Lecture 4: Computational Indistinguishability

Recall

- Distribution over Sample space

Recall

- Distribution over Sample space
- Distance between two distributions:

Recall

- Distribution over Sample space
- Distance between two distributions:
    - Prediction Advantage: Best strategy always outputs the most likely distribution

Recall

- Distribution over Sample space
- Distance between two distributions:
  - Prediction Advantage: Best strategy always outputs the most likely distribution
  - Total Variation Distance

Recall

- Distribution over Sample space
- Distance between two distributions:
  - Prediction Advantage: Best strategy always outputs the most likely distribution
  - Total Variation Distance
- Equivalent

Recall

- Distribution over Sample space
- Distance between two distributions:
  - Prediction Advantage: Best strategy always outputs the most likely distribution
  - Total Variation Distance
- Equivalent
- Think: Generalize to more than two distributions

Distinguisher

Given two probability ensembles $\{X_n\}$ and $\{Y_n\}$ and a n.u. TM D, if we have:

$$|\Pr[s \sim X_n \colon D(s) = 1] - \Pr[s \sim Y_n \colon D(s) = 1]| = \varepsilon(n)$$

then, we say that "D distinguishes $\{X_n\}$ and $\{Y_n\}$ with probability $\varepsilon(n)$."

Distinguisher

Given two probability ensembles $\{X_n\}$ and $\{Y_n\}$ and a n.u. TM D, if we have:

$$|\Pr[s \sim X_n : D(s) = 1] - \Pr[s \sim Y_n : D(s) = 1]| = \varepsilon(n)$$

then, we say that "D distinguishes $\{X_n\}$ and $\{Y_n\}$ with probability $\varepsilon(n)$."

- Predictive advantage: If n.u. TM D distinguishes $\{X_n\}$ and $\{Y_n\}$ with probability $\varepsilon(n)$ then D distinguishes $\{X_n\}$ and $\{Y_n\}$ with predictive advantage $\varepsilon(n)/2$

Distinguisher

Given two probability ensembles $\{X_n\}$ and $\{Y_n\}$ and a n.u. TM D, if we have:

$$|\Pr[s \sim X_n : D(s) = 1] - \Pr[s \sim Y_n : D(s) = 1]| = \varepsilon(n)$$

then, we say that "D distinguishes $\{X_n\}$ and $\{Y_n\}$ with probability $\varepsilon(n)$."

- Predictive advantage: If n.u. TM D distinguishes $\{X_n\}$ and $\{Y_n\}$ with probability $\varepsilon(n)$ then D distinguishes $\{X_n\}$ and $\{Y_n\}$ with predictive advantage $\varepsilon(n)/2$
- $\varepsilon(n)$-Indistinguishable: For all n.u. TM D we have

$$|\Pr[s \sim X_n : D(s) = 1] - \Pr[s \sim Y_n : D(s) = 1]| \leqslant \varepsilon(n)$$

Definition

Let $\{X_n\}$ and $\{Y_n\}$ be two probability distribution ensembles. For any n.u. PPT D, if we have:

$$|\Pr[s \sim X_n \colon D(s) = 1] - \Pr[s \sim Y_n \colon D(s) = 1]| \leqslant \varepsilon(n)$$

then, we say that "$\{X_n\}$ and $\{Y_n\}$ are $\varepsilon(n)$ computationally distinguishable." Represented by: $\{X_n\} \approx_{\varepsilon(n)} \{Y_n\}$.

Definition

Let $\{X_n\}$ and $\{Y_n\}$ be two probability distribution ensembles. For any n.u. PPT D, if we have:

$$|\Pr[s \sim X_n \colon D(s) = 1] - \Pr[s \sim Y_n \colon D(s) = 1]| \leqslant \varepsilon(n)$$

then, we say that "$\{X_n\}$ and $\{Y_n\}$ are $\varepsilon(n)$ computationally distinguishable." Represented by: $\{X_n\} \approx_{\varepsilon(n)} \{Y_n\}$.

If $\{X_n\}$ and $\{Y_n\}$ are $\nu(n)$ computationally indistinguishable, for some negligible function $\nu(\cdot)$, then we say that "$\{X_n\}$ and $\{Y_n\}$ are computationally indistinguishable" (represented by $\{X_n\} \approx \{Y_n\}$).

Data Processing Inequality

If $\{X_n\}$ and $\{Y_n\}$ are $\varepsilon(n)$ computationally indistinguishable, then for any n.u. PPT M we have: $\{M(X_n)\}$ and $\{M(Y_n)\}$ are $\varepsilon(n)$ computationally indistinguishable

Data Processing Inequality

If $\{X_n\}$ and $\{Y_n\}$ are $\varepsilon(n)$ computationally indistinguishable, then for any n.u. PPT M we have: $\{M(X_n)\}$ and $\{M(Y_n)\}$ are $\varepsilon(n)$ computationally indistinguishable

- Proof?

Data Processing Inequality

If $\{X_n\}$ and $\{Y_n\}$ are $\varepsilon(n)$ computationally indistinguishable, then for any n.u. PPT M we have: $\{M(X_n)\}$ and $\{M(Y_n)\}$ are $\varepsilon(n)$ computationally indistinguishable

- Proof?
- Special Case:
  $\{X_n\} \approx \{Y_n\} \implies \forall$n.u. PPT M: $\{M(X_n)\} \approx \{M(Y_n)\}$

Lemma (Hybrid Lemma)

Let $\{X_n^{(1)}\}, \{X_n^{(2)}\}, \ldots, \{X_n^{(m)}\}$ *be a set of probability ensembles. If there exists a n.u. PPT D (distinguisher) which distinguishes* $\{X_n^{(1)}\}$ *and* $\{X_n^{(m)}\}$ *with probability* $\varepsilon(n)$ *then there exists* $1 \leqslant i < m$ *such that n.u. PPT D distinguishes* $\{X_n^{(i)}\}$ *and* $\{X_n^{(i+1)}\}$ *with probability* $\varepsilon(n)/m$.

Lemma (Hybrid Lemma)

Let $\{X_n^{(1)}\}, \{X_n^{(2)}\}, \ldots, \{X_n^{(m)}\}$ *be a set of probability ensembles. If there exists a n.u. PPT D (distinguisher) which distinguishes* $\{X_n^{(1)}\}$ *and* $\{X_n^{(m)}\}$ *with probability* $\varepsilon(n)$ *then there exists* $1 \leqslant i < m$ *such that n.u. PPT D distinguishes* $\{X_n^{(i)}\}$ *and* $\{X_n^{(i+1)}\}$ *with probability* $\varepsilon(n)/m$.

- Proof?