Sequential Covariate Shift Detection Using Classifier Two-Sample Tests




Sooyong Jang[1] Sangdon Park[1,2] Insup Lee[1] Osbert Bastani[1]

1. PRECISE Center, University of Pennsylvania
2. Georgia Institute of Technology

ICML 2022

Penn Engineering
University of Pennsylvania

PRECISE
PENN RESEARCH IN EMBEDDED COMPUTING AND INTEGRATED SYSTEMS ENGINEERING

Georgia Tech

Motivation

Motivation

Assumption

Motivation

Assumption

$$\mathbb{P}[X_{train}, Y_{train}] = \mathbb{P}[X_{test}, Y_{test}]$$

Motivation

Assumption

$$\mathbb{P}[X_{train}, Y_{train}] \not= \mathbb{P}[X_{test}, Y_{test}]$$

Motivation

Assumption

$$\mathbb{P}[X_{train}, Y_{train}] \neq \mathbb{P}[X_{test}, Y_{test}]$$

For example,

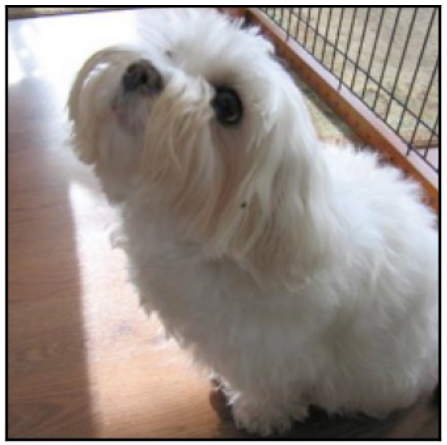

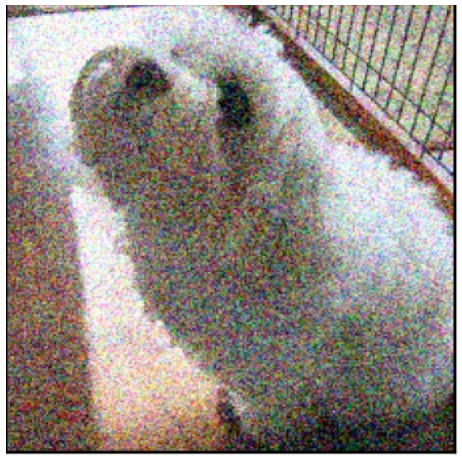

Covariate shift

Covariate shift

Brings bad consequences …

Covariate shift

Brings bad consequences …

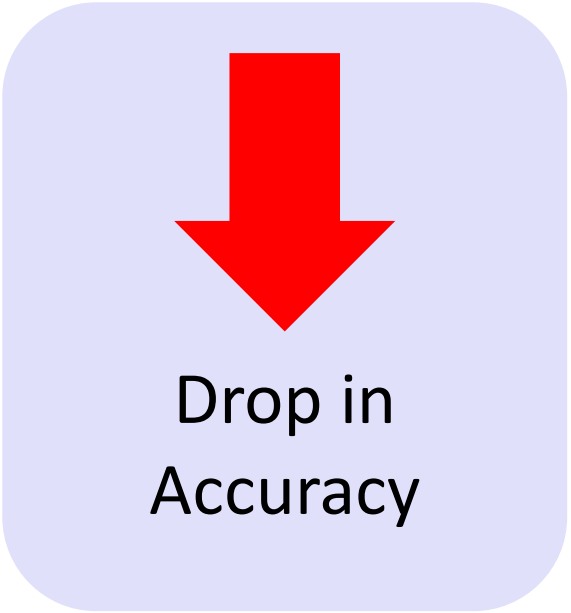

Covariate shift

Brings bad consequences …

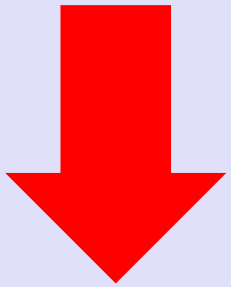

Drop in Accuracy

Invalidate Uncertainty Quantification

Covariate shift

Brings bad consequences …

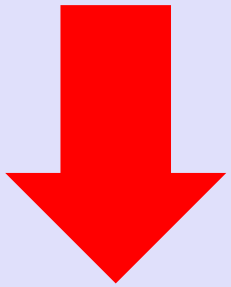

Drop in Accuracy

Invalidate Uncertainty Quantification

Need to detect covariate shift!

Sequential Setup

Sequential Setup

- Practically, data comes sequentially

Sequential Setup

- Practically, data comes sequentially

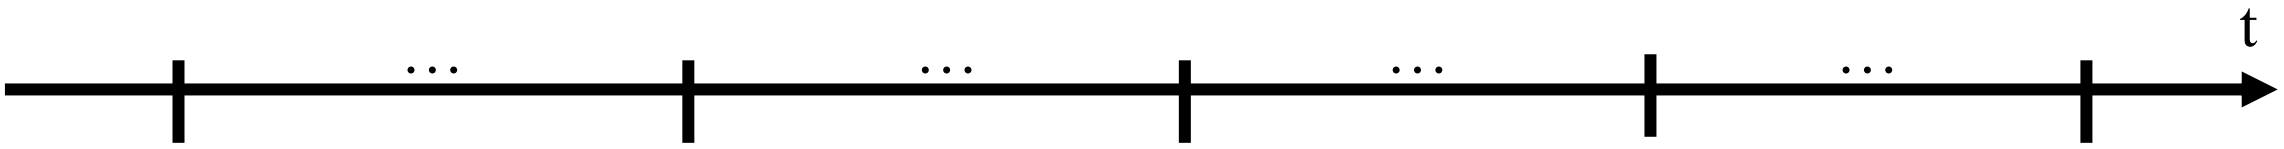

Sequential Setup

- Practically, data comes sequentially

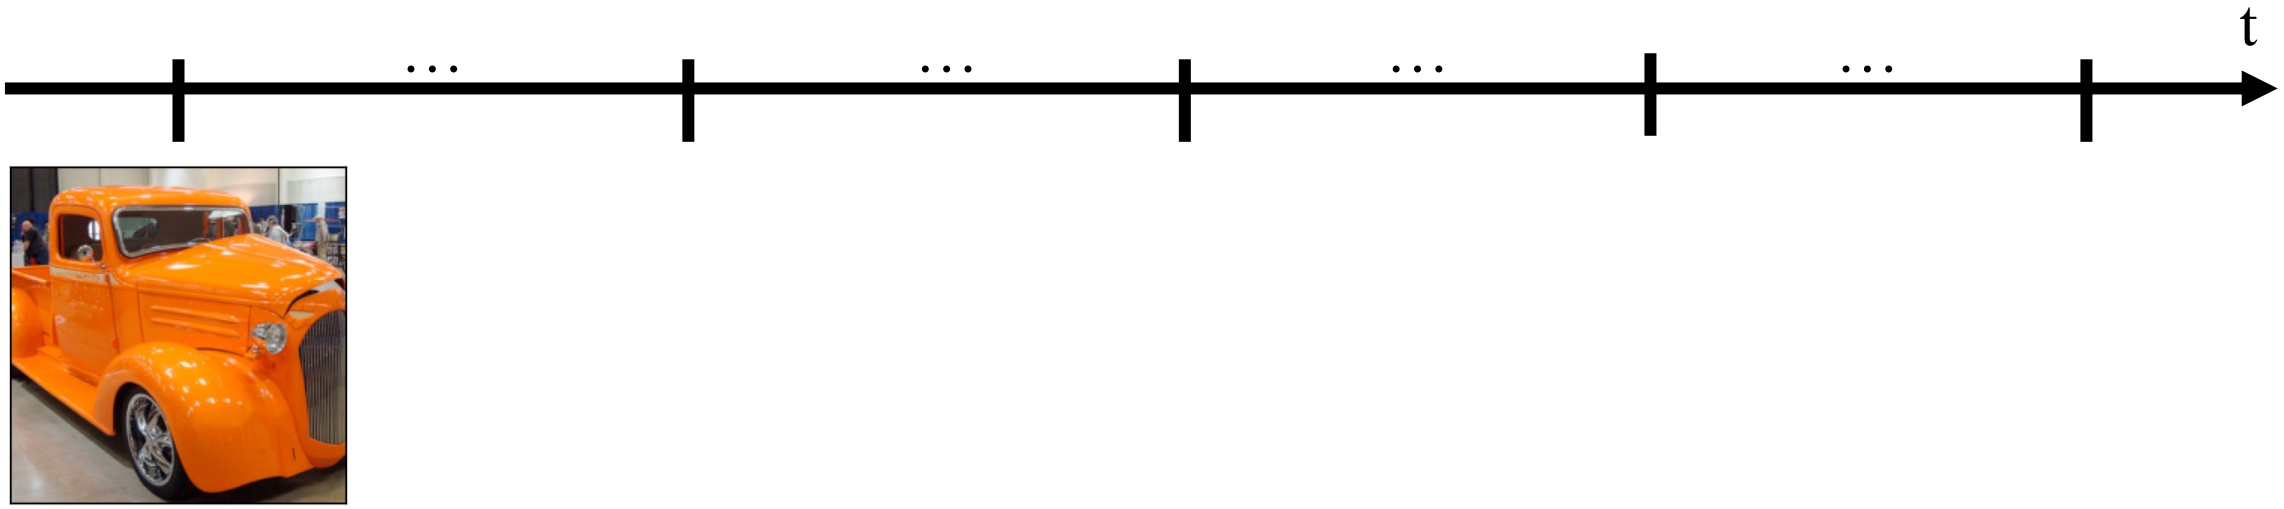

Sequential Setup

- Practically, data comes sequentially

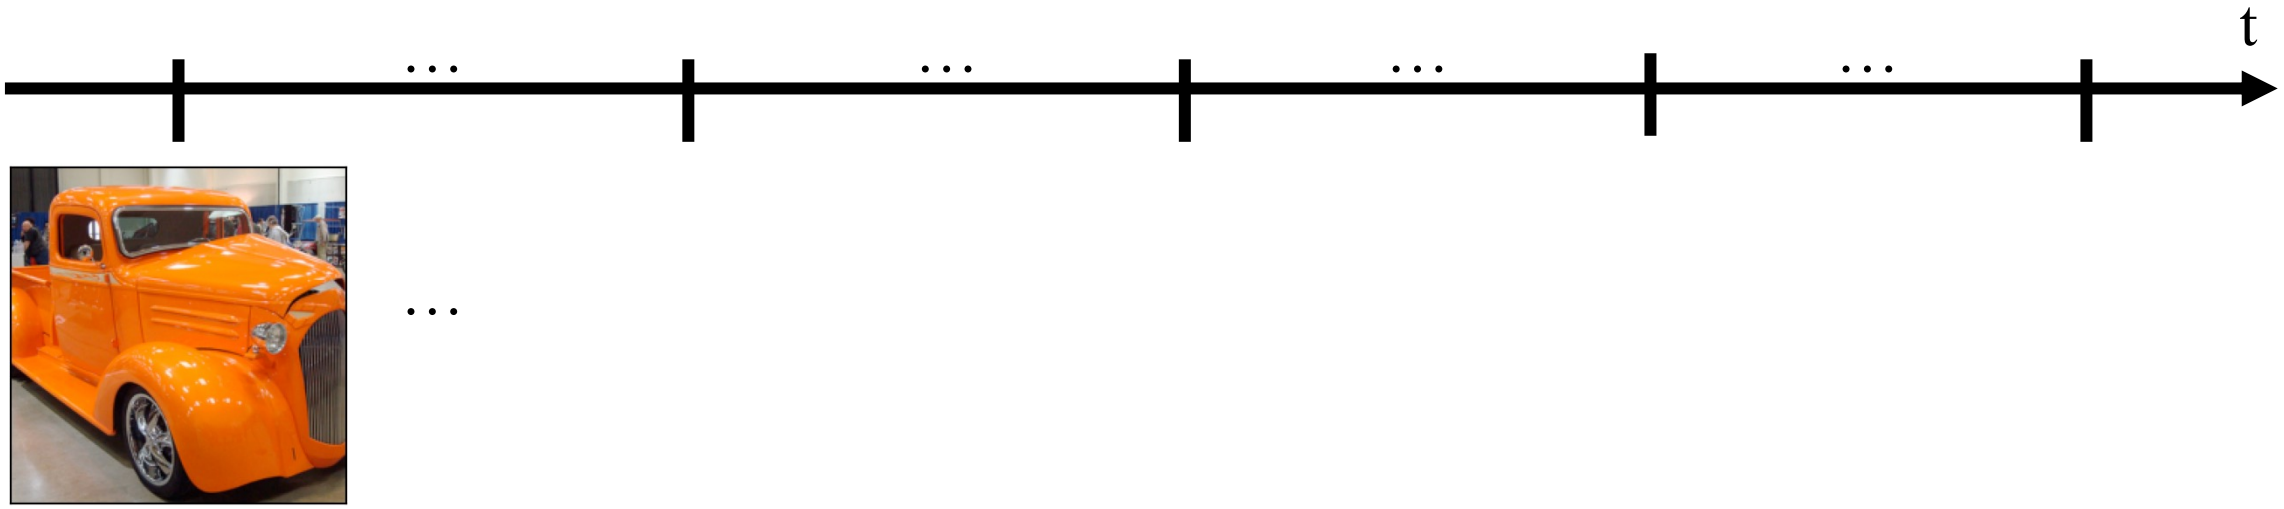

Sequential Setup

- Practically, data comes sequentially

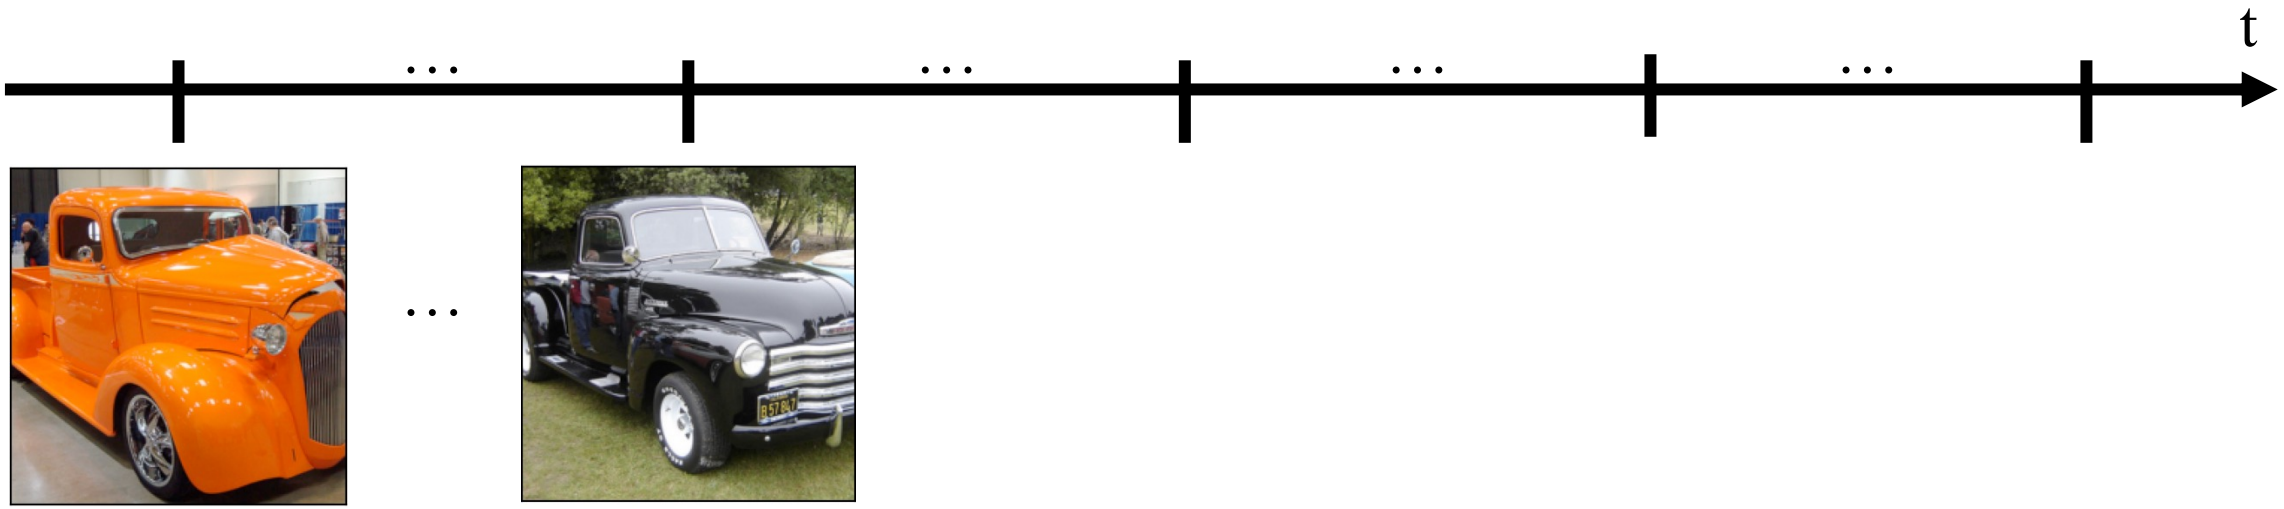

Sequential Setup

- Practically, data comes sequentially

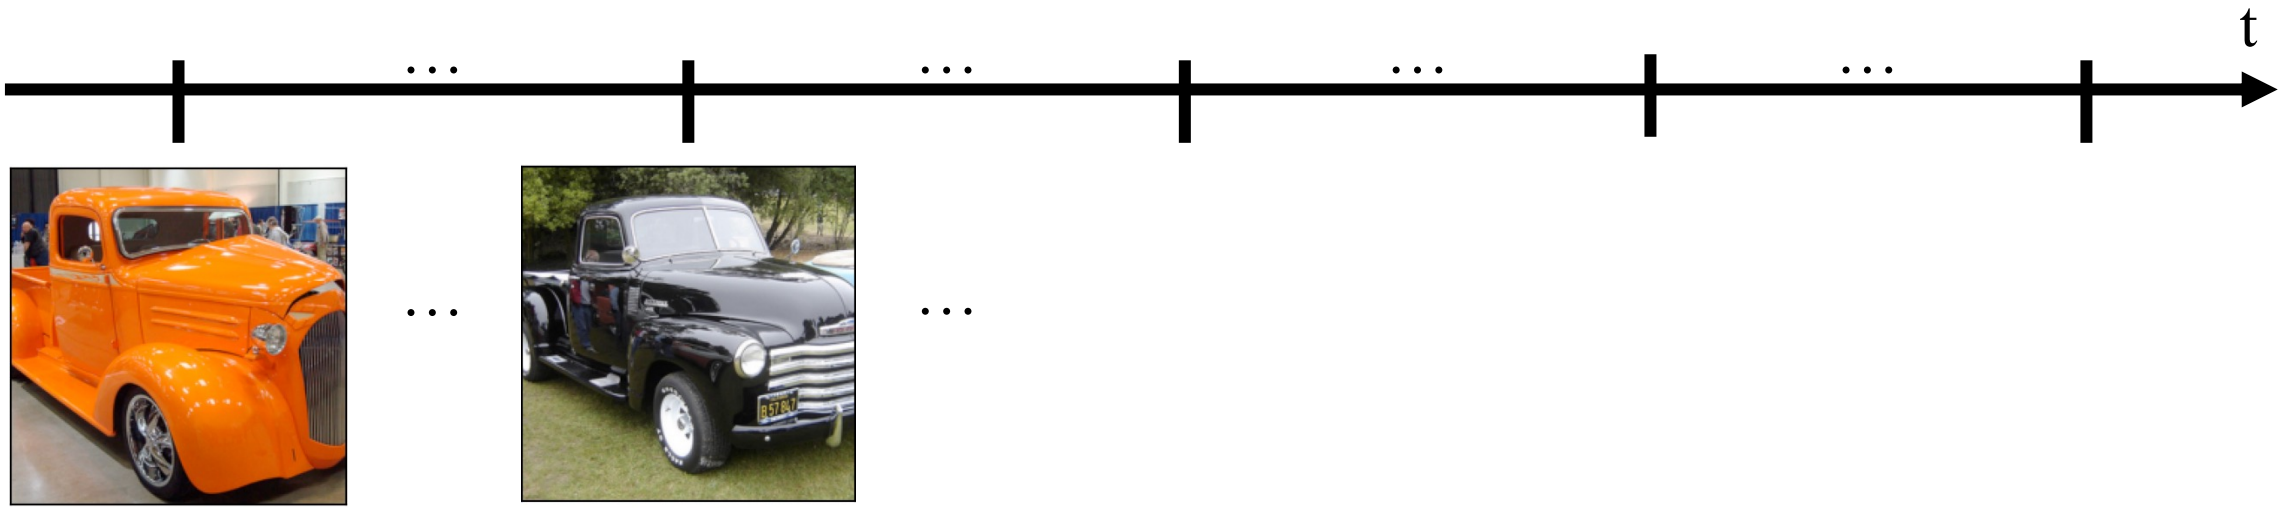

Sequential Setup

- Practically, data comes sequentially

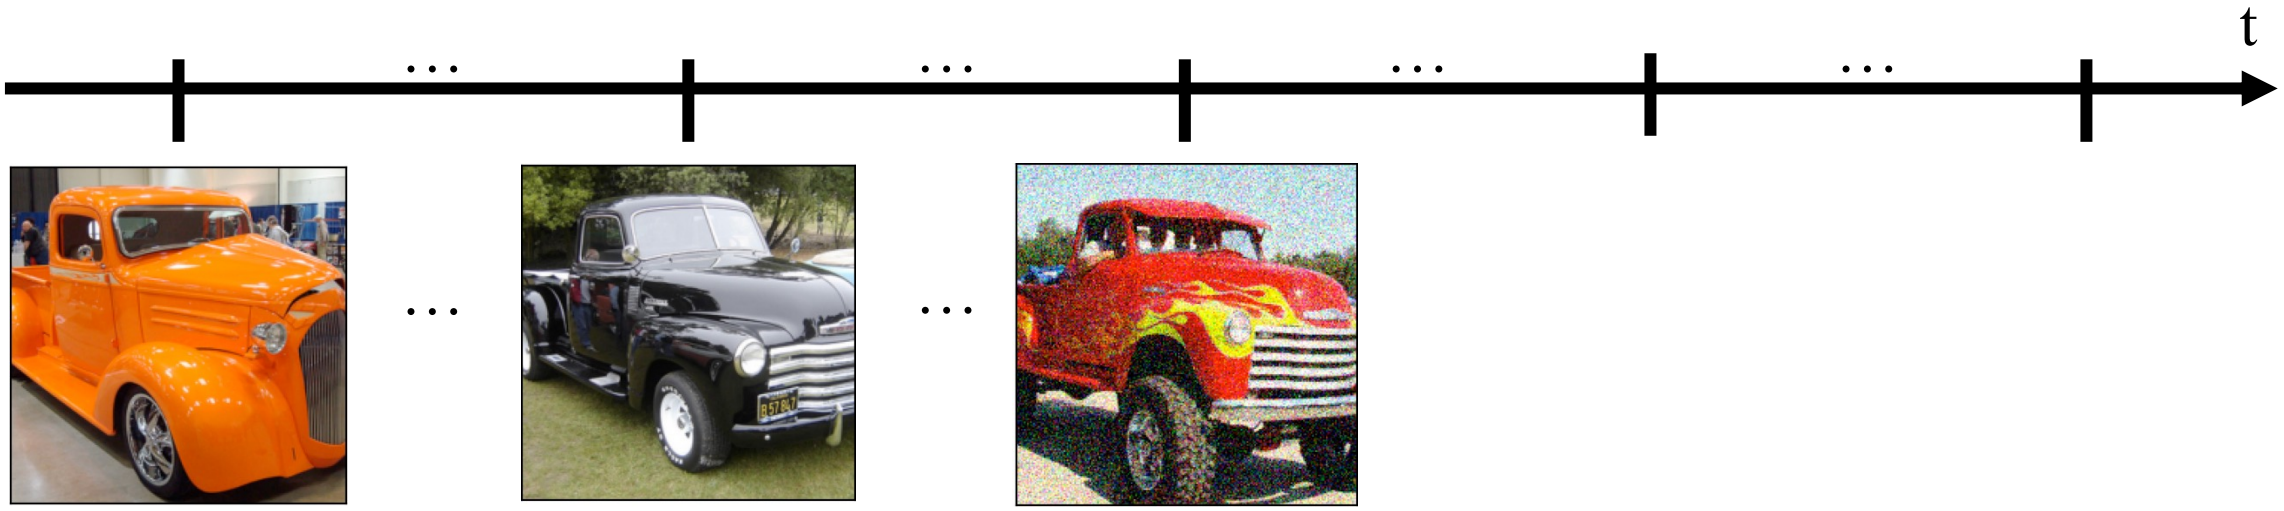

Sequential Setup

- Practically, data comes sequentially

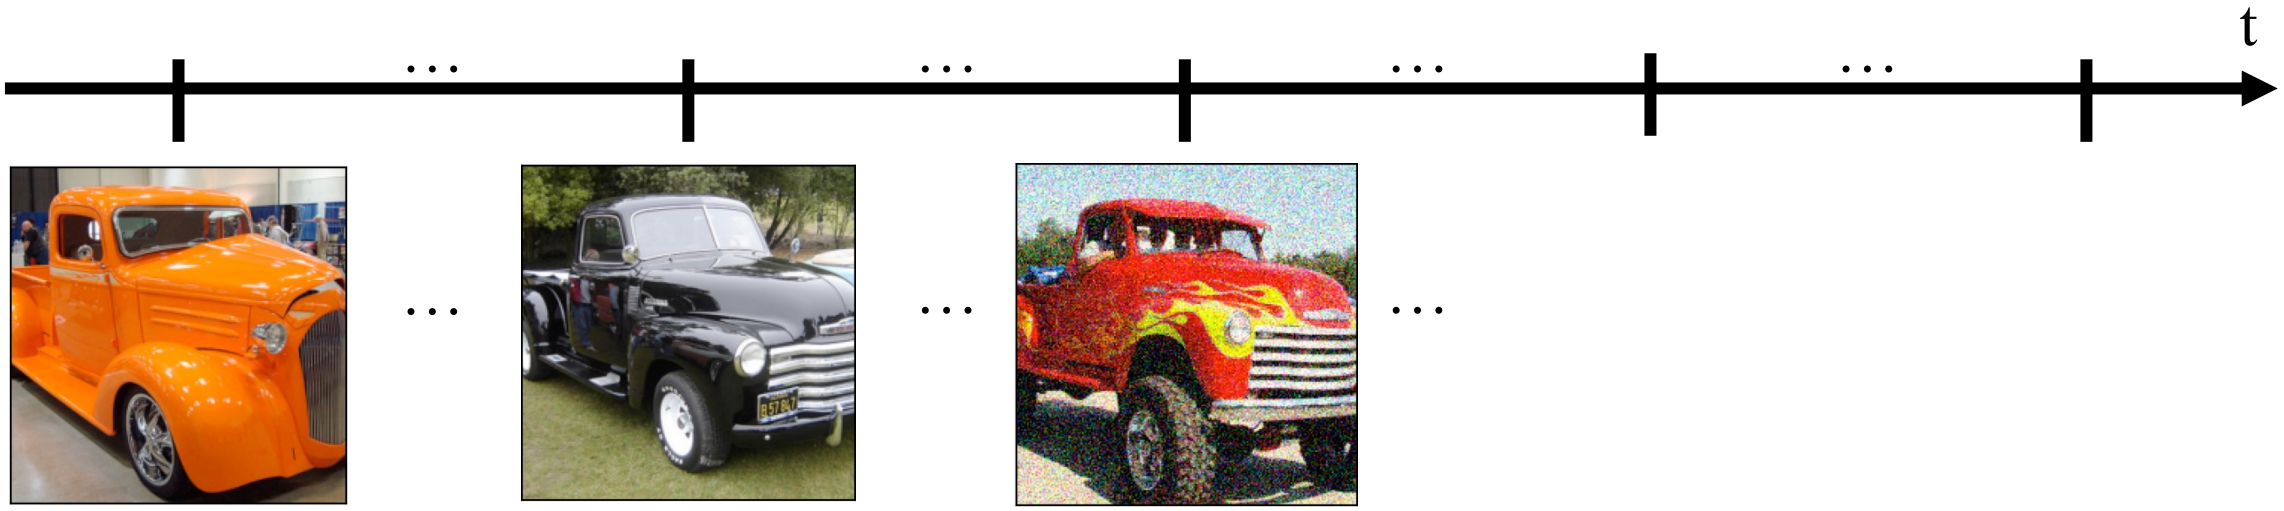

Sequential Setup

- Practically, data comes sequentially

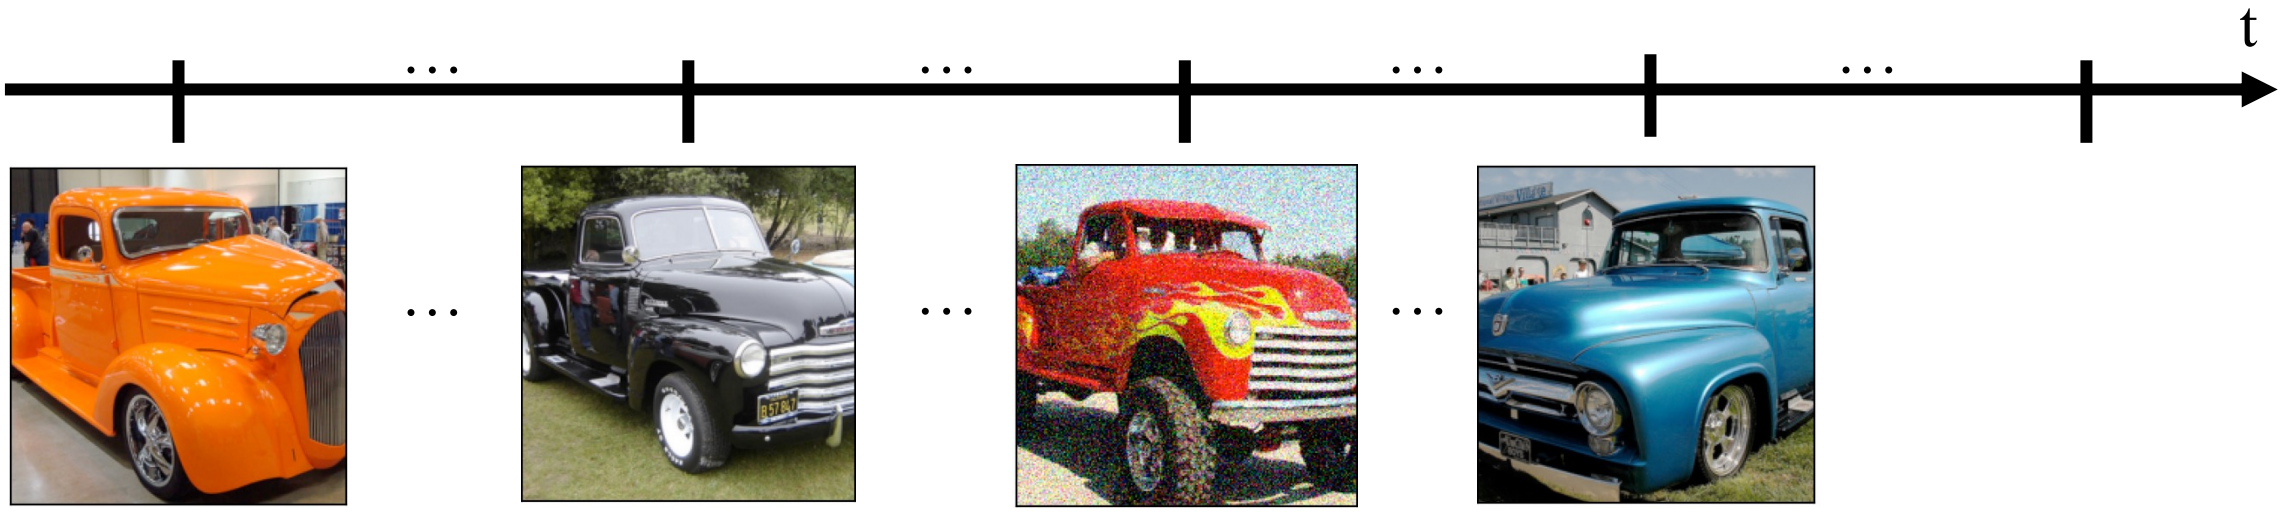

Sequential Setup

- Practically, data comes sequentially

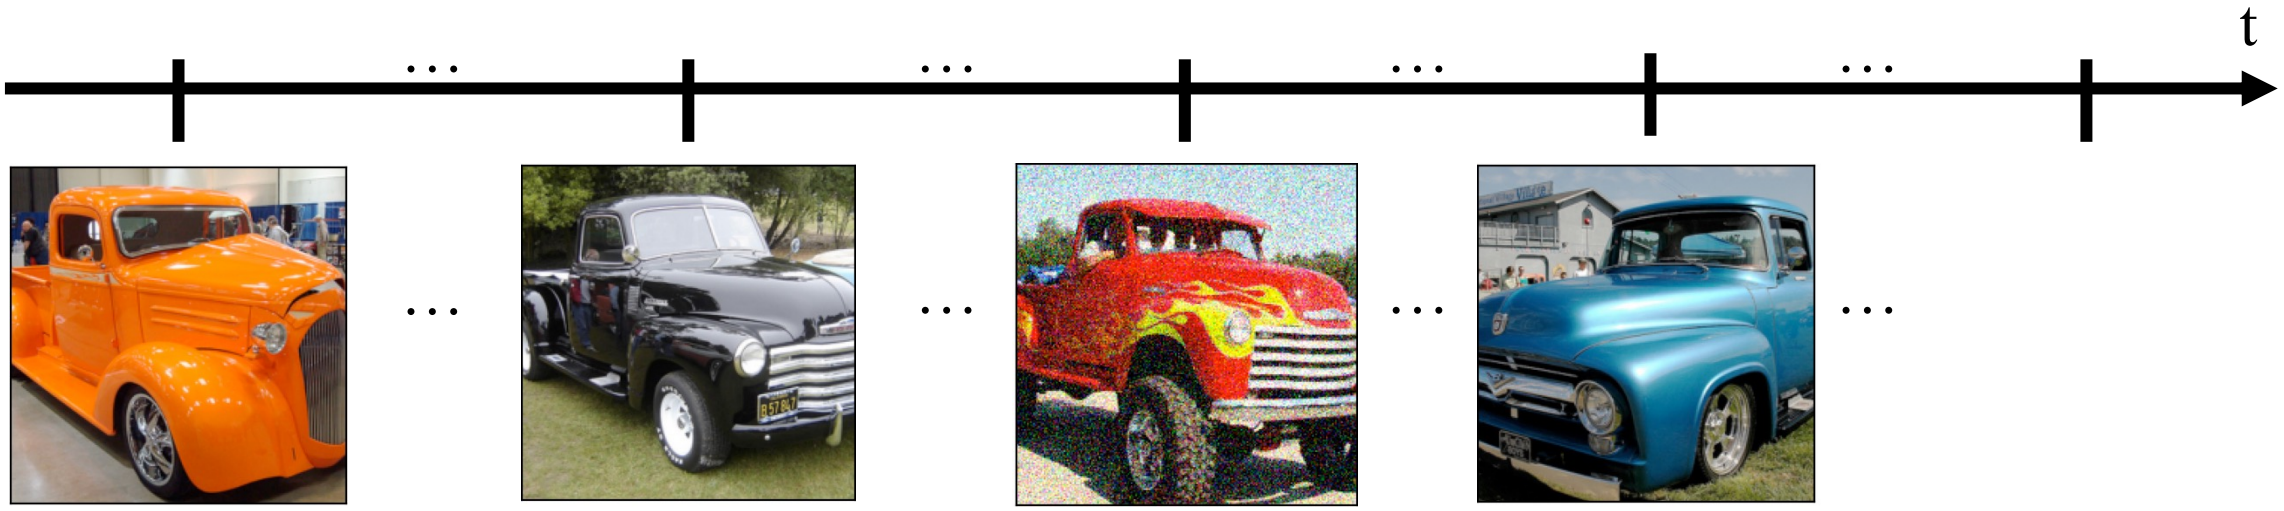

Sequential Setup

- Practically, data comes sequentially

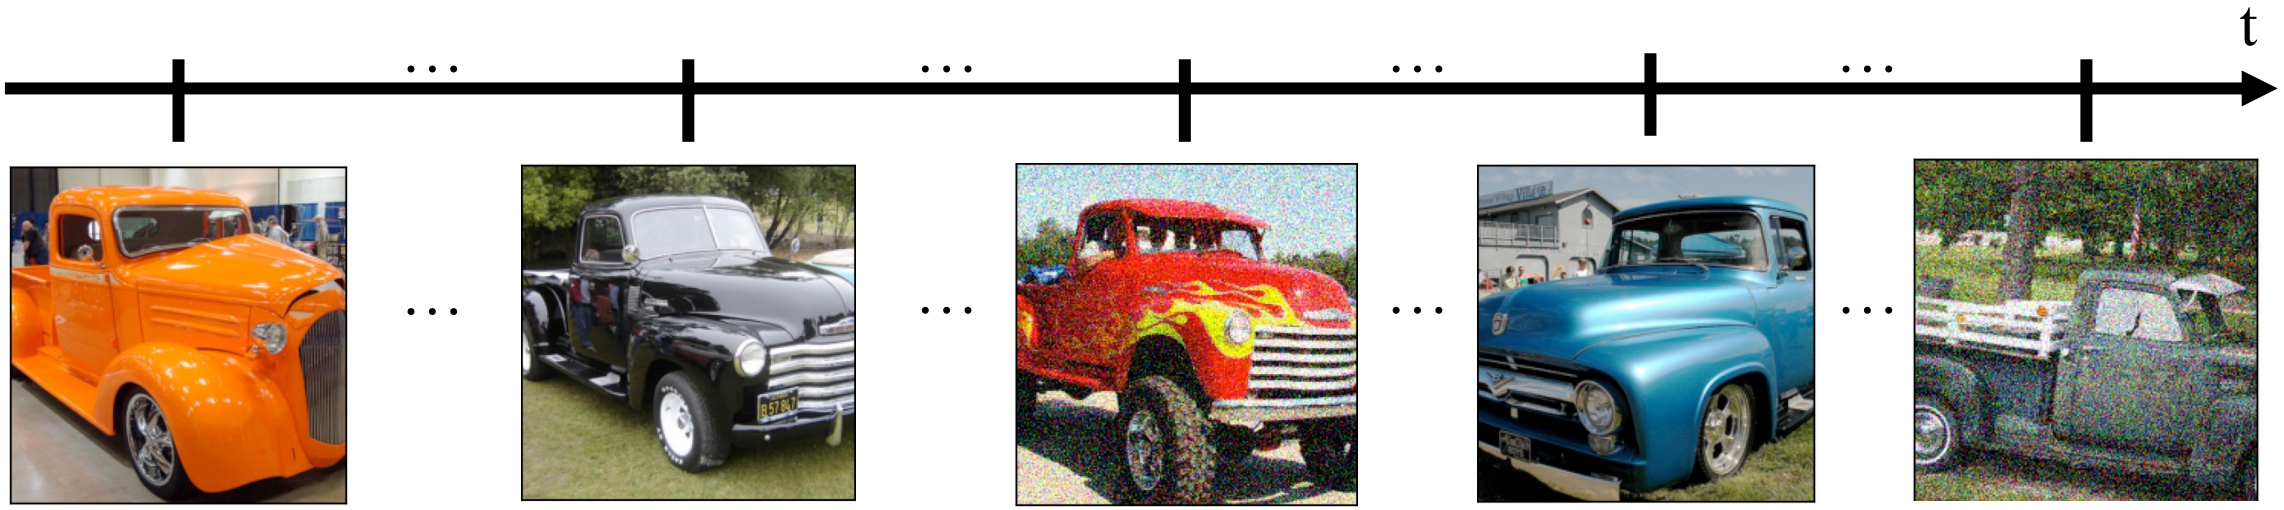

Sequential Setup

- Practically, data comes sequentially

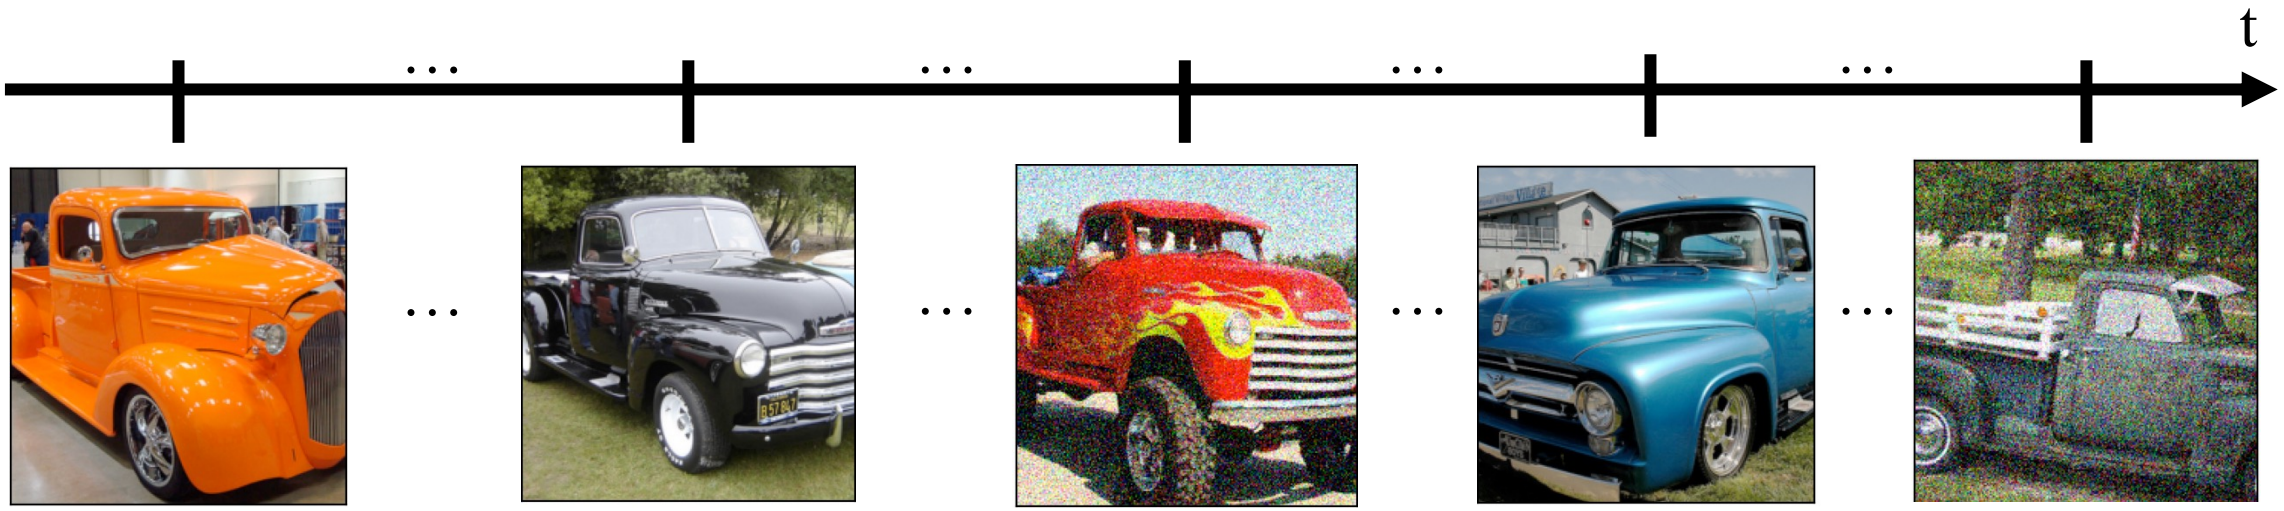

➡ Need to detect covariate shift at each time point

Problem: Sequential Covariate Shift

- Design a detector $\hat{f}$ for detecting covariate shift at each time step t

$$\hat{f}(S_{w,t}, T_{w,t}) \approx \begin{cases} 1 & \text{if } \mathcal{S} \neq \bar{\mathcal{T}}_{w,t} \\ 0 & \text{otherwise.} \end{cases}$$

Problem: Sequential Covariate Shift

- Design a detector $\hat{f}$ for detecting covariate shift at each time step t

$$\hat{f}(S_{w,t}, T_{w,t}) \approx \begin{cases} 1 & \text{if } \mathcal{S} \neq \bar{\mathcal{T}}_{w,t} \\ 0 & \text{otherwise.} \end{cases}$$

Source samples

Problem: Sequential Covariate Shift

- Design a detector $\hat{f}$ for detecting covariate shift at each time step t

$$\hat{f}(S_{w,t}, T_{w,t}) \approx \begin{cases} 1 & \text{if } \mathcal{S} \neq \bar{\mathcal{T}}_{w,t} \\ 0 & \text{otherwise.} \end{cases}$$

Source samples

Problem: Sequential Covariate Shift

- Design a detector $\hat{f}$ for detecting covariate shift at each time step t

$$\hat{f}(S_{w,t}, T_{w,t}) \approx \begin{cases} 1 & \text{if } \mathcal{S} \neq \bar{\mathcal{T}}_{w,t} \\ 0 & \text{otherwise.} \end{cases}$$

Source samples

target samples

Problem: Sequential Covariate Shift

- Design a detector $\hat{f}$ for detecting covariate shift at each time step t

$$\hat{f}(S_{w,t}, T_{w,t}) \approx \begin{cases} 1 & \text{if } \mathcal{S} \neq \bar{\mathcal{T}}_{w,t} \\ 0 & \text{otherwise.} \end{cases}$$

Source samples

target samples

Problem: Sequential Covariate Shift

- Design a detector $\hat{f}$ for detecting covariate shift at each time step t

$$\hat{f}(S_{w,t}, T_{w,t}) \approx \begin{cases} 1 & \text{if } \mathcal{S} \neq \bar{\mathcal{T}}_{w,t} \\ 0 & \text{otherwise.} \end{cases}$$

Source samples

target samples

Satisfying the following conditions

Problem: Sequential Covariate Shift

- Design a detector $\hat{f}$ for detecting covariate shift at each time step t

$$\hat{f}(S_{w,t}, T_{w,t}) \approx \begin{cases} 1 & \text{if } \mathcal{S} \neq \bar{\mathcal{T}}_{w,t} \\ 0 & \text{otherwise.} \end{cases}$$

Source samples

target samples

Satisfying the following conditions

- FPR Bound

Problem: Sequential Covariate Shift

- Design a detector $\hat{f}$ for detecting covariate shift at each time step t

$$\hat{f}(S_{w,t}, T_{w,t}) \approx \begin{cases} 1 & \text{if } \mathcal{S} \neq \bar{\mathcal{T}}_{w,t} \\ 0 & \text{otherwise.} \end{cases}$$

Source samples

target samples

Satisfying the following conditions

- FPR Bound
- FNR Bound

Idea: Source – Target classifier $\hat{g}$

Idea: Source – Target classifier $\hat{g}$

Source

Target

Idea: Source – Target classifier $\hat{g}$

Source

Target

Shift

Idea: Source – Target classifier $\hat{g}$

Source

Target

Shift

vs

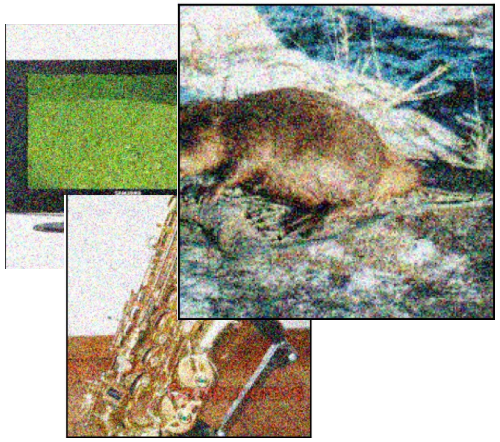

Idea: Source – Target classifier $\hat{g}$

Shift

Source

vs

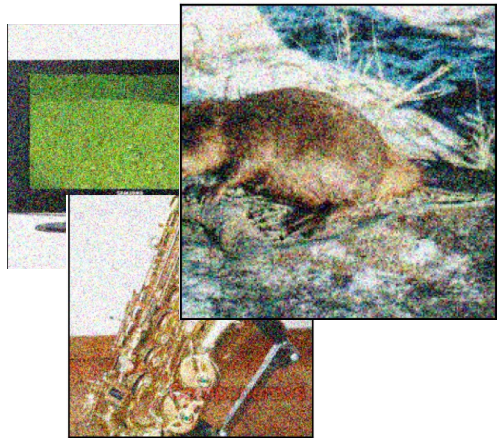

Target

Idea: Source – Target classifier $\hat{g}$

Shift

Source

vs

Target

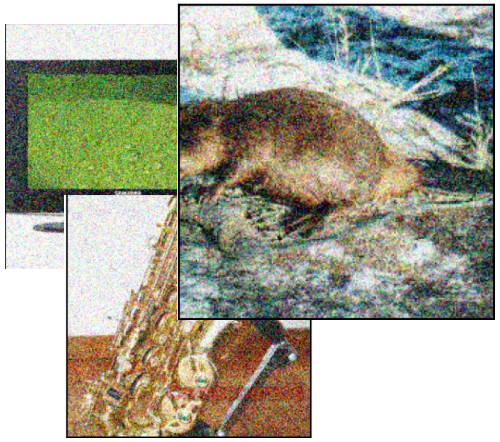

$\hat{g}$ **CAN** Differentiate

Idea: Source – Target classifier $\hat{g}$

Source

Shift

vs

Target

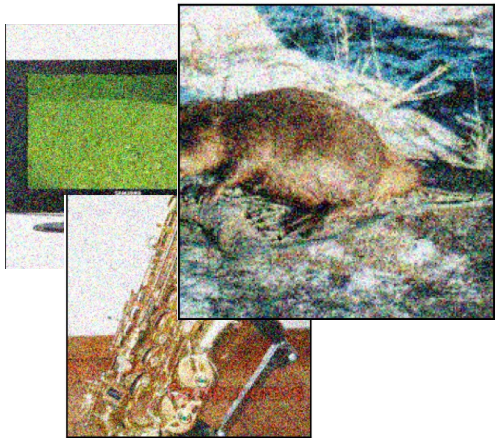

$\hat{g}$ **CAN** Differentiate

Accuracy > 0.5

Idea: Source – Target classifier $\hat{g}$

Source

Target

Shift

vs

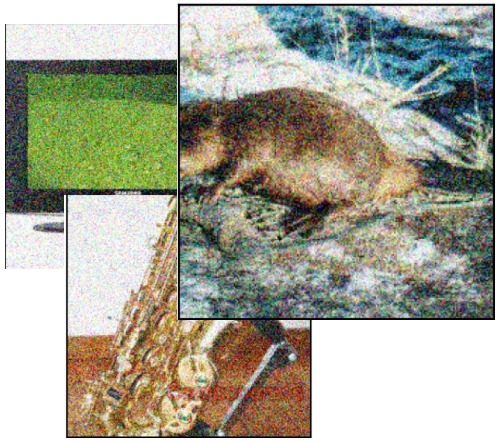

$\hat{g}$ **CAN** Differentiate

Accuracy > 0.5

Idea: Source – Target classifier $\hat{g}$

Source

Target

Shift

vs

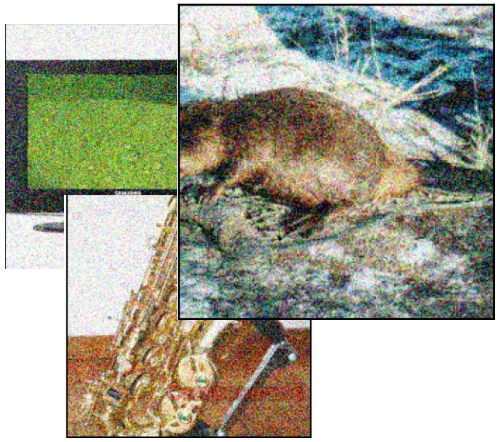

$\hat{g}$ **CAN** Differentiate

Accuracy > 0.5

No Shift

Idea: Source – Target classifier $\hat{g}$

Source

Target

Shift

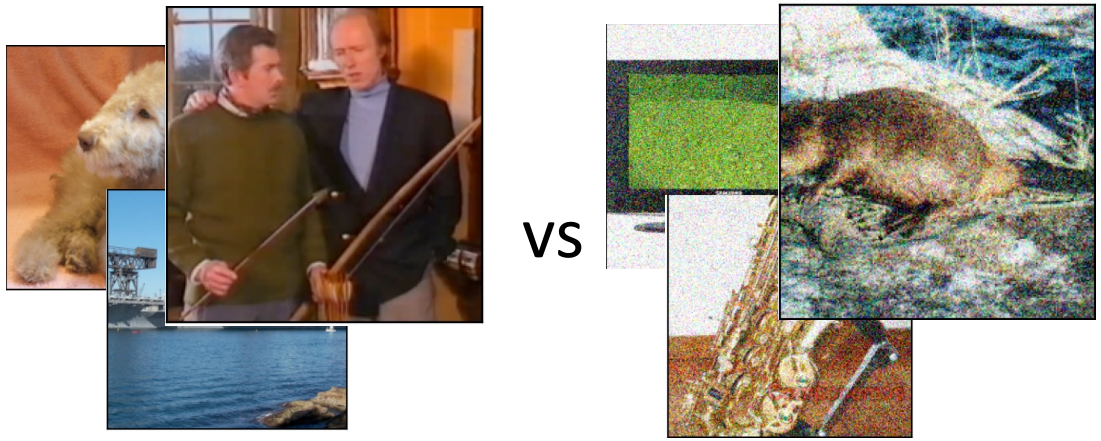

vs

$\hat{g}$ **CAN** Differentiate

Accuracy > 0.5

No Shift

vs

Idea: Source – Target classifier $\hat{g}$

Source

Target

Shift

vs

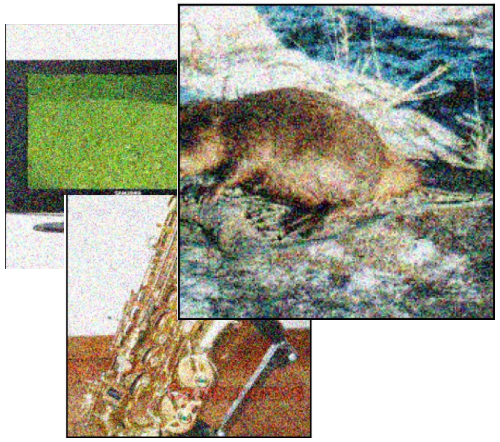

$\hat{g}$ **CAN** Differentiate

Accuracy > 0.5

No Shift

vs

Idea: Source – Target classifier $\hat{g}$

Source

Target

Shift

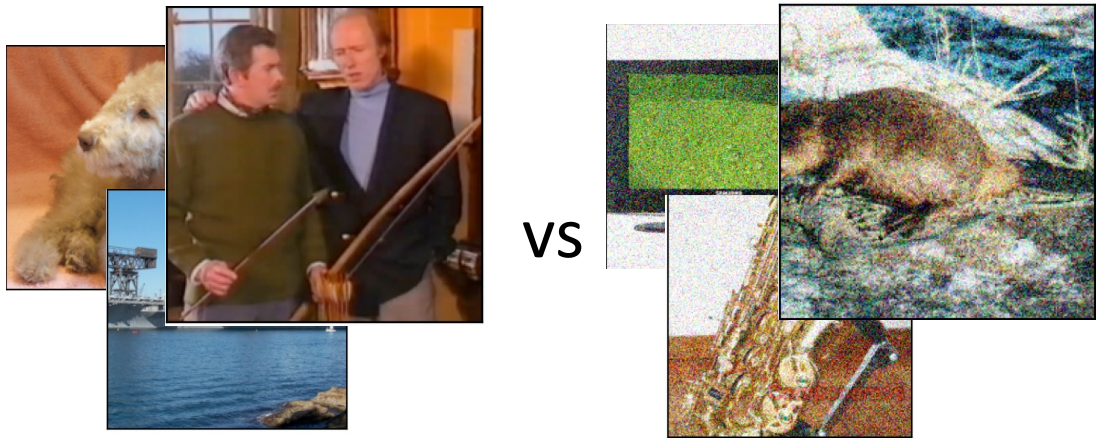

vs

$\hat{g}$ **CAN** Differentiate

Accuracy > 0.5

No Shift

vs

$\hat{g}$ **CANNOT** Differentiate

Idea: Source – Target classifier $\hat{g}$

Source

Target

Shift

vs

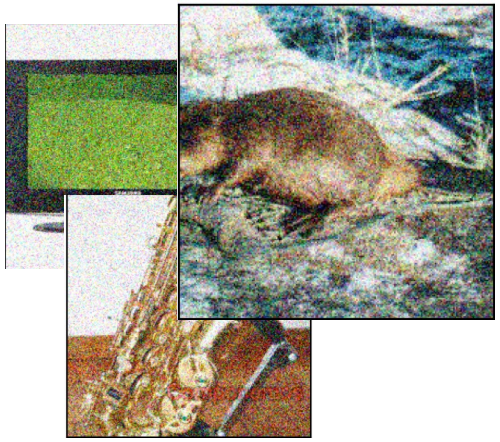

$\hat{g}$ **CAN** Differentiate

Accuracy > 0.5

No Shift

vs

$\hat{g}$ **CANNOT** Differentiate

Accuracy ≈ 0.5

Idea: Clopper-Pearson Interval

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

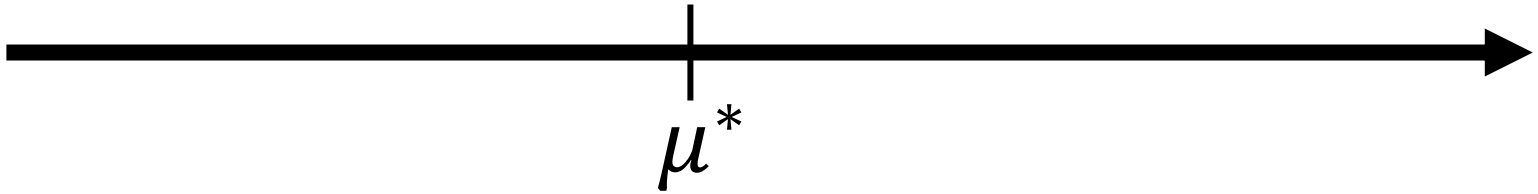

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

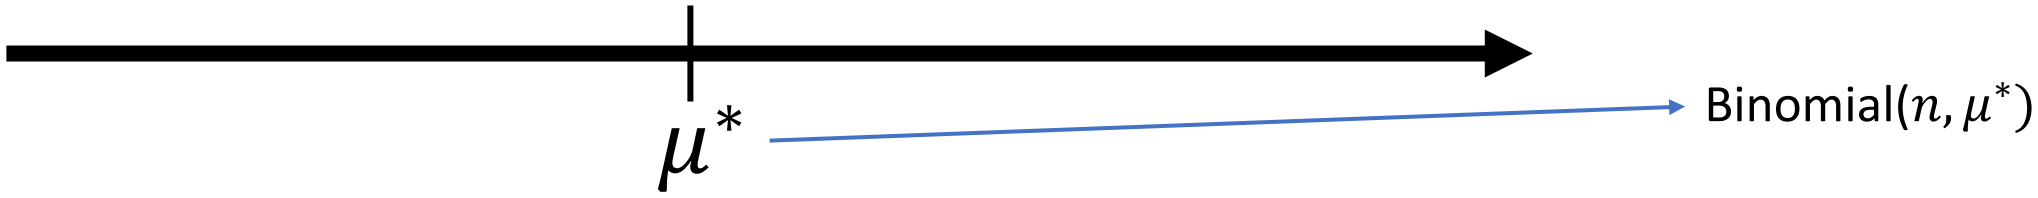

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

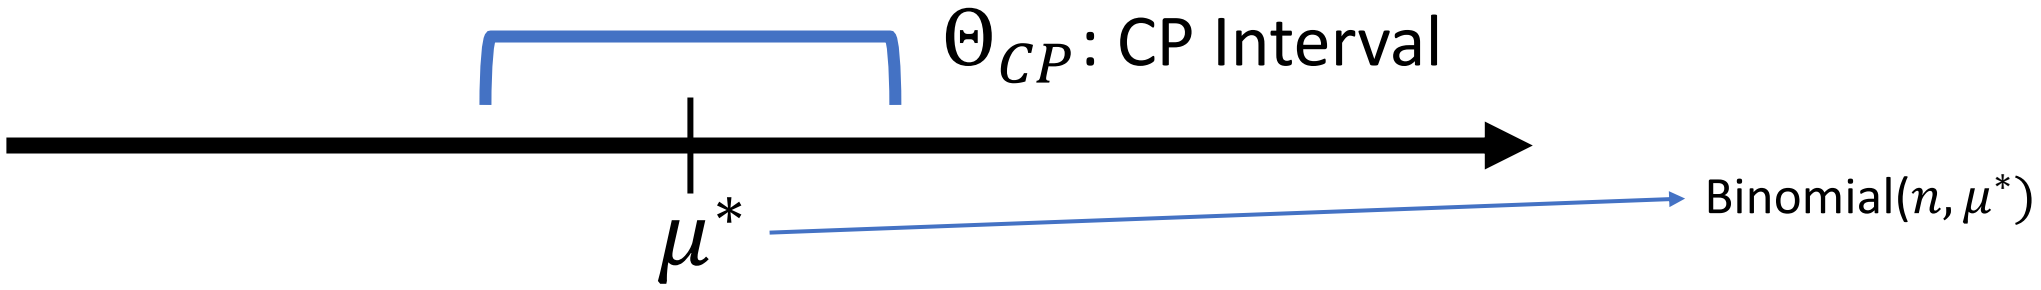

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

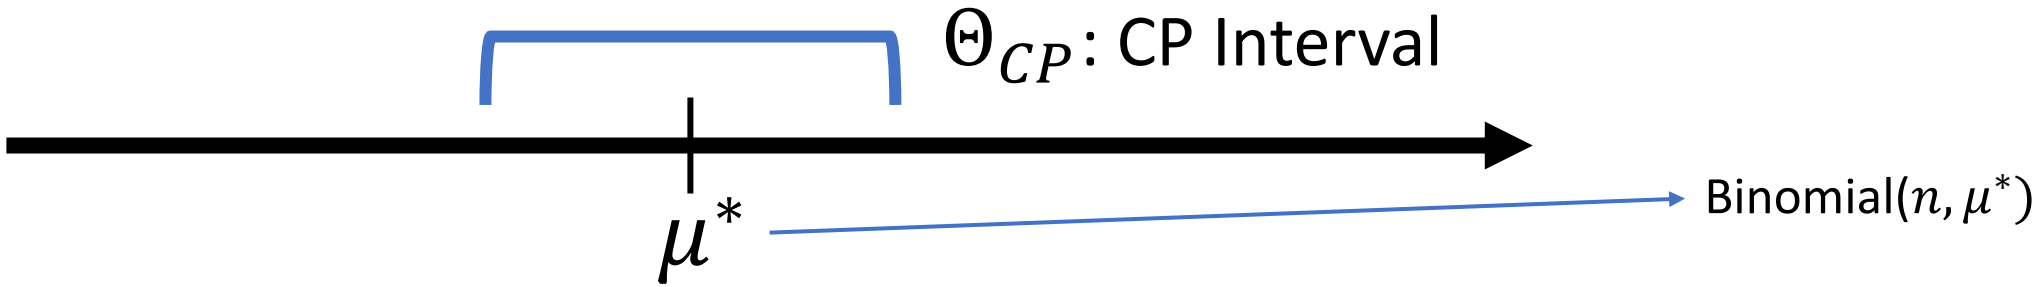

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

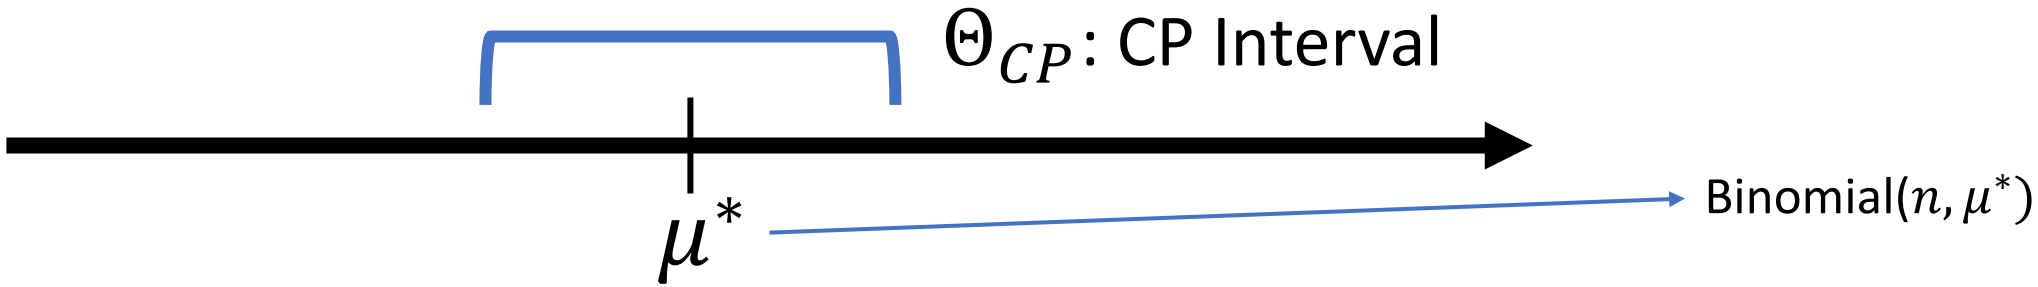

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

- Covariate Shift Detection

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

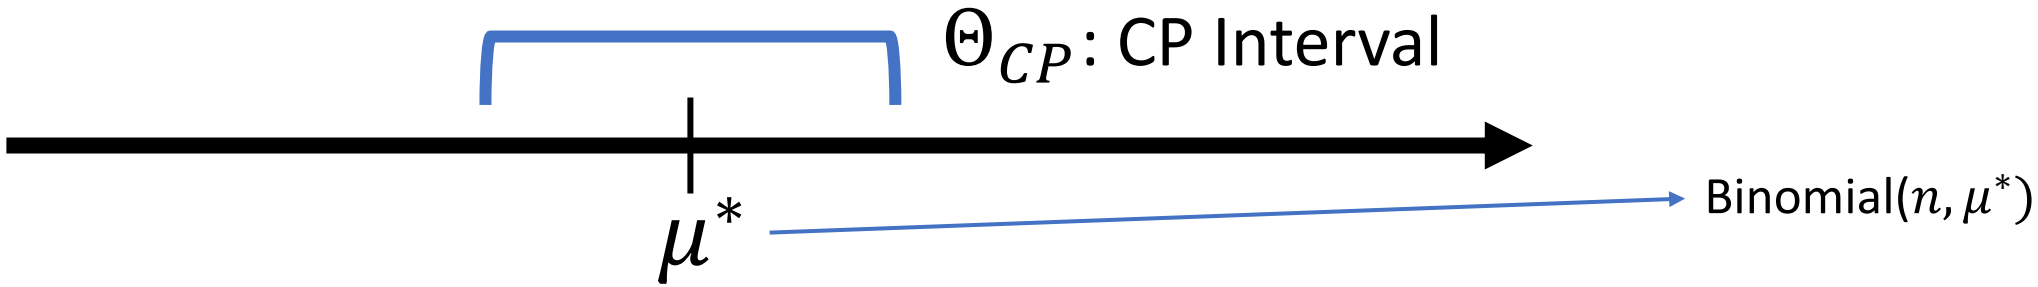

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

- Covariate Shift Detection

$\widehat{\boldsymbol{g}}(\boldsymbol{X}) = \boldsymbol{Y}$?

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

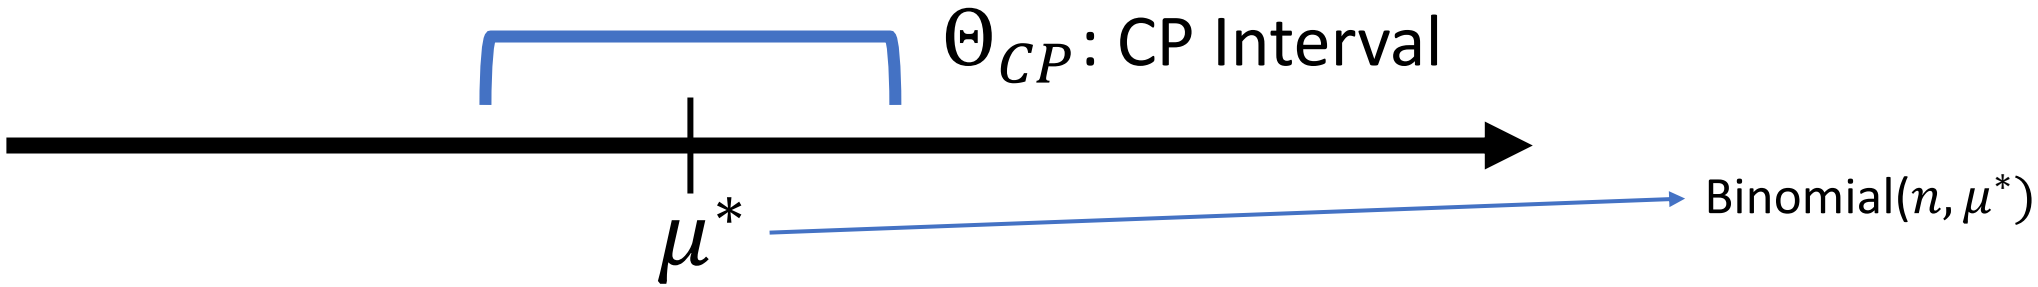

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

- Covariate Shift Detection

$\widehat{\boldsymbol{g}}(\boldsymbol{X}) = \boldsymbol{Y}$?

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

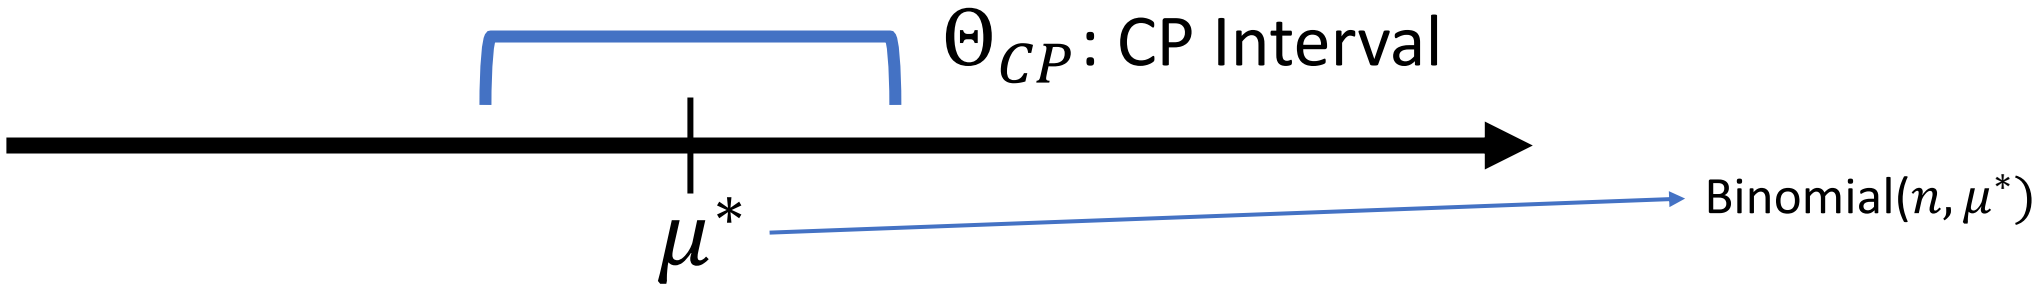

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

- Covariate Shift Detection

$\widehat{\boldsymbol{g}}(\boldsymbol{X}) = \boldsymbol{Y}$? ➡ Binomial$(n, \mu^*)$

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

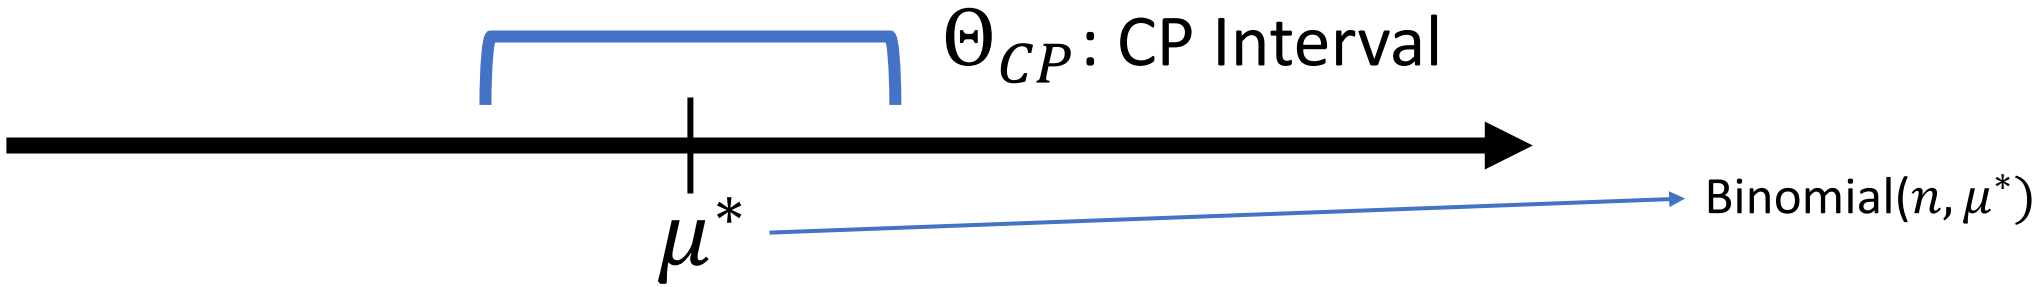

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

- Covariate Shift Detection

$\widehat{\boldsymbol{g}}(\boldsymbol{X}) = \boldsymbol{Y}$? ➡ Binomial$(n, \mu^*)$ ➡

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

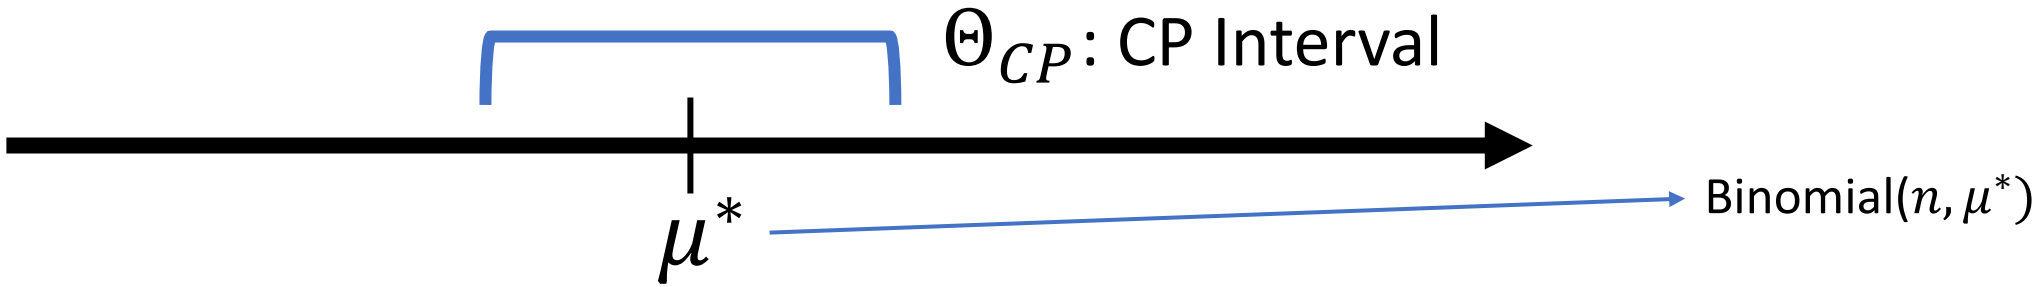

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

- Covariate Shift Detection

$\hat{\boldsymbol{g}}(\boldsymbol{X}) = \boldsymbol{Y}$? ➡ Binomial$(n, \mu^*)$ ➡ μ^*: Accuracy of $\hat{\boldsymbol{g}}$

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

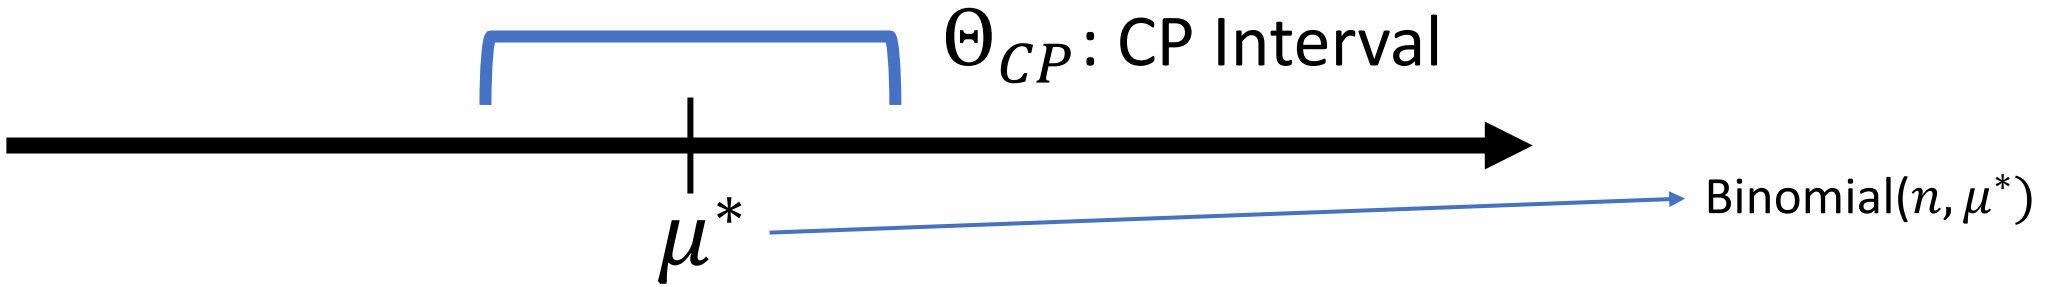

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

- Covariate Shift Detection

$\widehat{\boldsymbol{g}}(\boldsymbol{X}) = \boldsymbol{Y}$? ➡ Binomial$(n, \mu^*)$ ➡ μ^*: Accuracy of $\widehat{\boldsymbol{g}}$ ➡

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

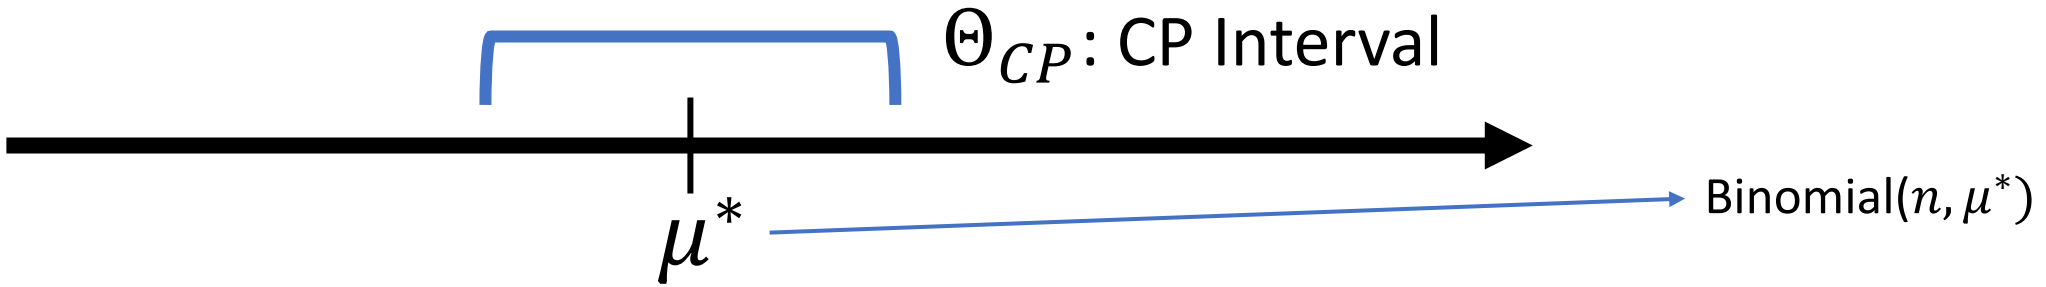

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

- Covariate Shift Detection

$\widehat{\boldsymbol{g}}(\boldsymbol{X}) = \boldsymbol{Y}$? ➡ Binomial$(n, \mu^*)$ ➡ μ^*: Accuracy of $\widehat{\boldsymbol{g}}$ ➡ $0.5 \notin \Theta_{CP}$

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

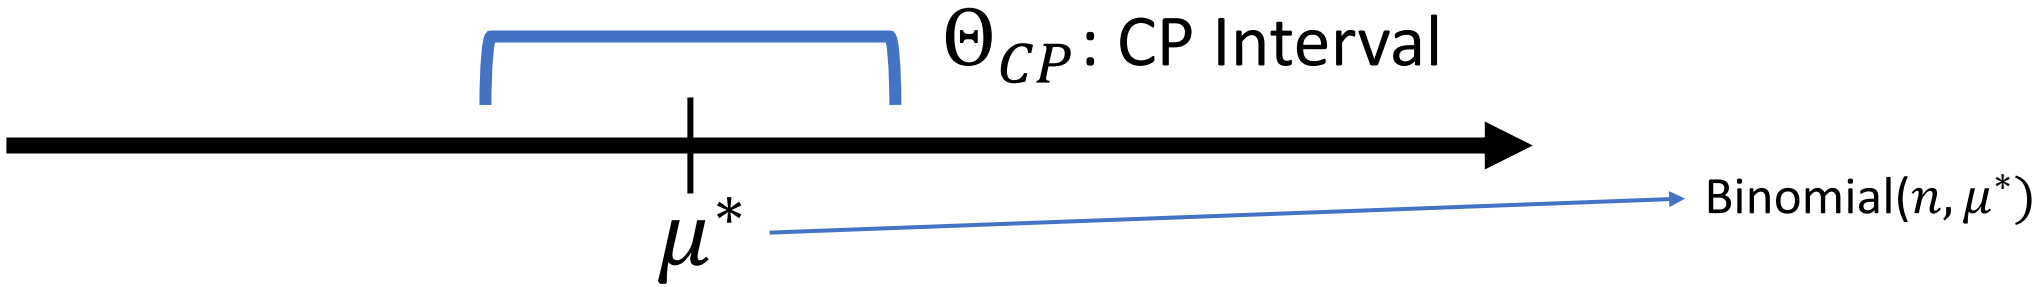

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

- Covariate Shift Detection

$\widehat{\boldsymbol{g}}(\boldsymbol{X}) = \boldsymbol{Y}$? ➡ Binomial$(n, \mu^*)$ ➡ μ^*: Accuracy of $\widehat{\boldsymbol{g}}$ ➡ $0.5 \notin \Theta_{CP}$ ➡ Shift

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

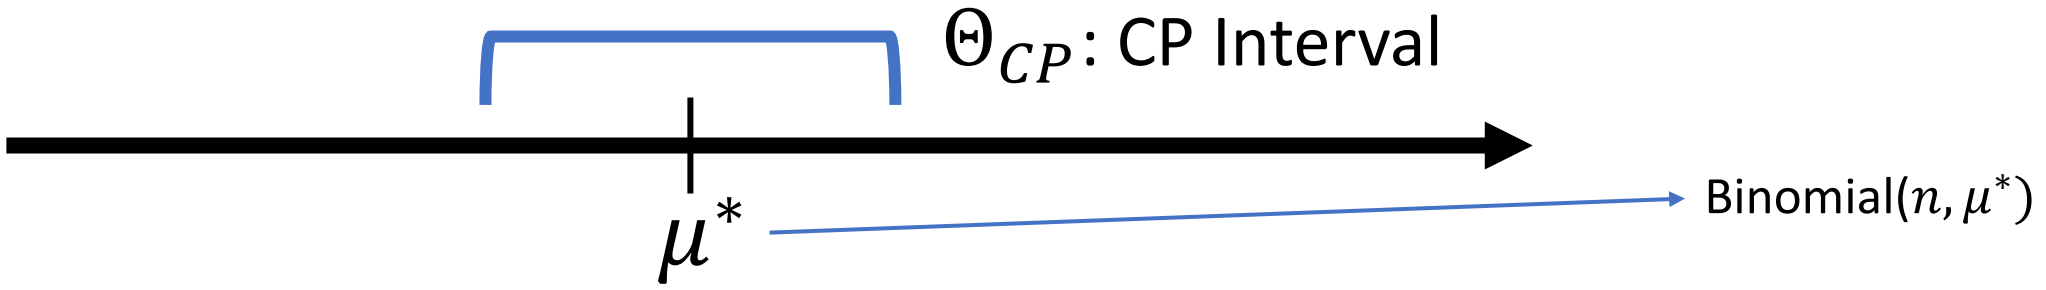

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

- Covariate Shift Detection

$\widehat{\boldsymbol{g}}(\boldsymbol{X}) = \boldsymbol{Y}$? ➡ Binomial$(n, \mu^*)$ ➡ μ^*: Accuracy of $\widehat{\boldsymbol{g}}$ ➡ $0.5 \notin \Theta_{CP}$ ➡ Shift

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

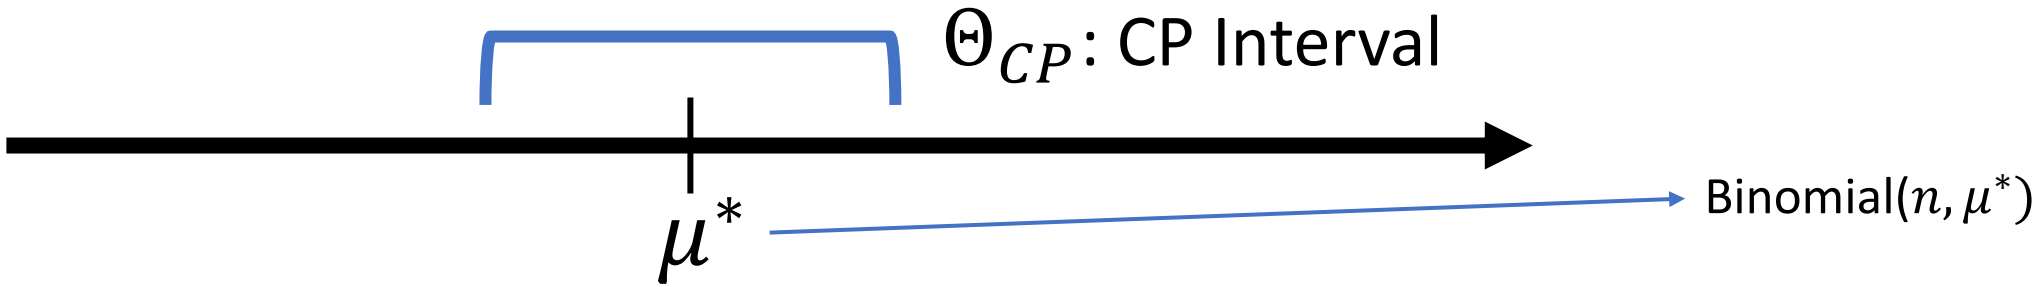

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

- Covariate Shift Detection

$\widehat{\boldsymbol{g}}(\boldsymbol{X}) = \boldsymbol{Y}$? ➡ Binomial$(n, \mu^*)$ ➡ μ^*: Accuracy of $\widehat{\boldsymbol{g}}$ ➡ $0.5 \notin \Theta_{CP}$ ➡ Shift

$0.5 \in \Theta_{CP}$

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

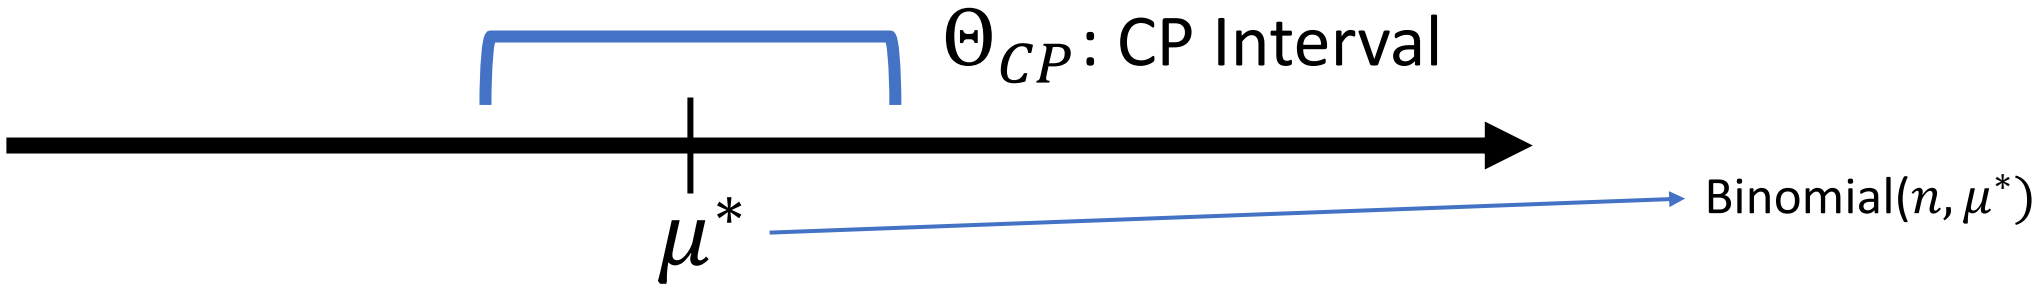

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

- Covariate Shift Detection

$\hat{\boldsymbol{g}}(\boldsymbol{X}) = \boldsymbol{Y}$? ➡ Binomial$(n, \mu^*)$ ➡ μ^*: Accuracy of $\hat{\boldsymbol{g}}$

- $0.5 \notin \Theta_{CP}$ ➡ Shift
- $0.5 \in \Theta_{CP}$ ➡ No - shift

Idea: Clopper-Pearson Interval

- Use the Clopper-Pearson (CP) Interval for covariate shift detection

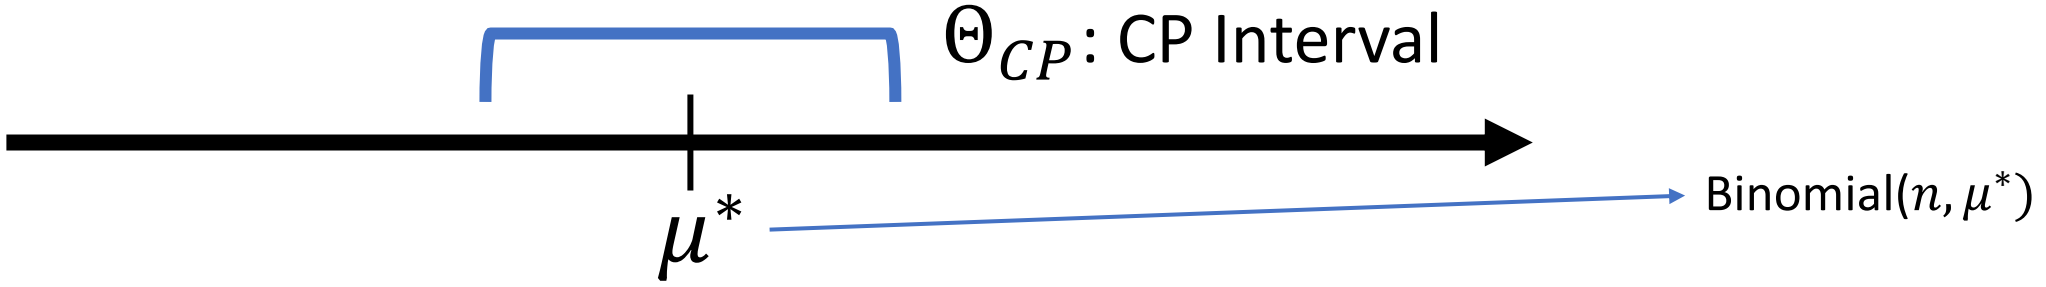

$$\mathbb{P}_{s\sim\text{Binomial}(n,\mu^*)}[\mu^* \in \Theta_{\text{CP}}(s, n; \alpha)] \geq 1 - \alpha,$$

- Covariate Shift Detection

$\hat{\boldsymbol{g}}(\boldsymbol{X}) = \boldsymbol{Y}$? ➡ Binomial$(n, \mu^*)$ ➡ μ^*: Accuracy of $\hat{\boldsymbol{g}}$

- $0.5 \notin \Theta_{CP}$ ➡ Shift
- $0.5 \in \Theta_{CP}$ ➡ No - shift

CP-Interval guarantees the finite-sample FPR bound

Idea: No Held-out Set for the CP Interval

- (Lemma 4.1) We can **reuse** samples for
  - source-target classifier update and
  - the CP interval construction

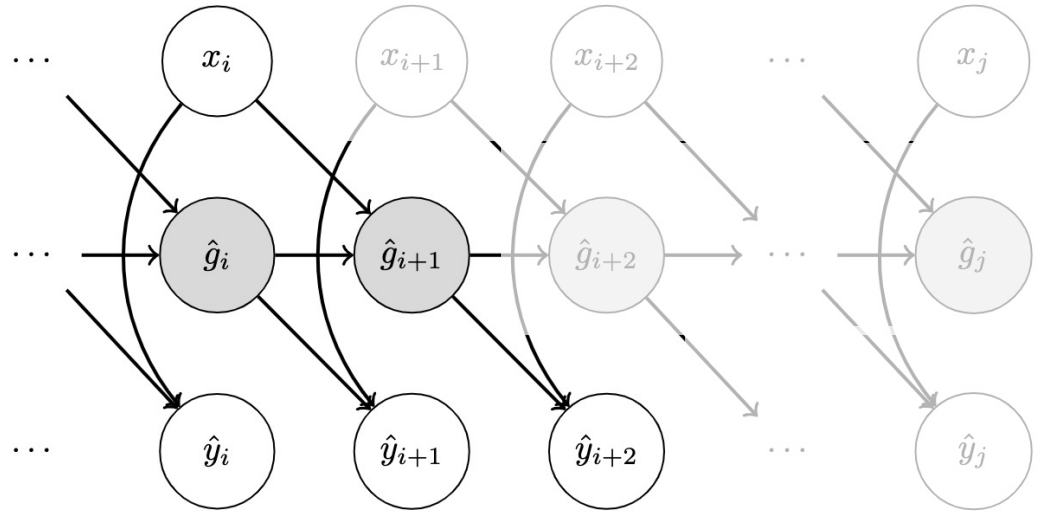

Idea: No Held-out Set for the CP Interval

- (Lemma 4.1) We can **reuse** samples for
  - source-target classifier update and
  - the CP interval construction

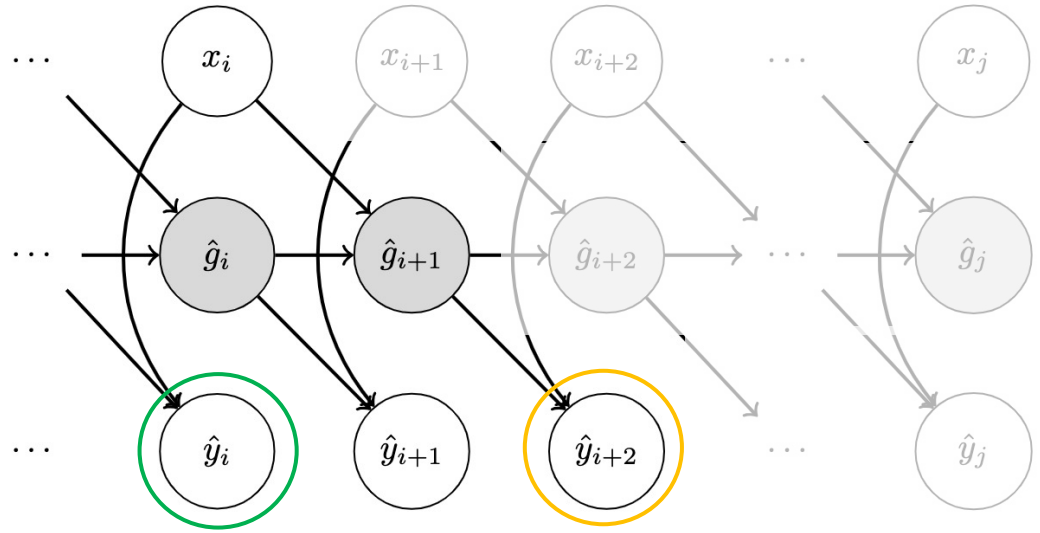

Idea: No Held-out Set for the CP Interval

- (Lemma 4.1) We can **reuse** samples for
  - source-target classifier update and
  - the CP interval construction

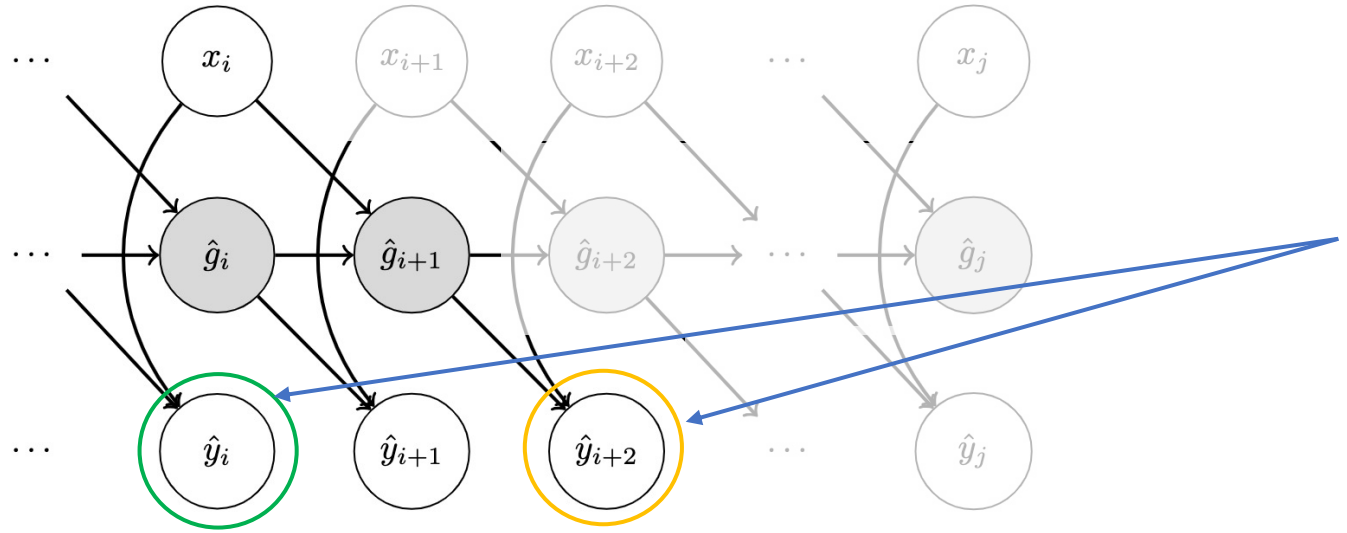

Conditionally Independent

Idea: No Held-out Set for the CP Interval

- (Lemma 4.1) We can **reuse** samples for
  - source-target classifier update and
  - the CP interval construction

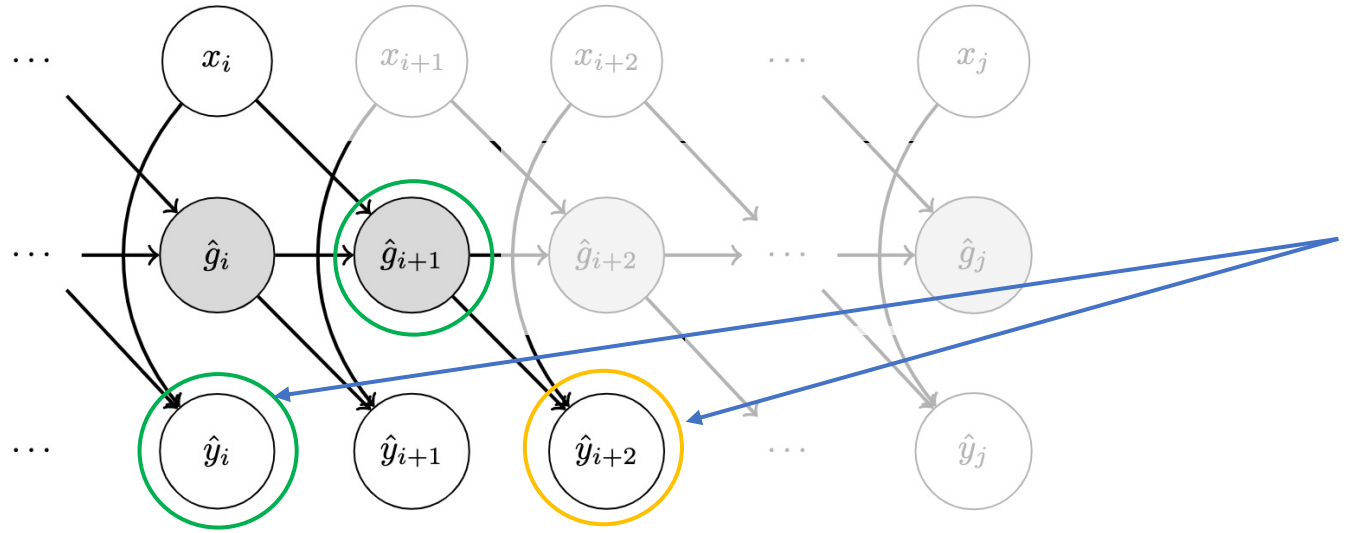

Conditionally Independent

Idea: No Held-out Set for the CP Interval

- (Lemma 4.1) We can **reuse** samples for
  - source-target classifier update and
  - the CP interval construction

Reuse

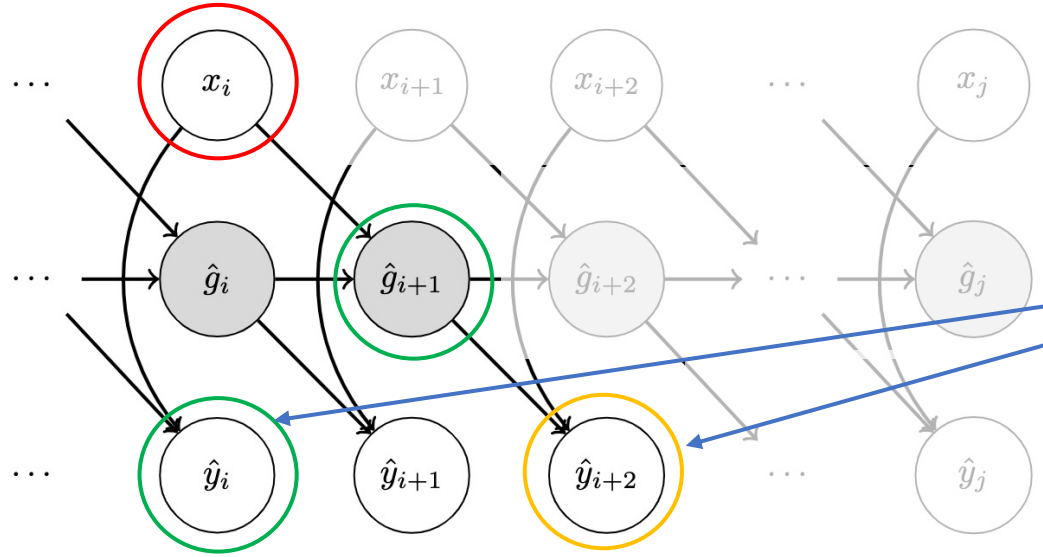

Conditionally Independent

Idea: No Held-out Set for the CP Interval

- (Lemma 4.1) We can **reuse** samples for
  - source-target classifier update and
  - the CP interval construction

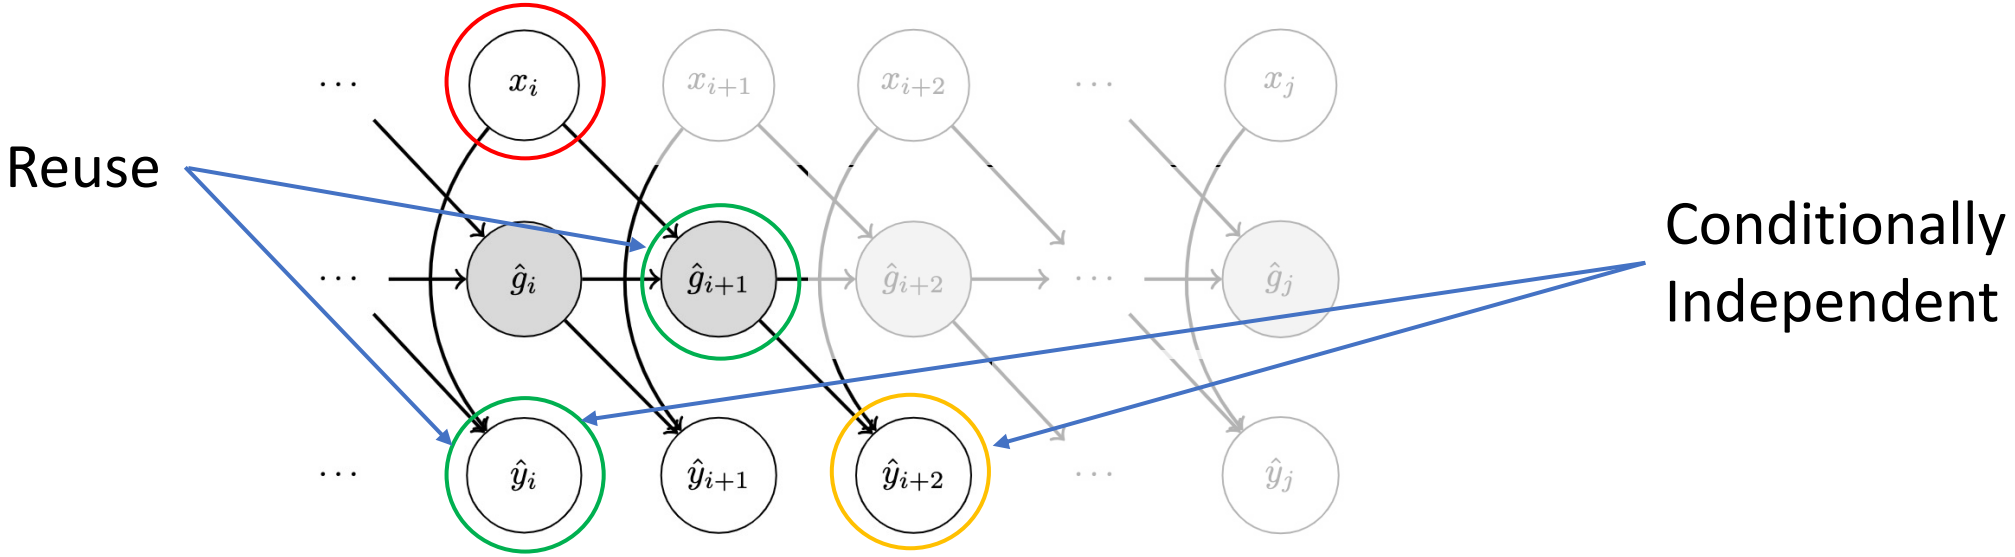

Idea: No Held-out Set for the CP Interval

- (Lemma 4.1) We can **reuse** samples for
  - source-target classifier update and
  - the CP interval construction

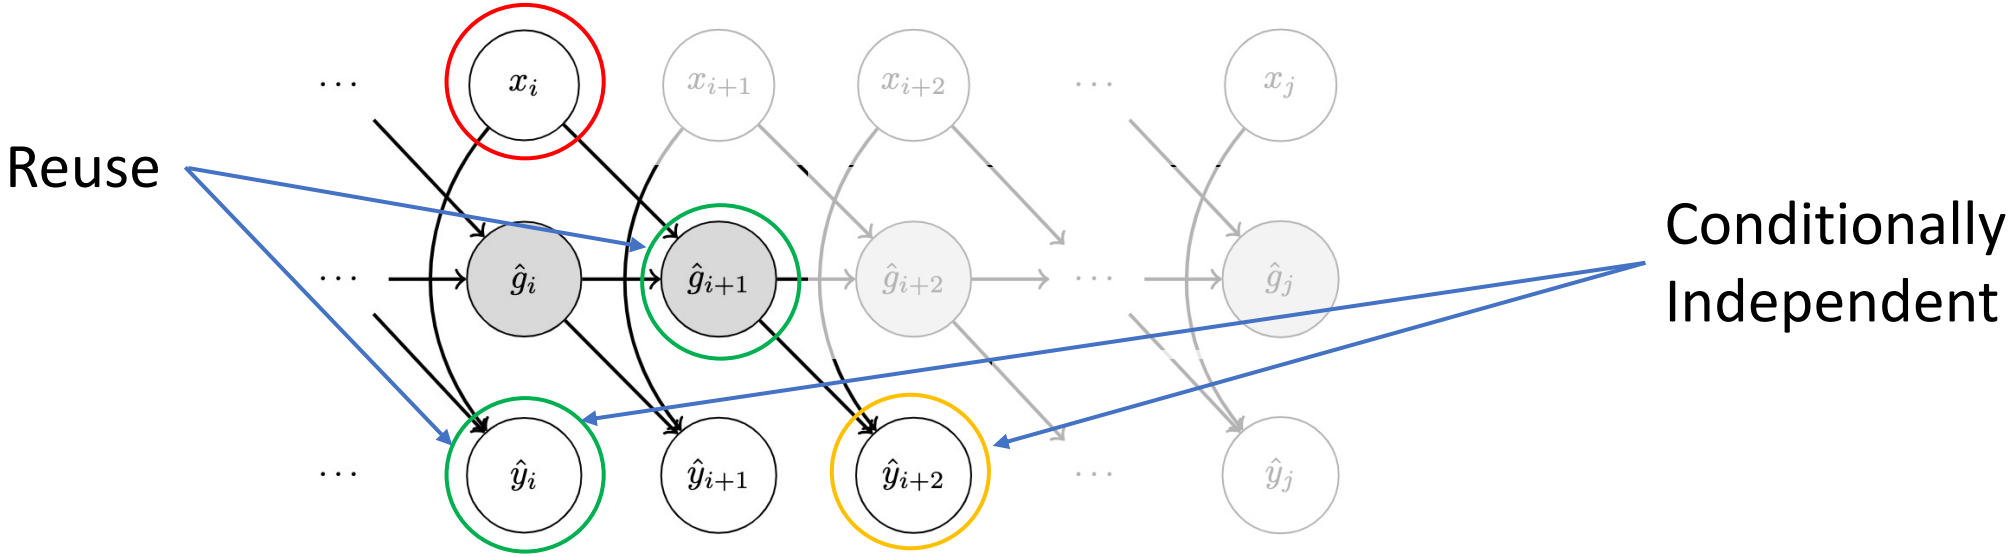

FPR Bound from CP interval is still valid

Idea: No Held-out Set for the CP Interval

- (Lemma 4.1) We can **reuse** samples for
  - source-target classifier update and
  - the CP interval construction

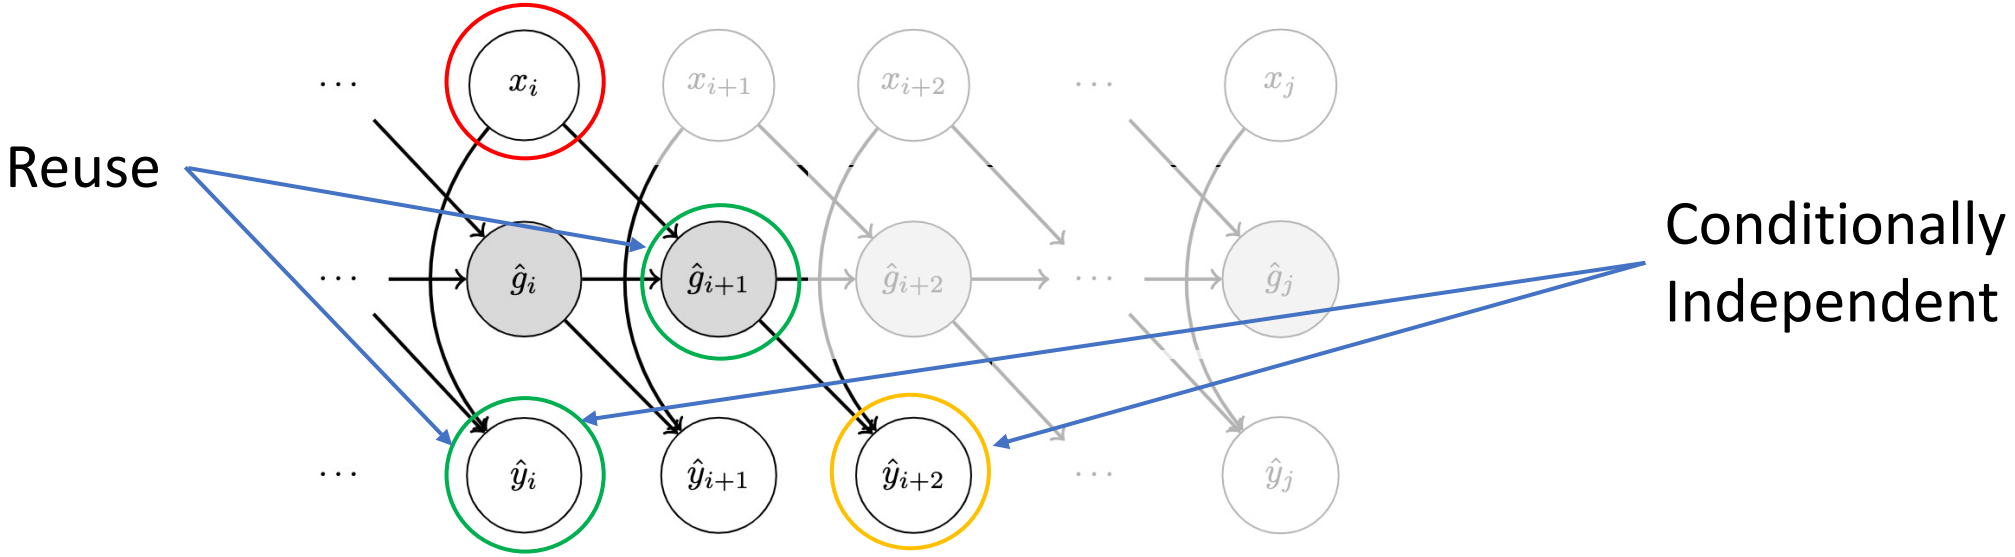

FPR Bound from CP interval is still valid

No held-out set is required: Sample Efficiency

Experiments

- Two shifts on ImageNet

Experiments

- Two shifts on ImageNet

Natural Shift
(Two sets of dog breeds)

Experiments

- Two shifts on ImageNet

Natural Shift
(Two sets of dog breeds)

Synthetic Shift
(Image Transformations)

Experiments

- Two shifts on ImageNet

Natural Shift
(Two sets of dog breeds)

vs

Synthetic Shift
(Image Transformations)

Experiments

- Two shifts on ImageNet

Natural Shift
(Two sets of dog breeds)

vs

Synthetic Shift
(Image Transformations)

vs

Shift Scenarios

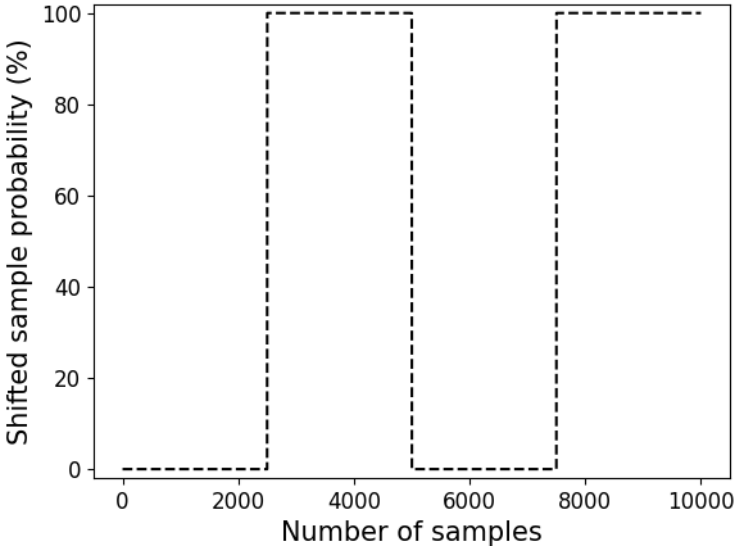

M-Shift

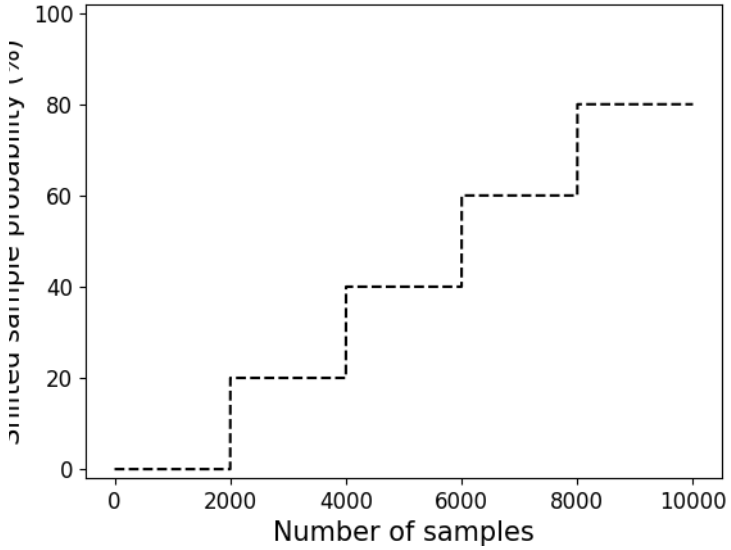

GI-Shift

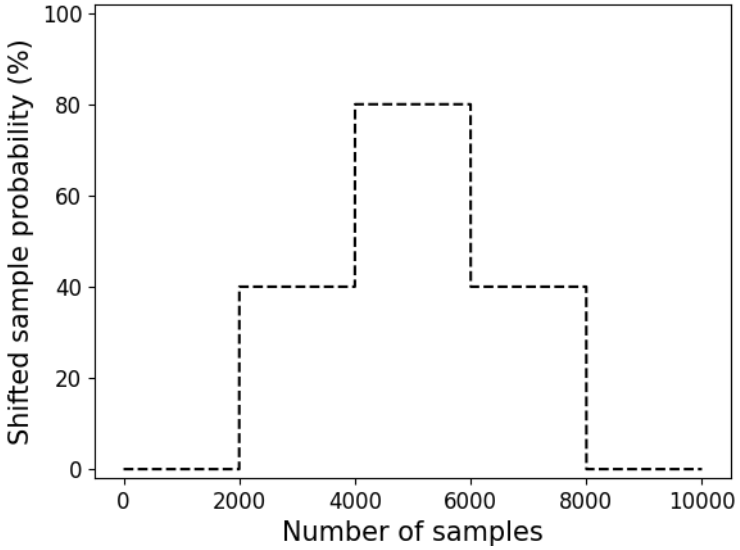

GID-Shift

Results - Natural shift

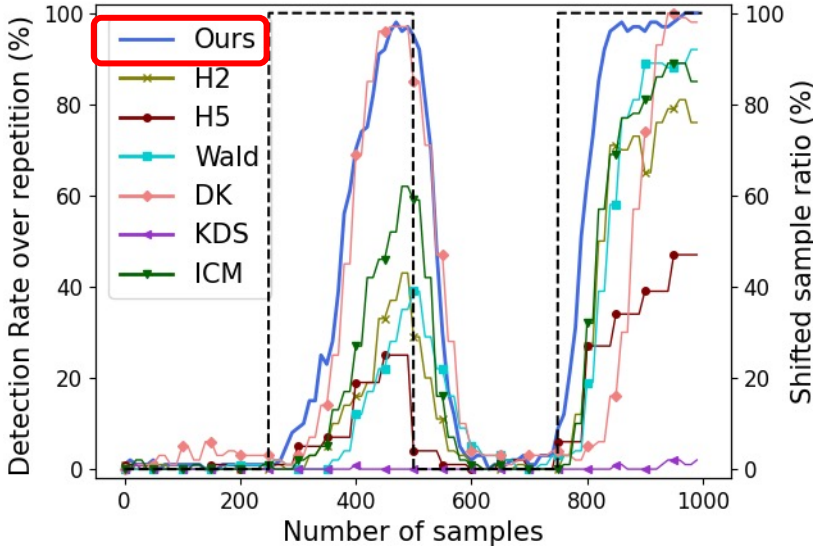

M-Shift

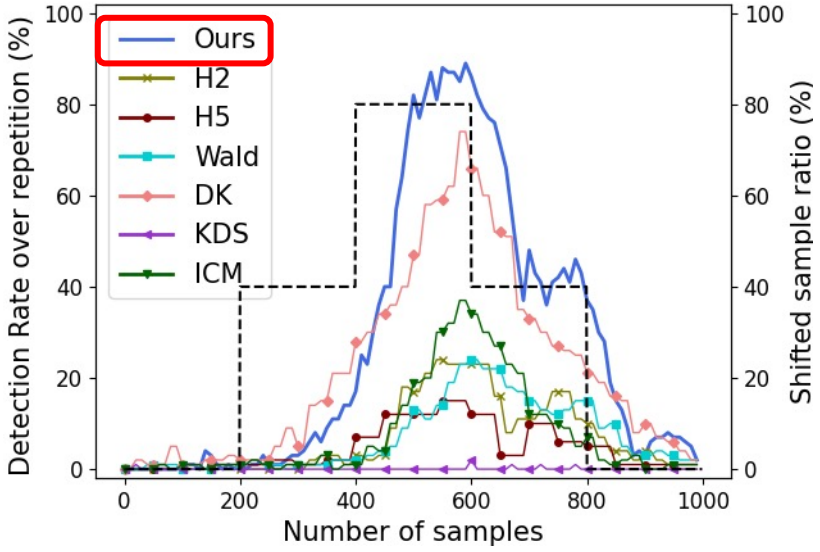

GI-Shift

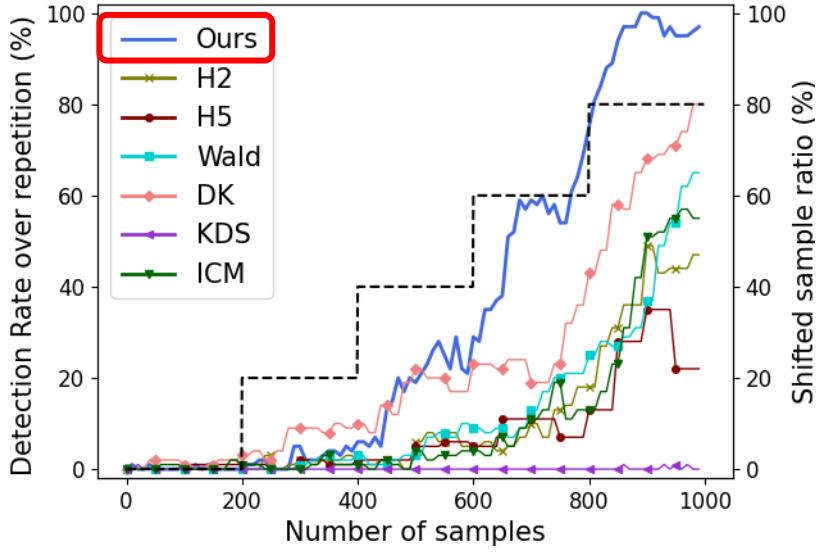

GID-Shift

Conclusion

Conclusion

- We propose a sequential covariate shift detection algorithm
  - Classifier Two-Sample tests and Clopper-Pearson Interval
  - FPR and FNR Bound
  - Sample efficient

Conclusion

- We propose a sequential covariate shift detection algorithm
  - Classifier Two-Sample tests and Clopper-Pearson Interval
  - FPR and FNR Bound
  - Sample efficient
- More experiments (IWildCam, Py150)

Conclusion

- We propose a sequential covariate shift detection algorithm
  - Classifier Two-Sample tests and Clopper-Pearson Interval
  - FPR and FNR Bound
  - Sample efficient
- More experiments (IWildCam, Py150)

https://github.com/sooyongj/sequential_covariate_shift_detection