Online System Identification through Active Causal Learning

Presentation at Ericsson Research
November 18, 2025

Kim Hammar
kimham@kth.se

Paper: *https://arxiv.org/pdf/2509.02130*
(*Kim Hammar and Rolf Stadler*)

Problem: Causal System Identification

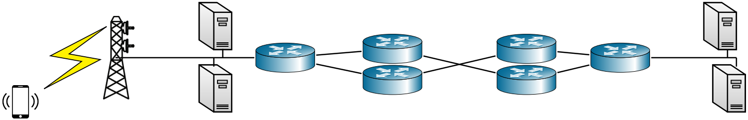

- **Networked system**: e.g., cloud system or mobile network.

Problem: Causal System Identification

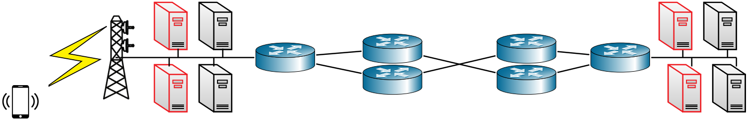

- **Networked system**: e.g., service mesh or cloud infrastructure.
- How does the resource allocation affect service response time?

Problem: Causal System Identification

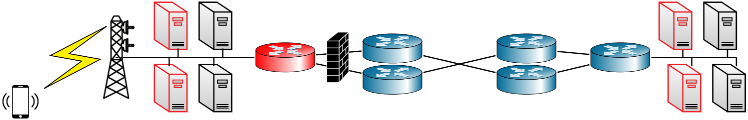

- **Networked system**: e.g., service mesh or cloud infrastructure.
- How does the resource allocation affect service response time?
- How does the network configuration affect service availability?

Problem: Causal System Identification

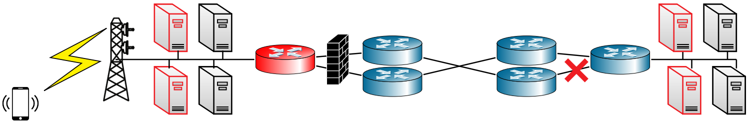

- **Networked system**: e.g., service mesh or cloud infrastructure.
- How does the resource allocation affect service response time?
- How does the network configuration affect service availability?
- How does a link failure impact network congestion?

Problem: Causal System Identification

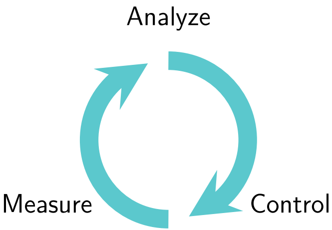

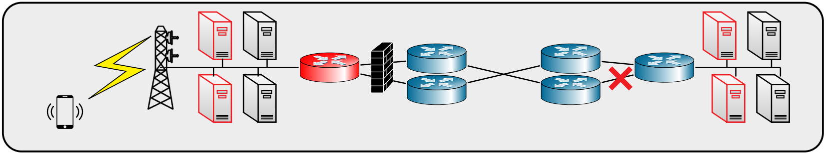

- **Networked system**: e.g., service mesh or cloud infrastructure.
- How does the resource allocation affect service response time?
- How does the network configuration affect service availability?
- How does a link failure impact network congestion?

Causal Reasoning is Needed for Reliable Operation

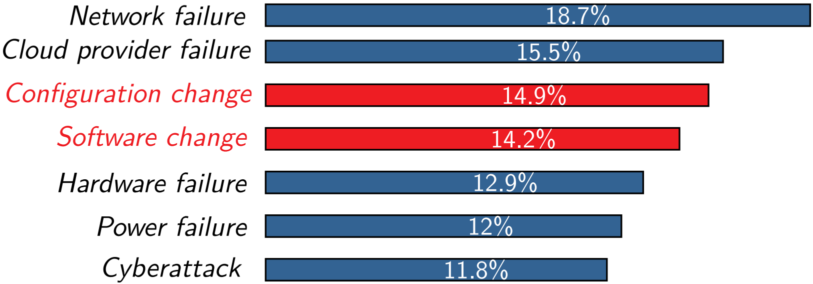

Causes of outages in networked systems. **Source: Dynatrace**.

Causal Reasoning is Needed for Reliable Operation

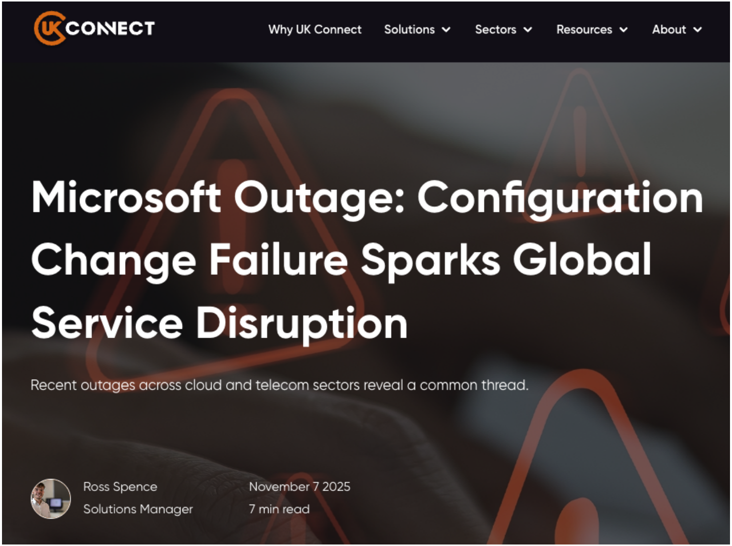

Causal Models Can Enable **Autonomous Management**

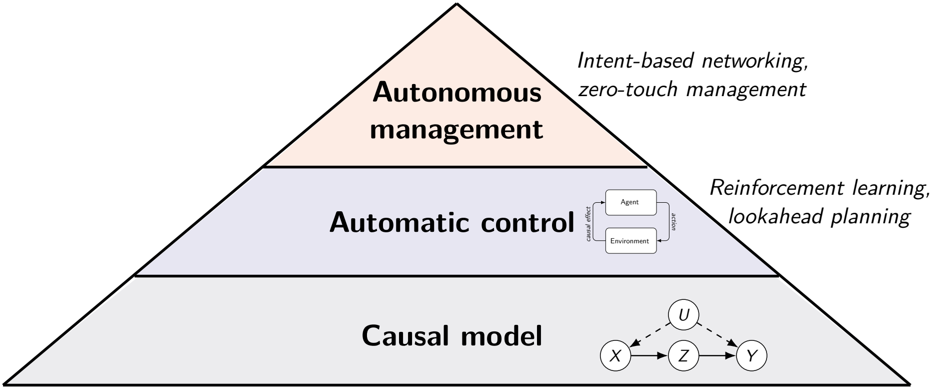

Our Method: (Online) Active Causal Learning

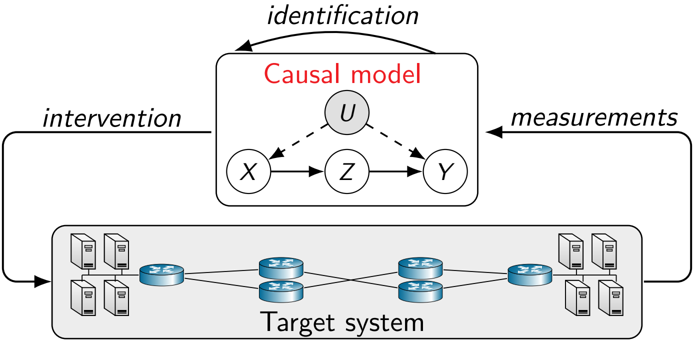

- **Target system**: e.g., service mesh or cloud infrastructure.
- **Measurements**: system metrics, e.g., CPU utilization.
- **Intervention**: change variables, e.g., resource allocation.
- **Causal model**: model of causal effects.
  - e.g., how does system load affect response time?

Outline

- **Structural Causal Models.**
  - *Theoretical background.*
- **Problem Formulation.**
  - *Online identification of an IT system.*
- **Active Causal Learning.**
  - *Our method for online system identification.*
- **Experimental Evaluation.**
  - *Target system: service mesh on our testbed.*
- **Conclusion.**
  - *Takeaways.*

Outline

- **Structural Causal Models.**
  - *Theoretical background.*
- **Problem Formulation.**
  - *Online identification of an IT system.*
- **Active Causal Learning.**
  - *Our method for online system identification.*
- **Experimental Evaluation.**
  - *Target system: service mesh on our testbed.*
- **Conclusion.**
  - *Takeaways.*

Outline

- **Structural Causal Models.**
  - *Theoretical background.*
- **Problem Formulation.**
  - *Online identification of an IT system.*
- **Active Causal Learning.**
  - *Our method for online system identification.*
- **Experimental Evaluation.**
  - *Target system: service mesh on our testbed.*
- **Conclusion.**
  - *Takeaways.*

Outline

- **Structural Causal Models.**
  - *Theoretical background.*
- **Problem Formulation.**
  - *Online identification of an IT system.*
- **Active Causal Learning.**
  - *Our method for online system identification.*
- **Experimental Evaluation.**
  - *Target system: service mesh on our testbed.*
- **Conclusion.**
  - *Takeaways.*

Outline

- **Structural Causal Models.**
  - *Theoretical background.*
- **Problem Formulation.**
  - *Online identification of an IT system.*
- **Active Causal Learning.**
  - *Our method for online system identification.*
- **Experimental Evaluation.**
  - *Target system: service mesh on our testbed.*
- **Conclusion.**
  - *Takeaways.*

The Causal Graph

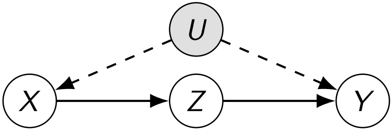

- Directed acyclic graph.
- Encodes **causal structure** of the system.
  - Nodes: system variables.
  - Edges: causal dependencies.
- Two types of variables: endogenous and exogenous (shaded).
  - Endogenous variables: *internal* properties (e.g., response time).
  - Exogenous variables: *external* factors (e.g., system load).

Causal Structure - Example

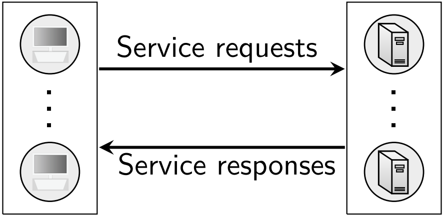

Causal Structure - Example

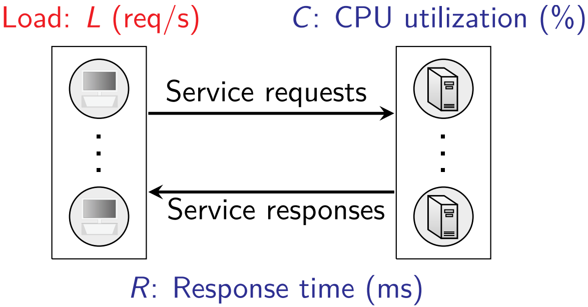

- Load (L) is exogenous (external to the system).
- CPU utilization (C) and response time (R) are endogenous.

Causal Structure - Example

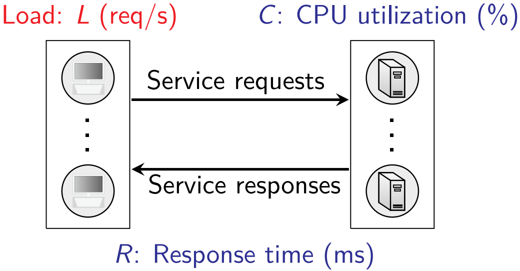

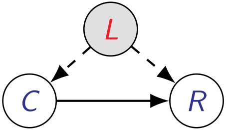

- Load (L) is exogenous (external to the system).
- CPU utilization (C) and response time (R) are endogenous.

Causal Functions

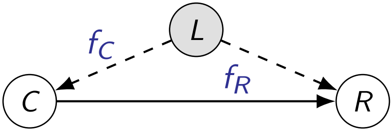

- Endogenous variables are associated with **causal functions**.
  - $C = f_C(L)$.
  - $R = f_R(C, L)$.

Causal Functions

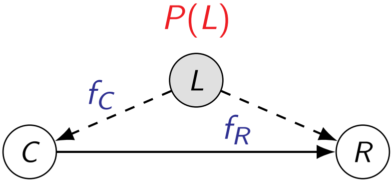

- Endogenous variables are associated with **causal functions**.
  - $C = f_C(L)$.
  - $R = f_R(C, L)$.

- Exogenous variables are associated with **distributions**.
  - $L \sim P(L)$.

Causal Functions

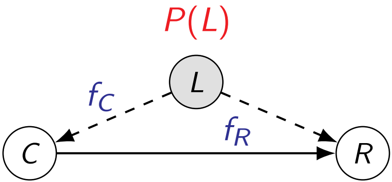

- Endogenous variables are associated with **causal functions**.
  - $C = f_C(L)$.
  - $R = f_R(C, L)$.

- Exogenous variables are associated with **distributions**.
  - $L \sim P(L)$.

Structural causal model (SCM)

- Graph $\mathcal{G}$, functions $\mathbf{F}$, distribution P.

Outline

- **Structural Causal Models.**
  - *Theoretical background.*
- **Problem Formulation.**
  - *Online identification of an IT system.*
- **Active Causal Learning.**
  - *Our method for online system identification.*
- **Experimental Evaluation.**
  - *Target system: service mesh on our testbed.*
- **Conclusion.**
  - *Takeaways.*

Problem Setup

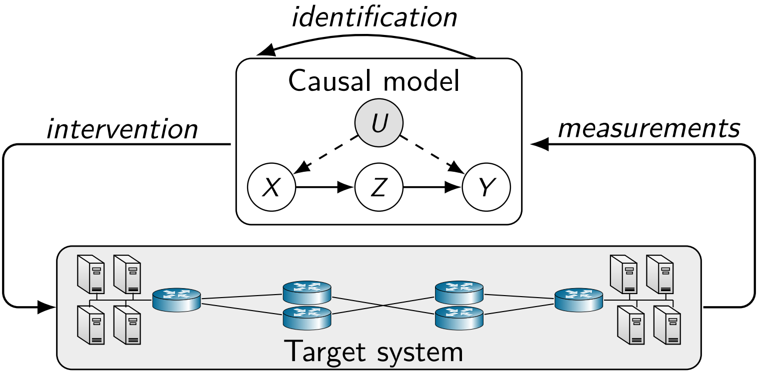

- We consider an IT system modeled by an SCM.
- The causal graph $\mathcal{G}$ and the distribution $P(\mathbf{U})$ are known.
- The causal functions $\mathbf{F}$ are unknown and time-varying.
- **Problem**: identify $\mathbf{F}$ from online system measurements.

Problem Setup

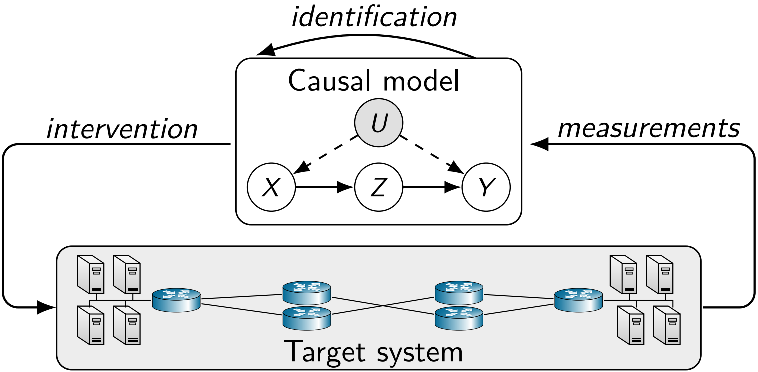

- We consider an IT system modeled by an SCM.
- The causal graph $\mathcal{G}$ and the distribution $P(\mathbf{U})$ are known.
- The causal functions $\mathbf{F}$ are unknown and time-varying.
- **Problem**: identify $\mathbf{F}$ from online system measurements.

Problem Setup

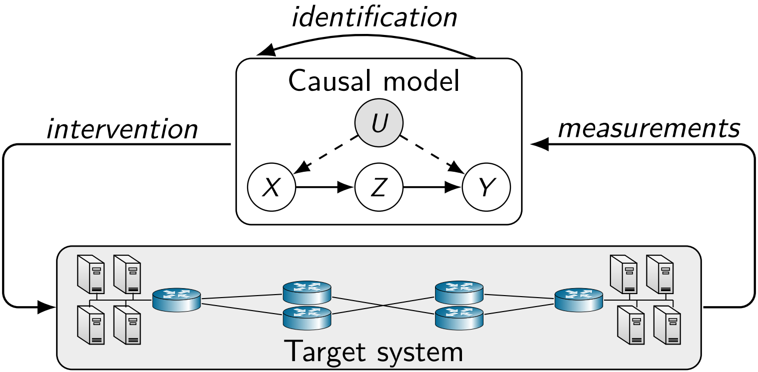

- We consider an IT system modeled by an SCM.
- The causal graph $\mathcal{G}$ and the distribution $P(\mathbf{U})$ are known.
- The causal functions $\mathbf{F}$ are unknown and time-varying.
- **Problem**: identify $\mathbf{F}$ from online system measurements.

Problem Setup

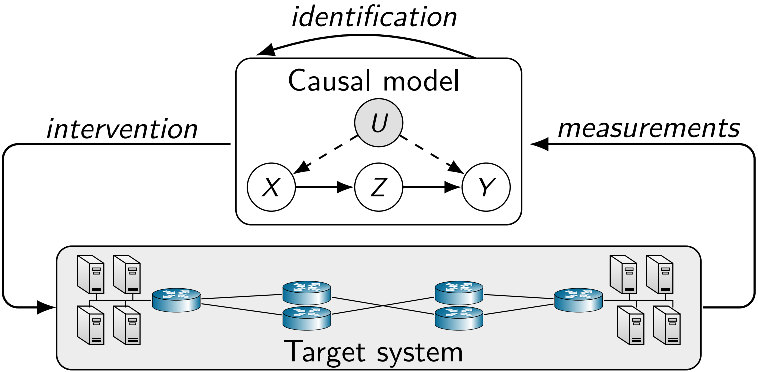

- We consider an IT system modeled by an SCM.
- The causal graph $\mathcal{G}$ and the distribution $P(\mathbf{U})$ are known.
- The causal functions $\mathbf{F}$ are unknown and time-varying.
- **Problem**: identify $\mathbf{F}$ from online system measurements.

Problem Statement

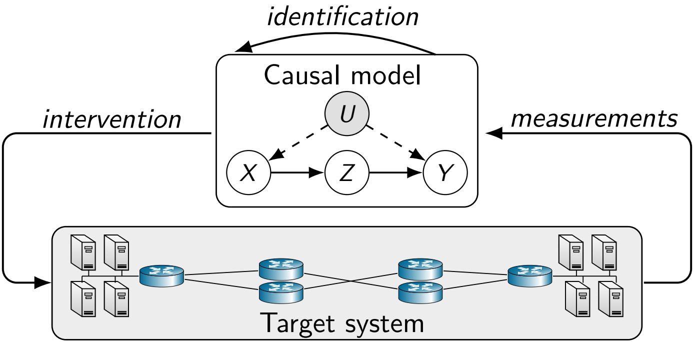

Select a sequence of **interventions** and use the resulting system measurements to estimate a sequence of functions

$$\hat{\mathbf{F}}_1, \hat{\mathbf{F}}_2, \ldots$$

that are as close as possible to the true sequence $\mathbf{F}_1, \mathbf{F}_2, \ldots$ while minimizing the operational cost of the interventions.

Interventions

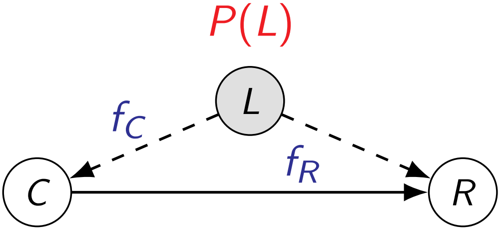

Interventions

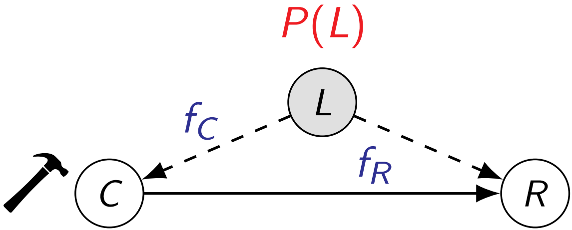

- Suppose that we intervene on the system and artifically increase the CPU utilization to $c\%$.

Interventions

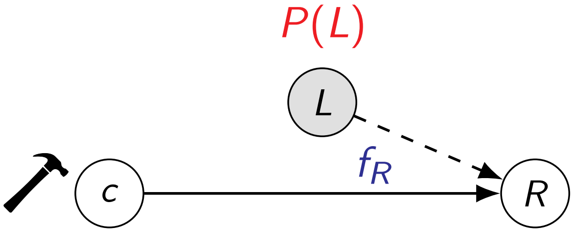

- Suppose that we intervene on the system and artifically increase the CPU utilization to c%.

Interventions

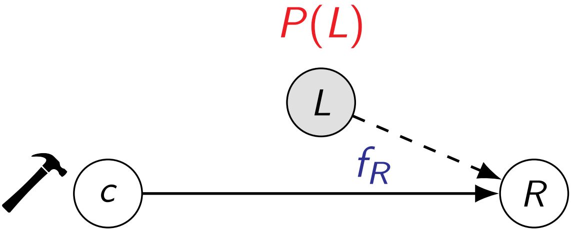

- Suppose that we intervene on the system and artifically increase the CPU utilization to $c\%$.

- We represent the intervention with the do-operator.
  - An **intervention** $\mathrm{do}(C = c)$ fixes the variable C to value c.

Interventions

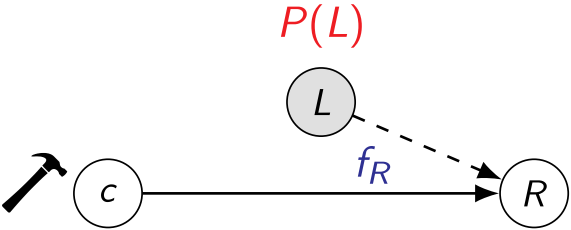

- Suppose that we intervene on the system and artifically increase the CPU utilization to c%.
- We represent the intervention with the do-operator.
  - An **intervention** $\mathrm{do}(C = c)$ fixes the variable C to value c.
- Allows to sample the causal function $f_R(C = c, L)$.

Example Causal Function

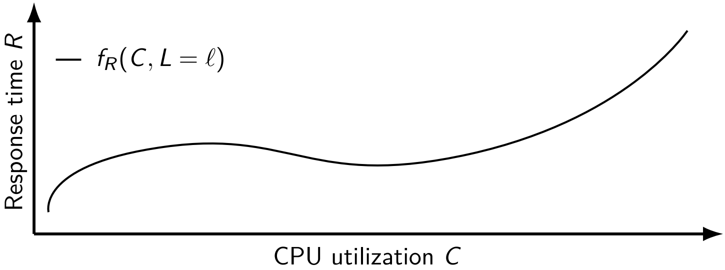

Estimating a Causal Function from Observational Data

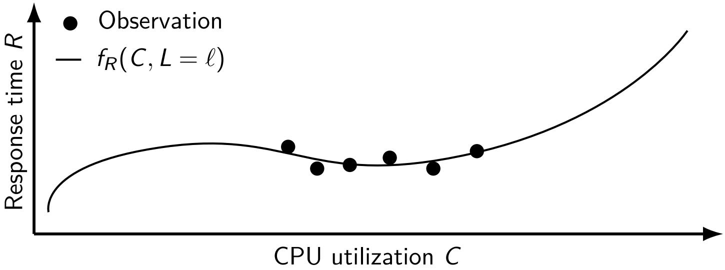

Estimating a Causal Function from Observational Data

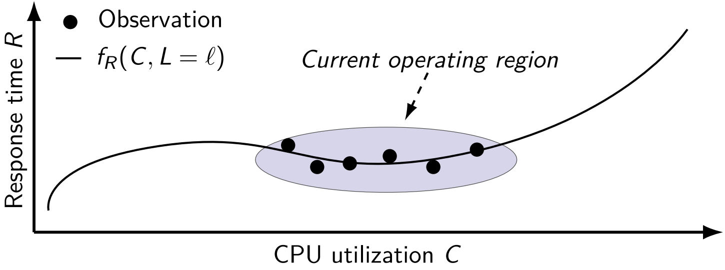

- Monitoring yields samples from the **current** operating region.
  - e.g., under current CPU utilization, response time is 10-15 ms.

Estimating a Causal Function from Interventional Data

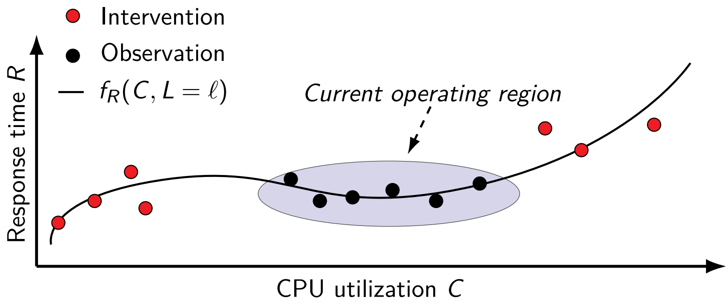

- Monitoring yields samples from the **current** operating region.
  - e.g., under current CPU utilization, response time is 10-15 ms.

- Intervention allows to explore the **complete** operating region.
  - e.g., under max CPU utilization, response time is 100-200 ms.

Problem Formalization

- We capture the cost of an intervention through a cost function $c(\mathrm{do}(\mathbf{X}' = \mathbf{x}'))$, which is system specific.
- We capture the difference between the estimate $\hat{\mathbf{F}}_t$ and the ground truth $\mathbf{F}_t$ through the loss function

$$\mathscr{L}(\hat{\mathbf{F}}_t, \mathbf{F}_t) = \sum_{V_i \in \mathbf{V}} \int_{\mathcal{R}(\mathrm{pa}_{\mathcal{G}}(V_i))} \Big(f_{V_i,t}(\mathbf{x}) - \hat{f}_{V_i,t}(\mathbf{x}) \Big)^2 \, \mathbb{P}[\mathrm{d}\mathbf{x}],$$

where $\mathbf{V}$ is the set of endogenous variables and $\mathrm{pa}_{\mathcal{G}}(V_i))$ is the set of parents of V_i in the causal graph $\mathcal{G}$.

Objective

$$\underset{(\mathrm{do}(\mathbf{X}'_t = \mathbf{x}'_t), \hat{\mathbf{F}}_t)_{t \geq 1}}{\text{minimize}} \sum_{t=1}^{\infty} \gamma^{t-1} \Big(\mathscr{L}(\hat{\mathbf{F}}_t, \mathbf{F}_t) + c(\mathrm{do}(\mathbf{X}'_t = \mathbf{x}'_t)) \Big),$$

where $\gamma \in (0, 1)$ is a discount factor.

Problem Formalization

- We capture the cost of an intervention through a cost function $c(\mathrm{do}(\mathbf{X}' = \mathbf{x}'))$, which is system specific.
- We capture the difference between the estimate $\hat{\mathbf{F}}_t$ and the ground truth $\mathbf{F}_t$ through the loss function

$$\mathscr{L}(\hat{\mathbf{F}}_t, \mathbf{F}_t) = \sum_{V_i \in \mathbf{V}} \int_{\mathcal{R}(\mathrm{pa}_{\mathcal{G}}(V_i))} \Big(f_{V_i,t}(\mathbf{x}) - \hat{f}_{V_i,t}(\mathbf{x}) \Big)^2 \, \mathbb{P}[\mathrm{d}\mathbf{x}],$$

where $\mathbf{V}$ is the set of endogenous variables and $\mathrm{pa}_{\mathcal{G}}(V_i))$ is the set of parents of V_i in the causal graph $\mathcal{G}$.

Objective

$$\underset{(\mathrm{do}(\mathbf{X}'_t = \mathbf{x}'_t), \hat{\mathbf{F}}_t)_{t \geq 1}}{\text{minimize}} \sum_{t=1}^{\infty} \gamma^{t-1} \Big(\mathscr{L}(\hat{\mathbf{F}}_t, \mathbf{F}_t) + c(\mathrm{do}(\mathbf{X}'_t = \mathbf{x}'_t)) \Big),$$

where $\gamma \in (0, 1)$ is a discount factor.

Problem Formalization

- We capture the cost of an intervention through a cost function $c(\mathrm{do}(\mathbf{X}' = \mathbf{x}'))$, which is system specific.
- We capture the difference between the estimate $\hat{\mathbf{F}}_t$ and the ground truth $\mathbf{F}_t$ through the loss function

$$\mathscr{L}(\hat{\mathbf{F}}_t, \mathbf{F}_t) = \sum_{V_i \in \mathbf{V}} \int_{\mathcal{R}(\mathrm{pa}_{\mathcal{G}}(V_i))} \left(f_{V_i,t}(\mathbf{x}) - \hat{f}_{V_i,t}(\mathbf{x}) \right)^2 \mathbb{P}[\mathrm{d}\mathbf{x}],$$

where $\mathbf{V}$ is the set of endogenous variables and $\mathrm{pa}_{\mathcal{G}}(V_i))$ is the set of parents of V_i in the causal graph $\mathcal{G}$.

Objective

$$\underset{(\mathrm{do}(\mathbf{X}'_t = \mathbf{x}'_t), \hat{\mathbf{F}}_t)_{t \geq 1}}{\text{minimize}} \sum_{t=1}^{\infty} \gamma^{t-1} \Big(\mathscr{L}(\hat{\mathbf{F}}_t, \mathbf{F}_t) + c\big(\mathrm{do}(\mathbf{X}'_t = \mathbf{x}'_t)\big) \Big),$$

where $\gamma \in (0, 1)$ is a discount factor.

Outline

- **Structural Causal Models.**
  - *Theoretical background.*
- **Problem Formulation.**
  - *Online identification of an IT system.*
- **Active Causal Learning.**
  - *Our method for online system identification.*
- **Experimental Evaluation.**
  - *Target system: service mesh on our testbed.*
- **Conclusion.**
  - *Takeaways.*

Active Causal Learning

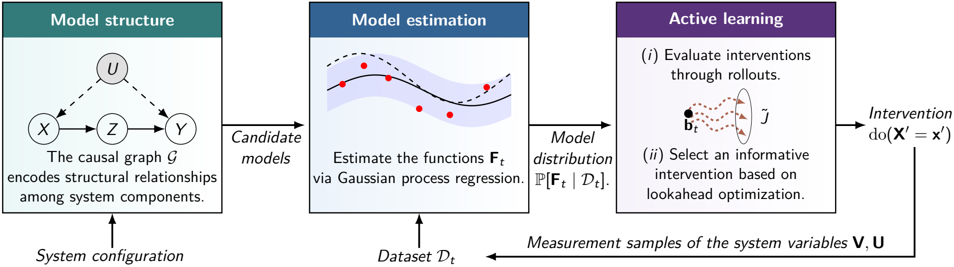

Active Causal Learning

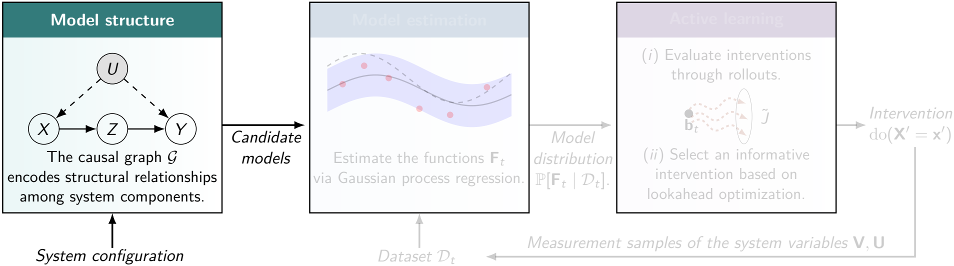

Active Causal Learning

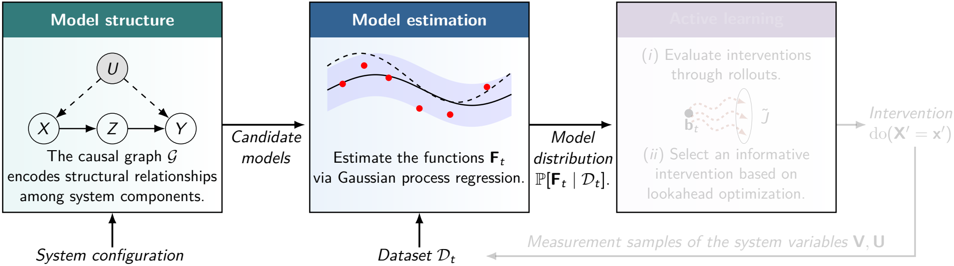

- Estimate the causal functions **F** with Gaussian processes.

Active Causal Learning

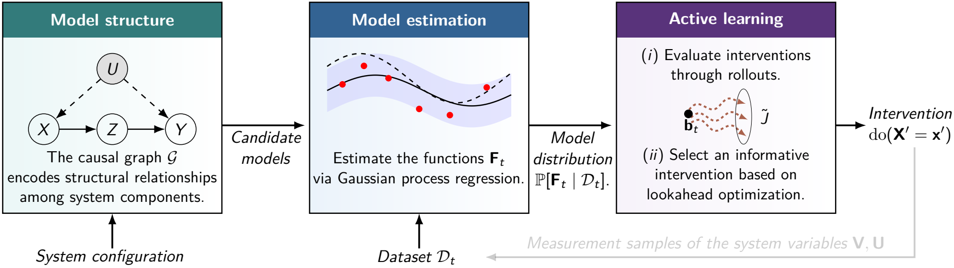

- Estimate the causal functions $\mathbf{F}$ with Gaussian processes.
- Select interventions through a rollout policy that balances
  - **Exploration**: reducing model uncertainty.
  - **Operational cost**: avoiding costly interventions.

Active Causal Learning

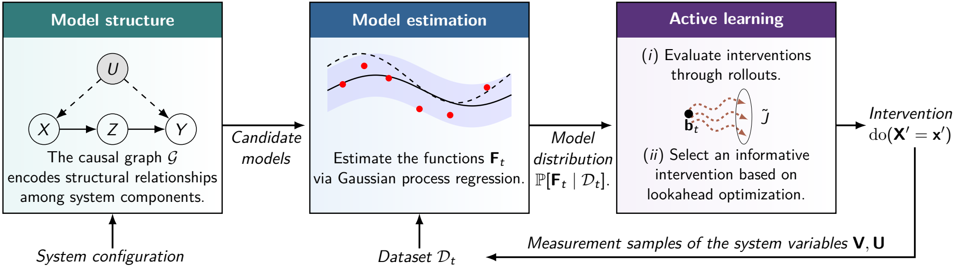

- Estimate the causal functions $\mathbf{F}$ with Gaussian processes.
- Select interventions through a rollout policy that balances
  - **Exploration**: reducing model uncertainty.
  - **Operational cost**: avoiding costly interventions.

Gaussian Process Regression

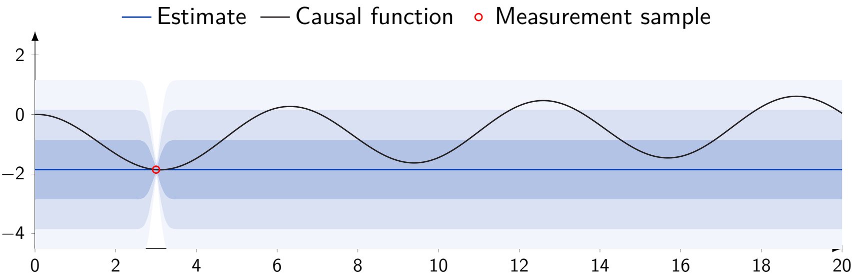

- We estimate each causal function with a Gaussian process.
- Sequentially update the estimate through Bayes rule.
- The Gaussian process provides **uncertainty quantification**.
  - Informs us where to intervene.

Gaussian Process Regression

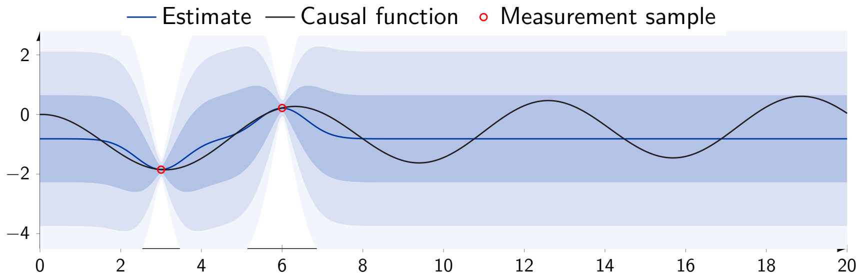

- We estimate each causal function with a Gaussian process.
- Sequentially update the estimate through Bayes rule.
- The Gaussian process provides **uncertainty quantification**.
  - Informs us where to intervene.

Gaussian Process Regression

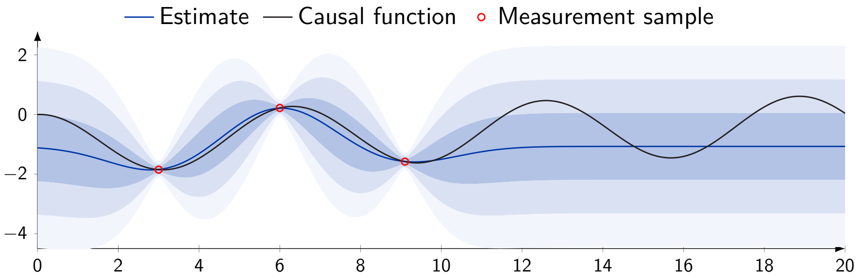

- We estimate each causal function with a Gaussian process.
- Sequentially update the estimate through Bayes rule.
- The Gaussian process provides **uncertainty quantification**.
  - Informs us where to intervene.

Gaussian Process Regression

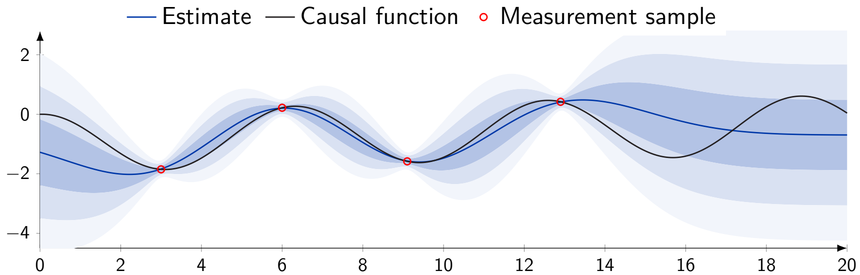

- We estimate each causal function with a Gaussian process.
- Sequentially update the estimate through Bayes rule.
- The Gaussian process provides **uncertainty quantification**.
  - Informs us where to intervene.

Gaussian Process Regression

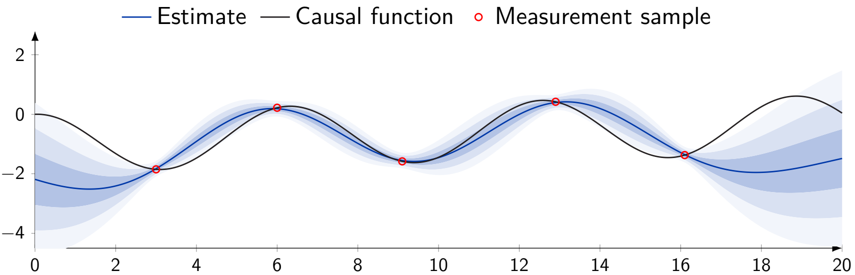

- We estimate each causal function with a Gaussian process.
- Sequentially update the estimate through Bayes rule.
- The Gaussian process provides **uncertainty quantification**.
  - Informs us where to intervene.

Gaussian Process Regression

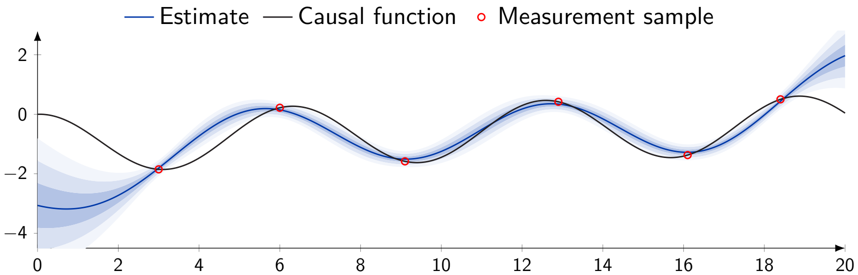

- We estimate each causal function with a Gaussian process.
- Sequentially update the estimate through Bayes rule.
- The Gaussian process provides **uncertainty quantification**.
  - Informs us where to intervene.

Gaussian Process Regression

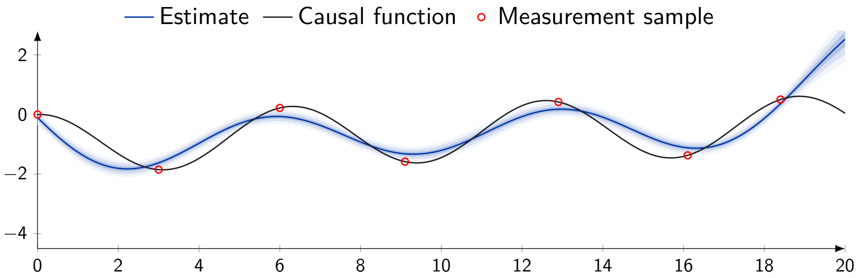

- We estimate each causal function with a Gaussian process.
- Sequentially update the estimate through Bayes rule.
- The Gaussian process provides **uncertainty quantification**.
  - Informs us where to intervene.

Active Learning through Rollout

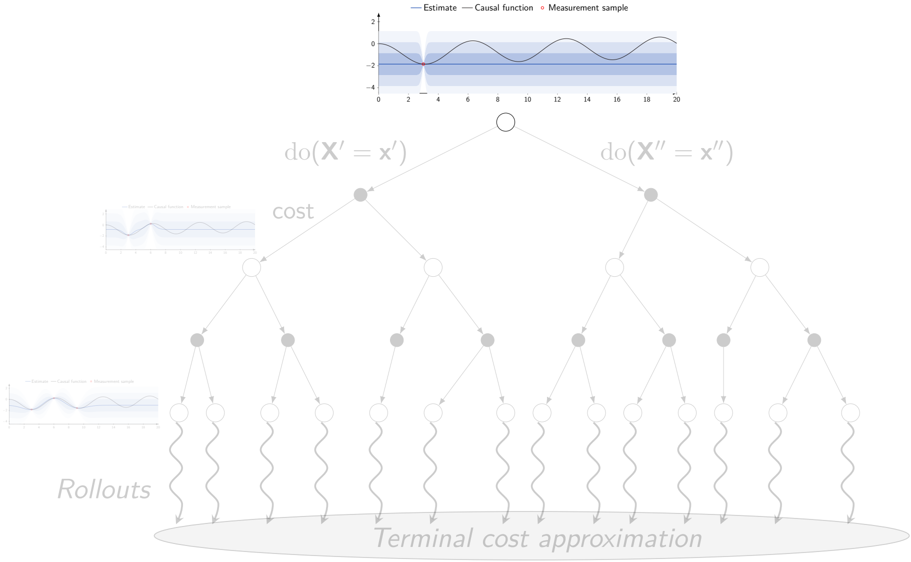

Active Learning through Rollout

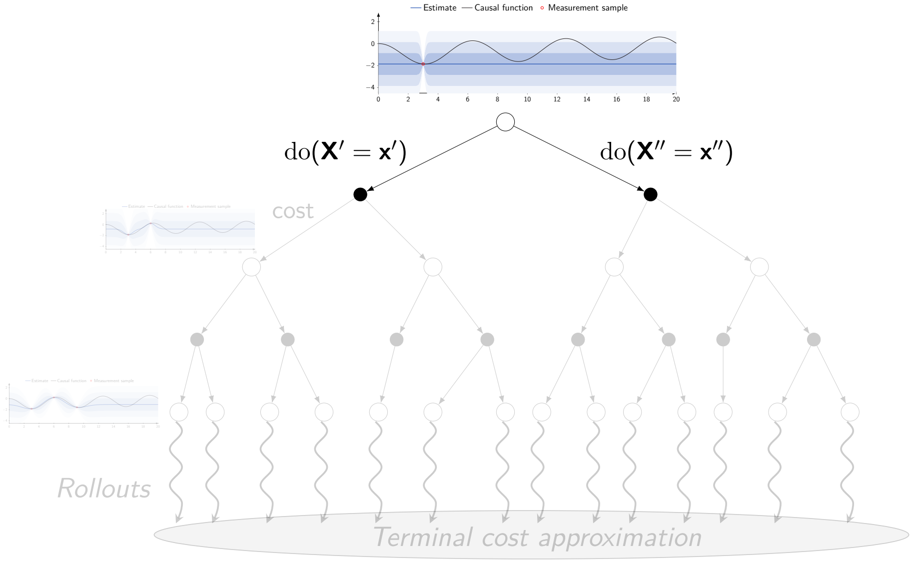

Active Learning through Rollout

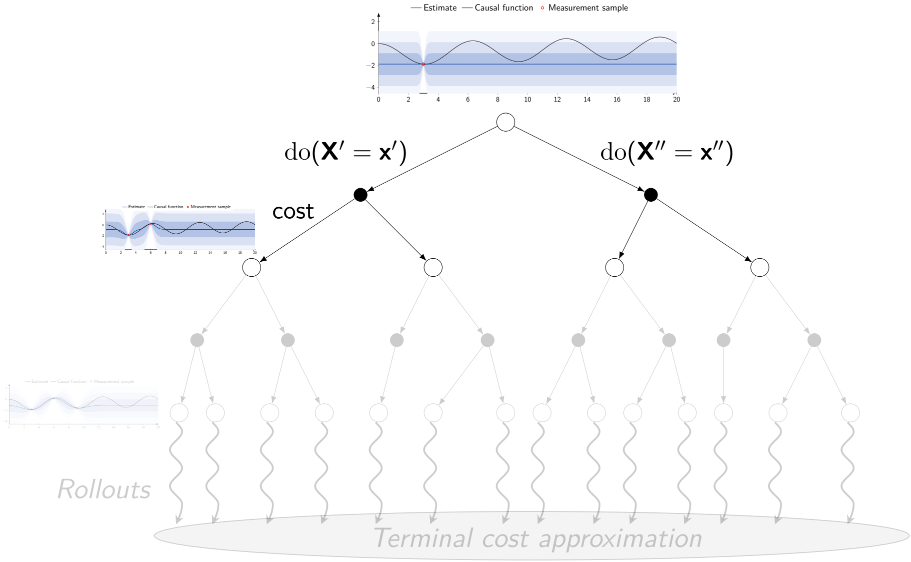

Active Learning through Rollout

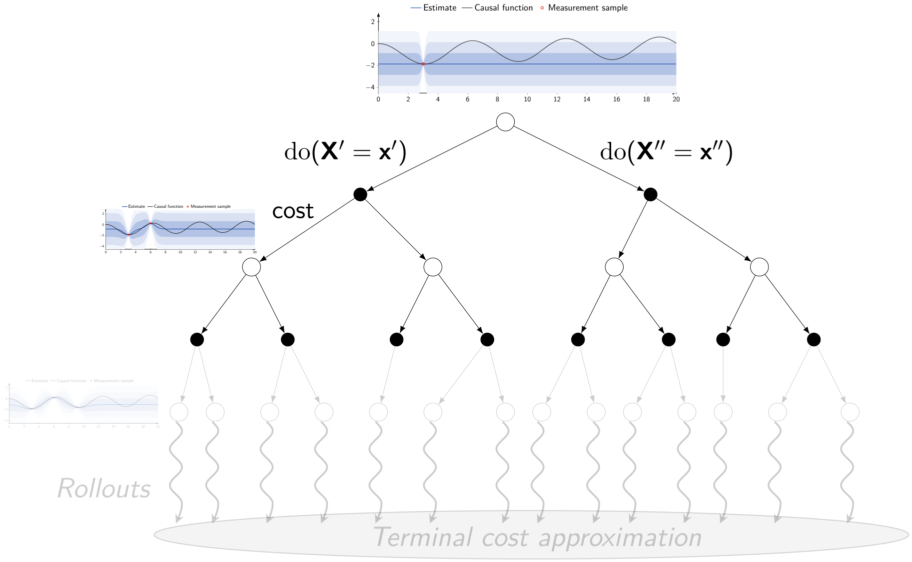

Active Learning through Rollout

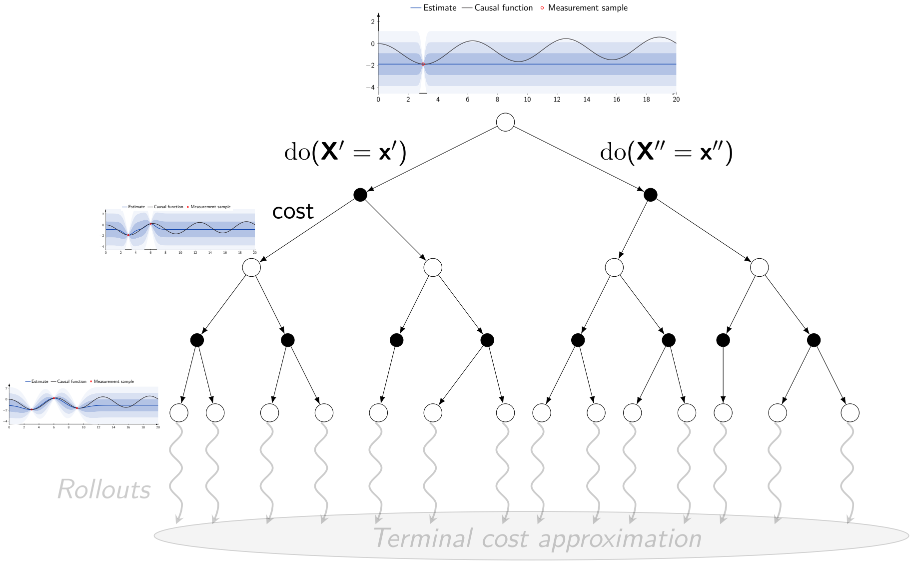

Active Learning through Rollout

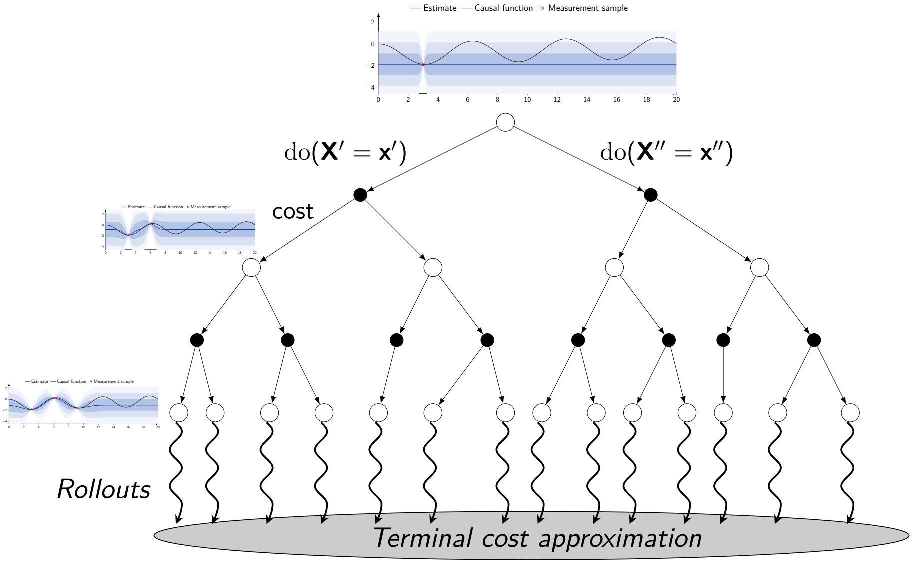

Active Learning through Rollout

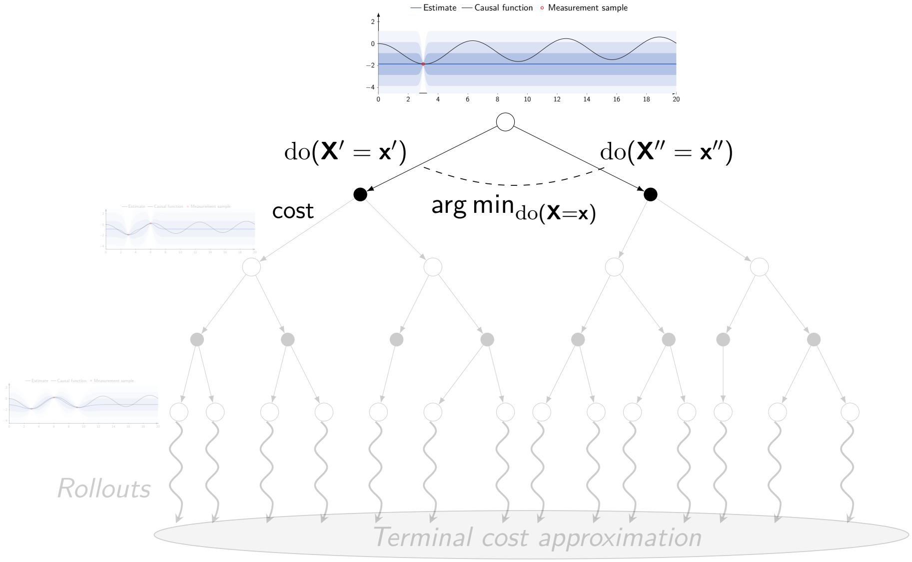

Illustrative Example

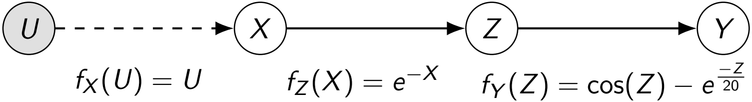

- Exogenous variables: $\mathbf{U} = \{U\}$.
- Endogenous variables: $\mathbf{V} = \{X, Z, Y\}$.
- Causal functions: $\mathbf{F} = \{f_X, f_Z, f_Y\}$.
- Probability distribution: $U \sim \mathcal{N}(0, 0.1)$.

Comparison with **Passive Learning** (1/2)

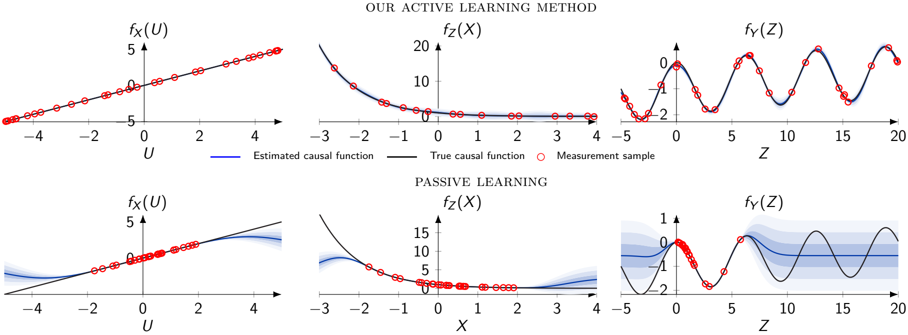

- **Passive learning**:
  - Use measurements from the current operating region.
  - No interventions.

Comparison with **Passive Learning** (1/2)

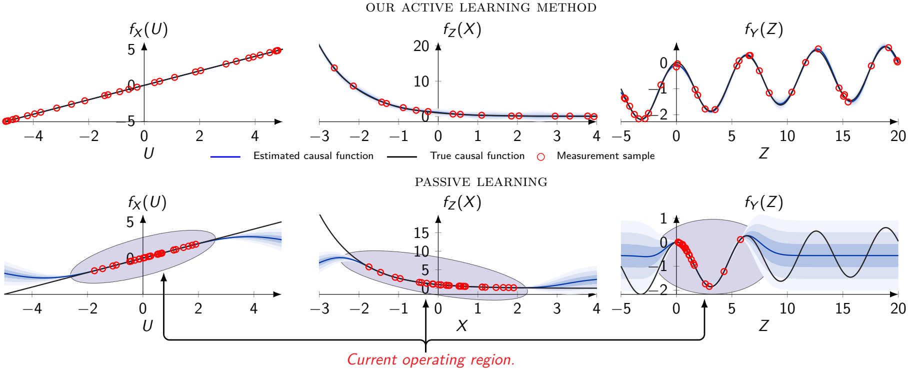

- **Passive learning**:
  - Use measurements from the current operating region.
  - No interventions.

Comparison with **Passive Learning** (2/2)

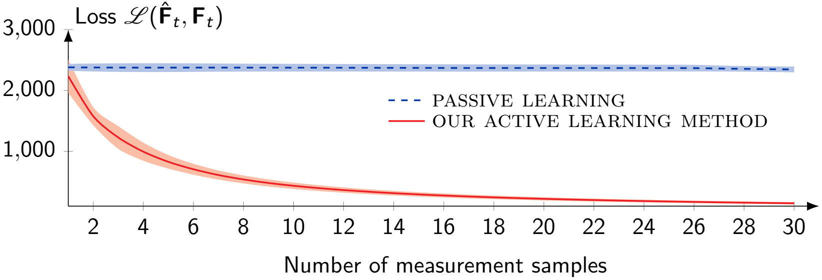

Takeaway.

Passive learning provides limited coverage of the system's operating region, leading to inaccurate model estimates.

Theoretical Properties (Informal)*

Proposition 1 (Asymptotic consistency)

If the causal function lies in the Reproducing kernel Hilbert space of the Gaussian process, then the mean of the Gaussian process converges to the causal function in the limit of infinite data.

Proposition 2 (Optimal Bayesian estimate)

Given the posterior distribution of the Gaussian process, then the mean of the Gaussian process minimizes the loss function $\mathscr{L}$.

Proposition 3 (Policy improvement)

The rollout policy improves the greedy intervention policy, i.e., the intervention policy without lookahead.

---

*For theoretical details, see our paper: "Online Identification of IT Systems through Active Causal Learning", Hammar and Stadler 2025, https://arxiv.org/pdf/2509.02130.

Outline

- **Structural Causal Models.**
  - *Theoretical background.*
- **Problem Formulation.**
  - *Online identification of an IT system.*
- **Active Causal Learning.**
  - *Our method for online system identification.*
- **Experimental Evaluation.**
  - *Target system: service mesh on our testbed.*
- **Conclusion.**
  - *Takeaways.*

Target System

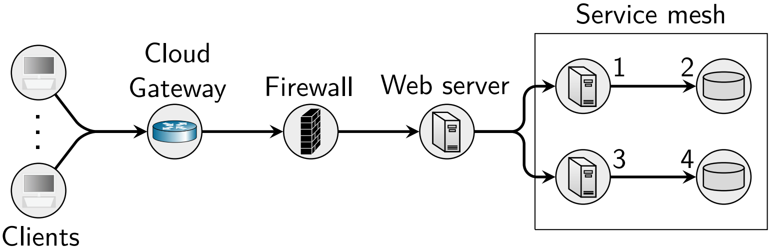

- A cloud-based web application with a service mesh backend.
- The web server is implemented using Flask.
- The service mesh is implemented with **Kubernetes and Istio**.

System Configuration

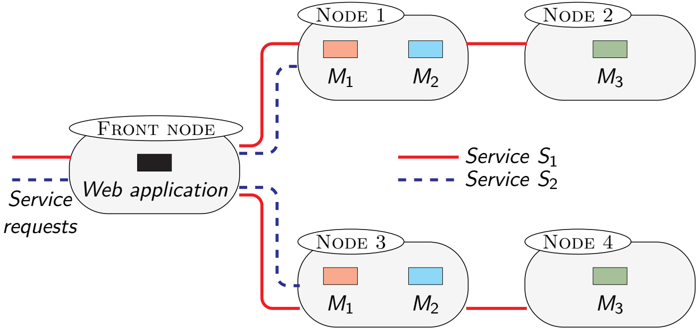

- The mesh consists of 4 physical nodes and 3 microservices.
- Nodes 1 and 3 run microservices M_1 and M_2.
- Nodes 2 and 4 run microservice M_3.
- The mesh runs **two services: S_1 and S_2.**
  - S_1 invokes M_1 and M_3.
  - S_2 invokes M_2.

Control Variables

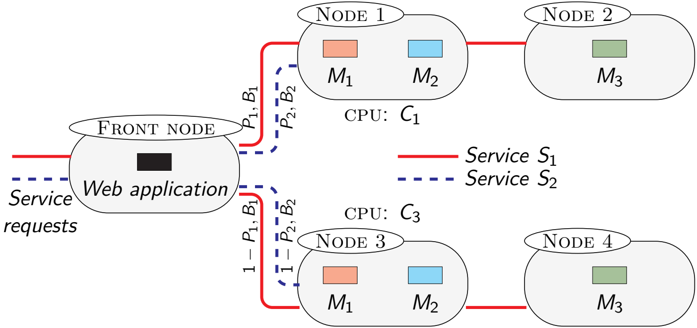

- The path of a request for service i is decided by the routing probability P_i.
- A request for service i is blocked at the front node with blocking probability B_i.
- The CPU allocated to nodes 1 and 3 are decided by the CPU allocations C_1 and C_3.

Causal Graph

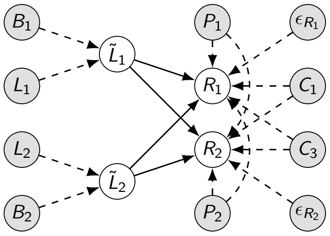

- **Exogenous variables**:
  - B_i: blocking probability of service S_i;
  - P_i: routing probability of service S_i to node 1;
  - C_j: CPU allocation to node j;
  - L_i: load of service S_i (requests per second); and
  - $\epsilon_{R_1}, \epsilon_{R_2}$: random noise variables.
- **Endogenous variables**:
  - $\tilde{L}_i$: carried (i.e., non-blocked) load of service S_i; and
  - R_i: response time (s) of service S_i.

Evaluation Scenarios

- **Evaluation:** Identify the causal model from data.
- **Scenario 1**:
  - Load service S_1 with $L_1 = 4$ requests per second.
  - Load service S_2 with $L_2 = 15$ requests per second.
  - Causal graph is as shown in the previous slide.
  - Causal functions are time-independent.
- **Scenario 2**:
  - Time interval $[1, 11)$: Run only S_2 with load $L_2 = 1$.
  - At time $t = 11$: Start S_1 in the background with $L_1 = 20$.
  - Causal functions are time-dependent.
  - Causal graph:

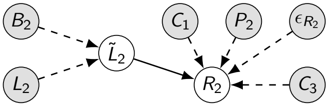

Ground Truth Model (1/2)

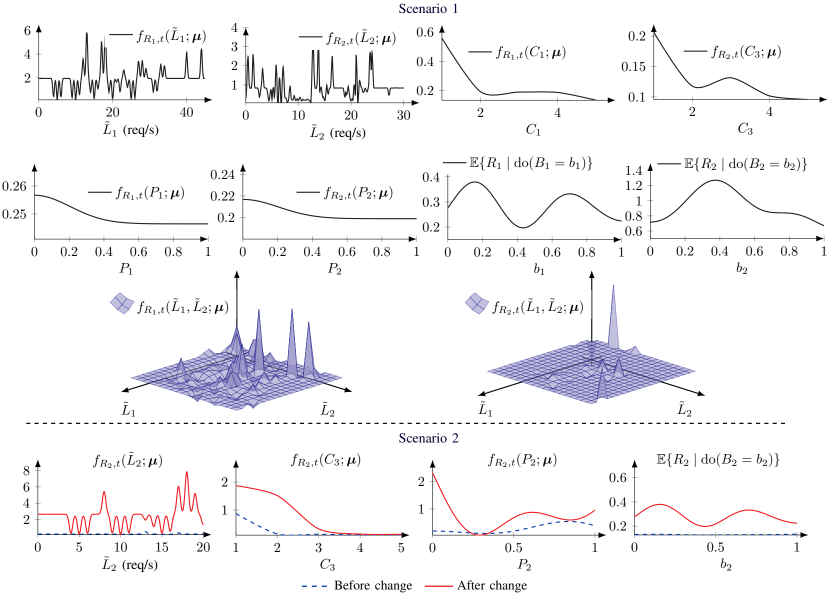

Ground Truth Model (2/2)

	C_1	P_1	$\tilde{L}_1$	R_1	C_3	P_2	$\tilde{L}_2$	R_2	B_1	B_2	L_1	L_2
C_1	1.00	0.00	0.08	-0.18	0.00	0.00	0.02	-0.11	0.08	0.02	0.03	0.00
P_1	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
$\tilde{L}_1$	0.08	0.00	1.00	0.00	0.08	0.00	0.17	-0.10	0.77	0.03	0.57	0.15
R_1	-0.18	0.00	0.00	1.00	-0.18	0.00	0.03	0.34	-0.01	0.03	0.00	0.00
C_3	0.00	0.00	0.08	-0.18	1.00	0.00	0.02	-0.10	0.08	0.02	0.02	0.00
P_2	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00
$\tilde{L}_2$	0.02	0.00	0.17	0.03	0.02	0.00	1.00	0.02	0.03	0.77	0.07	0.58
R_2	-0.11	0.00	-0.10	0.34	-0.10	0.00	0.02	1.00	-0.10	0.02	-0.04	0.00
B_1	0.08	0.00	0.77	-0.01	0.08	0.00	0.03	-0.10	1.00	0.04	0.00	0.00
B_2	0.02	0.00	0.03	0.03	0.02	0.00	0.77	0.02	0.04	1.00	0.00	0.02
L_1	0.03	0.00	0.57	0.00	0.02	0.00	0.07	-0.04	0.00	0.00	1.00	0.14
L_2	0.00	0.00	0.15	0.00	0.00	0.00	0.58	0.00	0.00	0.02	0.14	1.00

1.00

Pearson correlation

-1.00

Operational Costs of Interventions

Intervention	**Cost**
Emulating the service load L_i	3000
Adjusting the routing probability P_i	1000
Adjusting the blocking probability B_i	2000
Modifying the CPU allocation C_j	3000
Monitoring without intervening, i.e., $\mathrm{do}(\emptyset)$	1

Table 1: Intervention costs for defining the cost function c.

Passive Learning Baseline for Scenario 1

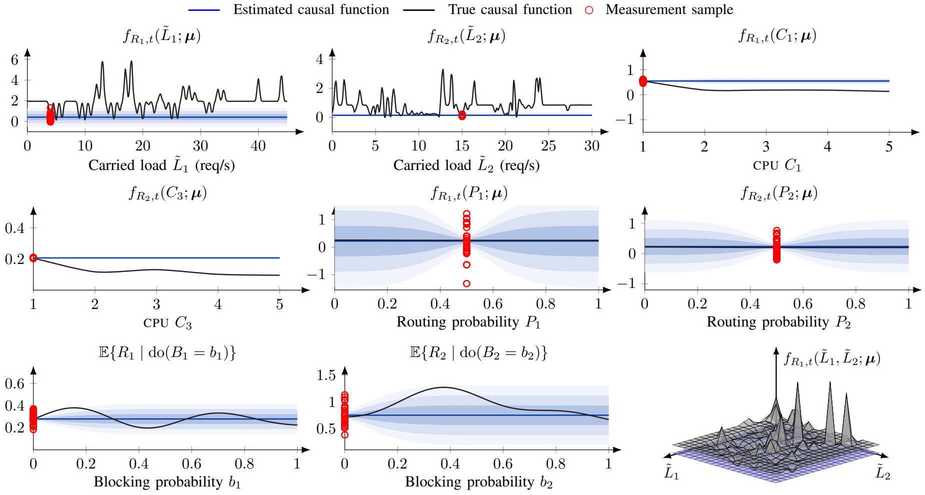

Causal Active Learning for Scenario 1

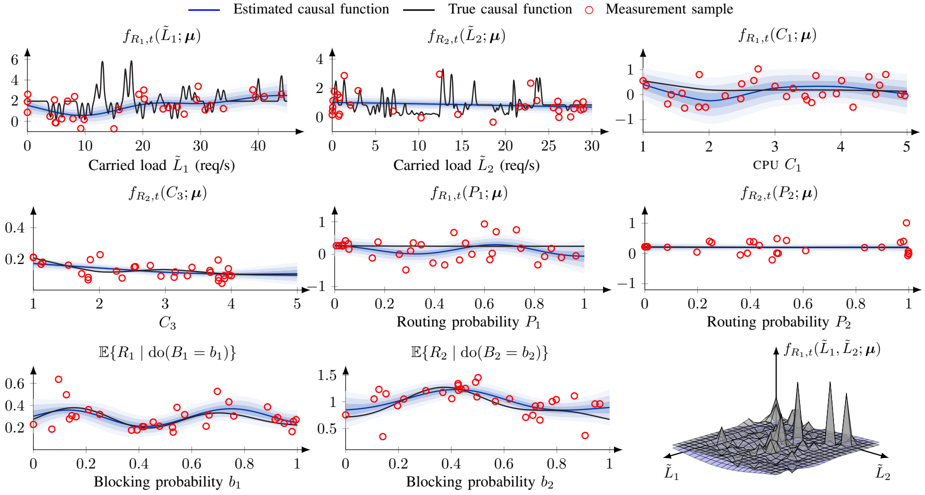

Causal Active Learning Dynamics for Scenario 1

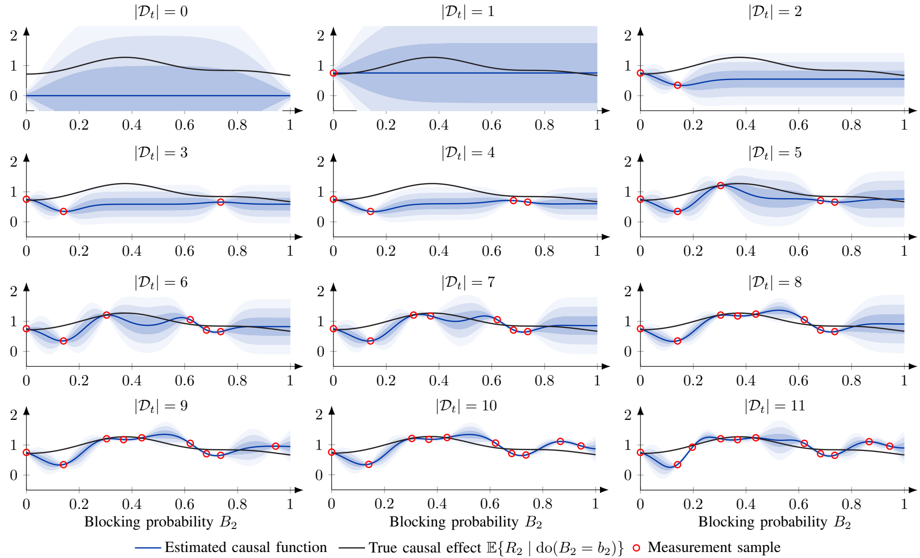

Estimation Loss for Scenario 1

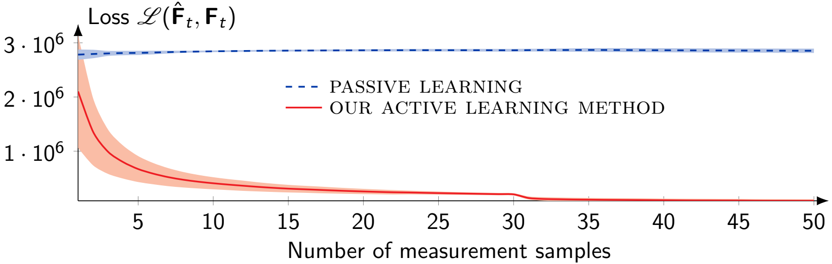

Causal Active Learning Dynamics for Scenario 2

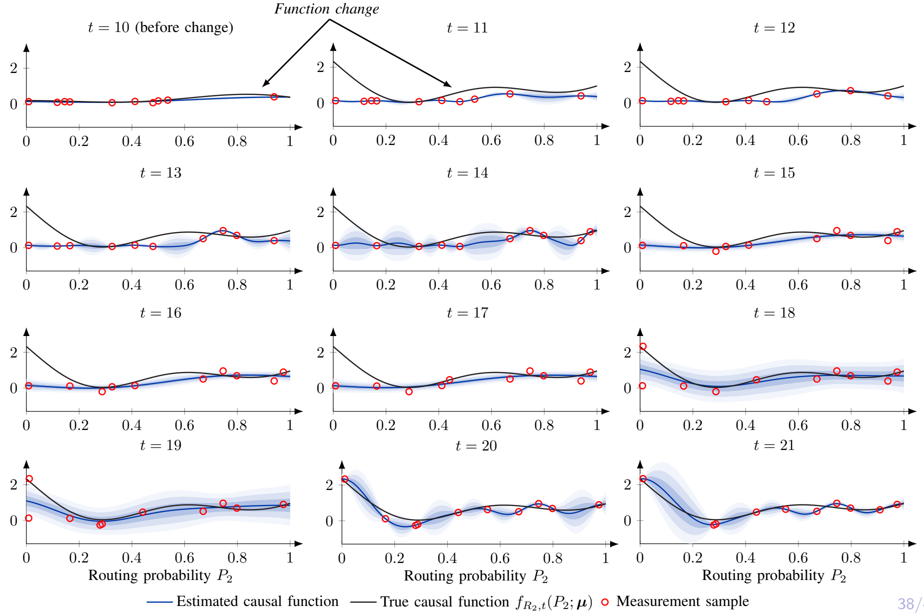

Estimation Loss for Scenario 2

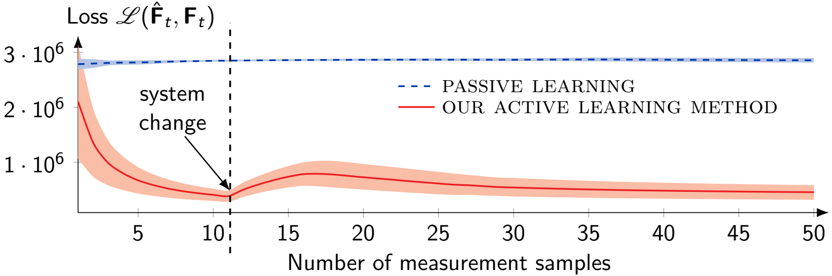

Takeaway

Unlike offline identification methods, our online method allows for quick adaptation of the model to system changes.

Conclusion

- **Causal models** are central to achieving the longstanding goal of automating network and system management tasks.

- We develop active causal learning.
  - The first online method for causal identification of IT systems.
  - We show that the method has appealing theoretical properties.
  - We validate the method experimentally on a service mesh.