Lecture 08: Computational Indistinguishability

Definition of Indistinguishability

- Let X and Y be probability distributions over the sample space Ω

Definition (Indistinguishability)

The distributions X and Y are ε-indistinguishable, represented by $X \approx_{\varepsilon} Y$, if for every adversary $\mathcal{A}\colon \Omega \to \{0,1\}$ the following holds:

$$\left|\Pr[\mathcal{A}(X) = 1] - \Pr[\mathcal{A}(Y) = 1]\right| \leqslant \varepsilon$$

Indistinguishability and Statistical Distance

Claim

$X \approx_\varepsilon Y \implies \mathrm{SD}\left(X, Y\right) \leqslant \varepsilon.$

Proof is left is an exercise. Try using various equivalent definitions of statistical distance as introduced in the previous lectures.

Computational Indistinguishability

- An algorithm $\mathcal{A}$ is *efficient* if its running time is bounded by a polynomial in its input-length

Definition (Computational Indistinguishability)

The distributions X and Y are ε-computationally indistinguishability, represented by $X \approx_{\varepsilon}^{(c)} Y$, if for every efficient $\mathcal{A}\colon \Omega \to \{0,1\}$ the following holds:

$$\left|\Pr[\mathcal{A}(X)=1] - \Pr[\mathcal{A}(Y)=1]\right| \leqslant \varepsilon$$

Claim

For any efficient $f : \Omega \to \Omega'$,

$$X \approx_{\varepsilon}^{(c)} Y \implies f(X) \approx_{\varepsilon}^{(c)} f(Y)$$

- We will prove the contrapositive, i.e. $\neg\left(f(X) \approx_{\varepsilon}^{(c)} f(Y)\right)$ implies $\neg\left(X \approx_{\varepsilon}^{(c)} Y\right)$
- The statement $\neg\left(f(X) \approx_{\varepsilon}^{(c)} f(Y)\right)$ implies that there exists an efficient $\mathcal{A} \colon \Omega' \to \{0,1\}$ such that

$$|\Pr[\mathcal{A}(f(X)) = 1] - \Pr[\mathcal{A}(f(Y)) = 1]| > \varepsilon$$

Proof Continued

- Let $\widetilde{\mathcal{A}}\colon \Omega \to \{0,1\}$ be the function defined as followed: $\widetilde{\mathcal{A}}(s) = \mathcal{A}(f(s))$
- Note that $\widetilde{\mathcal{A}}$ is efficient because $\mathcal{A}$ and f are both efficient
- Note that $\widetilde{\mathcal{A}}(X) \equiv \mathcal{A}(f(X))$ and $\widetilde{\mathcal{A}}(Y) \equiv \mathcal{A}(f(Y))$
- Then we have demonstrated that there exists an adversary $\widetilde{\mathcal{A}}$ such that:

$$\left|\Pr[\widetilde{\mathcal{A}}(X) = 1] - \Pr[\widetilde{\mathcal{A}}(Y) = 1]\right| > \varepsilon$$

- This shows $\neg\left(X \approx_{\varepsilon}^{(c)} Y\right)$

Claim

$$X^{(0)} \approx^{(c)}_{\varepsilon_1} X^{(1)} \approx^{(c)}_{\varepsilon_2} X^{(2)} \implies X^{(0)} \approx^{(c)}_{\varepsilon_1+\varepsilon_2} X^{(2)}$$

- We will prove the contrapositive
- Assume that there exists an efficient $\mathcal{A}$ such that:

$$\left|\Pr[\mathcal{A}(X^{(0)}) = 1] - \Pr[\Pr[\mathcal{A}(X^{(2)}) = 1]\right| > \varepsilon_1 + \varepsilon_2$$

- We want to construct two adversaries $\widetilde{\mathcal{A}}$ and $\widetilde{\mathcal{B}}$ such that: At least one of the following statements holds

$$\left|\Pr[\widetilde{\mathcal{A}}(X^{(0)}) = 1] - \Pr[\widetilde{\mathcal{A}}(X^{(1)}) = 1]\right| > \varepsilon_1$$
$$\left|\Pr[\widetilde{\mathcal{B}}(X^{(1)}) = 1] - \Pr[\widetilde{\mathcal{B}}(X^{(2)}) = 1]\right| > \varepsilon_2$$

- Proof is left as an exercise. Hint: Use $\widetilde{\mathcal{A}} = \widetilde{\mathcal{B}} = \mathcal{A}$.