



DeepRob
[Group 1] Lecture 04
Pranay, Aditya, Siddharth
Deformable Object Manipulation
University of Minnesota



DR

Objects



Rigid Objects (Credits: YCB Objects and Models)



DR

Objects



Rigid Objects (Credits: YCB Objects and Models)



Credits: PartNet-Mobility dataset



Credits: Dune 2021



Credits: GettyImages



DR

Deformable Objects

- Linear
- Planar
- 3D objects



Linear Object



Planar Objects



3D Object

“Feature sensing and robotic grasping of objects with uncertain information: A review.”, Wang, C., Zhang, X., Zang, X., Liu, Y., Ding, G., Yin, W., & Zhao, J. Sensors, 20(13), 3707 (2020).



Applications


Why do we even care about deformable objects?

- Healthcare
- Food Industry
- Textile
- Agriculture



Challenges

- Self Occlusion
- Complex dynamics
- High degrees of freedom





Human vs Robot



Credits: @DaveHax



Human vs Robot



Credits:@DaveHax



Credits: pi0

Which Sensors? Sensing which properties?
Actions?





Decoding Deformable Object Manipulation

Sensors	Property	Action
Camera	Shape	Lifting
LR Tactile	Size	Dragging
HR Tactile	Color	Pulling
Force/Torque	Material	Twisting
Spectrometer	Construction	Flinging
Auditory	Stiffness	Sliding
	Elasticity	Pressing
	Weight	
	Friction	



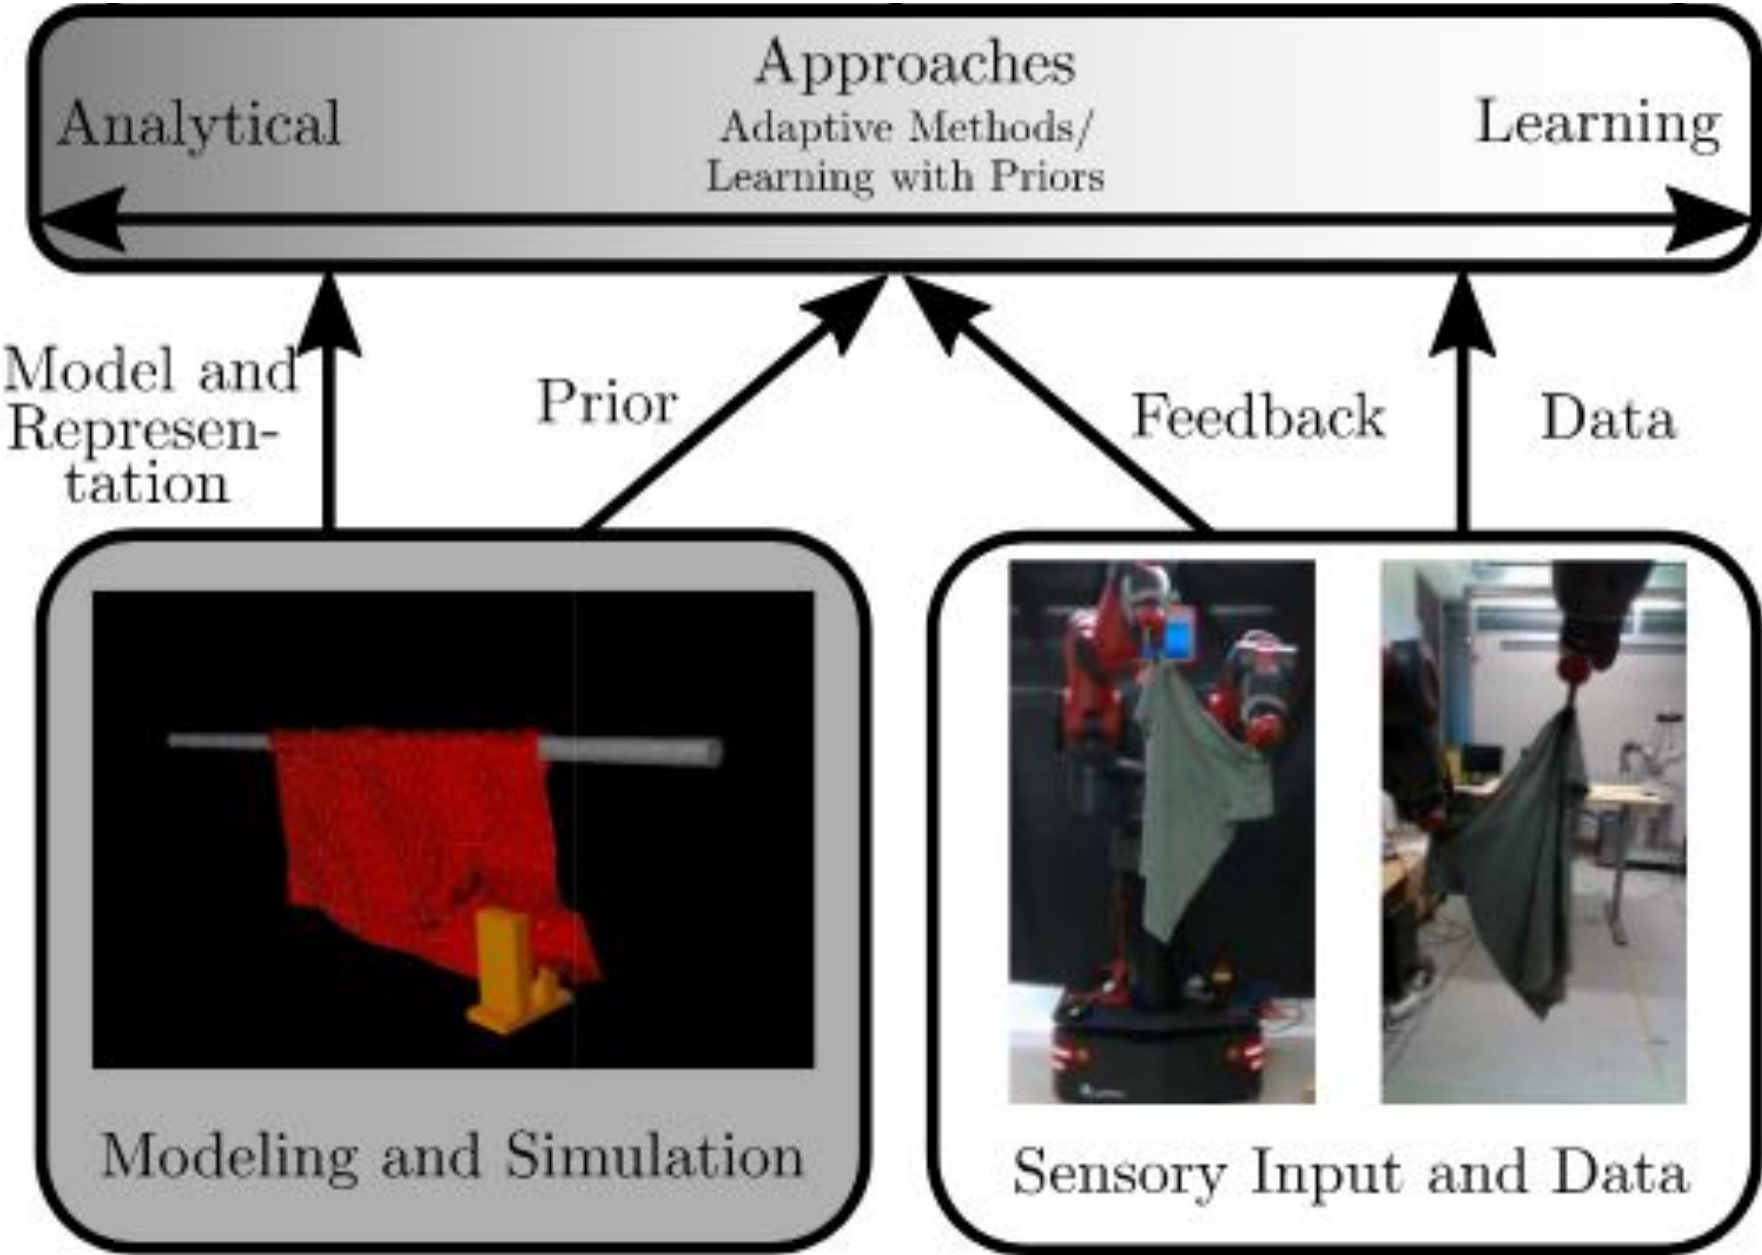


End-to-end Workflow

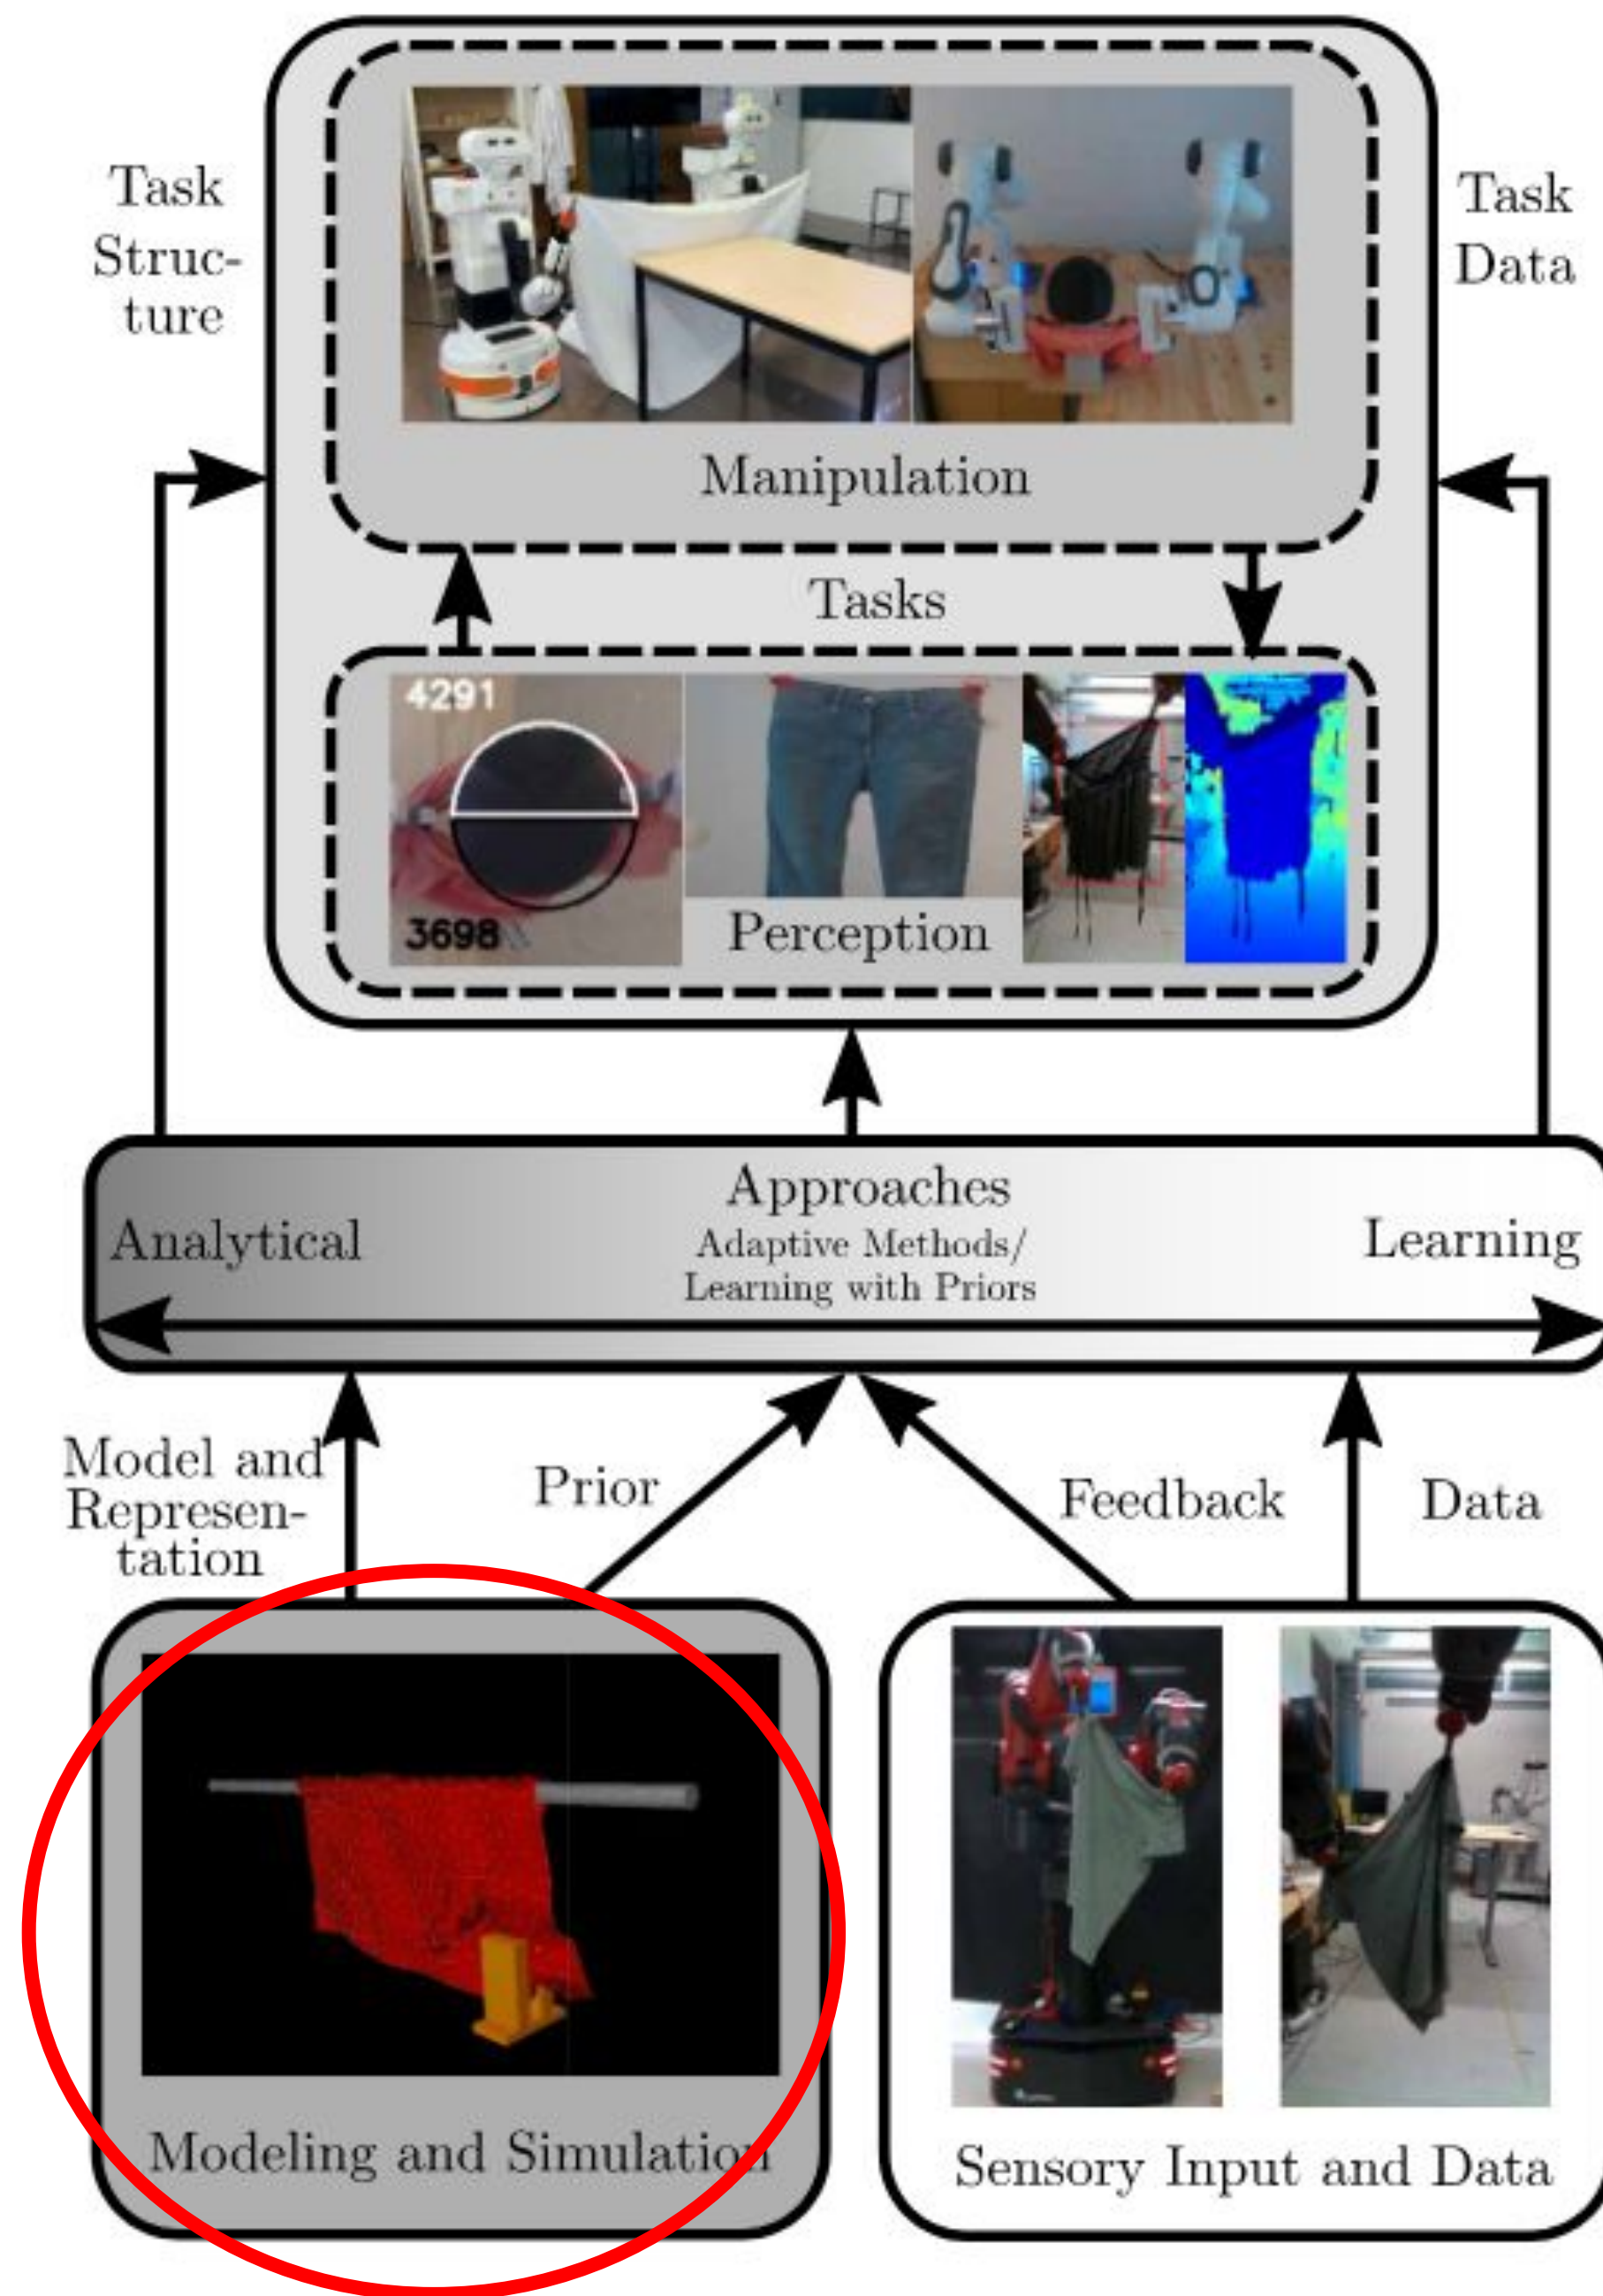




End-to-end Workflow

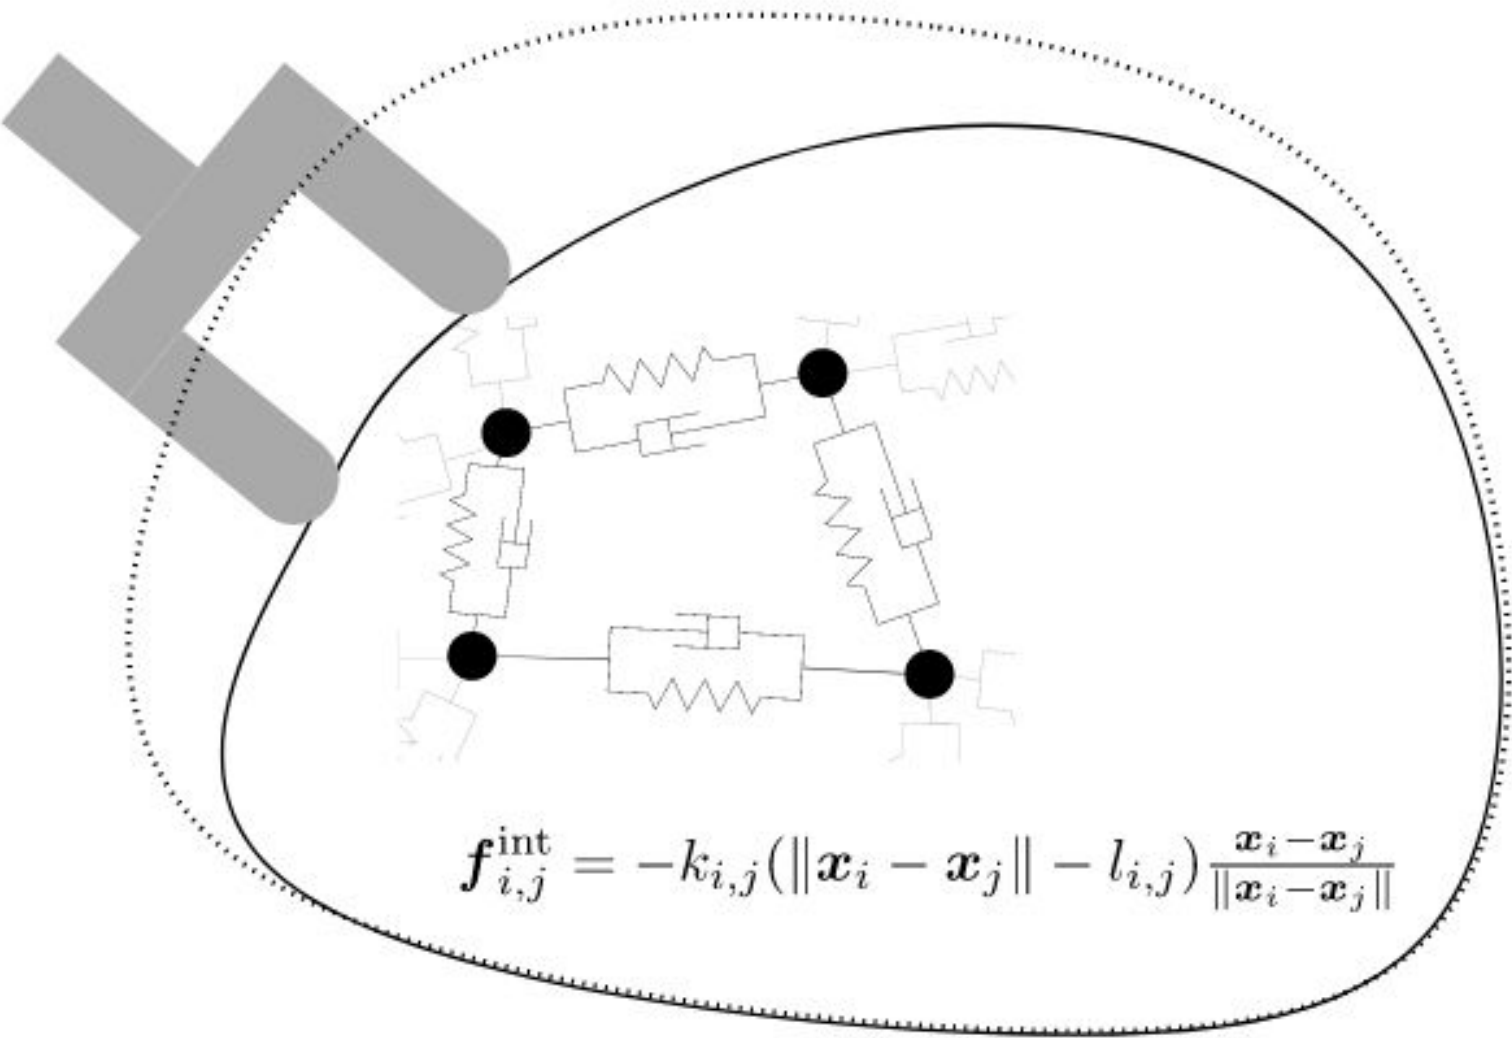


End-to-end Workflow




Physics-based Modeling of deformable objects


a) Mass-spring systems



b) Position-based dynamics




c) Continuum mechanics



Simulators

- Softgym
- MuJoCo
- SOFA
- PyBullet



NVIDIA Flex




Softgym built on Nvidia Flex bindings







Simulators




Real-world Demonstration




Bullet



MuJoCo



SOFA



Softgym



Traditional Methods: Before Deep Learning

- Contour Matching
- Template Matching
- Simple ML models

Lets understand this with an example work!





“Perception for the Manipulation of Socks”, Ping Chuan Wang, Stephen Miller, Mario Fritz, Trevor Darrell, Pieter Abbeel, 2011

- Input:
 - Single Image
- Output:
 - Structure(toe, ankle, etc)
 - Inside-out?
 - Match with candidates
 - Use of LBP, MR8 filter banks
 - Use of SVM classifier




Fig. 1. Given an initial image, we wish to recover the sock configuration.




Fig. 3. The MR8 filter bank consists of 6 gaussian derivative and 2 blob filters. A maximum operations is performed over different orientation variants in order to achieve robustness with respect to rotations.





“Perception for the Manipulation of Socks”, Ping Chuan Wang, Stephen Miller, Mario Fritz, Trevor Darrell, Pieter Abbeel, 2011





After Deep Learning





Data-Driven Models

- Advantages over Physics-based techniques:
 - Greater flexibility in defining state space





Data-Driven Models

- Advantages over Physics-based techniques:
 - Greater flexibility in defining state space
 - Learn dynamics from data (Images/3D Point clouds)



Data-Driven Models

- Advantages over Physics-based techniques:
 - Greater flexibility in defining state space
 - Learn dynamics from data (Images/3D Point clouds)
 - Image-based inputs allow partial observability



Data-Driven Models

- Advantages over Physics-based techniques:
 - Greater flexibility in defining state space
 - Learn dynamics from data (Images/3D Point clouds)
 - Image-based inputs allow partial observability
- Disadvantages:
 - Struggles with domain shifts (e.g., lighting, camera position)



Data-Driven Models

- Advantages over Physics-based techniques:
 - Greater flexibility in defining state space
 - Learn dynamics from data (Images/3D Point clouds)
 - Image-based inputs allow partial observability
- Disadvantages:
 - Struggles with domain shifts (e.g., lighting, camera position)

Solution?



Particle-based Representations

- Rely on 3D Geometric representations (Particles/Meshes) → Robust to changes in Visual conditions.
- Such representations require specific architectures → Capture local structures & Handle data sparsity efficiently.
- PointNet++ for unordered point sets & Graph Neural Networks (GNNs) for mesh-based representations.

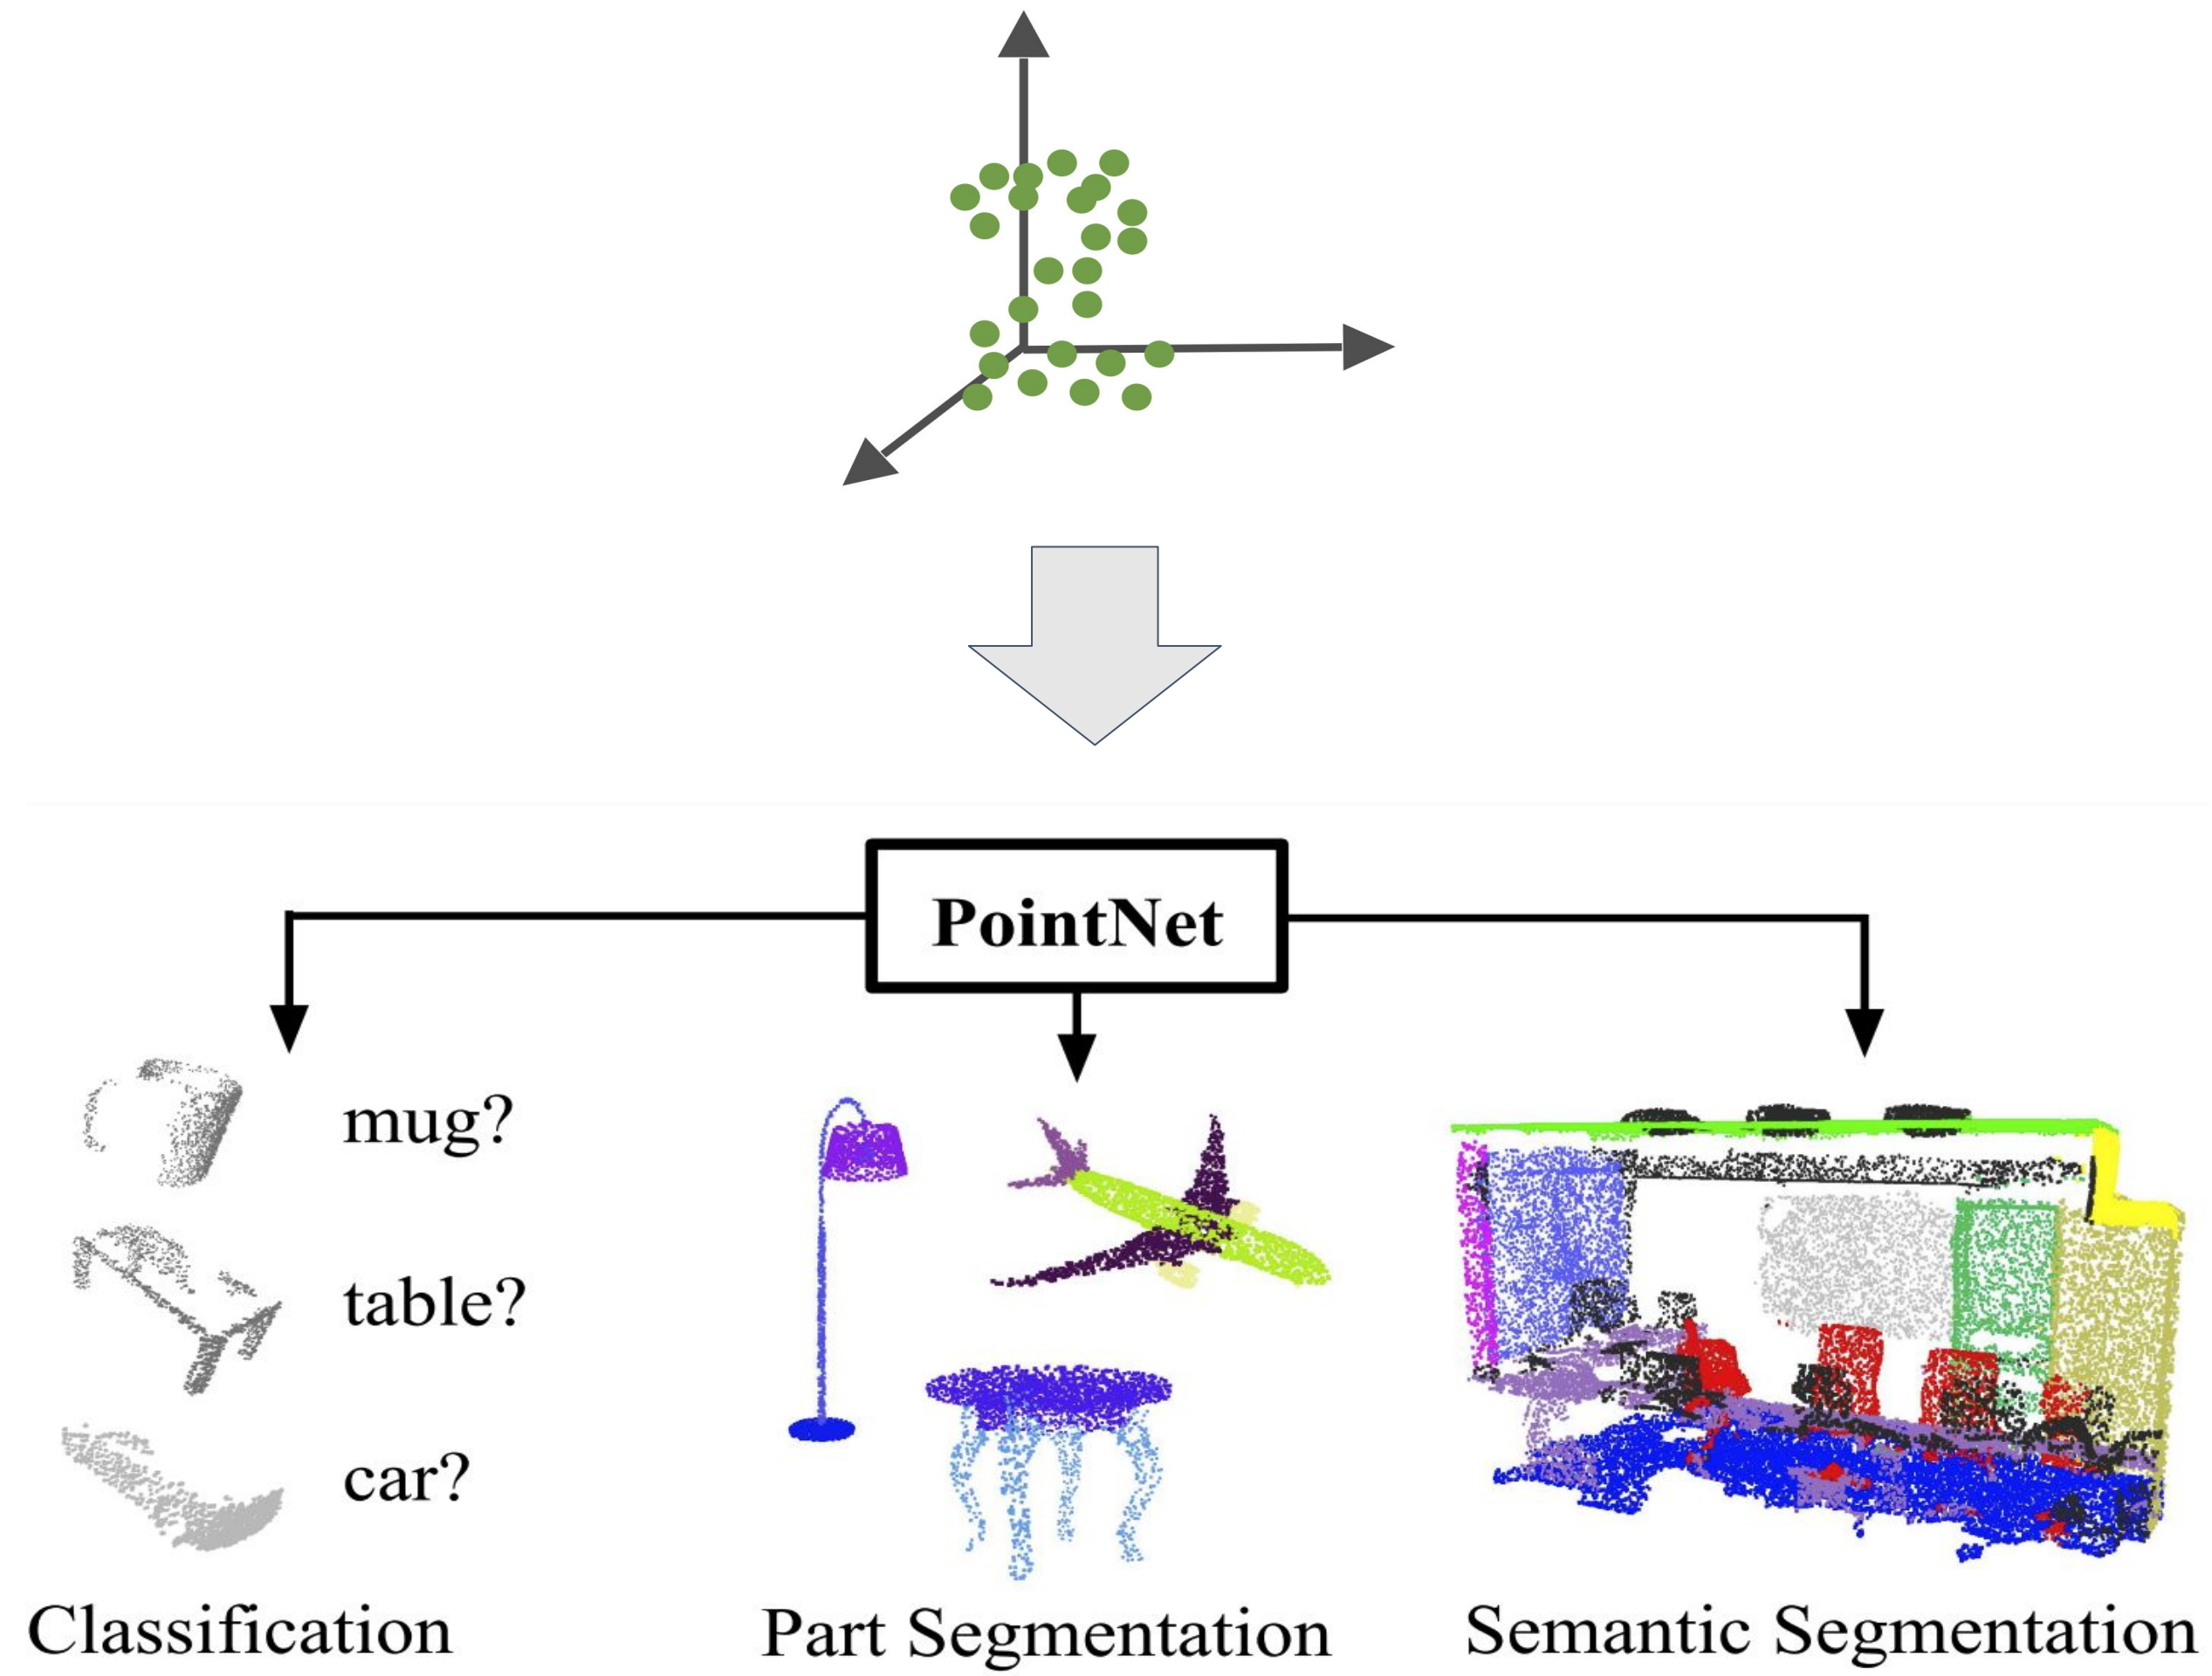


Point Cloud based
Representation of cloth




Mesh-grid based
Representation of cloth


PointNet



Graph Neural Networks




Graph Representation
of Problem

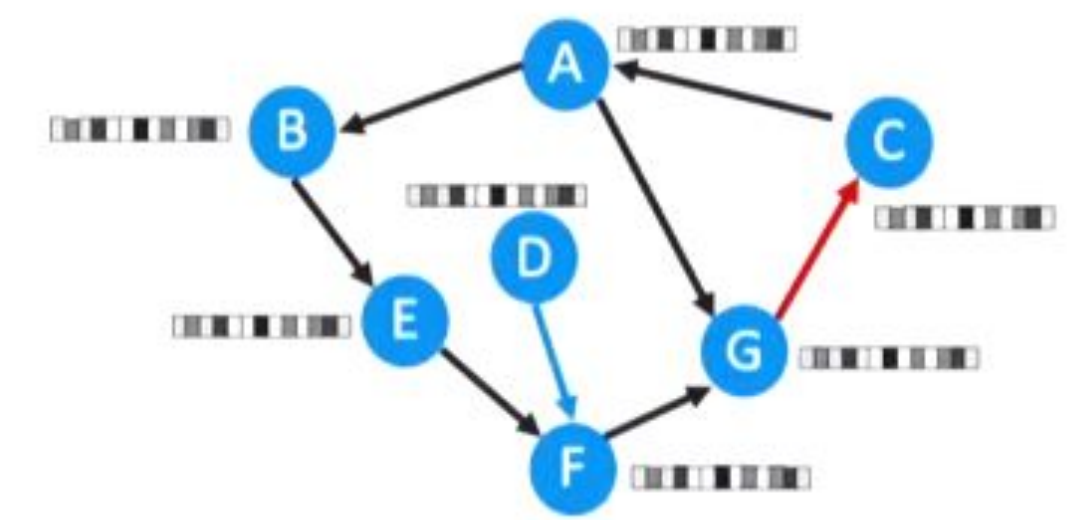


Initial Representation
of each node



Graph Neural Networks



Graph Representation of Problem




Initial Representation of each node



- Cloth mesh nodes
- Ball mesh nodes
- Mesh-space edges
- World-space edges
- Connectivity radius (Example for one ball mesh node)

Graph Neural Networks

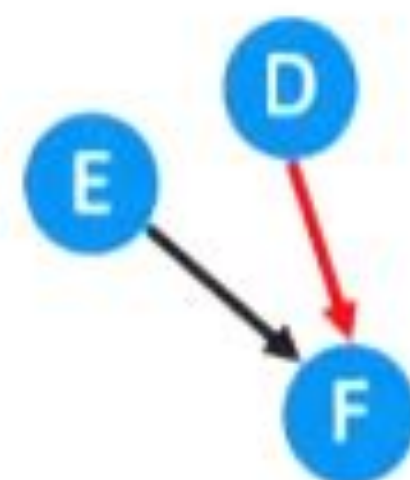


Initial Representation of each node

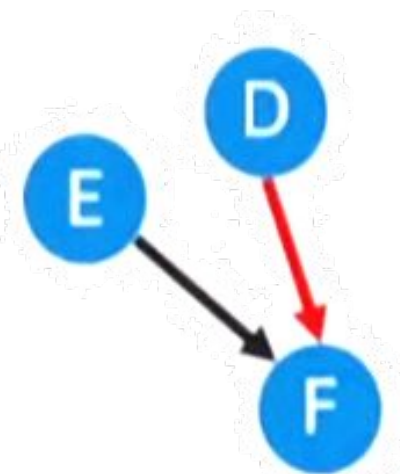

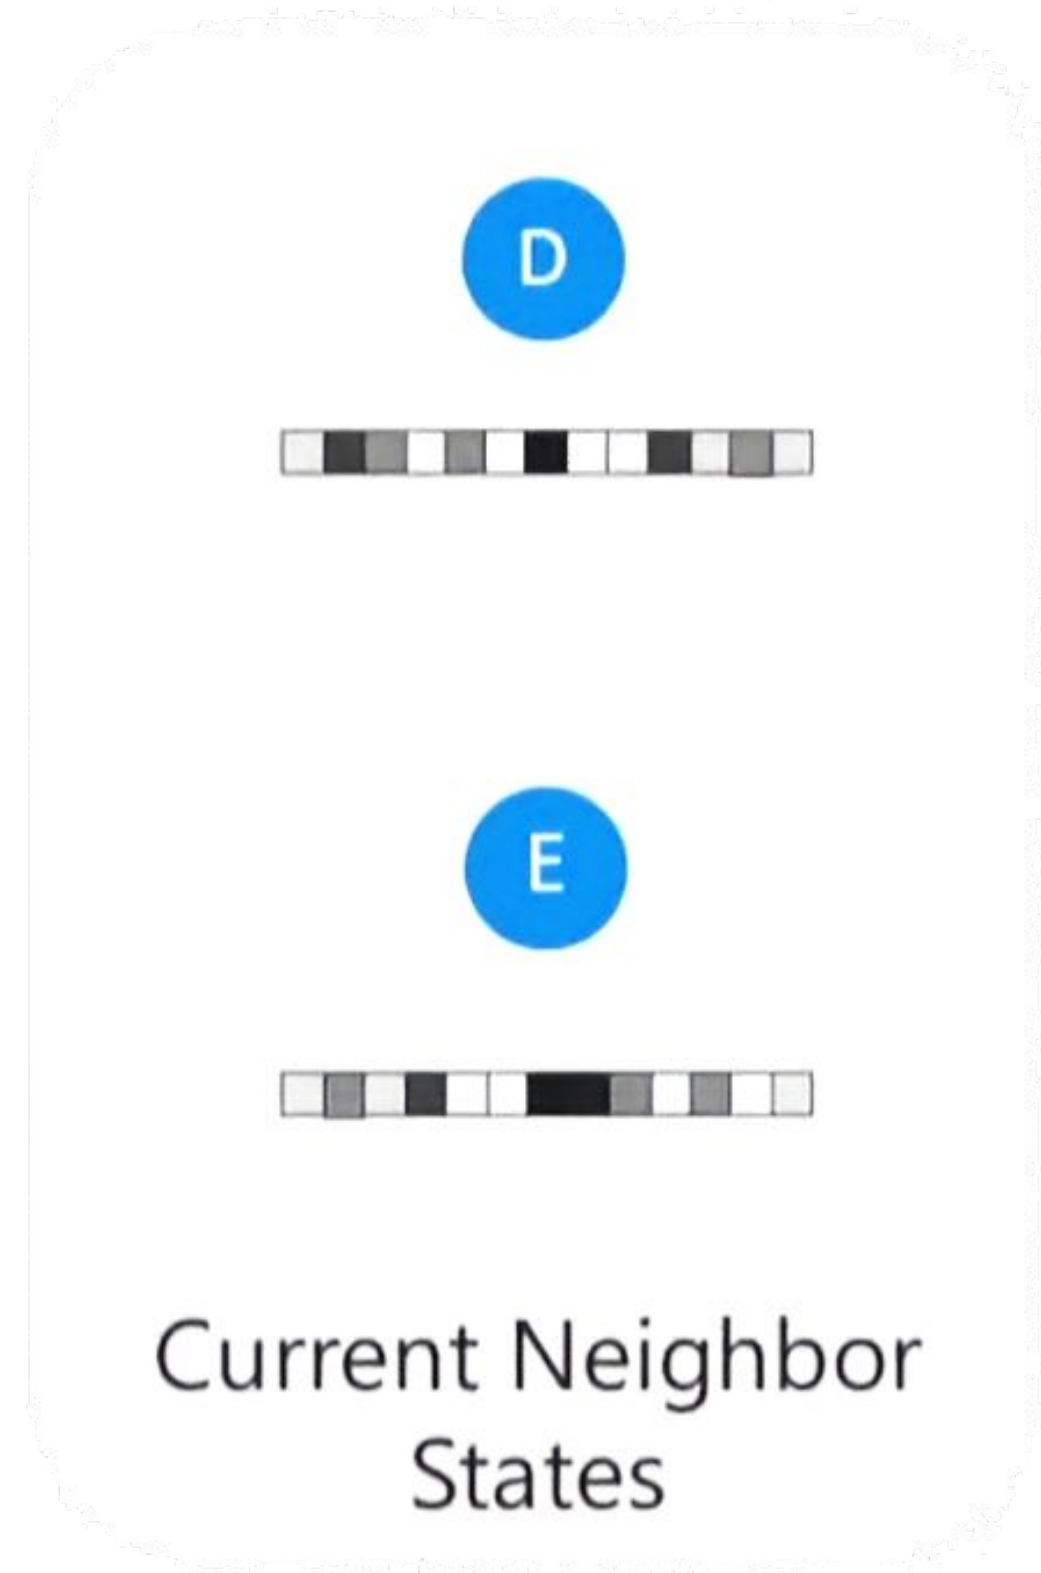
Output Representations of each Node




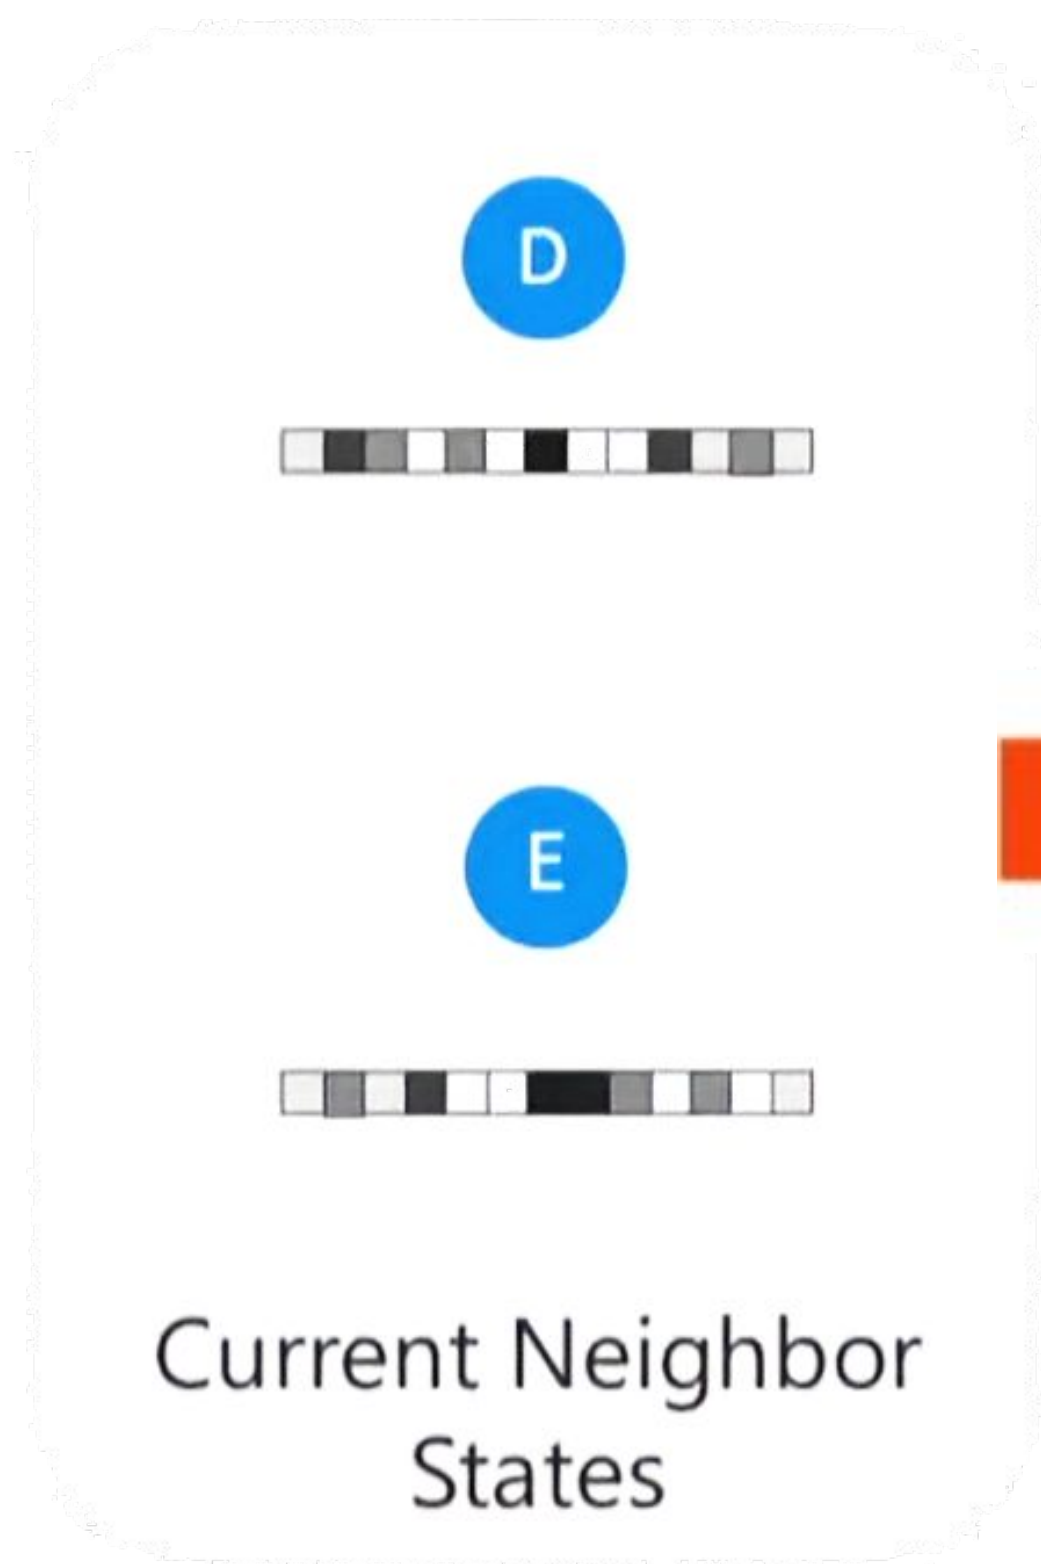
Neural Message Passing



Neural Message Passing



Neural Message Passing

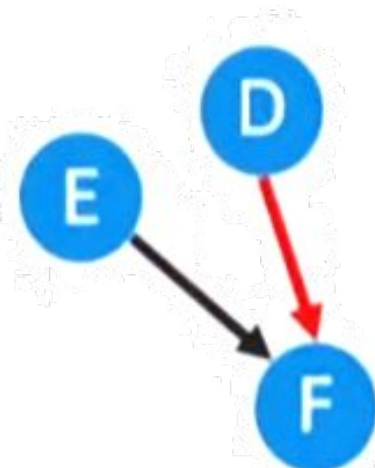
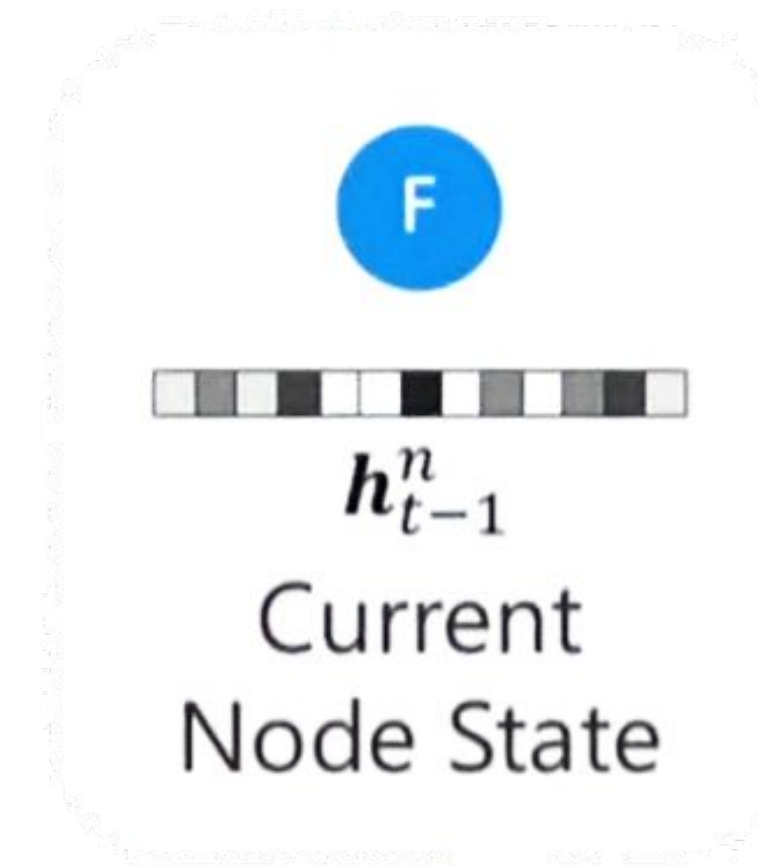


$$\text{Envelope} = f(\text{D} \rightarrow \text{F})$$



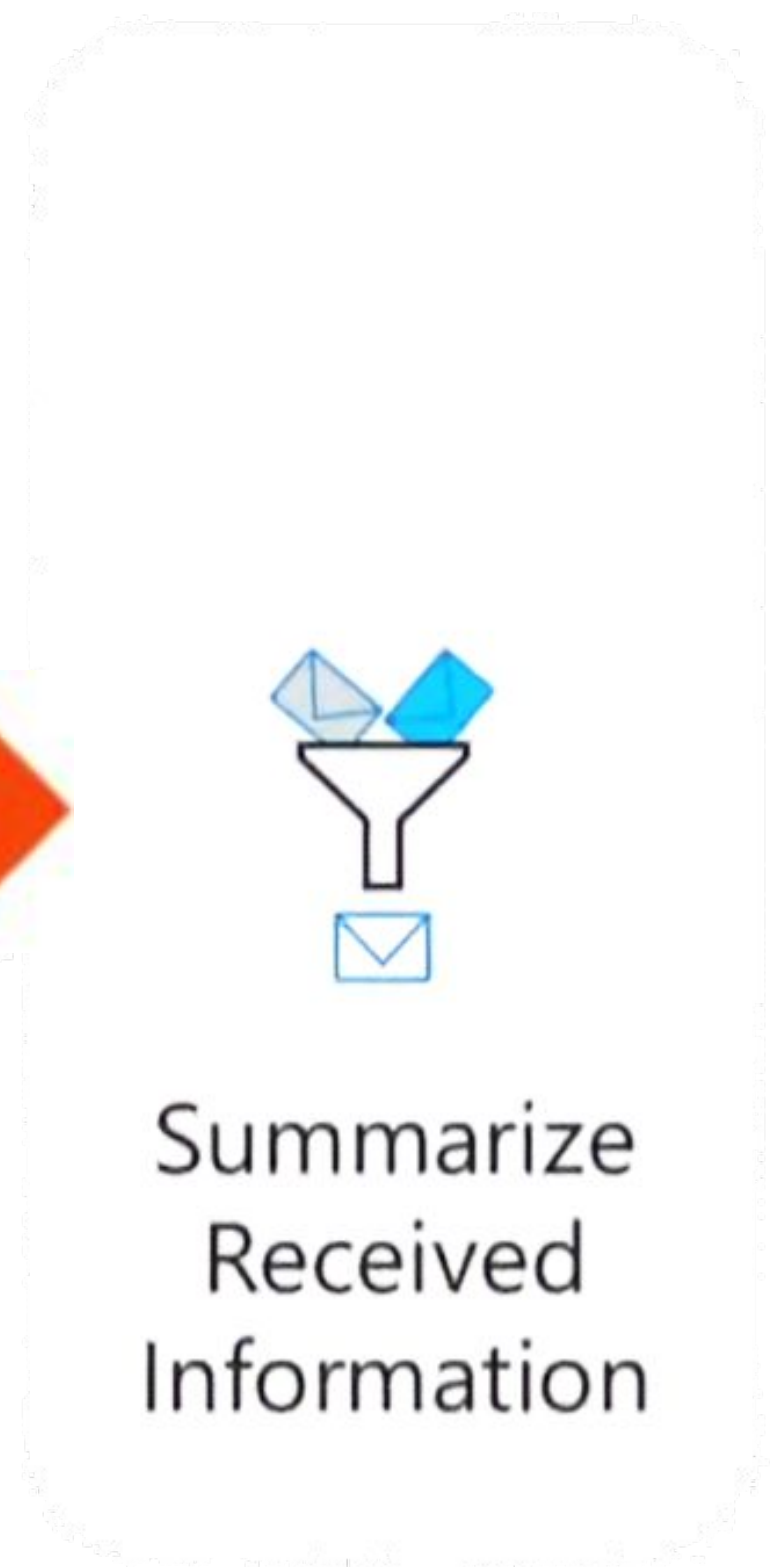
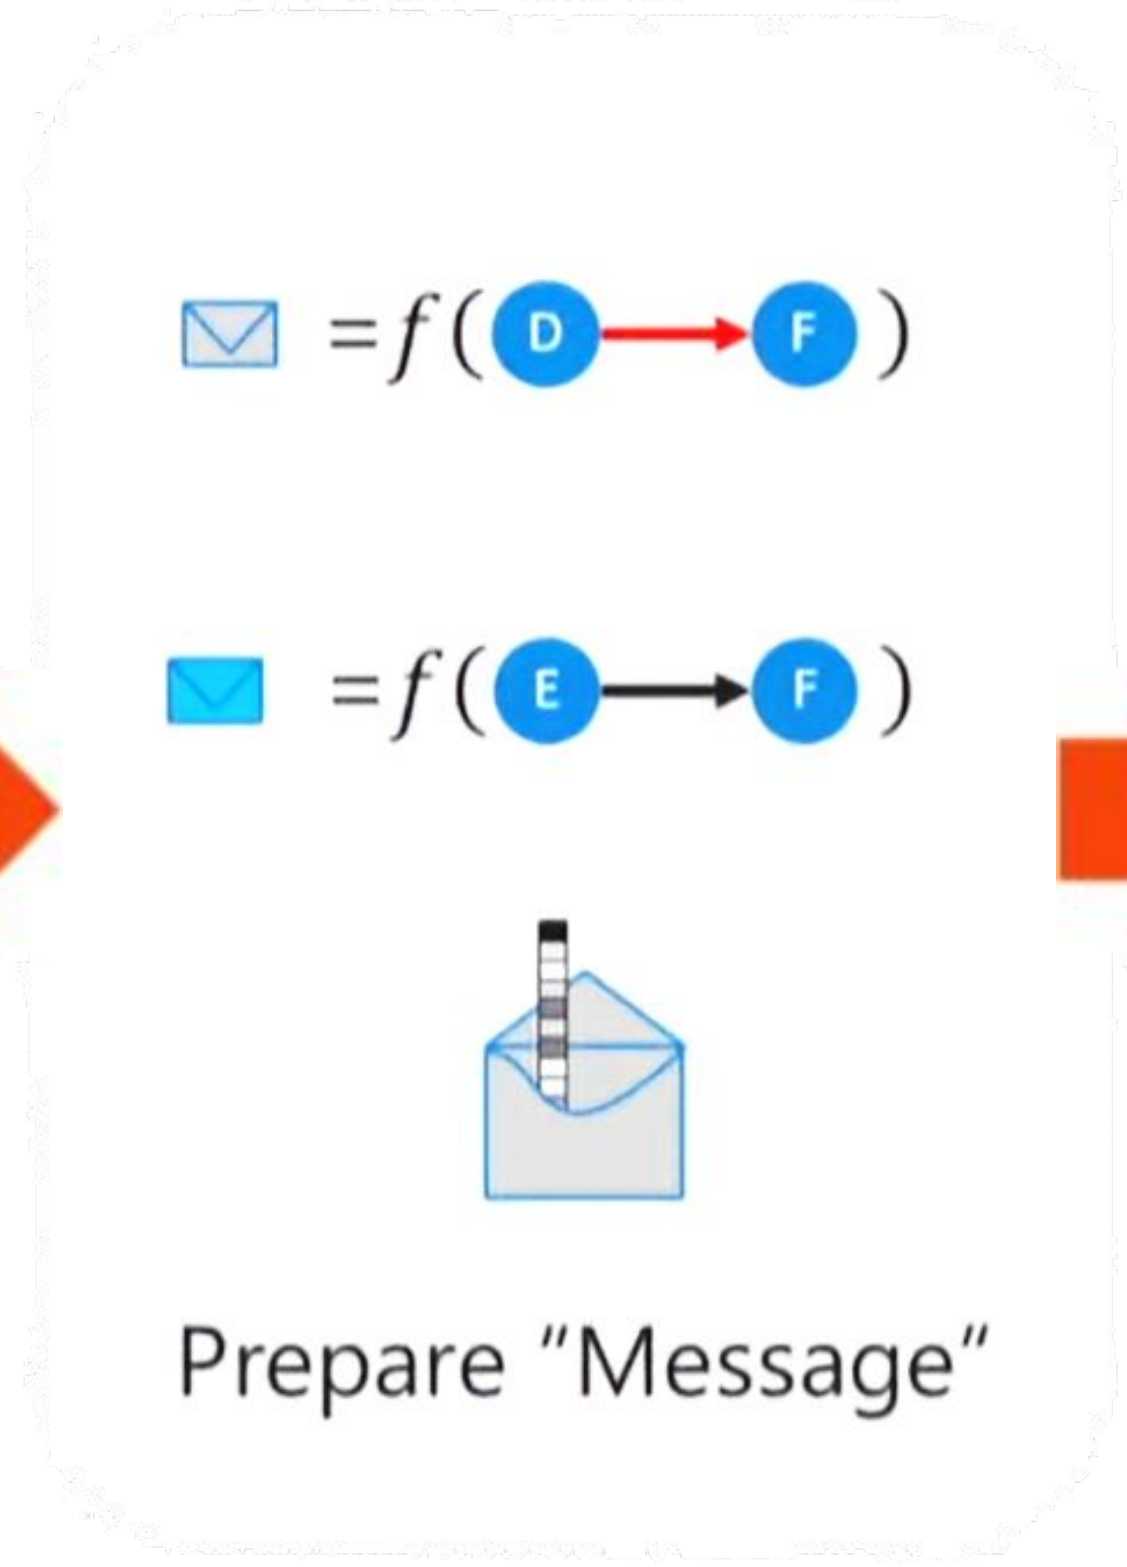
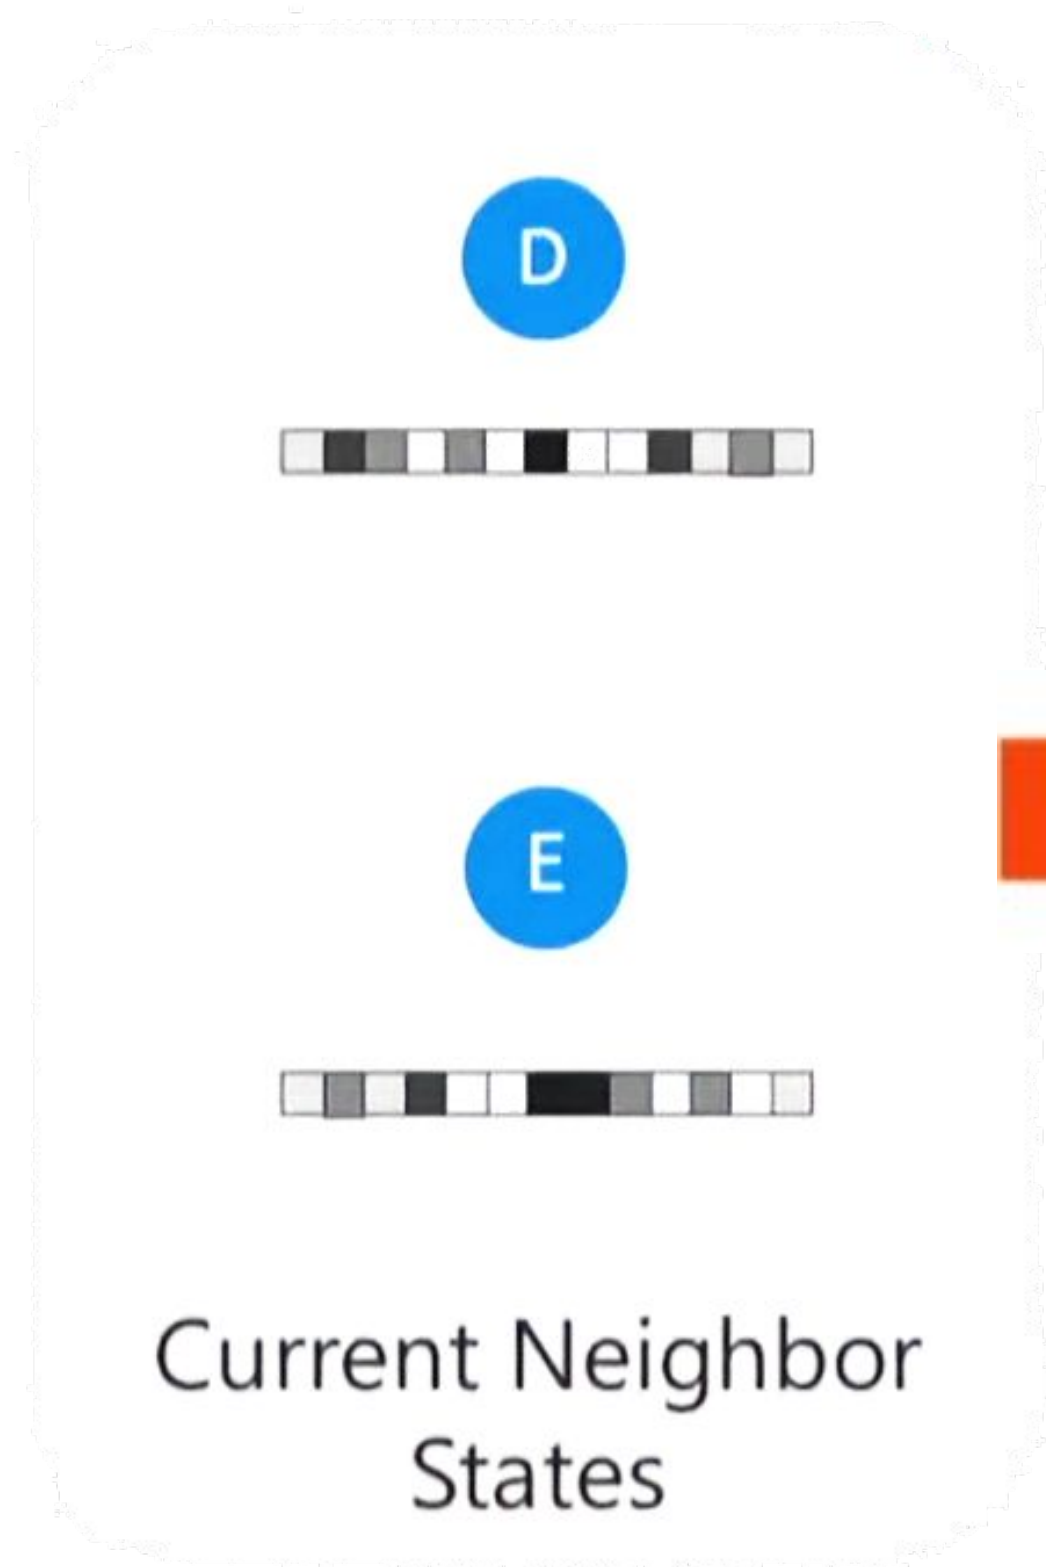
$$\text{Envelope} = f(\text{E} \rightarrow \text{F})$$




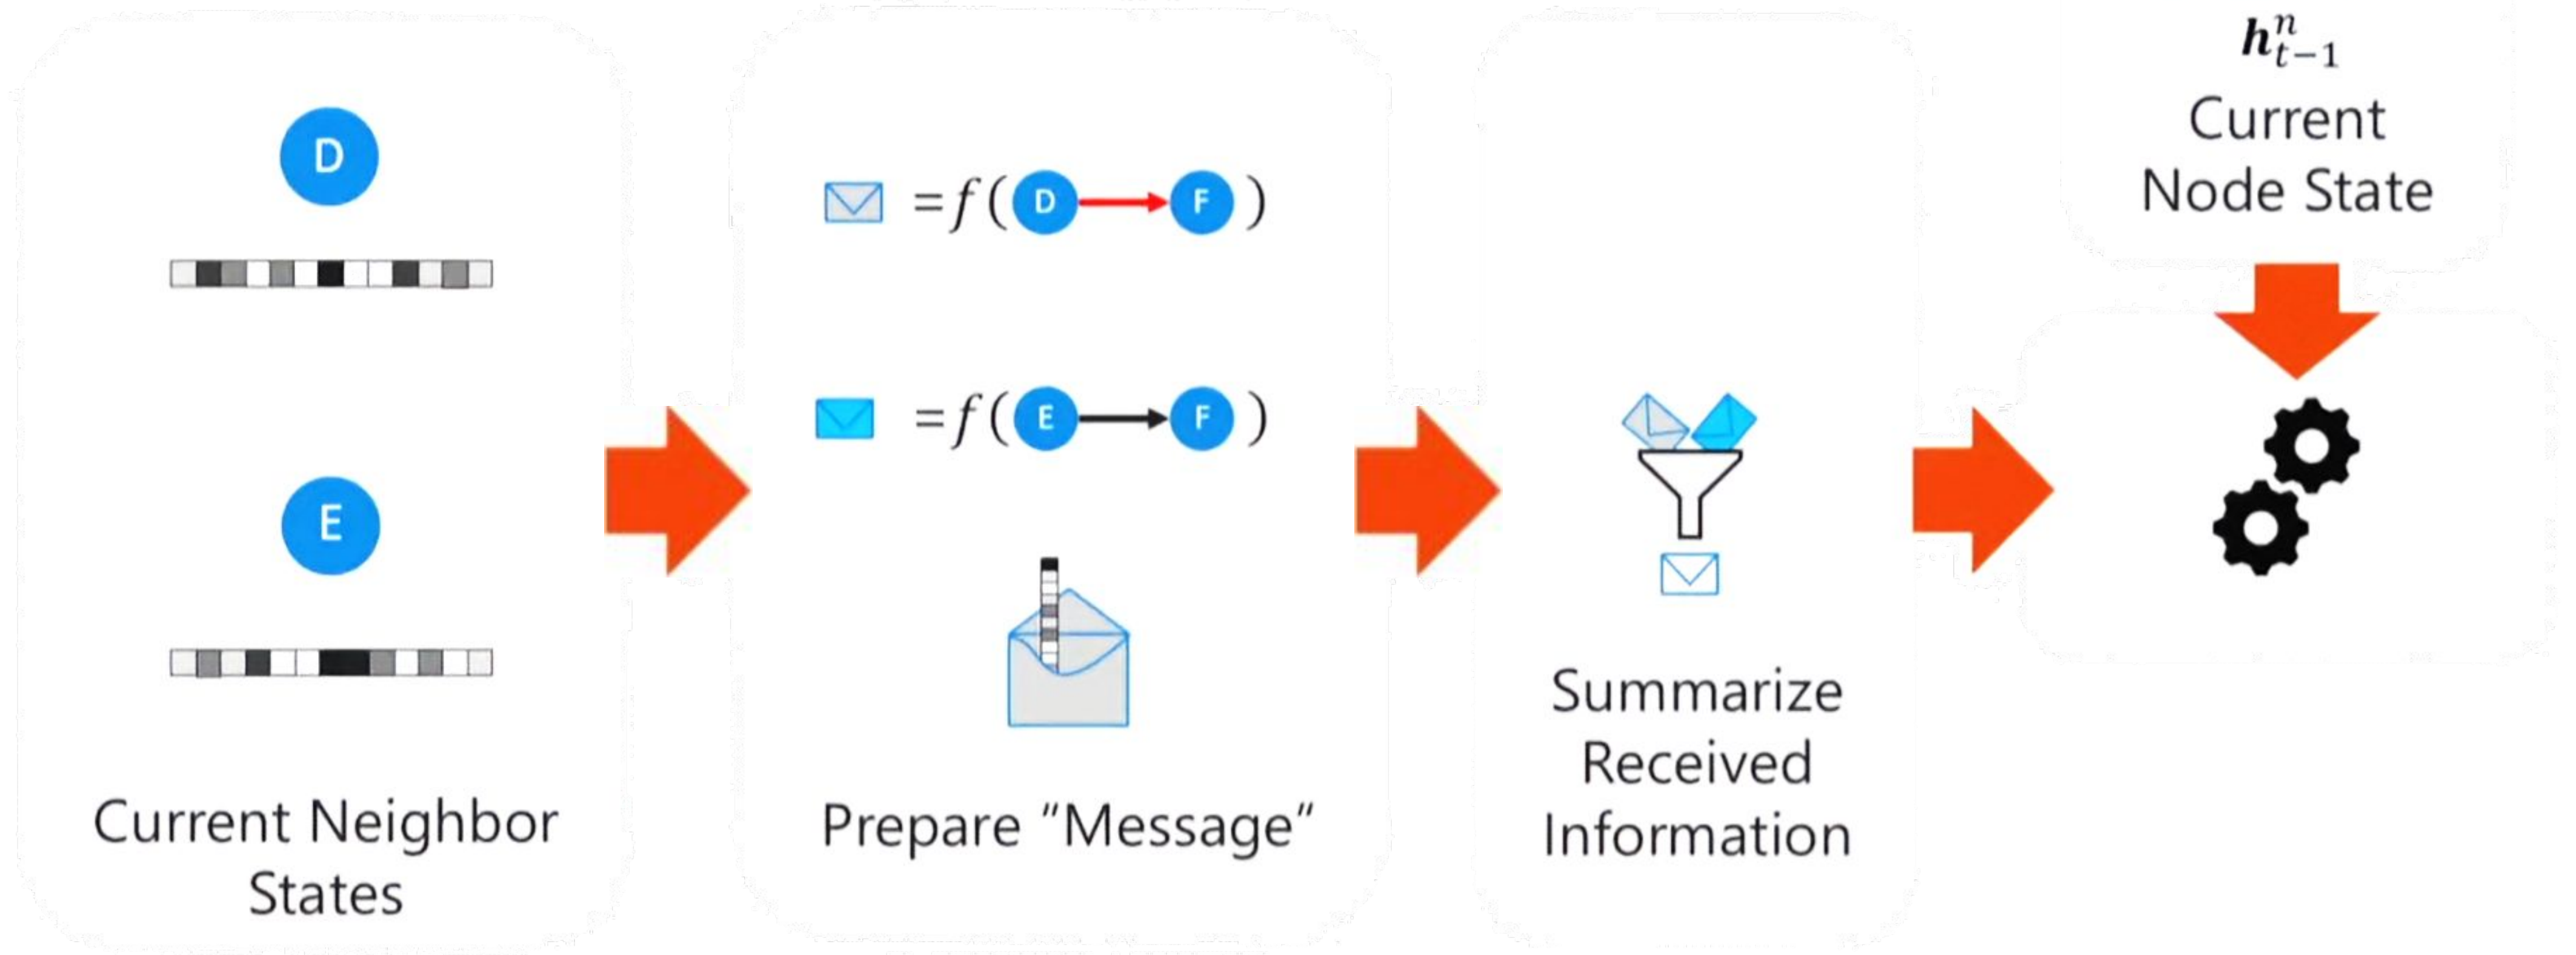
Prepare "Message"




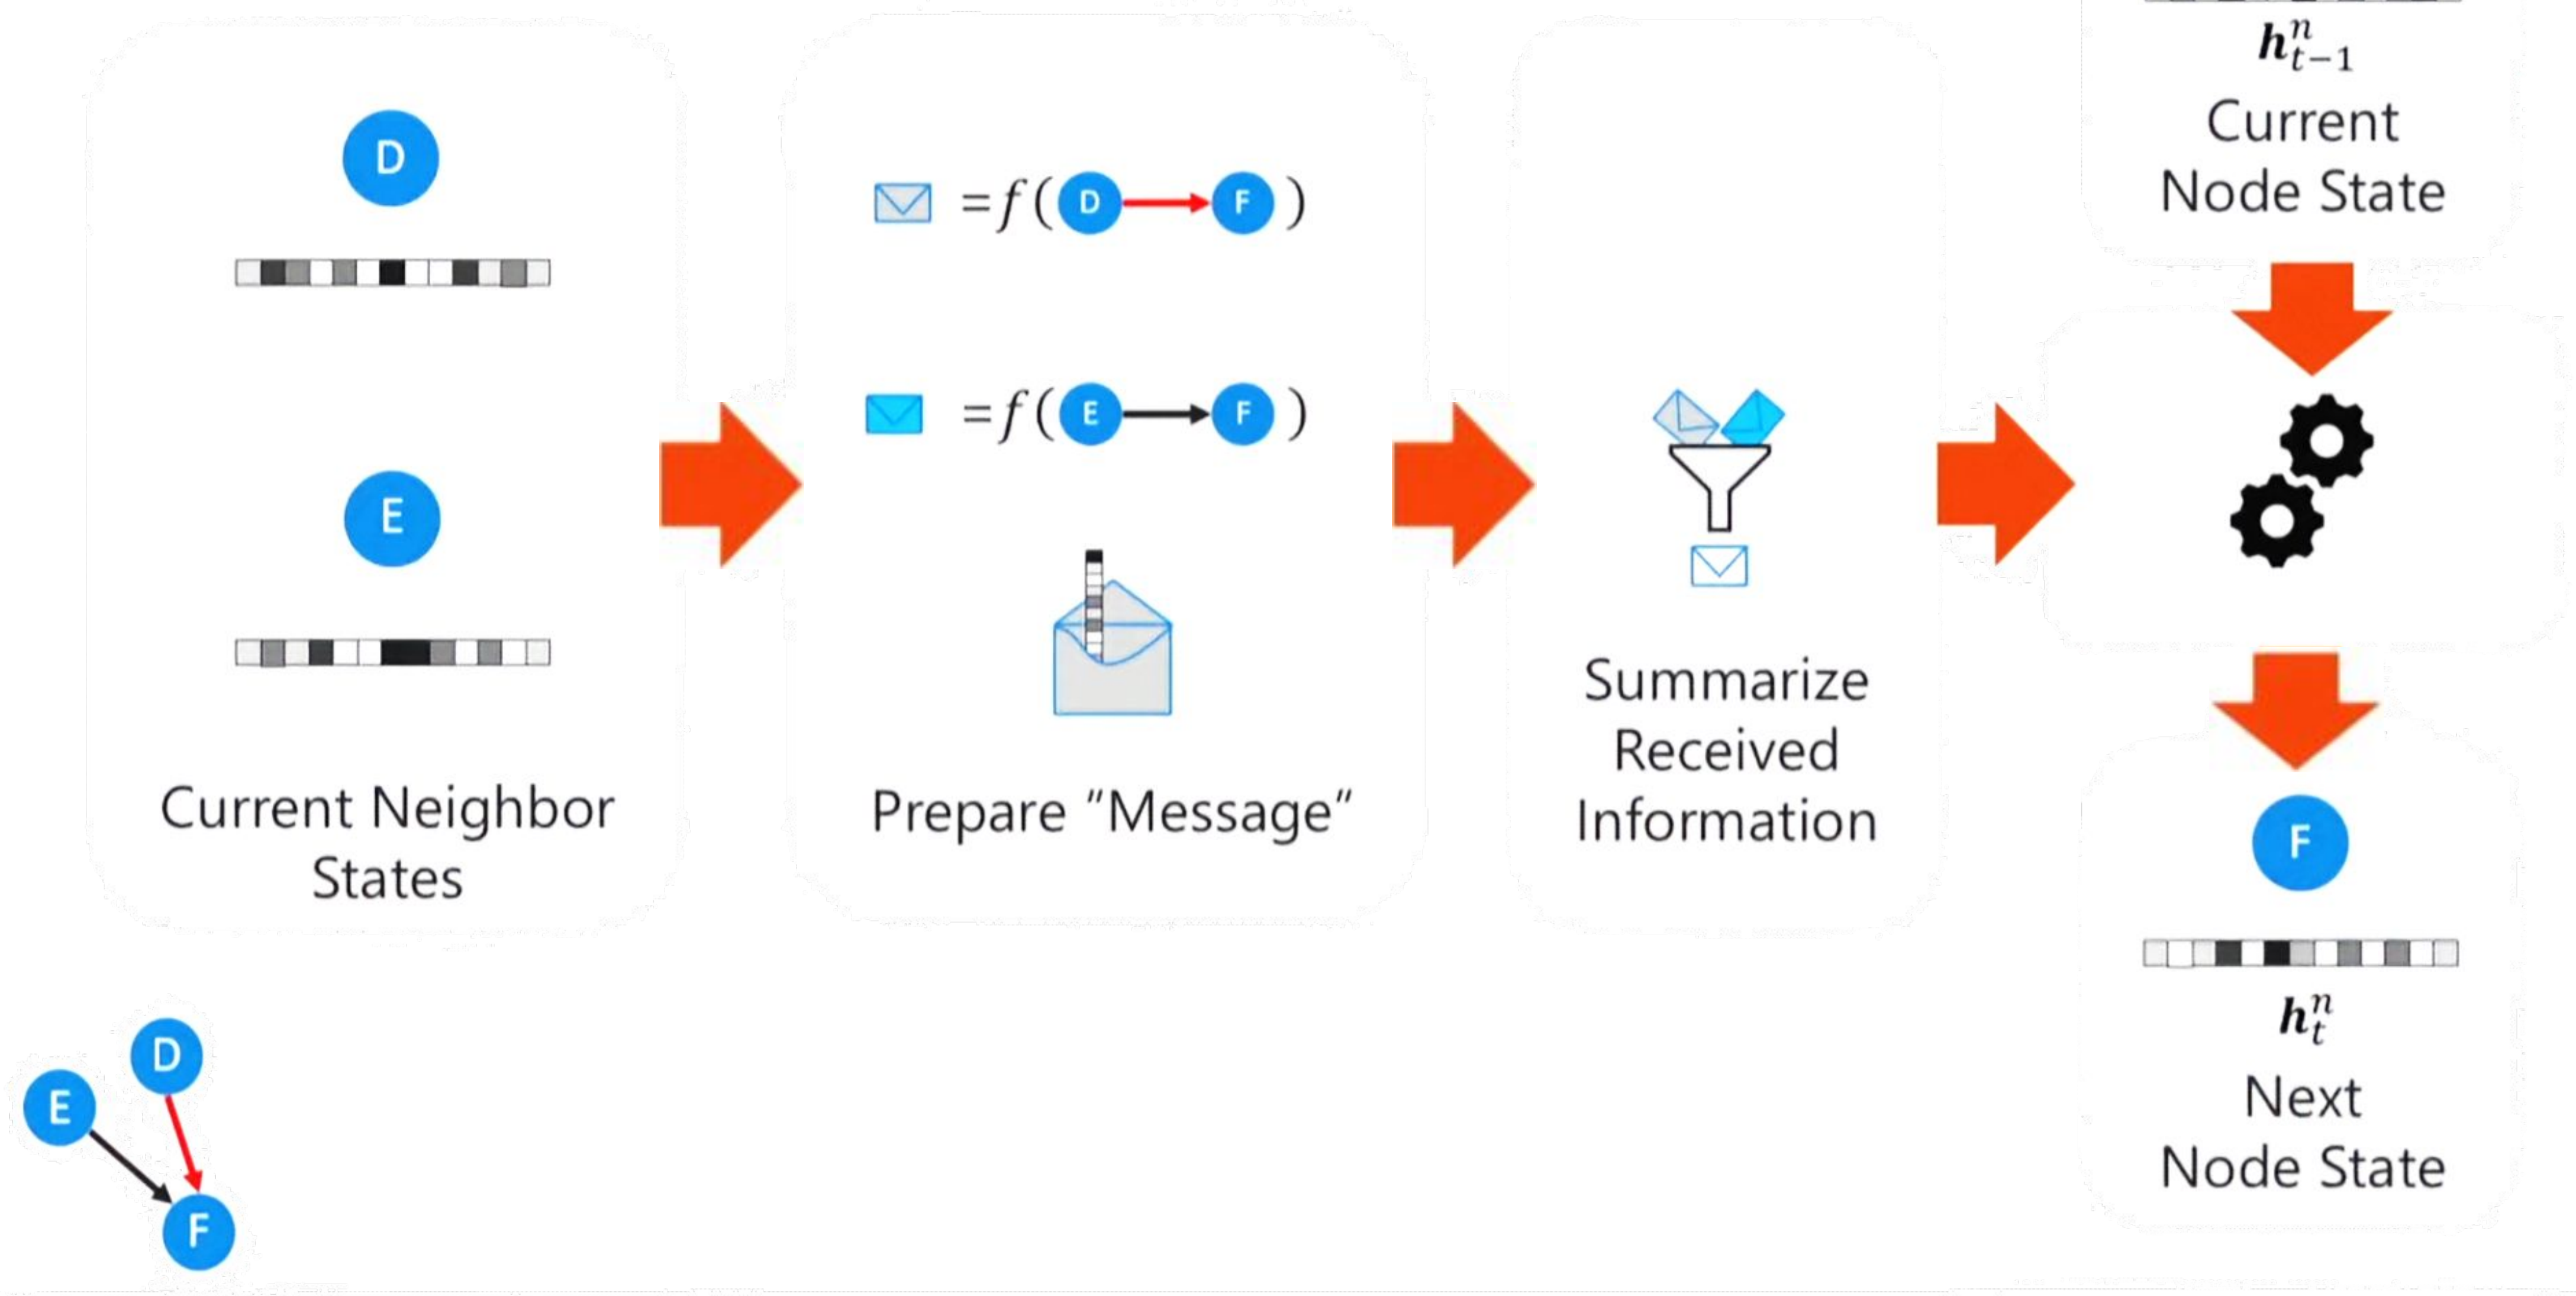
Neural Message Passing



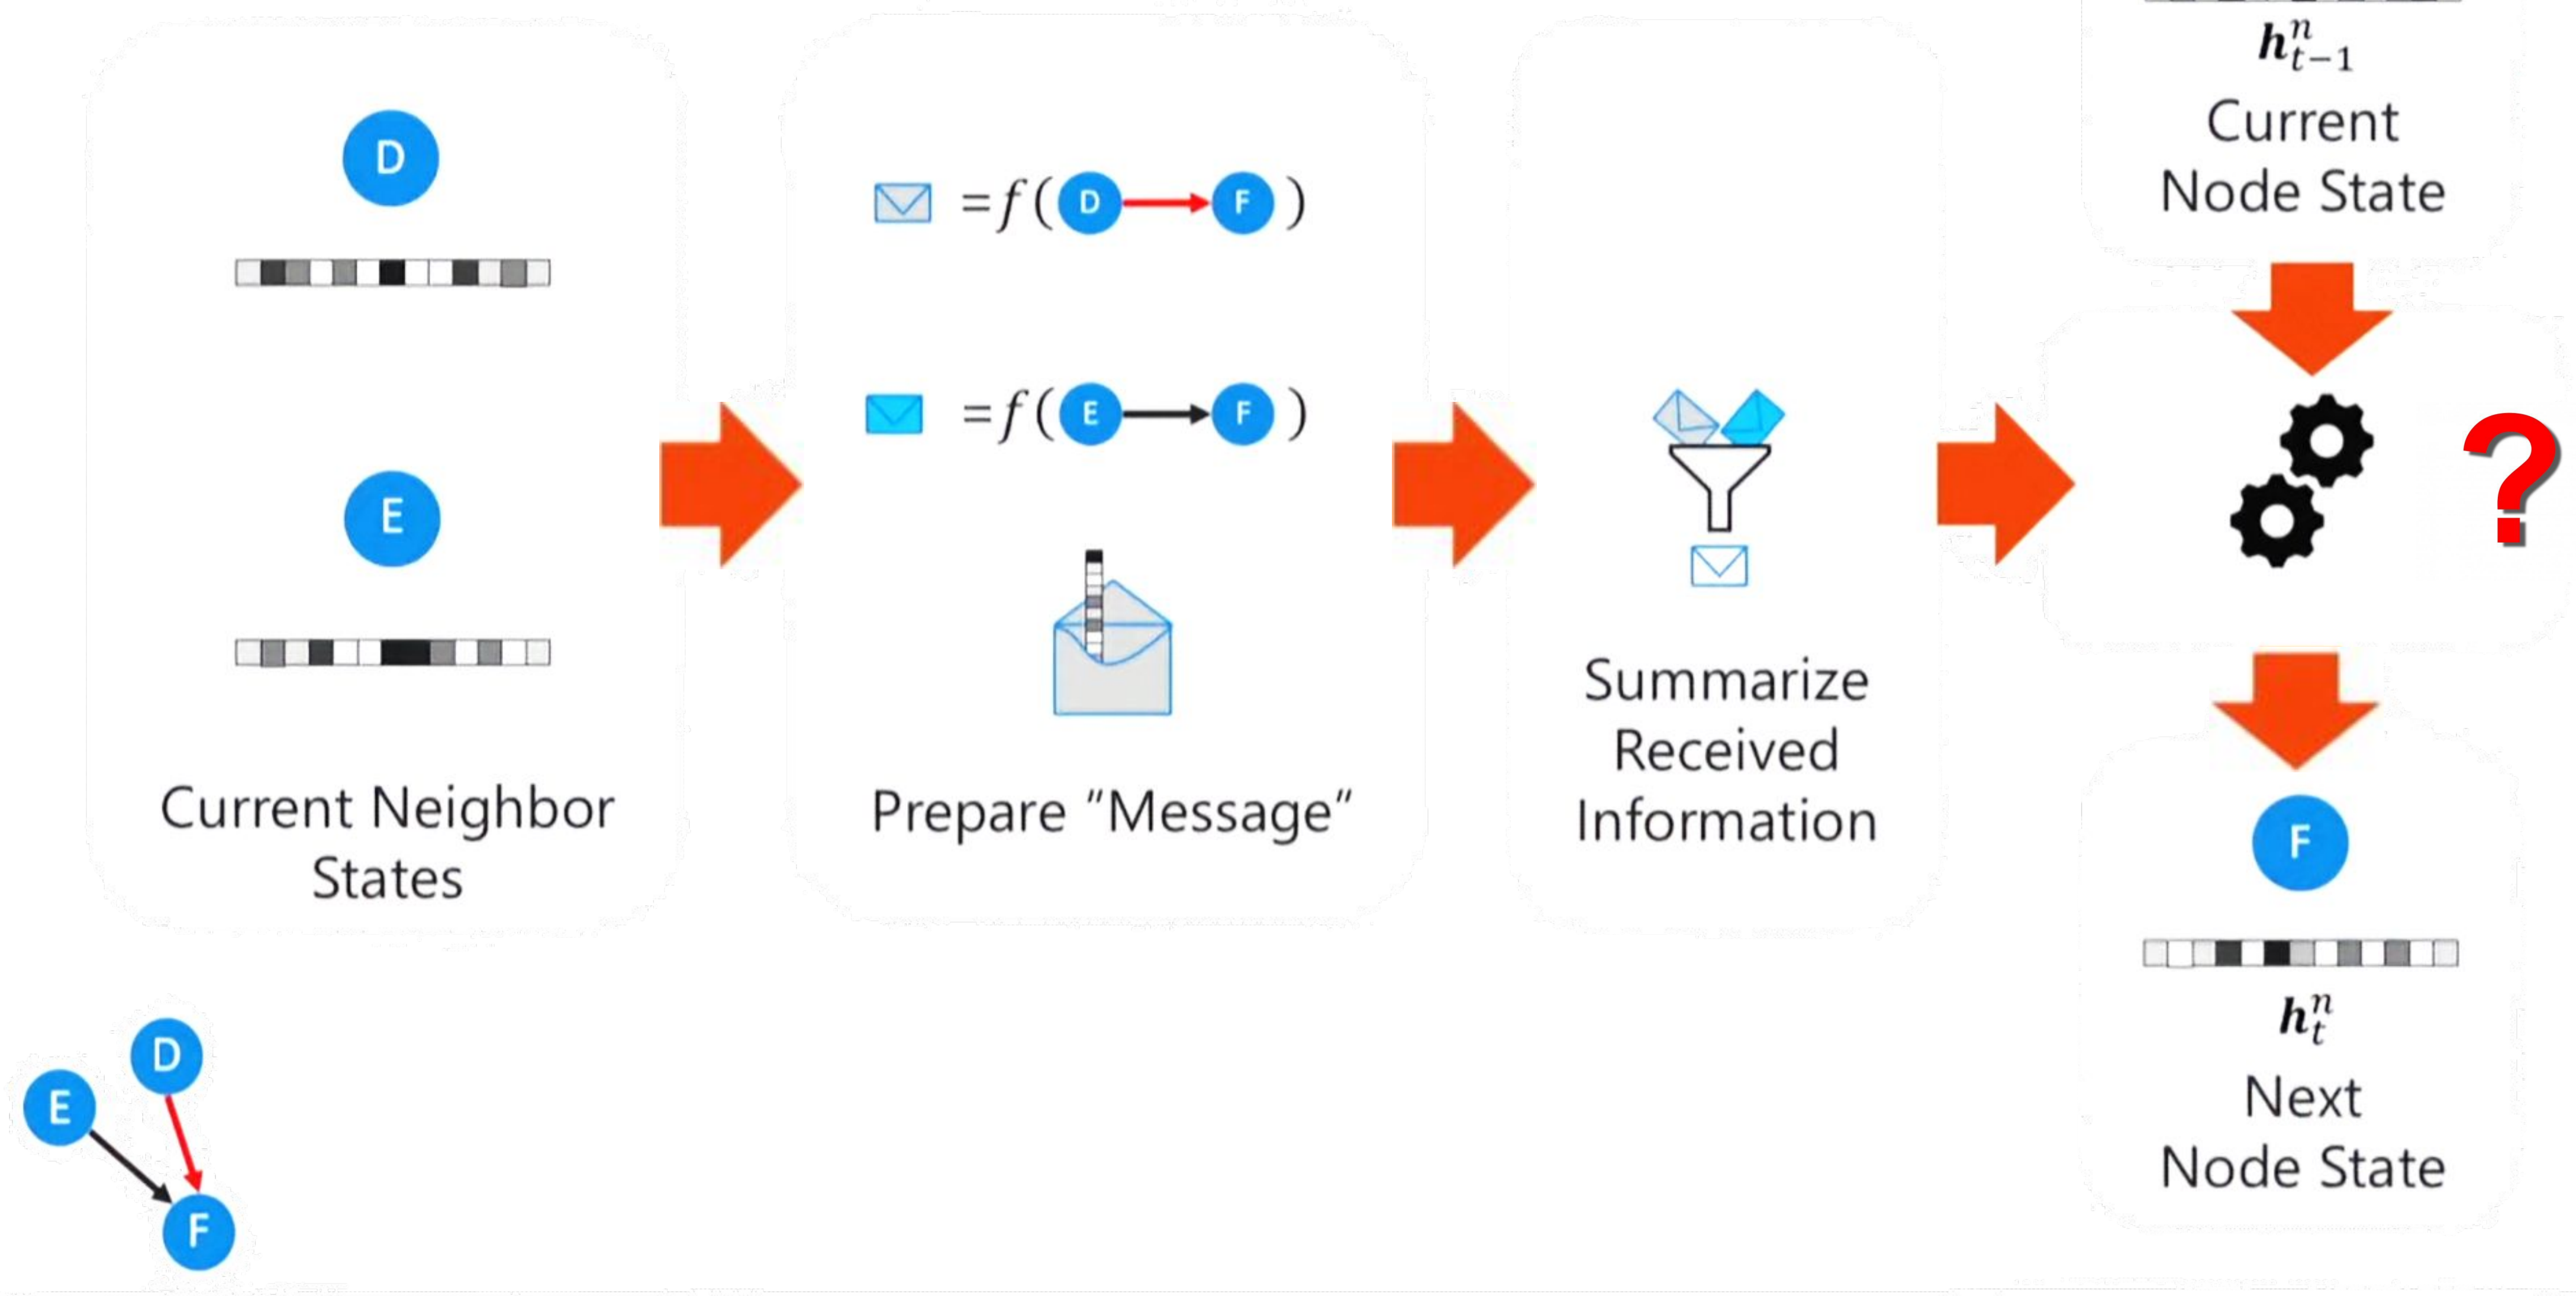
Neural Message Passing

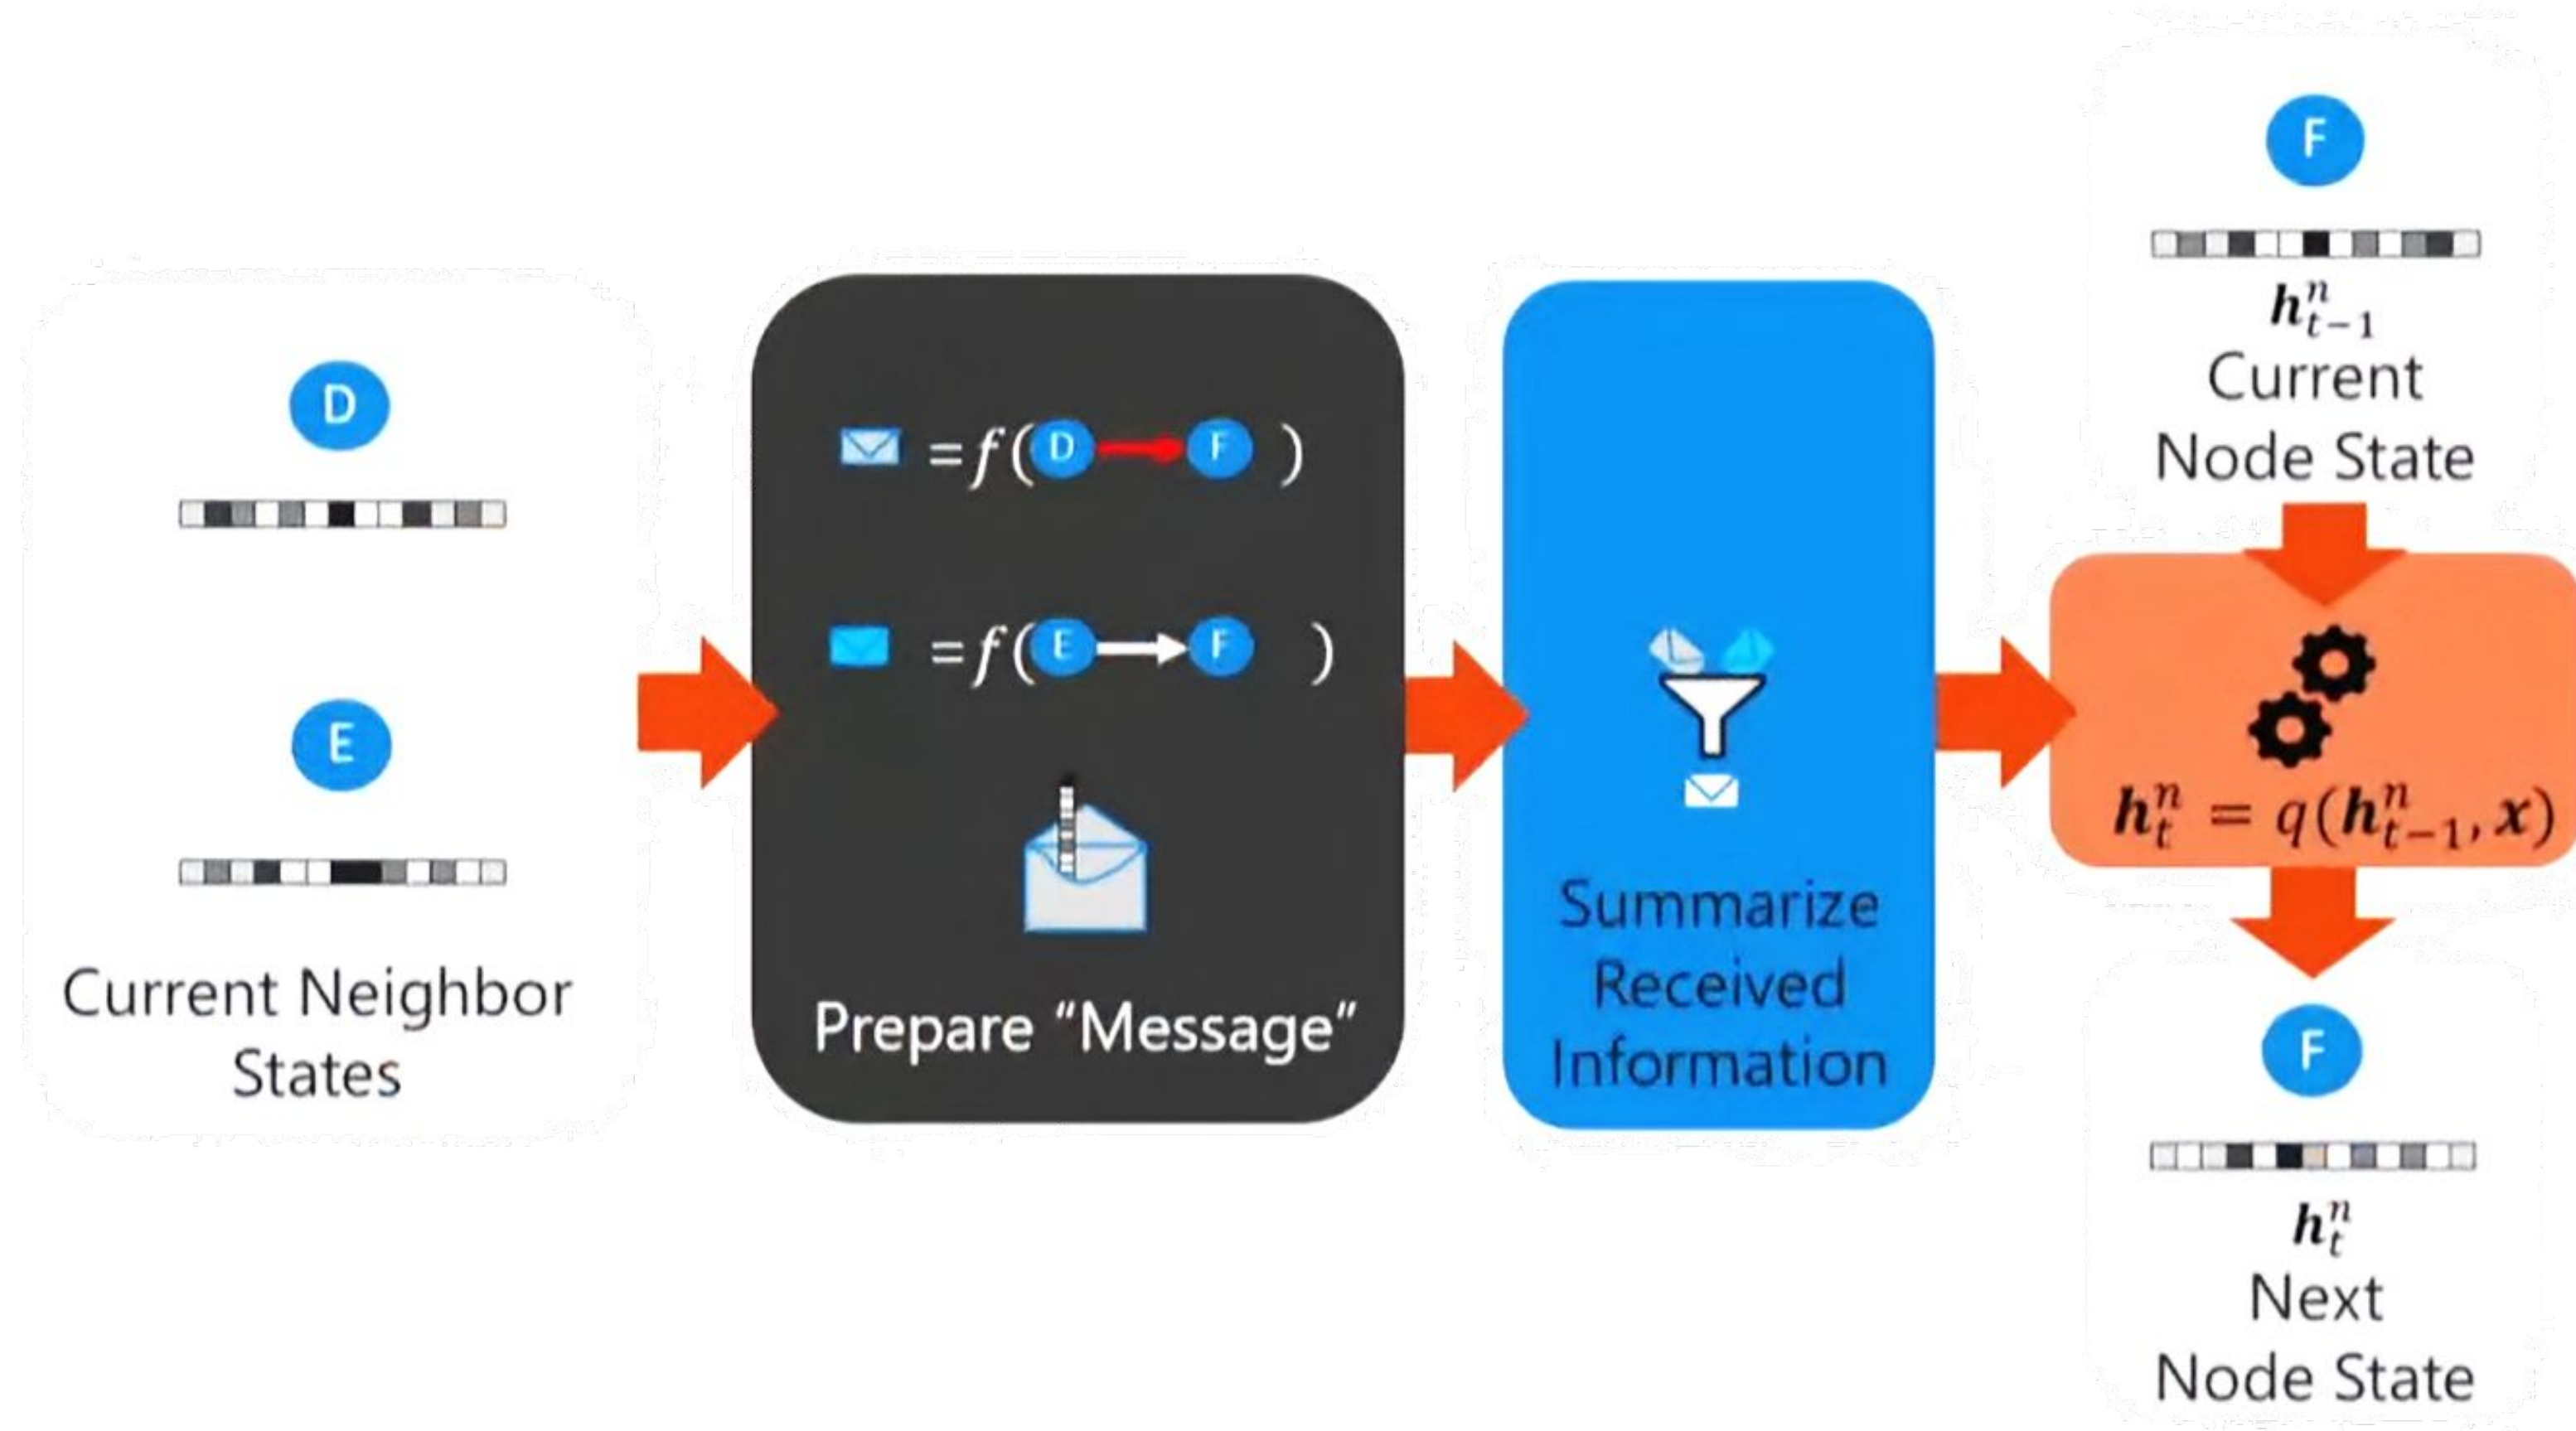


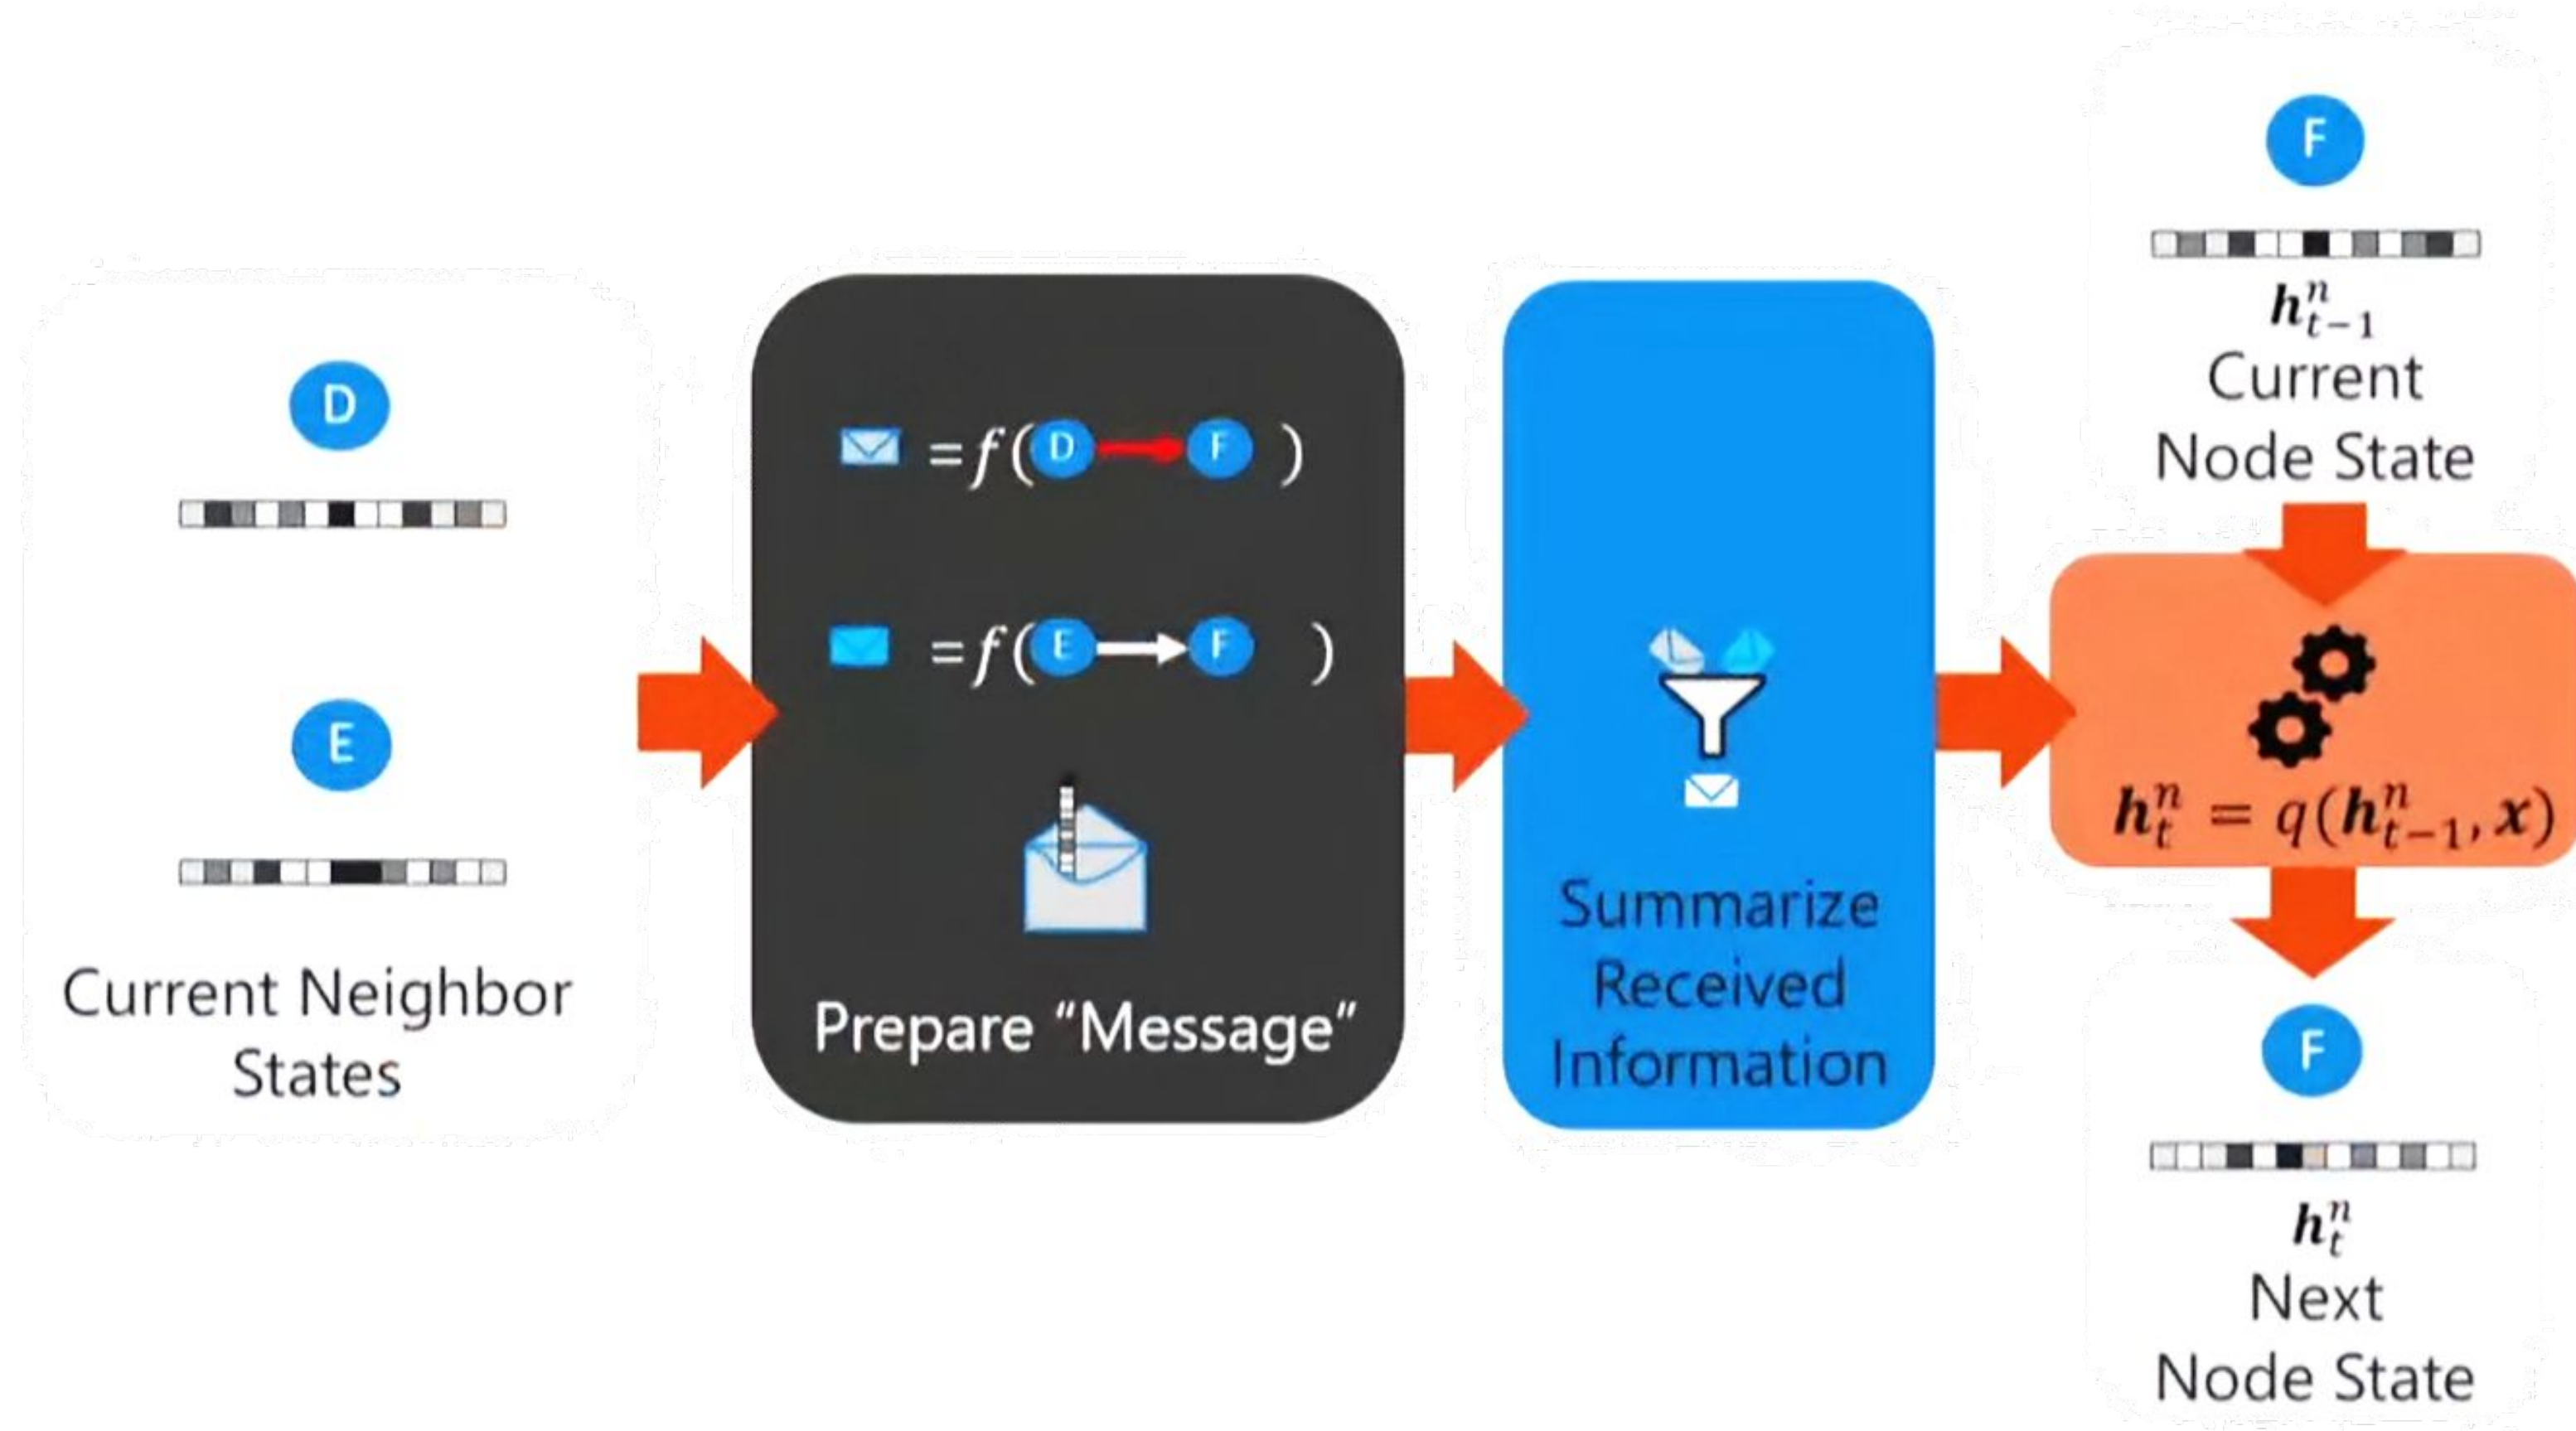
Neural Message Passing



Neural Message Passing







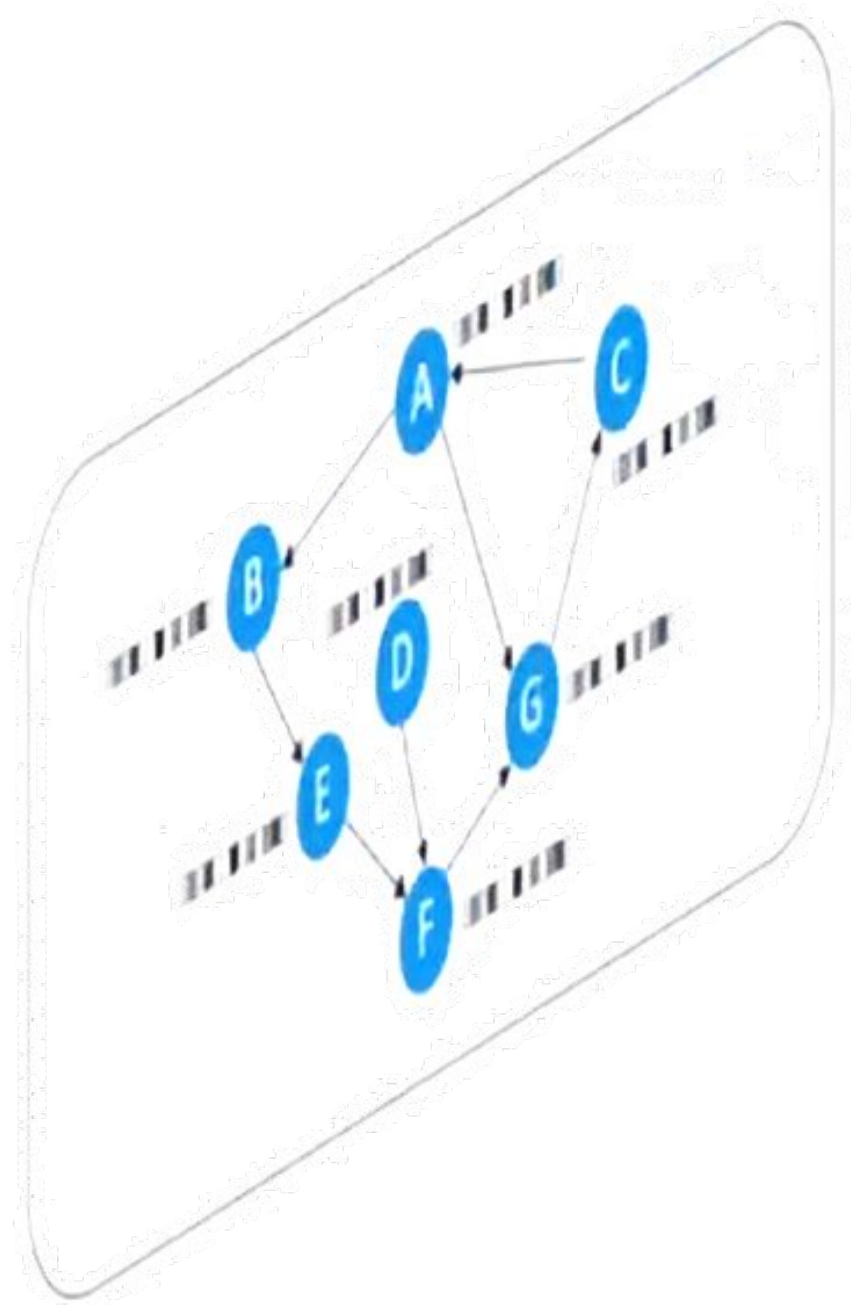


$$h_t^n = q \left(h_{t-1}^n, \bigcup_{\forall n_j: n \rightarrow n_j}^k f_t \left(h_{t-1}^n, k, h_{t-1}^{n_j} \right) \right)$$






Graph Neural Networks: Message Passing

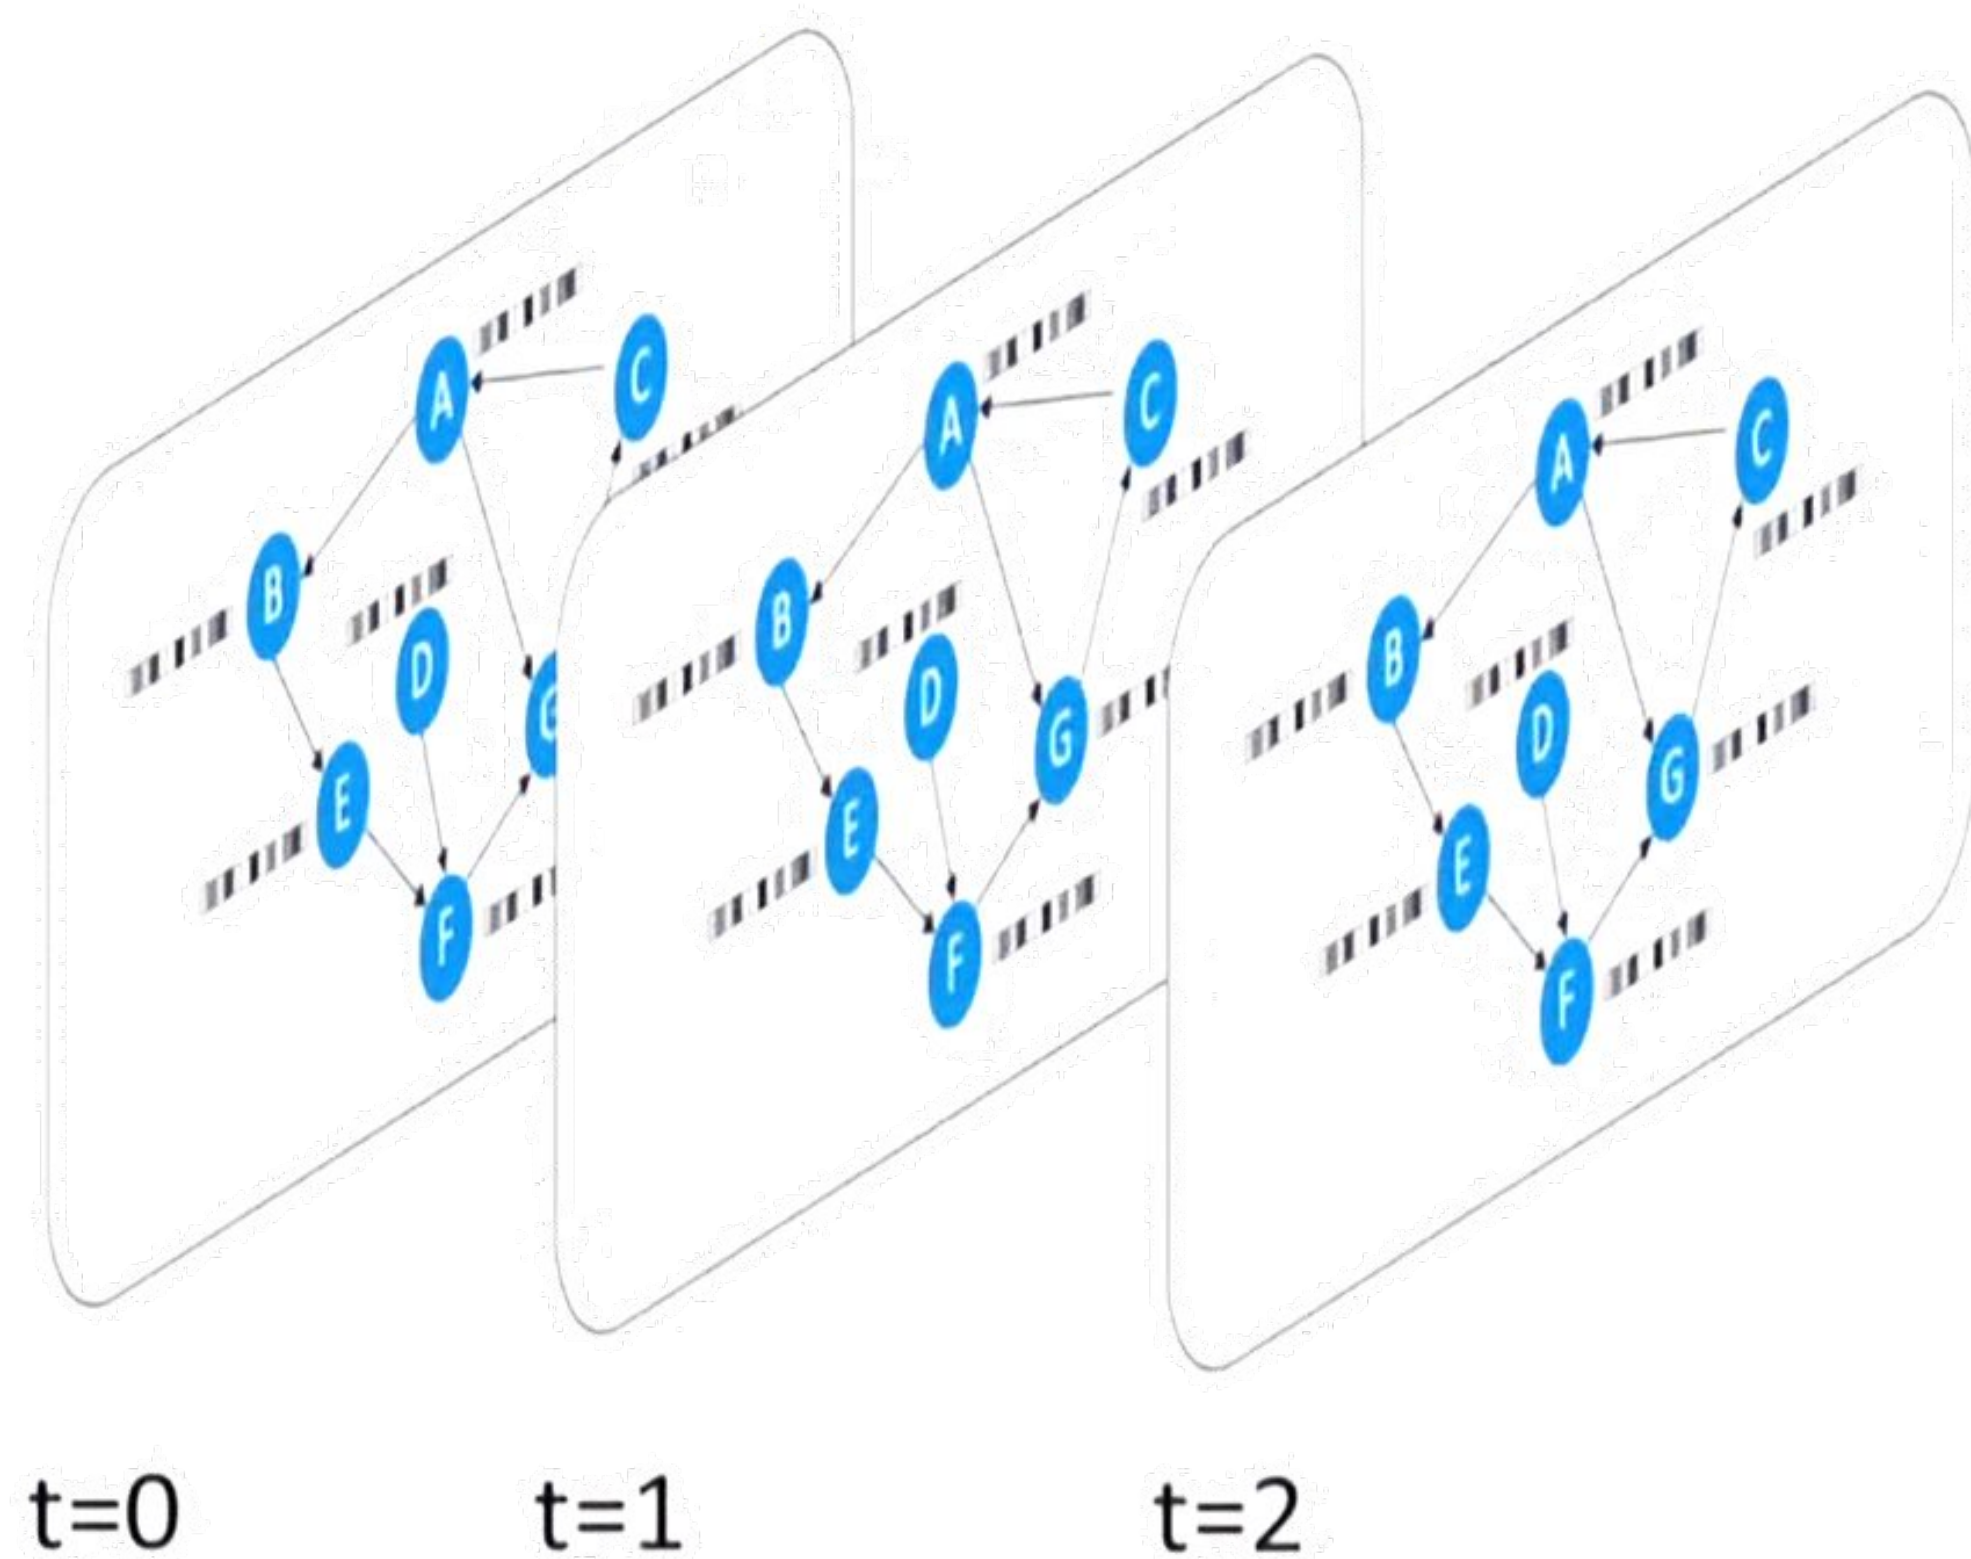


t=0

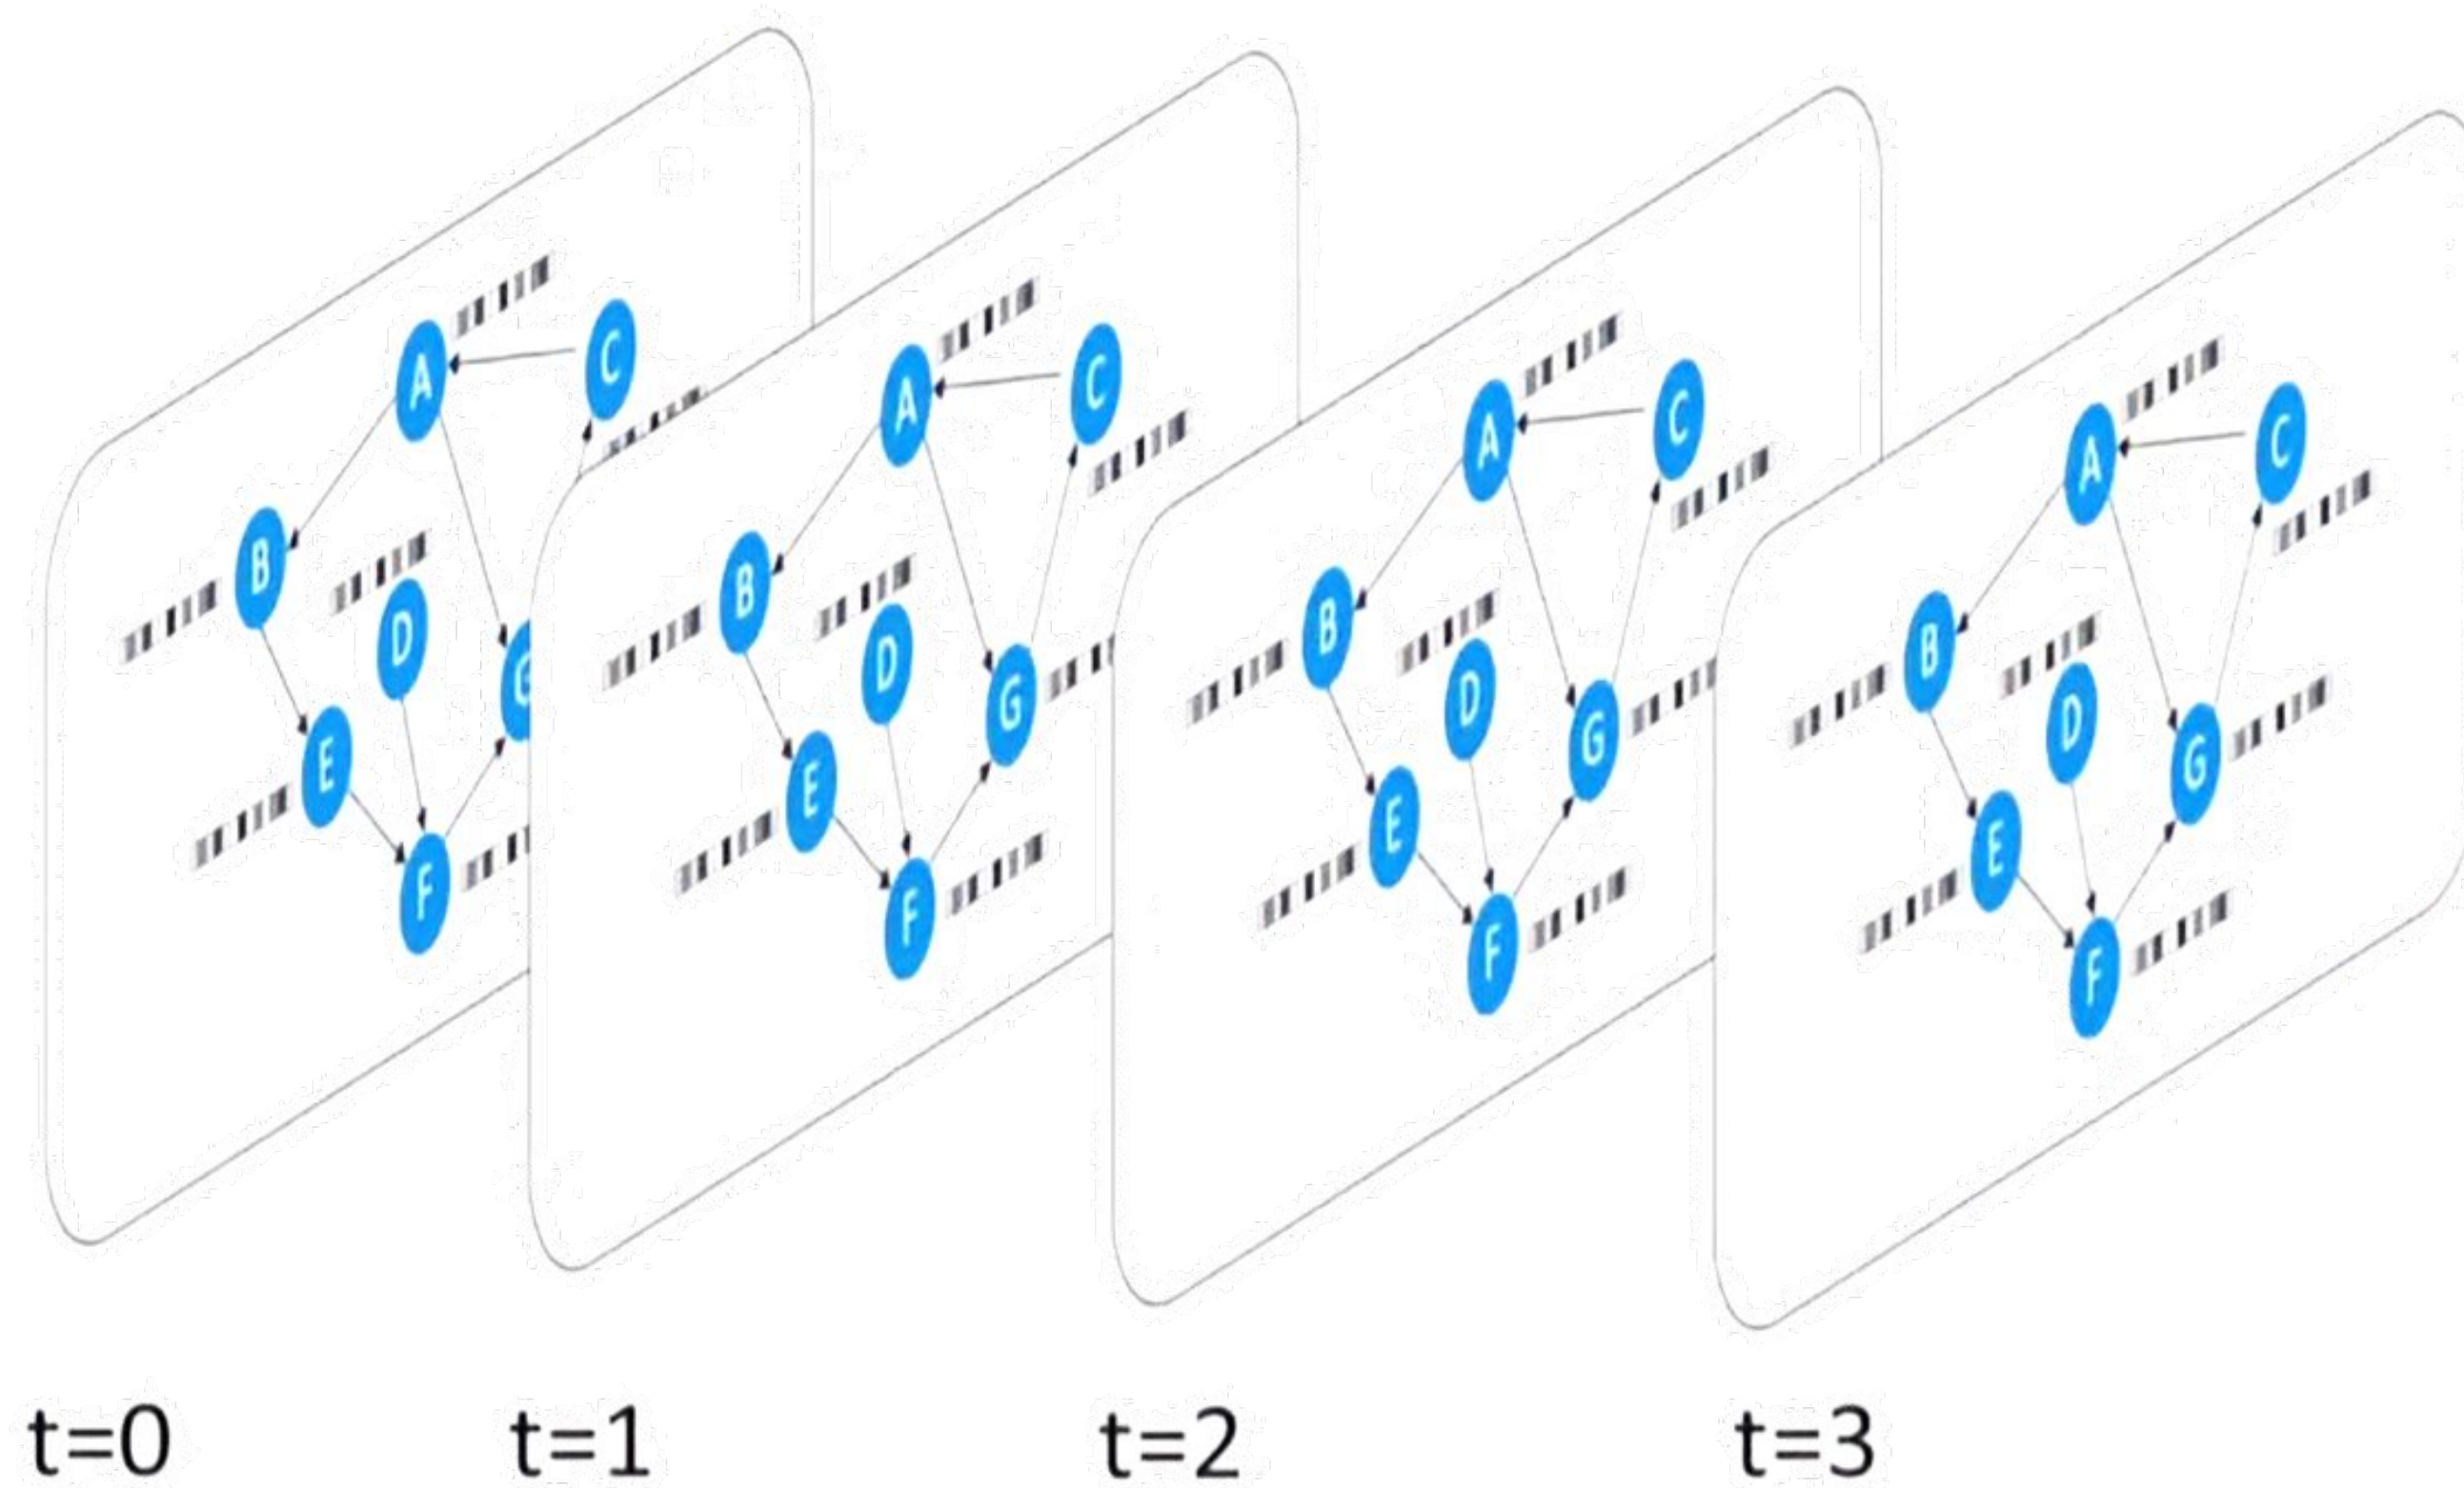
Graph Neural Networks: Message Passing




Graph Neural Networks: Message Passing




Graph Neural Networks: Message Passing




Graph Neural Networks: Message Passing



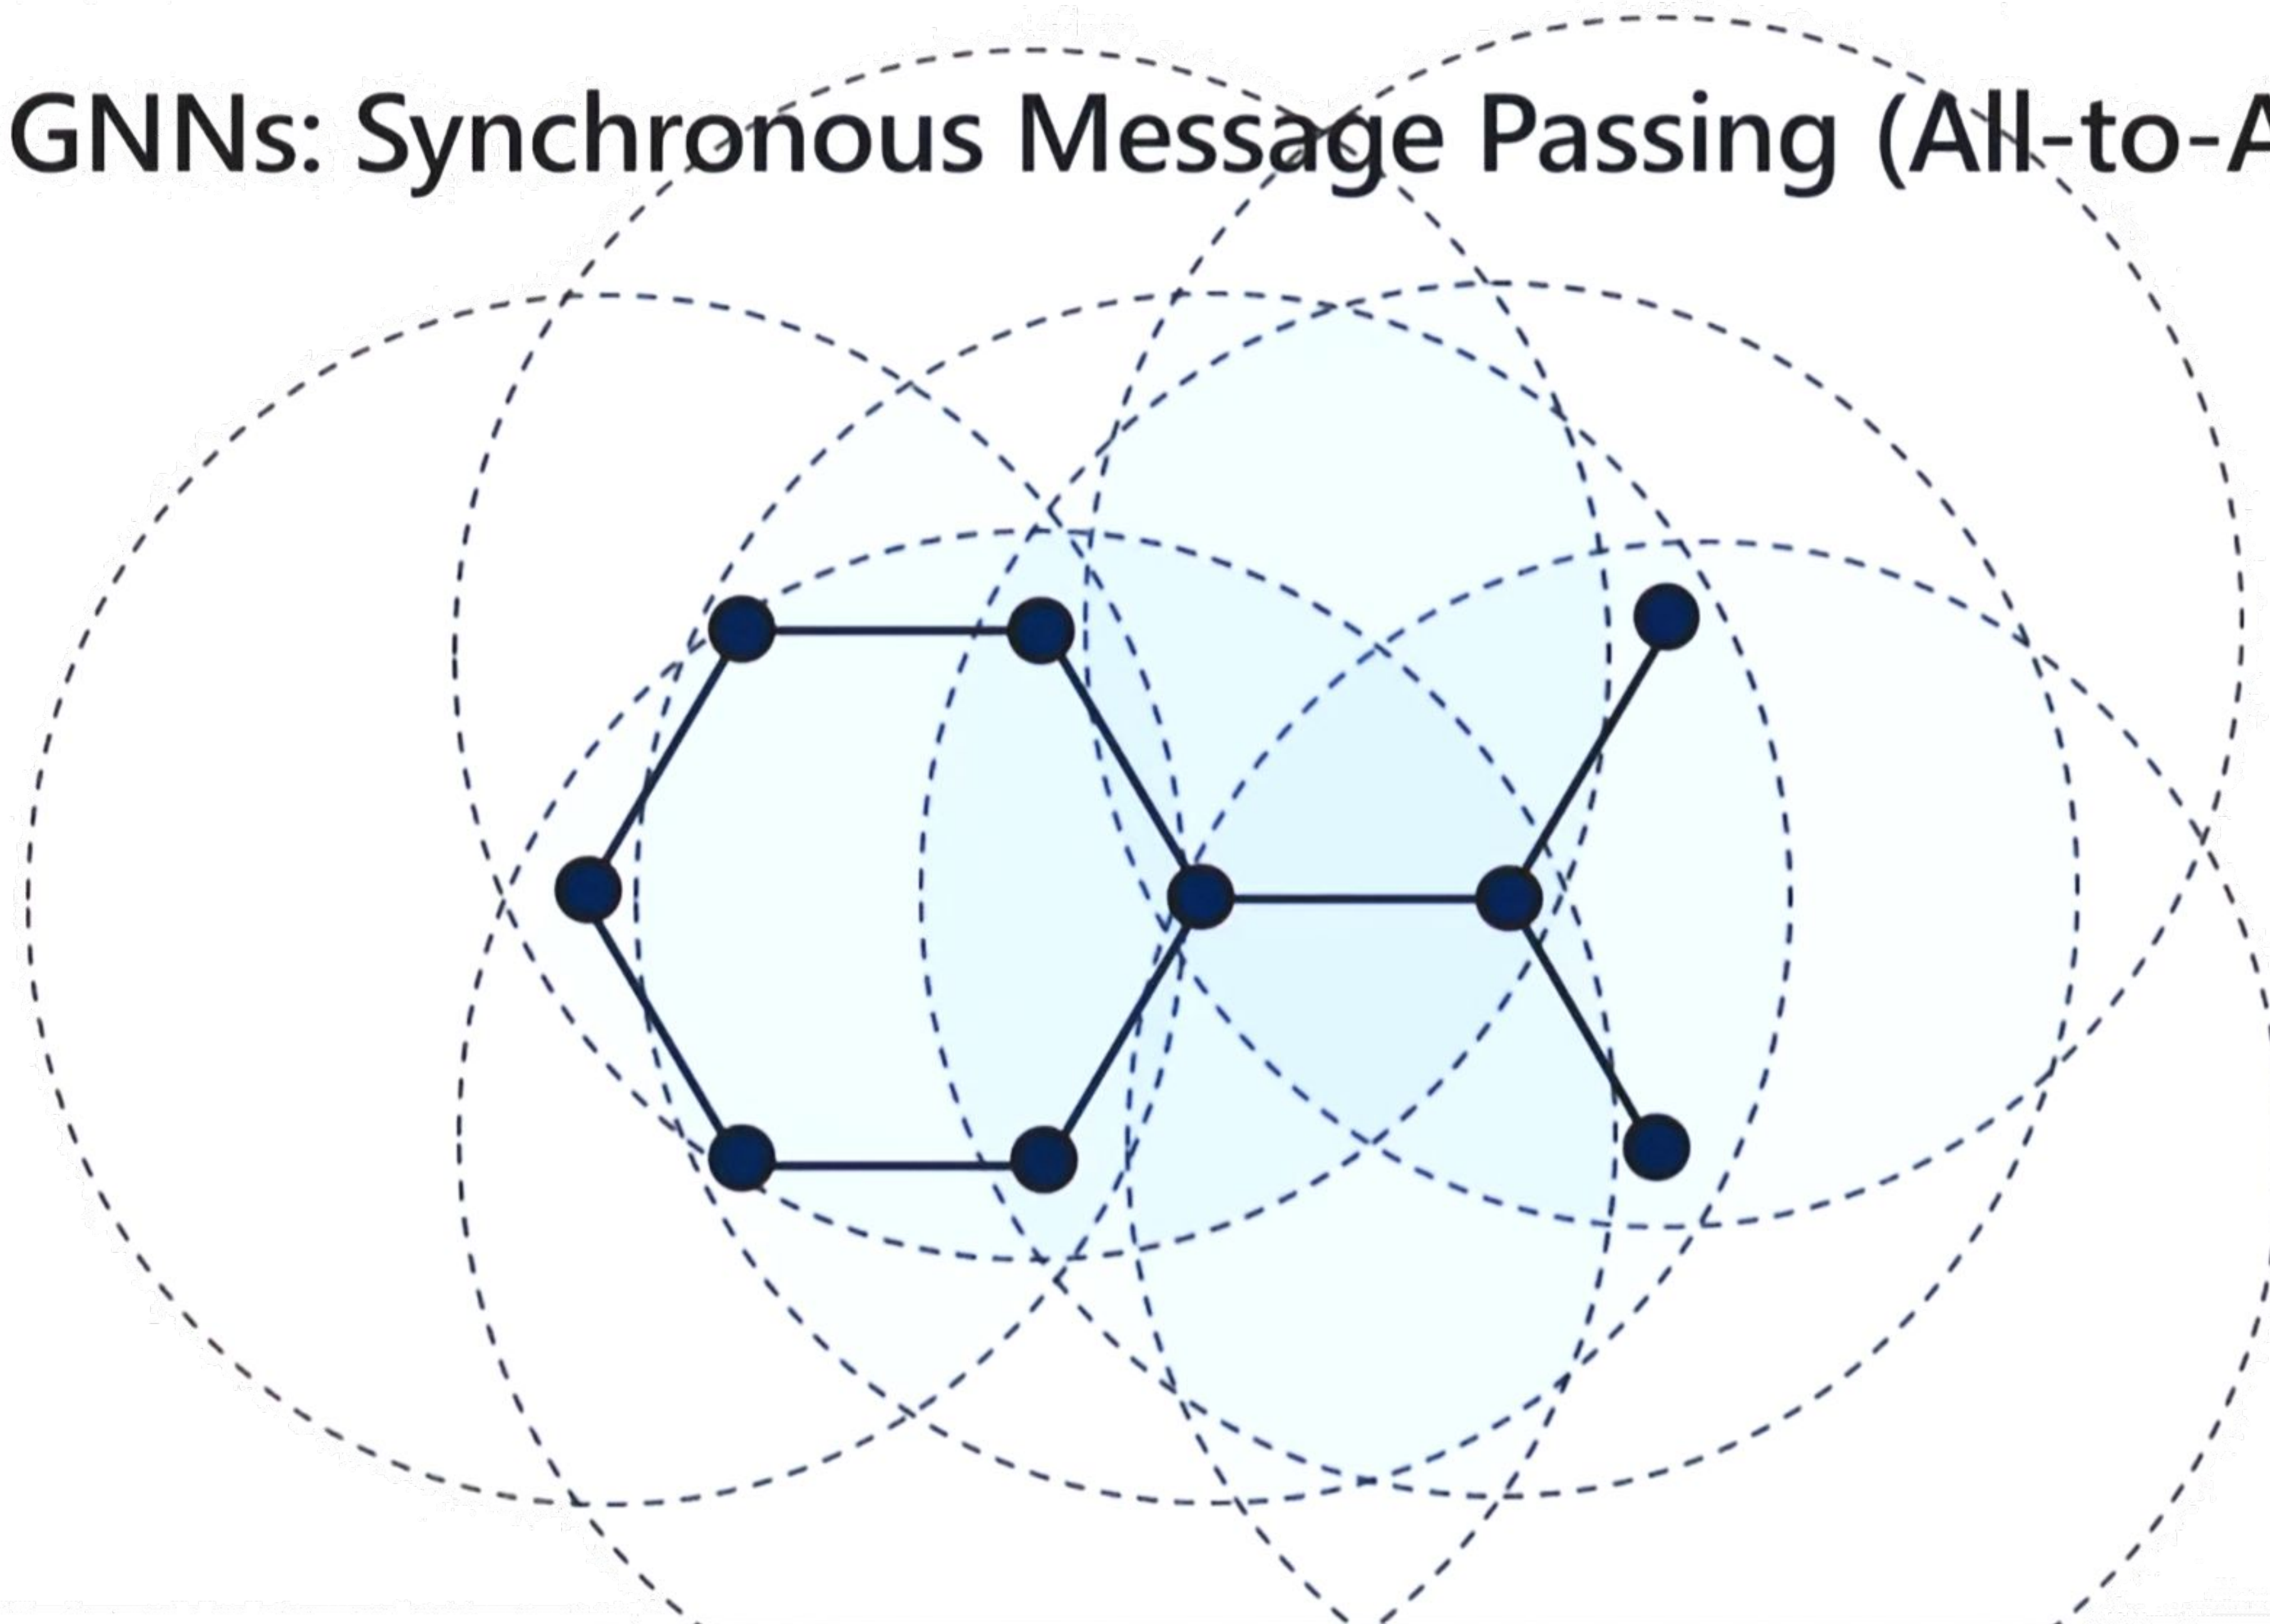
GNNs: Synchronous Message Passing (All-to-All)



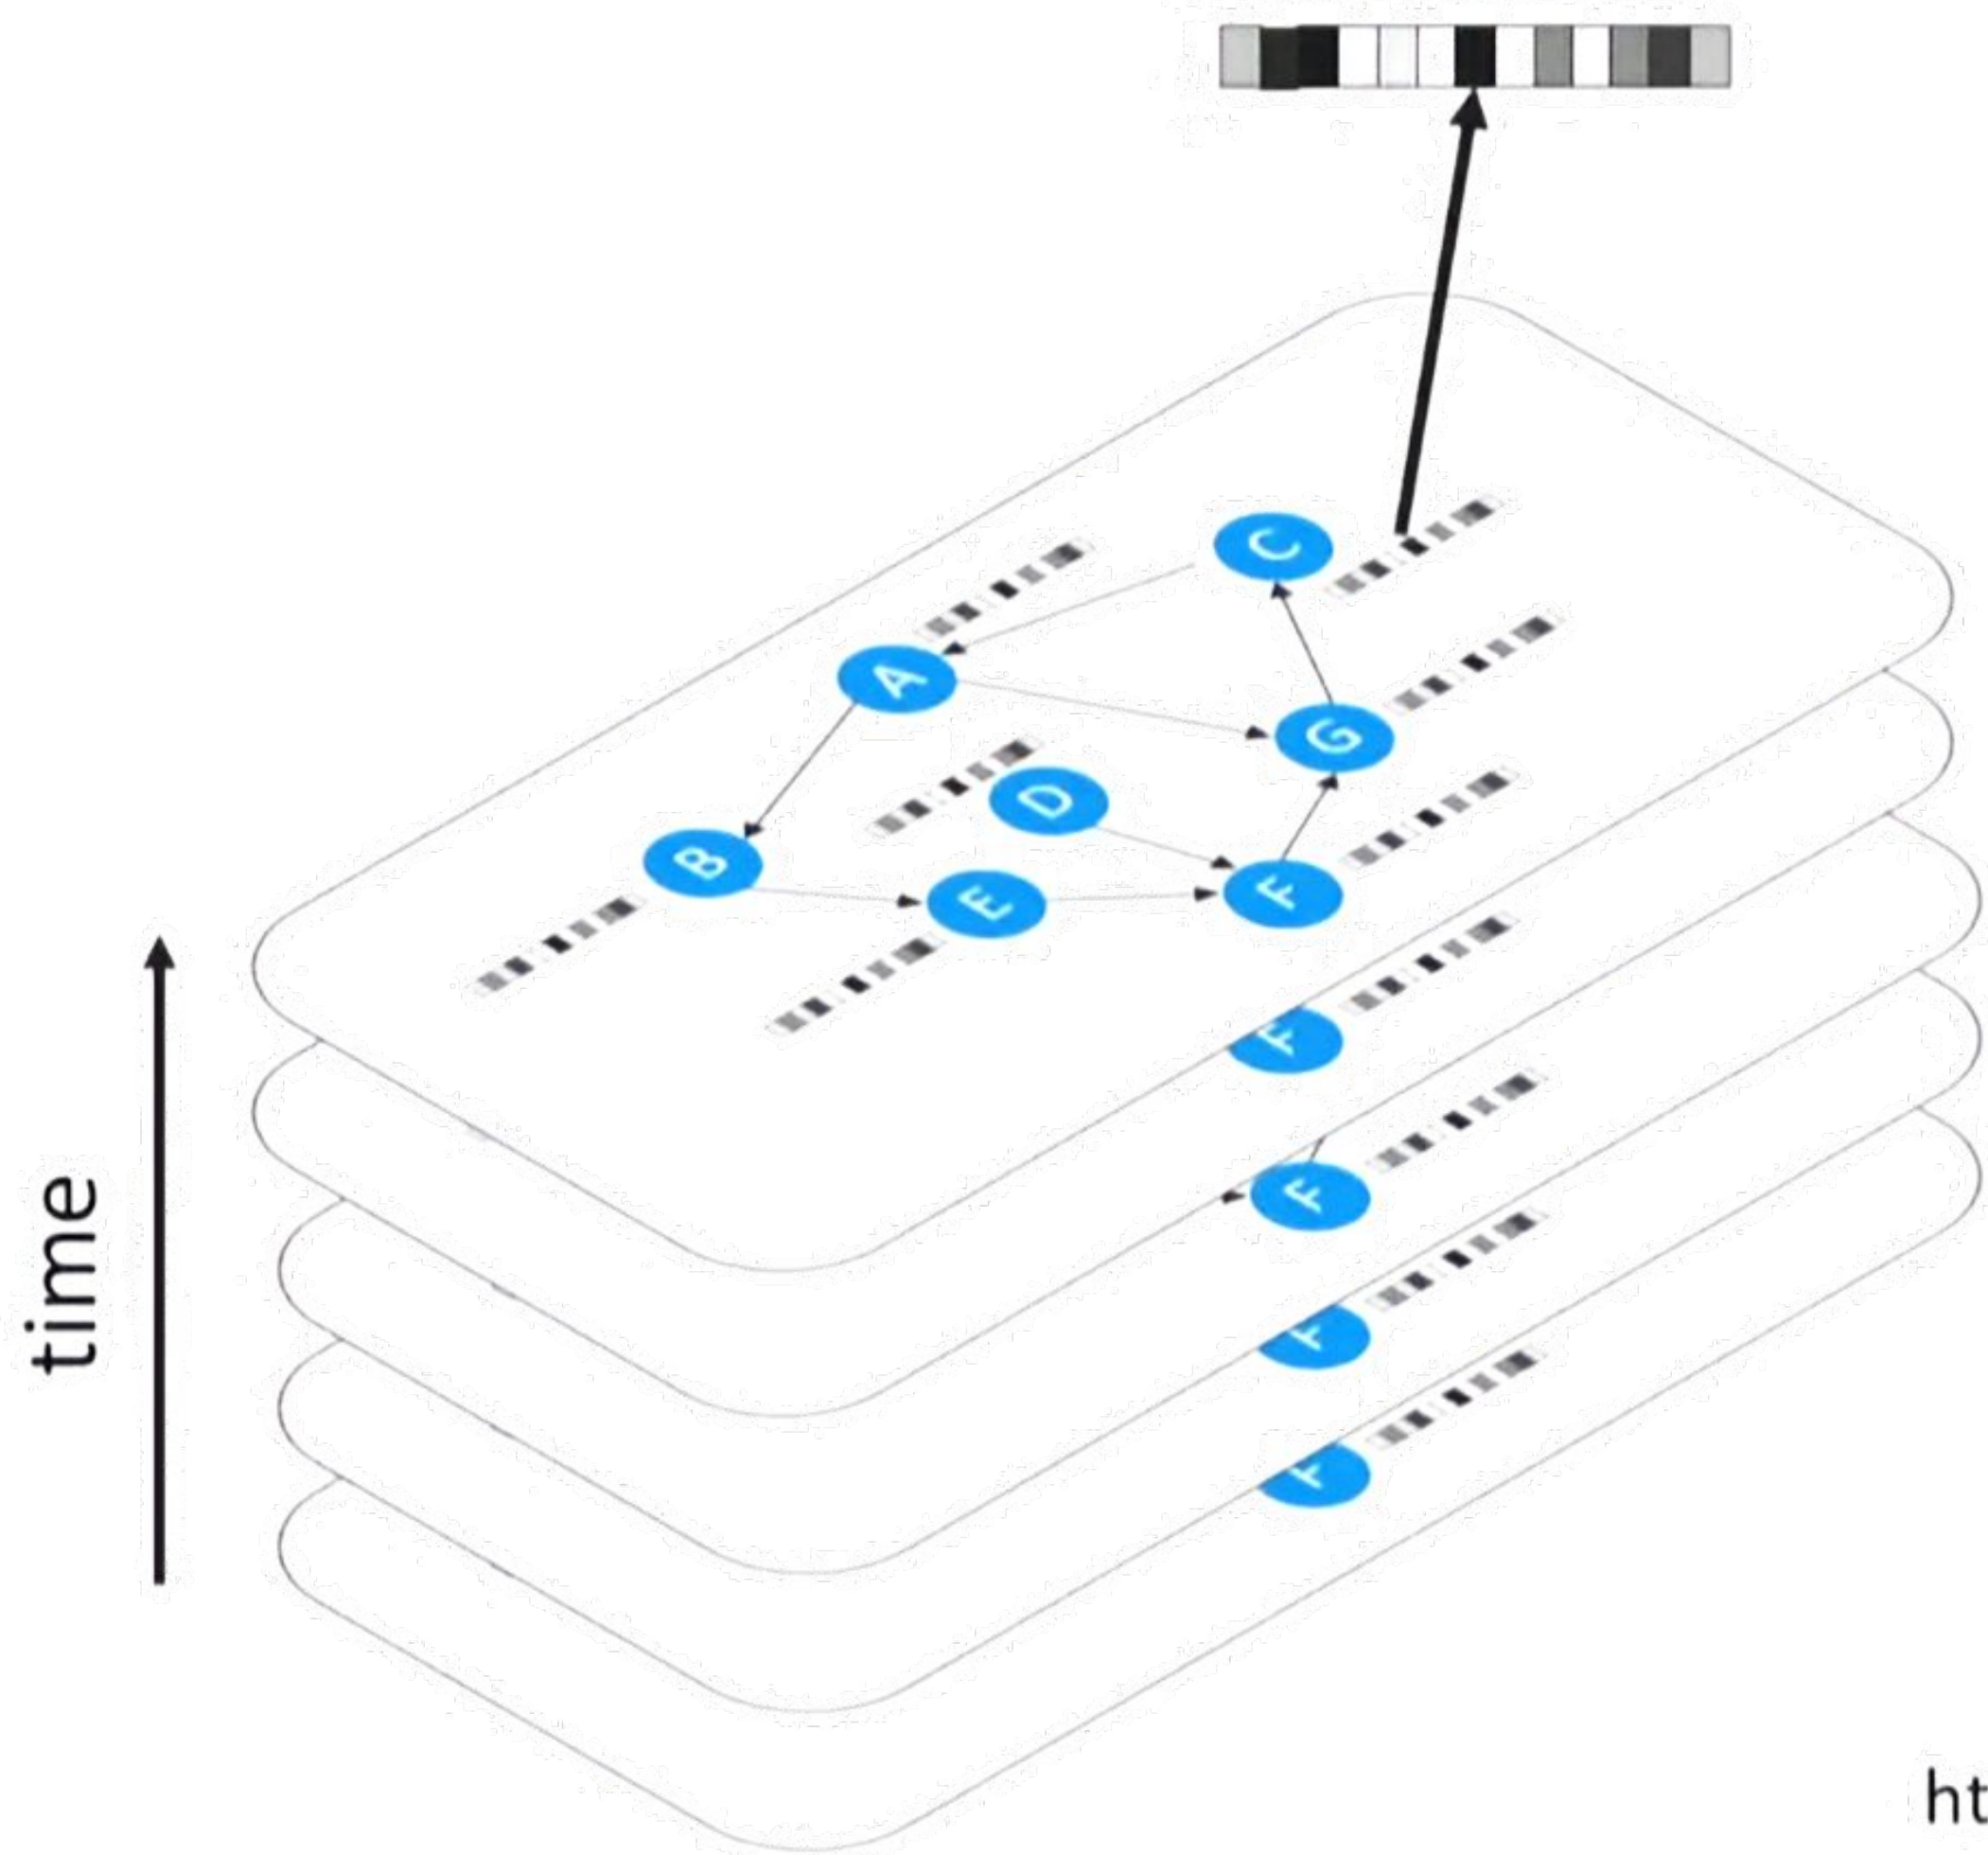
GNNs: Synchronous Message Passing (All-to-All)



GNNs: Synchronous Message Passing (All-to-All)



Graph Neural Networks: Output



- node selection
- node classification
- graph classification

<https://github.com/microsoft/tf-gnn-samples/>





Perception

The perceptual capabilities of robots encompass a variety of skills:

- State estimation
- Segmentation
- Tracking
- Recognition
- Classification



Perception

The perceptual capabilities of robots encompass a variety of skills:

- State estimation
- Segmentation
- Tracking
- Recognition
- Classification

How do we get these data?





Properties Perception

Sensors	Property	Action
Camera	Shape	Lifting
LR Tactile	Size	Dragging
HR Tactile	Color	Pulling
Force/Torque	Material	Twisting
Spectrometer	Construction	Flinging
Auditory	Stiffness	Sliding
	Elasticity	Pressing
	Weight	
	Friction	






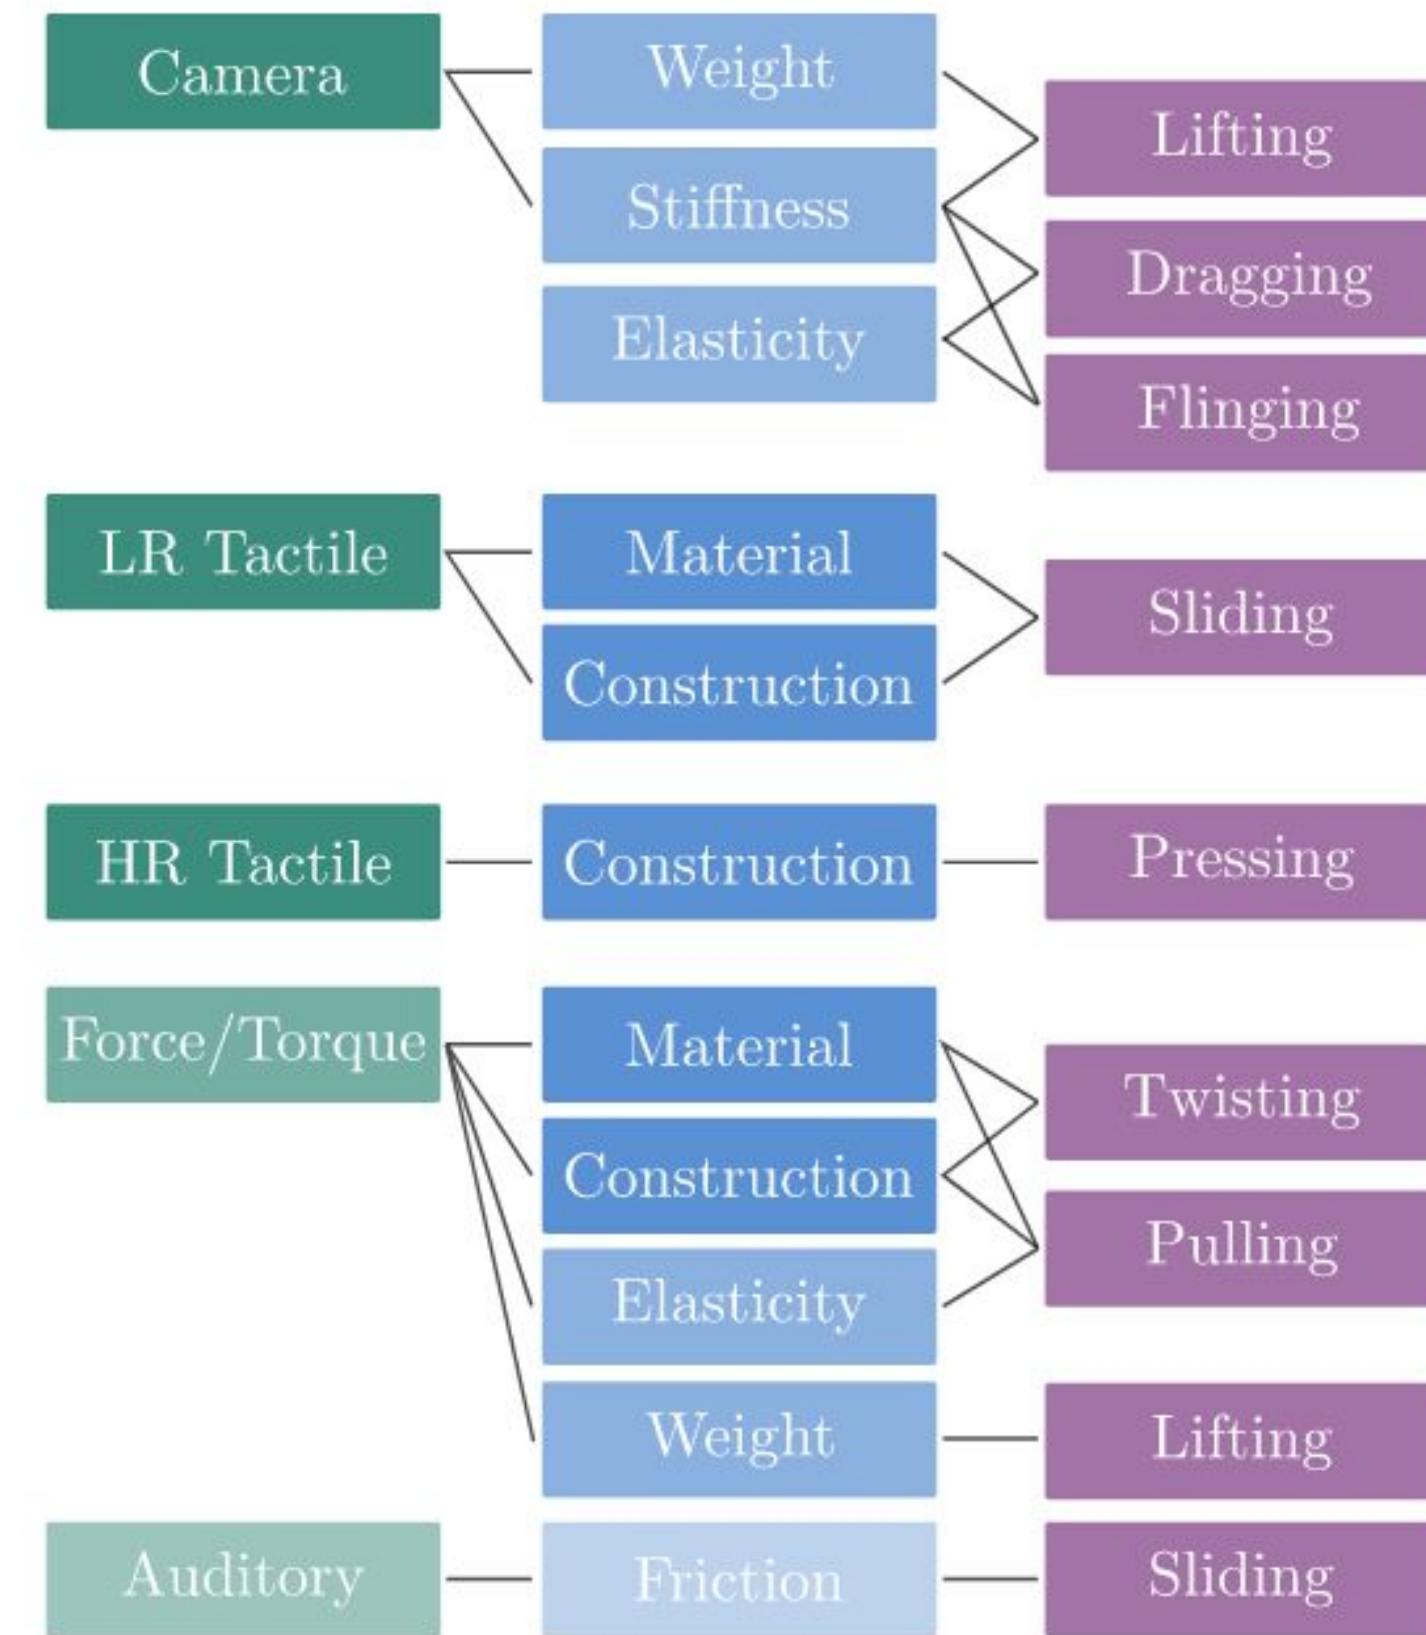
Properties Perception

Sensors	Property	Action
Camera	Shape	Lifting
LR Tactile	Size	Dragging
HR Tactile	Color	Pulling
Force/Torque	Material	Twisting
Spectrometer	Construction	Flinging
Auditory	Stiffness	Sliding
	Elasticity	Pressing
	Weight	
	Friction	

Passive Perception



Interactive Perception



Representation and State Estimation

State estimation of a deformable object x can be seen as

An optimization problem based on observations o and object representation $\mathbf{M}(\cdot)$

$$x^* = \arg \min_x \|o - \mathbf{M}(x)\|$$

$x \in \text{ObjectStates}$



Representation and State Estimation

State estimation of a deformable object x can be seen as

An optimization problem based on observations o and object representation $\mathbf{M}(\cdot)$

$$x^* = \arg \min_x \|o - \mathbf{M}(x)\|$$

$$x \in \text{ObjectStates}$$

The estimation problem may be formulated as


$$\theta^* = \arg \min_{\theta} \sum_t \|o_t - \mathbf{M}_t(\hat{x}_t)\|$$

$$\hat{x}_{t+1} = \text{ObjectDynamics}(\hat{x}_t, f_t, \theta)$$

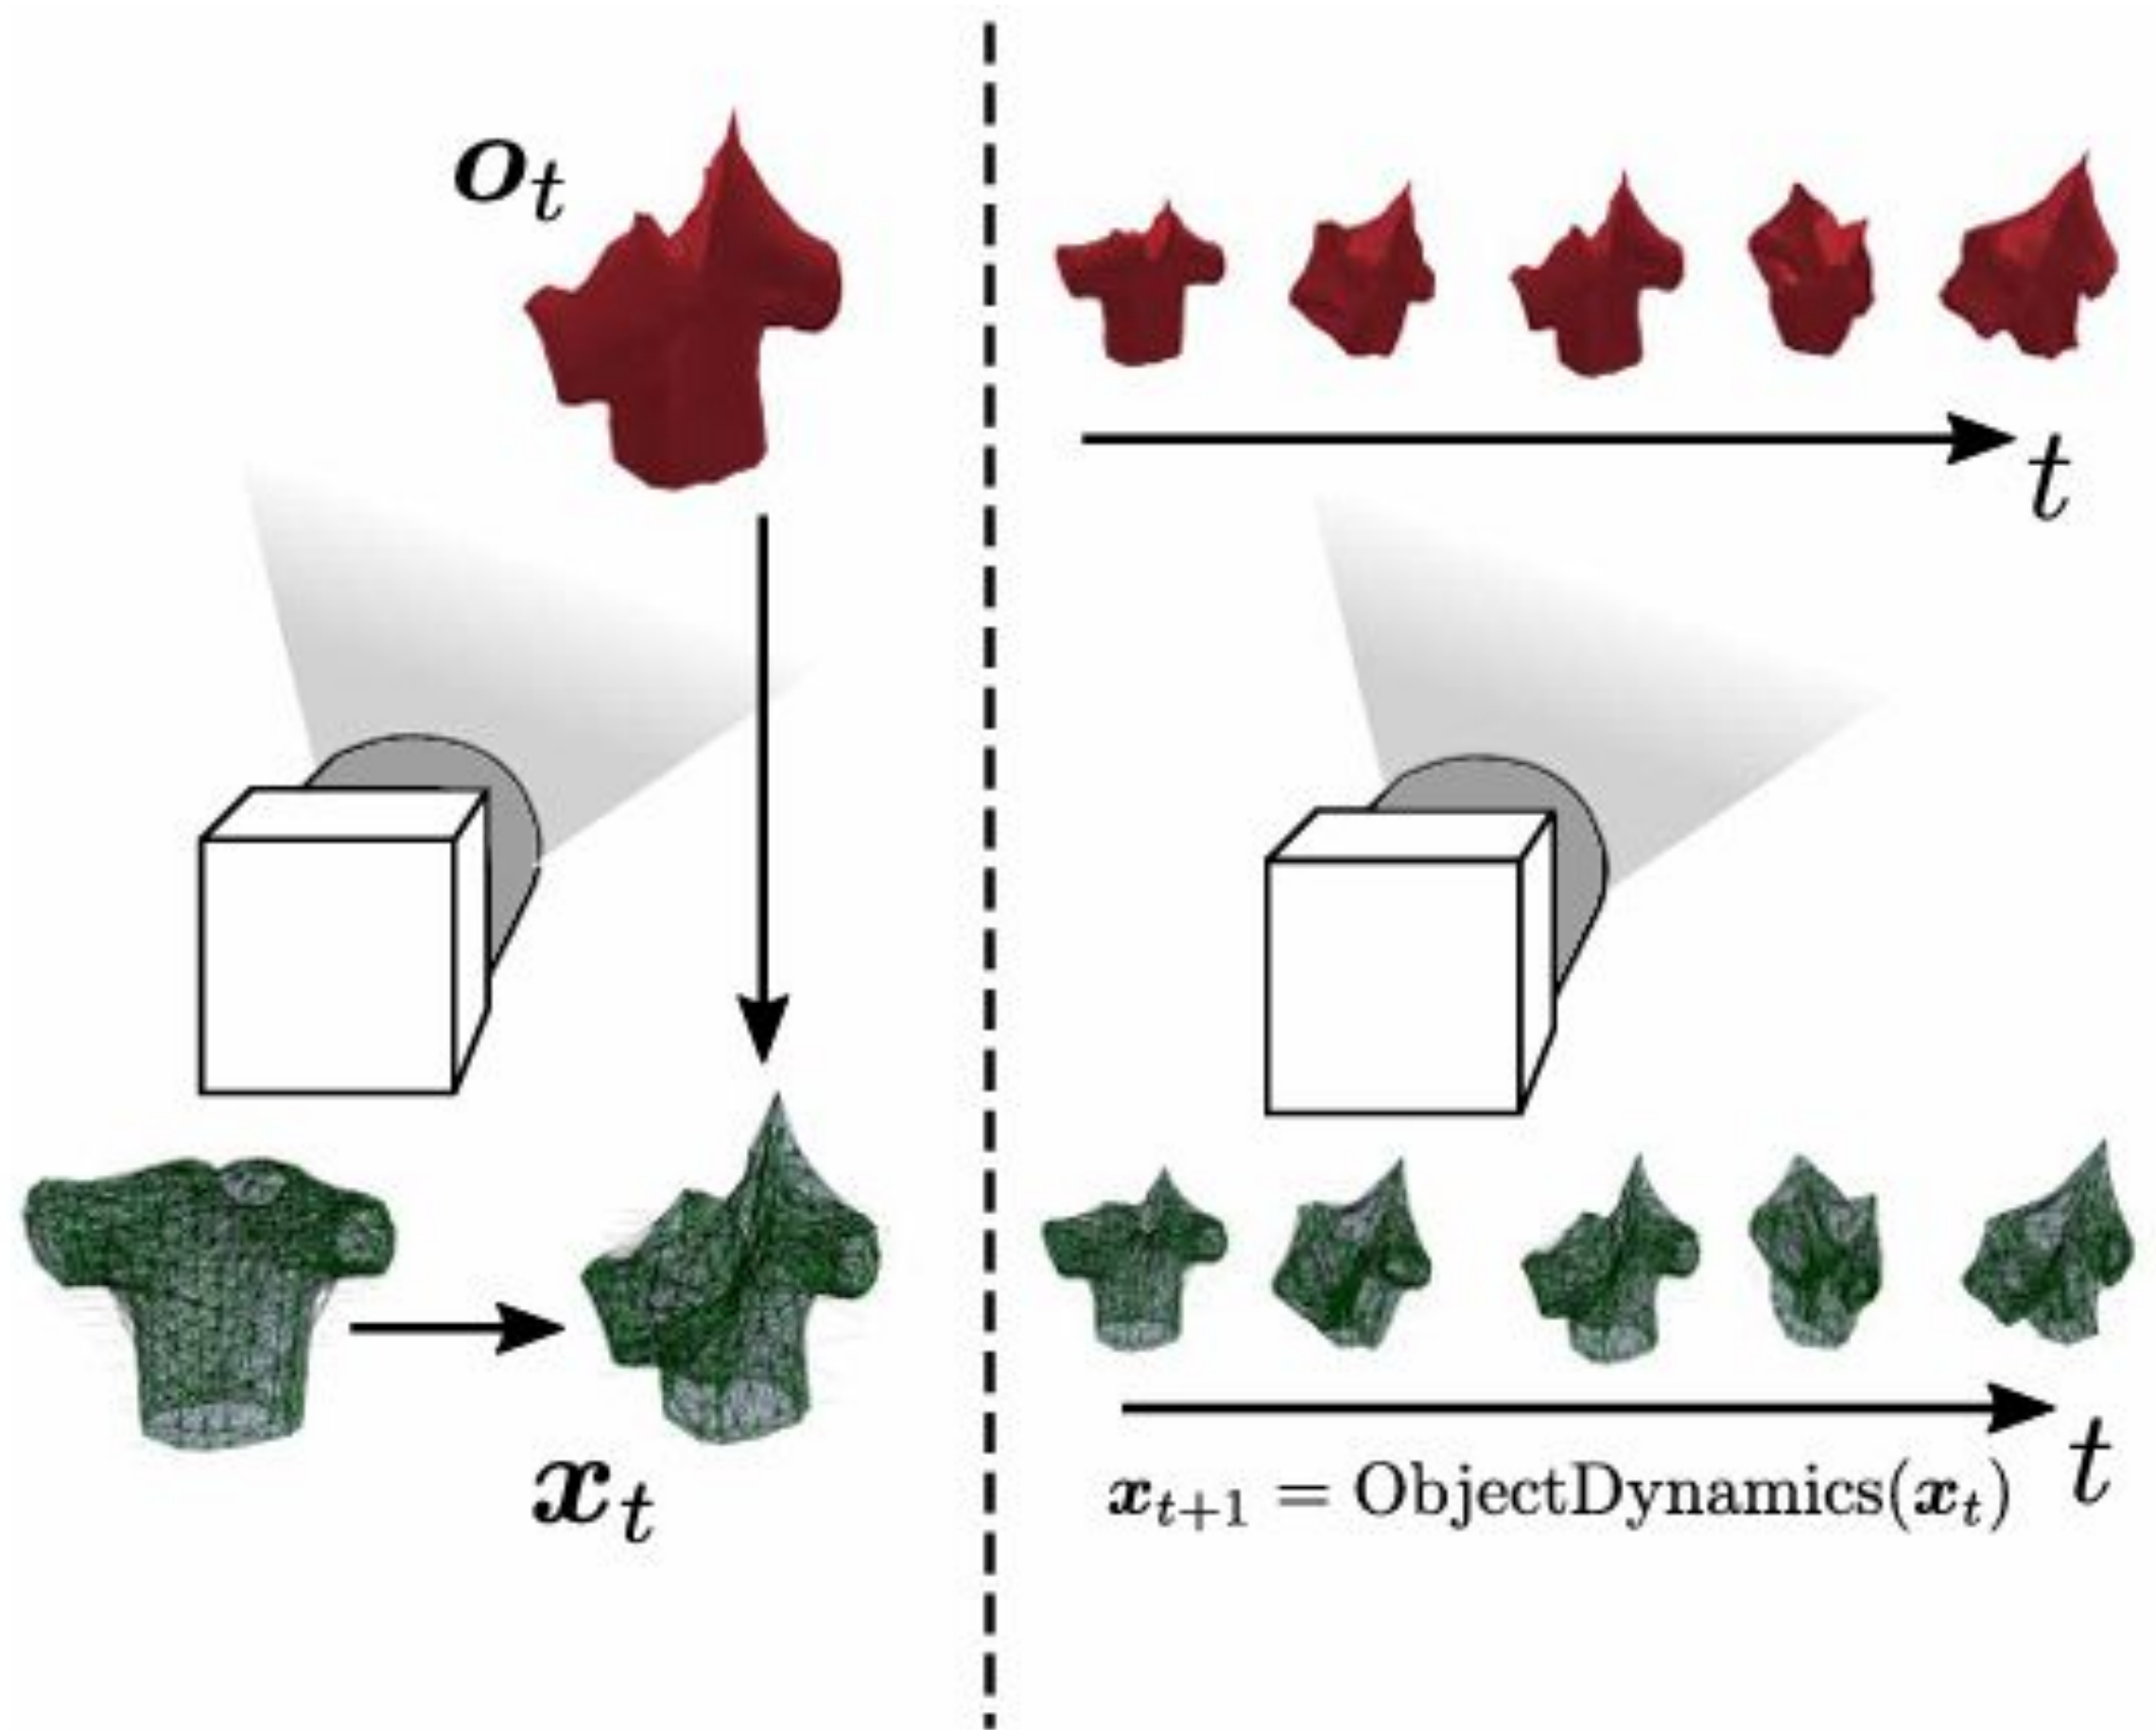
where \hat{x} is the predicted state that depends on object parameters such as material properties θ and applied forces f




Deformable Object Perception Tasks



Deformable Object Perception Tasks



Deformable Object Perception Tasks





Manipulation

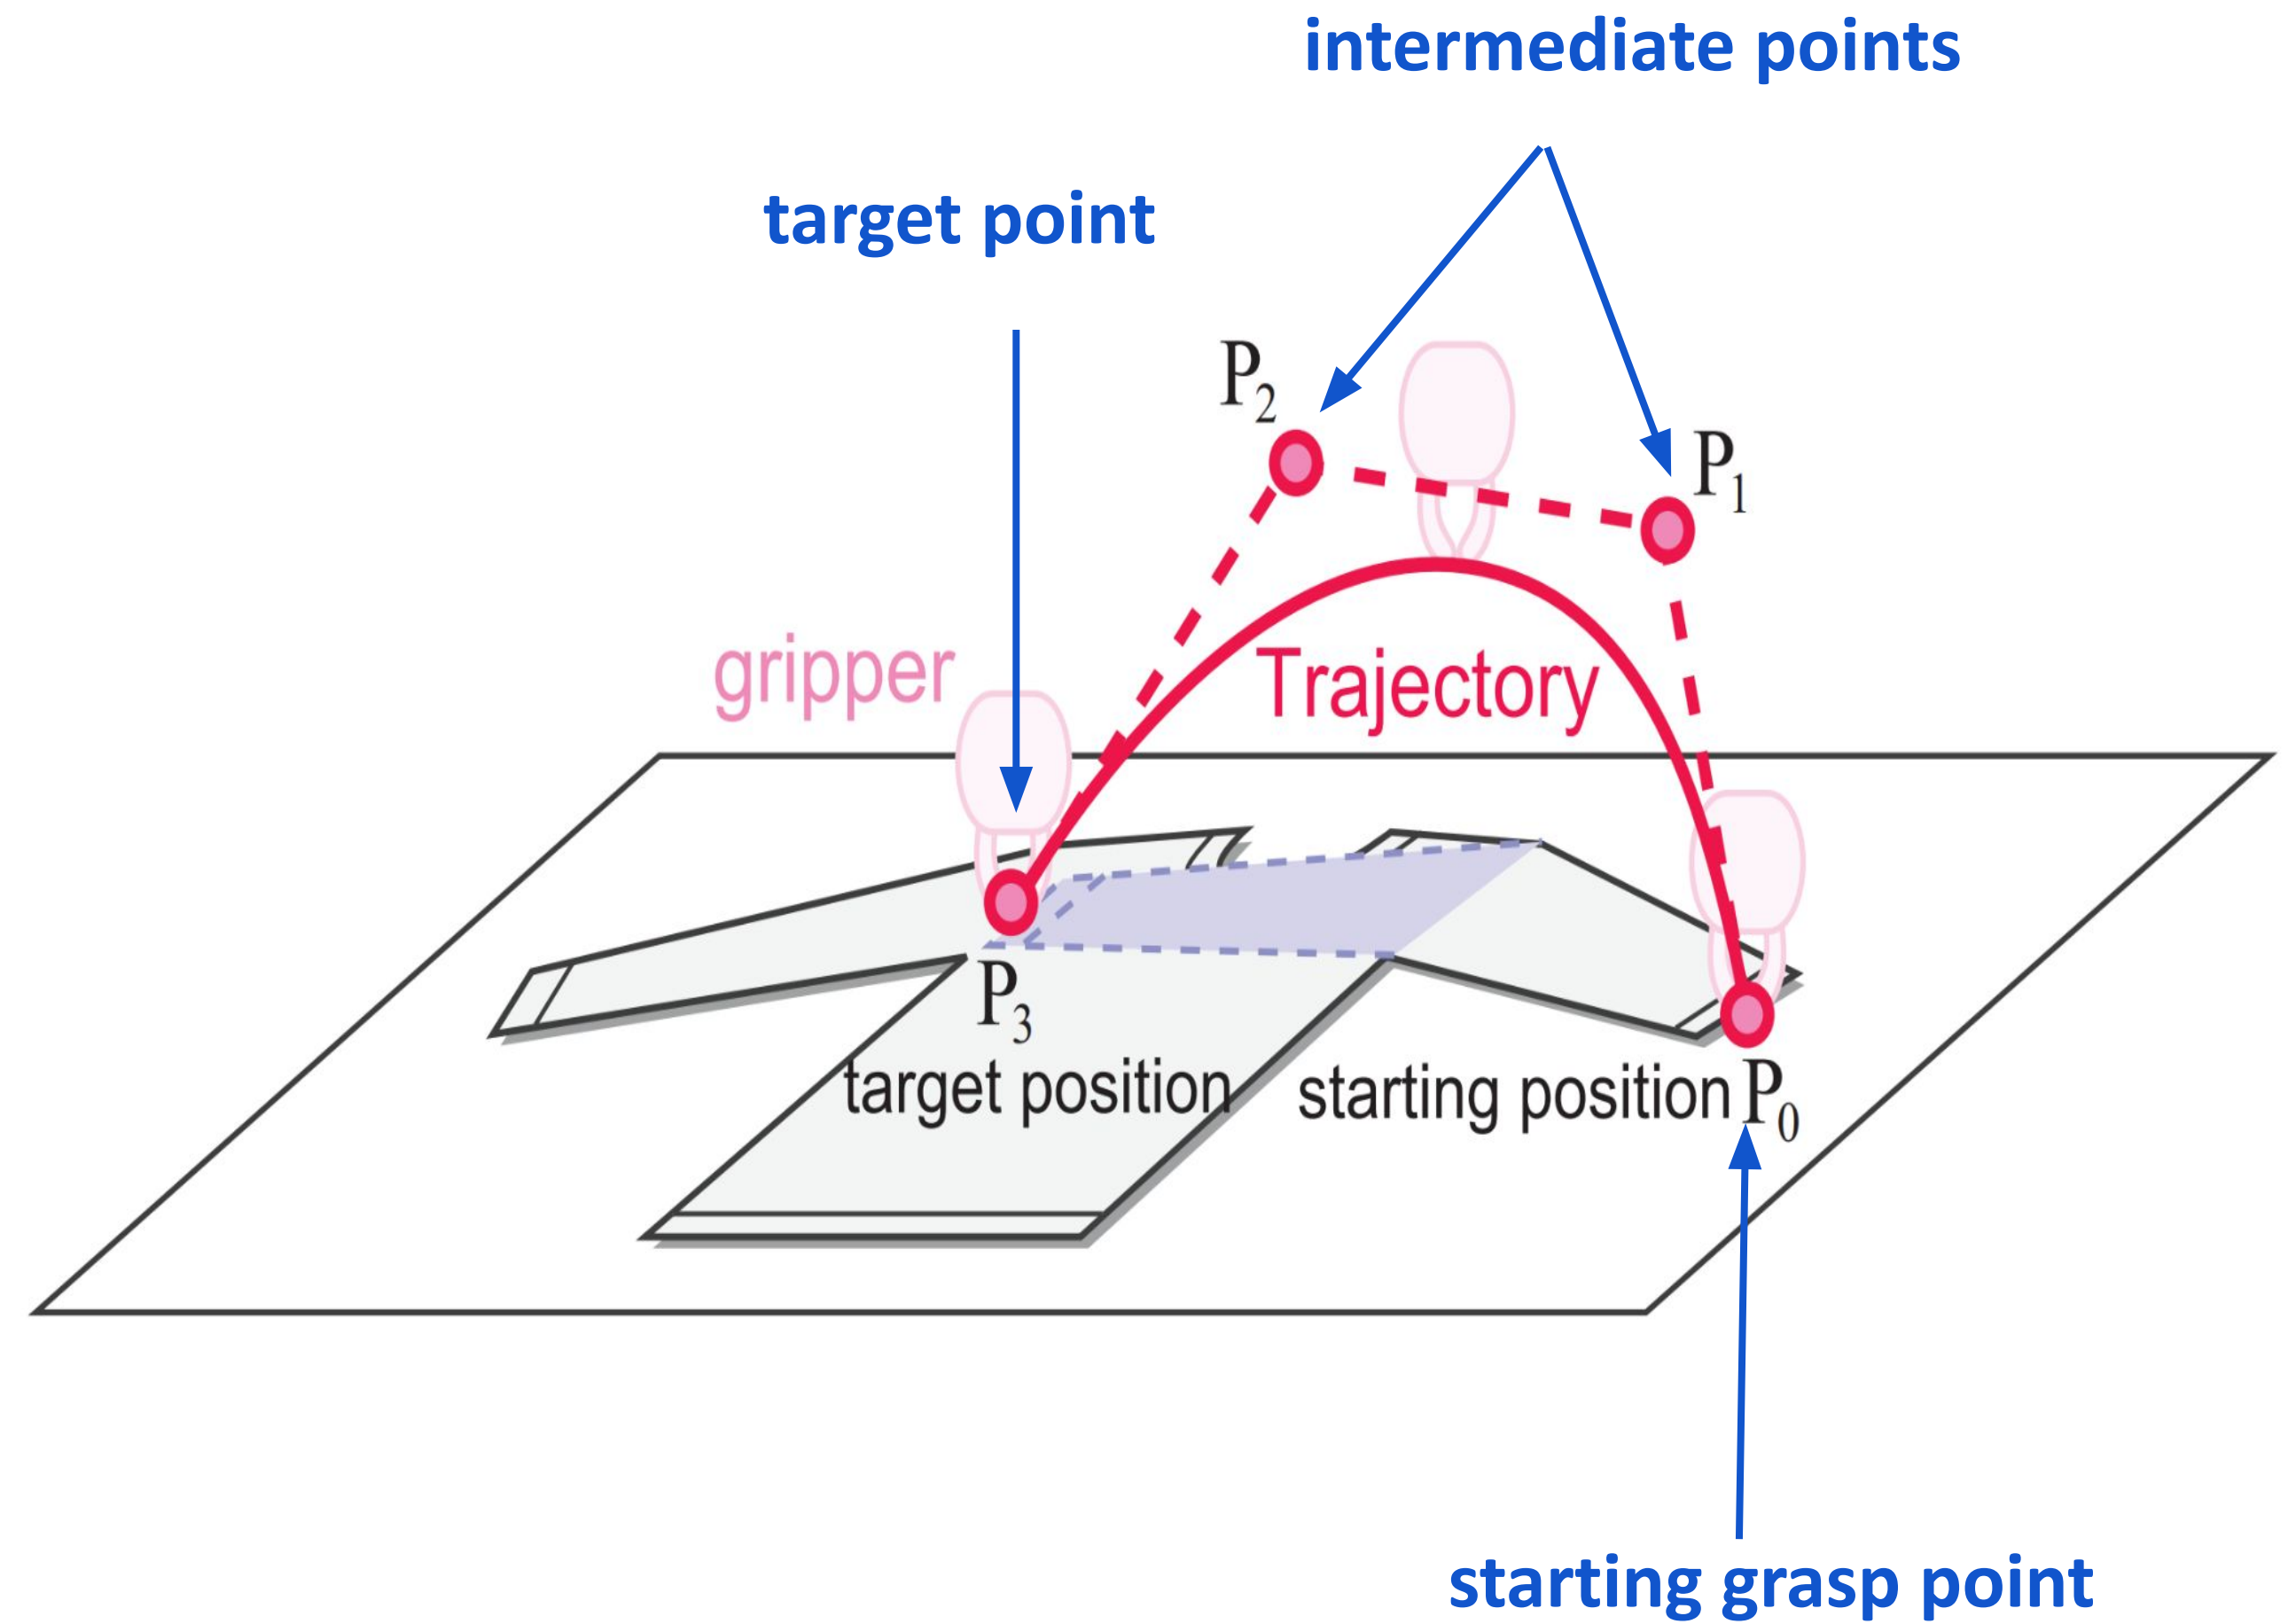




Problem Definition

Task: Fold the sleeve into blue target position

Use a robotic manipulator to grasp the sleeve at P_0 and move to target point P_3





Mathematical Formulation


Goal: To find

a) external force F_{ext} (acting on the cloth)

OR

b) motion of the grasp point (P_0, P_1, P_2, P_3)

in order to achieve the target position of sleeve



(4 points assumed for the sake of simplicity)



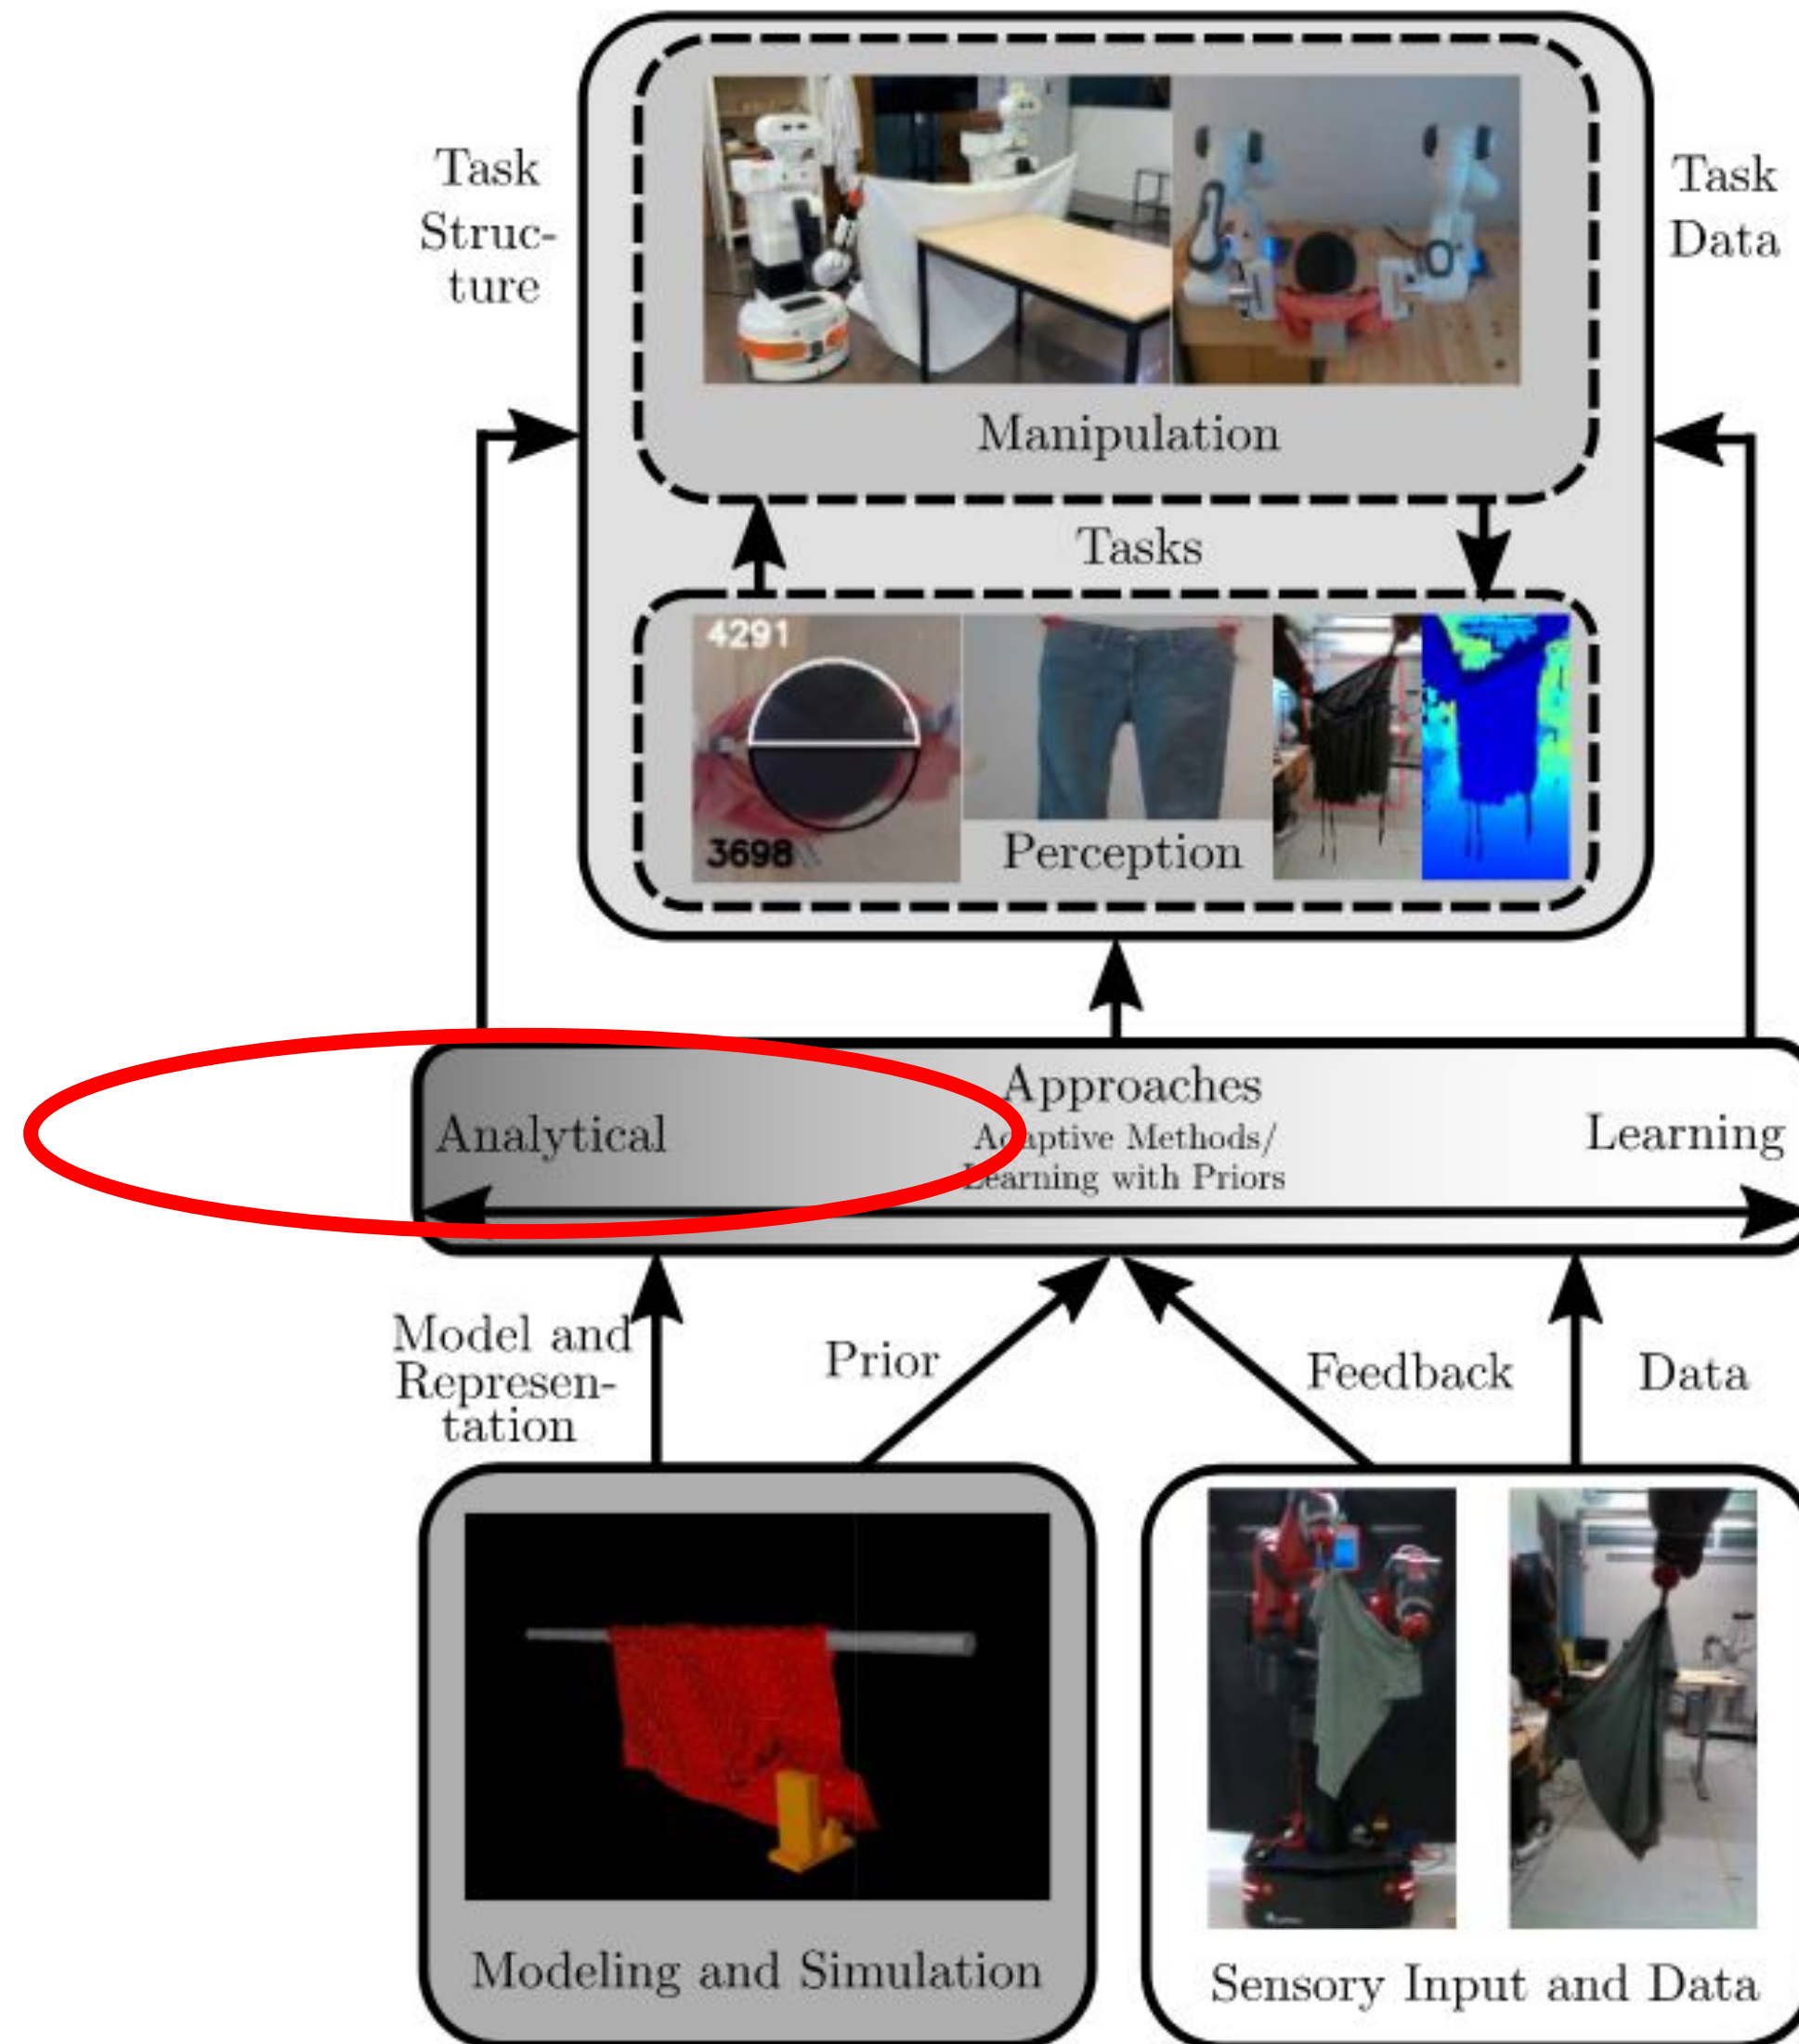


How can we solve this?





Recap





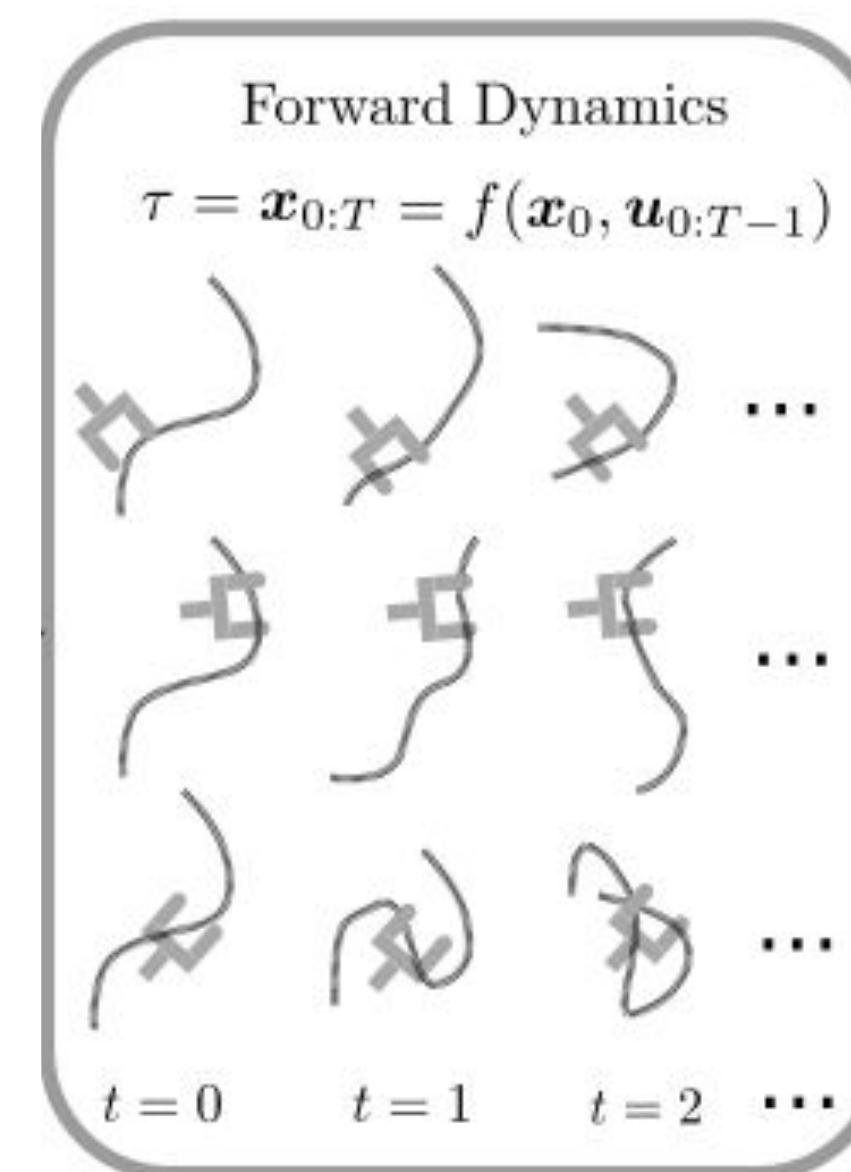


How can we solve this?

1) Analytical Method



1) Shooting in action space


Assumption: Dynamic model of the cloth is known



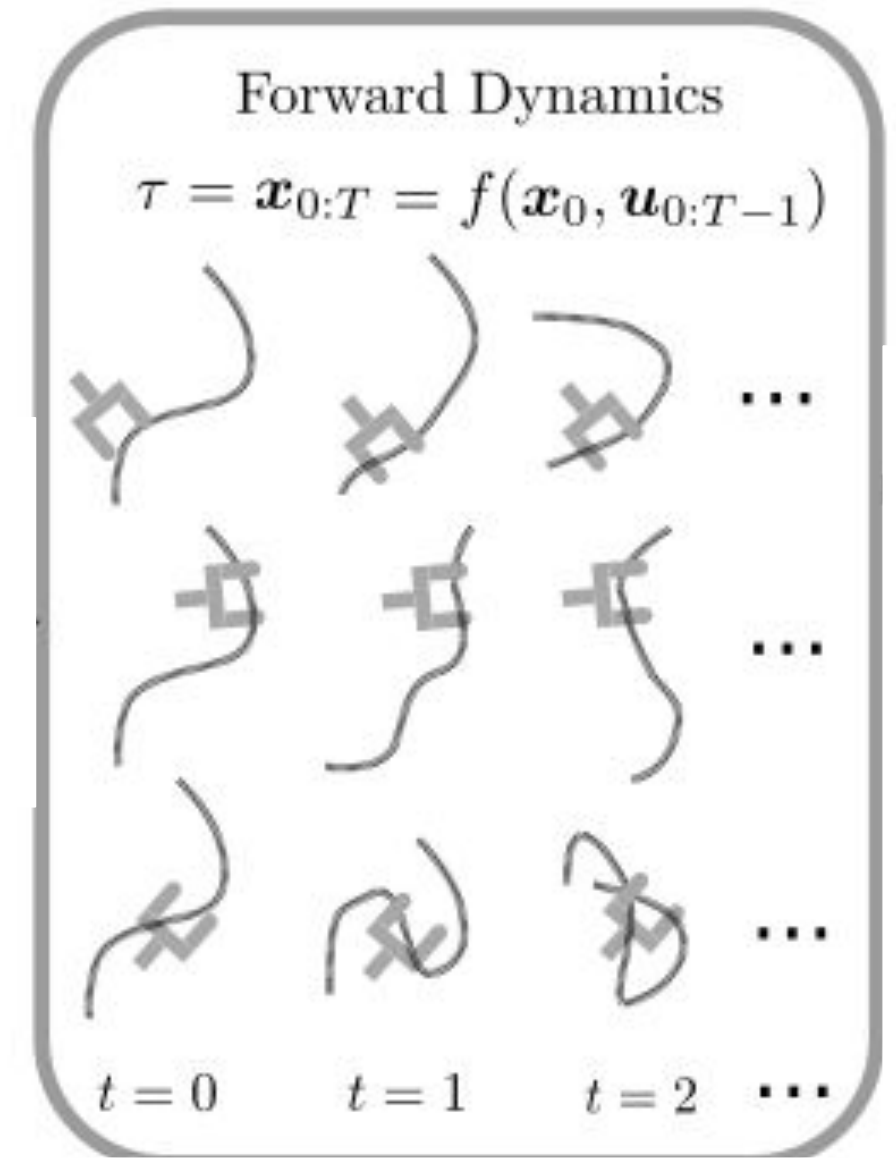




1) Shooting in action space

Assumption: Dynamic model of the cloth is known



Step 1) sample an action trajectory $u_{0:T-1}$ (random/heuristic)




$u_{0:T-1}$



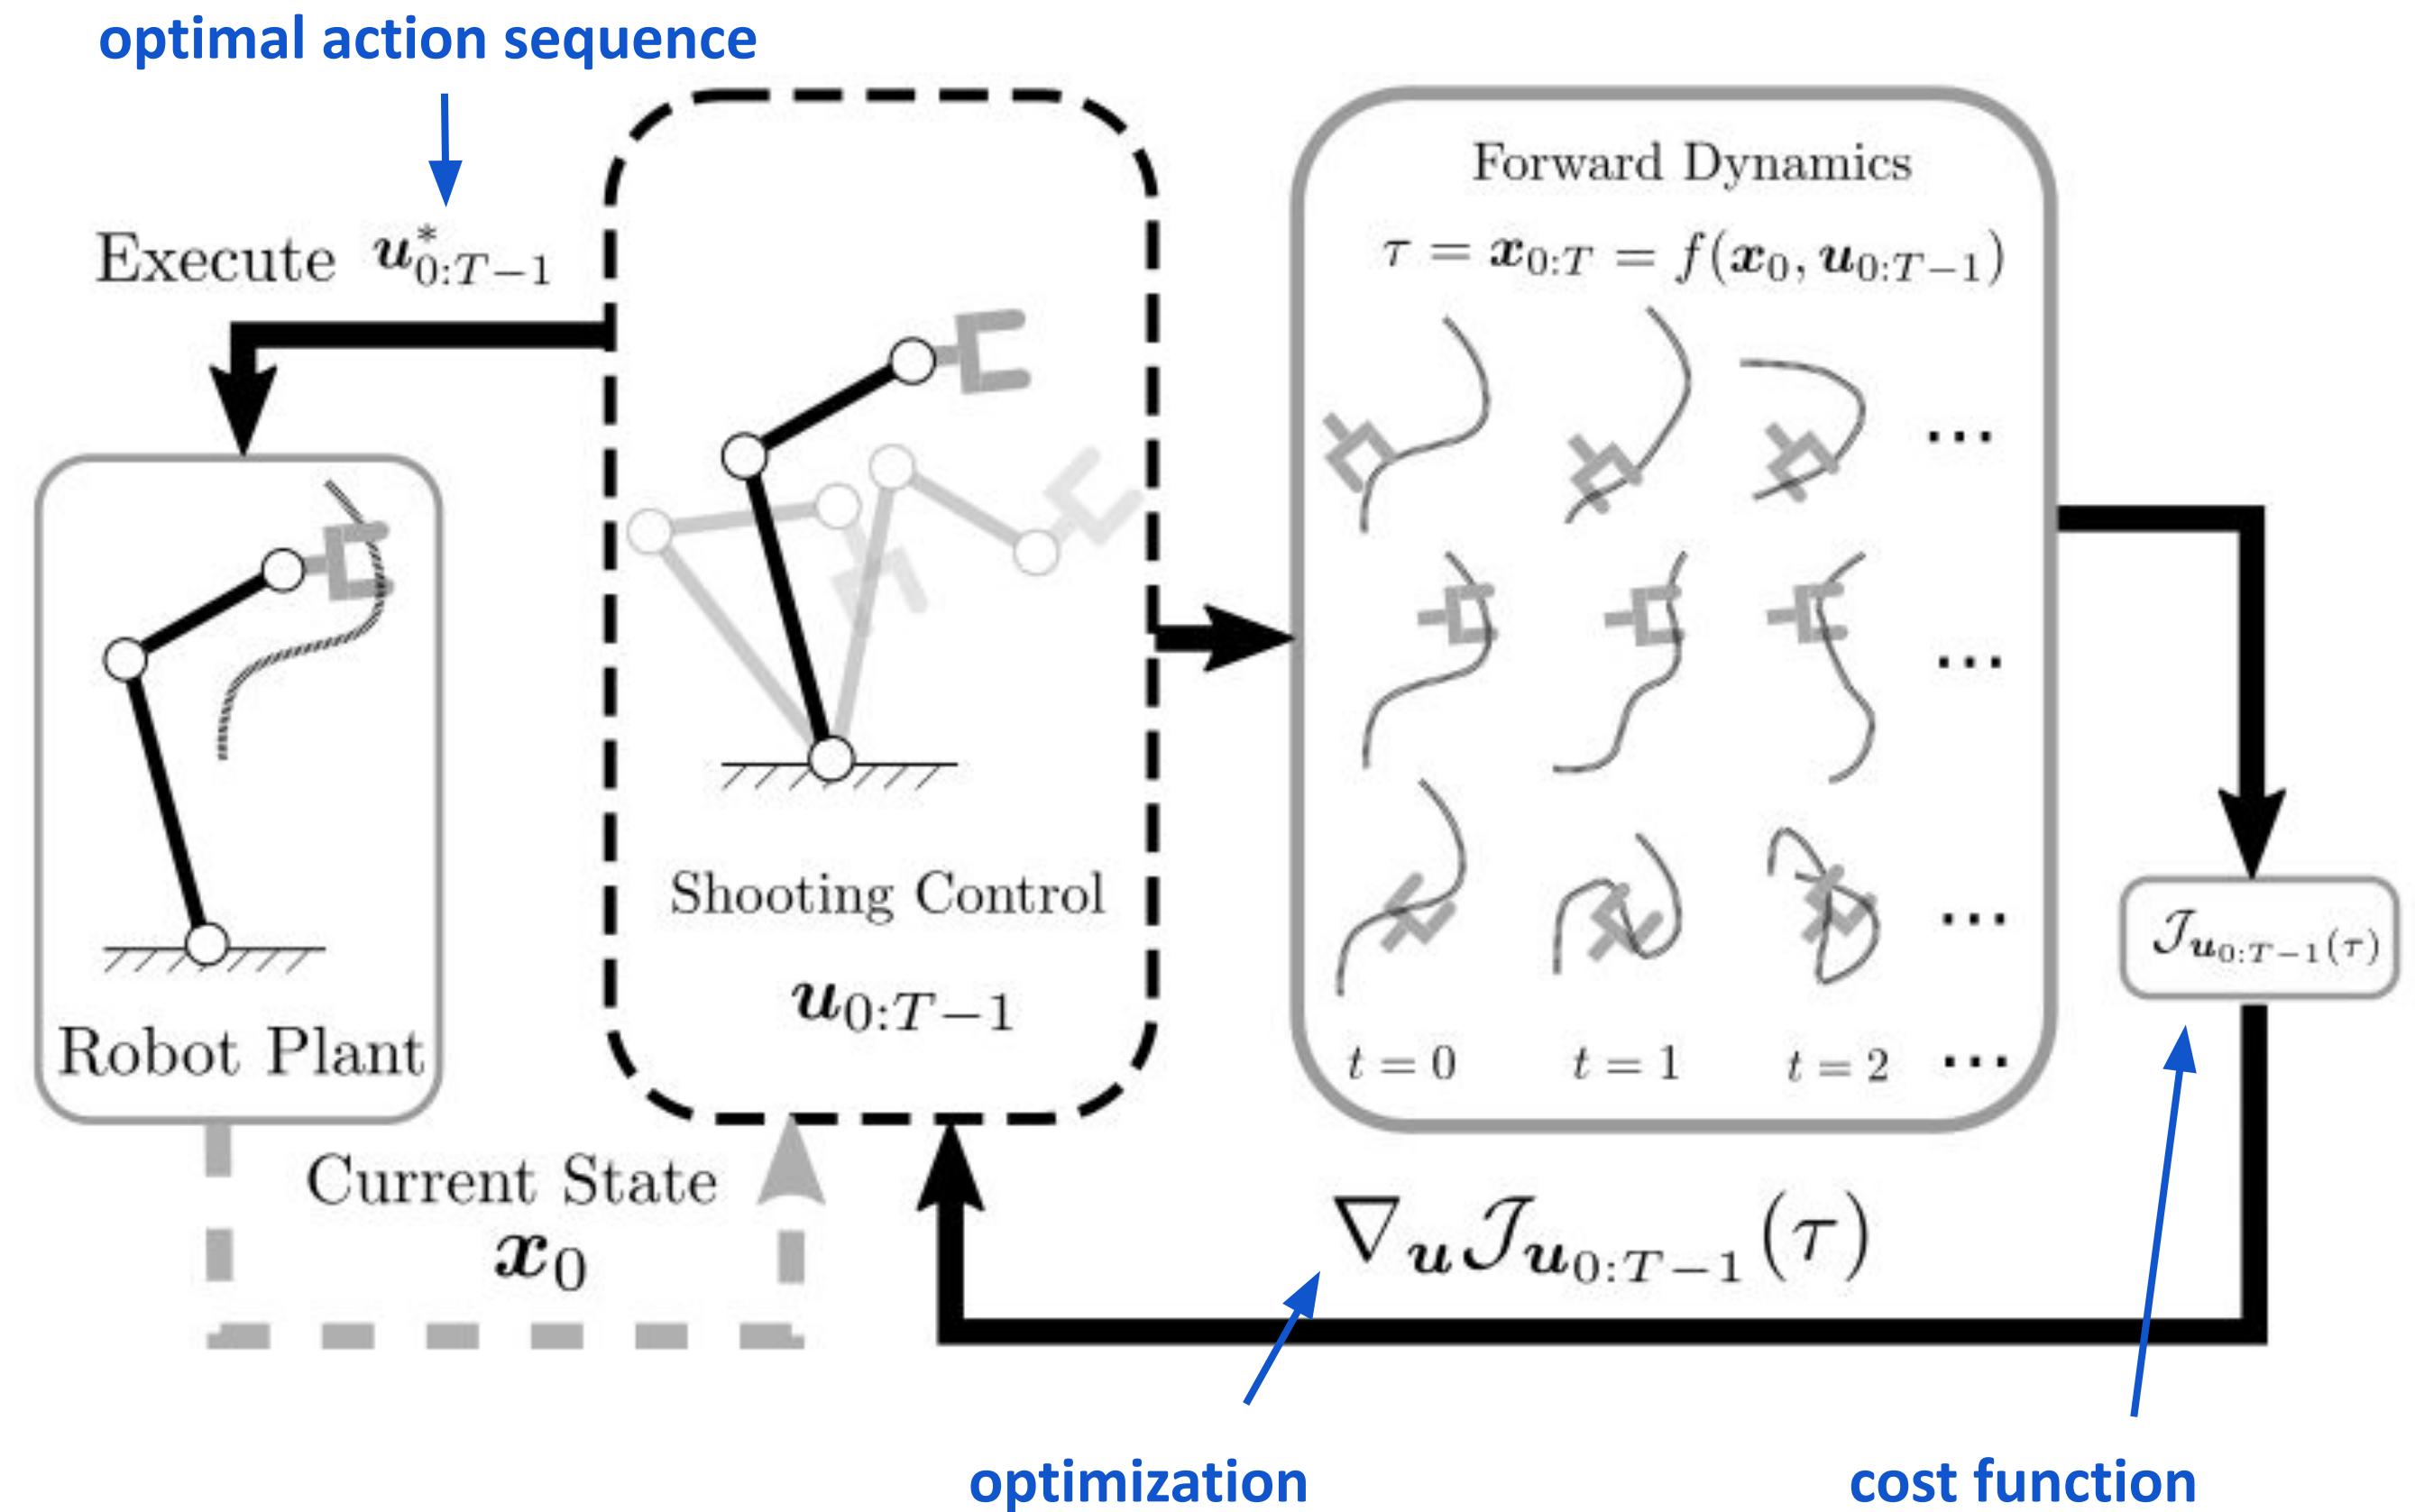
1) Shooting in action space

Assumption: Dynamic model of the cloth is known



Step 1) sample an action trajectory $u_{0:T-1}$ (random/heuristic)


Step 2) Update actions according to evaluated costs





1) Shooting in action space

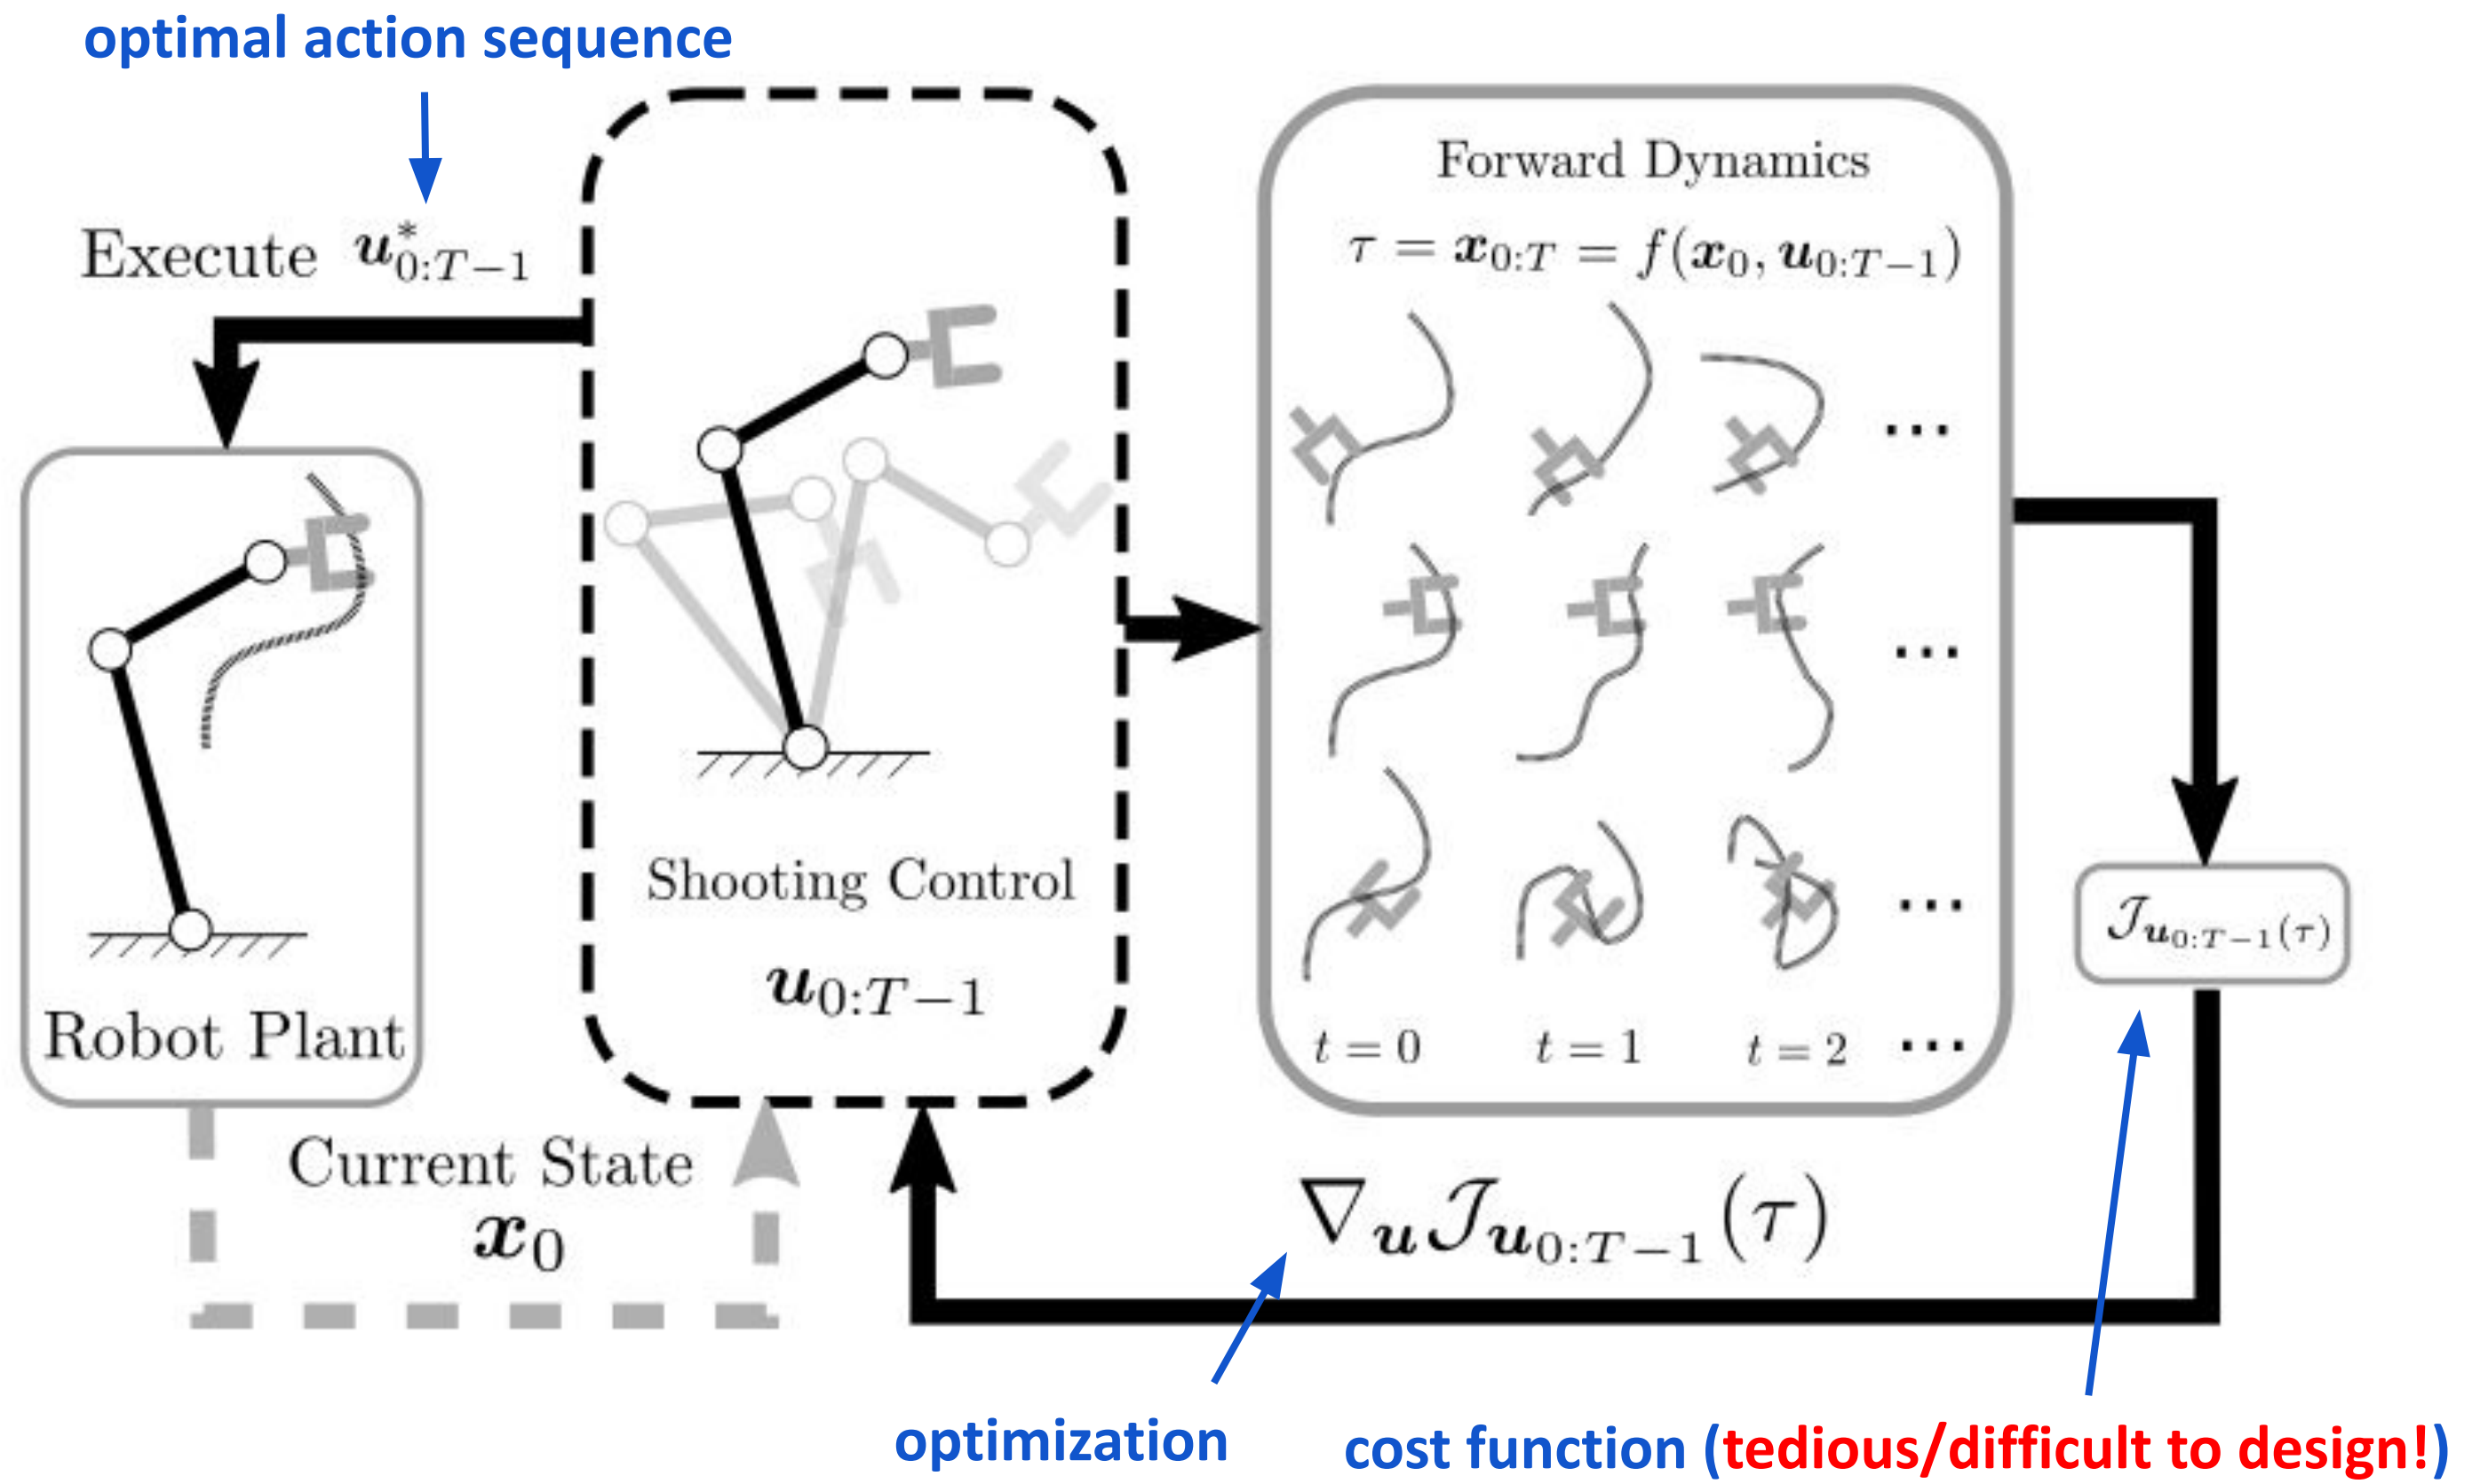
Assumption: Dynamic model of the cloth is known



Step 1) sample an action trajectory $u_{0:T-1}$ (random/heuristic)

Step 2) Update actions according to evaluated costs


Step 3) Resample after partial execution (receding horizon)





2) Search trajectories in object state space

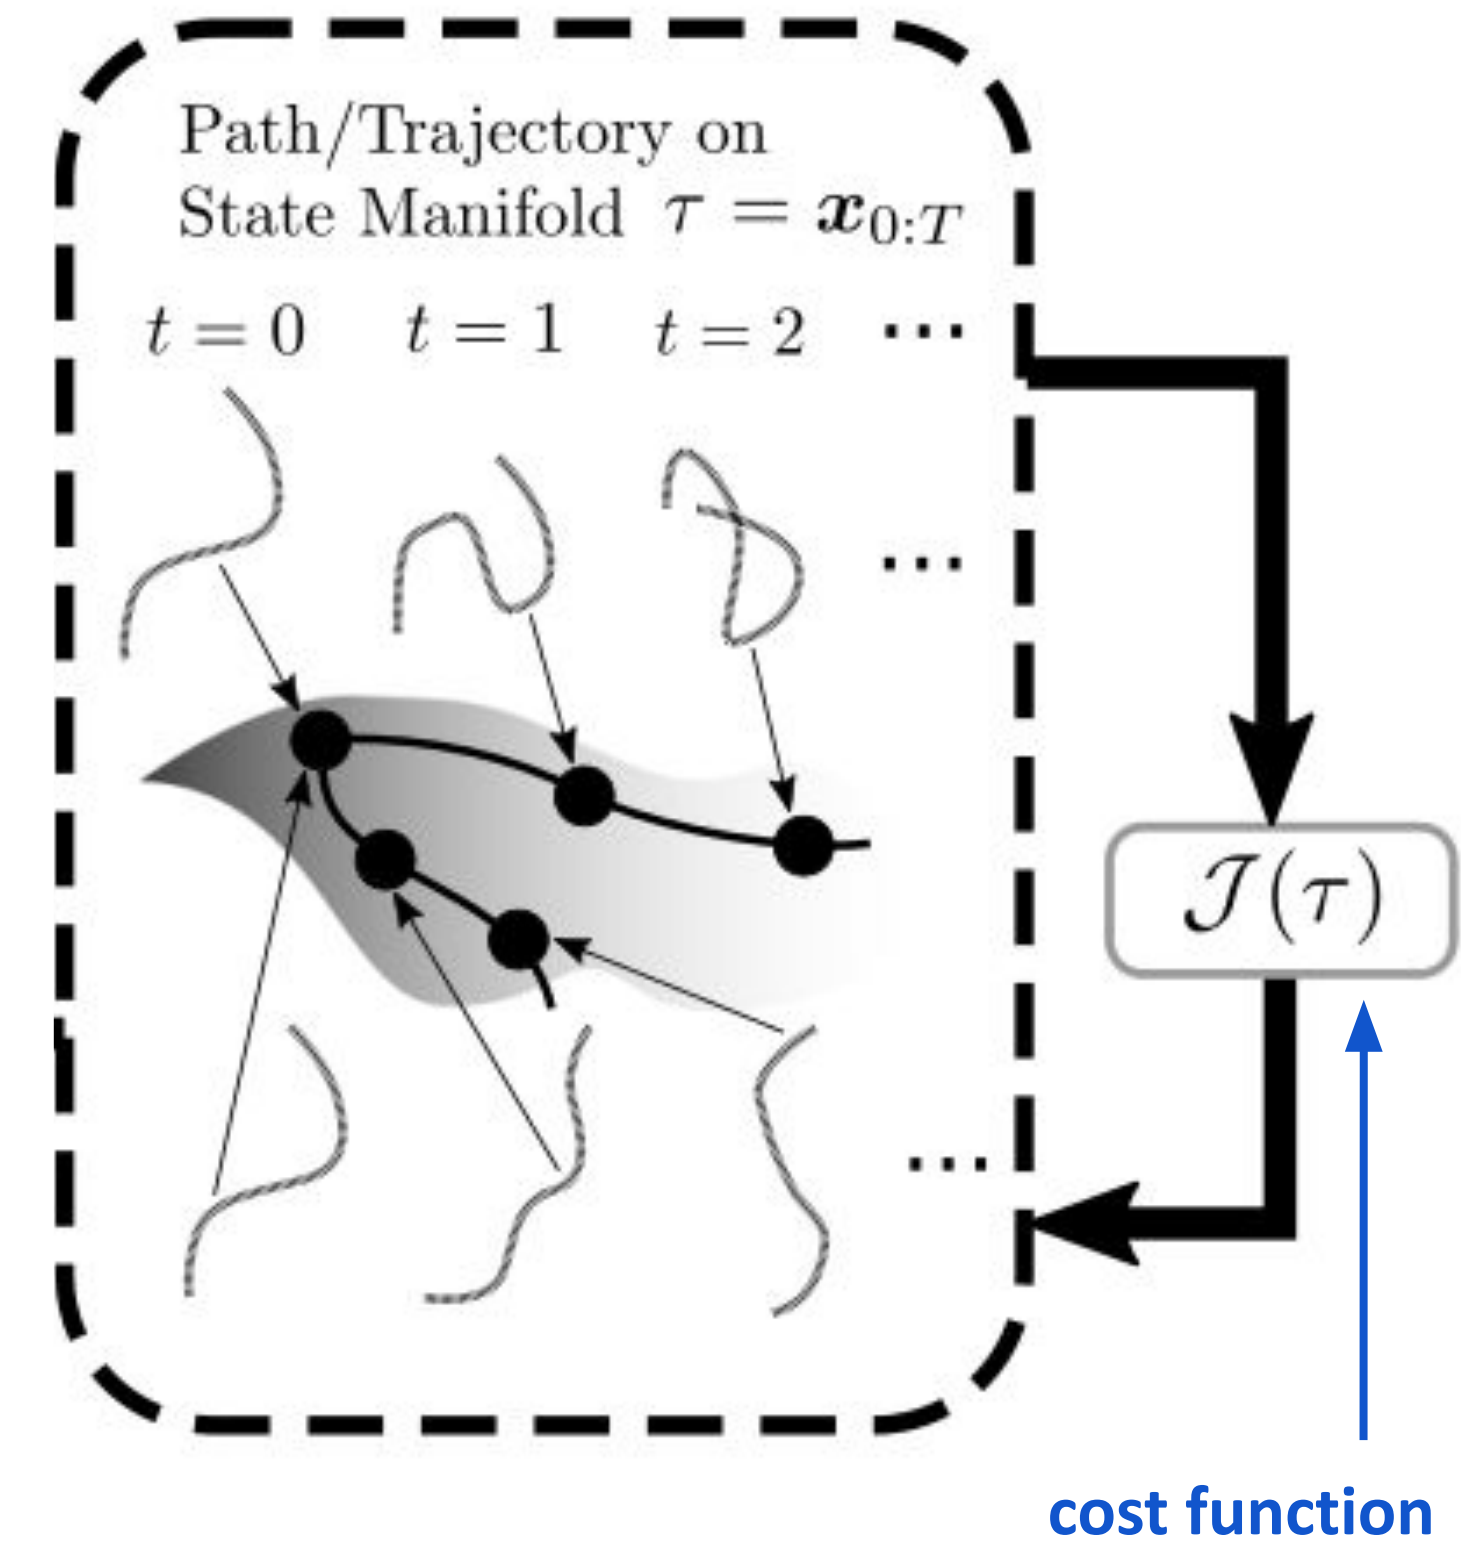

Assumption: Dynamic model of the cloth is known





2) Search trajectories in object state space

Step 1) find a low cost path $\tau = \mathbf{x}_{0:T}$ on the state manifold






2) Search trajectories in object state space

Step 1) find a low cost path $\tau = \mathbf{x}_{0:T}$ on the state manifold

Step 2) generate actions \mathbf{u}_t cause transitions at each t



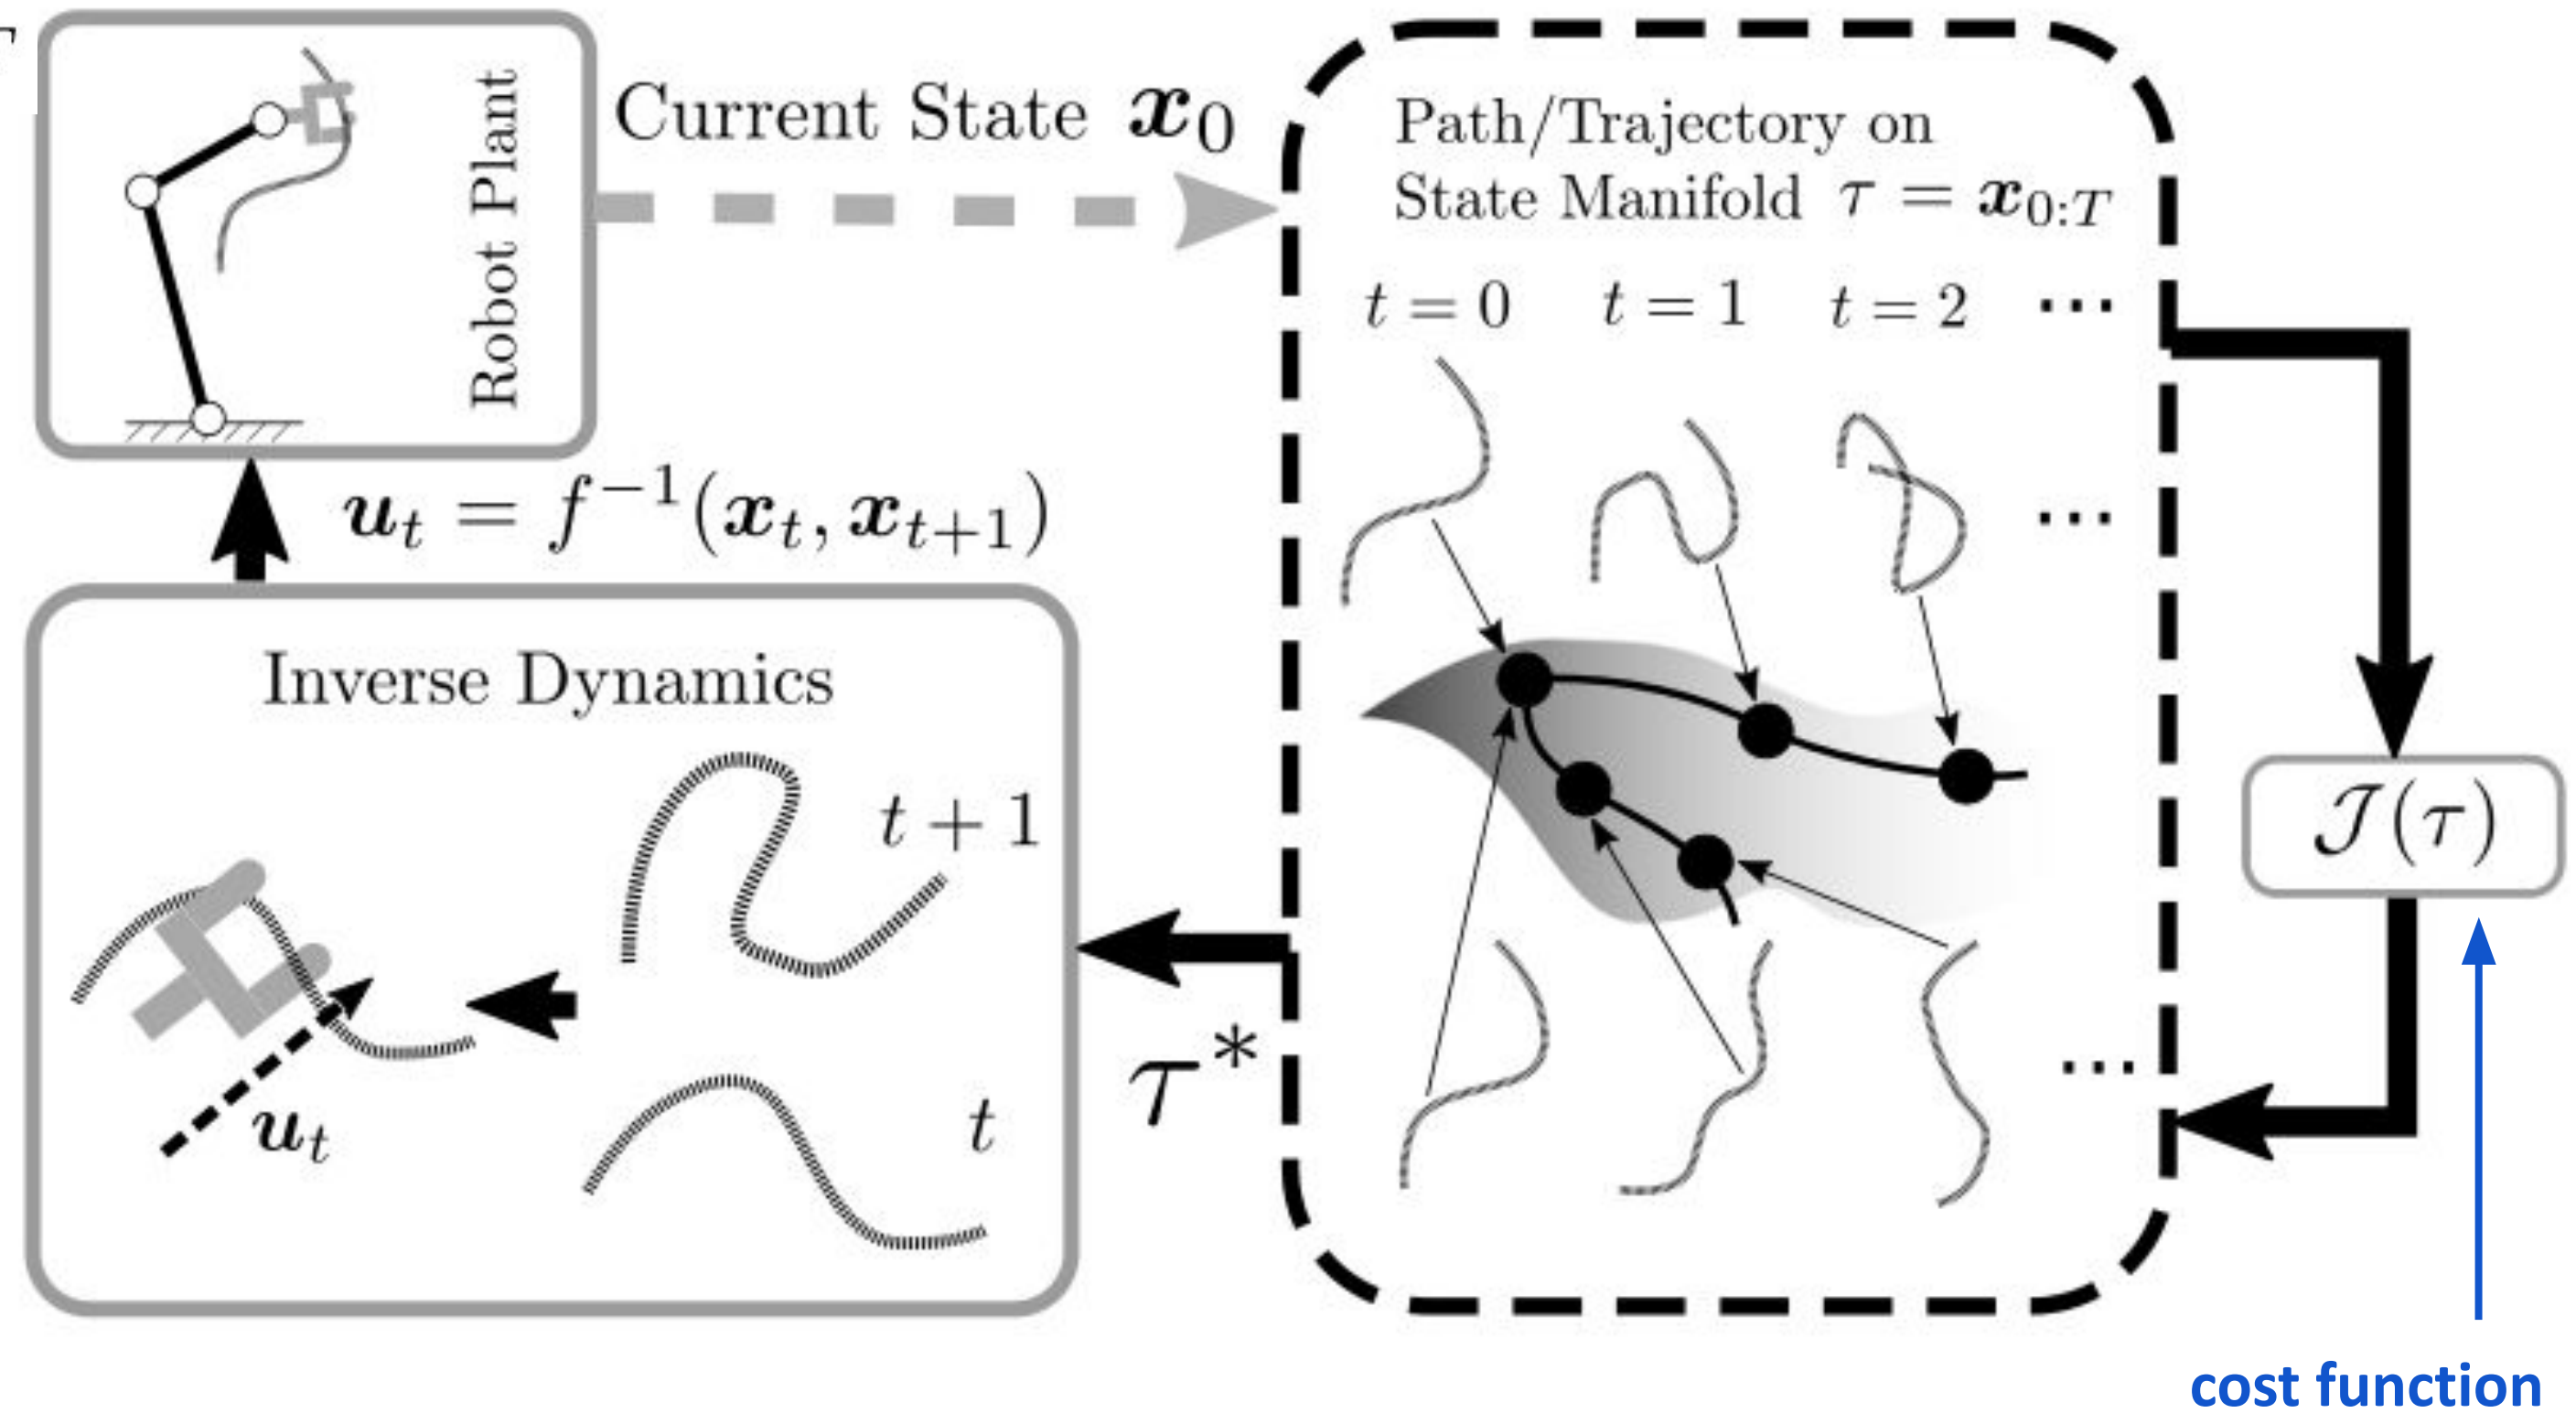


2) Search trajectories in object state space

Step 1) find a low cost path $\tau = \mathbf{x}_{0:T}$ on the state manifold

Step 2) generate actions \mathbf{u}_t cause transitions at each t

Issues:
Vast search space!
Tedious to design cost function!





In both the previous methods, we made an assumption that **we already know the dynamic model**





In both the previous methods, we made an assumption that **we already know the dynamic model**

But in reality, it is **very difficult** to find the dynamic model of the cloth due to

- high dimensionality
- nonlinear dynamics
- self collision

-
-
-





Can I solve this task without explicitly having a dynamic model?





Can I solve this task without explicitly having a dynamic model?

Can I learn to control in an end to end manner?





Can I solve this task without explicitly having a dynamic model?


Can I learn to control in an end to end manner?

Model free learning based approaches!



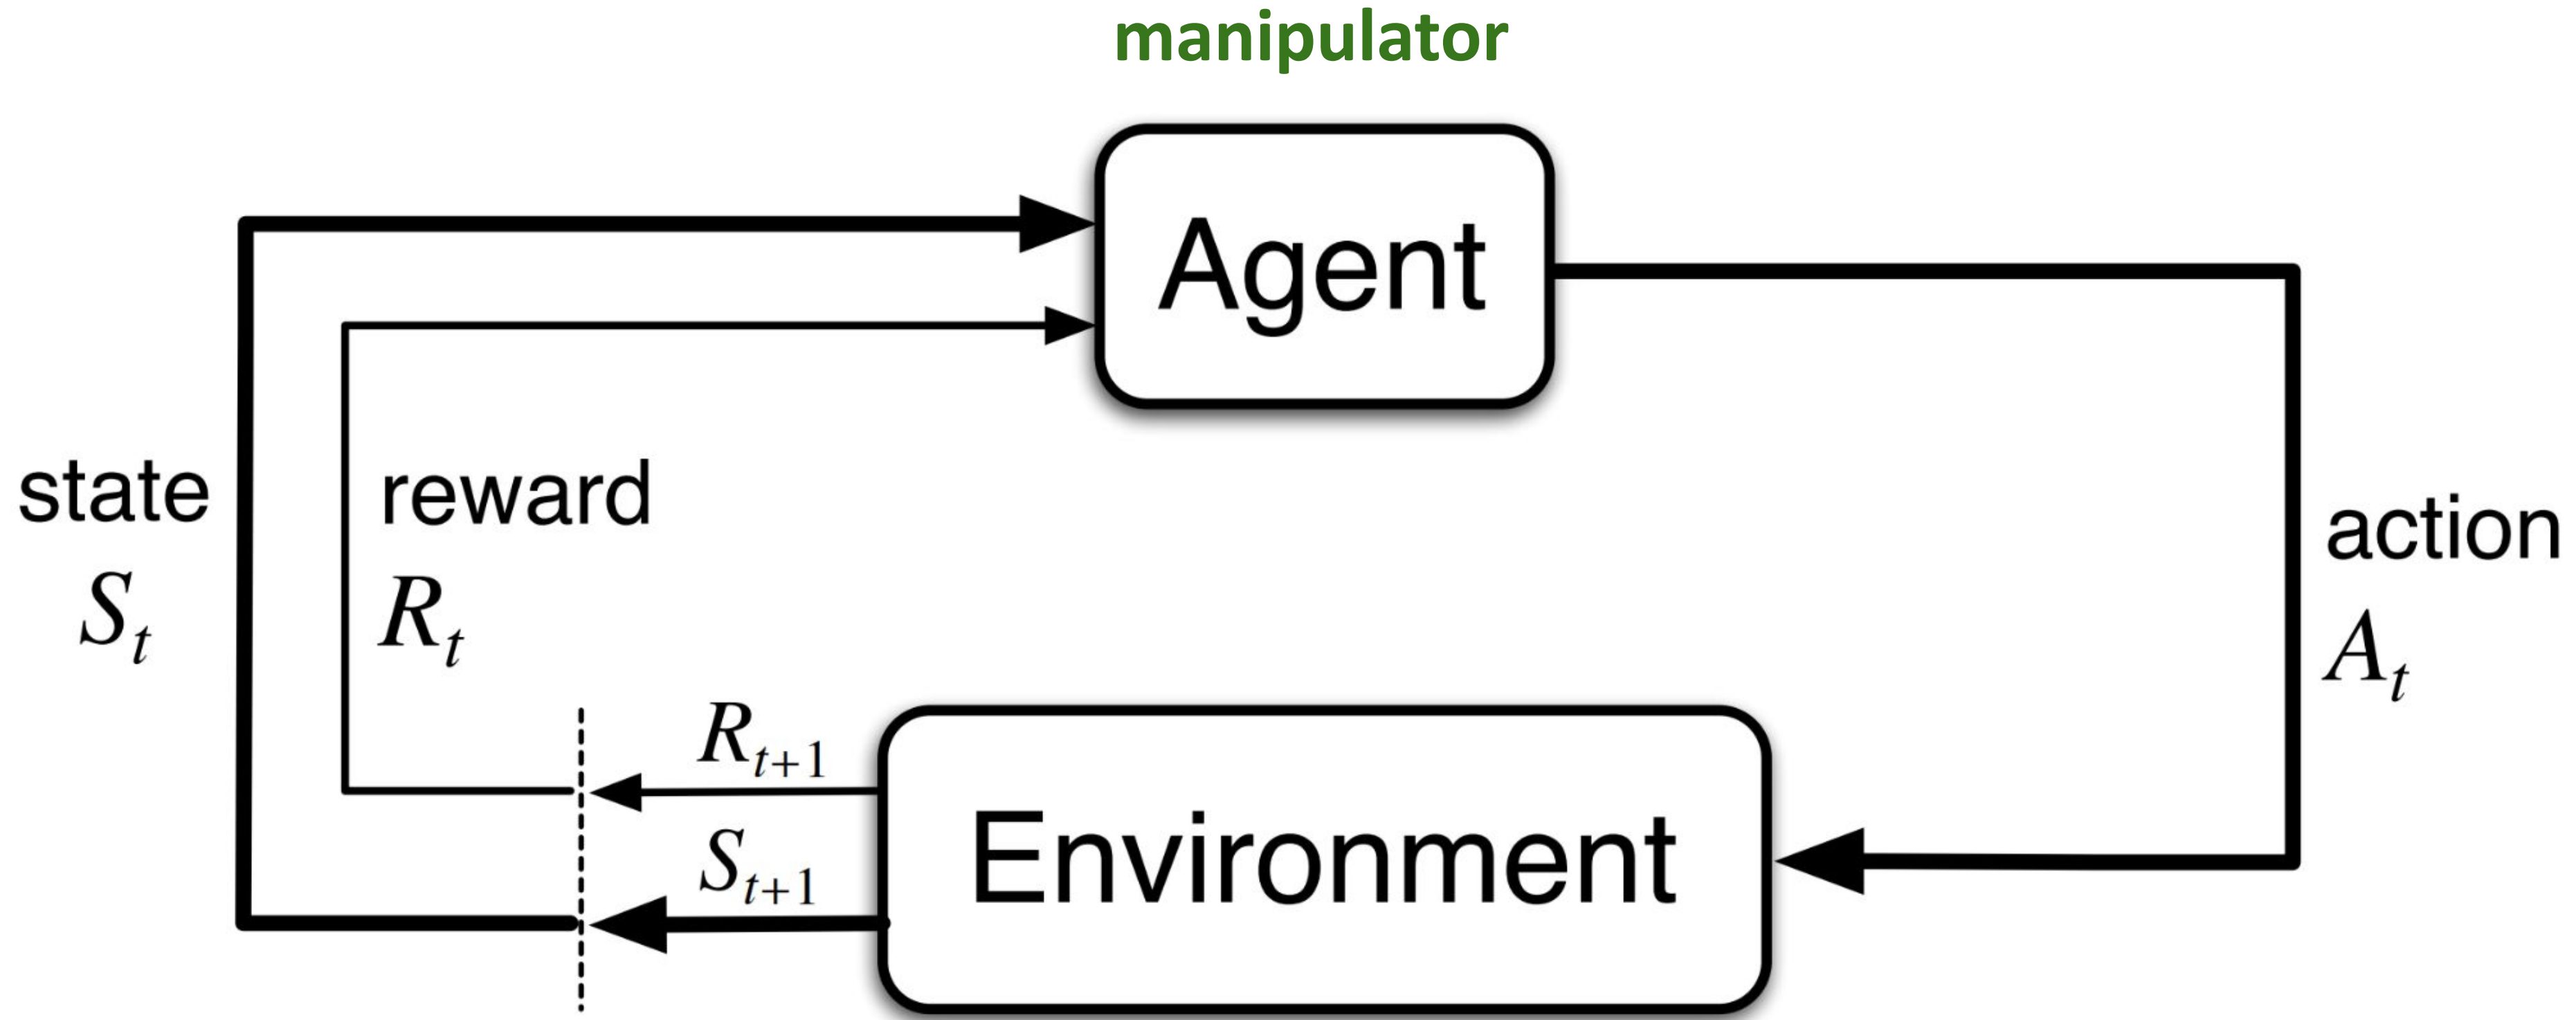


Reinforcement Learning



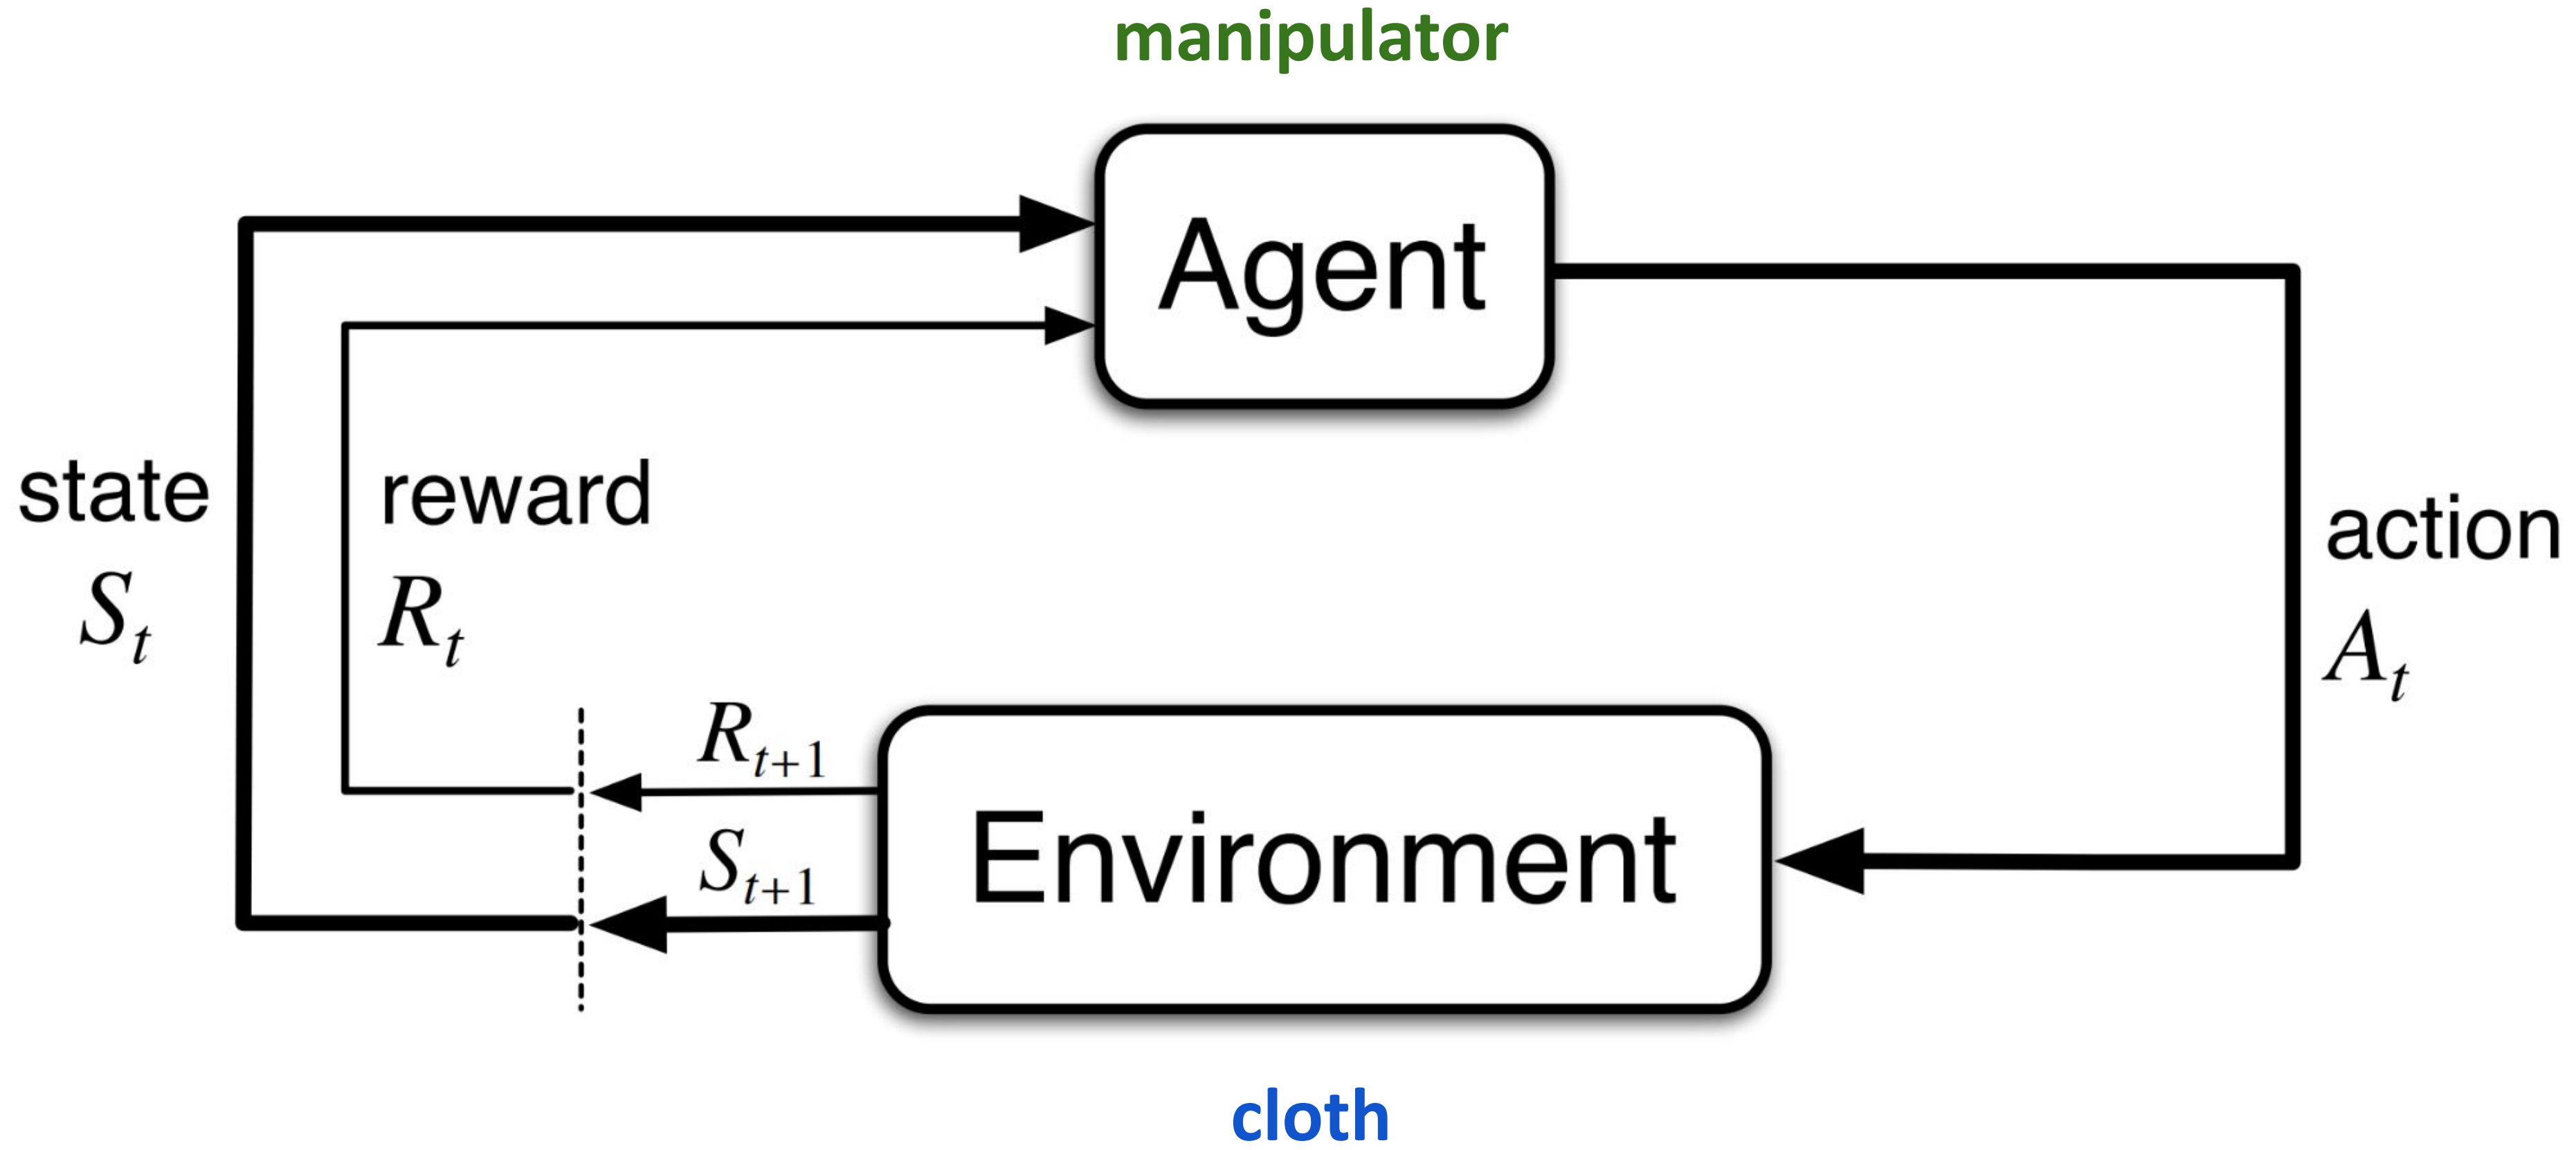


Reinforcement Learning



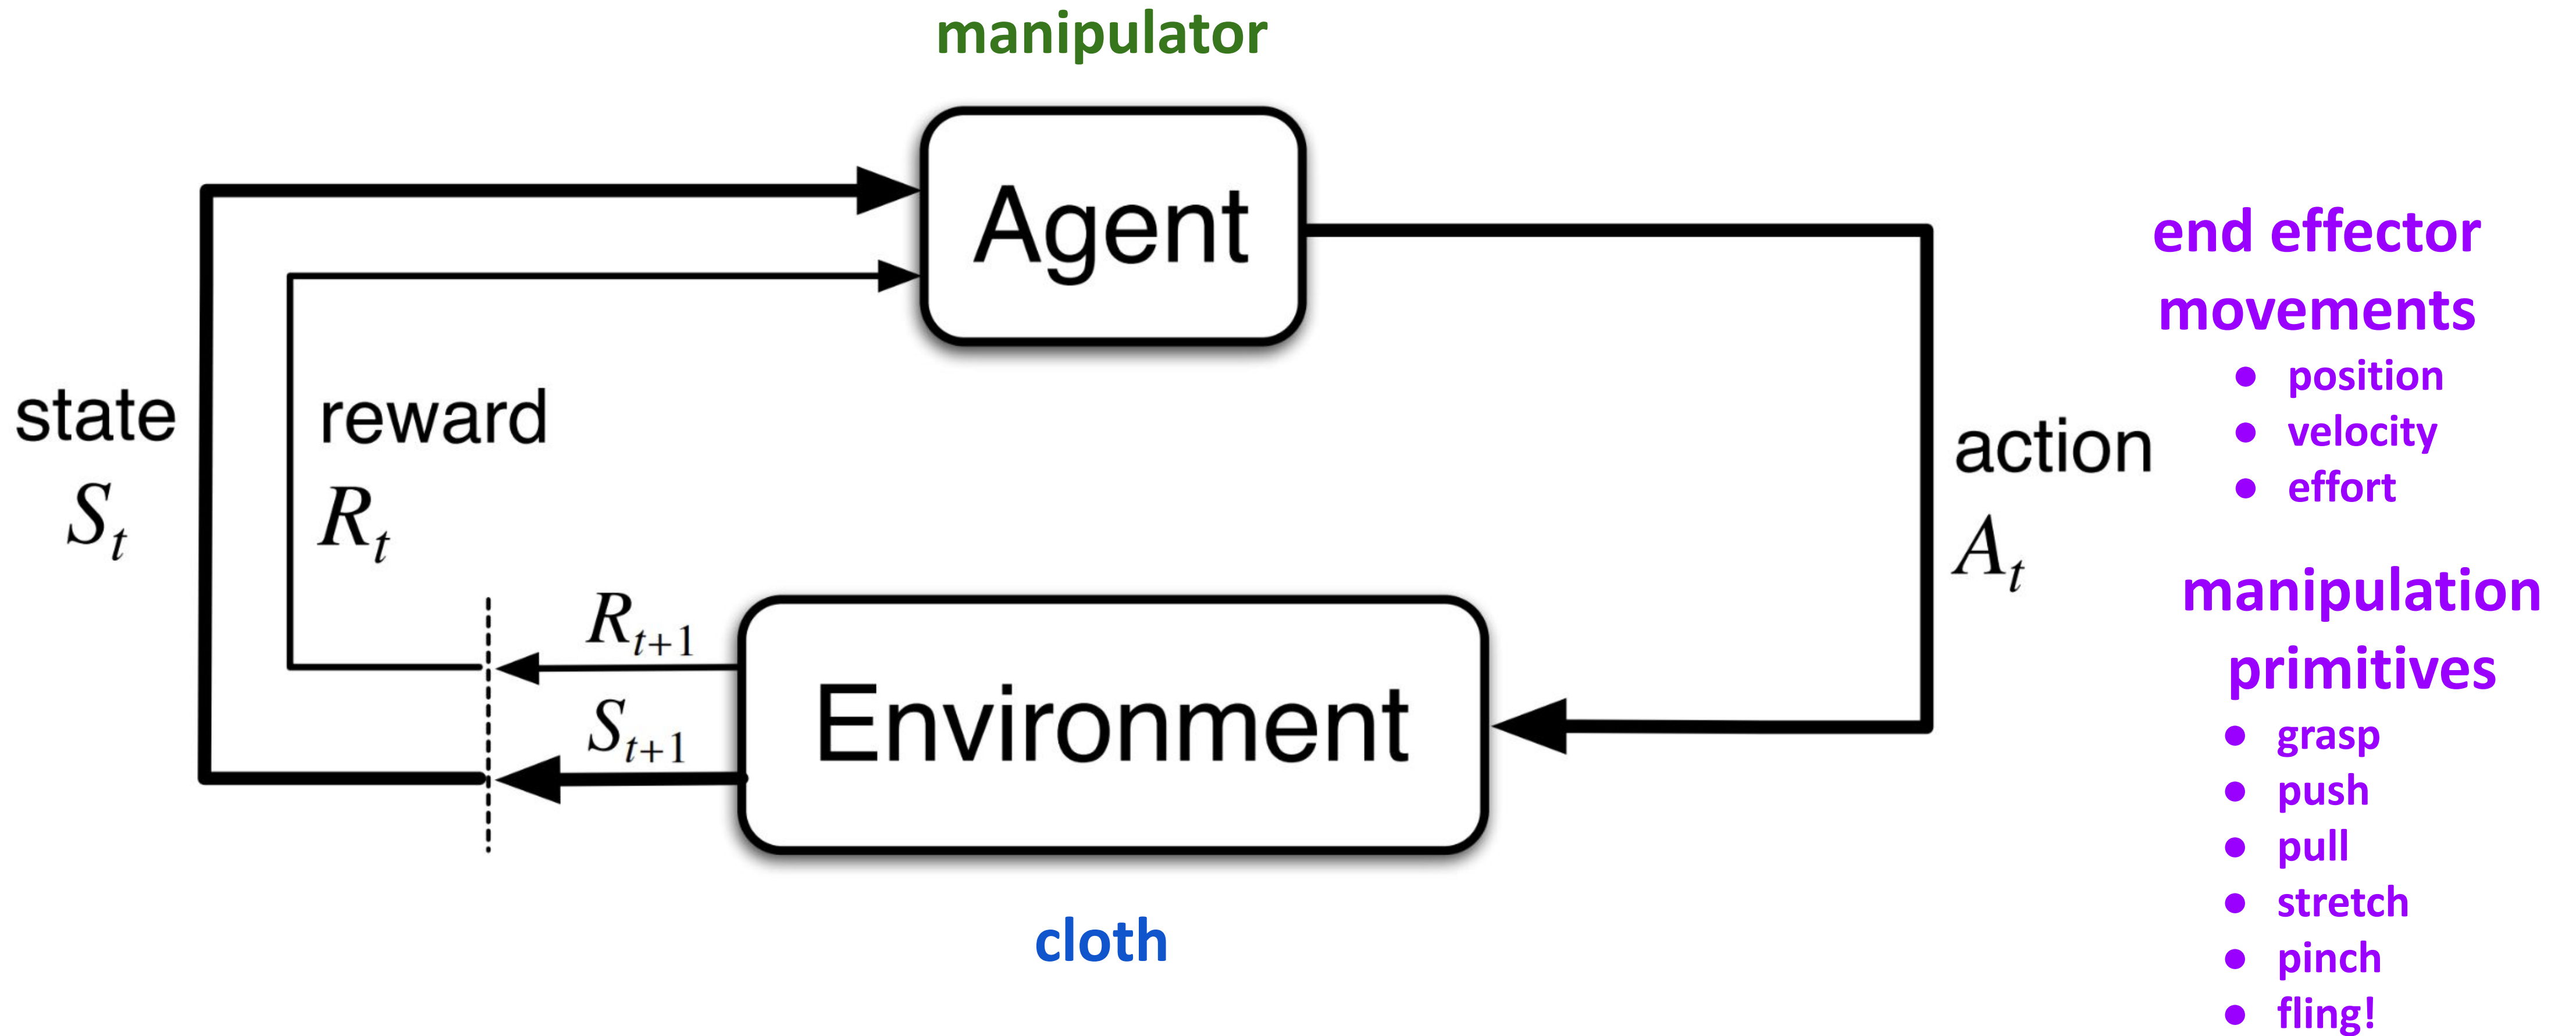


Reinforcement Learning





Reinforcement Learning






Reinforcement Learning

- robot state**
- joint angles
 - gripper status


- cloth state**
- topological
 - spherical

- end effector movements**
- position
 - velocity
 - effort
- manipulation primitives**
- grasp
 - push
 - pull
 - stretch
 - pinch
 - fling!






Example: Topological coordinates





Example: Spherical coordinates





Objective

“learn an action policy that maximizes cumulative reward G_t over time”

$$G_t = R_{t+1} + R_{t+2} + R_{t+3} + \dots = \sum_{k=0}^{\infty} \gamma^k R_{t+k+1}$$

↑
discount factor ($0 < \gamma < 1$)

“immediate rewards have more importance than future rewards”



RL can be computationally expensive and time consuming!





But can I get the data from simulators and use my learned method in real world?






But can I get the data from simulators and use my learned method in real world?

sim2real



Sim-to-Real Reinforcement Learning for Deformable Object Manipulation

Jan Matas, Stephen James, Andrew J. Davidson
Department of Computing
Imperial College London





Contributions

For deformable objects,

- learn manipulation policies through a combination of SOTA DRL algorithms
- learn policies in simulations that can be transferred to real world *without* additional training

Sim-to-Real Reinforcement Learning for Deformable Object Manipulation

Jan Matas

Department of Computing
Imperial College London
jm6214@imperial.ac.uk

Stephen James

Department of Computing
Imperial College London
slj12@imperial.ac.uk

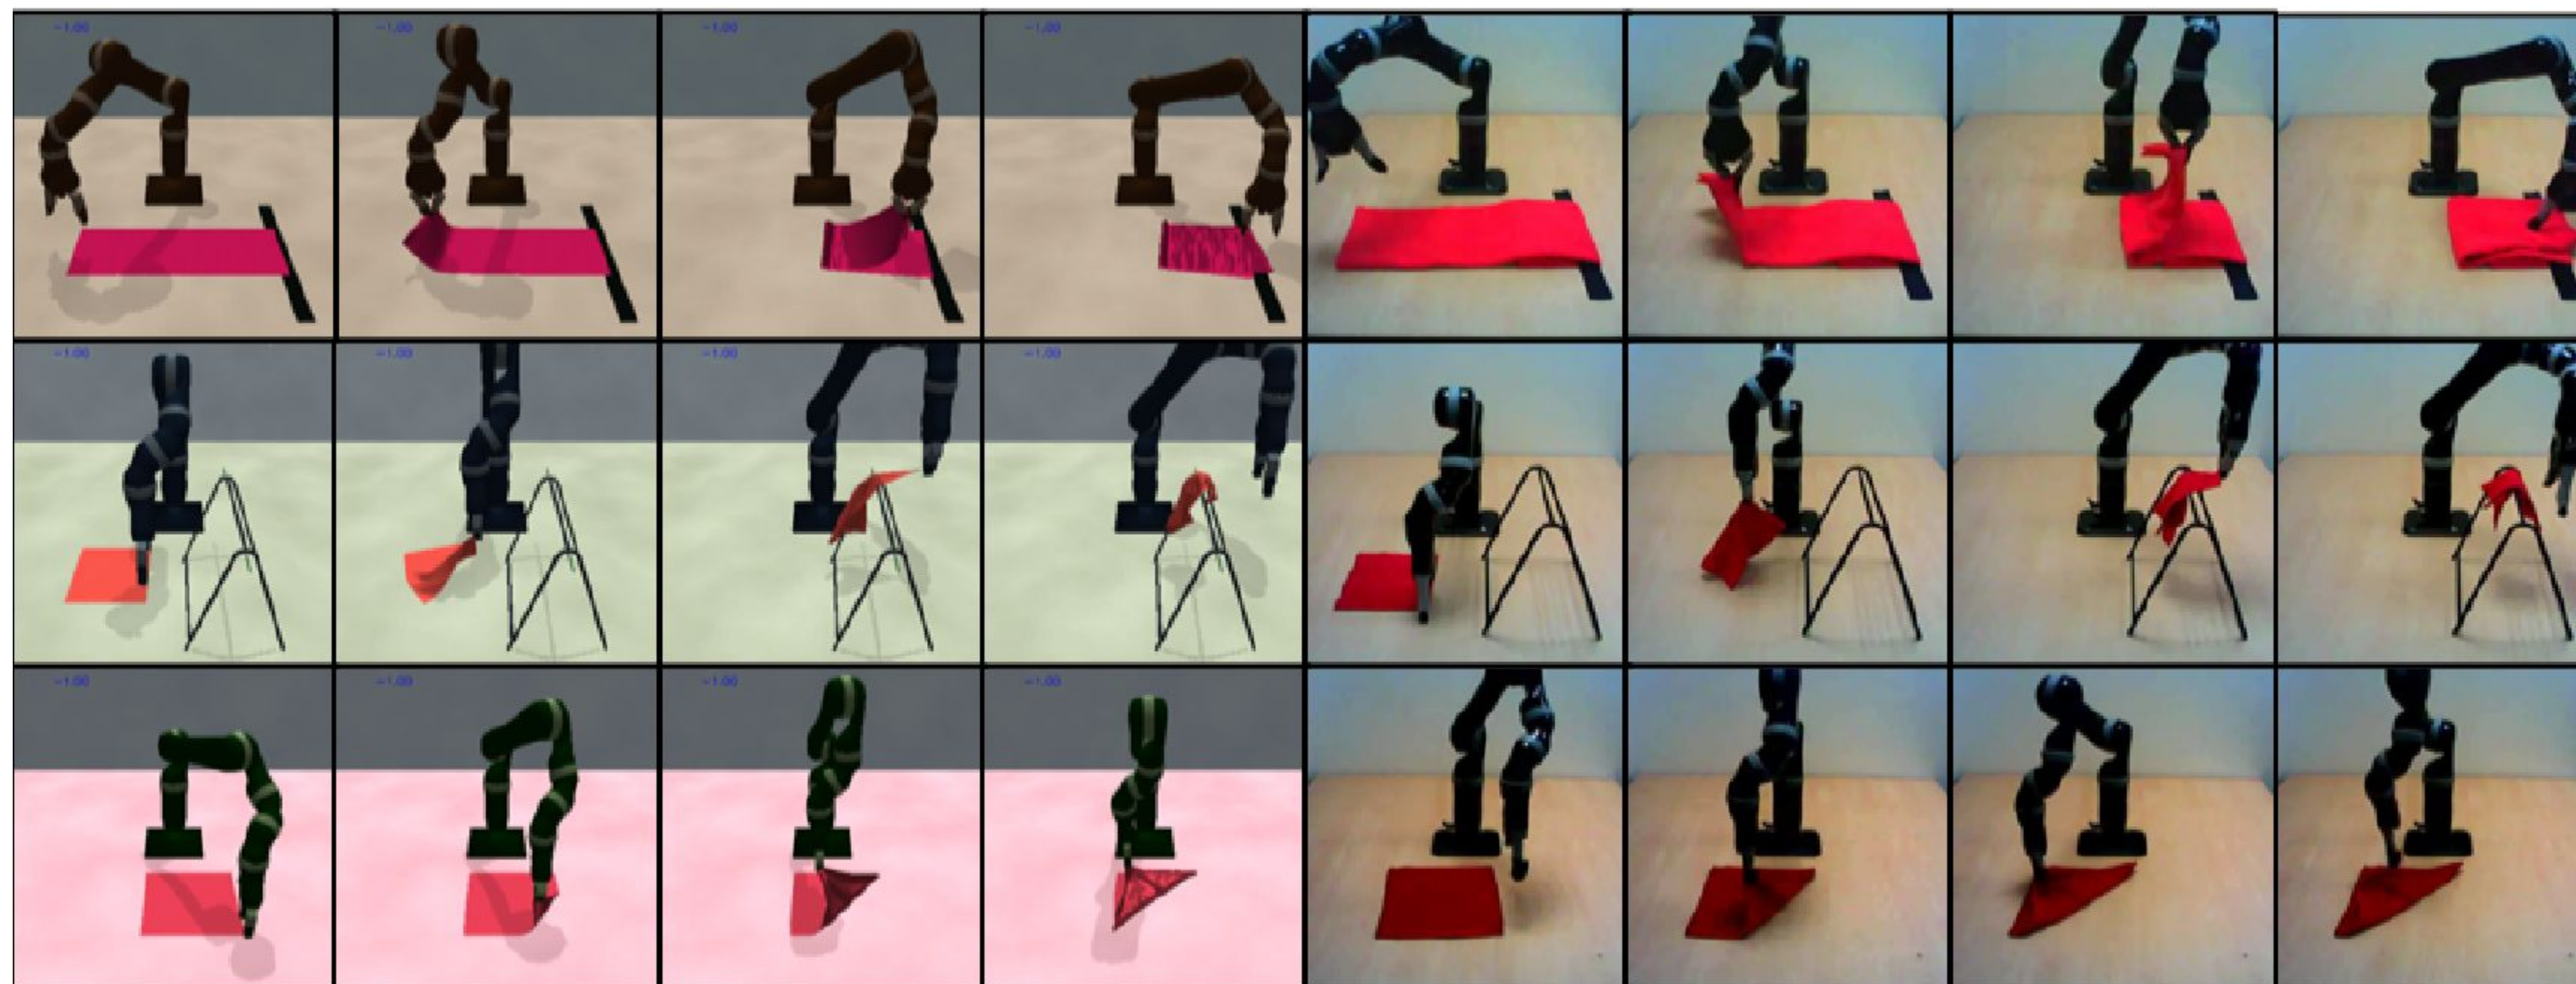
Andrew J. Davison

Department of Computing
Imperial College London
a.davison@imperial.ac.uk

Abstract: We have seen much recent progress in rigid object manipulation, but interaction with deformable objects has notably lagged behind. Due to the large configuration space of deformable objects, solutions using traditional modelling approaches require significant engineering work. Perhaps then, bypassing the need for explicit modelling and instead learning the control in an end-to-end manner serves as a better approach? Despite the growing interest in the use of end-to-end robot learning approaches, only a small amount of work has focused on their applicability to deformable object manipulation. Moreover, due to the large amount of data needed to learn these end-to-end solutions, an emerging trend is to learn control policies in simulation and then transfer them over to the real world. To-date, no work has explored whether it is possible to learn and transfer deformable object policies. We believe that if sim-to-real methods are to be employed further, then it should be possible to learn to interact with a wide variety of objects, and not only rigid objects. In this work, we use a combination of state-of-the-art deep reinforcement learning algorithms to solve the problem of manipulating deformable objects (specifically cloth). We evaluate our approach on three tasks — folding a towel up to a mark, folding a face towel diagonally, and draping a piece of cloth over a hanger. Our agents are fully trained in simulation with domain randomisation, and then successfully deployed in the real world without having seen any real deformable objects.



Fold until the tape




Hang towel on hanger

Diagonal cloth fold



Network Architecture

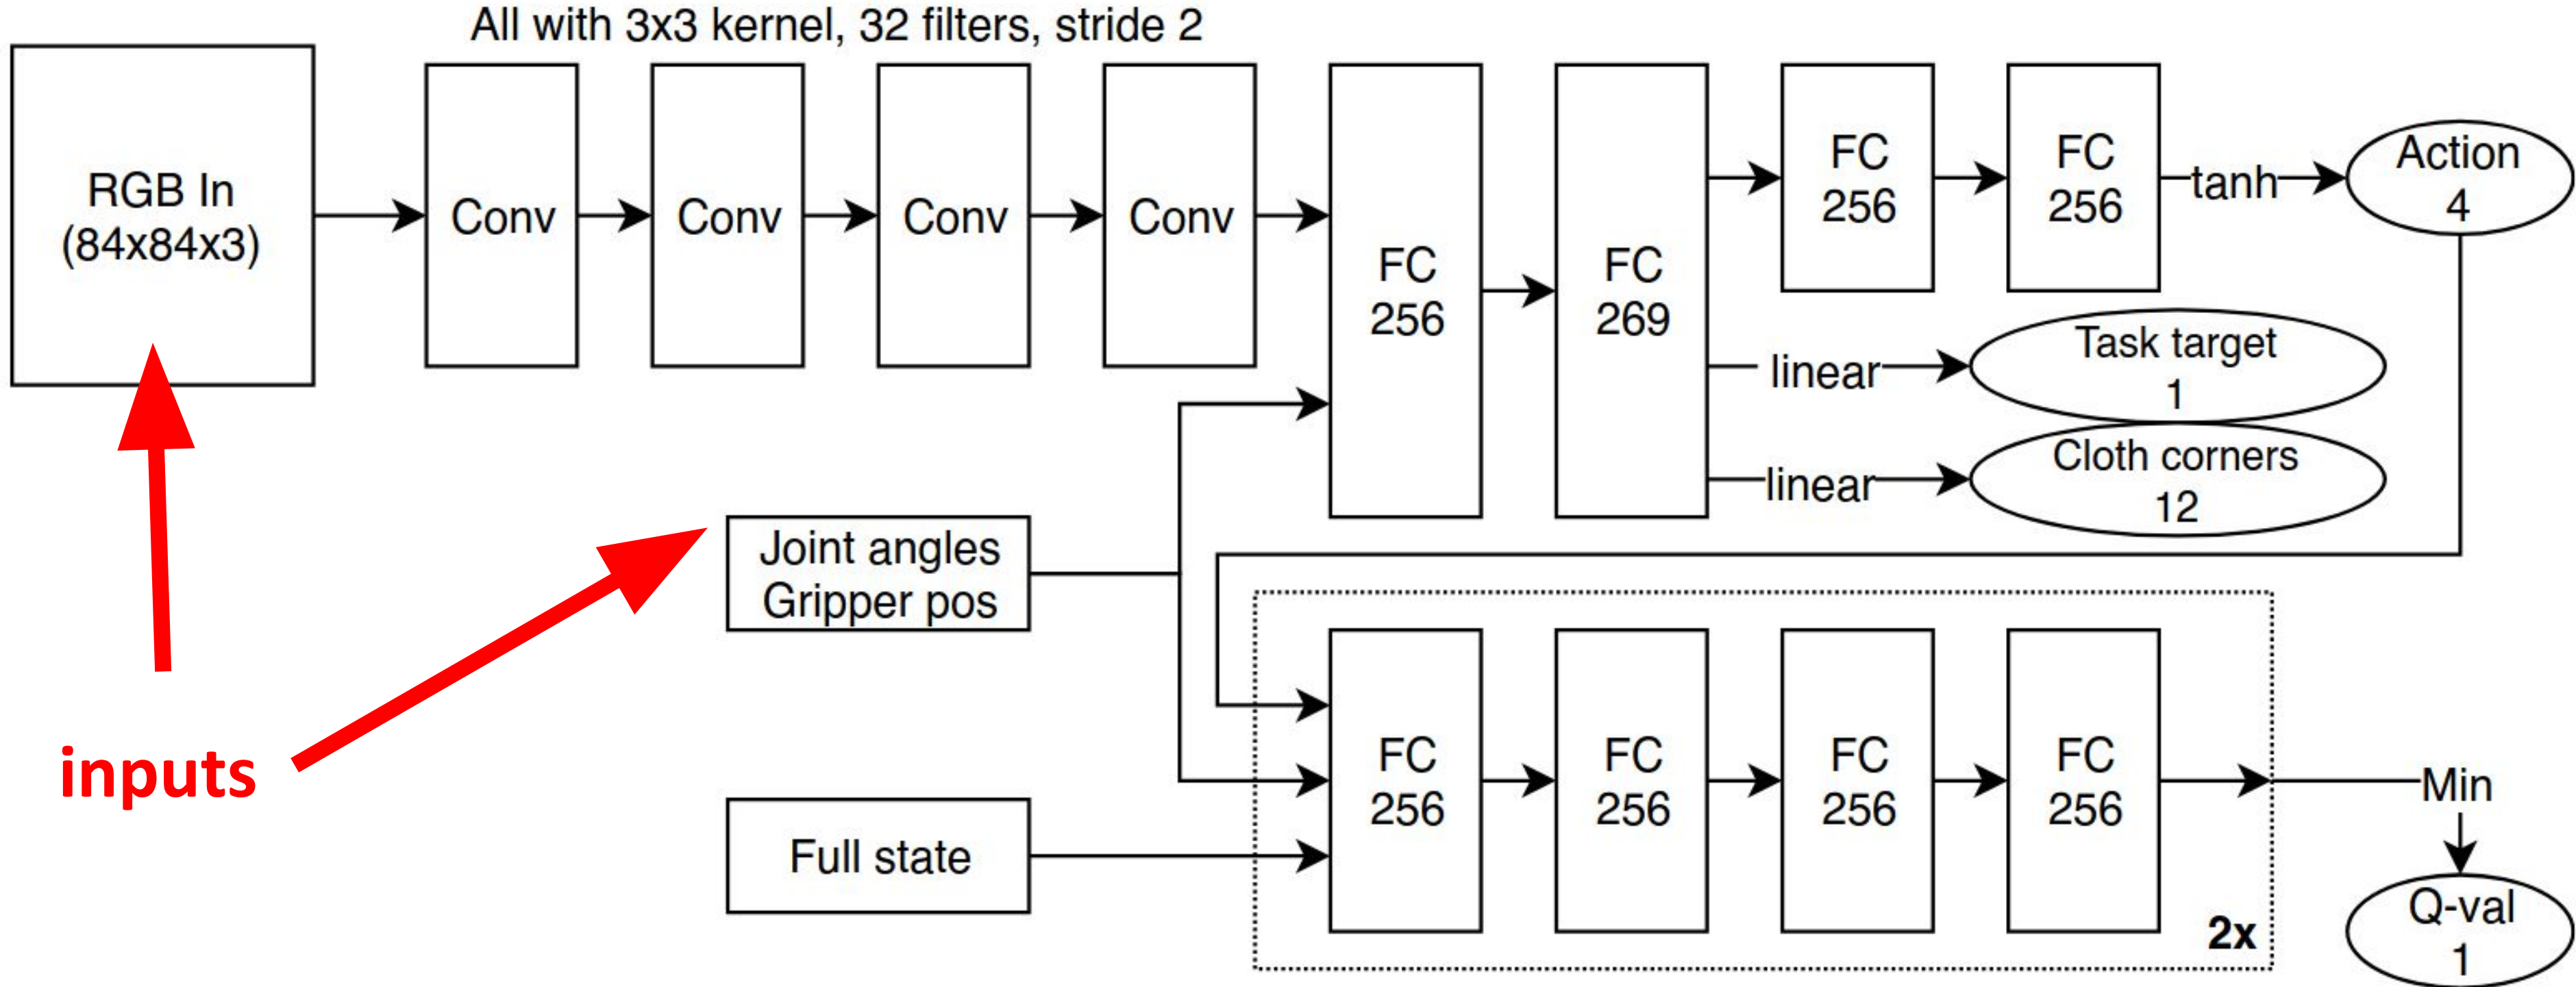


Tamei, T., Matsubara, T., Rai, A., & Shibata, T. (2011). Reinforcement learning of clothing assistance with a dual-arm robot. 2011 11th IEEE-RAS International Conference on Humanoid Robots, 733-738. doi:10.1109/Humanoids.2011.6100915





Network Architecture



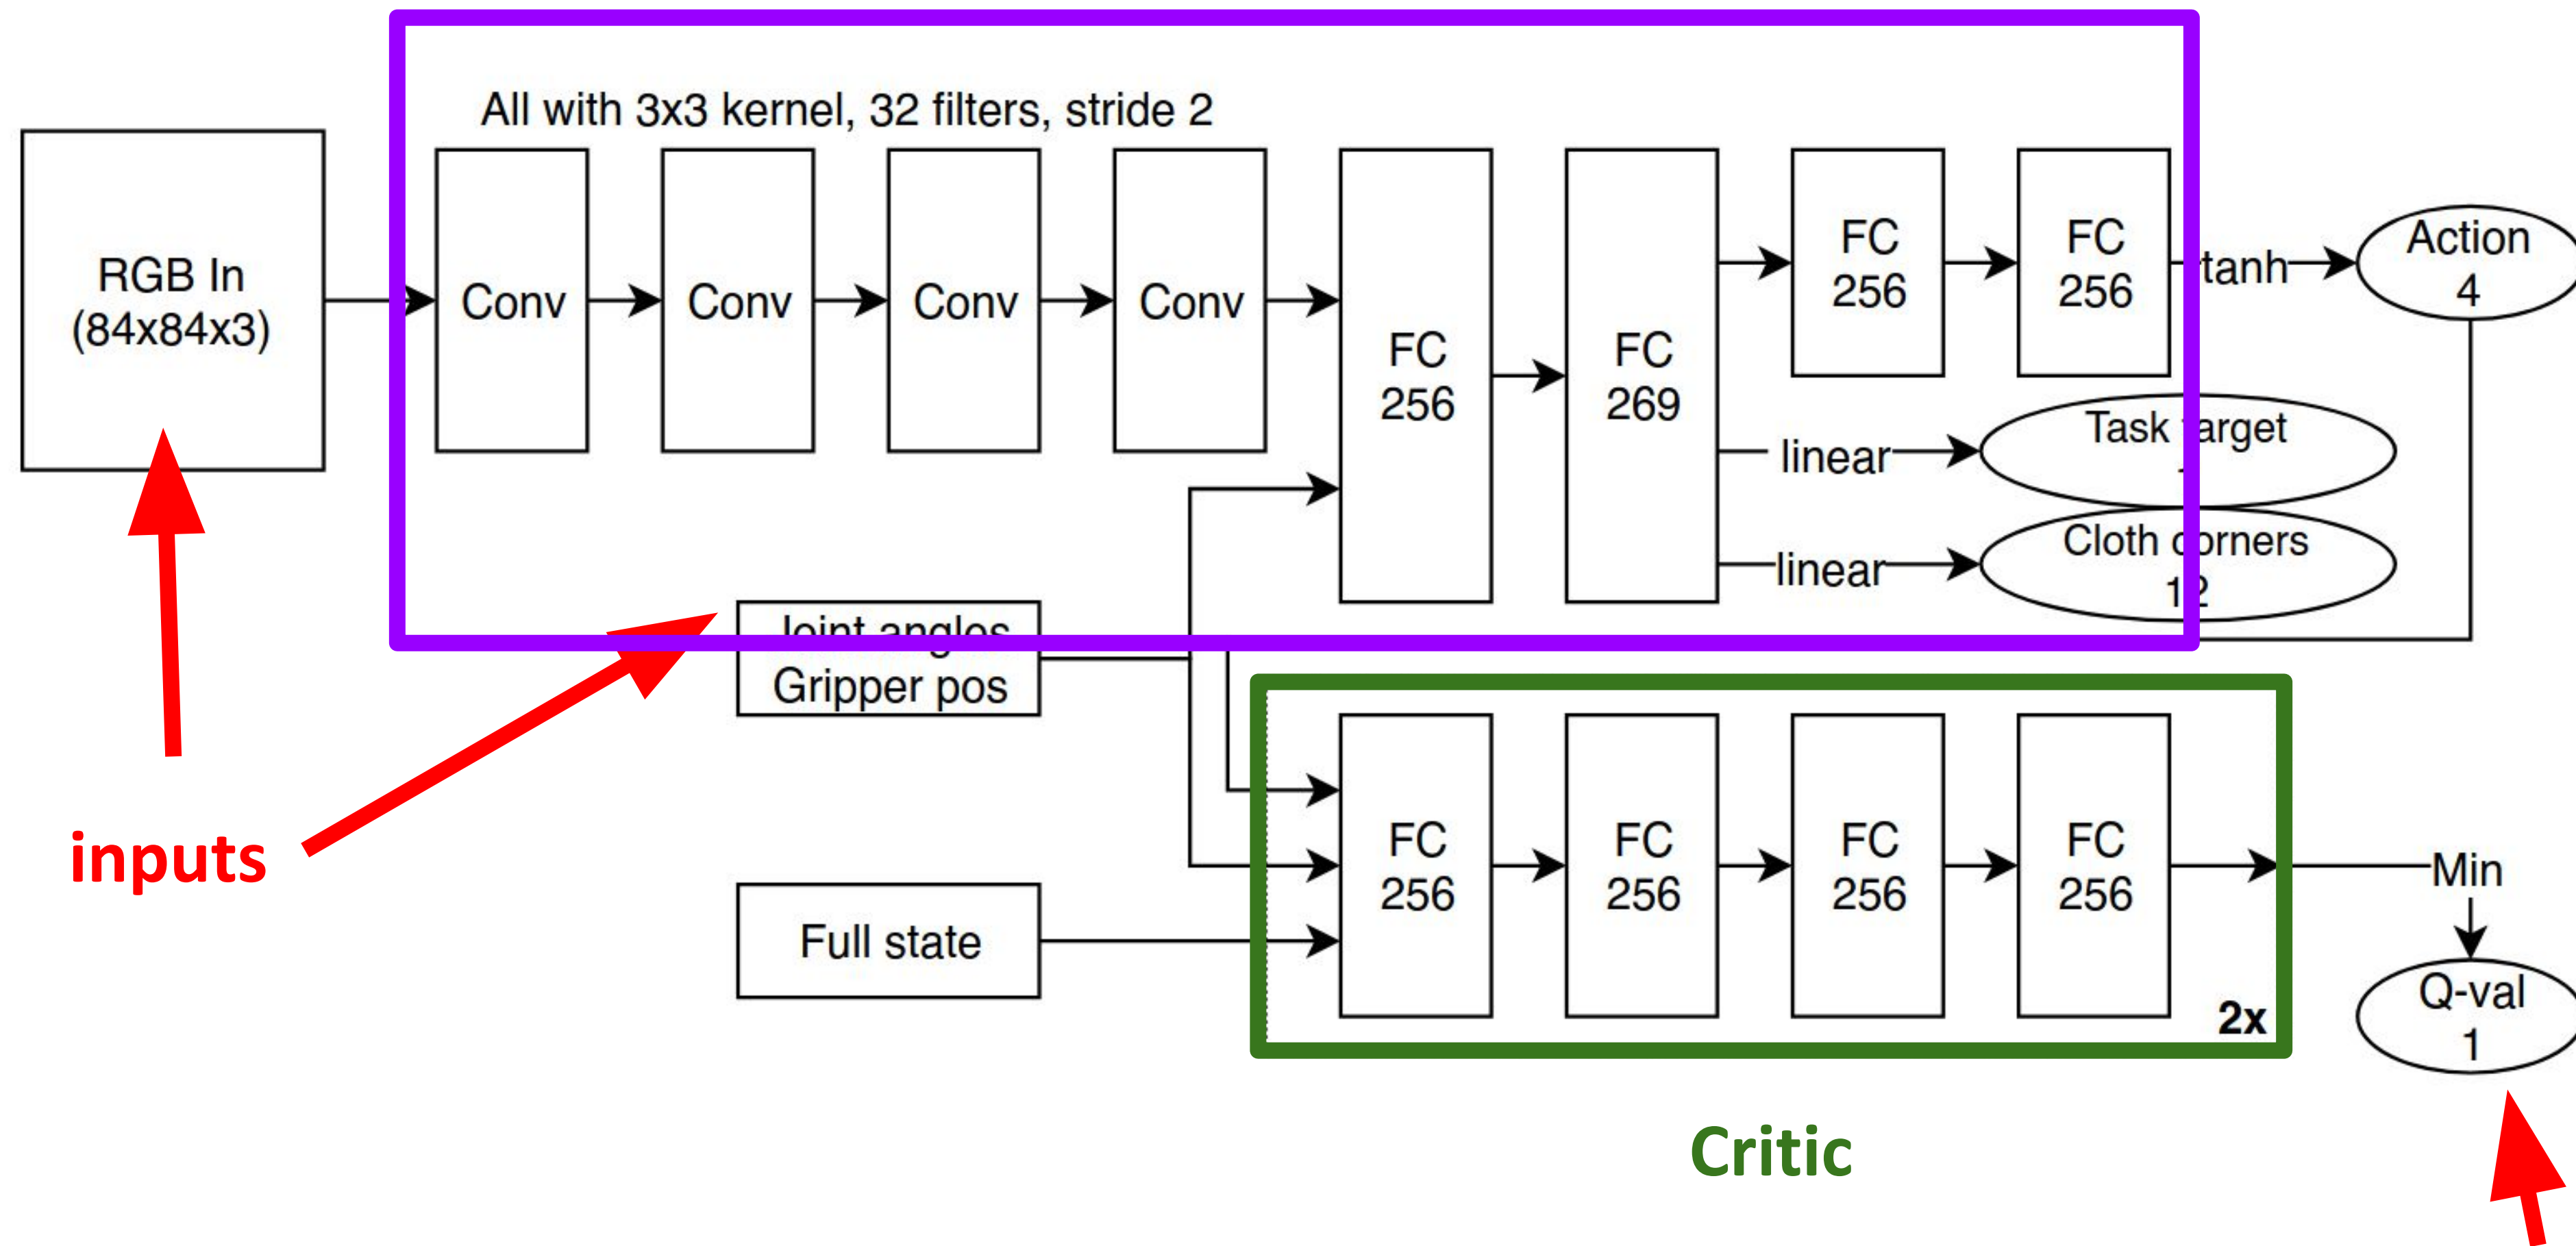
inputs





Network Architecture

Actor



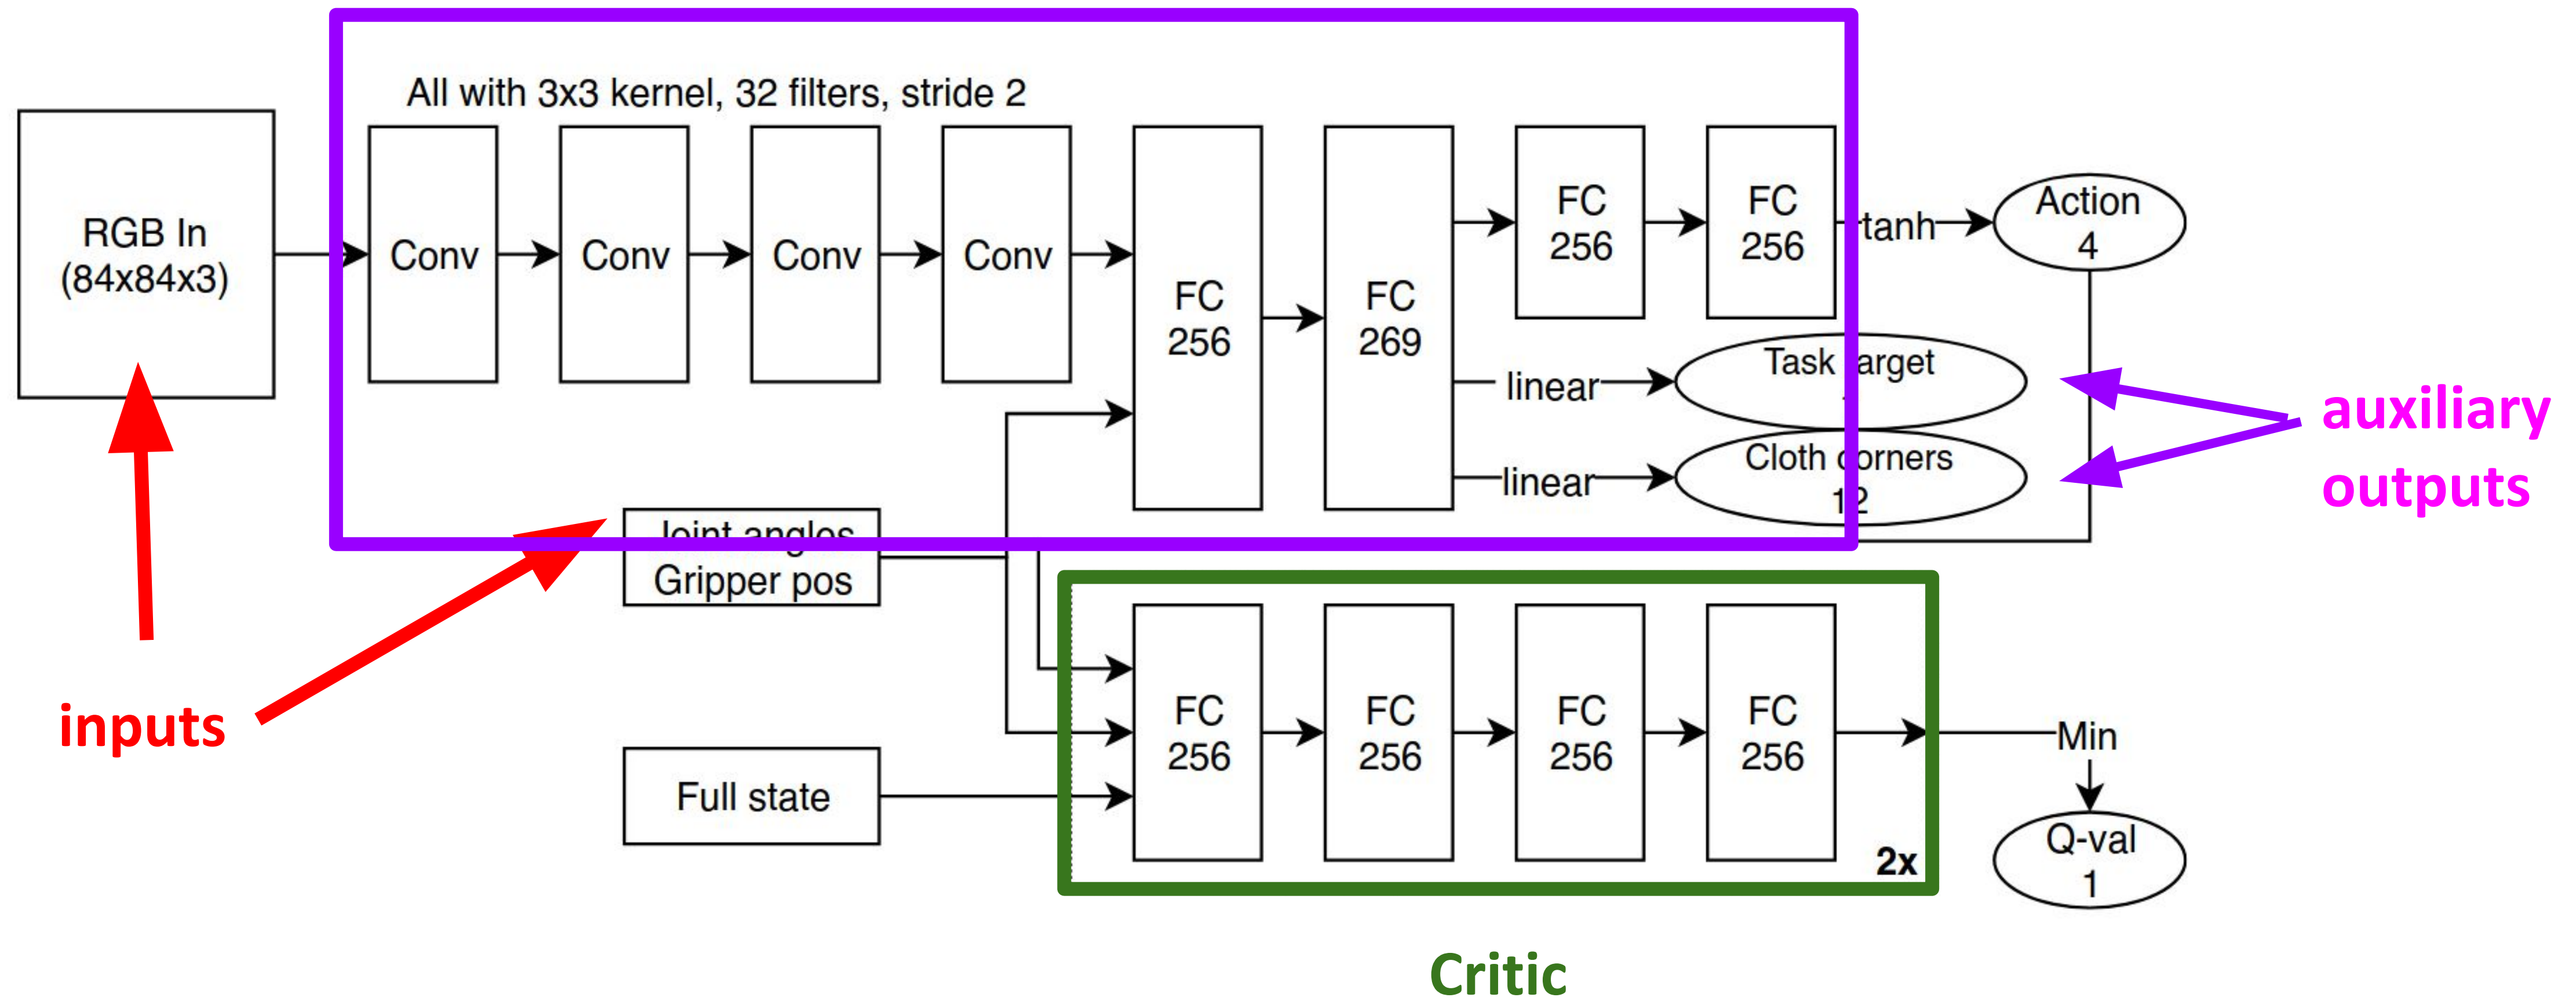
Critic





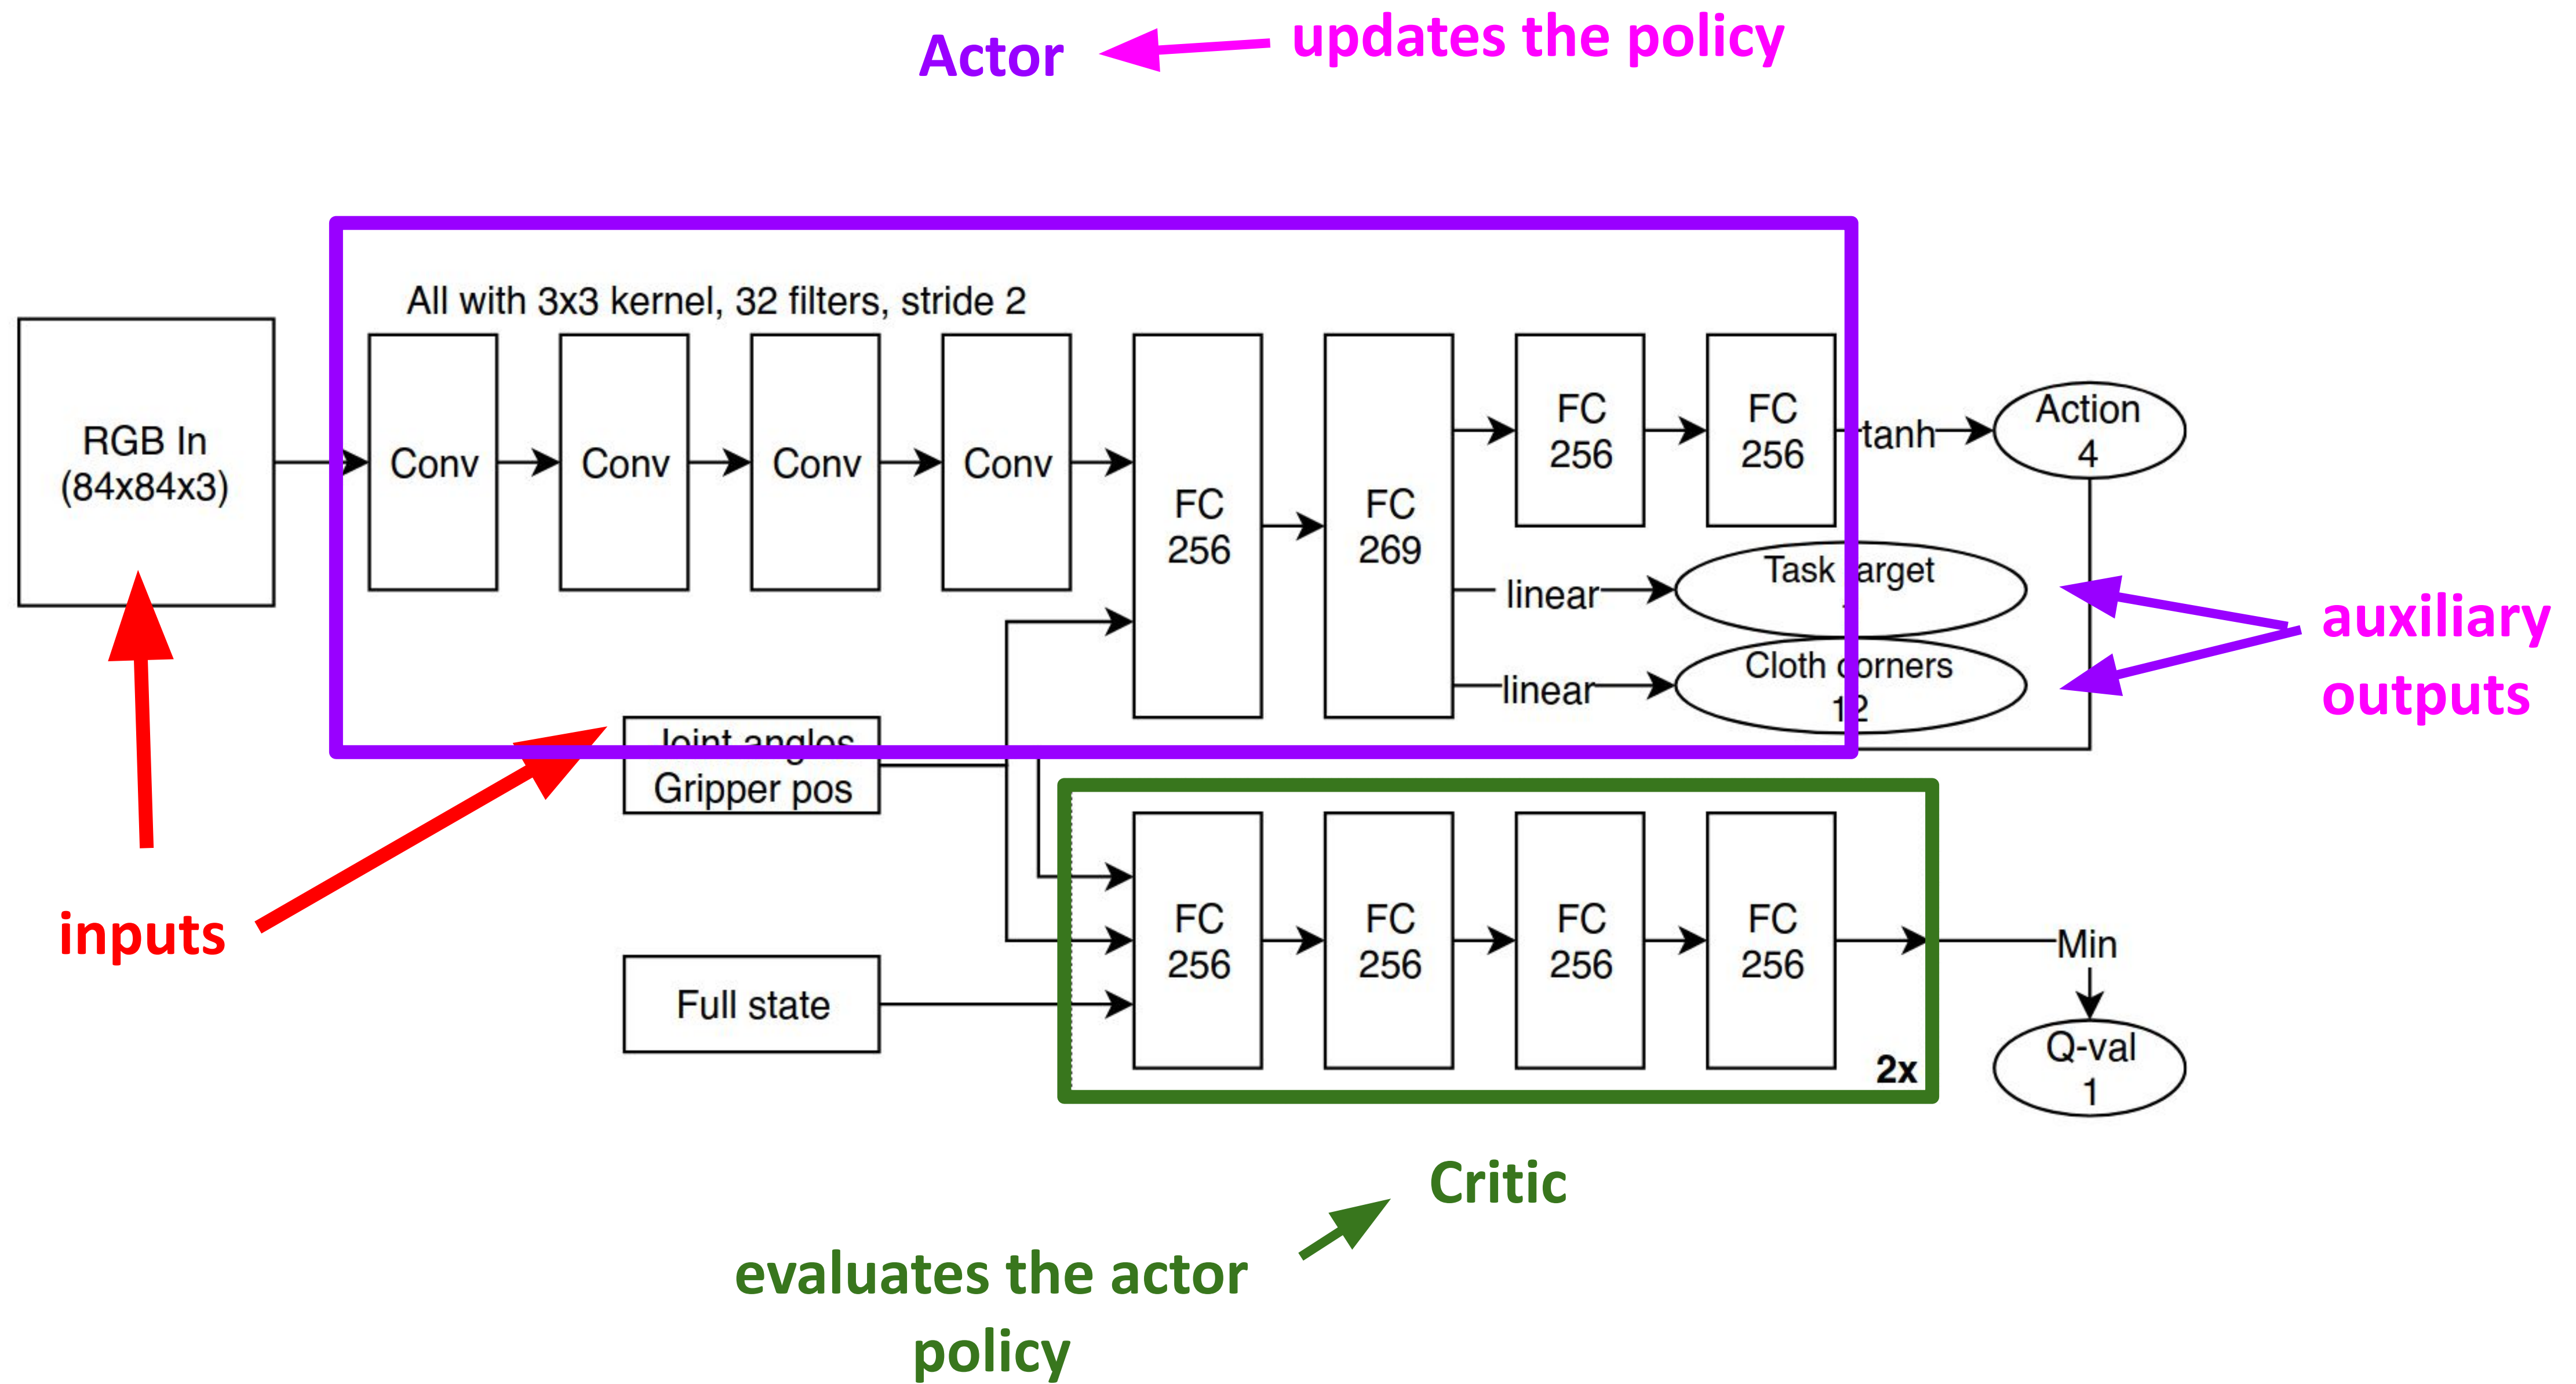
Network Architecture

Actor ← updates the policy






Network Architecture





Domain Randomization



randomized attributes:

- table textures
- cloth and arm colours
- light position
- camera position and orientation,
- cloth size and position,
- hanger size and position,
- initial arm position and size of arm base





sim



real



Policy trained in simulation transfers to real world **without further training**





Results

Simulations

Success rates (Sim)	
Diagonal folding	90%
Hanging	77%
Tape	86%

Real world

Hanging task	
Vicinity	100%
Grasp	76.6%
Drape over	70%
Full success	46.6%

Diagonal folding task	
Grasp	66.6%
Not crumpled	66.6%
$d \leq 0.15m$	53.3%
$d \leq 0.1m$	40%
$d \leq 0.05m$	20%

Tape folding task	
Grasp	90%
$d \leq 0.15m$	90%
$d \leq 0.1m$	76.6%
$d \leq 0.05m$	43%


- vicinity:** gripper being from 5 cm from the cloth
- drape over:** cloth touching top part of the hanger
- full success:** cloth does not fall after released
- not crumpled:** adjacent corners are more than 15 cms from each other
- d:** distance between the diagonal (folded) corners (lower the better)




Failure Modes

Common **failure** modes are

Grasp above/bellow the towel



Crumpling the towel



Weak Grasp





Open Problems

- 1) High Fidelity Simulation
- 2) Sim2Real gap
- 3) Robust Perception in Dynamic environments
- 4) Multi stage manipulation
- 5) Dataset and Benchmark standardization





Datasets

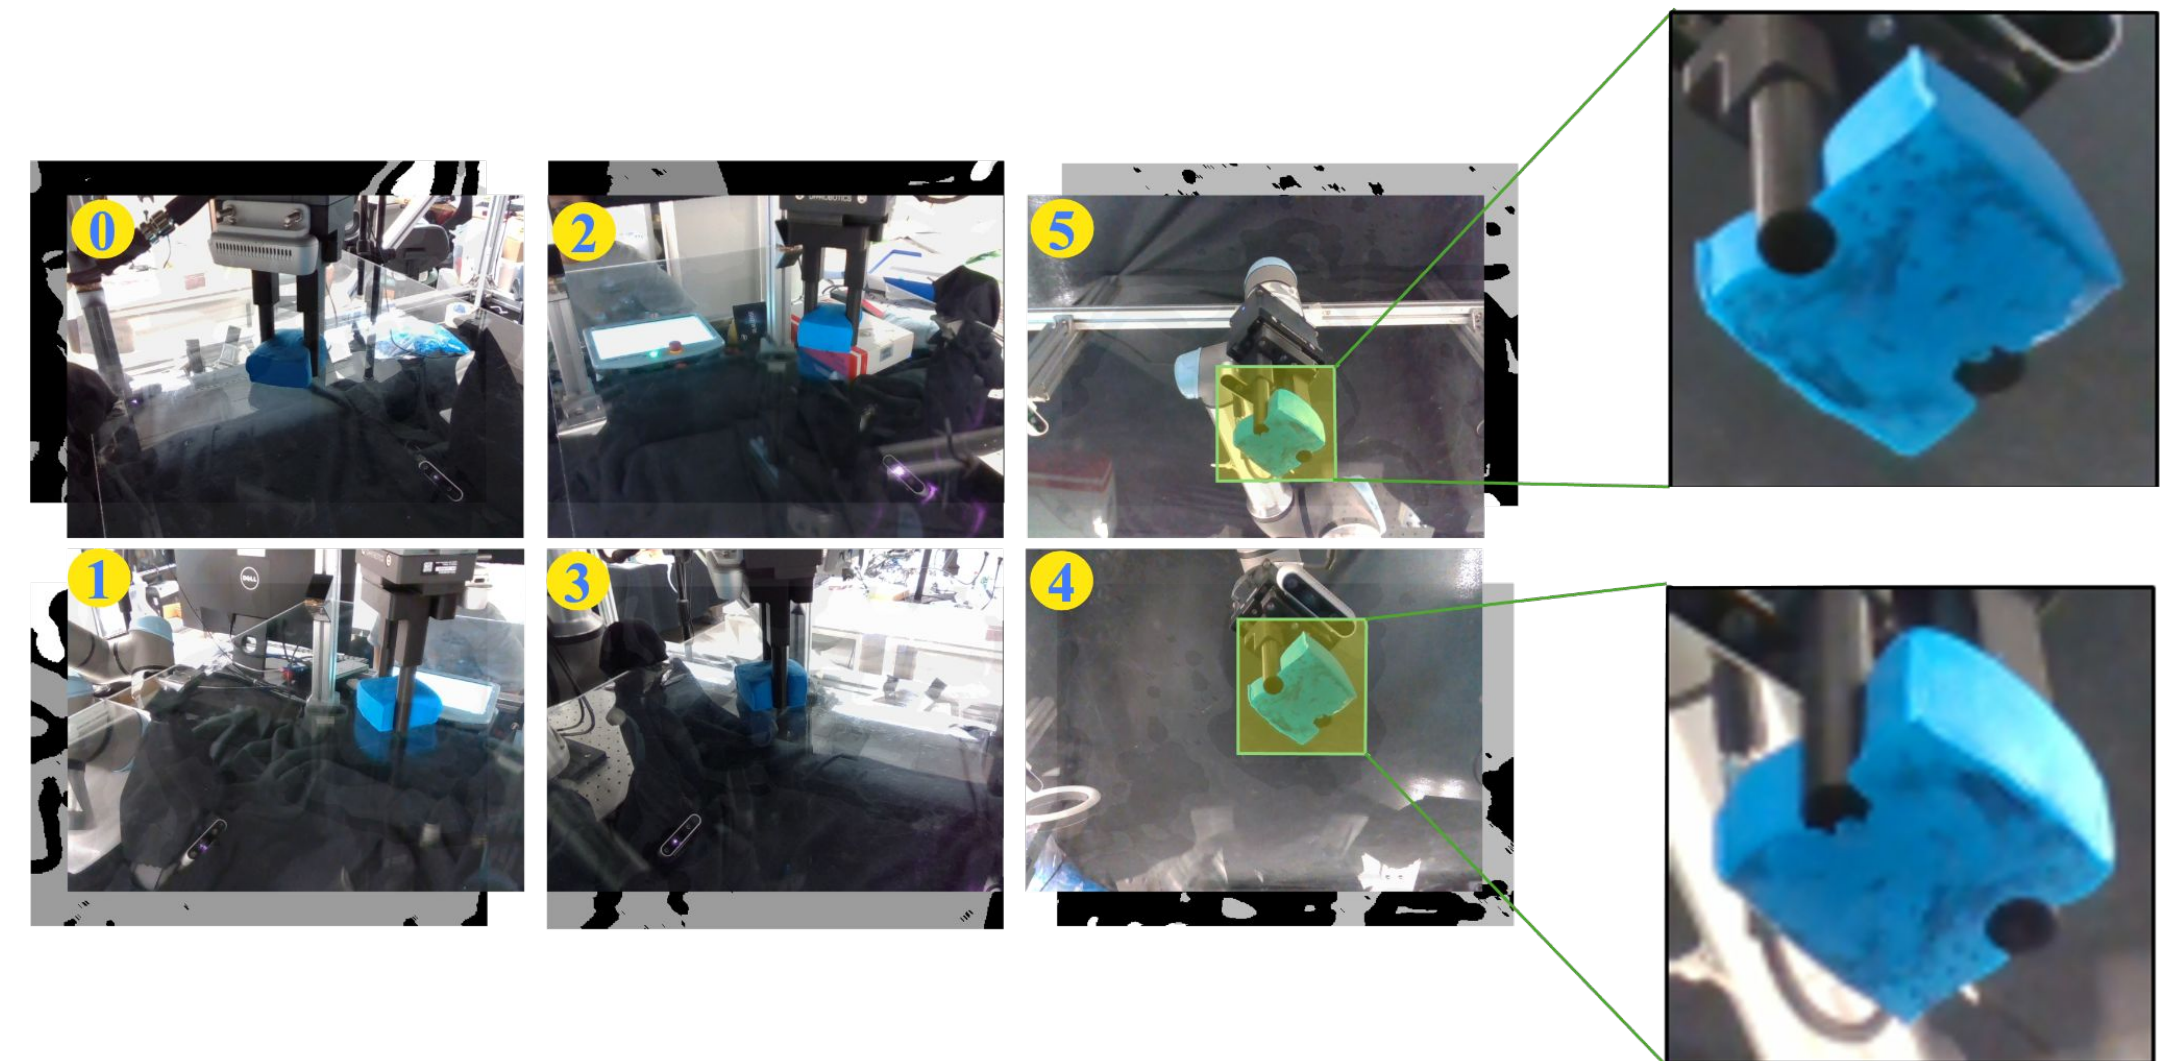


[1] Cloth3D



[2] PokeFlex

[3] DOFS



[1] Liu Z, Luo P, Qiu S, Wang X, Tang X. 2016. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In 2016 IEEE Conference on Computer Vision and Pattern Recognition, pp. 1096–104. Piscataway, NJ: IEEE

[2] Obrist, J., Zamora, M., Zheng, H., Zarate, J., Katzschmann, R. K., & Coros, S. (2024). PokeFlex: Towards a Real-World Dataset of Deformable Objects for Robotic Manipulation. arXiv preprint arXiv:2409.17124.

[3] Zhang, Z., Chu, X., Yunxi, T., & Au, K. W. S. (2024). DOFS: A Real-world 3D Deformable Object Dataset with Full Spatial Information for Dynamics Model Learning. CoRL Workshop on Learning Robot Fine and Dexterous Manipulation: Perception and Control.

Retrieved from <https://openreview.net/forum?id=QADznDIGM4>





Thank you!



Next Lecture:
Student Lecture 5
**Multisensory and Multimodal
Learning + Manipulation**



DeepRob

[Group 1] Lecture 04

Pranay, Aditya, Siddharth

Deformable Object Manipulation
University of Minnesota

