TD 1: Play with definitions (corrected version)

Exercise 1. *Statistical distance*

Definition 1 (Statistical distance). *Let X and Y be two discrete random variables over a countable set A. The statistical distance between X and Y is the quantity*

$$\Delta(X,Y) = \frac{1}{2}\sum_{a\in A} |\Pr[X=a] - \Pr[Y=a]|.$$

The statistical distance verifies usual properties of distance function, i.e., it is a positive definite binary symmetric function that satisfies the triangle inequality:

- $\Delta(X,Y) \geq 0$, with equality if and only if X and Y are identically distributed,
- $\Delta(X,Y) = \Delta(Y,X)$,
- $\Delta(X,Z) \leq \Delta(X,Y) + \Delta(Y,Z)$.

1. Show that if $\Delta(X,Y) = 0$, then for any deterministic adversary $\mathcal{A}$, we have $\mathrm{Adv}_{\mathcal{A}}(X,Y) = 0$.

☞ By definition, $\mathrm{Adv}_{\mathcal{A}}(X,Y) = |\Pr_{a\leftarrow X}[\mathcal{A}(a)=1] - \Pr_{a\leftarrow Y}[\mathcal{A}(a)=1]|$. Since $\Delta(X,Y)=0$, we directly obtain that $\Pr[X=a] = \Pr[Y=a]$ for all $a \in S$, or in other words, X and Y are identically distributed. As a result, $\Pr_{a\leftarrow X}[\mathcal{A}(a)=1] = \Pr_{a\leftarrow Y}[\mathcal{A}(a)=1]$ and thus $\mathrm{Adv}_{\mathcal{A}}(X,Y) = 0$.

In the next question, we will prove the *data processing inequality* for the statistical distance.

2. Let X, Y be two random variables over a common set A.

(a) Let $f : A \to S$ be a deterministic function with domain S. Show that

$$\Delta(f(X), f(Y)) \leq \Delta(X,Y).$$

☞ We write the definition of Δ.

$$\Delta(f(X),f(Y)) = \frac{1}{2}\sum_{s\in S}|\Pr(f(X)=s) - \Pr(f(Y)=s)|$$

Then decompose the event $f(X)=s$ into something more explicit.

$$\Delta(f(X),f(Y)) = \frac{1}{2}\sum_{s\in S}\left|\sum_{a\in f^{-1}(s)}\Pr(X=a) - \sum_{a\in f^{-1}(s)}\Pr(Y=a)\right|$$

Now use the triangle inequality.

$$\Delta(f(X),f(Y)) \leq \frac{1}{2}\sum_{s\in S}\sum_{a\in f^{-1}(s)}|\Pr(X=a) - \Pr(Y=a)|$$

Finally, recall that $\sqcup_{s\in S} f^{-1}(s) = A$, and this ends the proof.

(b) Let Z be another random variable with domain $\mathcal{Z}$, statistically independent from X and Y. Show that

$$\Delta((X,Z),(Y,Z)) = \Delta(X,Y).$$

☞ Once again, we write the definition of the statistical distance.

$$\begin{aligned}\Delta((X,Z),(Y,Z)) &= \sum_{(a,z)\in A\times\mathcal{Z}}|\Pr(X=a\wedge Z=z) - \Pr(Y=a\wedge Z=z)|\\ &= \sum_{(a,z)\in A\times\mathcal{Z}}|\Pr(Z=z)\cdot(\Pr(X=a) - \Pr(Y=a))|\\ &= \sum_{z\in\mathcal{Z}}\Pr(Z=z)\cdot\sum_{a\in A}|\Pr(X=a) - \Pr(Y=a)|.\end{aligned}$$

And this is exactly $\Delta(X,Y)$.

(c) Let f be a (possibly probabilistic) function with domain S. Define f' a deterministic function and R a random variable independent from X and Y such that for any input x, we have $f'(x,R) = f(x)$. The random variable R is the internal randomness of f. Using f' and R, show that $\Delta(f(X), f(Y)) = \Delta(f'(X,R), f'(Y,R)) \leq \Delta(X,Y)$.

☞ We apply the two previous results: $\Delta(f(X), f(Y)) \leq \Delta((X,R),(Y,R)) = \Delta(X,Y)$.

3. Show that for any (possibly probabilistic) adversary $\mathcal{A}$, we have $\mathrm{Adv}_{\mathcal{A}}(X,Y) \leq \Delta(X,Y)$.

☞ This follows from the definition of the advantage, and from the above property ($\mathcal{A}$ is a function):

$$\mathrm{Adv}_{\mathcal{A}}(X,Y) = |\Pr[\mathcal{A}(X)=1] - \Pr[\mathcal{A}(Y)=1]| = \frac{1}{2}\sum_{b\in\{0,1\}} |\Pr[\mathcal{A}(X)=b] - \Pr[\mathcal{A}(Y)=b]| = \Delta(\mathcal{A}(X),\mathcal{A}(Y)) \leq \Delta(X,Y).$$

4. Assuming the existence of a secure PRG $G : \{0,1\}^s \to \{0,1\}^n$, show that $\Delta(G(U(\{0,1\}^s)), U(\{0,1\}^n))$ can be much larger than $\max_{\mathcal{A}} \mathrm{Adv}_{\mathcal{A}}(G(U(\{0,1\}^s)), U(\{0,1\}^n))$.

☞ By definition,

$$\begin{aligned}
\Delta(G(U(\{0,1\}^s)), U(\{0,1\}^n)) &= \frac{1}{2}\sum_{a\in\{0,1\}^n} |\Pr[G(U(\{0,1\}^s))=a] - \Pr[U(\{0,1\}^n)=a]| \\
&= \frac{1}{2}\left(\sum_{\substack{a\in\{0,1\}^n \\ a\notin G(\{0,1\}^s)}} \left|0 - \frac{1}{2^n}\right| + \sum_{\substack{a\in\{0,1\}^n \\ a\in G(\{0,1\}^s)}} \left|\Pr[G(U(\{0,1\}^s))=a] - \frac{1}{2^n}\right|\right) \\
&= \frac{1}{2} - \frac{\#G(\{0,1\}^s)}{2\cdot 2^n} + \frac{1}{2}\sum_{\substack{a\in\{0,1\}^n \\ a\in G(\{0,1\}^s)}} \left(\Pr[G(U(\{0,1\}^s))=a] - \frac{1}{2^n}\right) \\
&= 1 - \frac{\#G(\{0,1\}^s)}{2^{n+1}} - \frac{\#G(\{0,1\}^s)}{2^{n+1}} \\
&\geq 1 - 2^{s-n}.
\end{aligned}$$

At line 3, we use the fact that for $a \in G(\{0,1\}^s)$, we have $\Pr[G(U(\{0,1\}^s))] = a \geq 1/2^s \geq 1/2^n$ (because at least one element b is such that $G(b) = a$ and as b is chosen uniformly in $\{0,1\}^s$, this happens with probability at least 2^s). We also use the fact that $\sum_{\substack{a\in\{0,1\}^n \\ a\in G(\{0,1\}^s)}} \Pr[G(U(\{0,1\}^s))=a] = 1$.

For the last inequality, we use the fact that $\#G(\{0,1\}^s) \leq 2^s$. As in the lecture we assumed $n \gg s$, then in particular as soon as $n > s+1$, the statistical distance will be greater than $1/2$.

On the contrary, as G is a secure PRG, then by definition $\max_{\mathcal{A}} \mathrm{Adv}_{\mathcal{A}}(G(U(\{0,1\}^s)), U(\{0,1\}^n))$ is negligible, i.e. much smaller than $1/2$.

Exercise 2. *A weird distinguisher...*

We consider two distributions D_0 and D_1 over $\{0,1\}^n$.

1. Recall the definitions that were given in class for the notions of *distinguisher*, *advantage* and *indistinguishability* of D_0 and D_1.

☞ To sum up the behavior of a distinguisher $\mathcal{A}$, two experiments $Exp_b, b \in \{0,1\}$ can be defined as follows.

$\mathcal{C}$	$\mathcal{A}$
sample $x \hookleftarrow D_b$ send x to $\mathcal{A}$	
	compute a bit b' output b'

Then, we consider the advantage $\mathrm{Adv}(\mathcal{A}) = |\Pr[\mathcal{A} \xrightarrow{Exp_0} 1] - Pr[\mathcal{A} \xrightarrow{Exp_1} 1]|$; the distributions D_0 and D_1 are said to be indistinguishable if $\mathrm{Adv}(\mathcal{A})$ is negligible for any PPT $\mathcal{A}$.

You found a distinguisher $\mathcal{A}$ on internet. However, you cannot find anywhere in the documentation its performances!

2. Assuming that you have access to as many samples as you like from D_0 and D_1 (you can for instance assume that you can sample yourself from these distributions), how would you estimate the advantage of $\mathcal{A}$? *Hint: use the Chernoff Bound:* $\Pr(|X - np| \geq nt) \leq 2\exp(-2nt^2)$, *where* X *follows a binomial distribution with parameters* (n, p). ☞ Run N times Exp 0 and Exp 1 for a number N to be determined later. This gives us $b_1^{(1)}, \dots b_1^{(N)}$ and $b_2^{(1)}, \dots, b_2^{(N)}$, $2N$ results. Define

$$\bar{b}_1 := \frac{\sum_{i=1}^{N} b_1^{(i)}}{N} \text{ and } \bar{b}_2 := \frac{\sum_{i=1}^{N} b_2^{(i)}}{N}.$$

Then let p_b be the probability that $\mathcal{A}$ outputs 1 at the end of Exp b. The Chernoff bound gives

$$\Pr(|\bar{b}_i - p_i| \geq \varepsilon) \leq 2\exp(-2N\varepsilon^2),$$

for any accuracy $\varepsilon > 0$. Then notice the following sequence of inequalities:

$$\mathsf{Adv}(\mathcal{A}) = |p_1 - p_0| \leq |p_1 - \bar{b}_1| + |\bar{b}_1 - \bar{b}_0| + |\bar{b}_0 - p_0| \leq 2\varepsilon + |\bar{b}_1 - \bar{b}_0|,$$

where the last inequality holds with probability at least $1 - 4\exp(-2N\varepsilon^2)$. The same sequence can be written by reversing the roles of p_b and $\bar{b}_p$. This gives us $|\mathsf{Adv}(\mathcal{A}) - |\bar{b}_1 - \bar{b}_0|| \leq 2\varepsilon$ with probability at least $1 - 4\exp(-2N\varepsilon^2)$.

Assuming that you want to compute the advantage with accuracy $\frac{1}{\lambda^c}$ and probability 0.95, set $\varepsilon := \frac{1}{2\lambda^c}$ and N such that $1 - 4\exp(N/(2\lambda^{2c})) \geq 0.95$ i.e. $N/\lambda^{2c} \geq 2\ln(80) \approx 8.76$.

By convention, you want to design a distinguisher such that it outputs 1 when it thinks the sample comes from D_1 and 0 otherwise. However, because of the definition of advantage, it is also possible to design distinguishers that do the reverse, and still have the same advantage. For instance, the above distinguisher $\mathcal{A}$ may often be "wrong". This could be troublesome if your aim is to use its output to do further computations. Luckily, there exists a way to transform $\mathcal{A}$ into a distinguisher that is more often right than wrong, whatever it previously did.

3. The definition of advantage from question 1 may be called Absolute Advantage, for the purpose of this exercise. In this question, we define the Positive Advantage of $\mathcal{A}$ as

$$\mathsf{Adv}_P(\mathcal{A}) := \Pr(\mathcal{A} \xrightarrow{Exp_1} 1) - \Pr(\mathcal{A} \xrightarrow{Exp_0} 1).$$

Given a distinguisher $\mathcal{A}$ with Absolute Advantage ε, we build a distinguisher $\mathcal{A}'$ that does the following:

1. Upon receiving a sample $y \hookleftarrow D_b$, it runs $b' \leftarrow \mathcal{A}(y)$.
2. It samples $x_0 \hookleftarrow D_0$ and $x_1 \hookleftarrow D_1$ and runs $b_0 \leftarrow \mathcal{A}(x_0)$ and $b_1 \leftarrow \mathcal{A}(x_1)$.
3. It returns b' if $b_0 = 0$ and $b_1 = 1$. It returns $1 - b'$ if $b_0 = 1$ and $b_1 = 0$.
4. In any other cases, it returns a uniform bit.

Prove that the Positive Advantage of $\mathcal{A}'$ is ε^2.

☞ The probability that $\mathcal{A}$ outputs 1 in experience Exp b is $p_1(1-p_0)p_b + p_0(1-p_1)(1-p_b) + \frac{1}{2}(p_0p_1 + (1-p_0)(1-p_1))$. The positive advantage of $\mathcal{A}'$ is then:

$$\begin{aligned}\mathsf{Adv}_P(\mathcal{A}') &= p_1(1-p_0)(p_1-p_0) + p_0(1-p_1)(p_0-p_1)\\ &= (p_1-p_0)\cdot(p_1(1-p_0) - p_0(1-p_1))\\ &= (p_1-p_0)\cdot(p_1 - p_0p_1 - p_0 + p_0p_1)\\ &= \varepsilon^2.\end{aligned}$$

Exercise 3. *Bit-flip of a PRG*

Let G a pseudo-random generator (PRG) of input range $\{0,1\}^s$ and output range $\{0,1\}^n$. We define $\bar{G}$ as follows:

$$\forall x \in \{0,1\}^s, \bar{G}(x) := 1^n \oplus G(x),$$

where $\oplus$ denotes the XOR operation. This corresponds to flipping every bit of the output of G.

1. Prove that $\bar{G}$ is secure if and only if G is secure.

   ☞ Assume that G is secure. We will prove that $\bar{G}$ is secure. Assume by contradiction that there exists an adversary $\mathcal{A}$ that distinguishes between $\bar{G}(U(\{0,1\}^s))$ and $U(\{0,1\}^n)$ with non-negligible advantage. We build $\mathcal{A}'$ a distinguisher between $G(U(\{0,1\}^s))$ and $U(\{0,1\}^n)$ the following way: on input a sample y, $\mathcal{A}'$ calls $\mathcal{A}$ on the sample $1^n \oplus y$. It outputs the same value.

   Notice the following: if y is uniformly distributed, then so is $1^n \oplus y$. If y follows the distribution $G(U(\{0,1\}^s))$, then $1^n \oplus y$ follows the distribution $\bar{G}(U(\{0,1\}^s))$. Then $\mathcal{A}$'s view is exactly as intended. It guesses from which distribution is sampled $1^n \oplus y$ with non-negligible advantage, and the advantage of $\mathcal{A}'$ is equal to the advantage of $\mathcal{A}$, which contradicts the assumption that G is secure.

   Finally, we notice that the flipped version of $\bar{G}$ is G, and the previous proof also shows that $\bar{G}$ secure implies G secure.

Exercise 4. *(Bonus) Variable-length OTP is not secure*

A *variable length one-time pad* is a cipher (E, D), where the keys are bit strings of some fixed length L, while messages and ciphertexts are variable length bit strings, of length at most L. Thus, the cipher (E, D) is defined over $(\mathcal{K}, \mathcal{M}, \mathcal{C})$, where

$$\mathcal{K} := \{0,1\}^L \text{ and } \mathcal{M} := \mathcal{C} = \{0,1\}^{\leq L}$$

for some parameter L. Here, $\{0,1\}^{\leq L}$ denotes the set of all bit strings of length at most L (including the empty string). For a key $k \in \{0,1\}^L$ and a message $m \in \{0,1\}^{\leq L}$ of length ℓ, the encryption function is defined as follows:

$$E(k, m) := k[0 \,..\, \ell - 1] \oplus m$$

1. Provide a counter-example showing that the variable length OTP is not secure for perfect secrecy.

   ☞ Assume by contradiction that the variable-length OTP is secure for some message distribution $\mathcal{M}$. Let $k \in \{0,1\}^L$ be some key and for any $m \in \{0,1\}^{\leq L}$, let $c \in \{0,1\}^{\leq L}$ such that $|m| \neq |c|$. We then have $\Pr_{M \hookleftarrow \mathcal{M}}(M = m | E(k, M) = c) = 0 = \Pr_{M \hookleftarrow \mathcal{M}}(M = m)$. This means that any message has zero probability: this contradicts the fact that $\mathcal{M}$ is a probability distribution over $\{0,1\}^{\leq L}$. From A Graduate Course in Applied Cryptography, Dan Boneh and Victor Shoup Example 2.5 p9-10