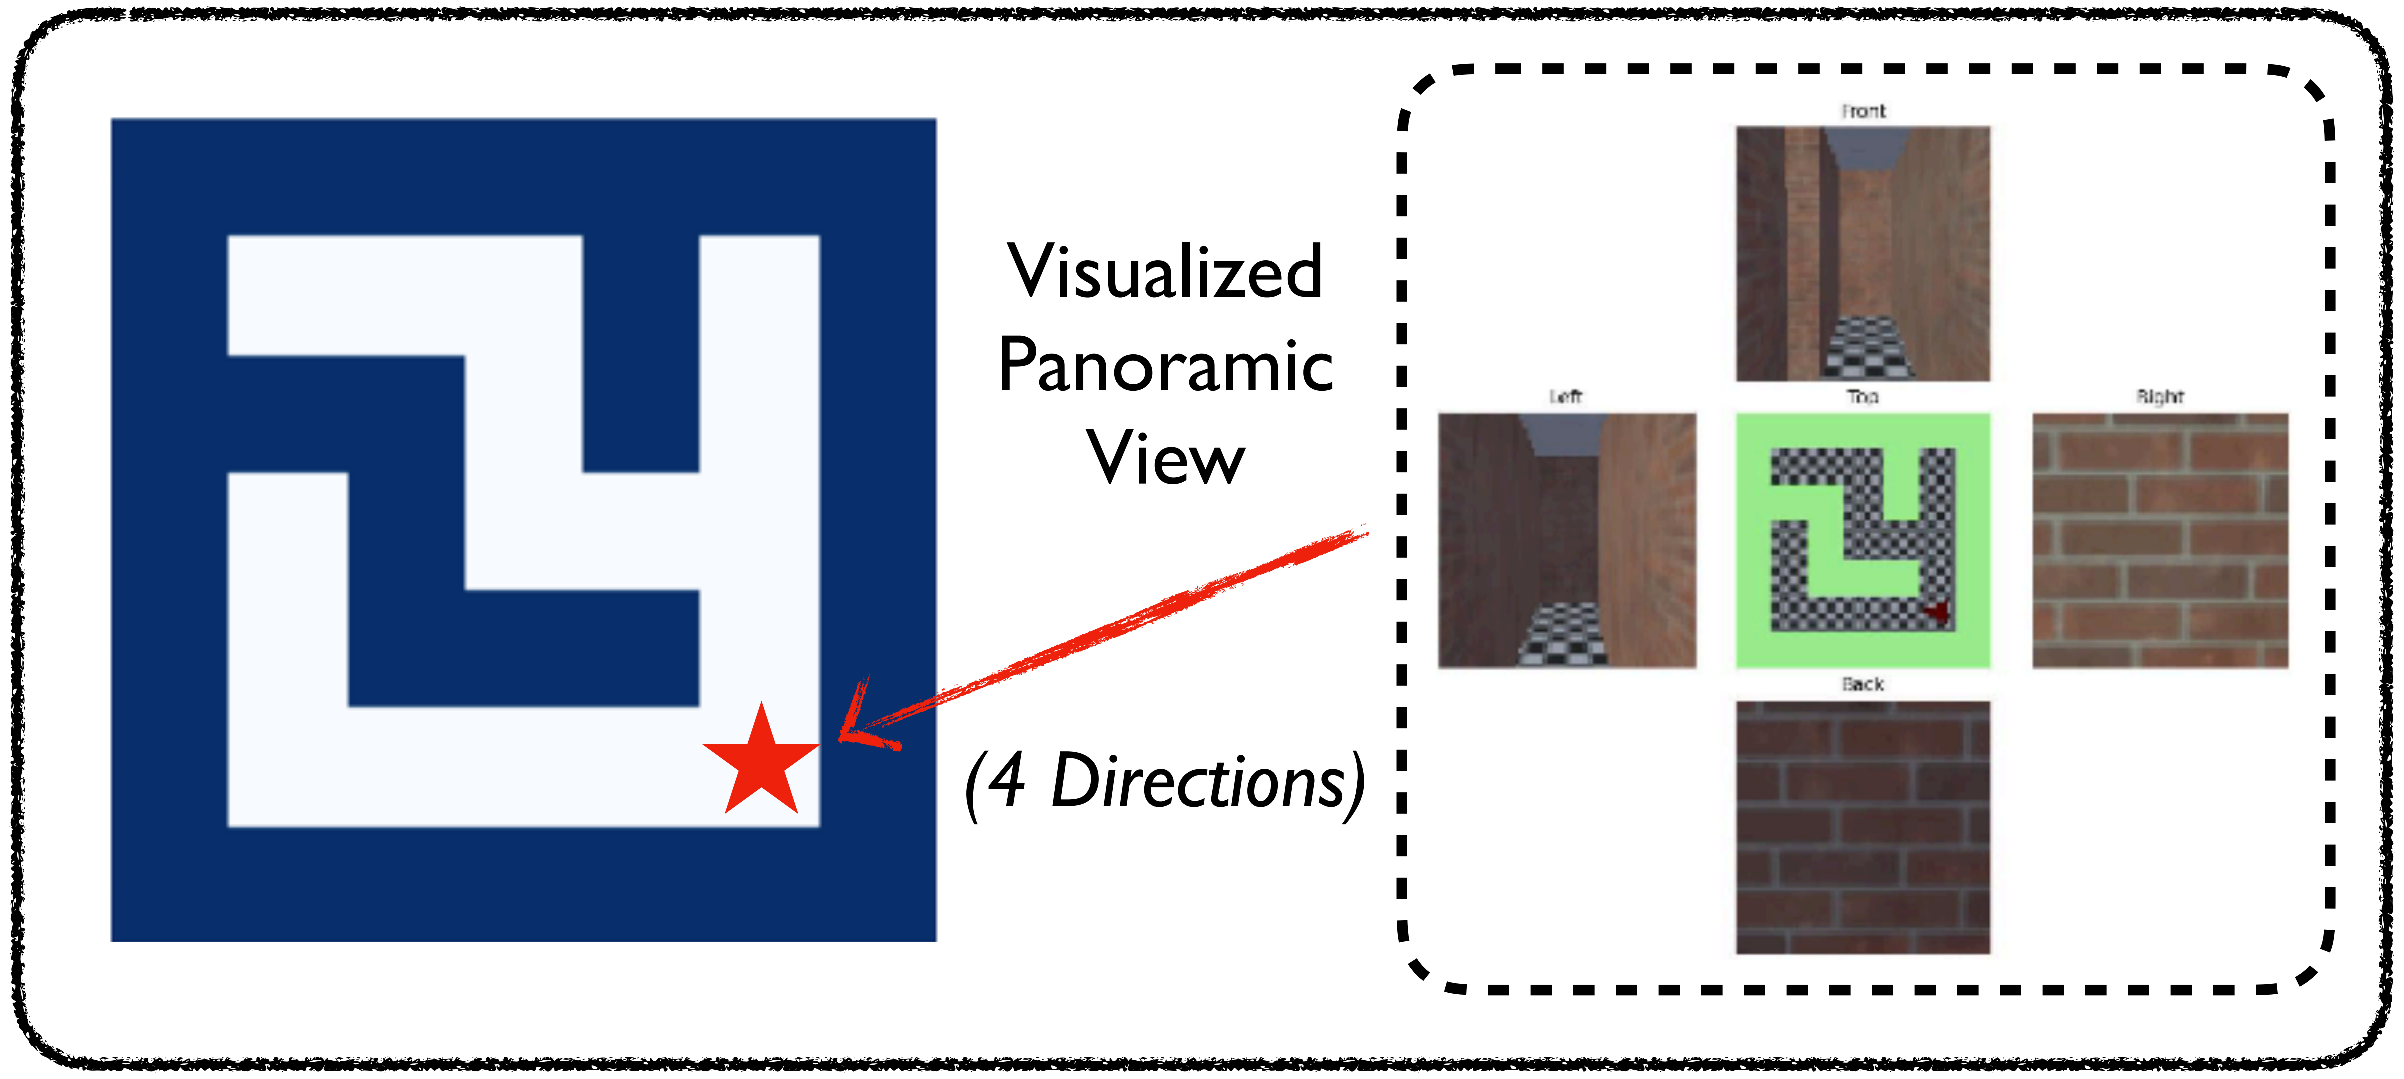

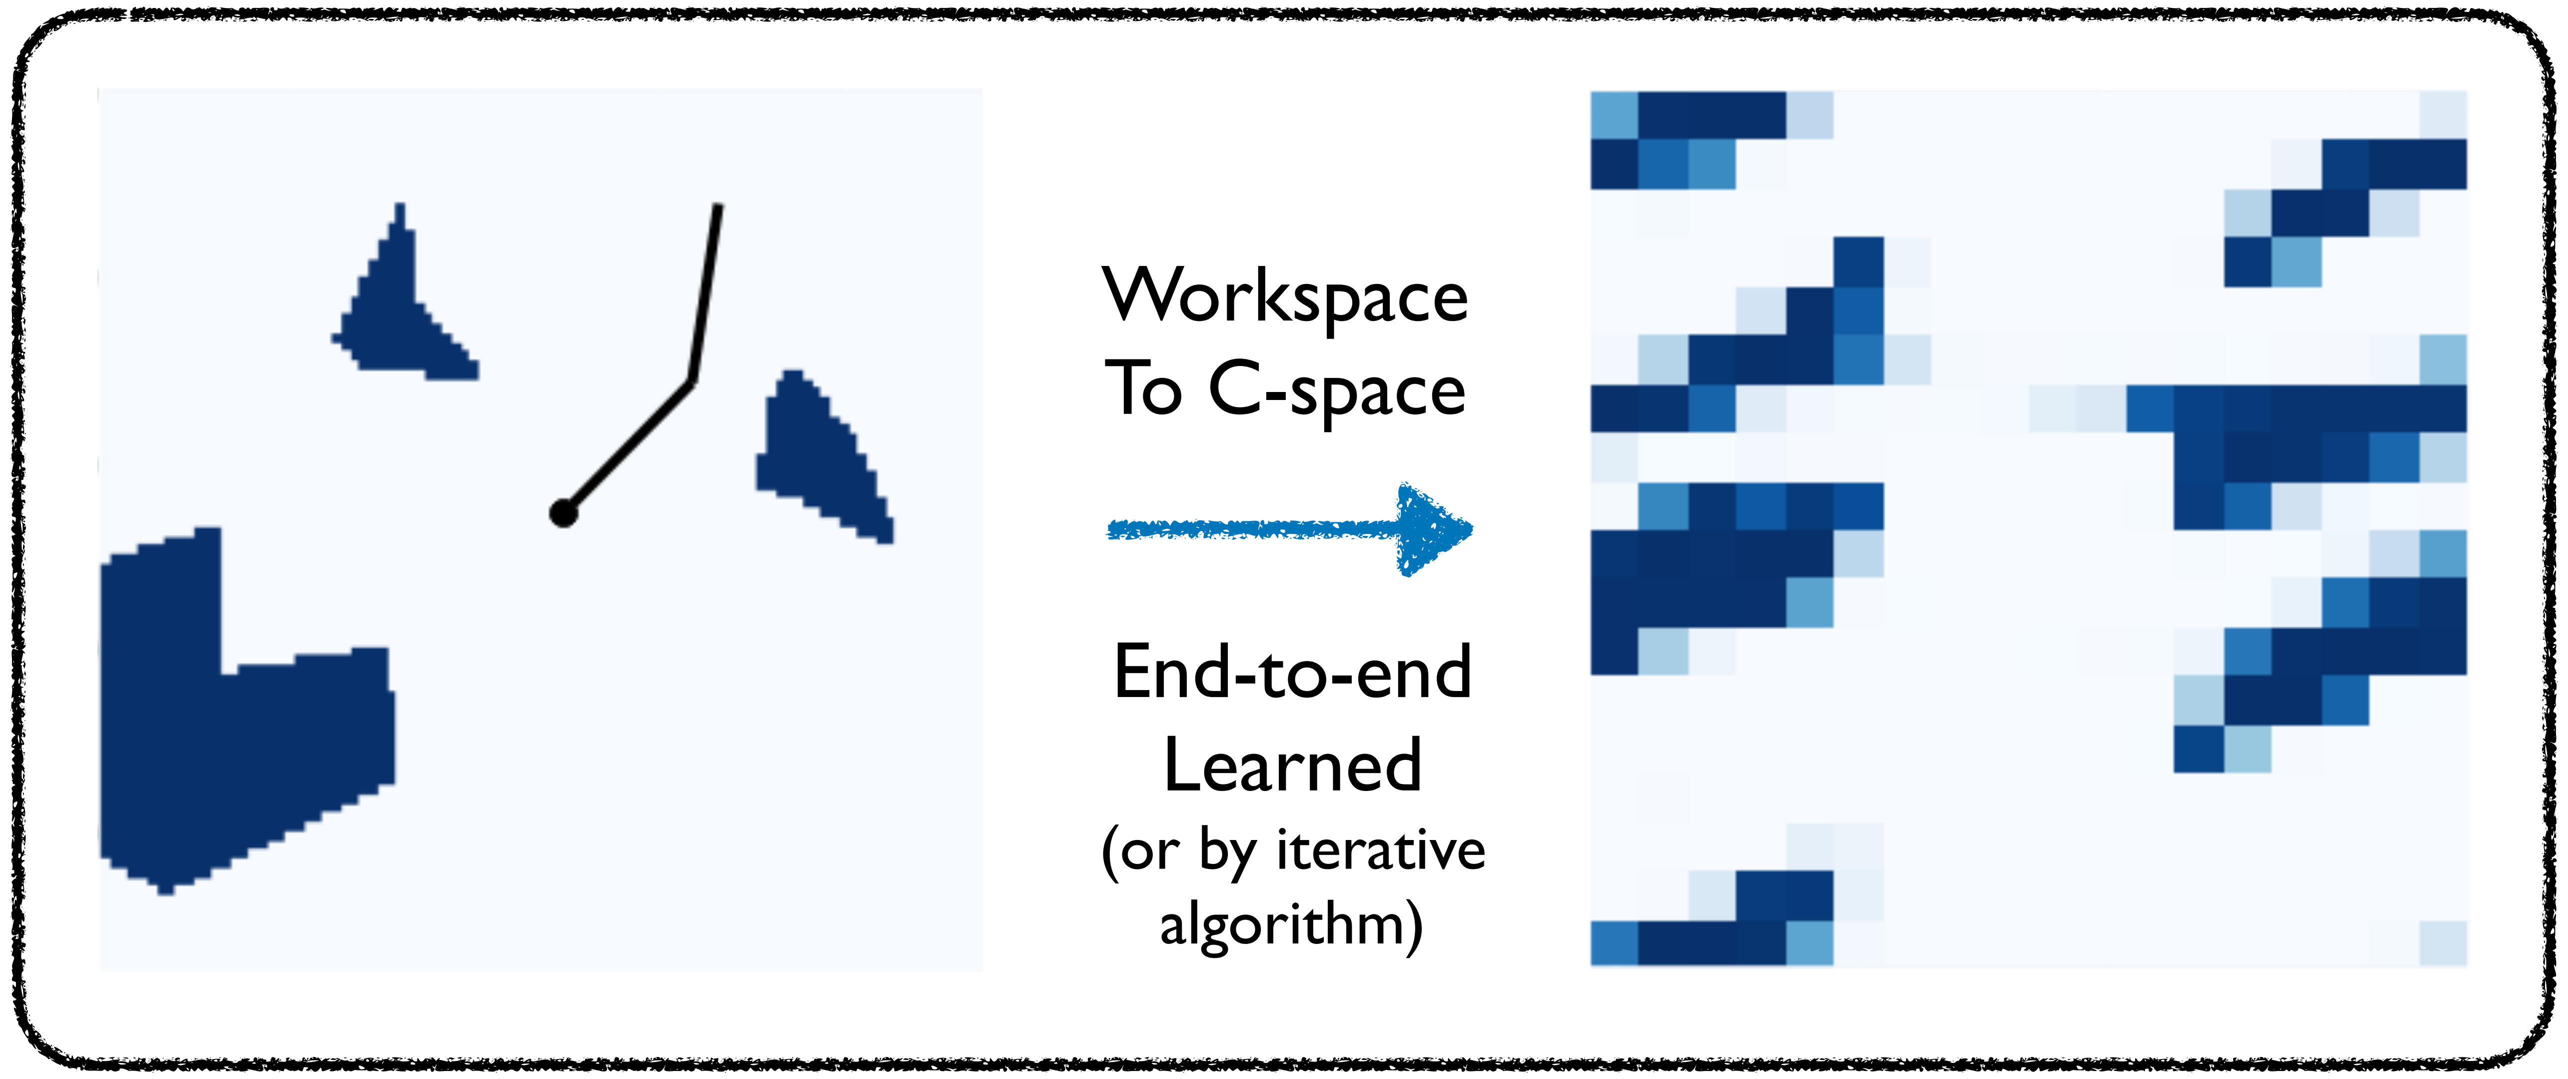

2D and Visual
Maze Navigation

2-DOF Manipulation
In Workspace and C-space

Results: Training

2D Maze Navigation

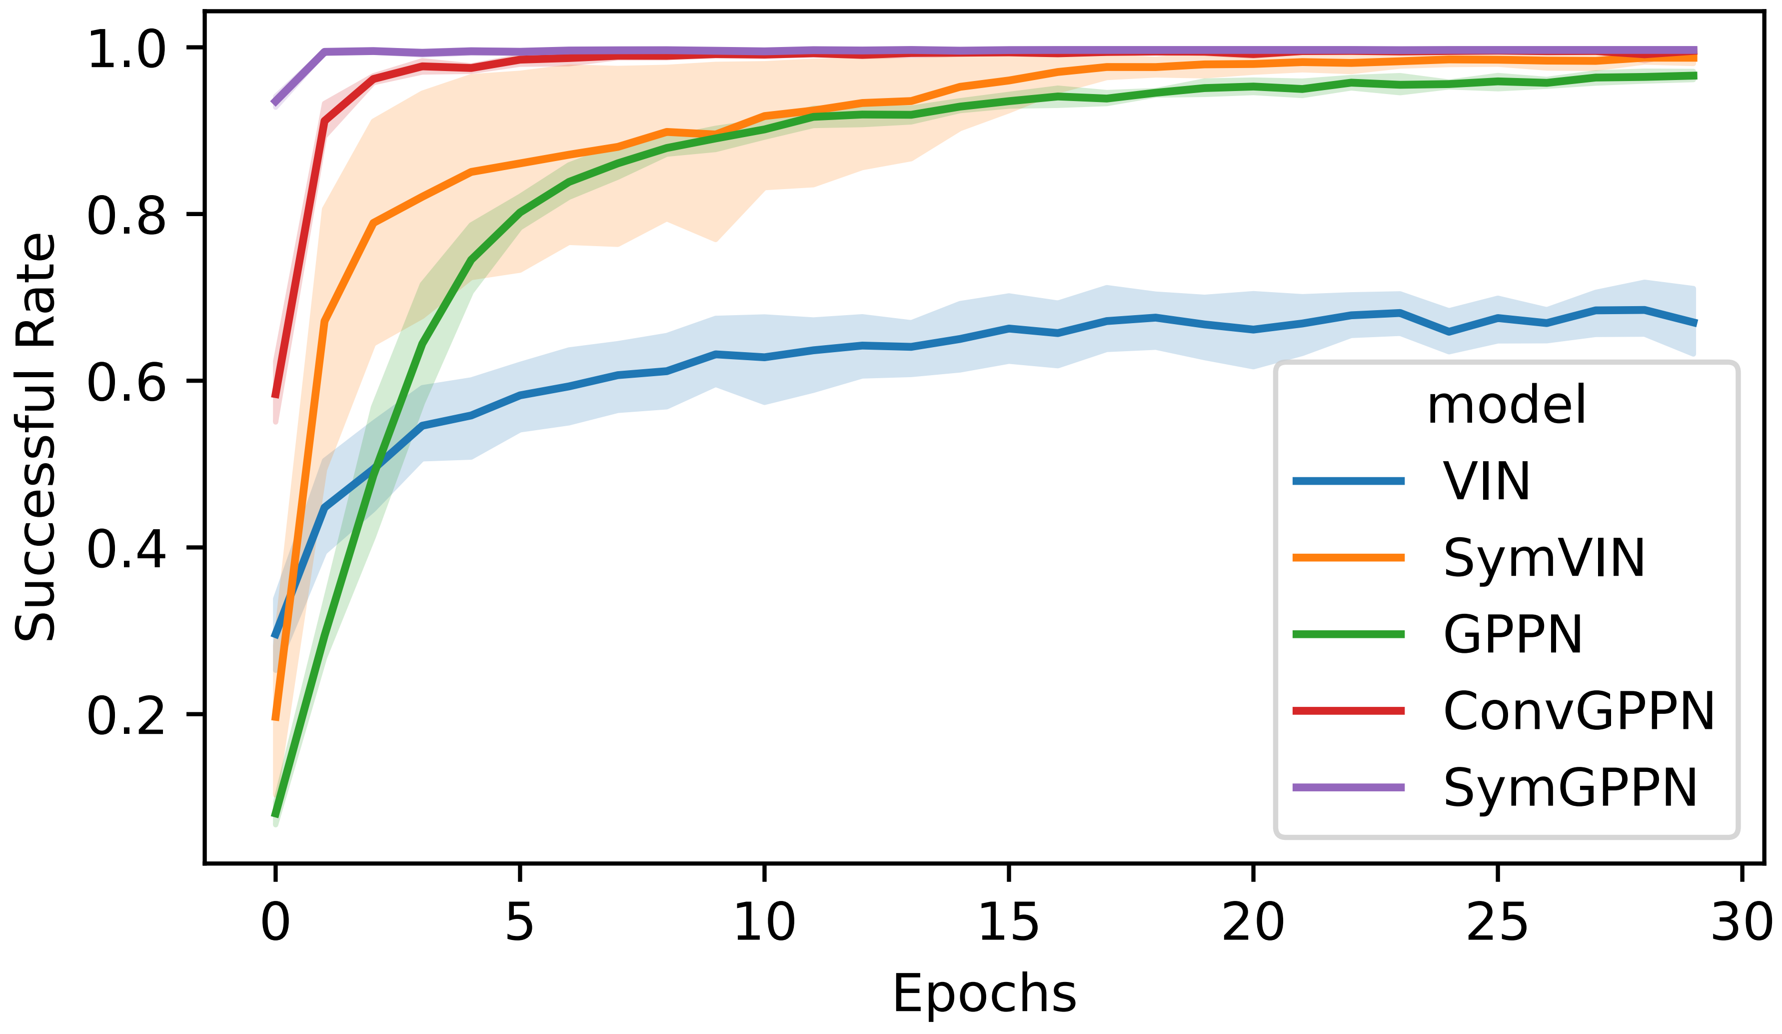

2-DOF Manipulation

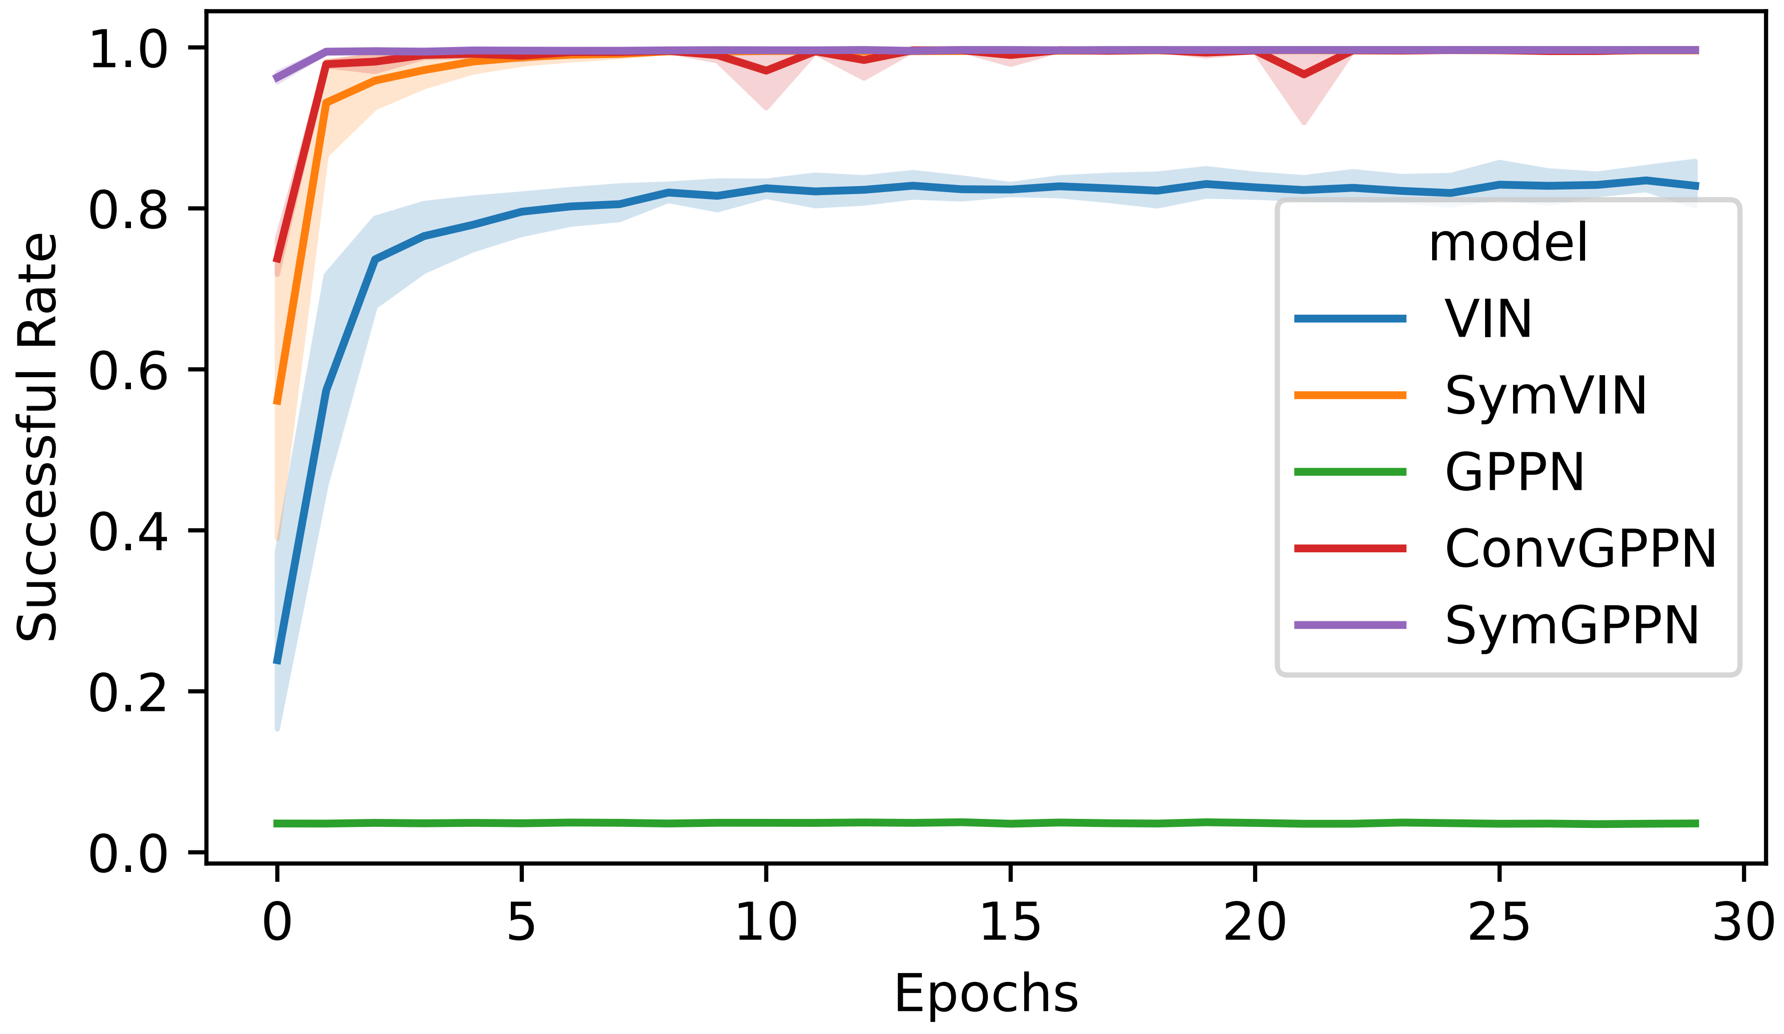

- Training curves on 15x15 maps; Use 3x3 filters, fixed 30 iterations (shared layers)
- More efficient training; Higher asymptotic performance; Better generalization

Results: Evaluation on test maps

Method (10K Data)	Navigation 15 × 15	28 × 28	50 × 50	Visual	Manipulation 18 × 18	36 × 36	Workspace
VIN	66.97	67.57	57.92	50.83	77.82	84.32	80.44
SymVIN	98.99	98.14	86.20	95.50	99.98	99.36	**91.10**
GPPN	96.36	95.77	91.84	93.13	2.62	1.68	3.67
ConvGPPN	99.75	99.09	97.21	98.55	99.98	99.95	89.88
SymGPPN	**99.98**	**99.86**	**99.49**	**99.78**	**100.00**	**99.99**	90.50

- Better generalization on novel maps
- Test novel maps are not necessarily rotated version of training maps

Visualization: VIN

Feed in M and $\circlearrowleft 90° \circ M$

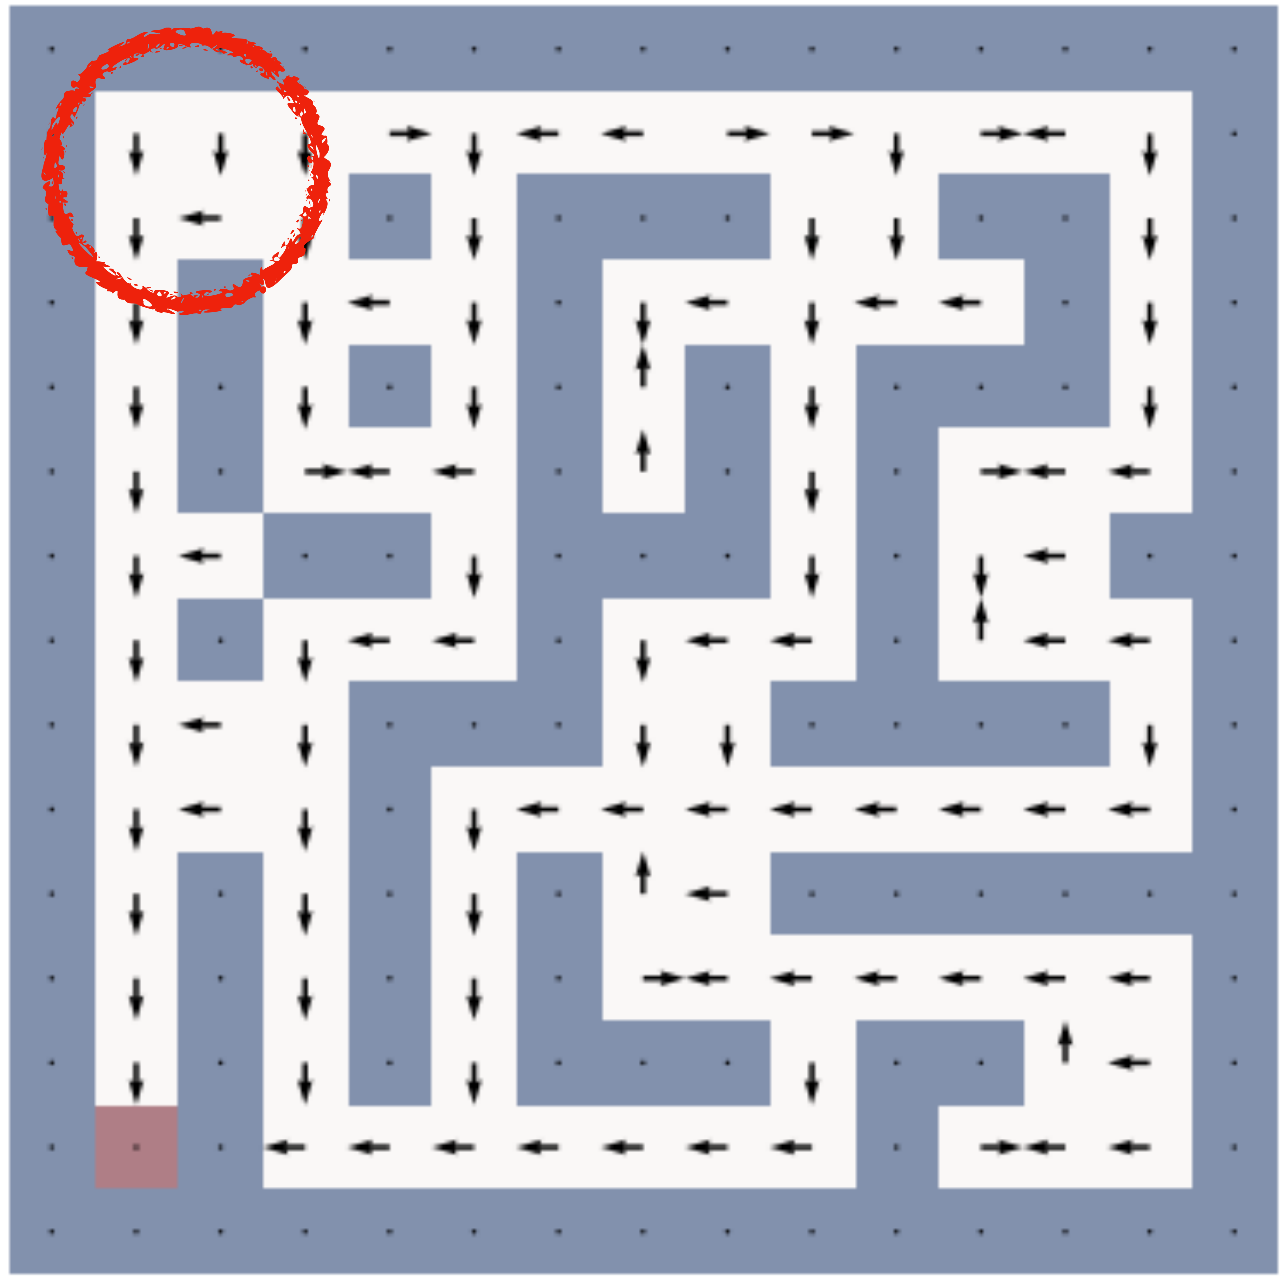

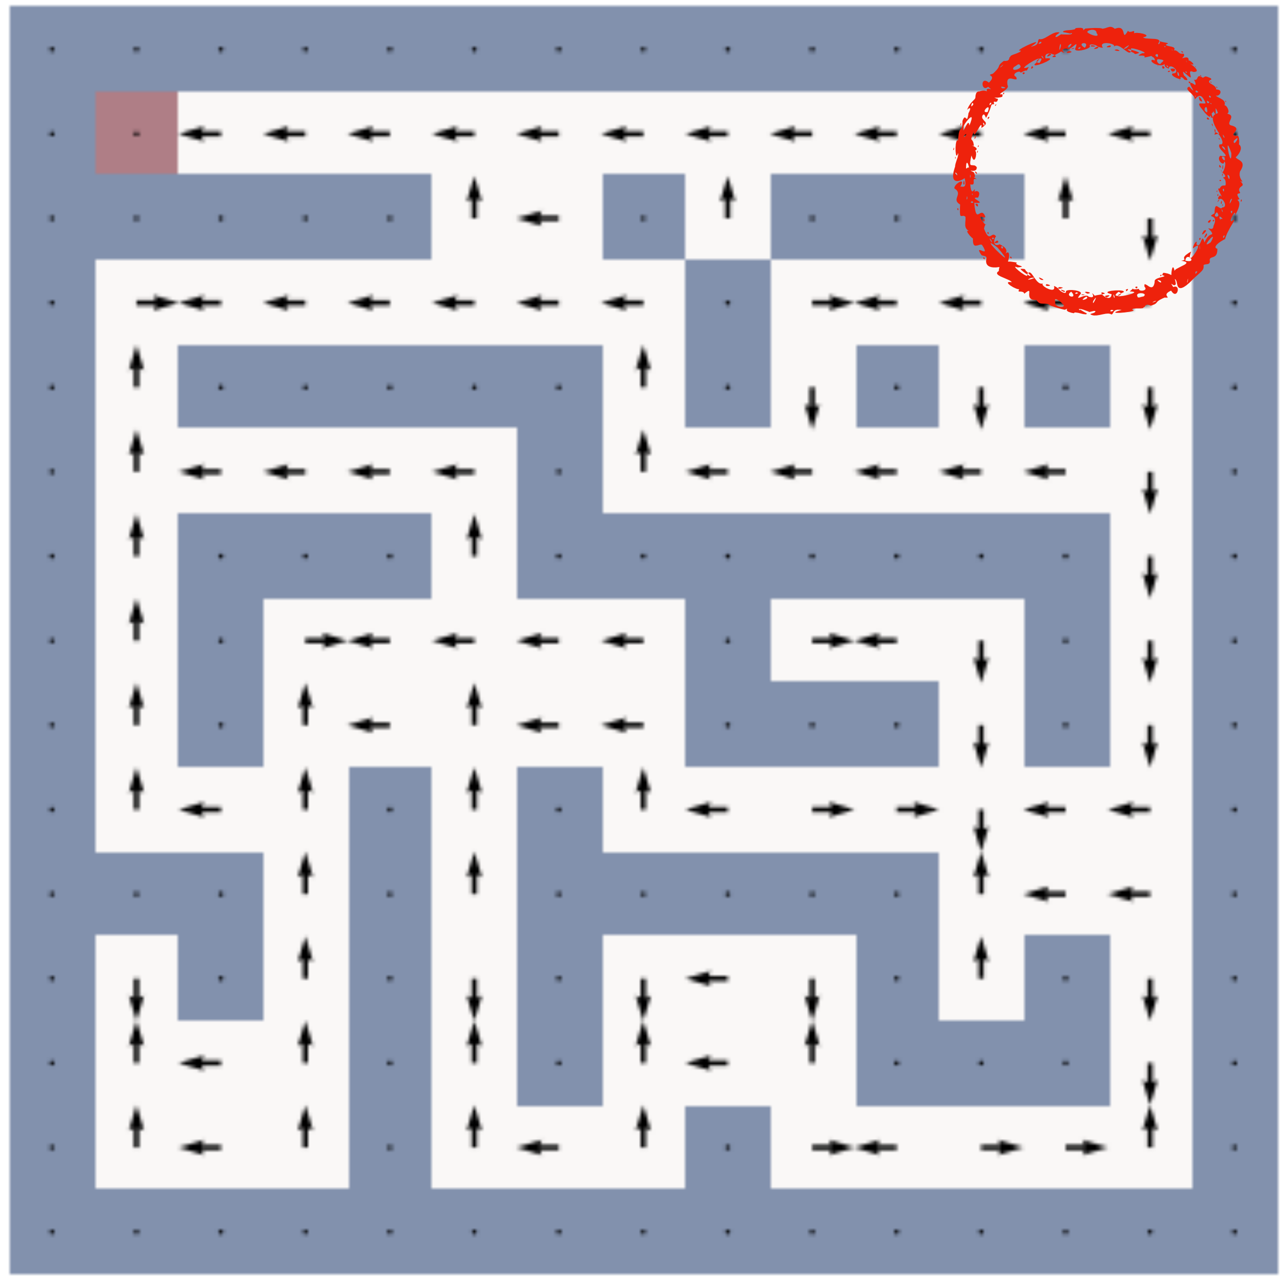

VIN output doesn't satisfy equivariance

Visualization: SymVIN

Feed in M and $\circlearrowleft 90^\circ \circ M$

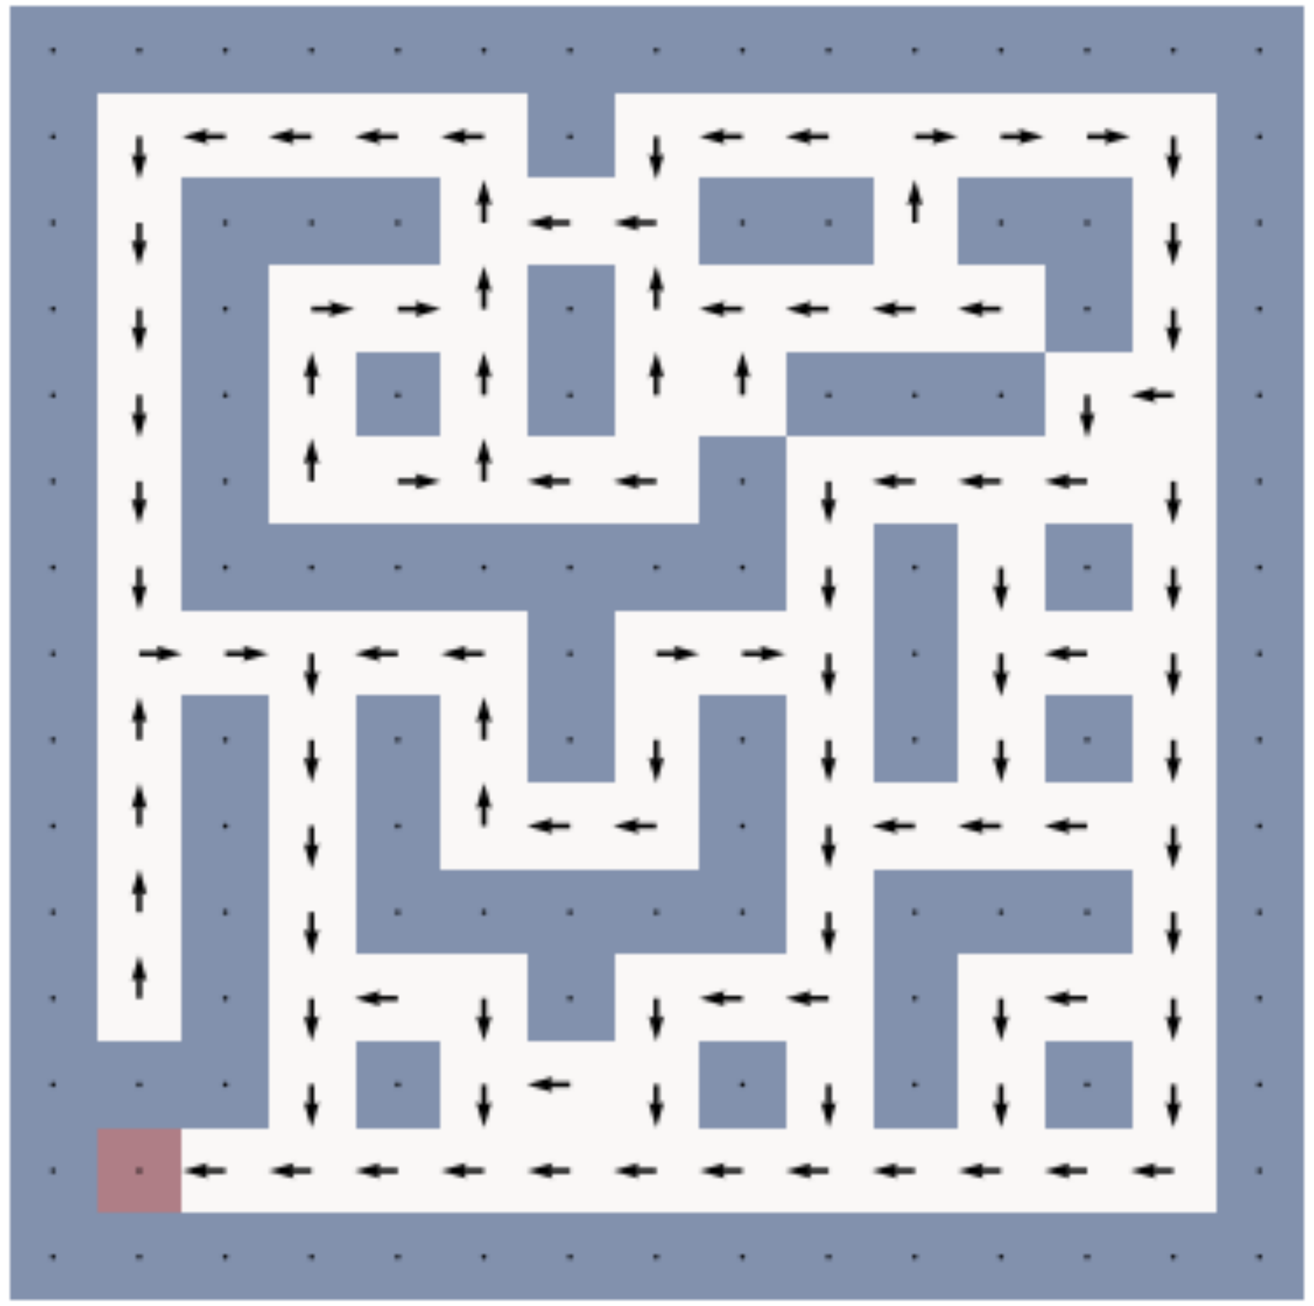

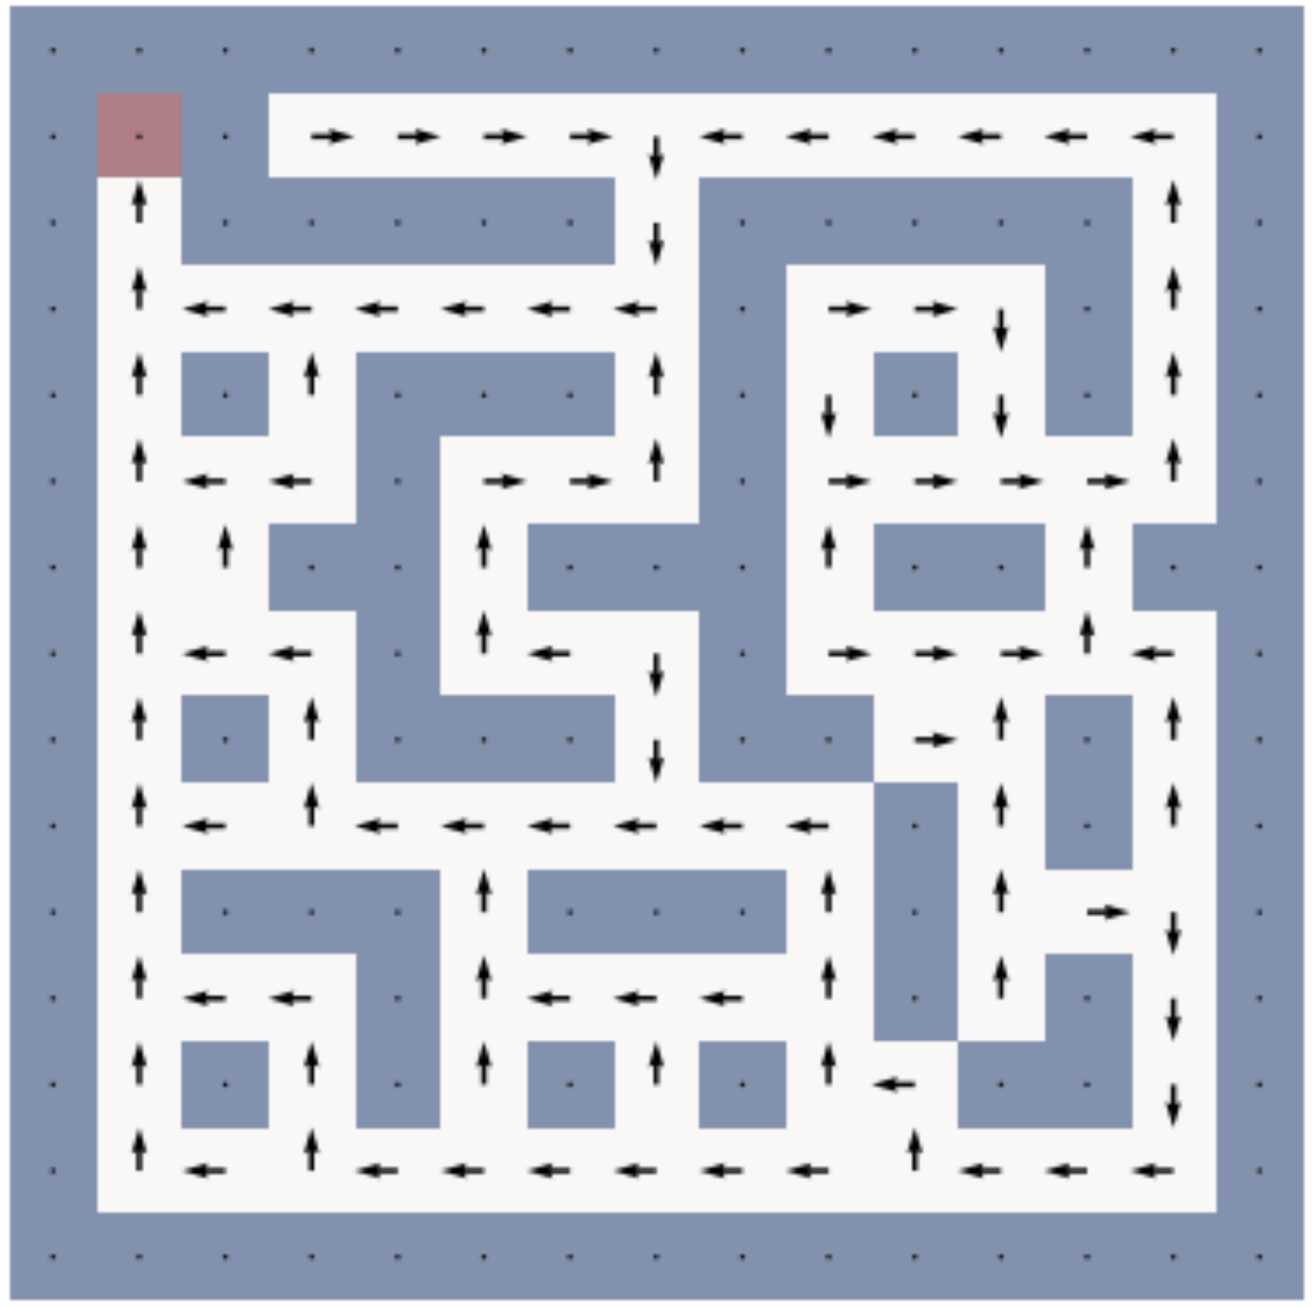

SymVIN guarantees output is equivariant

Summary

- Introduce a framework for incorporating symmetry into path-planning problems
- Prove that value iteration for path planning can be treated as a steerable CNN
- Show that Symmetric Planners improve in training efficiency and generalization

Extension: Navigation on Graph

E(2)-Equivariant Graph Planning for Navigation

Submission to RA-L (2023)

Linfeng Zhao*, Hongyu Li*, Taskin Padir, Huaizu Jiang^, Lawson L.S. Wong^




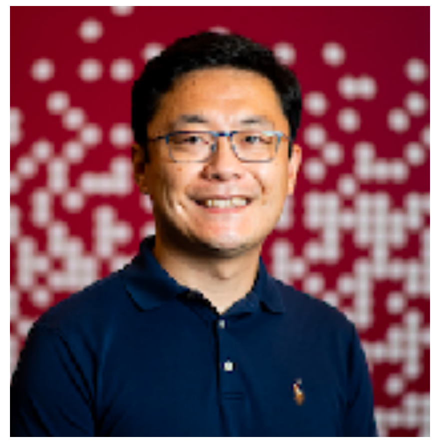


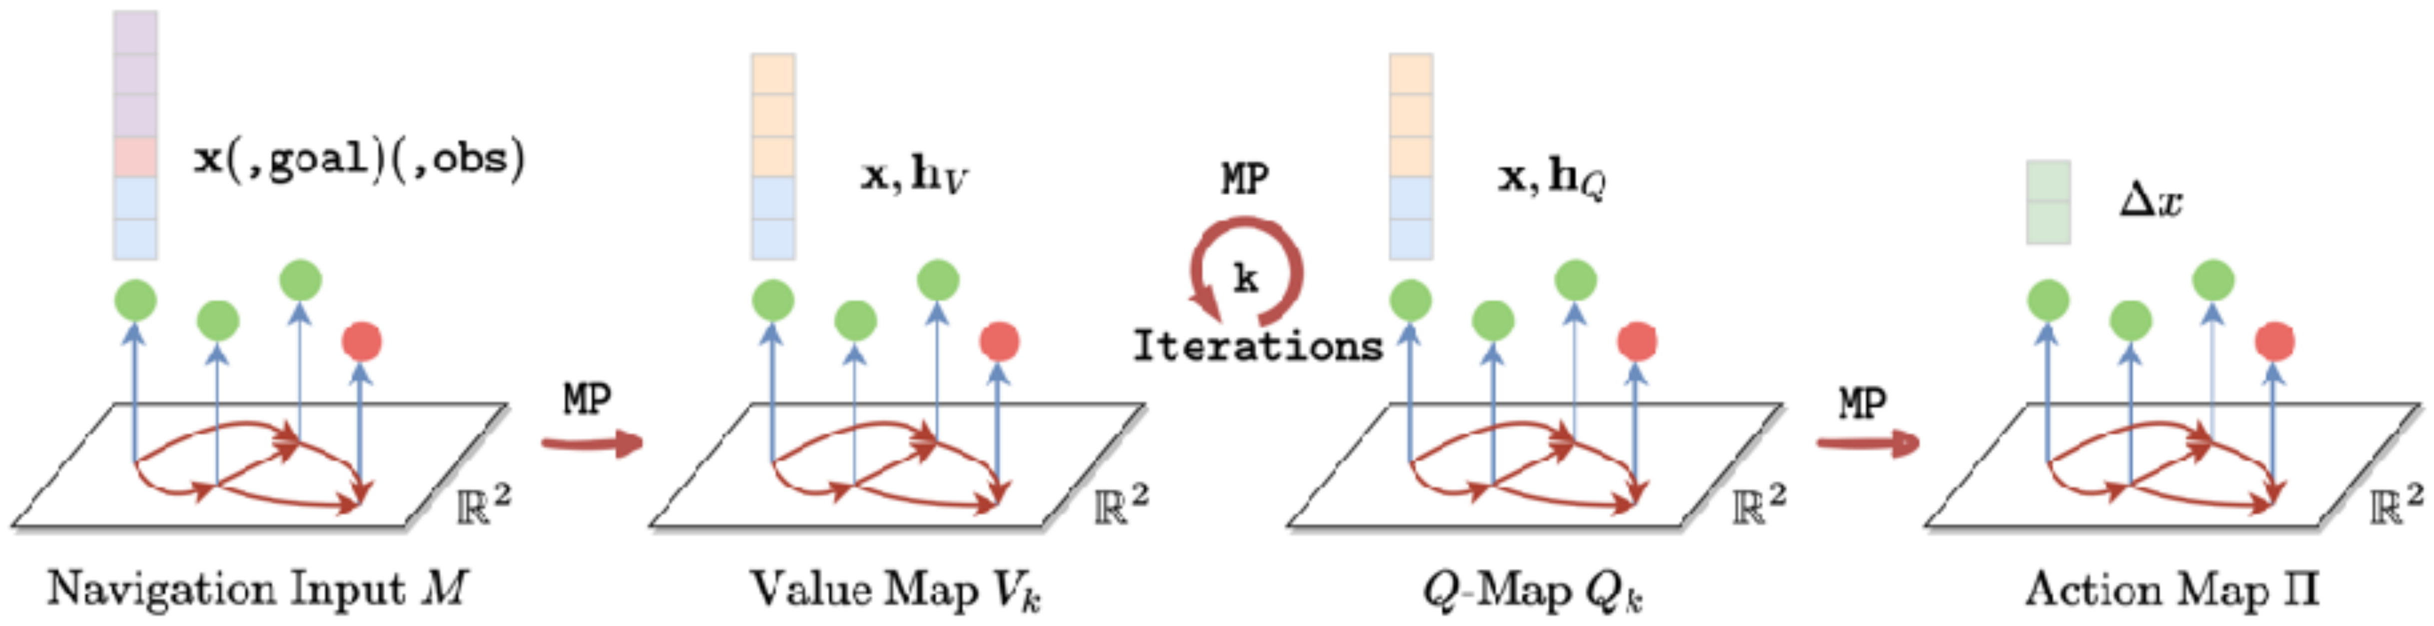

Extension: Navigation on Graph

E(2)-Equivariant Graph Planning for Navigation

Submission to RA-L (2023)

Linfeng Zhao*, Hongyu Li*, Taskin Padir, Huaizu Jiang^, Lawson L.S. Wong^

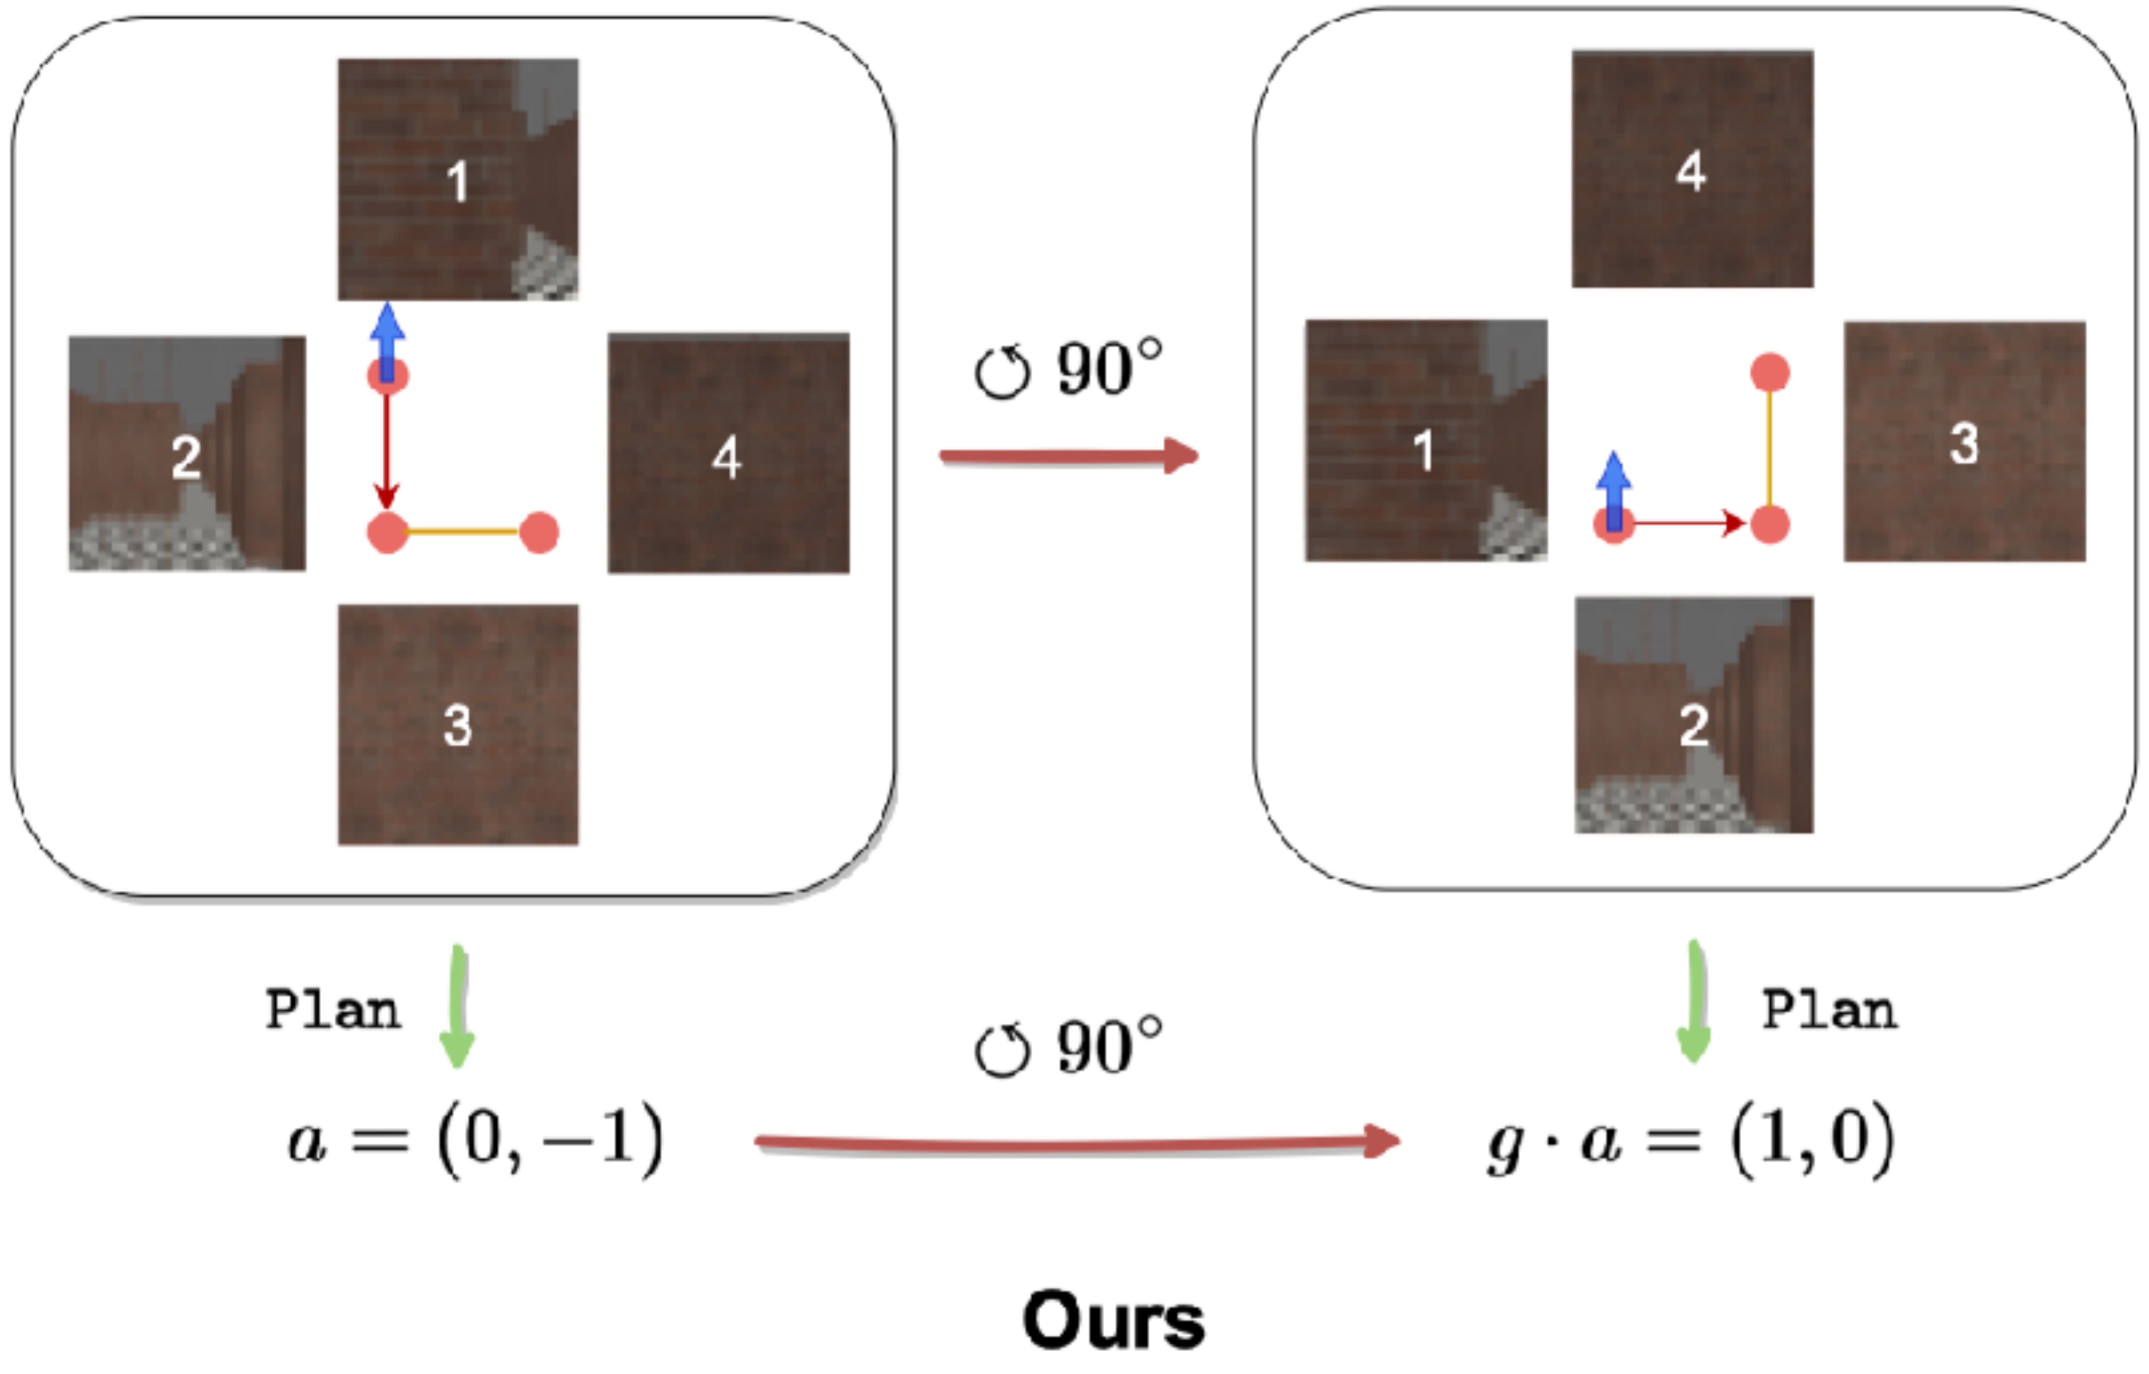

Setup: Path Planning on *Graph*

- Path planning on graph: finding shortest path on graphs given goal

Challenge 1: Grid to Graph

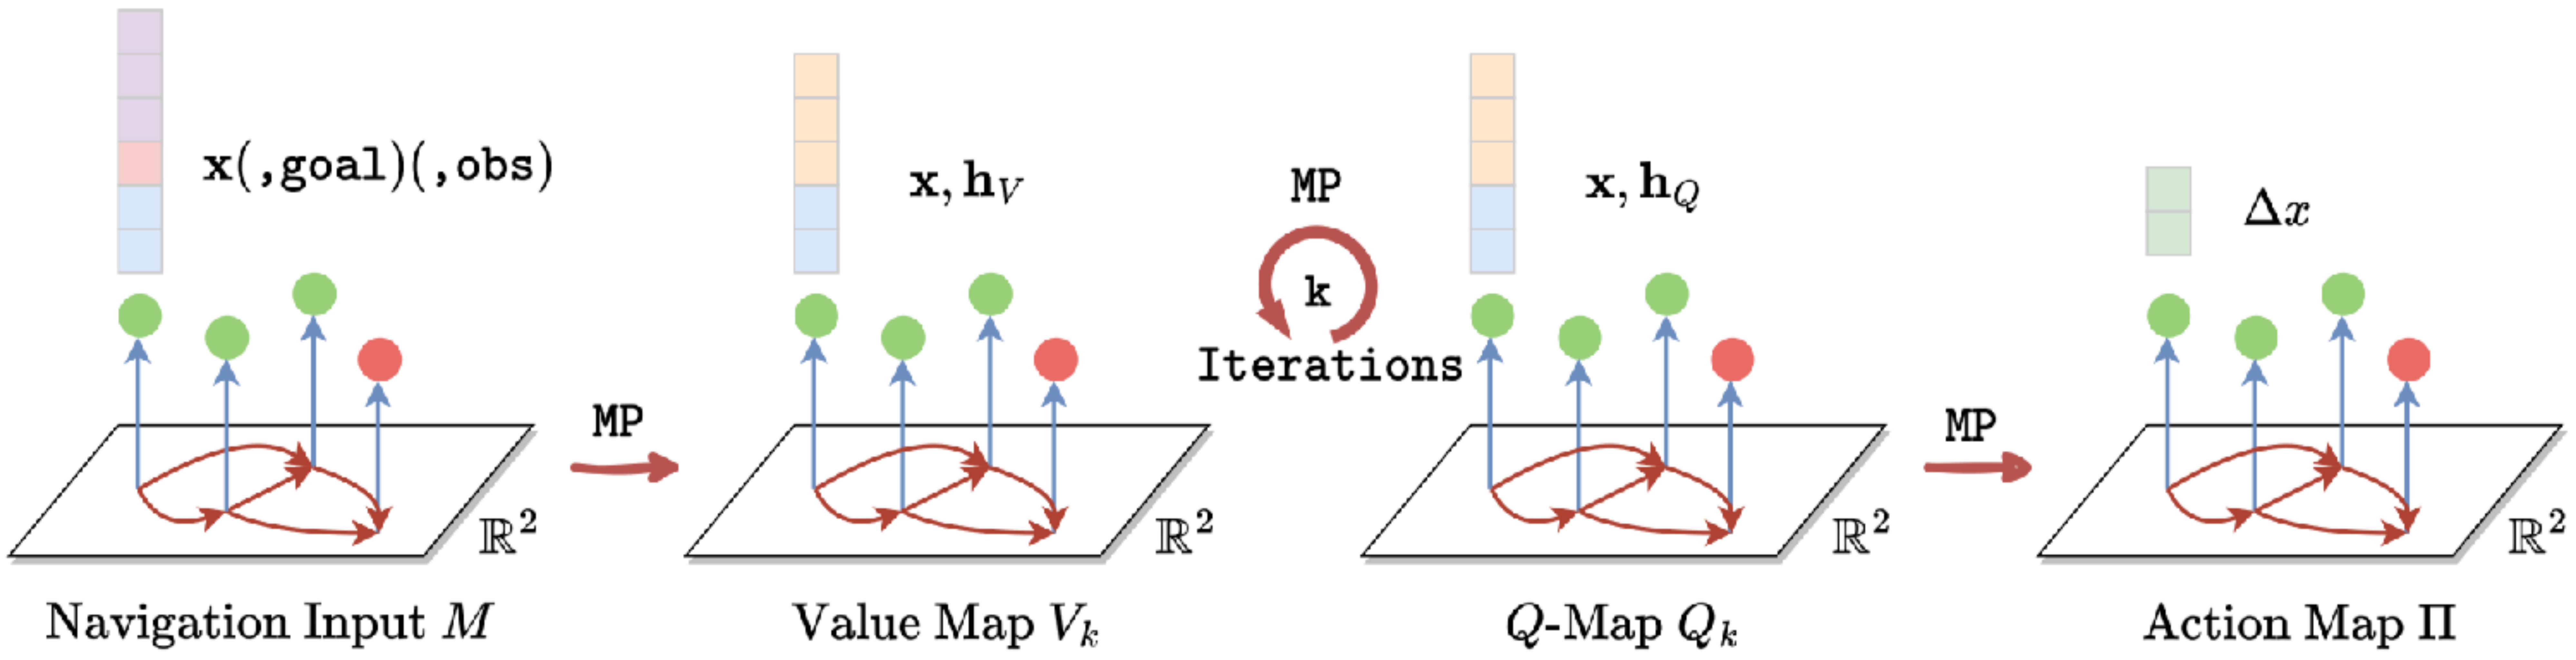

- SymVIN only allows planning on 2D grid and 2D discrete symmetry
- We extend to graphs: performing value iteration via message passing GNN and supporting E(2) continuous symmetry

Challenge 2: Camera/View Layout

Robots may only have K views

- Naive equivariance only allow C_K $(360°/K)$ rotation symmetry

- We lift it to $\mathrm{SO}(2)$ to potentially allow *continuous* symmetry in downstream planning network

Commutative diagram of the ***lift*** layer:

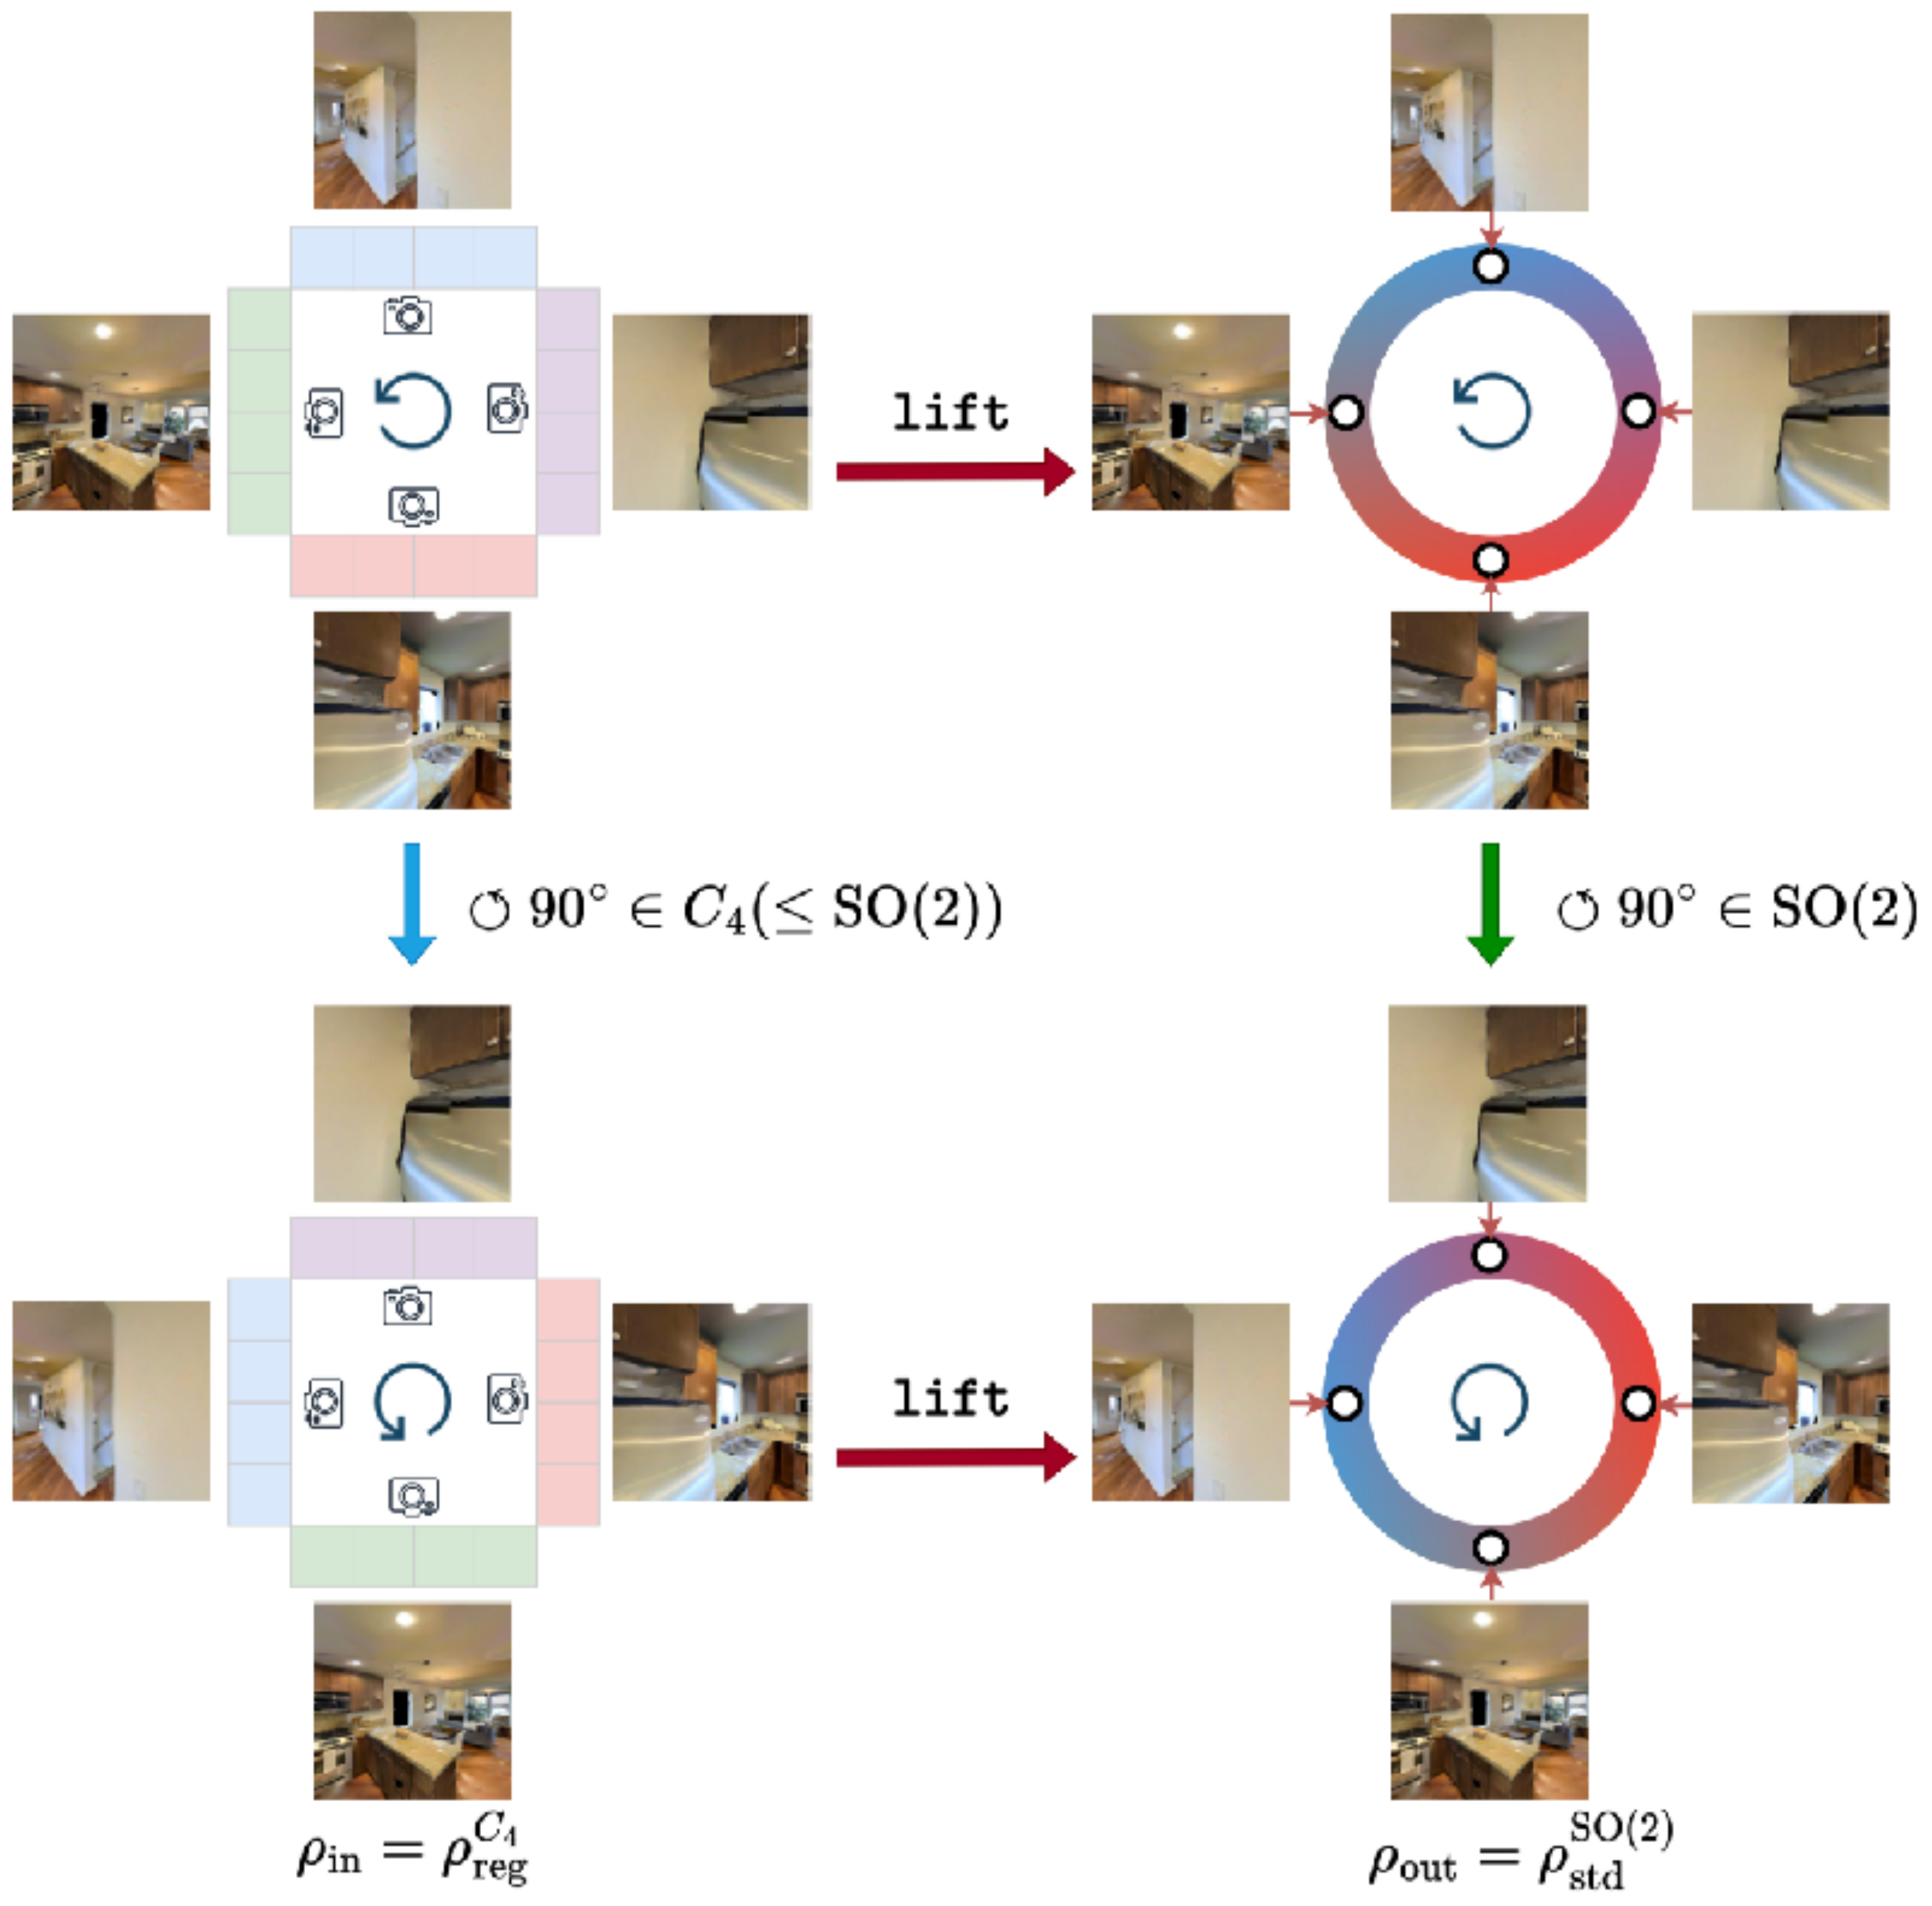

Pipeline: E(2) Message Passing VIN

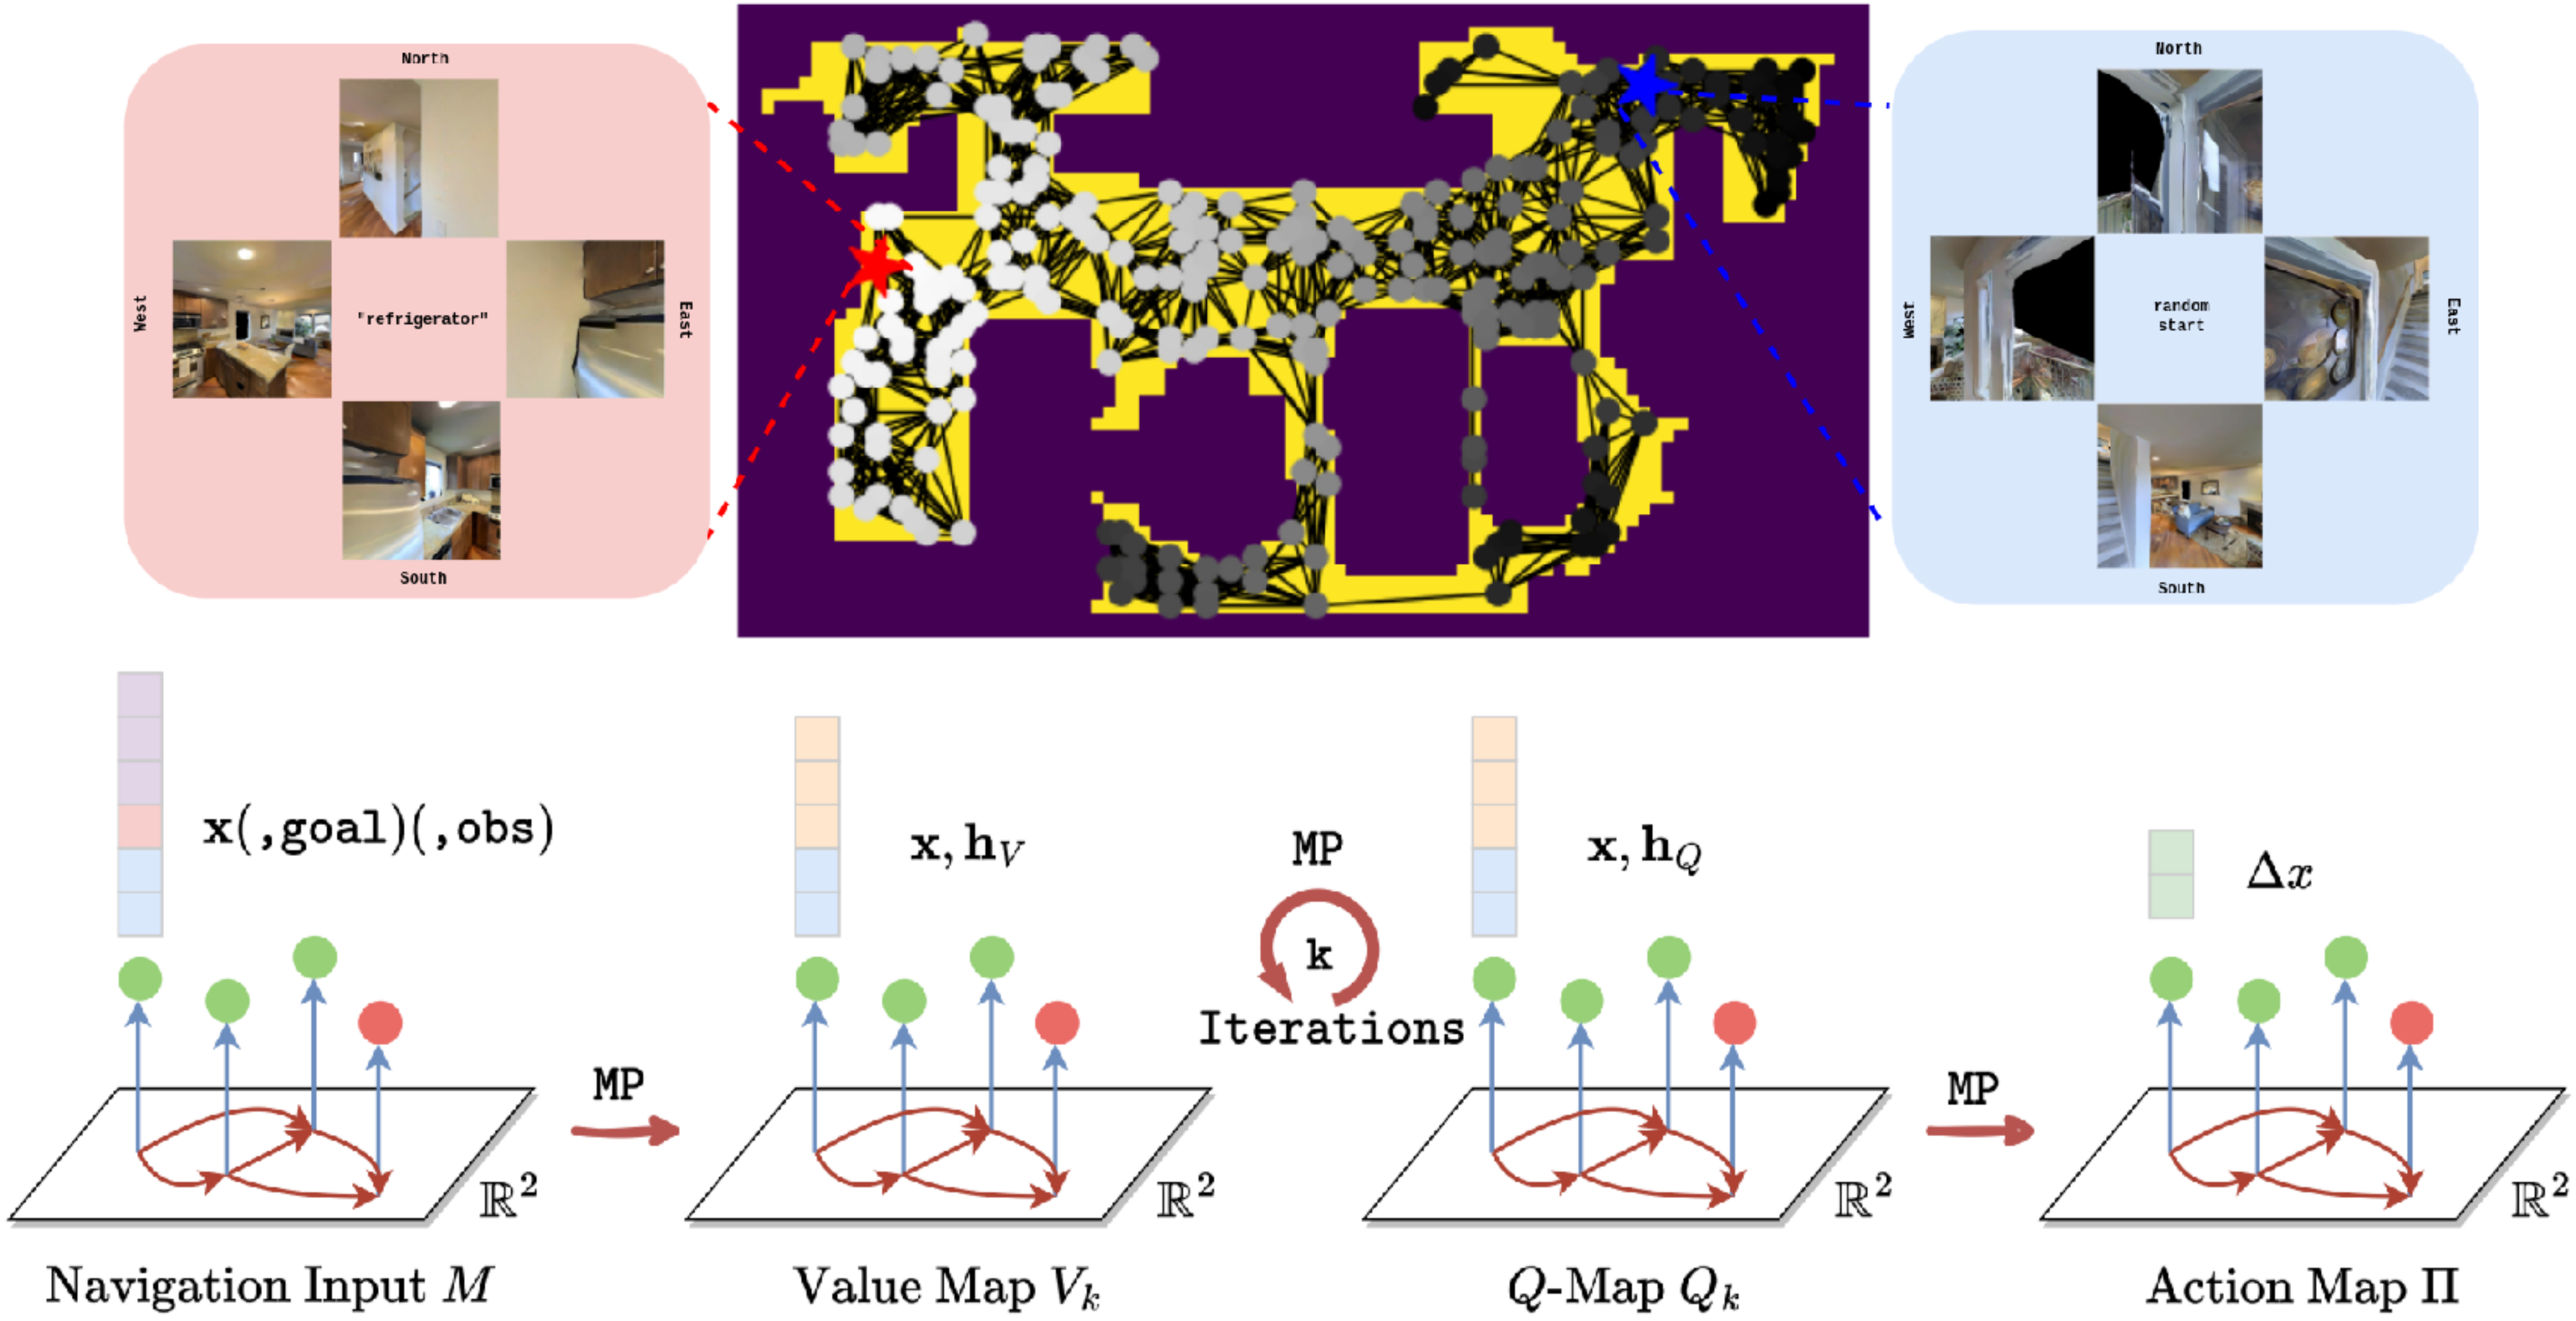

Compositional Structures

Toward Compositional Generalization in Object-Oriented World Modeling

ICML 2022 Long Oral Presentation

Linfeng Zhao, *Lingzhi Kong, Robin Walters, Lawson Wong*

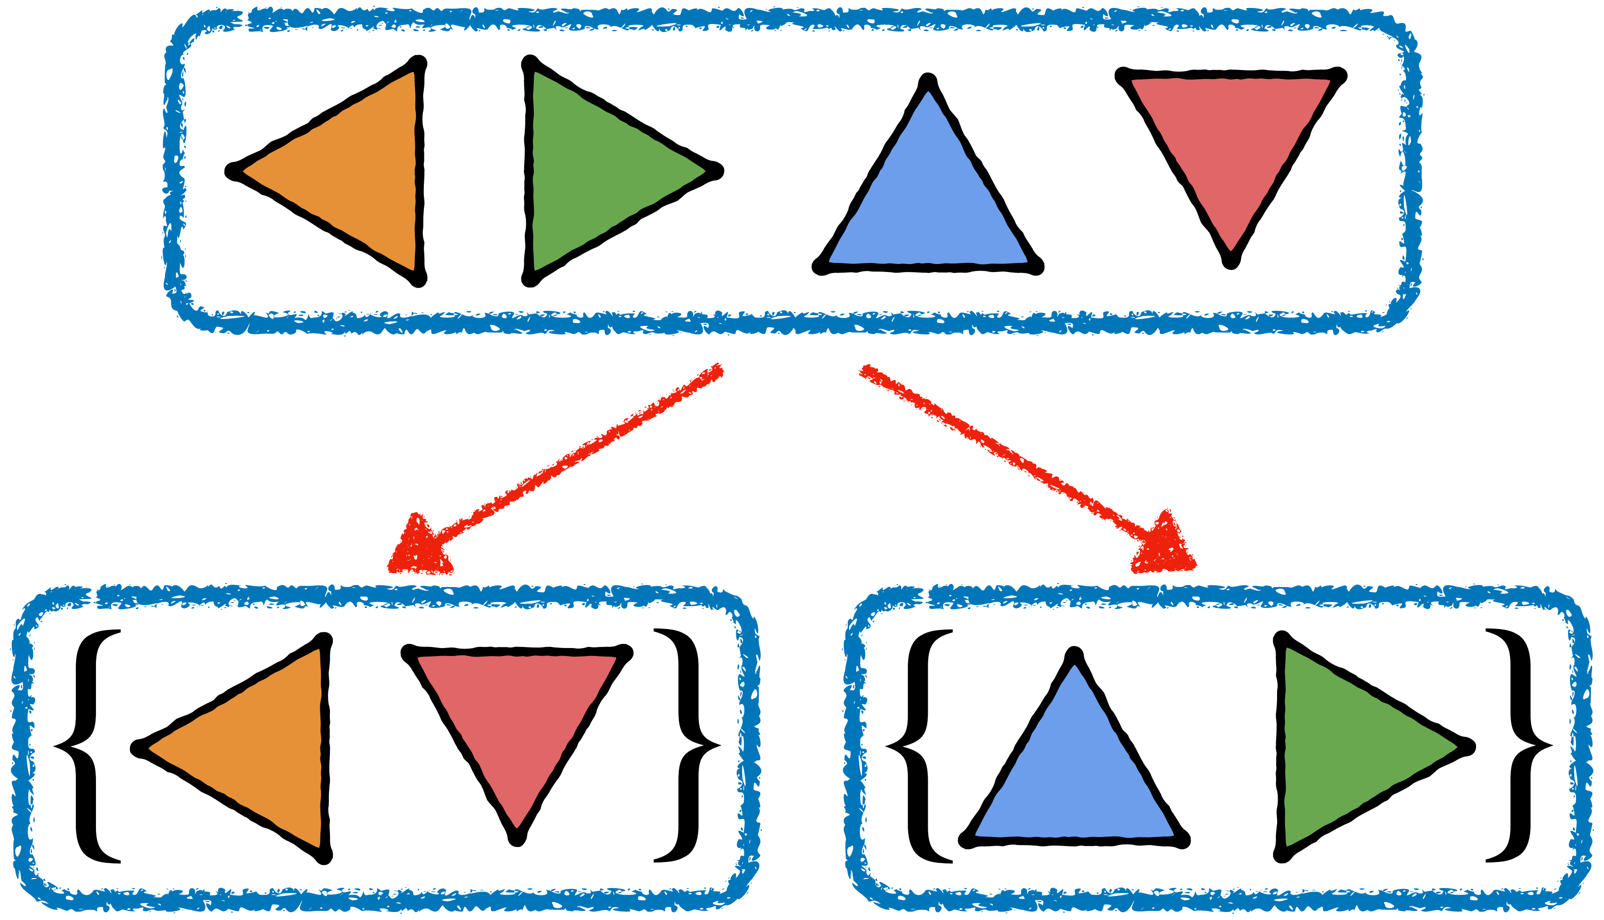

Object Interchangeability

Permutation Symmetry S_N

Motivation

Proposed Setup: Object Library

Motivation: sampling words from vocabulary to form sentences

Object Library $\mathbb{L}$
“Vocabulary”
All possible objects

Scenes $\mathbb{O}_i \subset \mathbb{L}$
(Ordered) “Sentences”
A combination of objects

Scene MDPs $\mathscr{M}_{\mathbb{O}_i}$
Generated by $\mathbb{O}$
Moving objects on a table

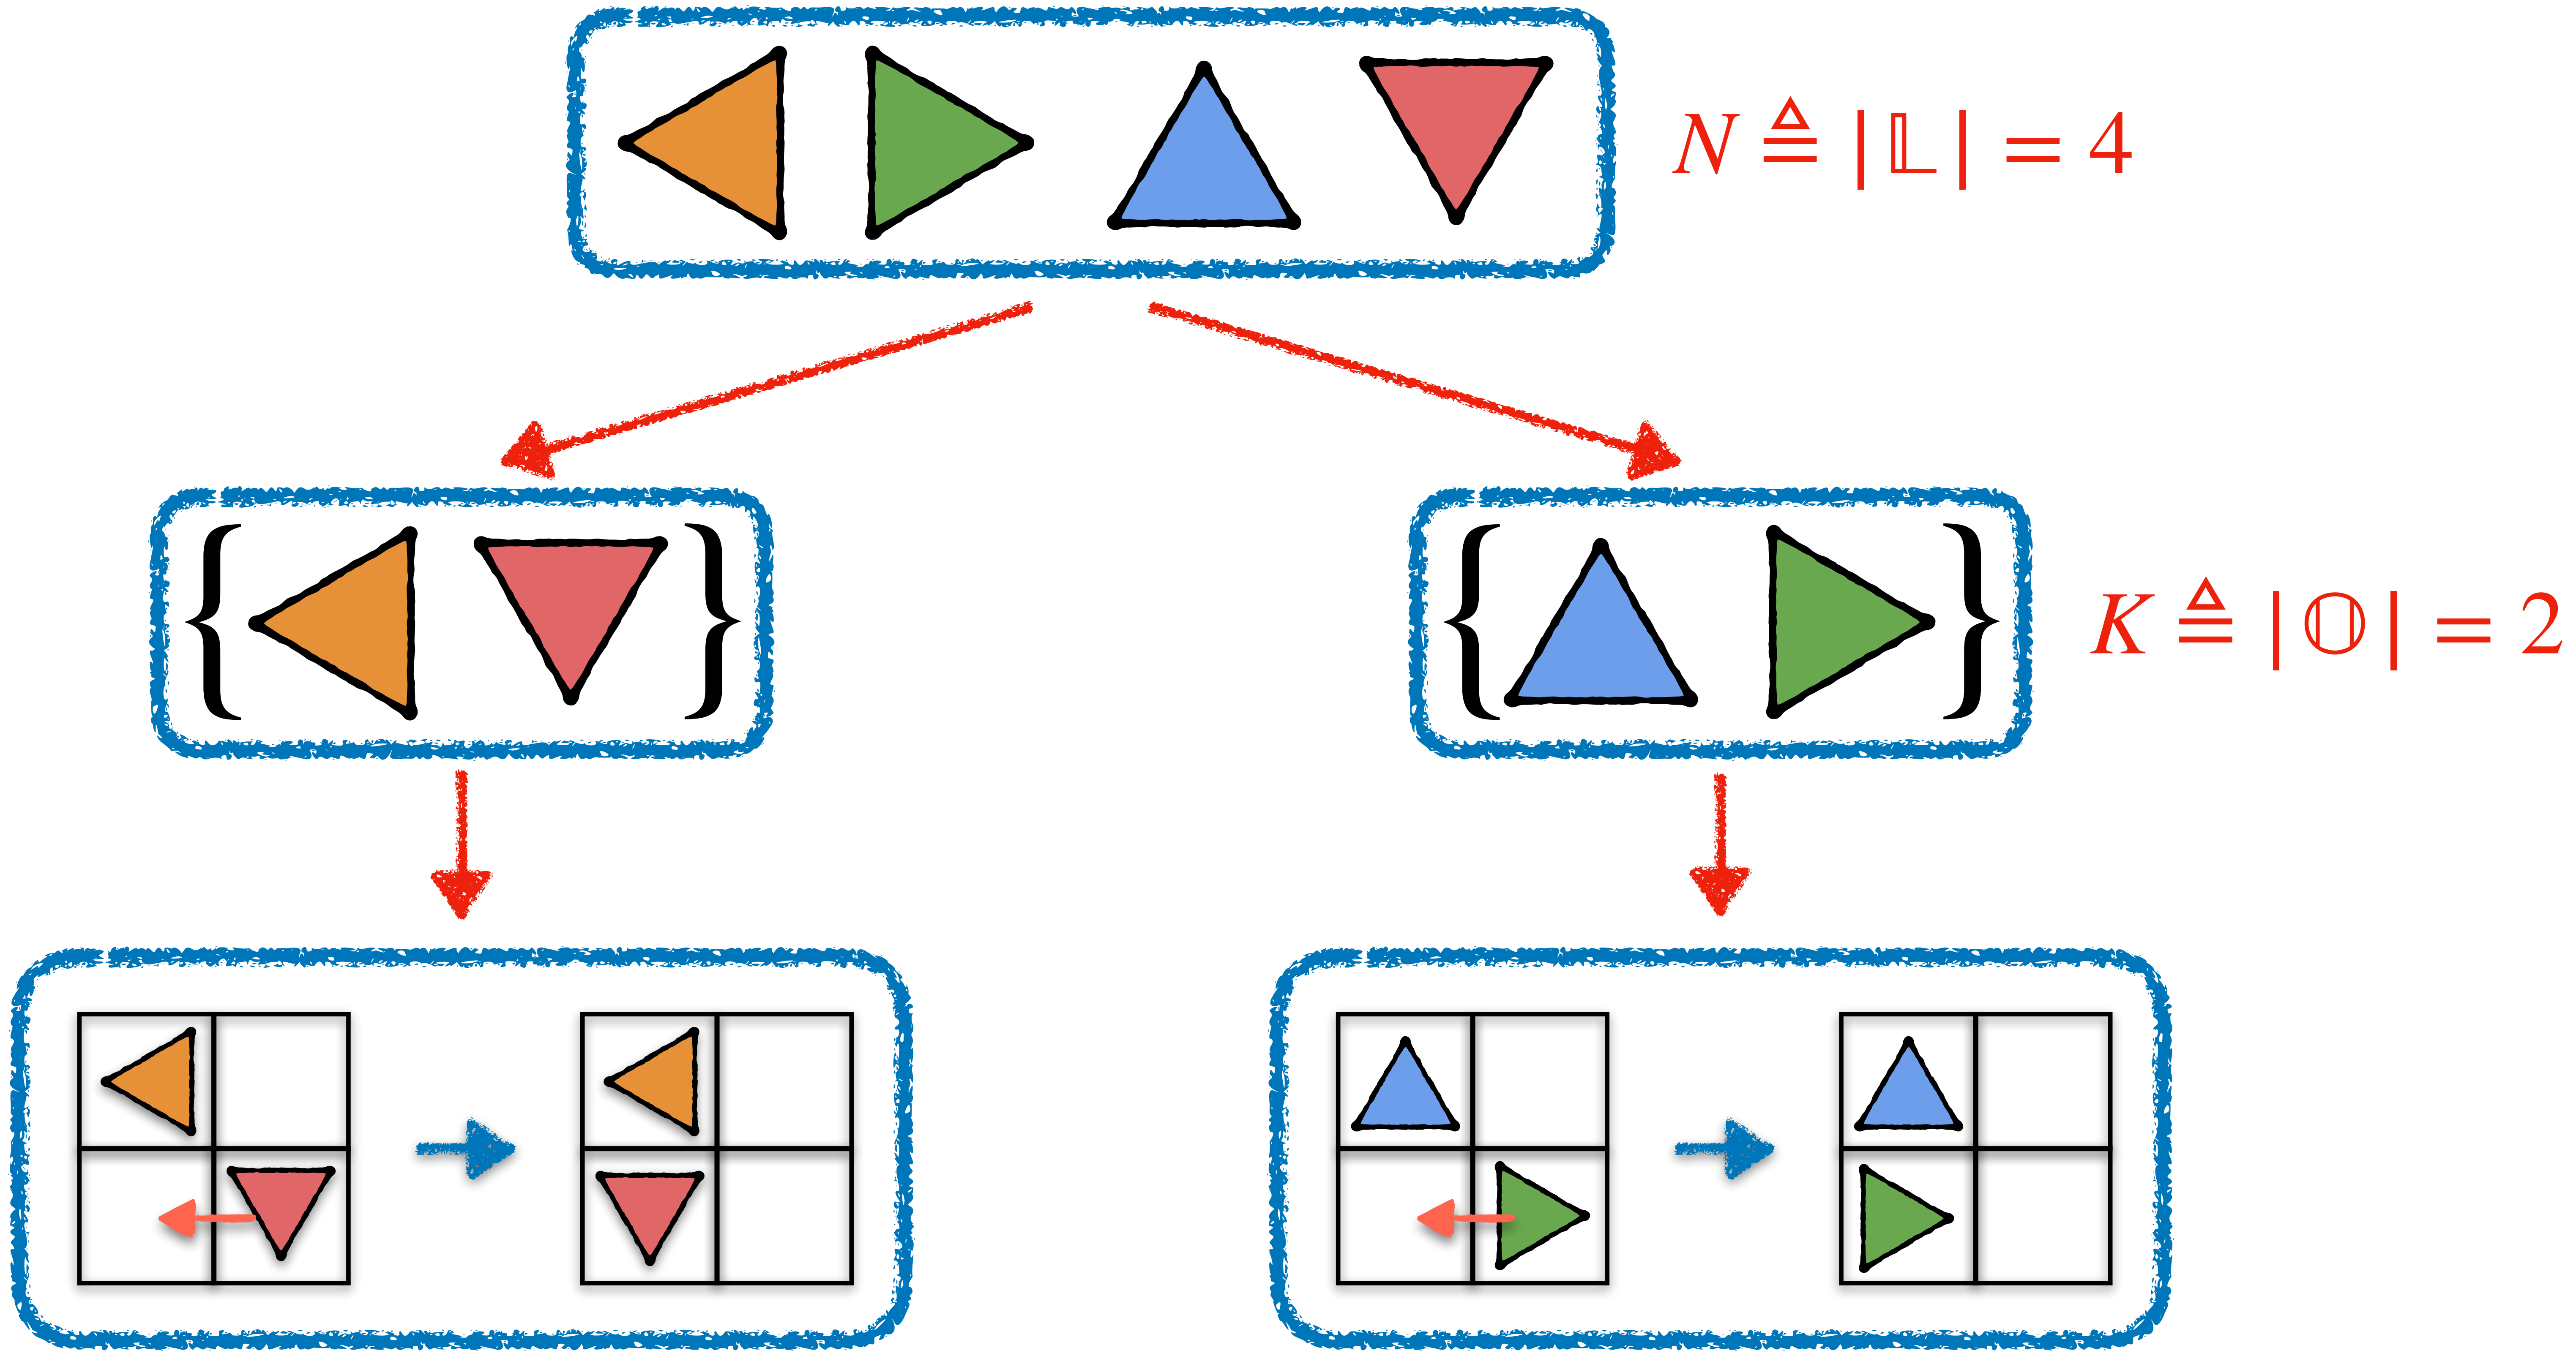

Proposed Setup: Object Library

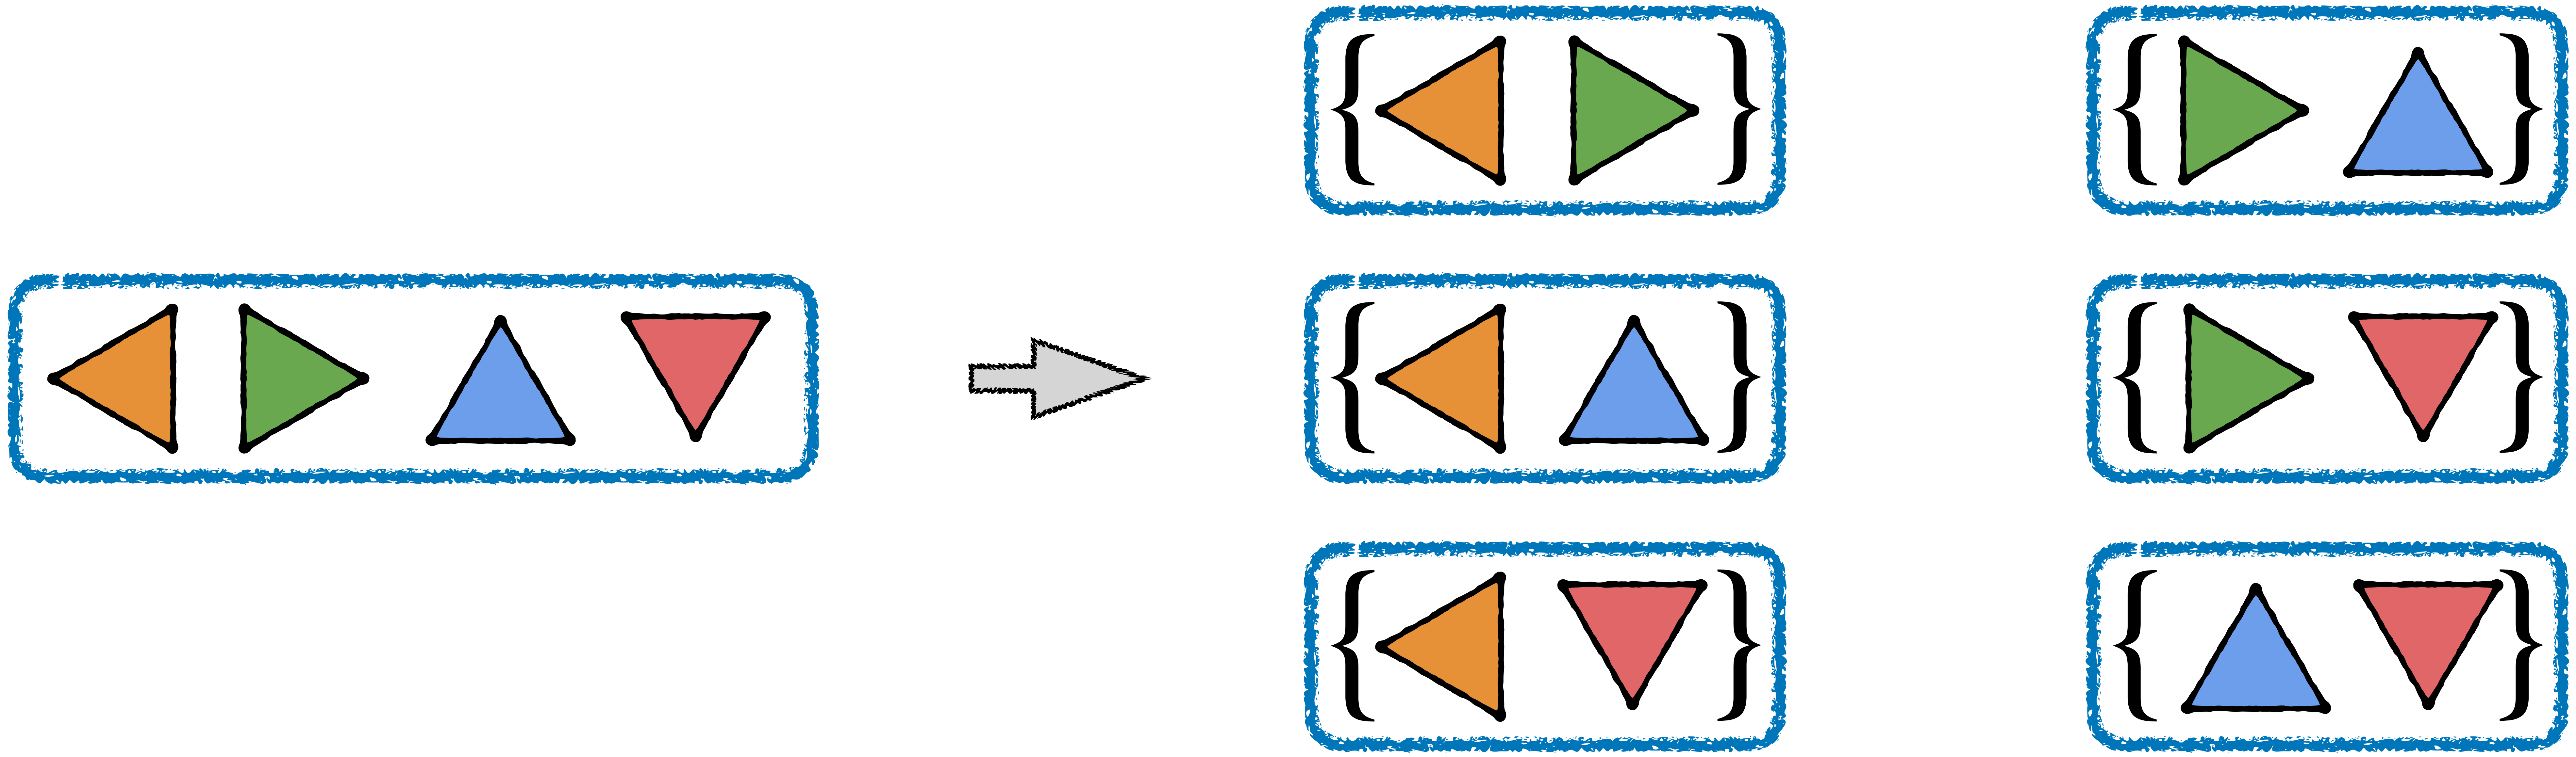

All $\binom{N}{K} = \binom{4}{2} = 6$ combinations

Binding Visualization

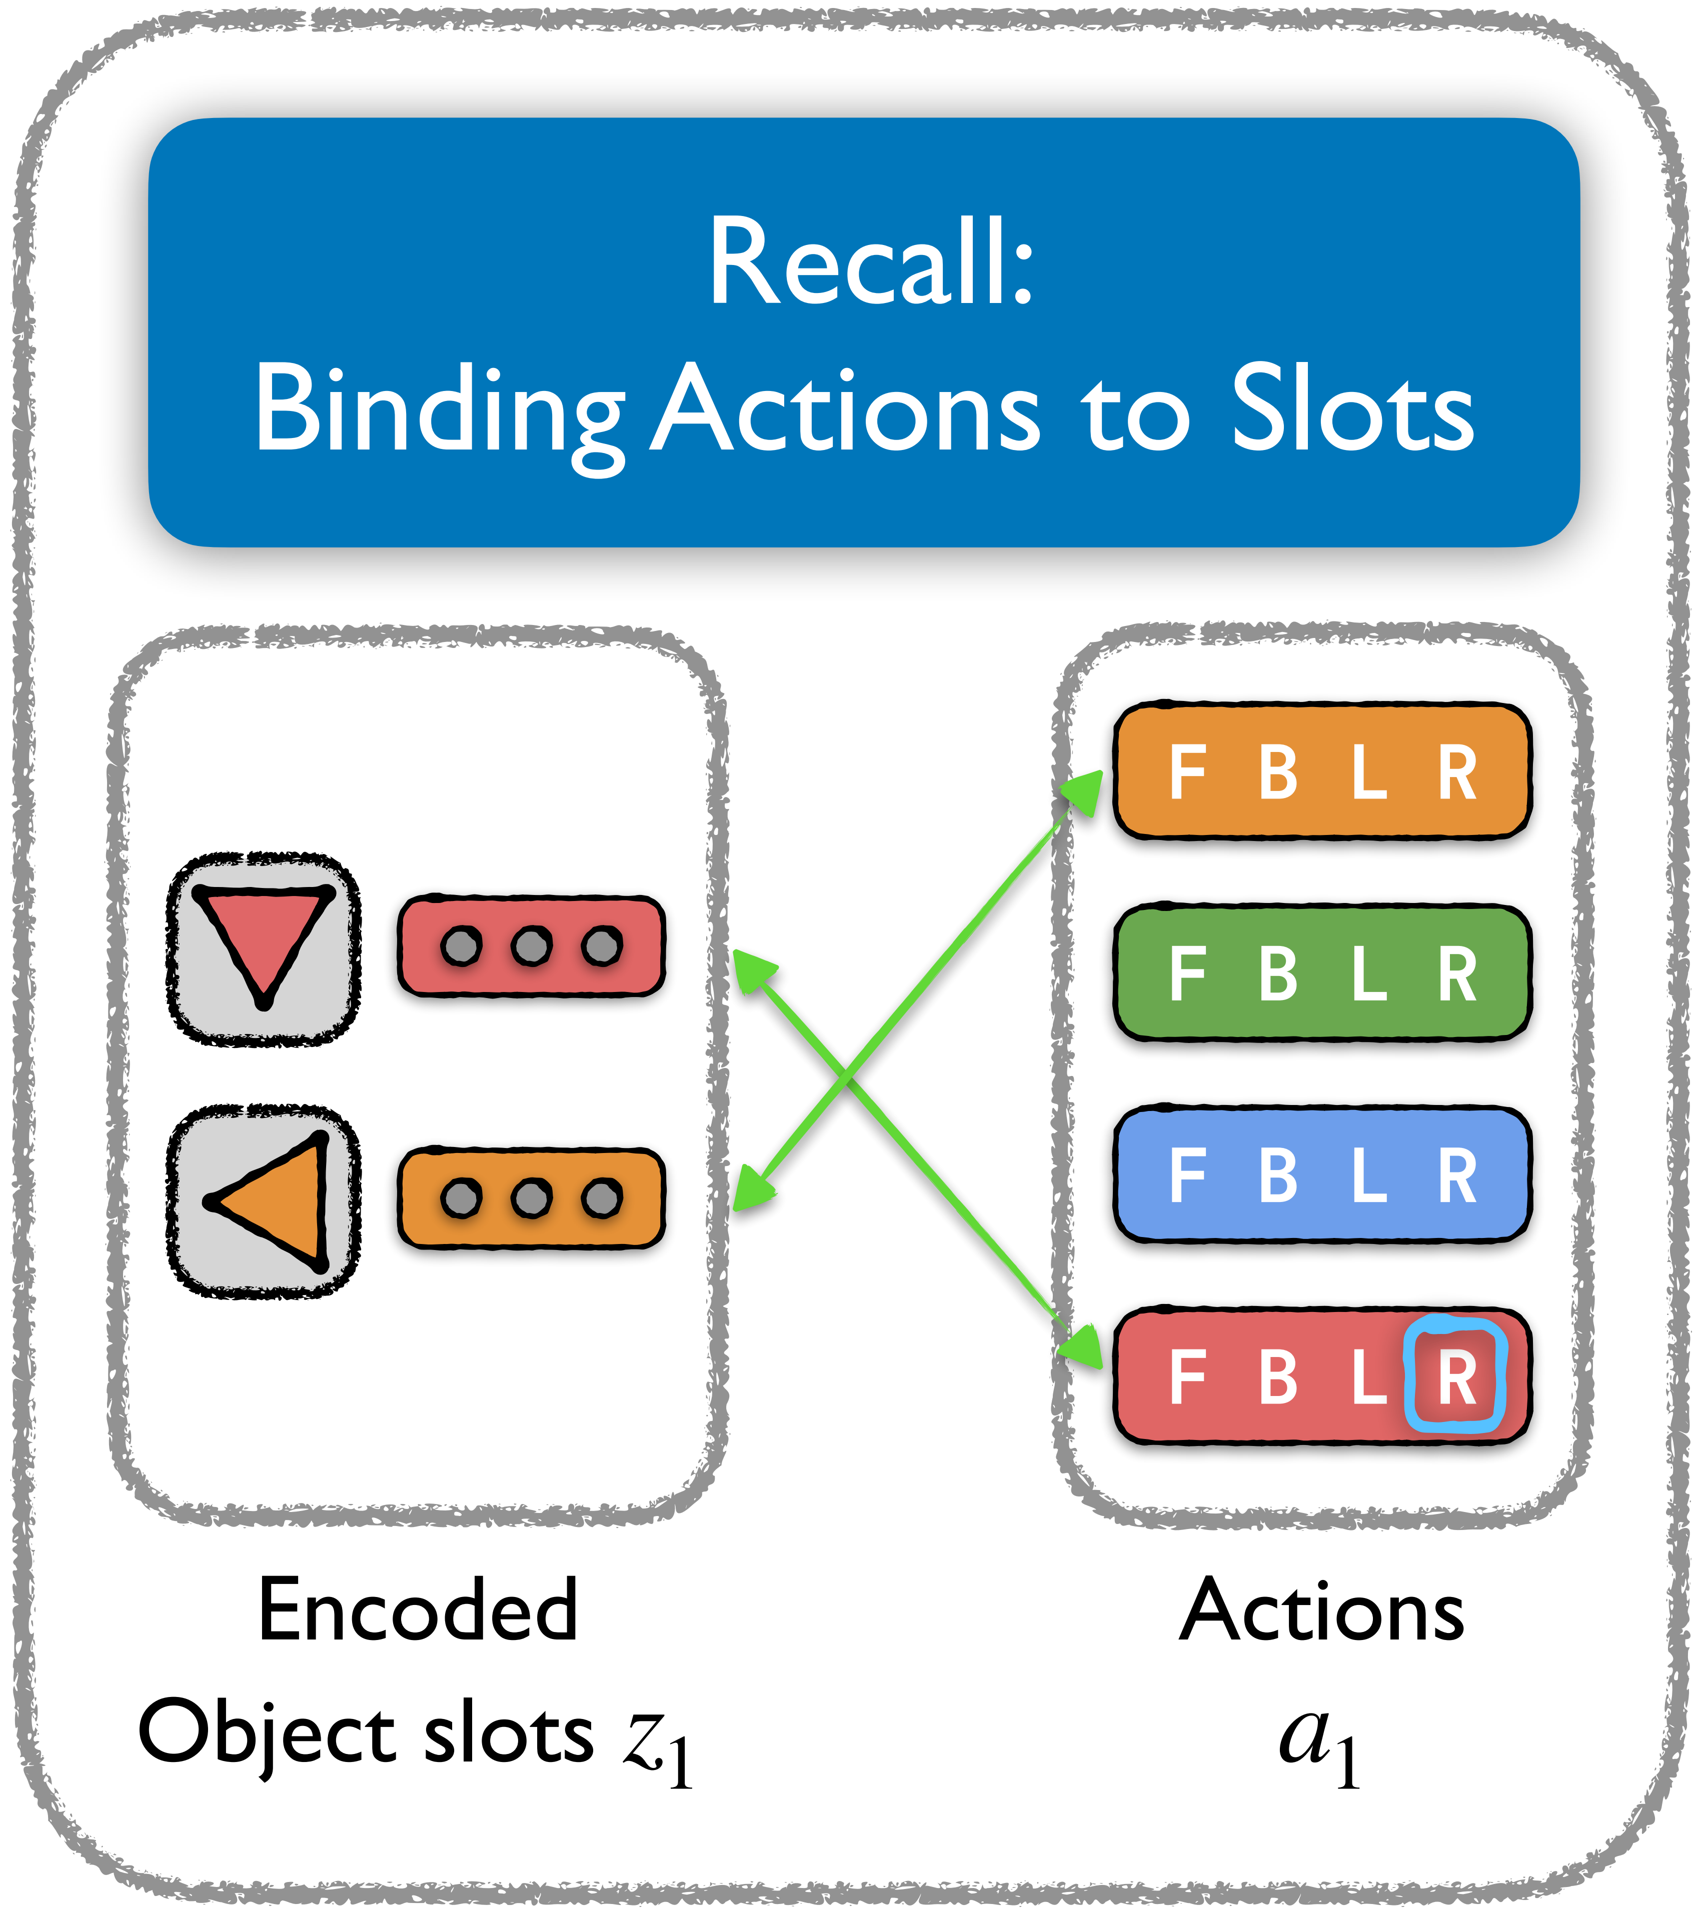

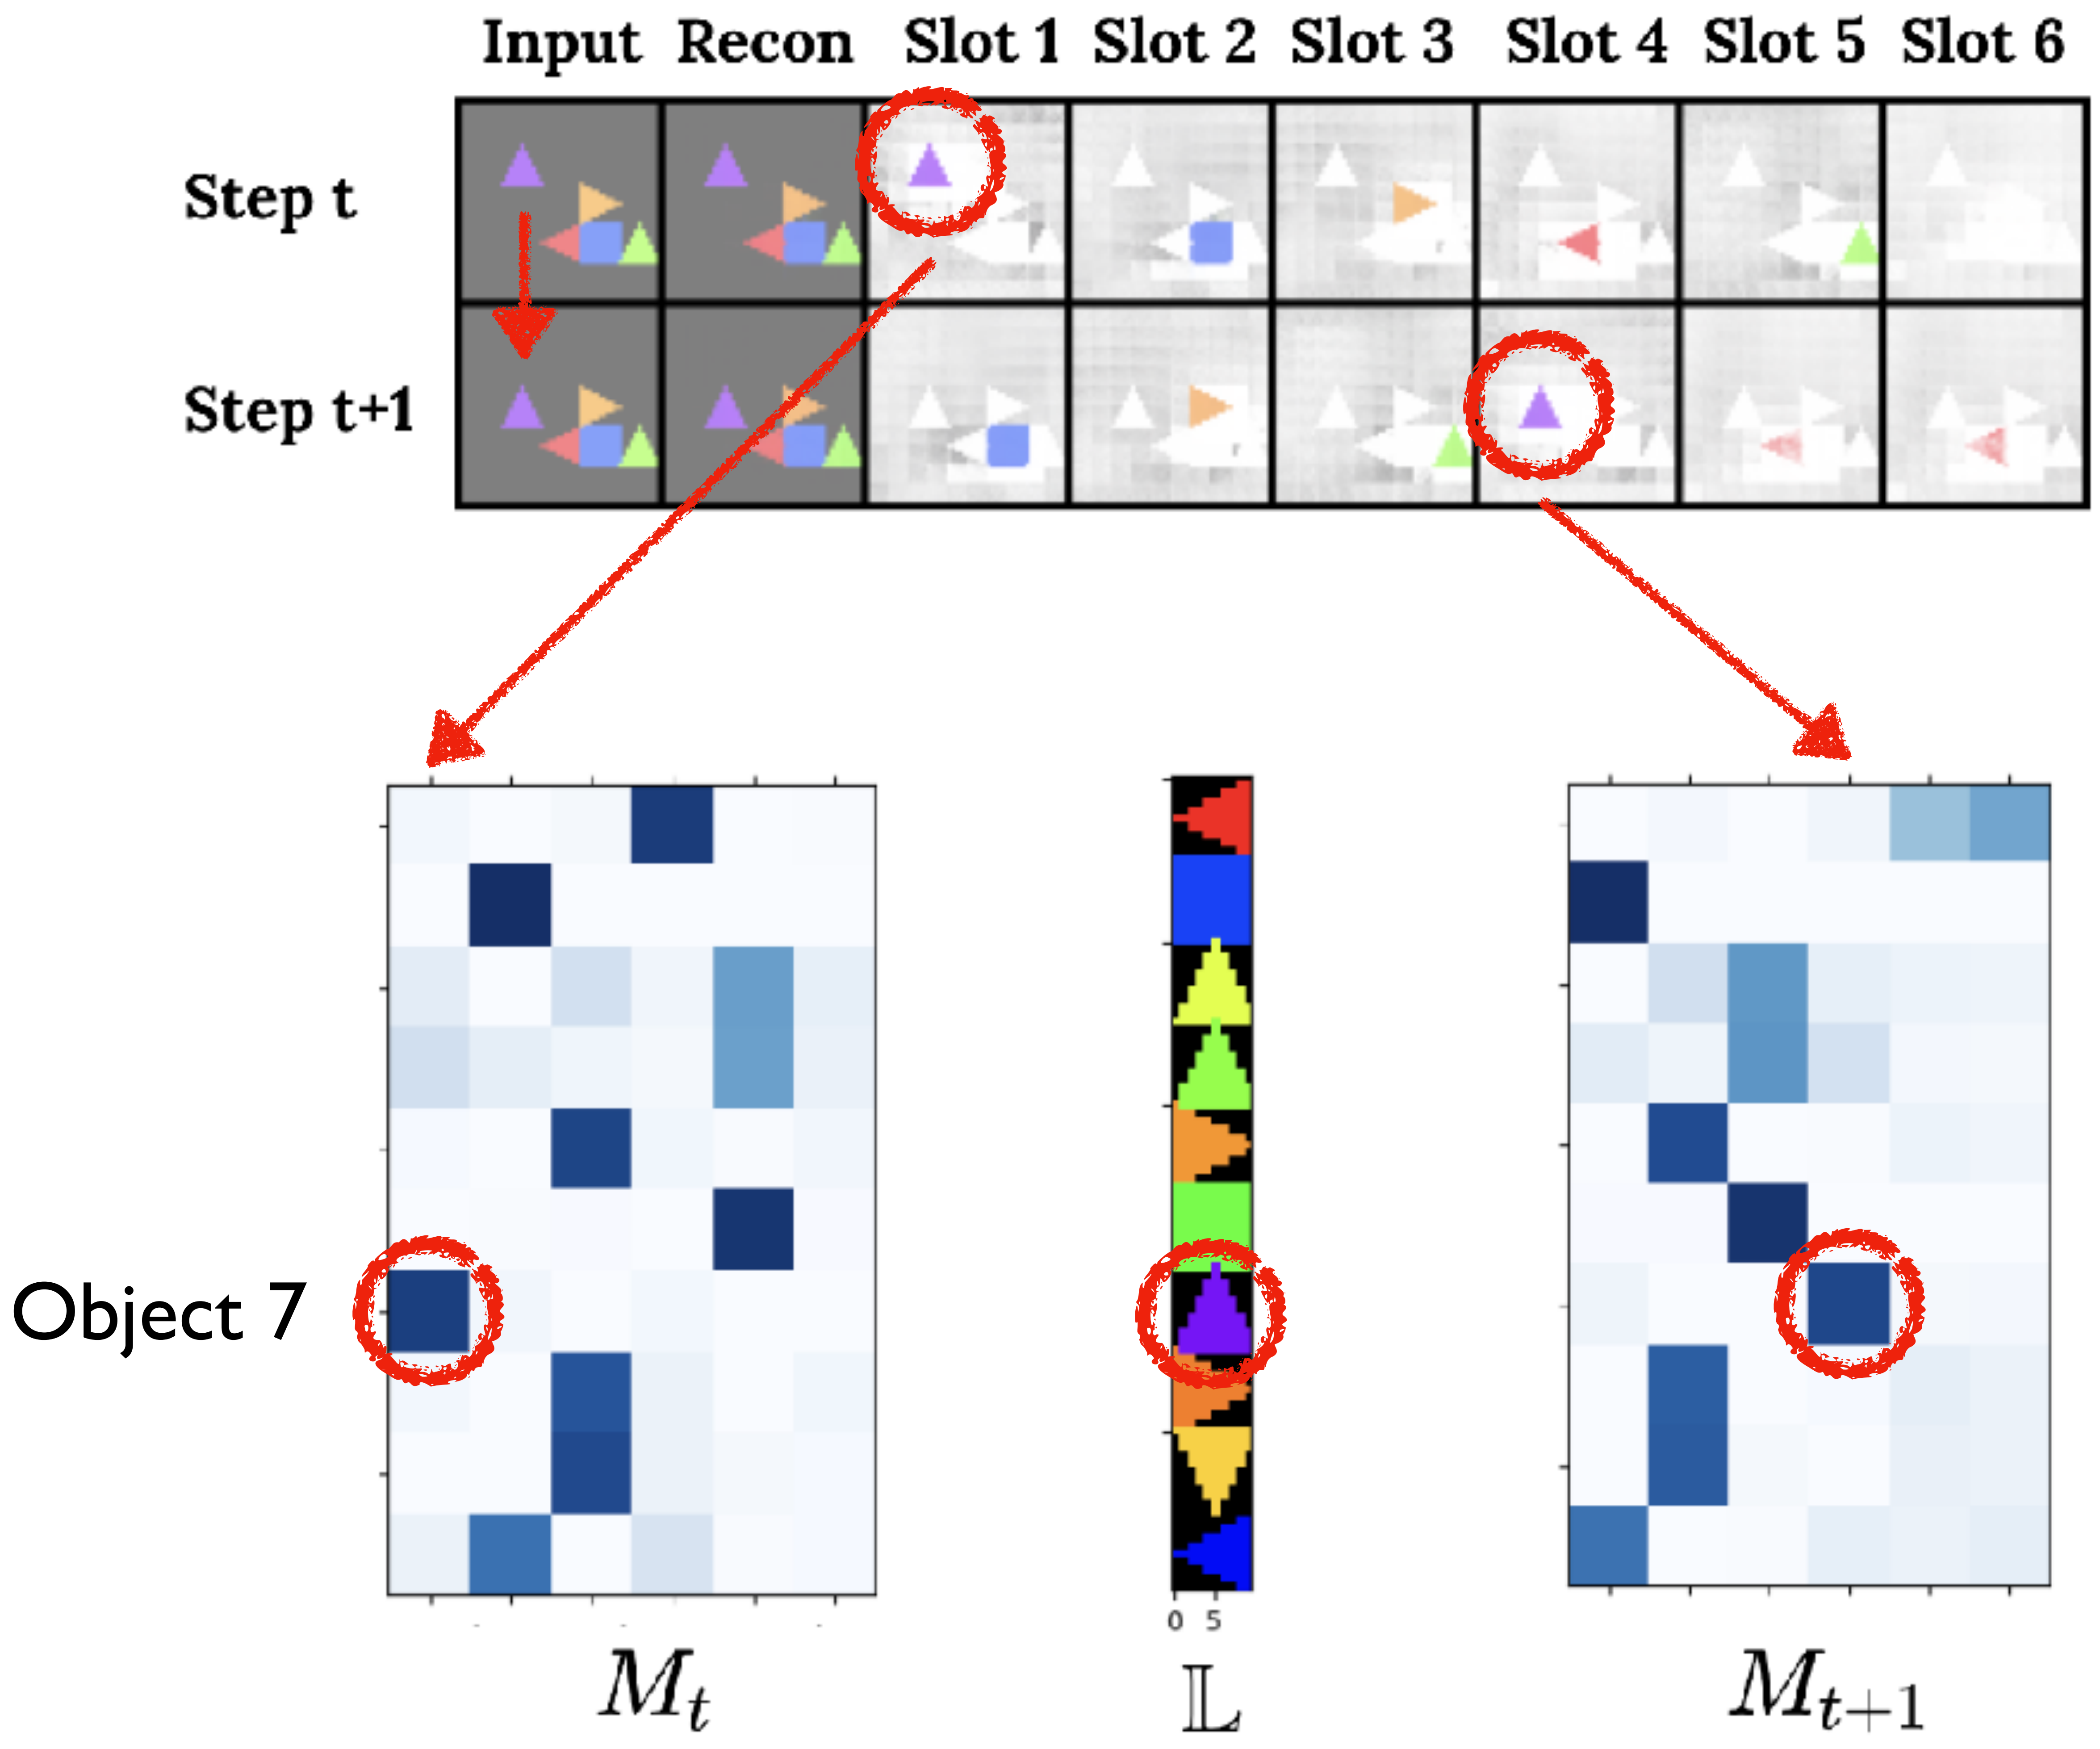

K=5 slots
5+1 rows
(+ 1 background)

N=10 objects
10 columns

Found object identity through actions (unknown identity)

Takeaways

- +: Symmetry and object compositionality are useful

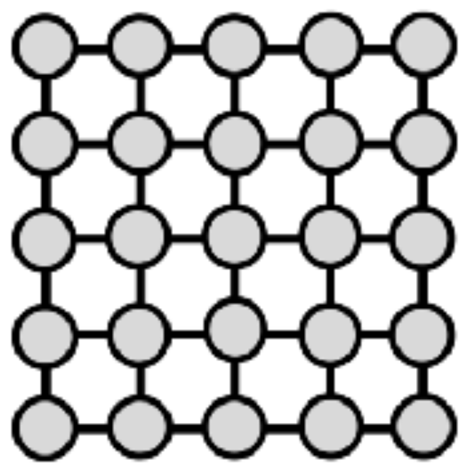

- -: Still assume we know them *a priori*

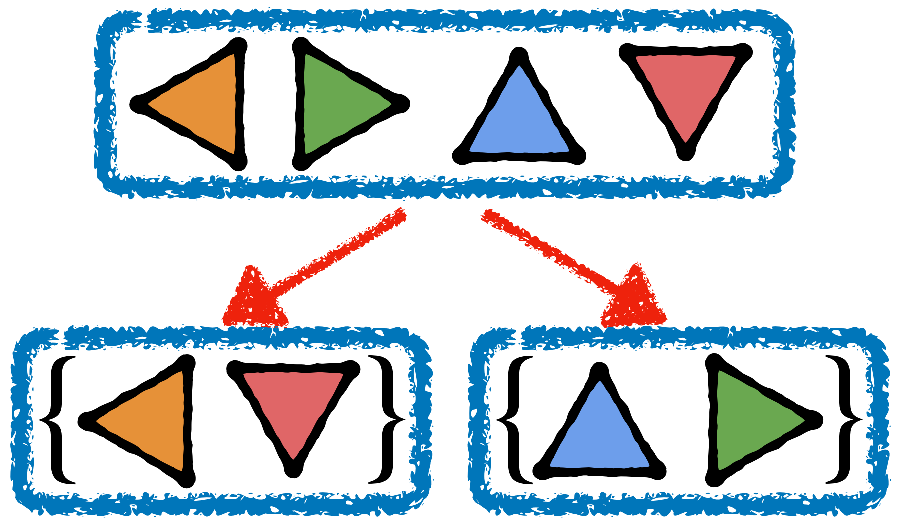

- -: Still need complicated architecture to use them and need more memory+time to train

- Future: More efficient architectures; Applications on more diverse problems, Less additional efforts on identifying/using geometric structure

Topics

- *Learning for* End-to-end Planning
  — How to learn to plan and improve stability

- *Symmetric (and Compositional)* Structure in Robot Planning
  — How does symmetry improve (path) planning and world modeling

- *Symmetric* Structure in Robot Learning
  — How could symmetry improve learning robot skills

Grey items are secondary

Motivation

- In learning robotic skills, such as grasping, need data efficiency and generalizability
- Need to be aware of geometric structures like symmetry

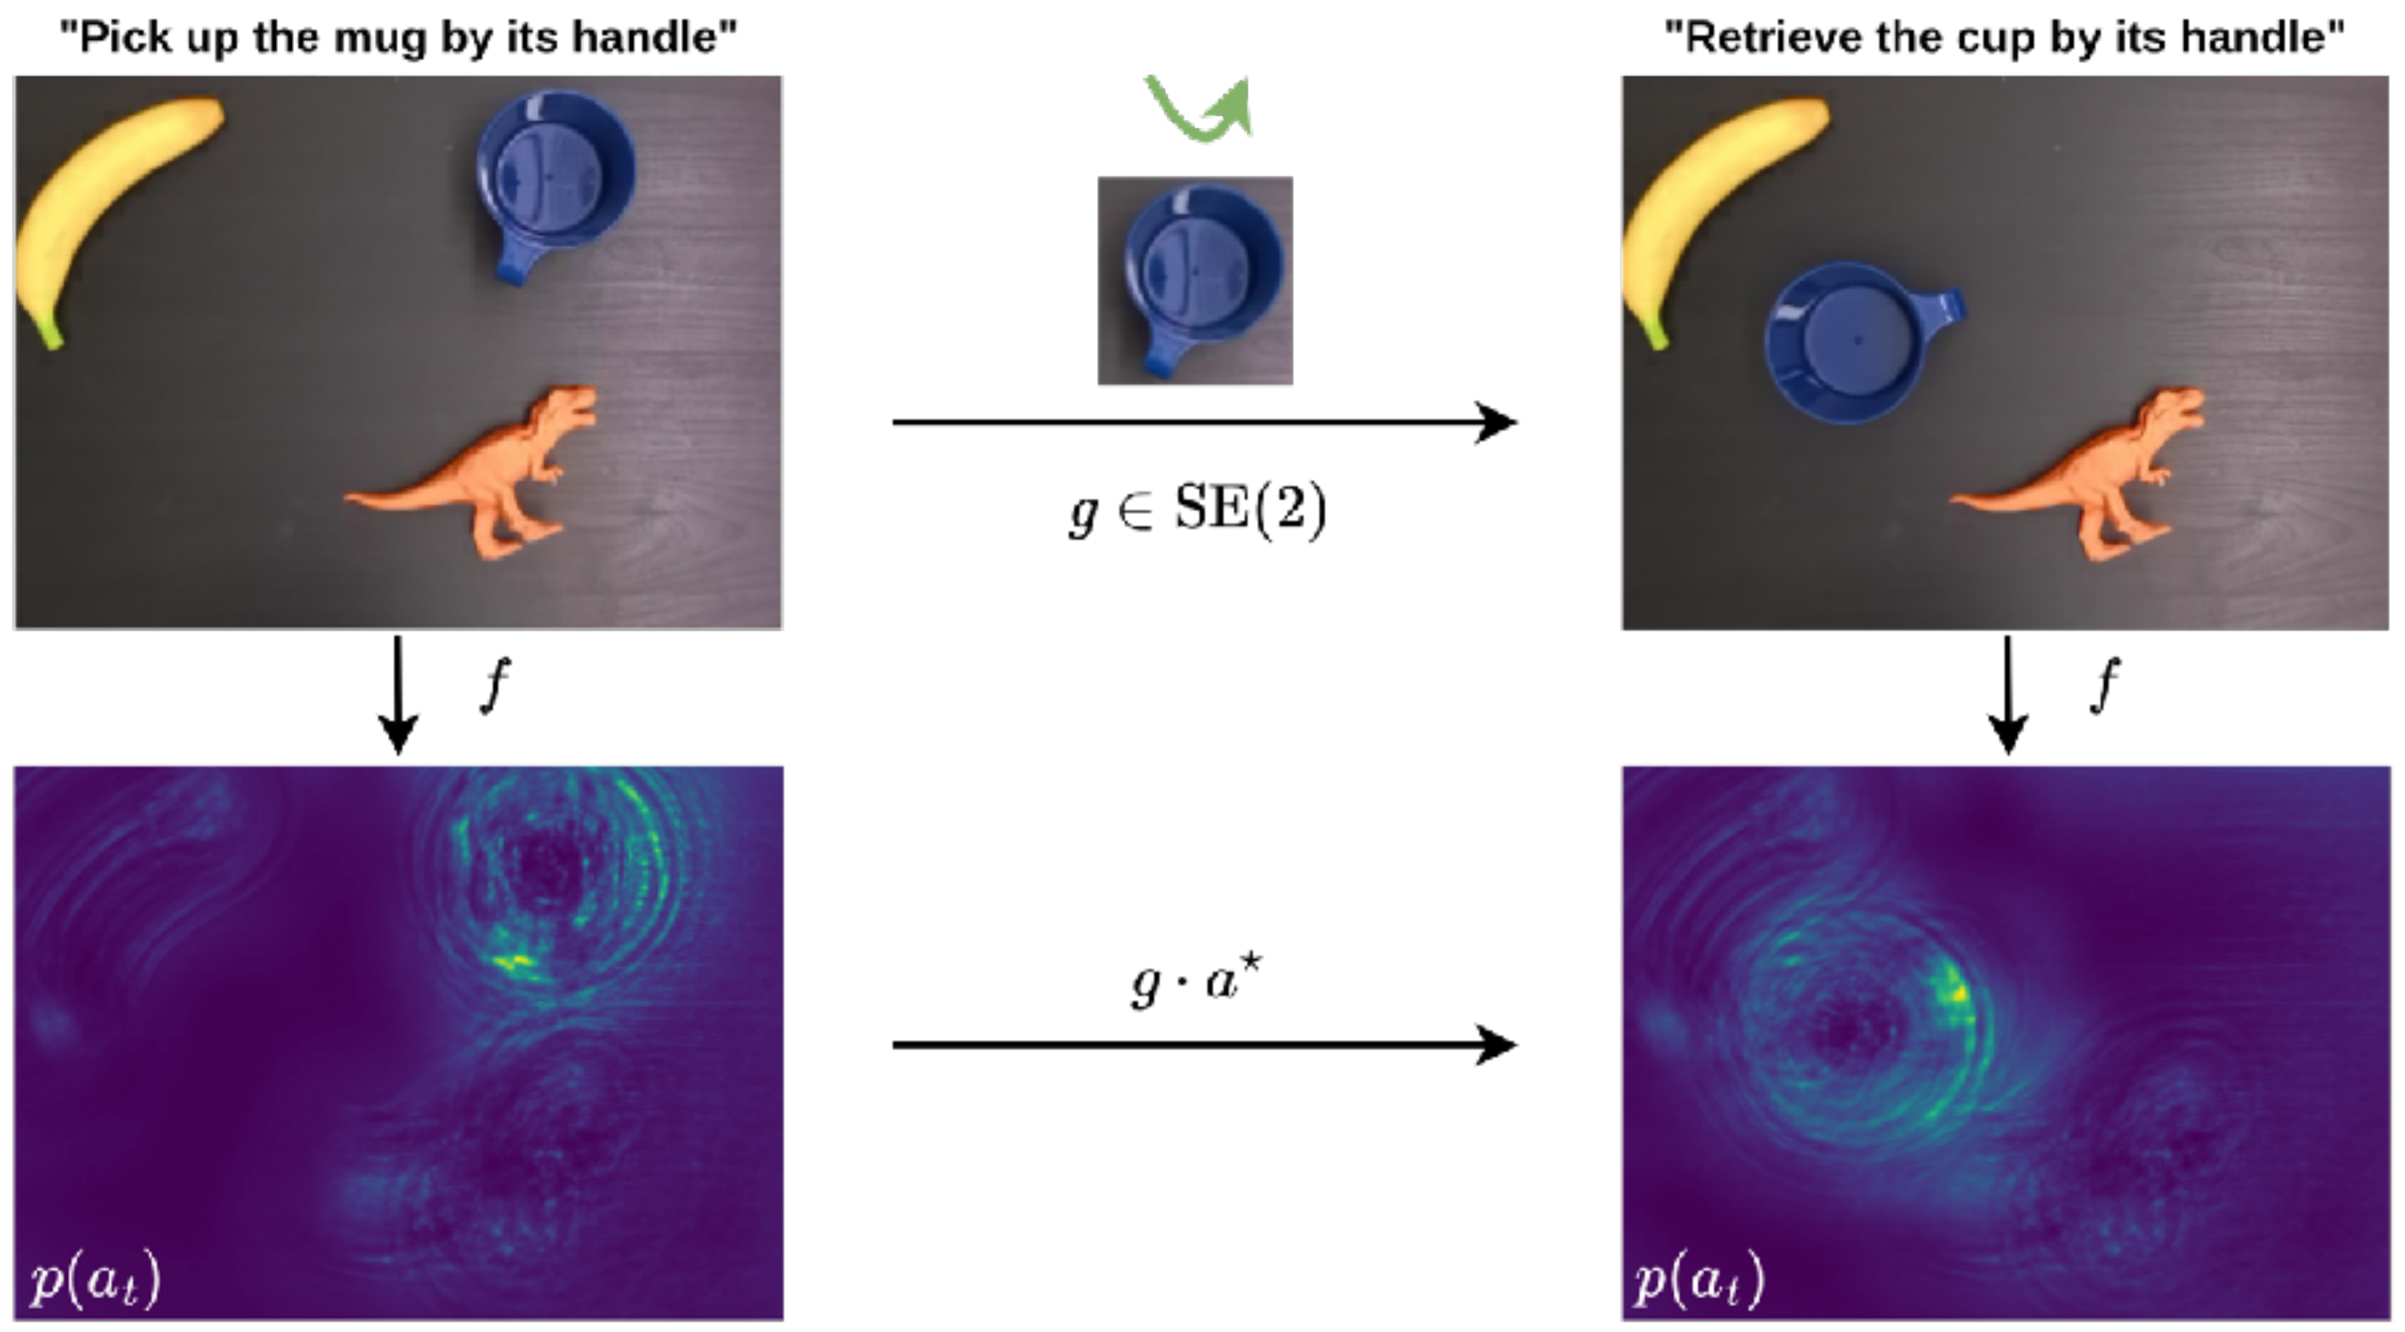

Symmetry in Language-Conditioned Grasping

Language-Conditioned Equivariant Grasp

In Submission, 2023

Haojie Huang, Mingxi Jia, Zhewen Zhang, Ondrej Biza, Linfeng Zhao, Robin Walters, Robert Platt

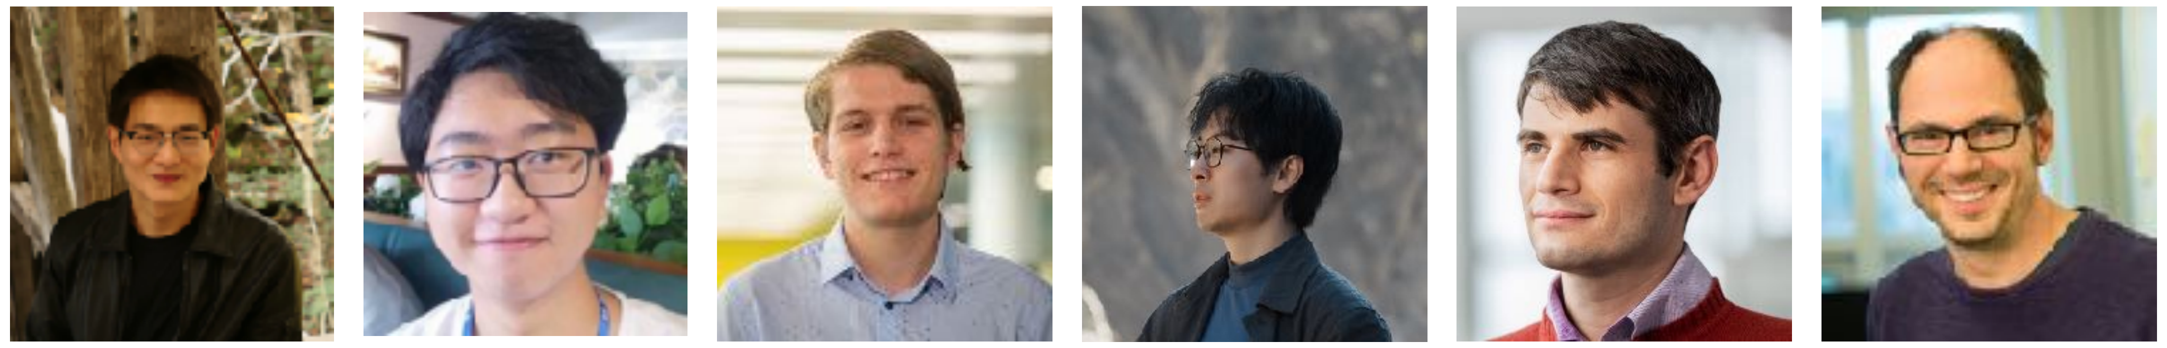

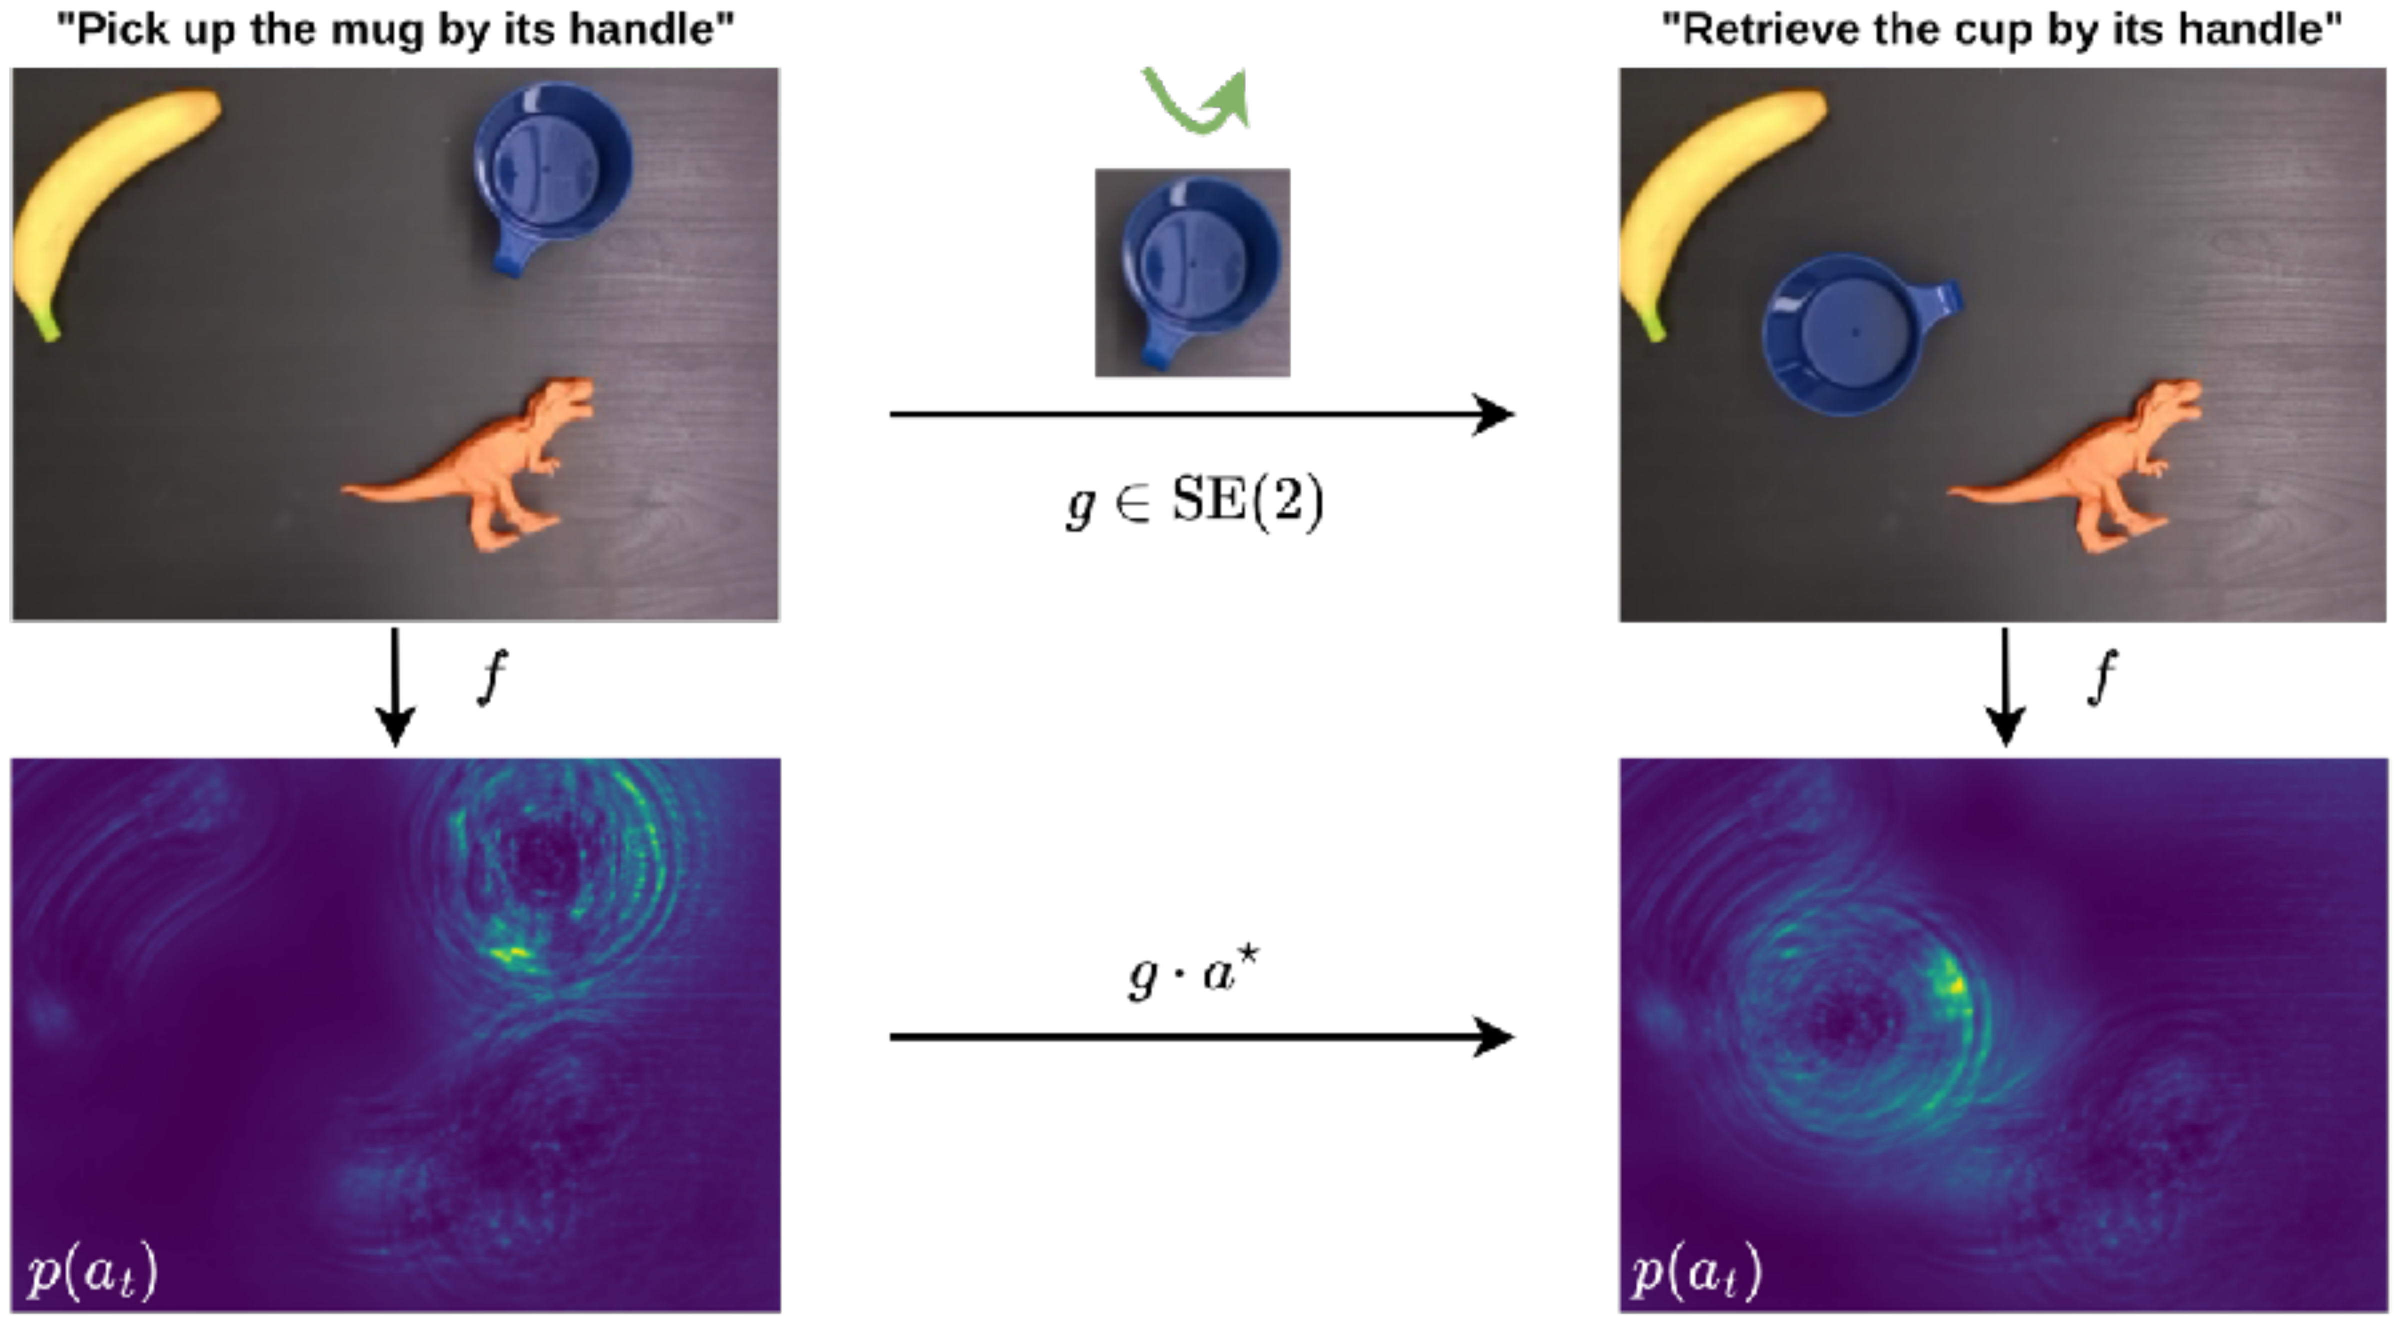

Demo: Training Objects/Parts

Demo: Unseen Objects

Takeaways

- How to ground language while preserving geometric structure is a challenge for the literature
- The “steerable kernel” → sample efficiency + accuracy for rotation
- Object-Part grasp benchmark: Shows language conditioning enables some degree of object-part compositionality

Other Related Papers

- Learning Symmetric Embeddings for Equivariant World Models. ICML 2022.
  Jung Yeon Park*, Ondrej Biza*, ***Linfeng Zhao***, Jan Willem van de Meent, Robin Walters.

- Equivariant Single View Pose Prediction Via Induced and Restriction Representations. NeurIPS 2023.
  Owen Howell, David Klee, Ondrej Biza, ***Linfeng Zhao***, Robin Walters

- Can Euclidean Symmetry Help in Reinforcement Learning and Planning?. arXiv 2023.
  Linfeng Zhao, Owen Howell, Jung Yeon Park, Xupeng Zhu, Robin Walters, Lawson L.S. Wong

Takeaways

- Learning representation for planning from data is a challenge.
- Geometric structures may enhance learning and planning on robotic tasks like navigation and manipulation.
  - For example, considering symmetry and compositionality improves efficiency, generalization, scalability etc.
- Deciding which structures to inject as inductive biases is hidden behind the scene and hugely impacts the performance.
- Other challenges still remain in the path towards generally intelligent robots.

Future directions

Happy to further chat and collaborate!

- Structured representation for planning
    - how to utilize equivariance in continuous environment and their algorithms
    - how to make 3D equivariance from 2D multi-view
    - how to make pretrained foundation models equivariant
- Paradigms for learning representations for planning
    - how to exploit and improve planning in foundation models

Acknowledgments

Lawson Wong

Robin Walters

Ondrej Biza

Jung Yeon Park

Hongyu Li

Xupeng Zhu

Huazhe Xu

Robert Platt

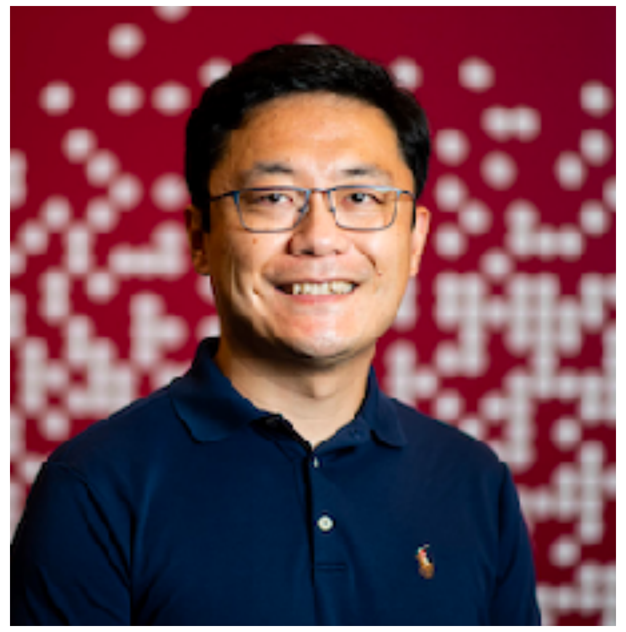

Huaizu Jiang

Owen Howell

Haojie Huang

Lingzhi Kong