

The Promise and Perils of AI: Part I

Introduction to Machine Intelligence

Patrick M. Pilarski

*Canada Research Chair in Machine Intelligence for Rehabilitation
Principal Investigator, Alberta Innovates Centre for Machine Learning
Assistant Professor, Department of Medicine*



Learning Objectives

- Be able to **define** and **discuss** machine intelligence.
- Be able to state **why** machine intelligence is important to society.
- Understand **what** machines might learn (representation, prediction, and control learning.)
- Understand **how** machines can learn about their world. (three learning approaches.)
- Understand **when** machines might learn. (online/real-time vs offline learning.)

Cheat-sheet: P.M. Pilarski, *Alberta ICT Magazine*, 2nd Ed., pp. 31
<http://www.ualberta.ca/~pilarski/docs/papers/Pilarski-Learning-AlbertaICTMagazine2012.pdf>





P.M. Pilarski, 2008. Dettifoss, Iceland.



P.M. Pilarski, 2007. Iguazu Falls, Argentina.

KEY IDEA

Data is now ubiquitous; it flows between connected systems at high volume and with great diversity.

Why Machine Intelligence?

- **Enhanced control** over a changing and increasingly complex world.
- **Anticipation** of future events and outcomes.
- **General tools** for solving hard problems.
- “Optimizing the control of complex systems and extracting knowledge from massive amounts of data.”
- Examples: finance, healthcare, energy, resources, transport, information processing.

Alternate Identities

- **Artificial Intelligence:**
does it need to learn?

- **Machine Learning:**
is it truly intelligent?

- **Pattern Recognition and Analysis:**
are they more than just
deterministic processes?



Intelligent Systems: One Possible Definition

- A system that can:

Perceive and Represent its world.

Predict its world.

Control its world.

- “The **Pursuance of future ends** and the **choice of means** for their attainment, are thus the mark and criterion of the presence of mentality in a phenomenon” (James, 1890)
- **Purposeful**: to have, seek & achieve goals (Sutton, 2001).


KEY IDEA

Intelligence revolves around
maintaining and using knowledge
(representation, prediction, control)
in a purposeful way.

Why Learning?


End

- **Things are Unknown:**
known ends but unclear means.
- **Things are Complex:**
scaling up is demanding or impossible.
- **Things Change:**
systems need to adapt!




Why Learning?

- **Things are Unknown:**
known ends but unclear means.
- **Things are Complex:**
scaling up is demanding or impossible.
- **Things Change:**
systems need to adapt!



Why Learning?

- **Things are Unknown:**
known ends but unclear means.
- **Things are Complex:**
scaling up is demanding or impossible.
- **Things Change:**
systems need to adapt!



KEY IDEA


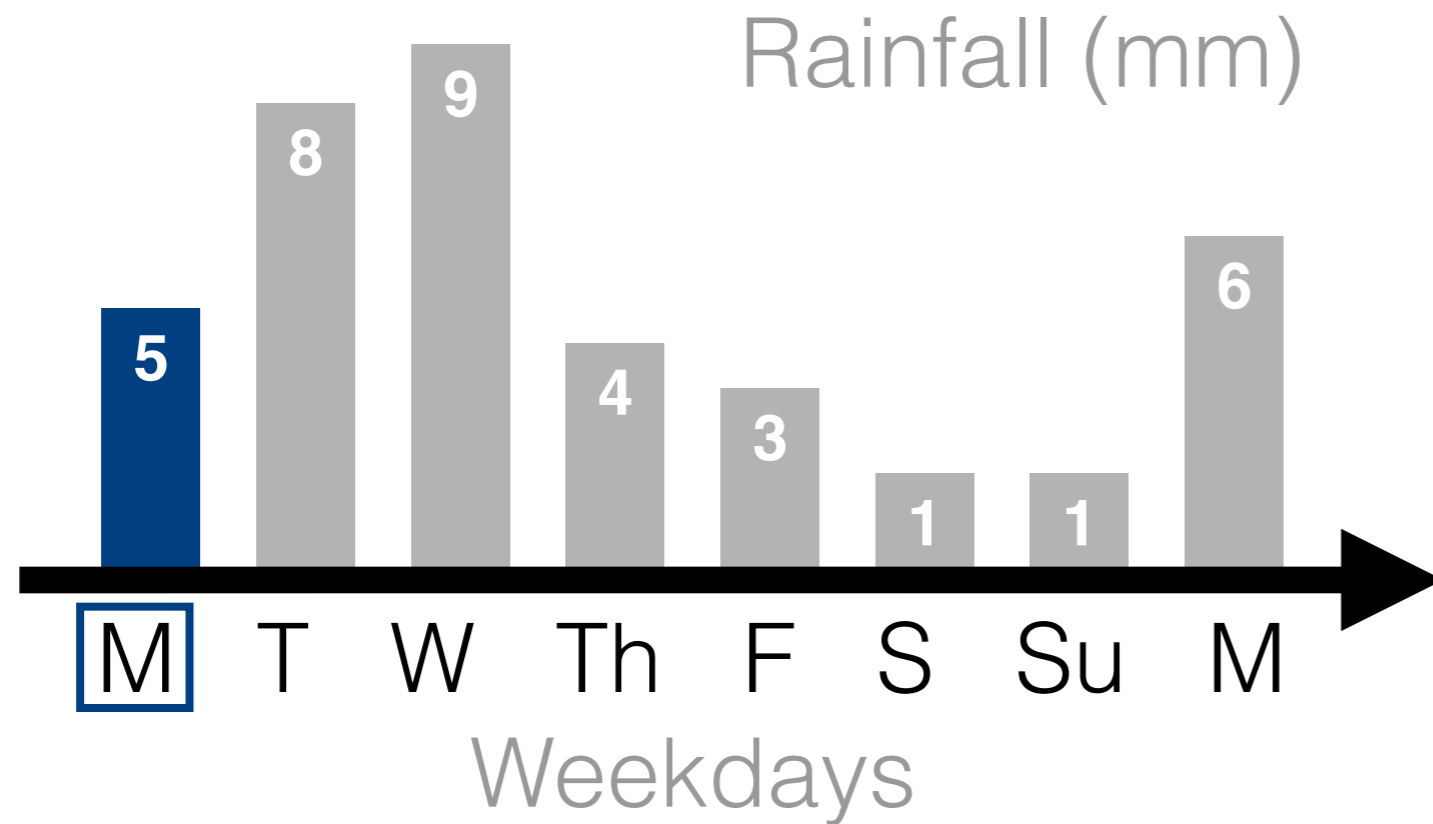
Our ability to directly engineer an intelligent system no longer scales up to our goals or to the complexity of the digital world.

What to Learn



- **Prediction Learning:** building up knowledge.
- **Control Learning:** using knowledge to act.
- **Representation Learning:** structuring knowledge.

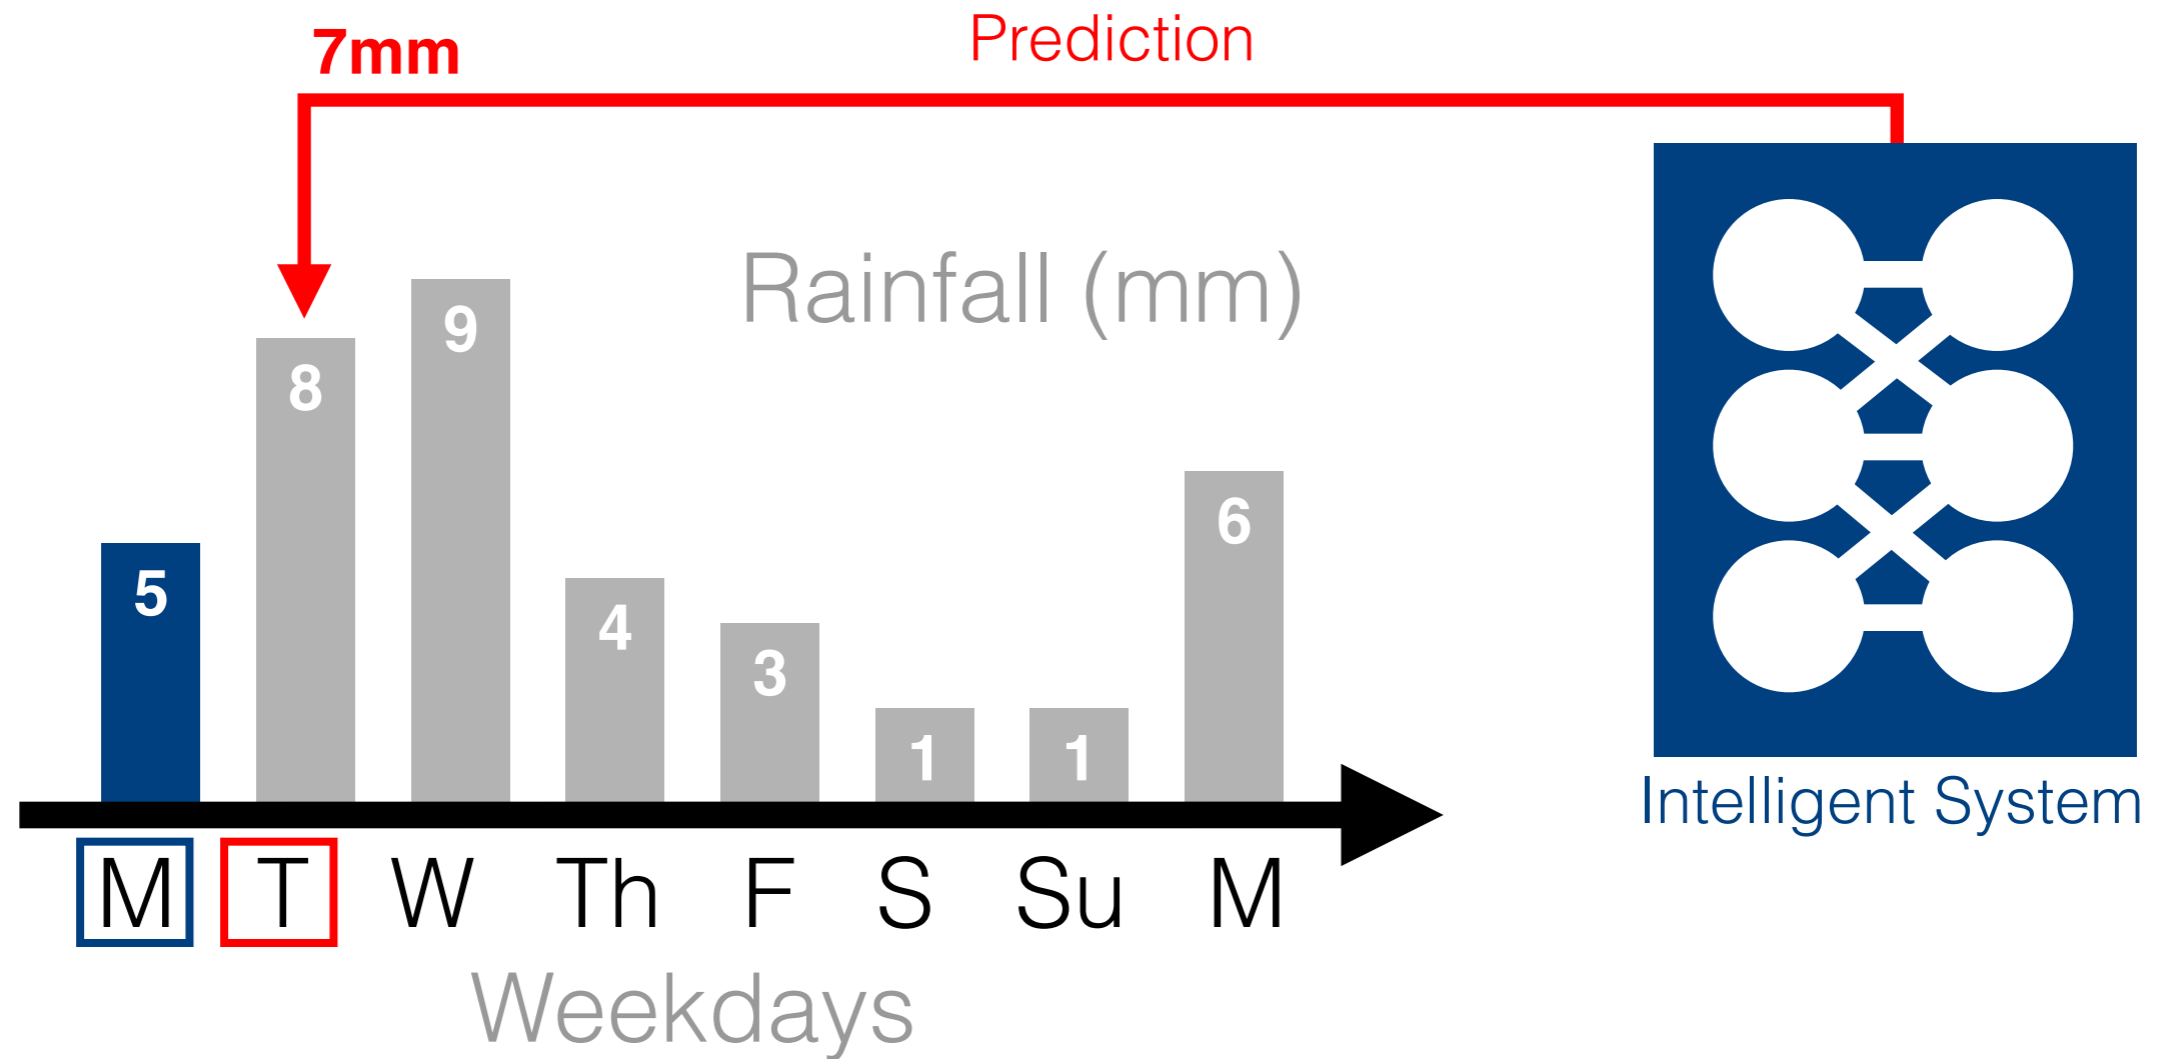
Prediction Learning



One time-step prediction.



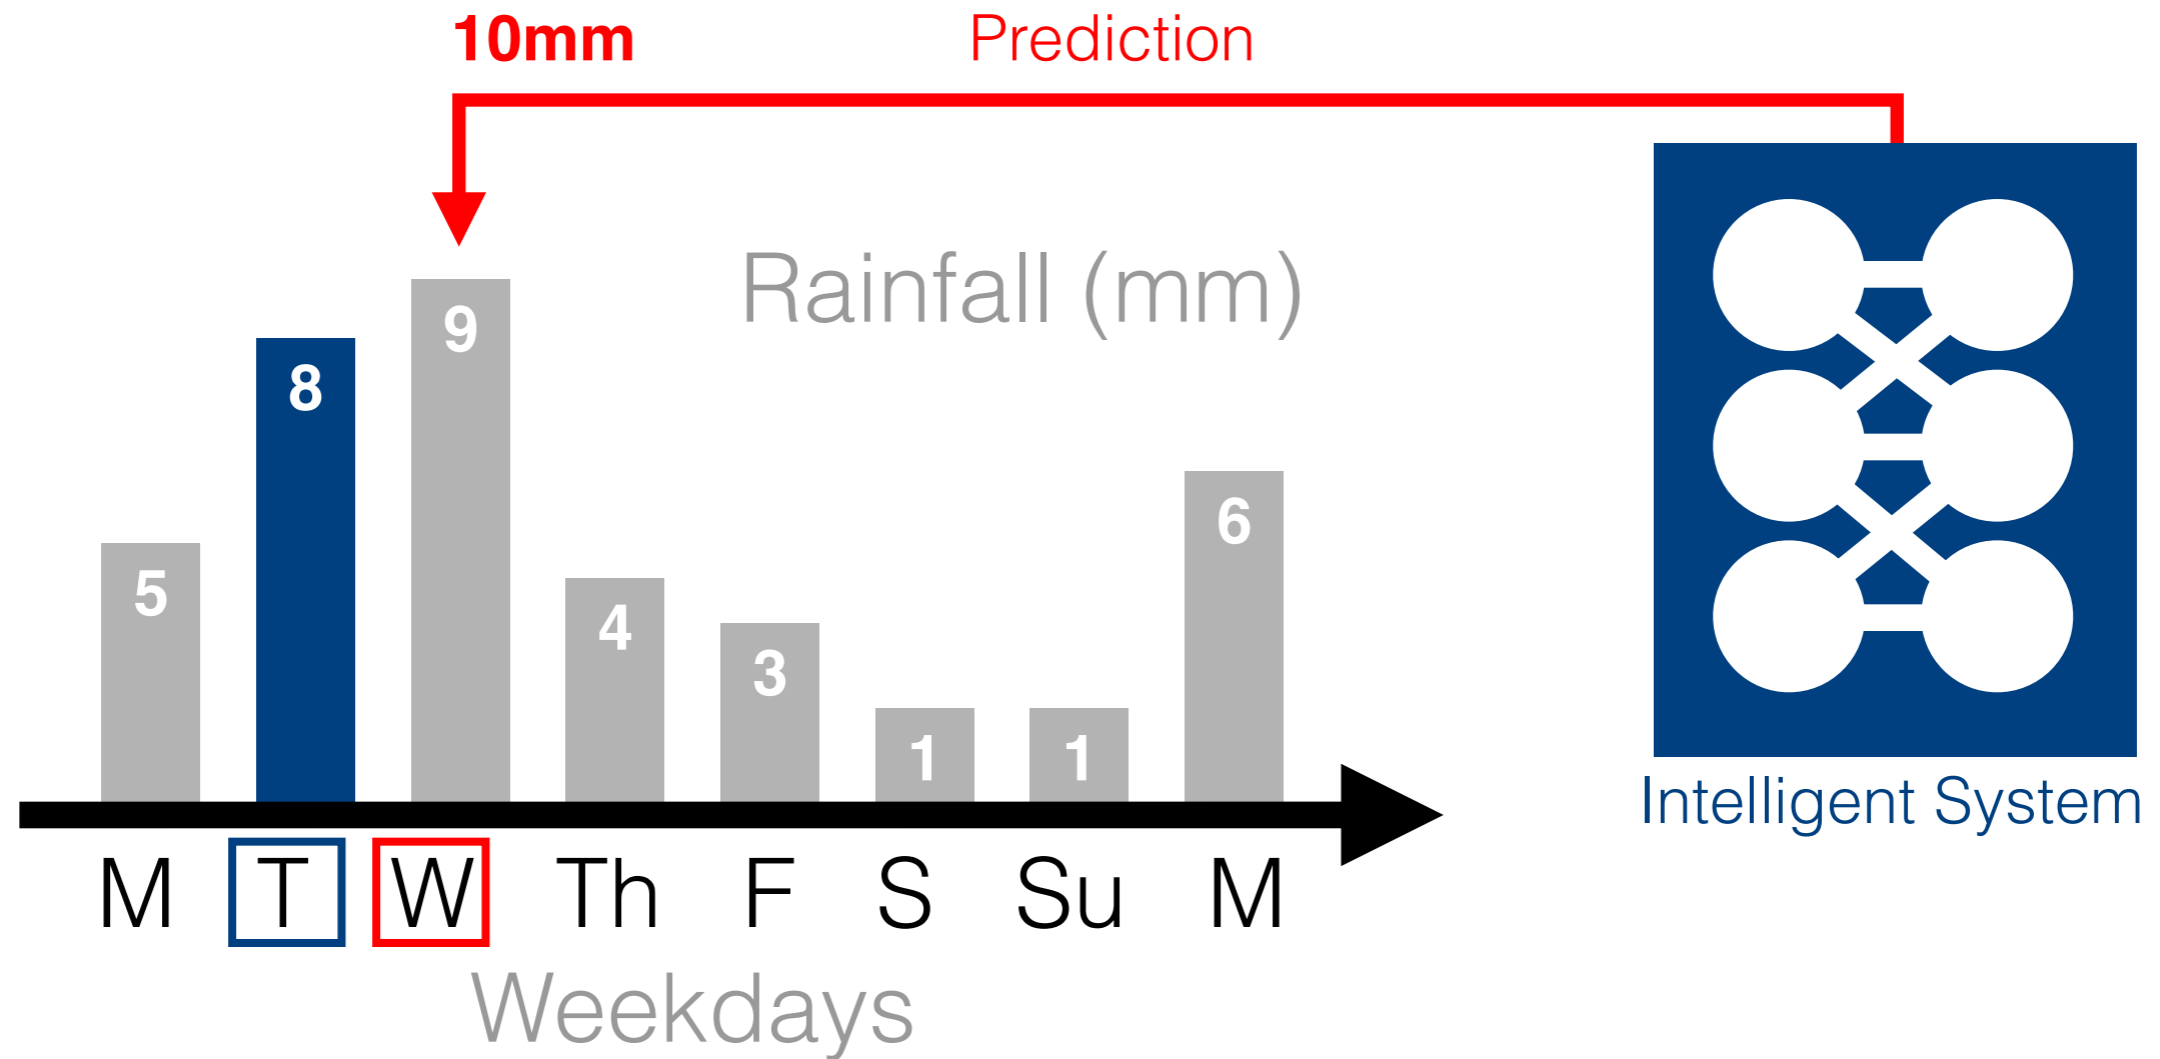
Prediction Learning



One time-step prediction.



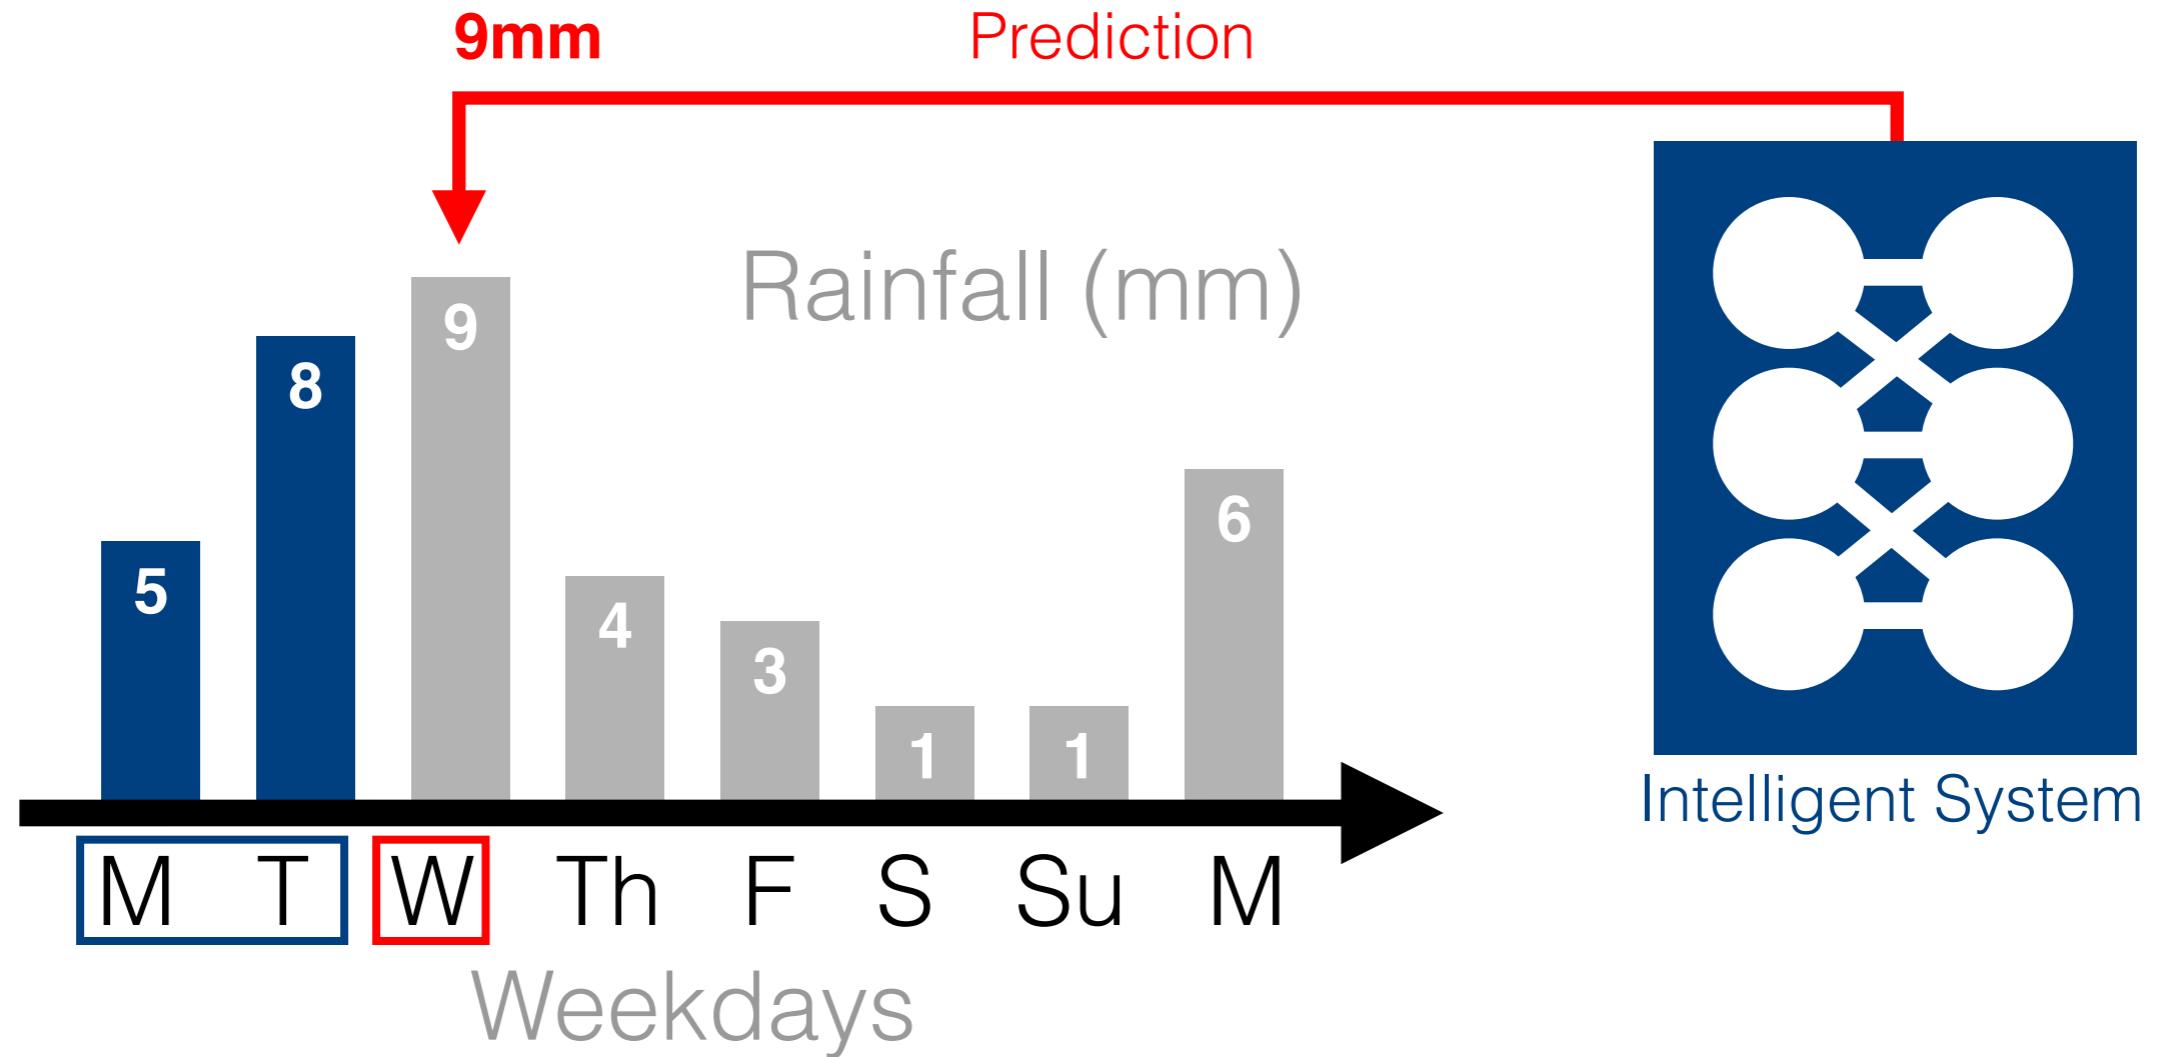
Prediction Learning



One time-step prediction.



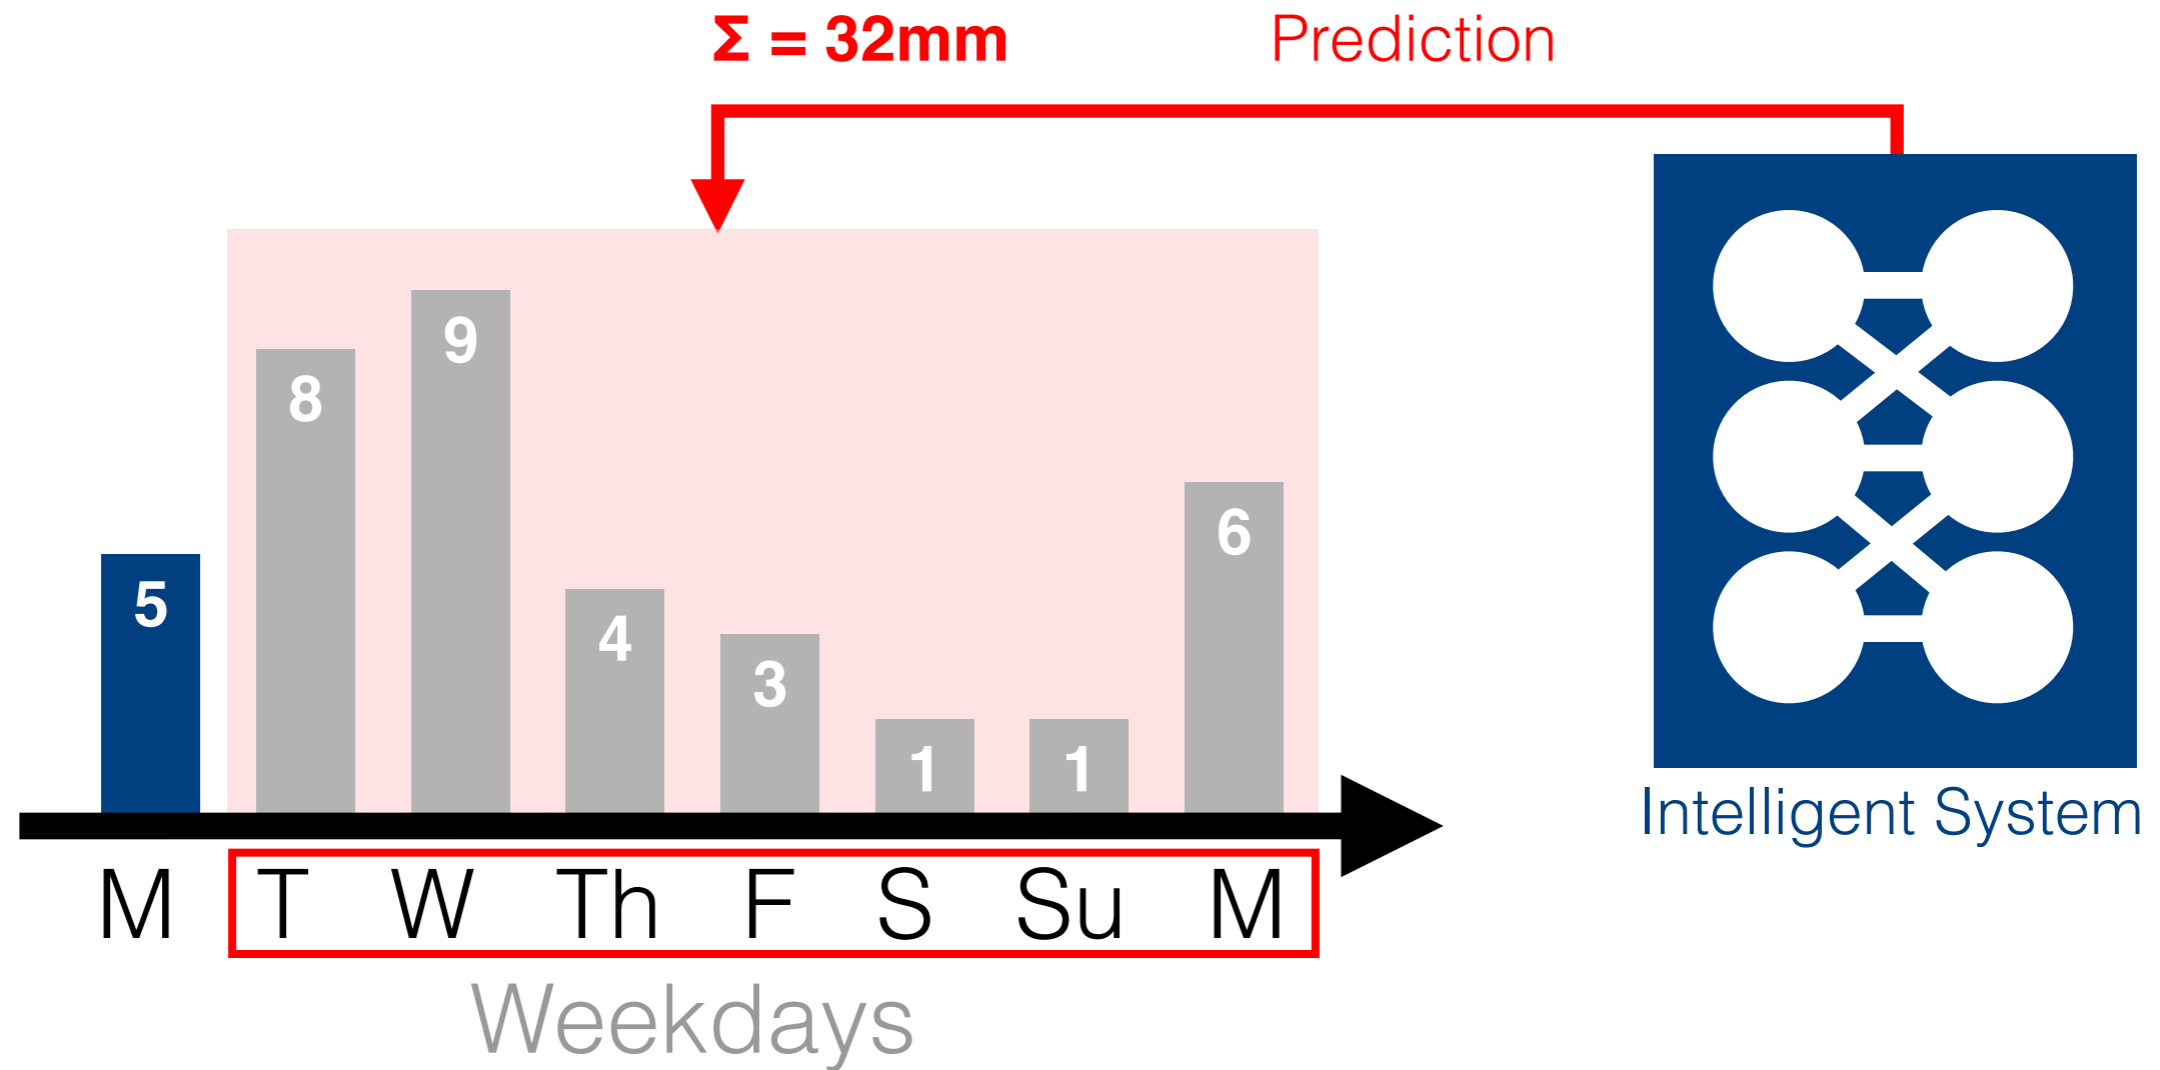
Prediction Learning



One time-step prediction.




Prediction Learning



Temporally extended prediction.




Prediction Learning



Temporally extended prediction.




Control Learning



Using a learned model or values.



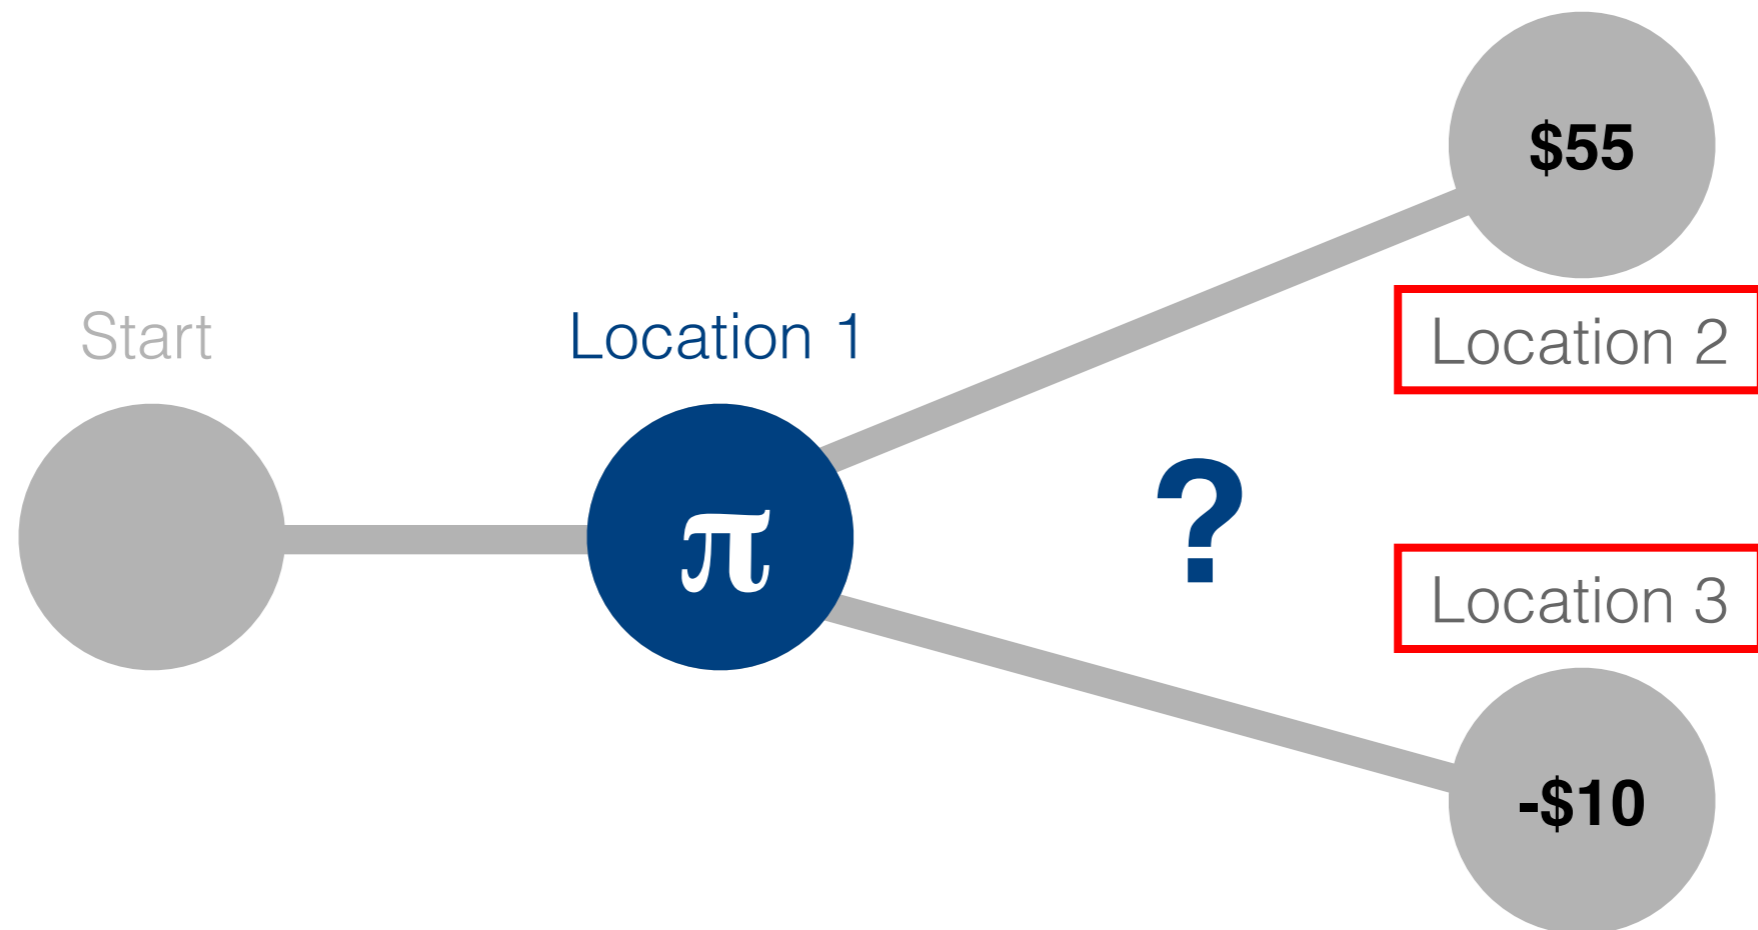
Control Learning



Using a learned model or values.



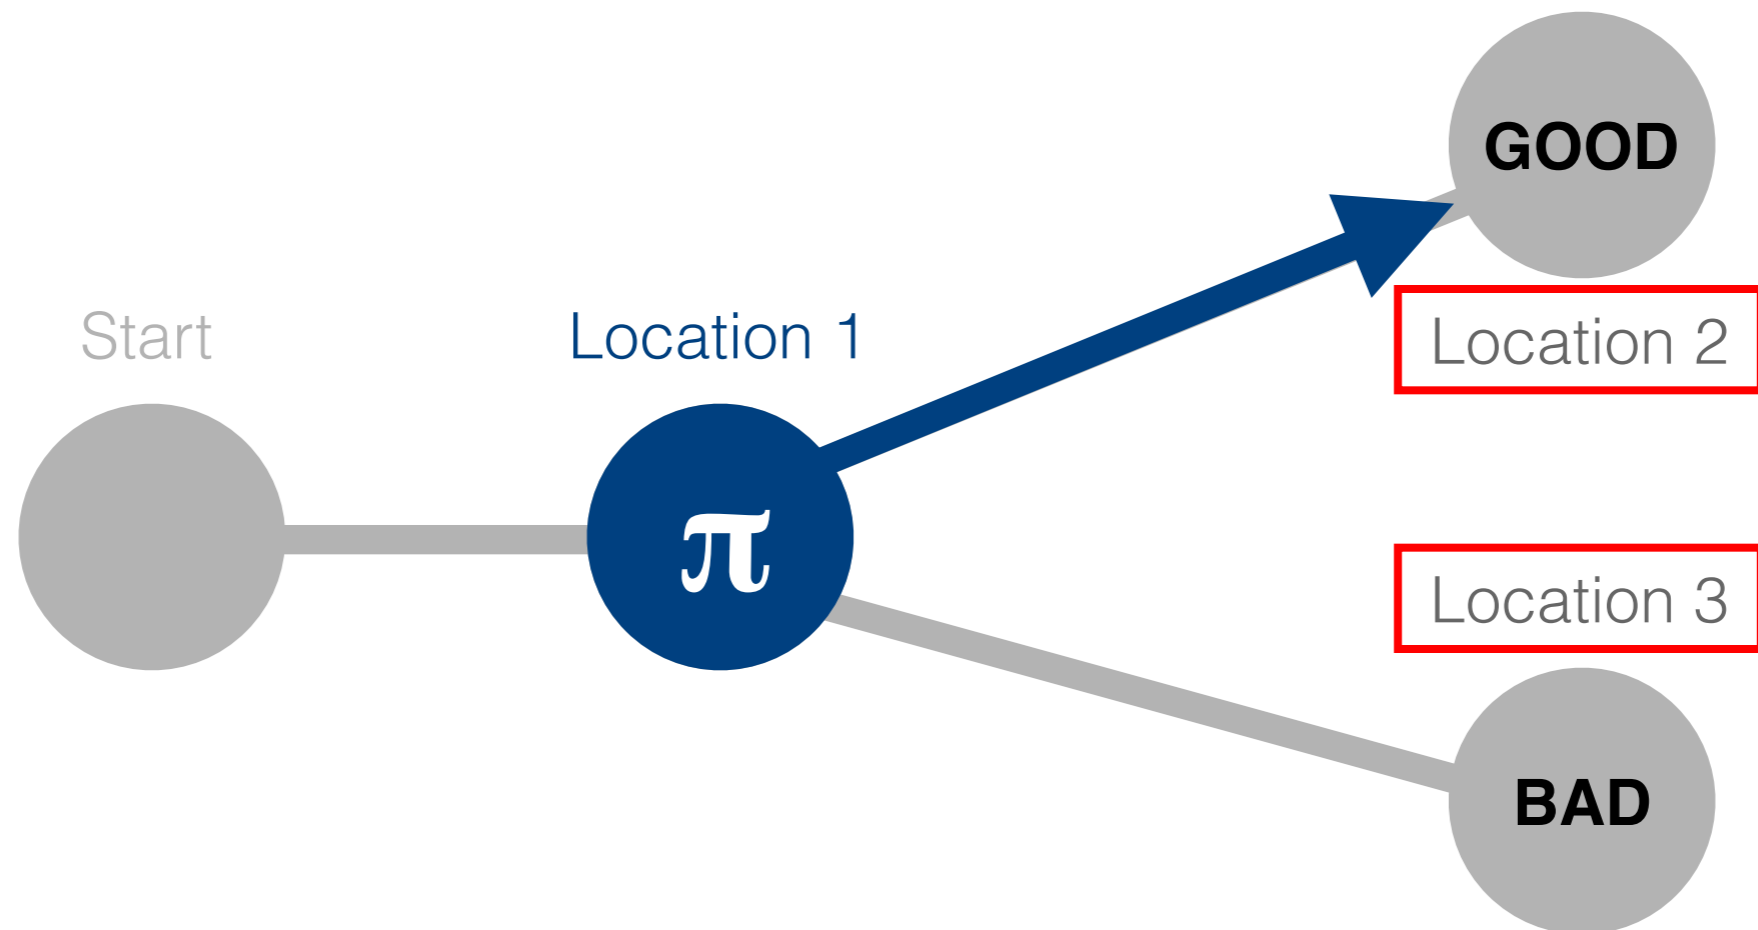
Control Learning



Using a learned model or values.




Control Learning



Using a learned model or values.




Control Learning



Using a learned model or values.



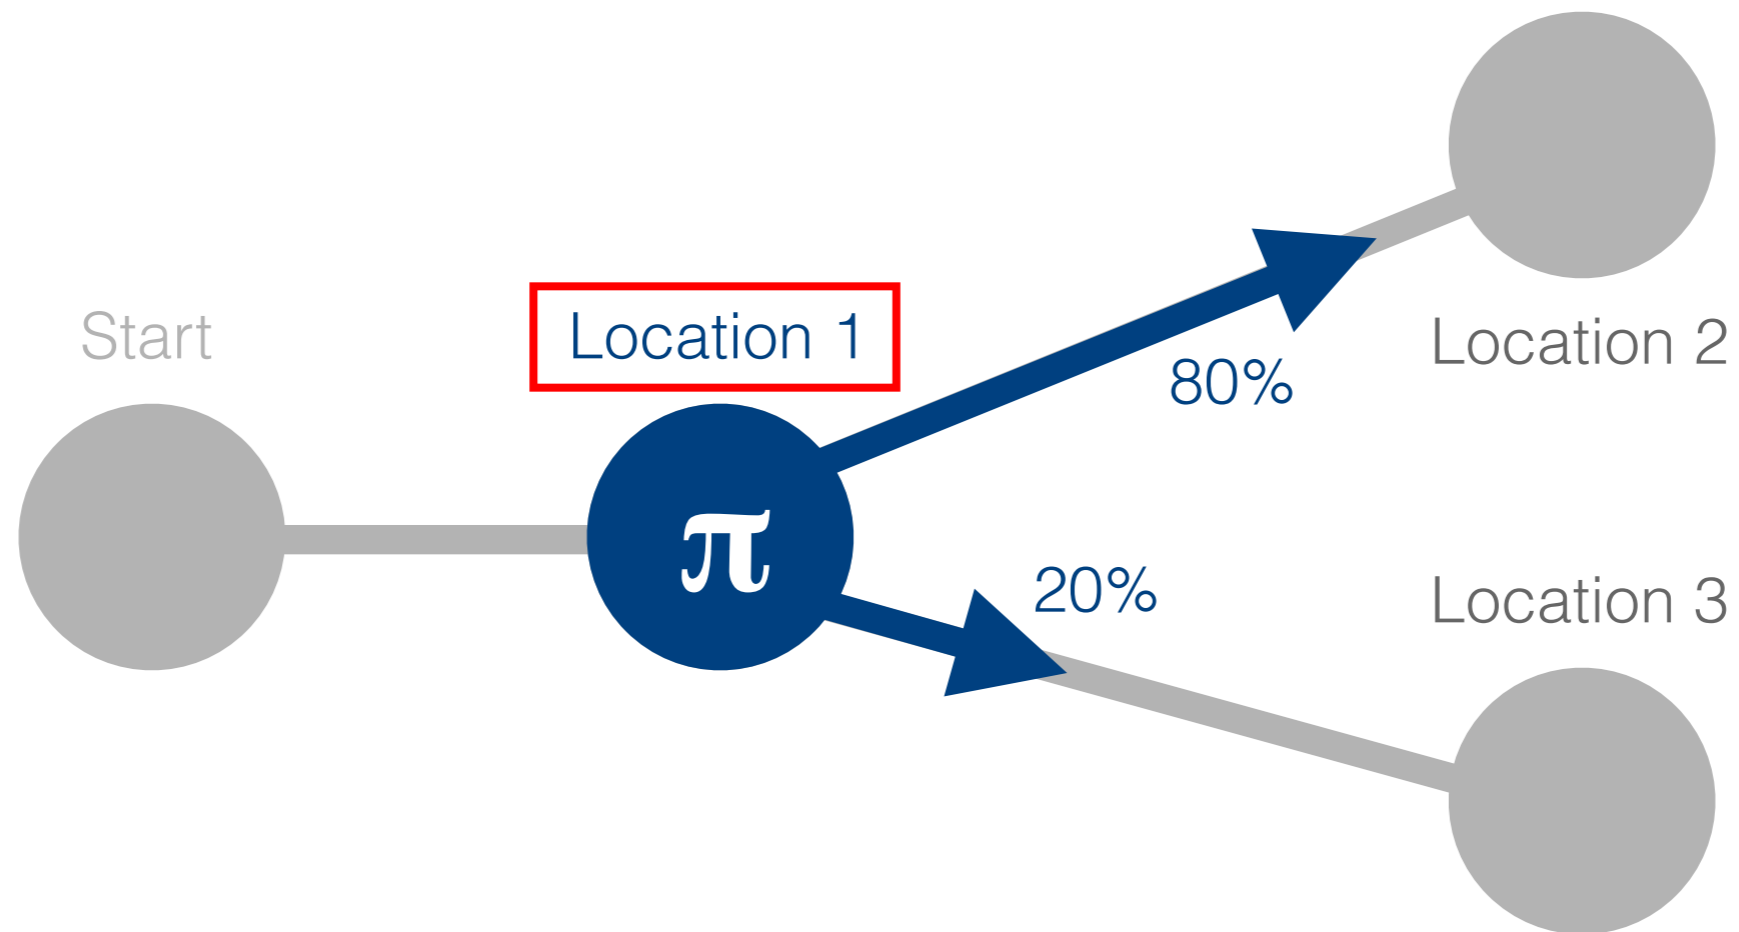
Control Learning



Altering and improving a policy.



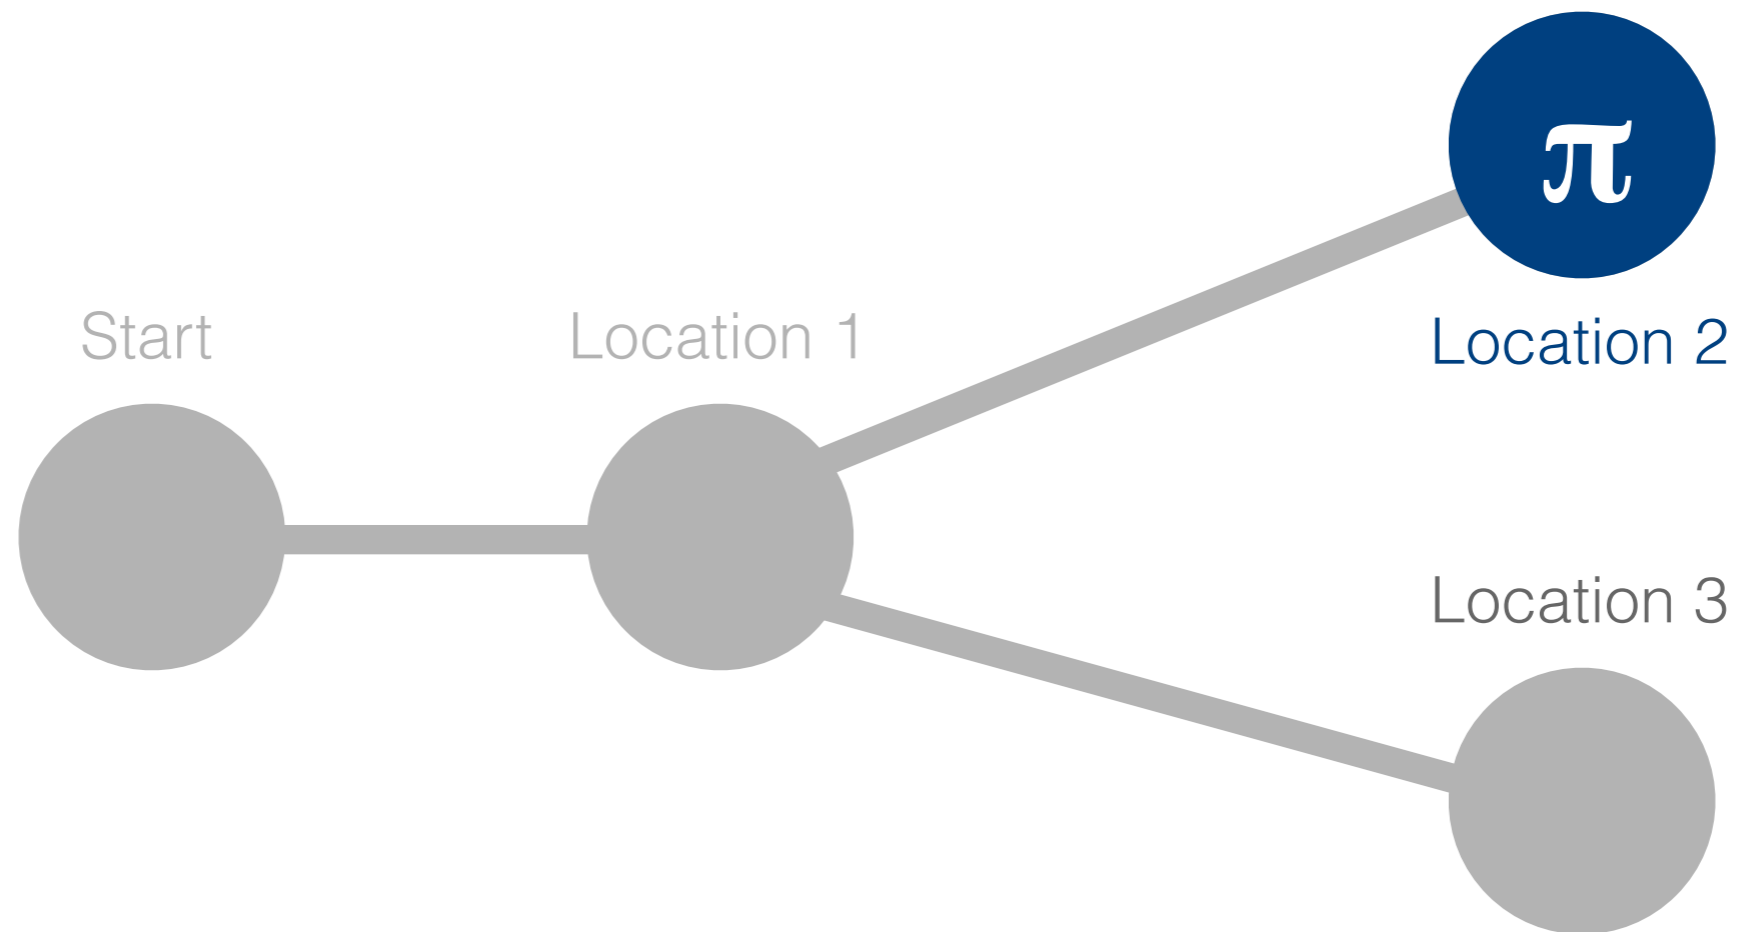
Control Learning



Altering and improving a policy.




Control Learning



Altering and improving a policy.




Representation Learning



Simplifying the data.




Representation Learning



Simplifying the data.




Representation Learning



Simplifying the data.




Representation Learning




Simplifying the data.



Representation Learning



Input Signals



Intelligent System


ABAACABAACA

Output Symbols

Simplifying the data.




How to Learn



- **From labeled examples:** e.g., prediction learning.
- **Finding structure in the data:** e.g., representation learning.
- **Through trial and error:** e.g., control learning.


From Labeled Examples



Prediction Learning.




From Labeled Examples



Prediction Learning.




From Labeled Examples



Prediction Learning.




From Labeled Examples



Prediction Learning.




From Labeled Examples




Prediction Learning.



From Labeled Examples



Prediction Learning.



Finding Structure in the Data




A



B

Make two groupings.

Representation Learning.



Finding Structure in the Data



A




B



C

Make two groupings.

Representation Learning.



Finding Structure in the Data



A



B




C



D

Make two groupings.

Representation Learning.



Finding Structure in the Data

Beards

No Beards



A


B

C

D

Make two groupings.

Representation Learning.



Finding Structure in the Data

No Visible Teeth

Visible Teeth

No Visible Teeth



A



B




C



D

Make two groupings.

Representation Learning.



Finding Structure in the Data

Not Patrick's Mother

Patrick's Mother

Not Patrick's Mother



A

B

C

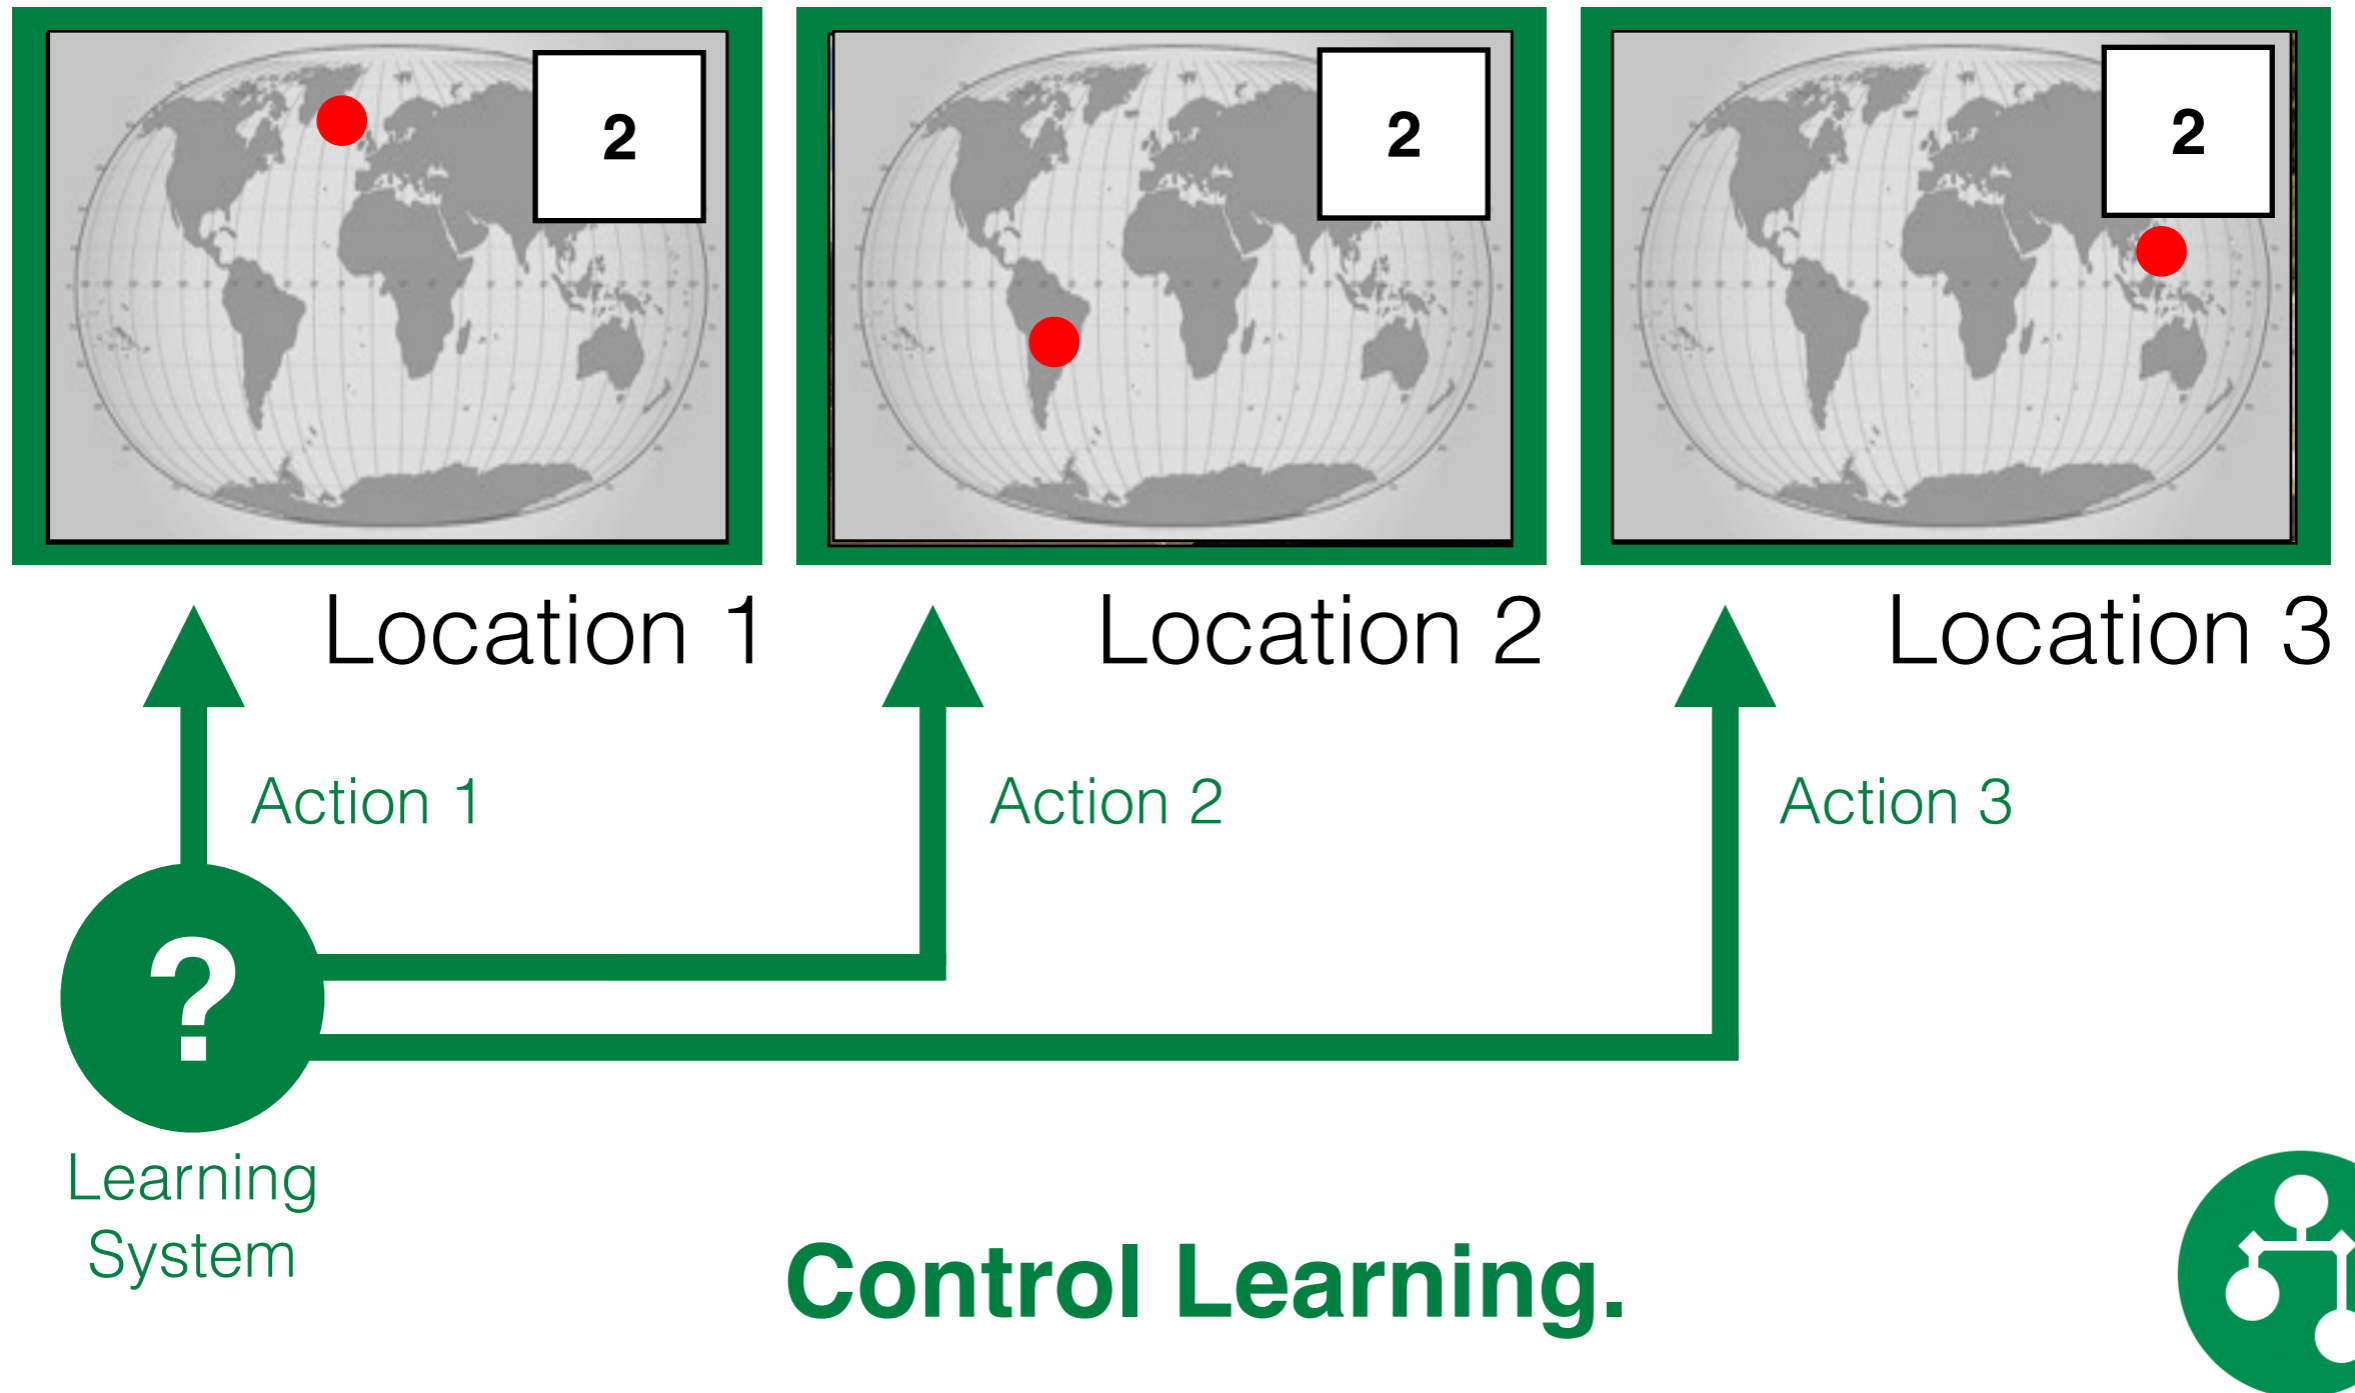
D

Make two groupings.

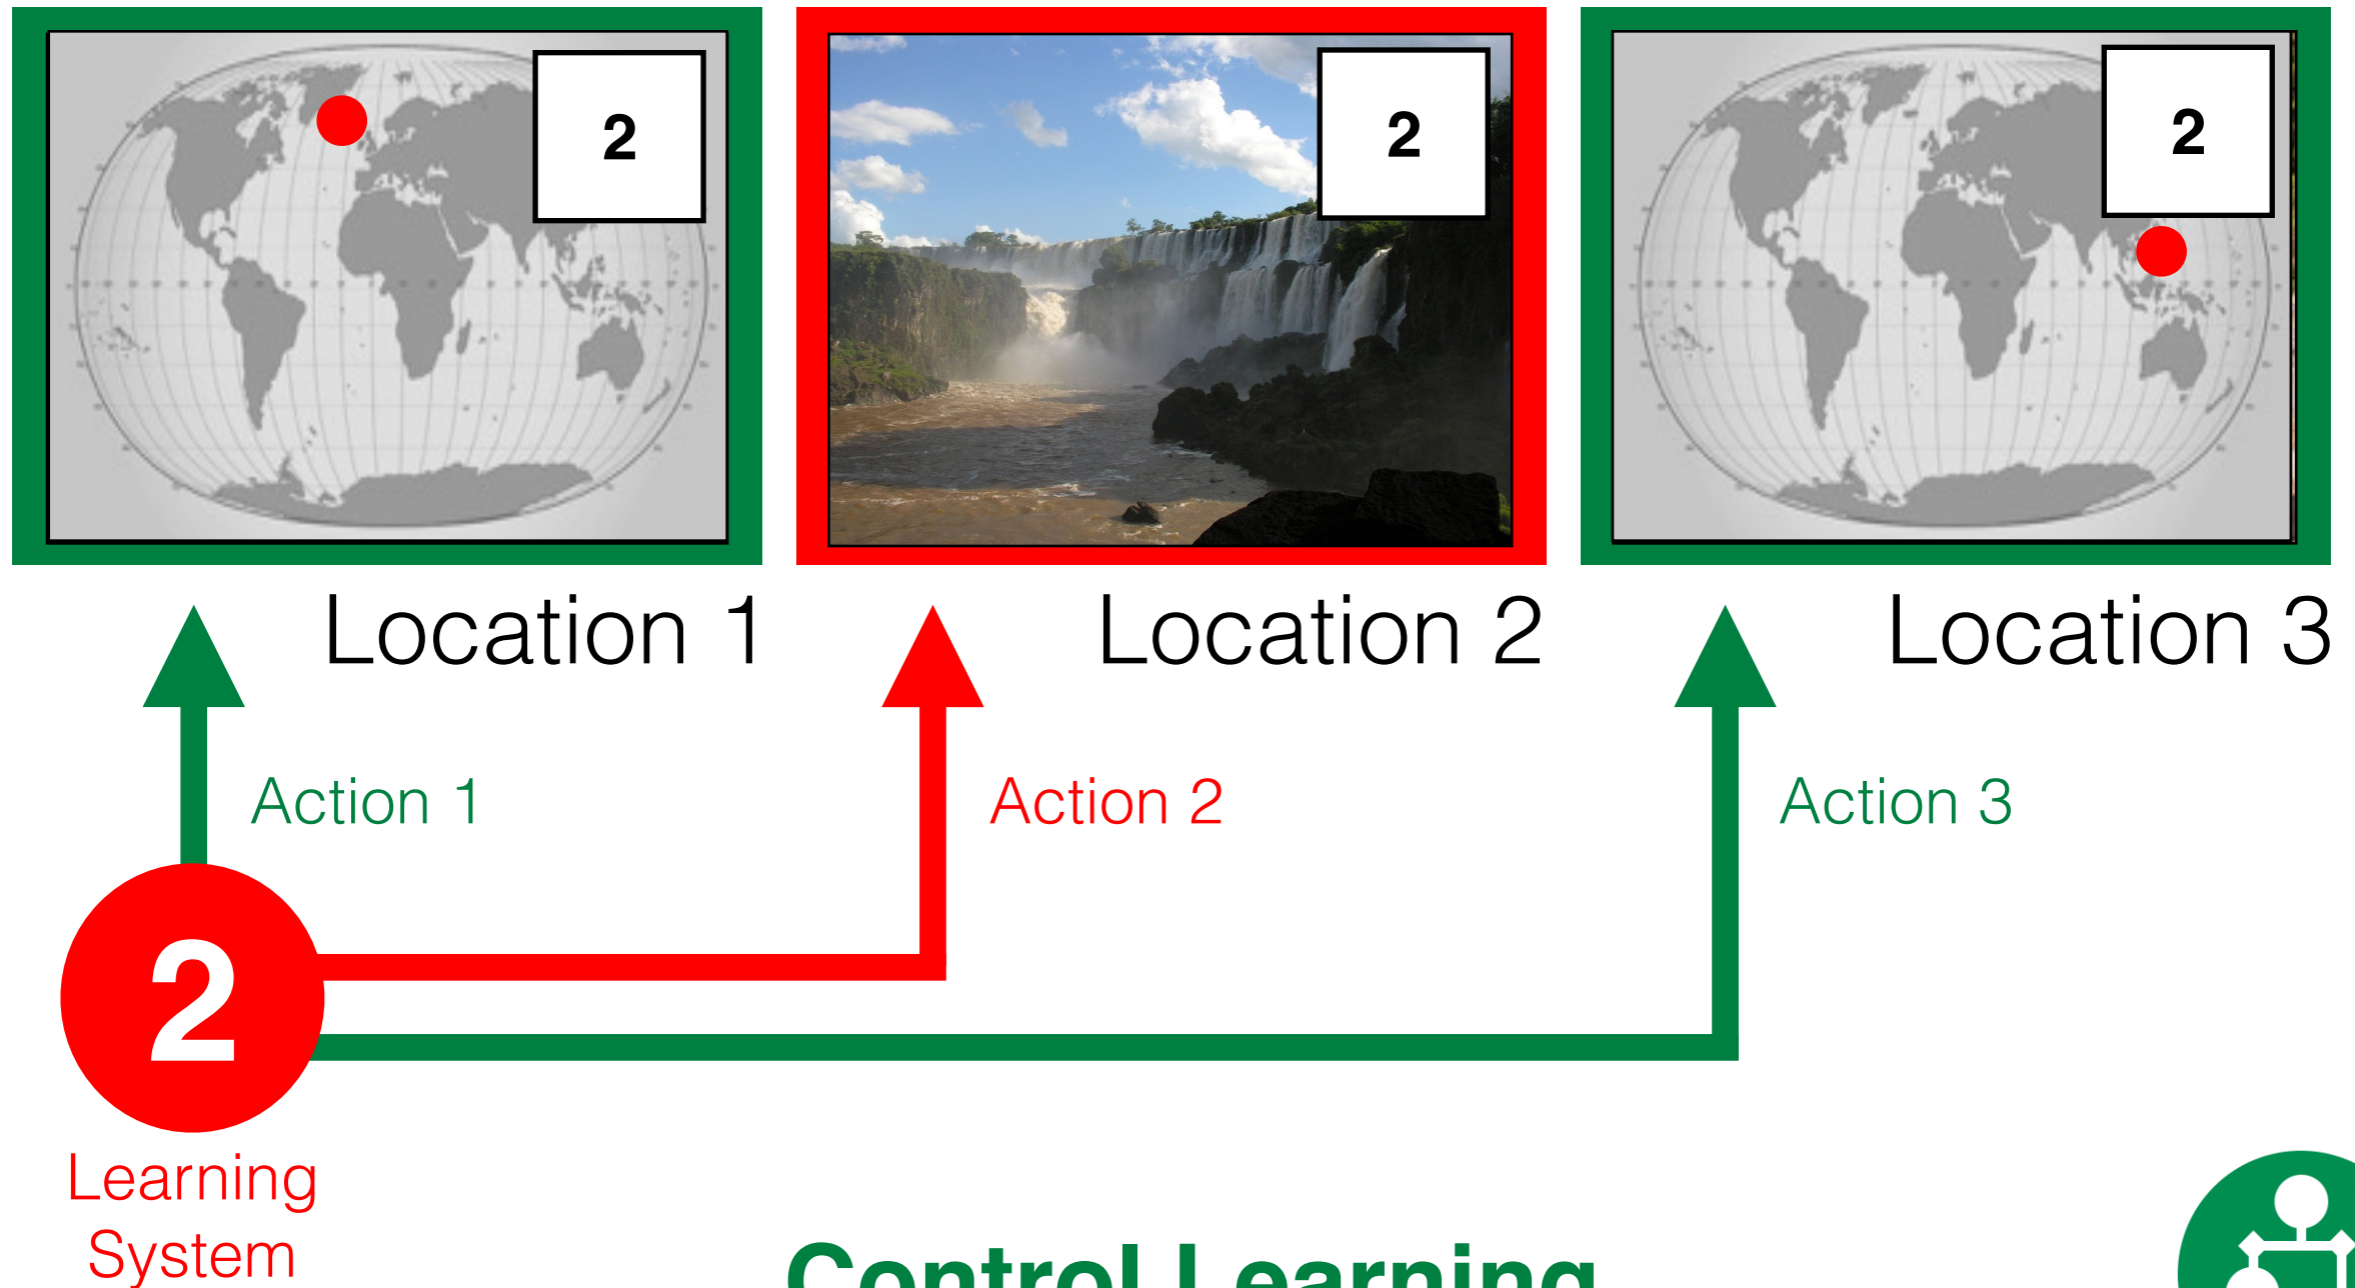
Representation Learning.



Through Trial and Error




Through Trial and Error



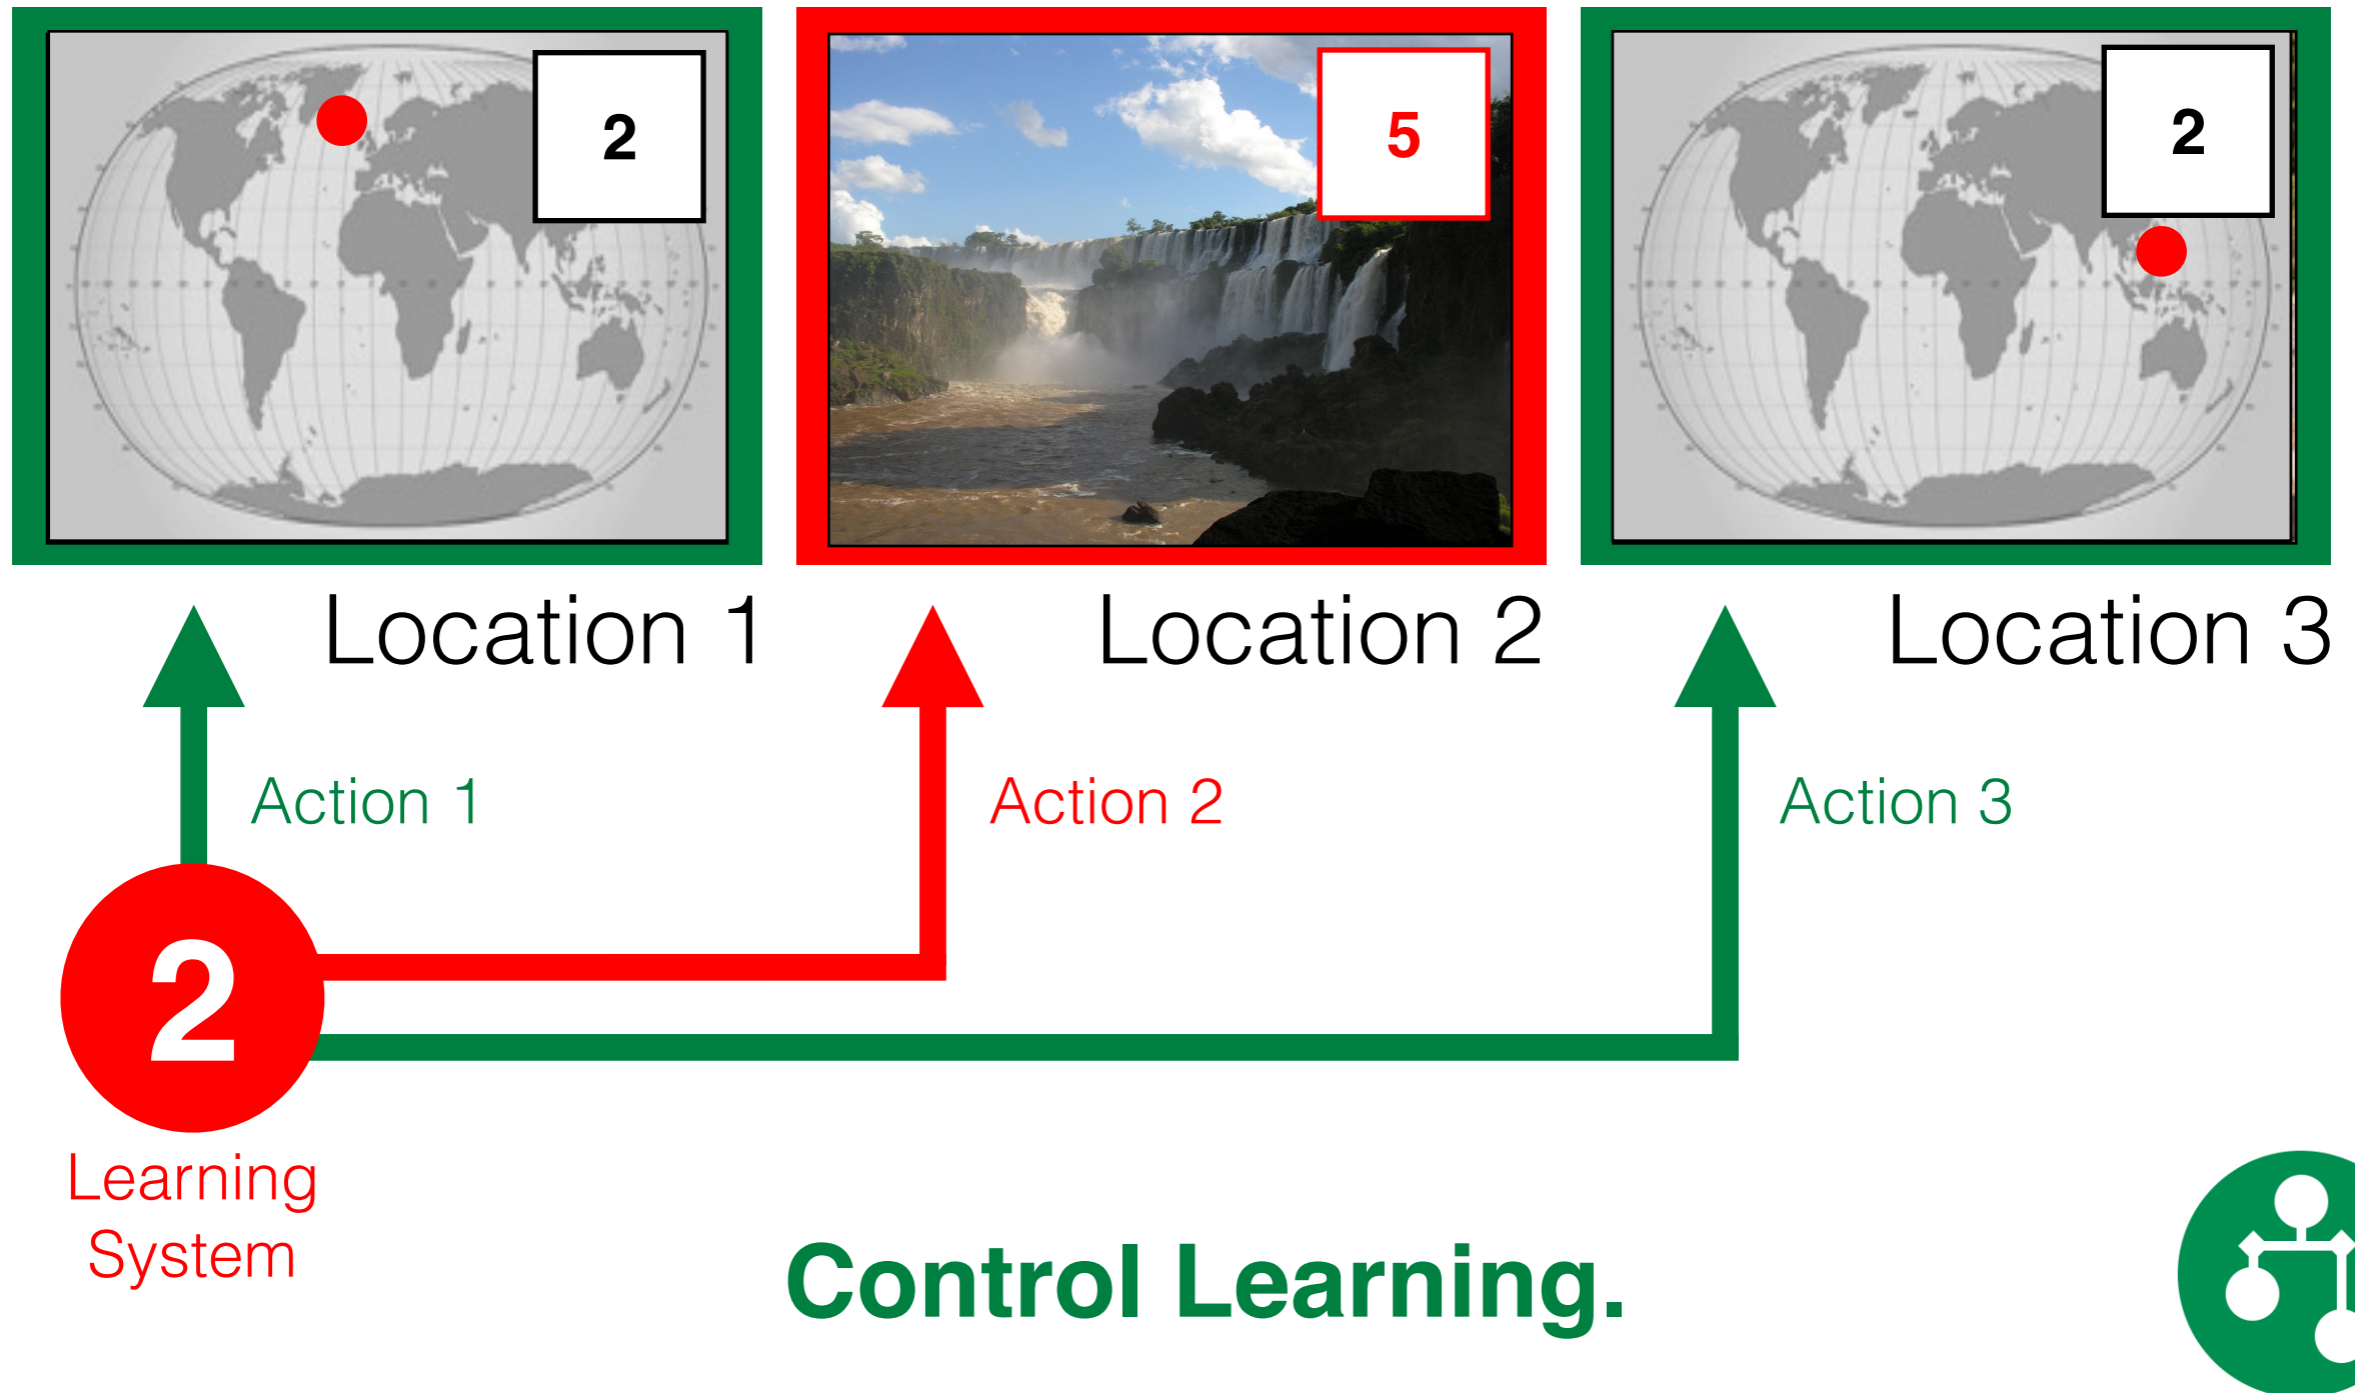
Control Learning.



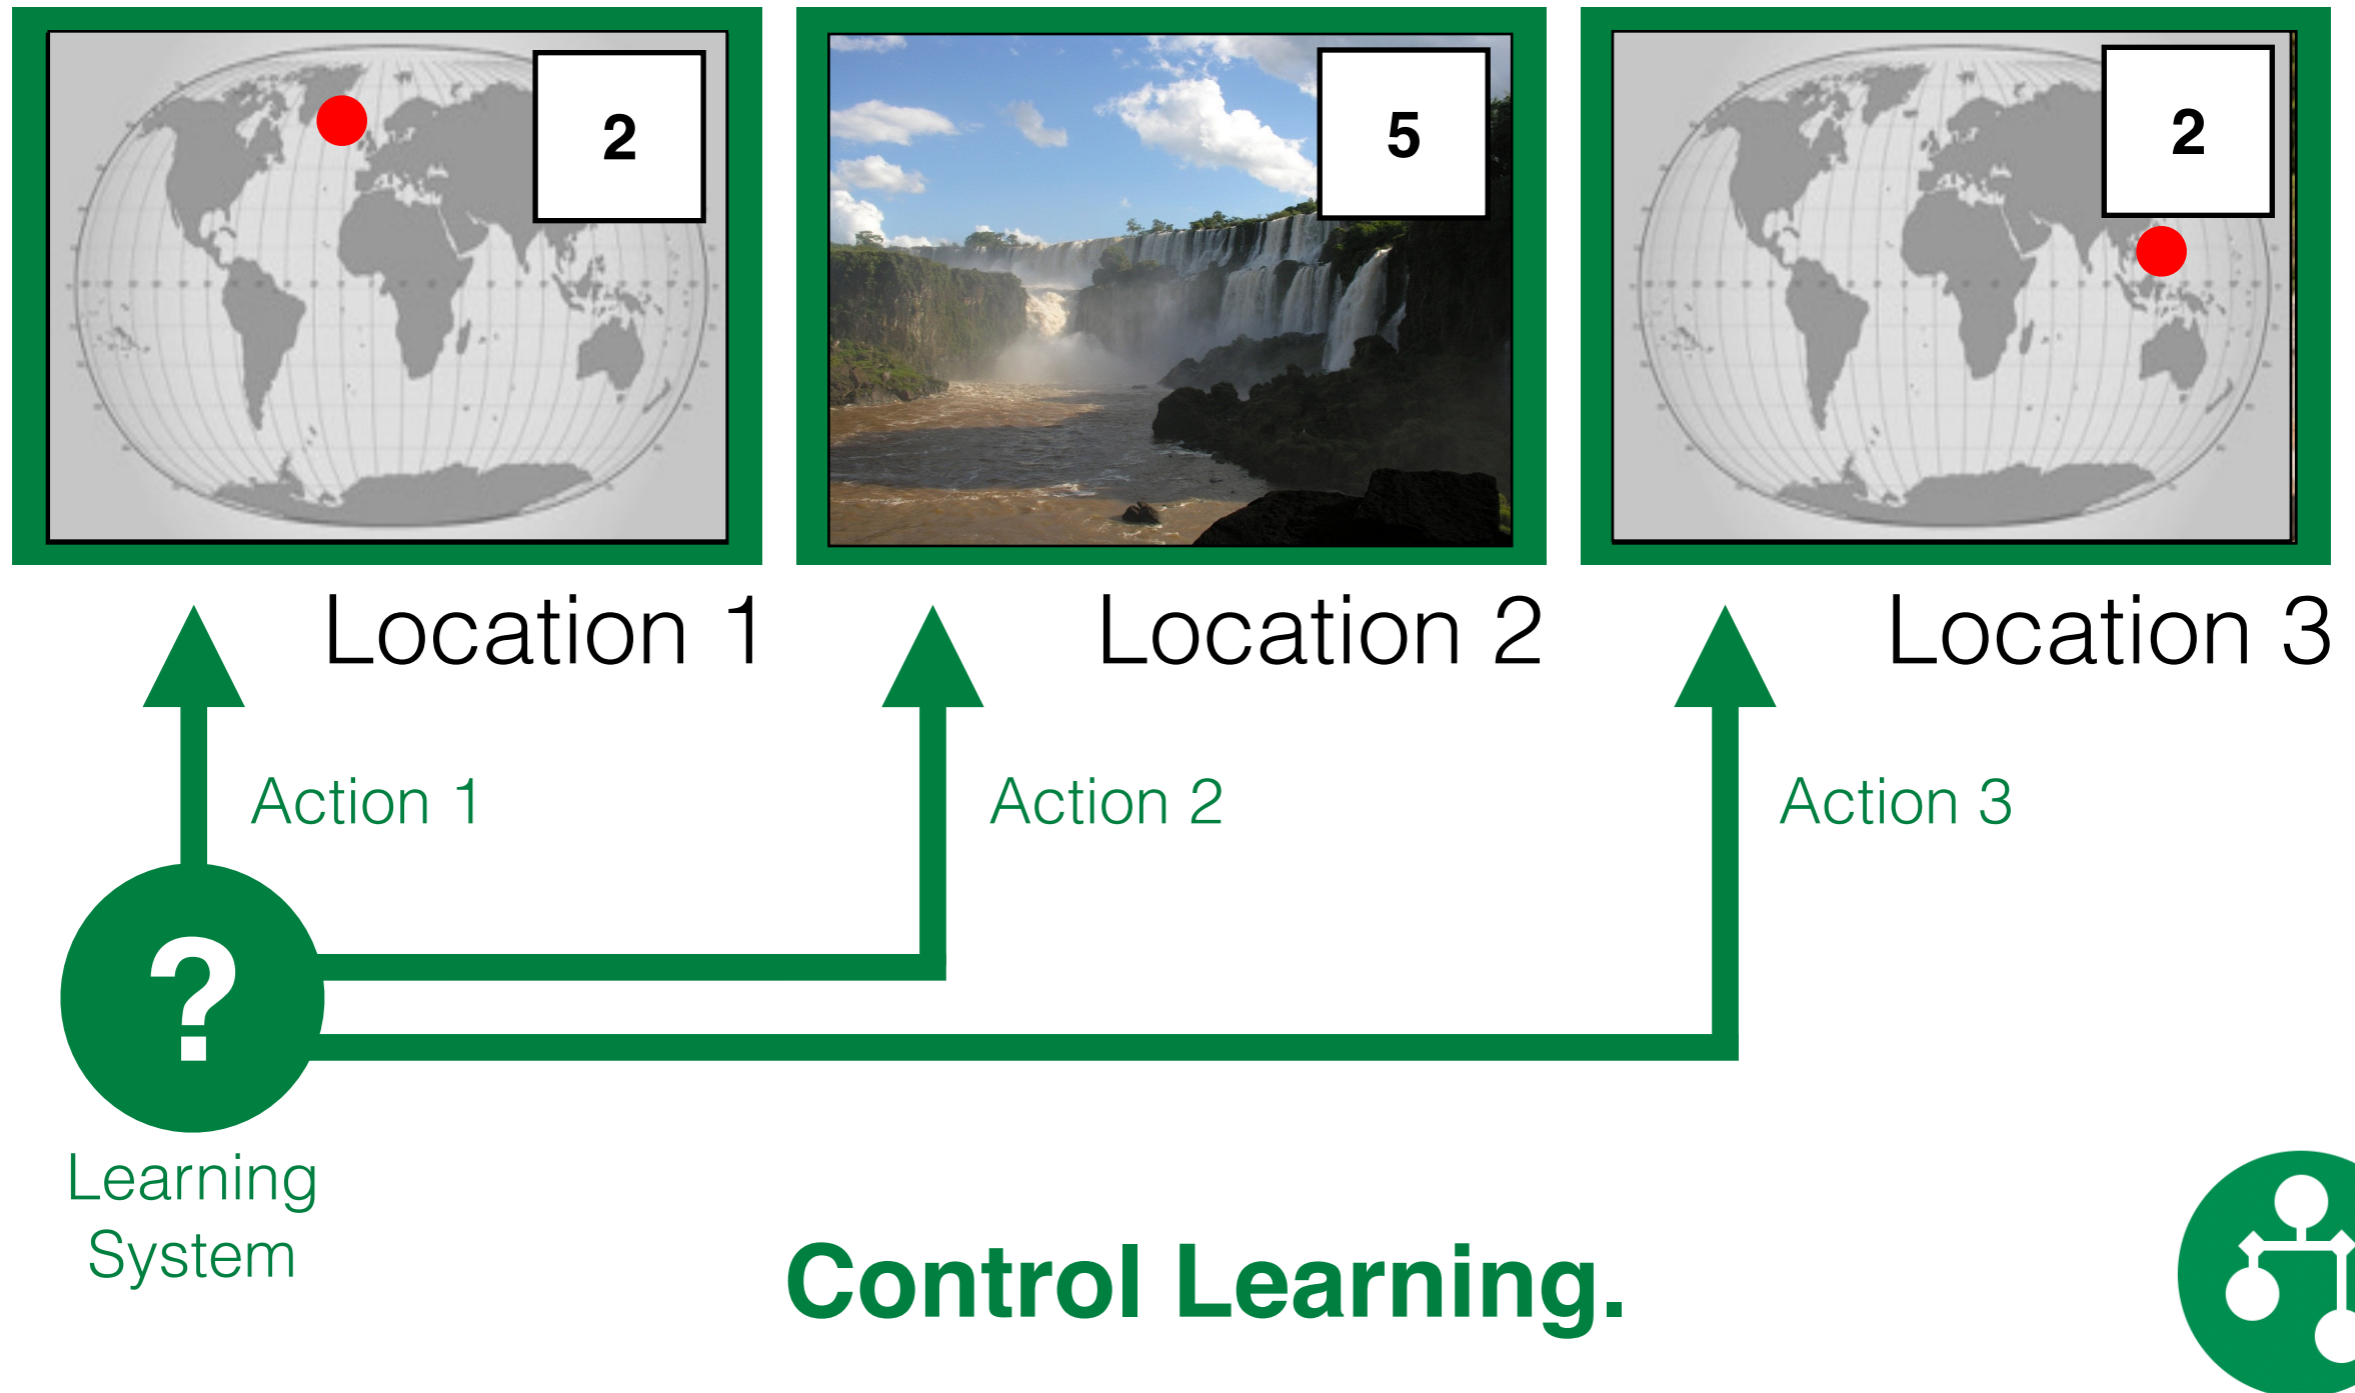
Through Trial and Error



Through Trial and Error




Through Trial and Error



Through Trial and Error



Control Learning.




Through Trial and Error



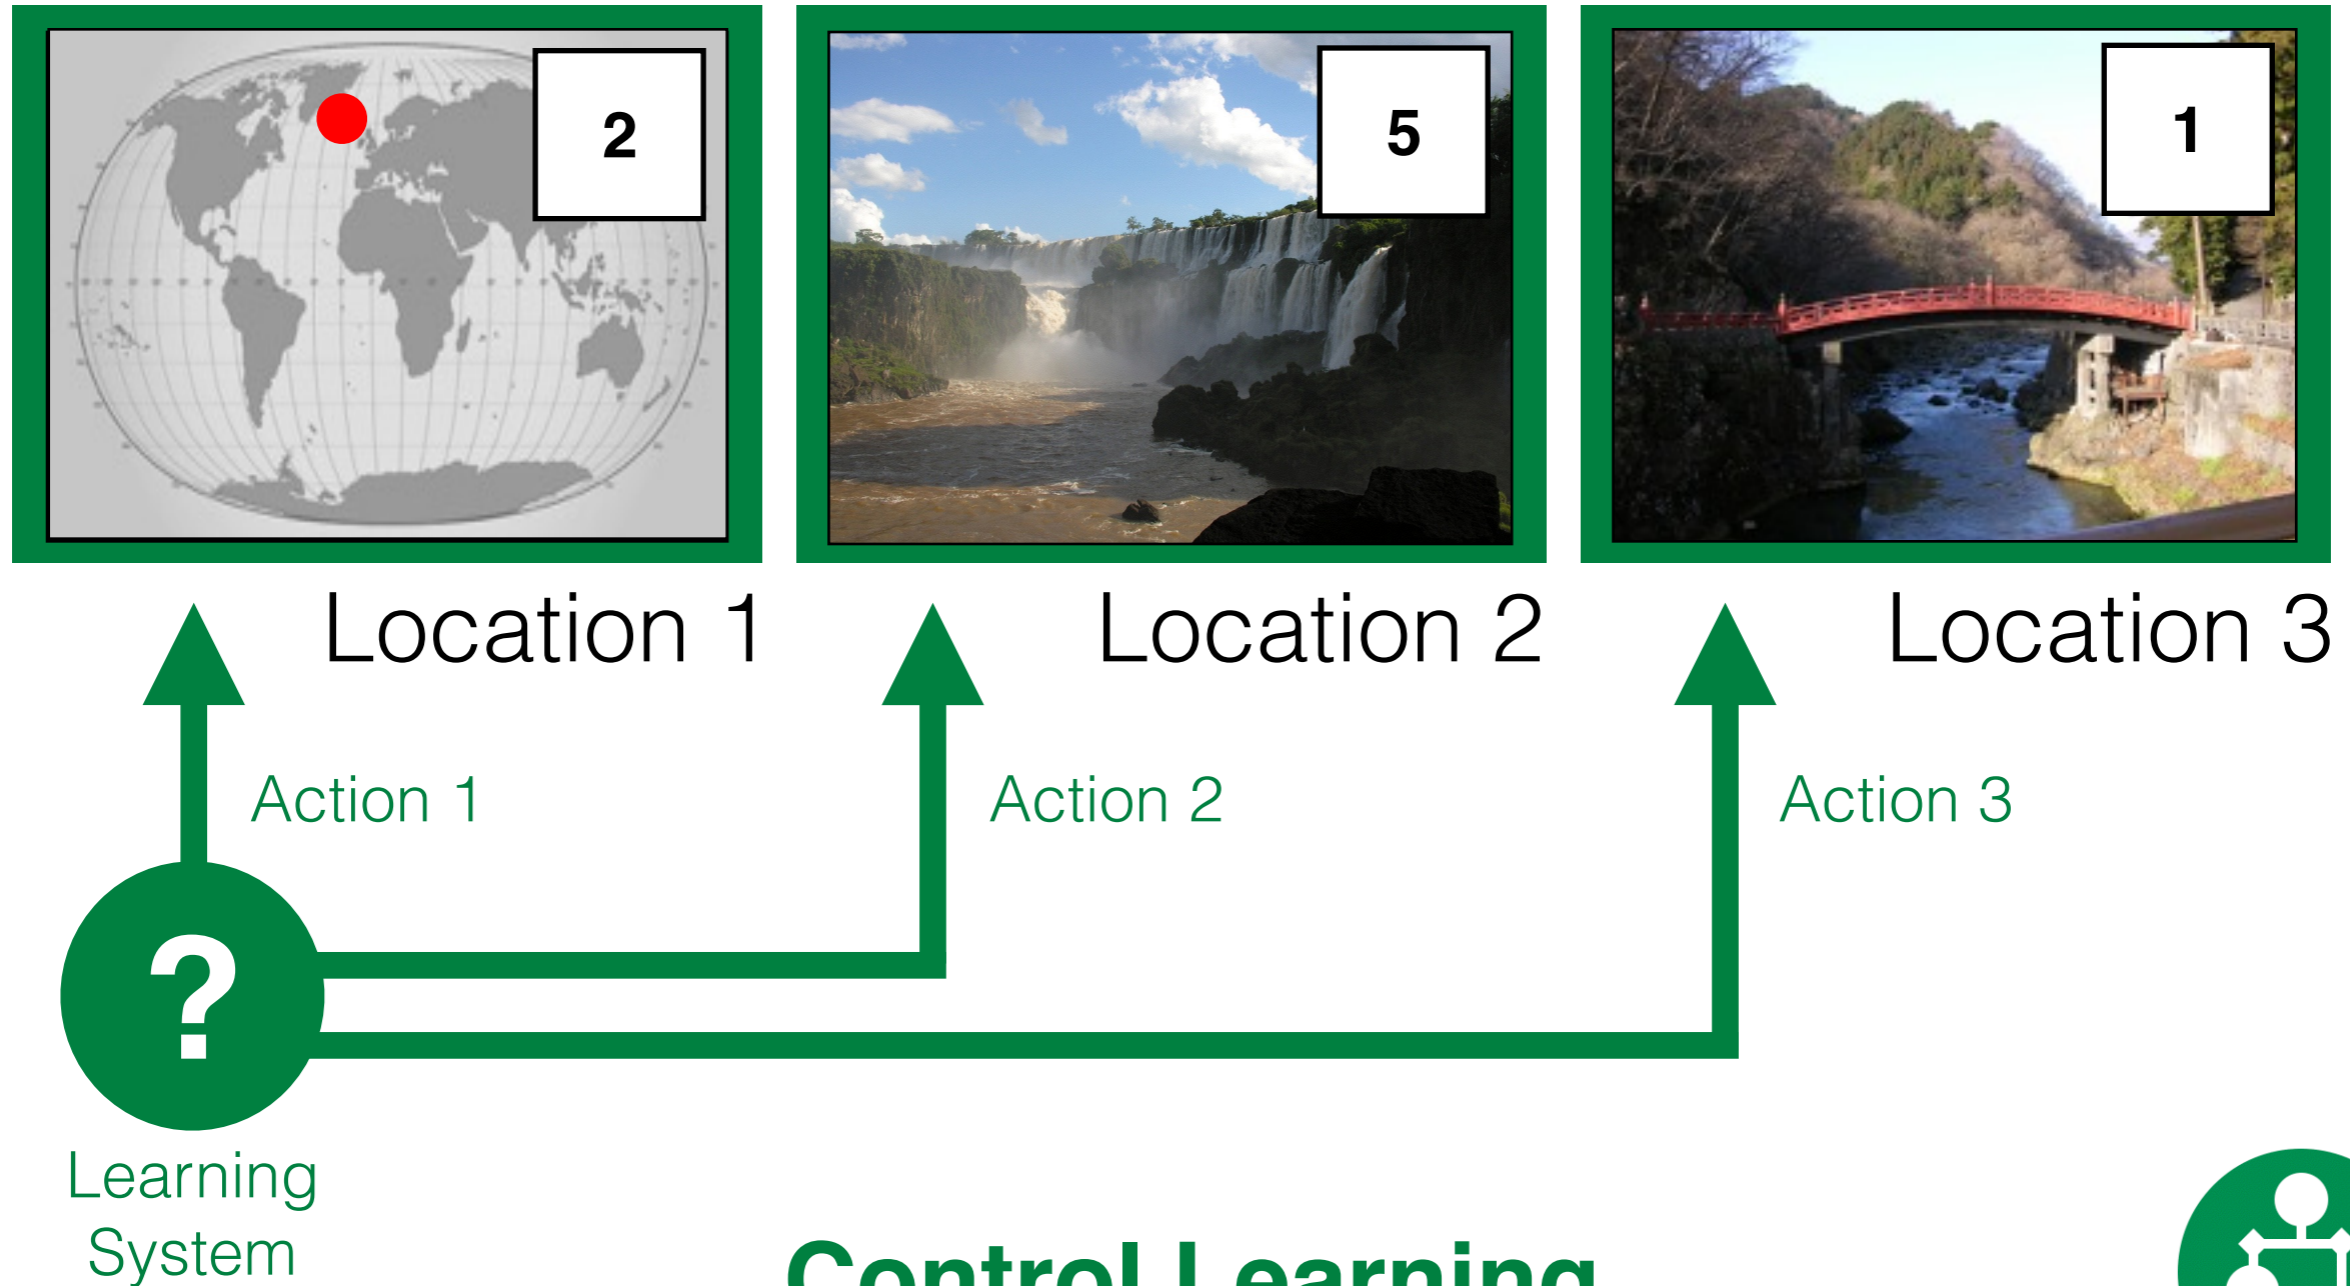
Through Trial and Error




Control Learning.



Through Trial and Error



Control Learning.



Through Trial and Error



Control Learning.



Through Trial and Error



Location 1



Location 2




Location 3

-\$10

Action 1


Action 2

Action 3



Learning System


Control Learning.



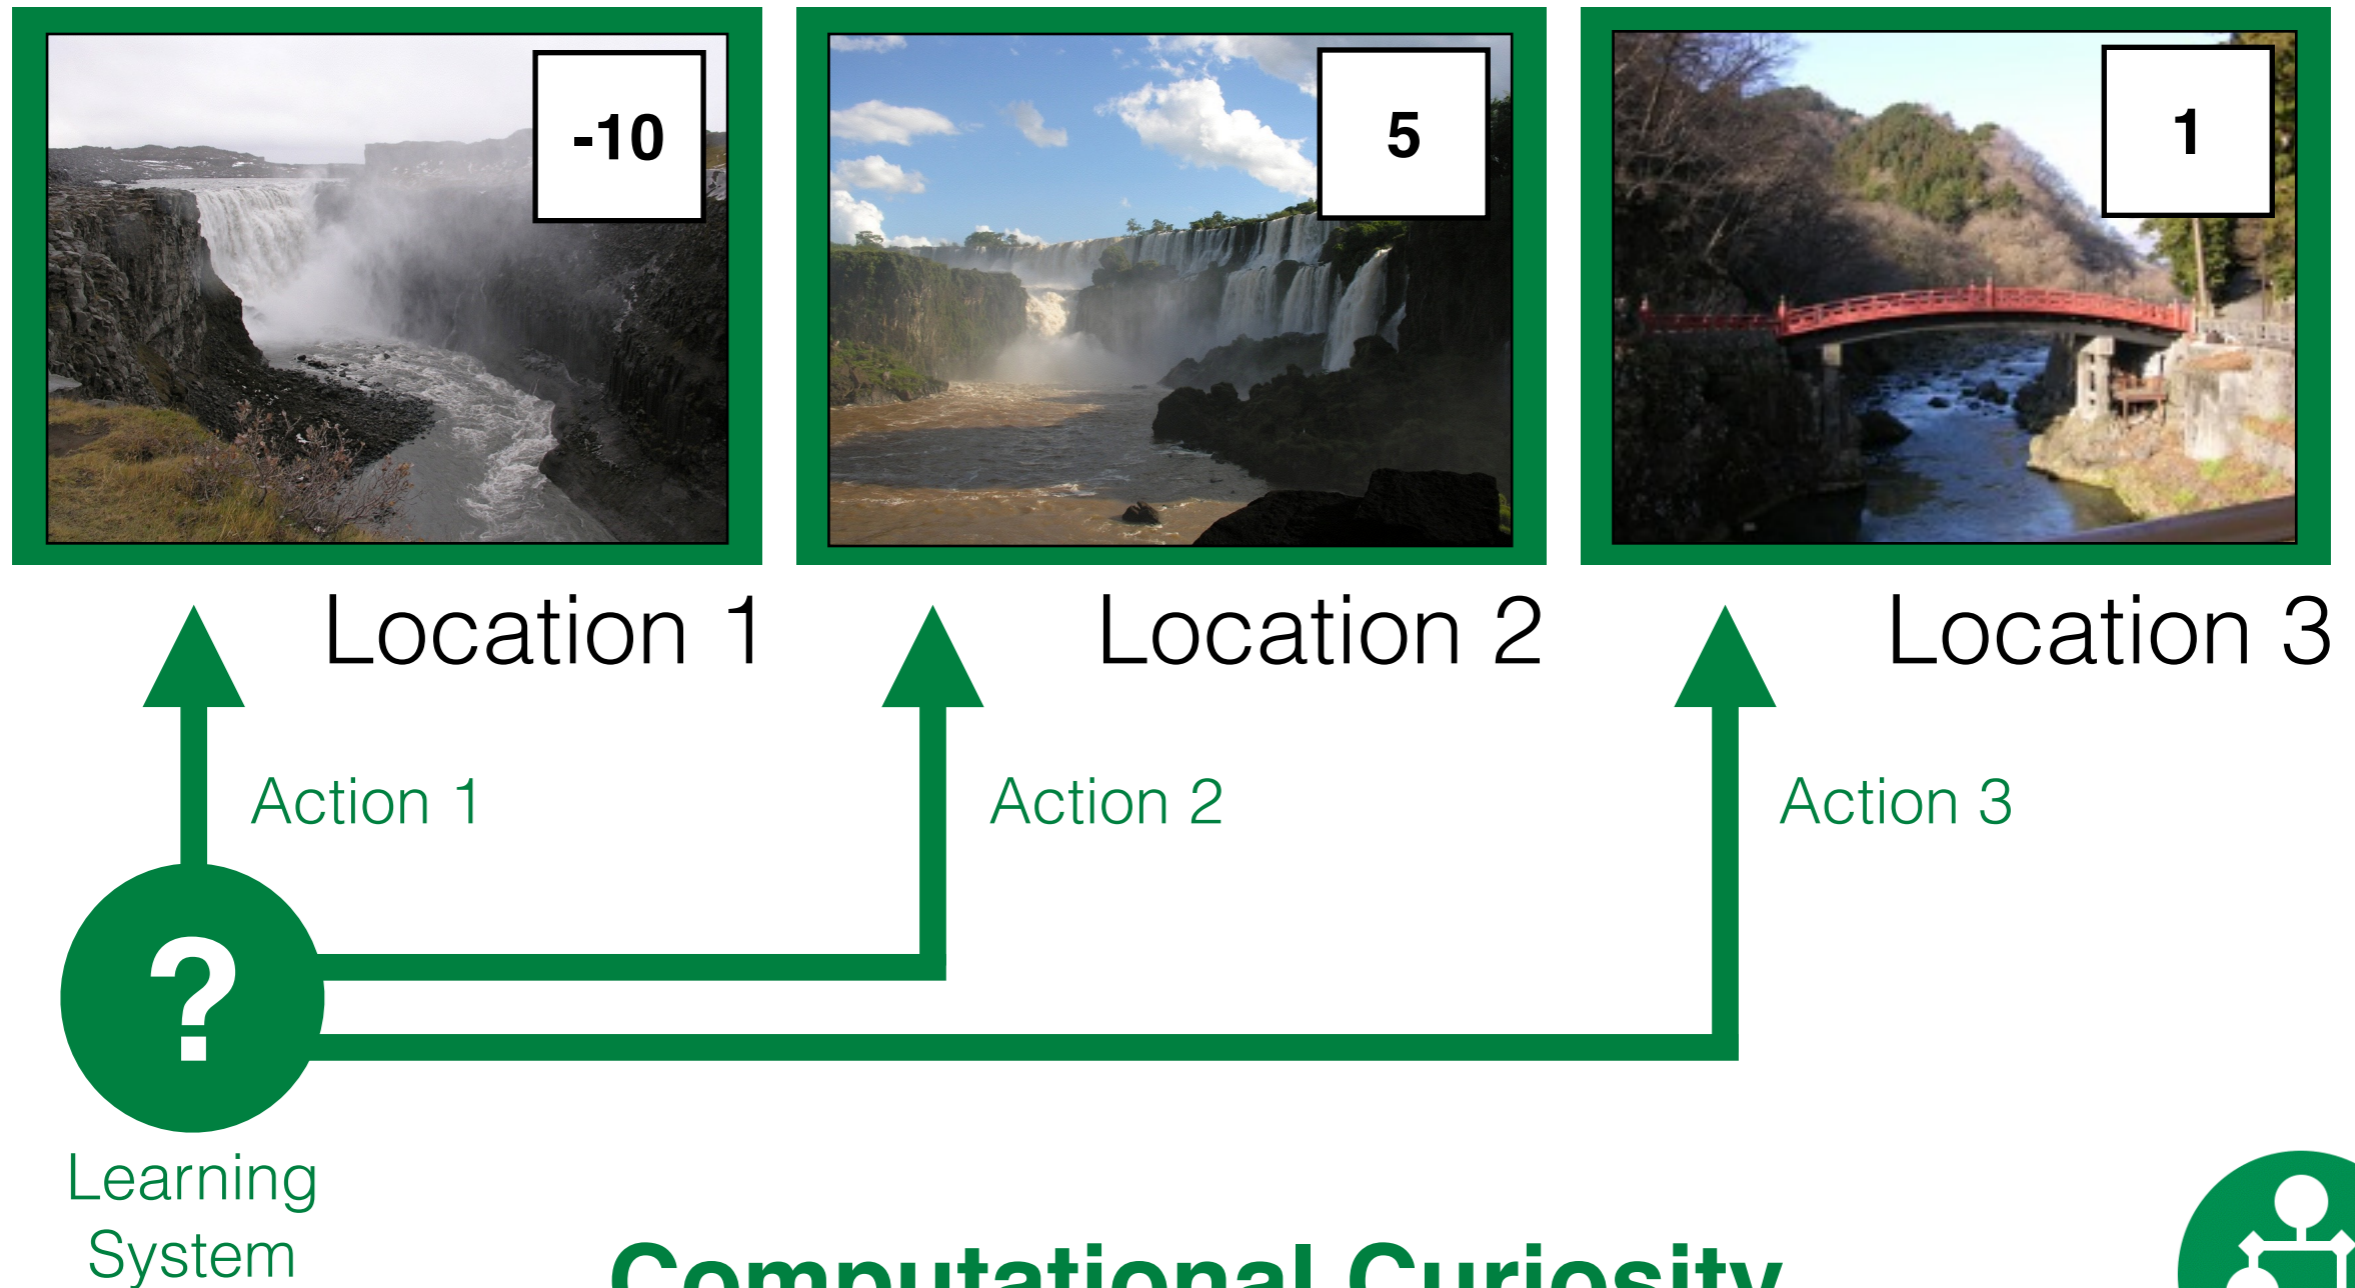
Through Trial and Error



Control Learning.



Extrinsic Motivation



Computational Curiosity.



Intrinsic Motivation?



Computational Curiosity.



Intrinsic Motivation?



Computational Curiosity.





When to Learn



- **In real time:** online learning.
- **From past experience:** offline or batch learning.



Offline Learning



Learning from past experience.



Offline Learning


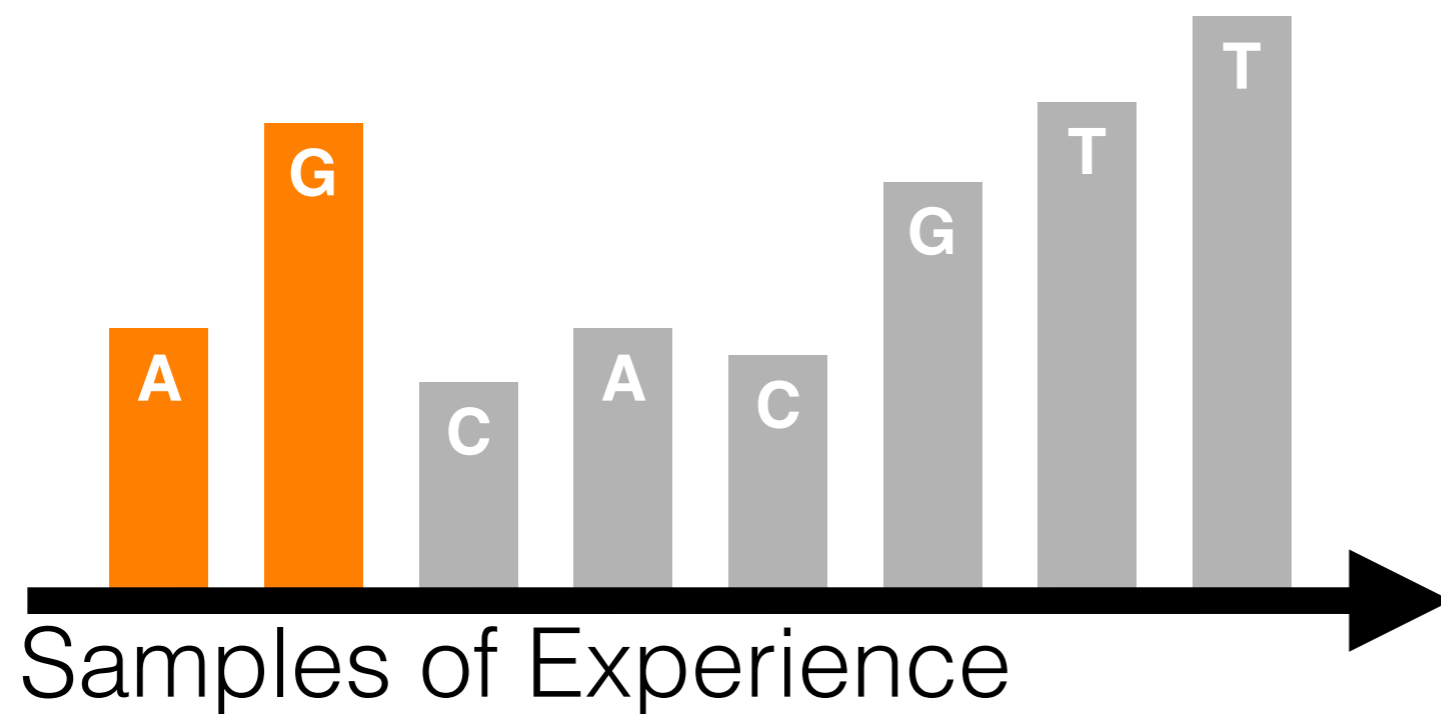


Acquire samples.

Learning from past experience.



Offline Learning





Acquire samples.

Learning from past experience.



Offline Learning





Acquire samples.

Learning from past experience.



Offline Learning




Acquire samples.

Learning from past experience.




Offline Learning

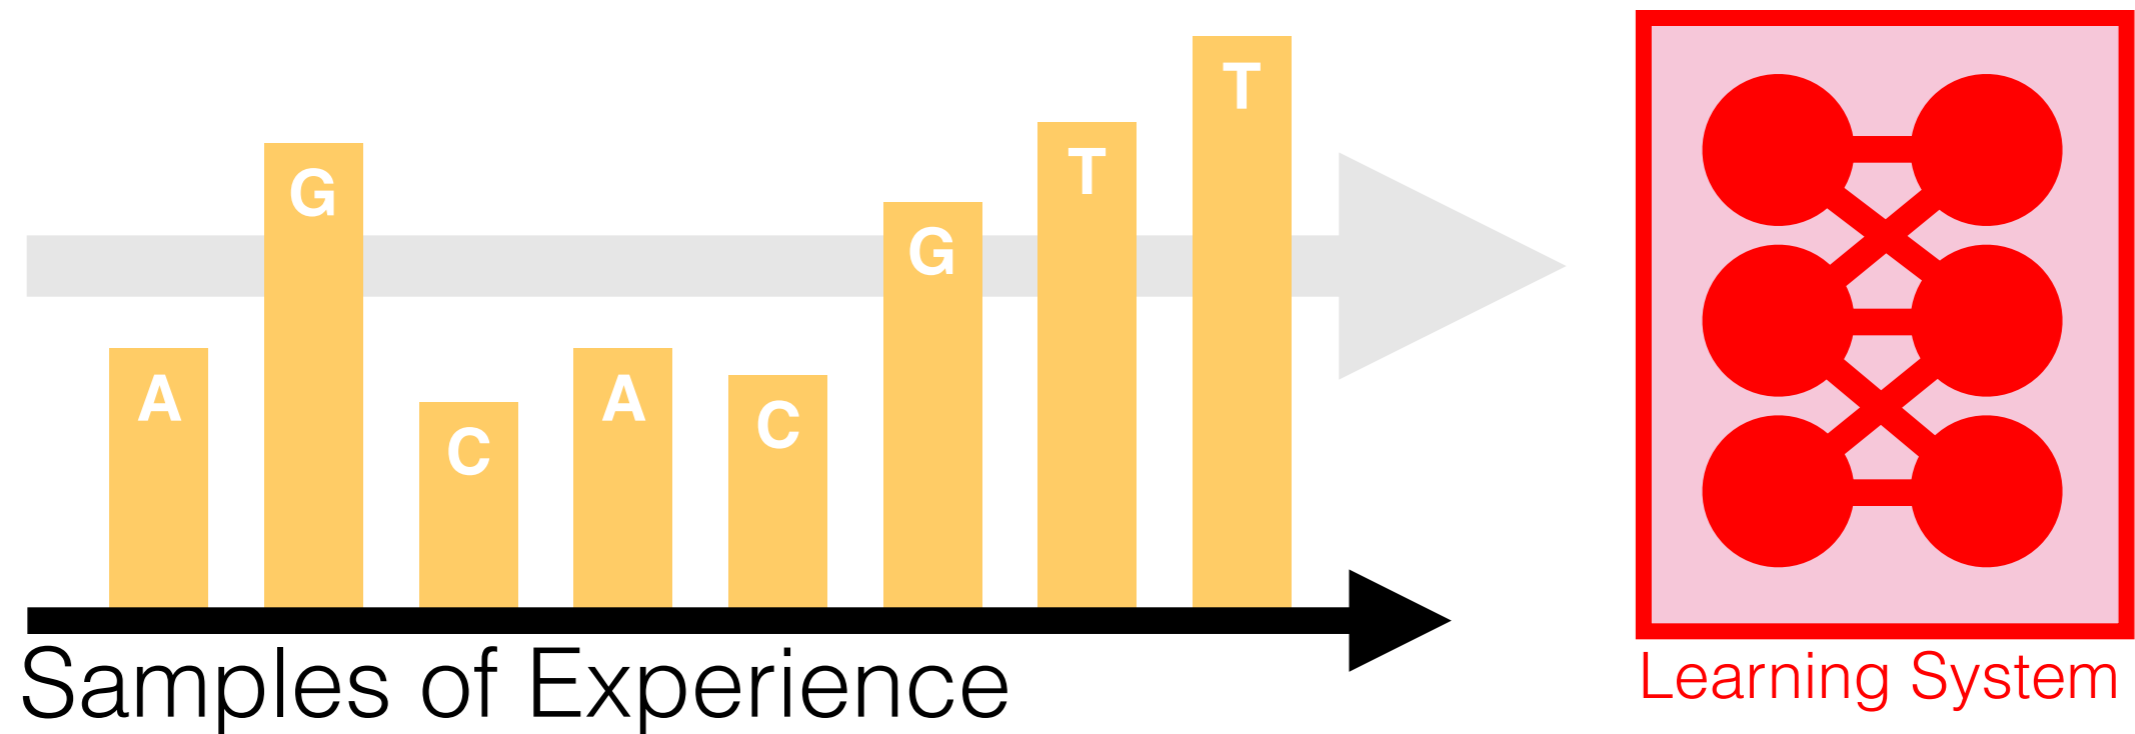


Acquire samples.

Learning from past experience.




Offline Learning

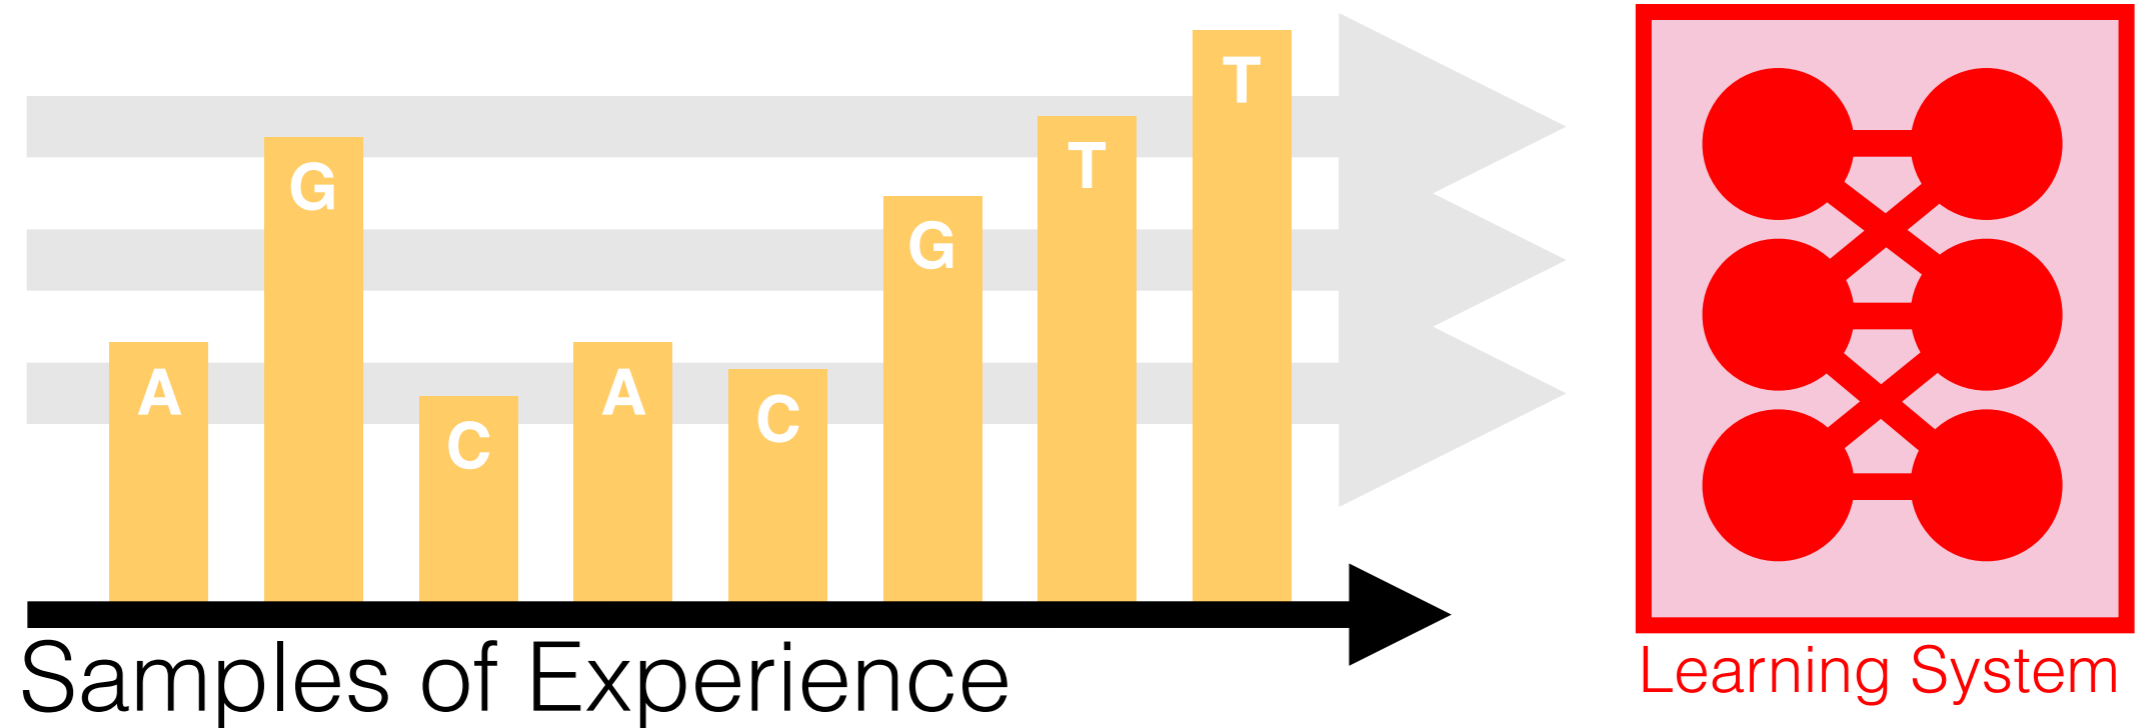


Update the learning system.

Learning from past experience.



Offline Learning

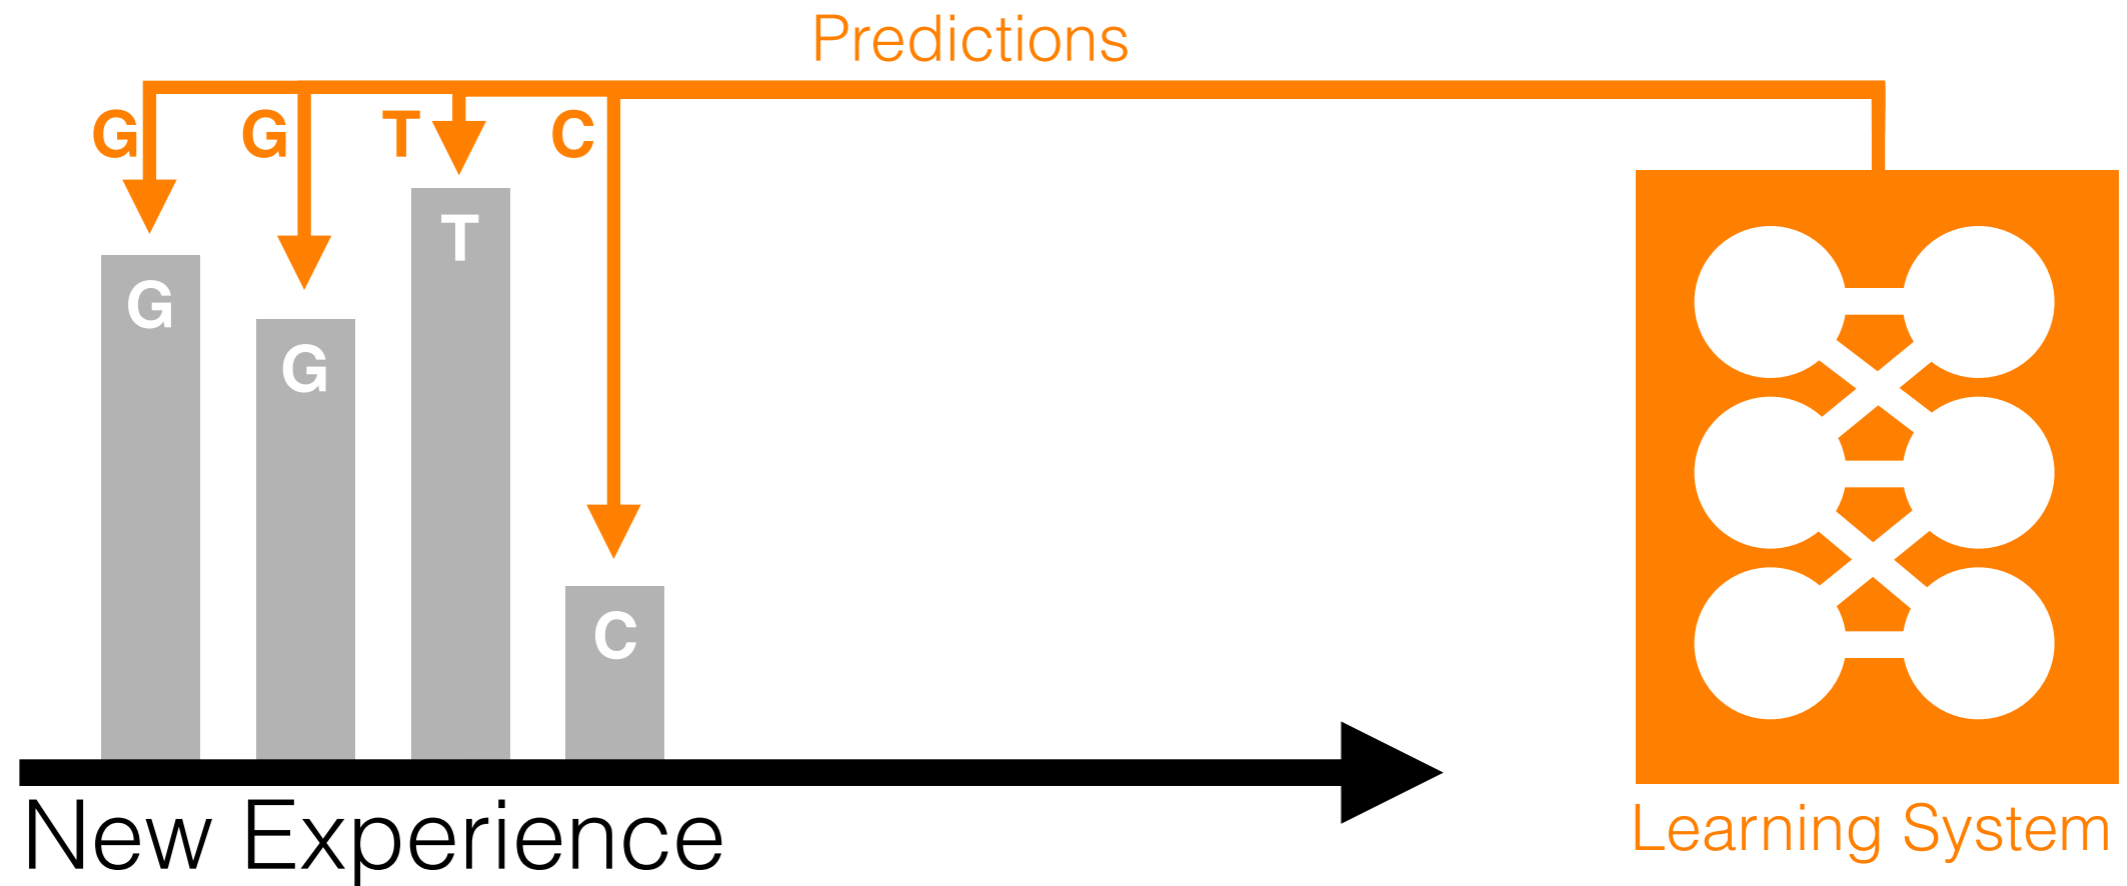


... often more than once.

Learning from past experience.



Offline Learning





Make predictions.


Learning from past experience.



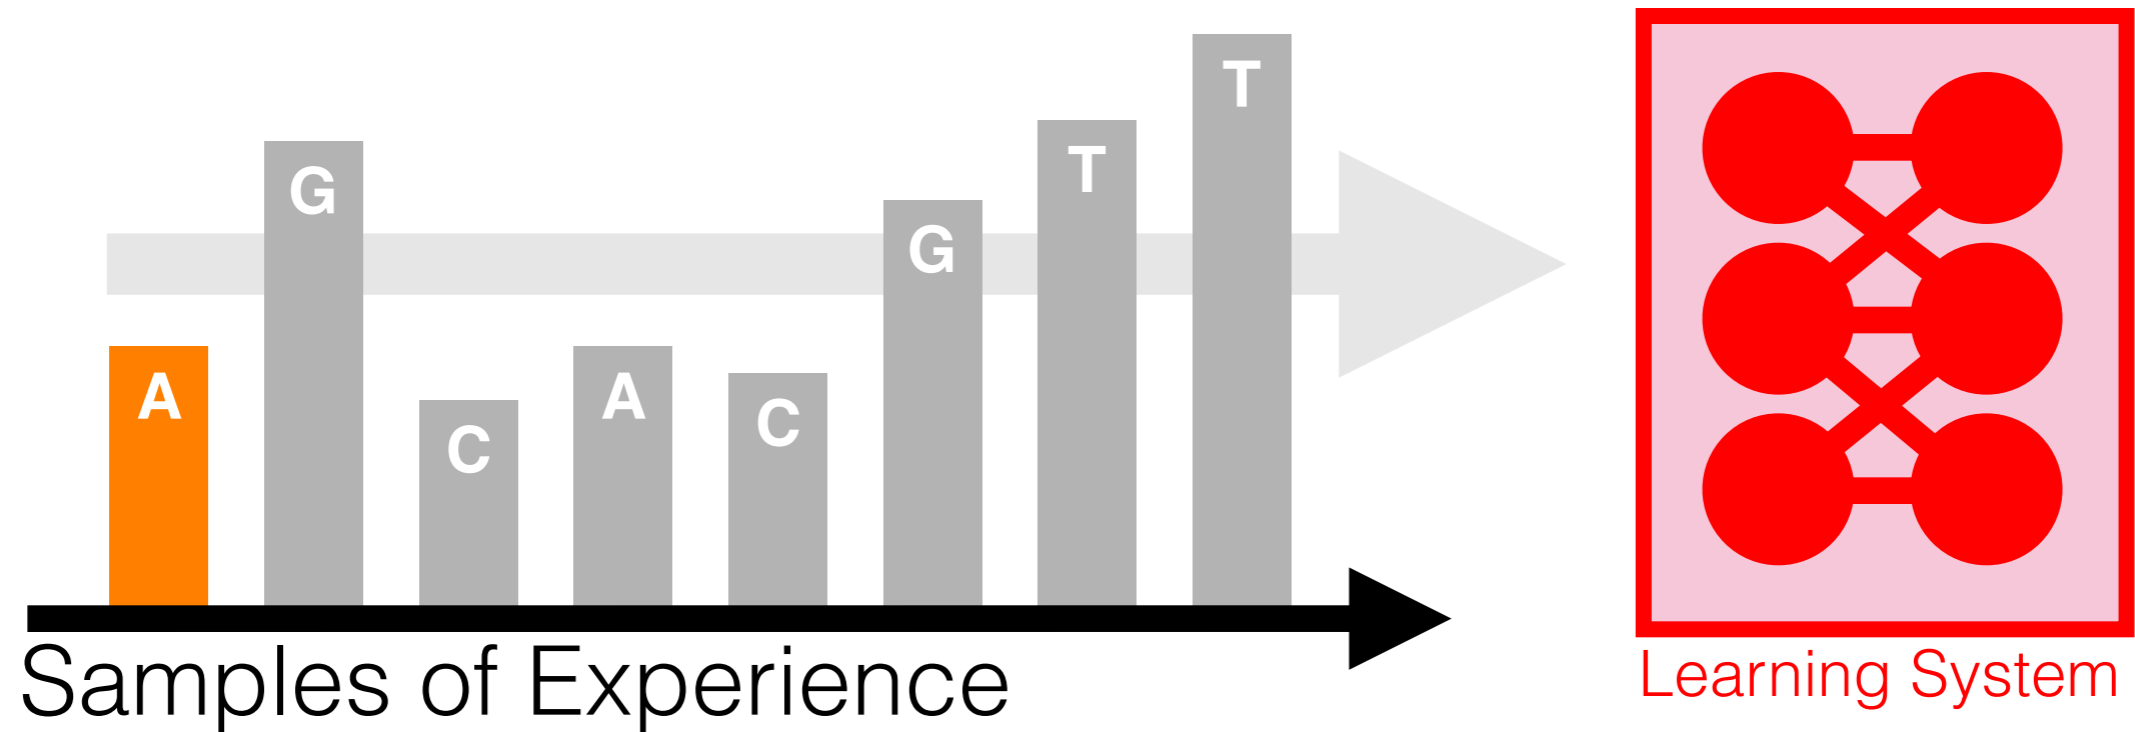
Real-time Learning



Learning from ongoing experience.




Real-time Learning

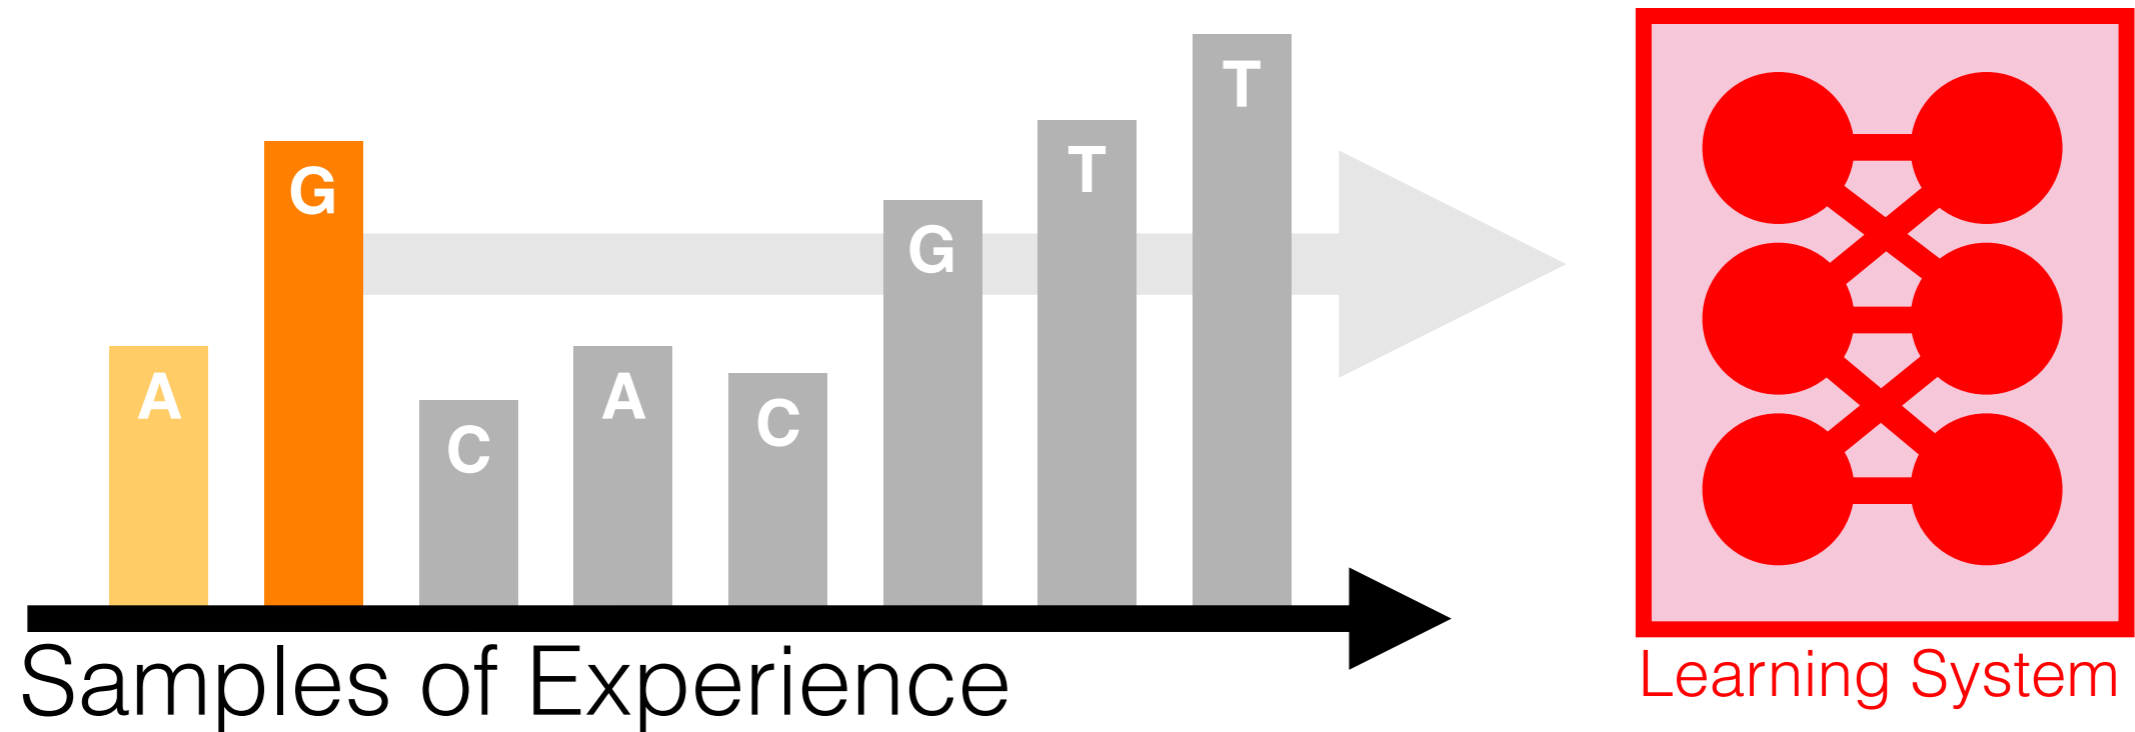


Acquire and update.

Learning from ongoing experience.



Real-time Learning

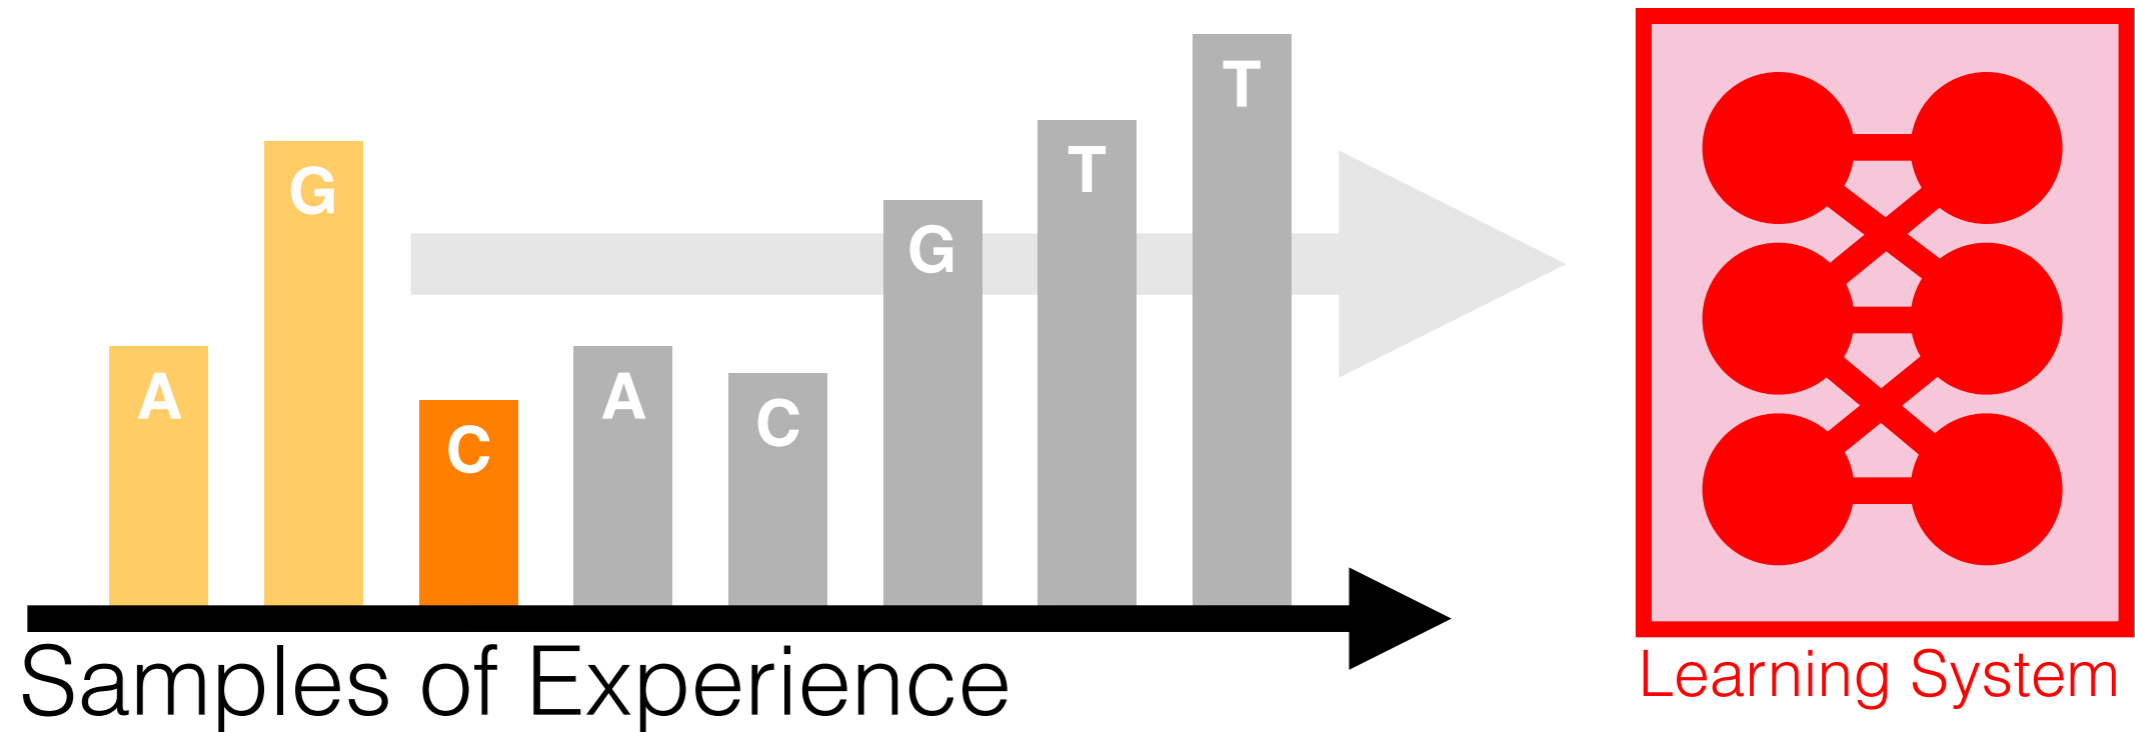


Acquire and update.

Learning from ongoing experience.



Real-time Learning




Acquire and update.

Learning from ongoing experience.




Real-time Learning

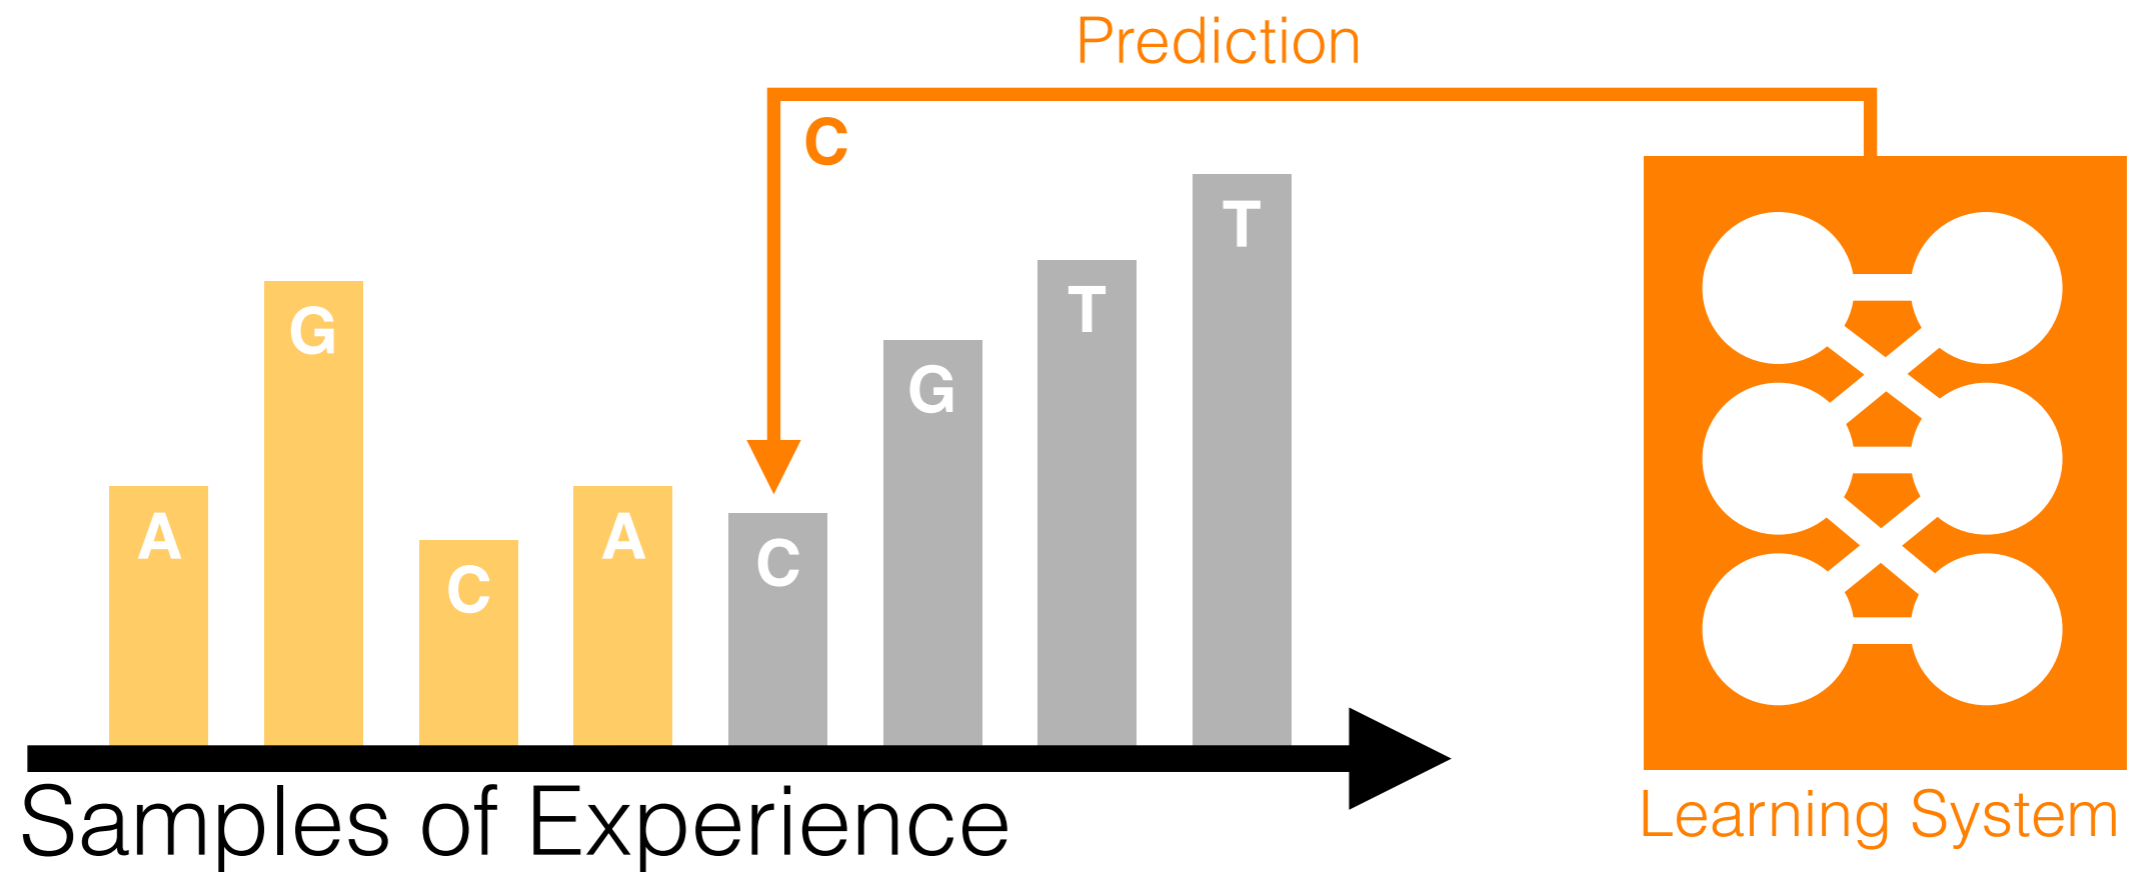


Acquire and update.

Learning from ongoing experience.




Real-time Learning




... while making predictions.

Learning from ongoing experience.



KEY IDEA

Many possible (*compatible and interchangeable*) ways for a machine to approach the acquisition and utilization of knowledge.



Learning Summary

- Be able to **define** and **discuss** machine intelligence, and be able to state **why** machine intelligence is important to society.
- Understand **what** machines might learn (representation, prediction, and control learning.)
- Understand **how** machines can learn about their world. (three learning approaches.)
- Understand **when** machines might learn. (online/real-time vs offline learning.)
- **Imagine** the promise and perils of intelligent machines.

Cheat-sheet: P.M. Pilarski, *Alberta ICT Magazine*, 2nd Ed., pp. 31
<http://www.ualberta.ca/~pilarski/docs/papers/Pilarski-Learning-AlbertaICTMagazine2012.pdf>
















The field of AI *needs* poetry.
And it *needs* perception.



c.f., the geopoetics
of **Don McKay**



e.g., TFAATJ on tour in London with **Kim Solez**

e.g., **Dennis Lee** and **Christopher Dewdney**

QUESTIONS

pilarski@ualberta.ca

<http://www.ualberta.ca/~pilarski/>