

A Connection between Indistinguishability and Semantic Security

Sven Laur

University of Tartu

1 Formal Definitions

Intuitively, if objects are indistinguishable then we cannot determine their personal properties which vary among the population. Indeed, if we could reliably determine whether an object is green or not then we could easily distinguish green objects from yellow ones. Hence, indistinguishability indirectly implies that one can reliably detect only trivial properties that either hold or do not hold for the entire population. This basic argumentation template forms the cornerstone of contemporary cryptography. In the following, we fill out all details that are needed to convert this informal argumentation to a formal proof.

Let \mathcal{S}_0 be a distribution of secret values s and let $\text{supp}(\mathcal{S}_0)$ be the corresponding support. Then we can define indistinguishability of states $\text{supp}(\mathcal{S}_0)$ w.r.t. a function $f : \mathcal{S}_0 \rightarrow \mathcal{X}$. We say that states from $\text{supp}(\mathcal{S}_0)$ are (t, ε) -*indistinguishable* if for any $s_0, s_1 \in \text{supp}(\mathcal{S}_0)$ and for any t -time algorithm \mathcal{A} :

$$\text{Adv}_{s_0, s_1}^{\text{ind}}(\mathcal{A}) = |\Pr[x \leftarrow f(s_0) : \mathcal{A}(x) = 1] - \Pr[x \leftarrow f(s_1) : \mathcal{A}(x) = 1]| \leq \varepsilon .$$

To define semantic security, we have to formalise which properties of a hidden state are plausible or relevant. Essentially, we can talk about semantic security w.r.t. *computable* functions $g : \mathcal{S}_0 \rightarrow \mathcal{Y}$. Of course, for fixed input and output domains \mathcal{S}_0 and \mathcal{Y} all functions are computable and thus this restriction is only a cosmetic addition that represents our intent. In some scenarios, there are obvious restrictions to the function g . For instance, we might require that the output of g must be computable during the next one hundred years.

Now note that for any function g there is a trivial predictor algorithm \mathcal{A}_* that outputs the most probable outcome of g over the distribution \mathcal{S}_0 and thus the advantage of an algorithm \mathcal{A} is defined as the difference

$$\Pr[s \leftarrow \mathcal{S}_0, x \leftarrow f(s) : \mathcal{A}(x) = g(s)] - \Pr[s \leftarrow \mathcal{S}_0, x \leftarrow f(s) : \mathcal{A}_*(x) = g(s)] .$$

More formally, we say that states $\text{supp}(\mathcal{S}_0)$ are (t, ε) -*semantically secure* w.r.t. functions f and g if for any t -time adversary \mathcal{A} the corresponding advantage

$$\text{Adv}_{f, g}^{\text{sem}}(\mathcal{A}) = \Pr[\mathcal{A}(x) = g(s)] - \max_{y_* \in \mathcal{Y}} \Pr[s \leftarrow \mathcal{S}_0 : g(s) = y_*] \leq \varepsilon .$$

2 Indistinguishability Implies Semantic Security

The main aim of this section is to prove the classical theorem about semantic security and introduce basic concepts of game-playing proofs.

Theorem 1. *If states from $\text{supp}(\mathcal{S}_0)$ are $(2t, \varepsilon)$ -indistinguishable w.r.t. the function f , then states $\text{supp}(\mathcal{S}_0)$ are also (t, ε) -semantically for all functions g , i.e., $\text{Adv}_{f,g}^{\text{sem}}(\mathcal{A}) \leq \varepsilon$ for all t -time adversaries \mathcal{A} .*

Proof. In order to present the proof modularly in easily understandable steps, we start from the basic semantic security game

$$\mathcal{G}^{\mathcal{A}} \left[\begin{array}{l} s \leftarrow \mathcal{S}_0 \\ x \leftarrow f(s) \\ y \leftarrow \mathcal{A}(x) \\ \mathbf{return} [y \stackrel{?}{=} g(s)] \end{array} \right]$$

and then gradually rewrite the game until we obtain the desired bound on the success probability

$$\Pr [\mathcal{G}^{\mathcal{A}} = 1] \leq \Pr [\mathcal{G}^{\mathcal{A}^*} = 1] + \varepsilon .$$

COIN FIXING. Note that g does not have to be a deterministic function. However, if $\text{Adv}_{f,g}^{\text{sem}}(\mathcal{A}) > \varepsilon$ for some randomised function $g : \mathcal{S}_0 \times \Omega \rightarrow \mathcal{Y}$, then there exists a deterministic function g_* such that $\text{Adv}_{f,g_*}^{\text{sem}}(\mathcal{A}) > \varepsilon$. Indeed, by definition

$$\begin{aligned} \text{Adv}_{f,g}^{\text{sem}}(\mathcal{A}) &= \Pr [\mathcal{G}^{\mathcal{A}} = 1] - \Pr [\mathcal{G}^{\mathcal{A}^*} = 1] \\ &= \sum_{\omega_0 \in \Omega} \Pr [\omega \leftarrow \Omega : \omega = \omega_0] \cdot (\Pr [\mathcal{G}_{\omega}^{\mathcal{A}} = 1] - \Pr [\mathcal{G}_{\omega}^{\mathcal{A}^*} = 1]) \\ &\leq \max_{\omega \in \Omega} \{ \Pr [\mathcal{G}_{\omega}^{\mathcal{A}} = 1] - \Pr [\mathcal{G}_{\omega}^{\mathcal{A}^*} = 1] \} \end{aligned}$$

where the game \mathcal{G}_{ω} is defined as follows

$$\mathcal{G}_{\omega}^{\mathcal{A}} \left[\begin{array}{l} s \leftarrow \mathcal{S}_0 \\ x \leftarrow f(s) \\ y \leftarrow \mathcal{A}(x) \\ \mathbf{return} [y \stackrel{?}{=} g(s; \omega)] \end{array} \right]$$

Since $g_{\omega}(\cdot) = g(\cdot; \omega)$ is a deterministic function, we have obtained

$$\text{Adv}_{f,g}^{\text{sem}}(\mathcal{A}) \leq \max_{\omega \in \Omega} \text{Adv}_{f,g_{\omega}}^{\text{sem}}(\mathcal{A})$$

and without loss of generality we can assume that g is deterministic. This kind of coin-fixing argument is common in many cryptographic proofs.

SAMPLING IDIOM. As a next step, we will split the domain of \mathcal{S}_0 into a set of disjoint sub-domains that cover the entire domain

$$\mathcal{S}_y \doteq \{s \in \mathcal{S}_0 : g(s) = y\} \ .$$

Since g is deterministic the partition is indeed a well defined. Now a distribution \mathcal{S}_0 naturally defines a distribution over indices y . Let \mathcal{Y}_0 be the distribution over elements of \mathcal{Y} such that for all $y_0 \in \mathcal{Y}$

$$\Pr [y \leftarrow \mathcal{Y}_0 : y = y_0] = \Pr [s \leftarrow \mathcal{S}_0 : s \in \mathcal{S}_{y_0}] \ .$$

Secondly, let \mathcal{S}_y denote also a conditional distribution that emerges if we restrict the set of outputs to the domain \mathcal{S}_y , that is, for all $s_0 \in \mathcal{S}_0$ and $y \in \mathcal{Y}$:

$$\Pr [s \leftarrow \mathcal{S}_y : s = s_0] = \Pr [s \leftarrow \mathcal{S}_0 : s = s_0 | g(s) = y] \ .$$

Then the sampling procedure $s \leftarrow \mathcal{S}_0$ can be rewritten in two steps

$$\mathcal{G}_*^{\mathcal{A}} \left[\begin{array}{l} y \leftarrow \mathcal{Y}_0 \\ s \leftarrow \mathcal{S}_y \\ x \leftarrow f(s) \\ y \leftarrow \mathcal{A}(x) \\ \mathbf{return} [y \stackrel{?}{=} g(s)] \end{array} \right]$$

so that the success probability does not change for any adversary \mathcal{A} . Indeed

$$\begin{aligned} \Pr [\mathcal{G}^{\mathcal{A}} = 1] &= \sum_{y \in \mathcal{Y}} \Pr [s \leftarrow \mathcal{S}_0 : g(s) = y_0] \Pr [\mathcal{G}^{\mathcal{A}} = 1 | g(s) = y_0] \\ &= \sum_{y \in \mathcal{Y}} \Pr [y \leftarrow \mathcal{Y}_0 : y = y_0] \Pr [\mathcal{G}_{y_0}^{\mathcal{A}} = 1] \end{aligned}$$

where the game \mathcal{G}_{y_0} is defined as follows

$$\mathcal{G}_{y_0}^{\mathcal{A}} \left[\begin{array}{l} s \leftarrow \mathcal{S}_{y_0} \\ x \leftarrow f(s) \\ y \leftarrow \mathcal{A}(x) \\ \mathbf{return} [y \stackrel{?}{=} g(s)] \end{array} \right]$$

To be punctual, one has to use the total probability formula and the definition of conditional probabilities to formally prove

$$\Pr [\mathcal{G}^{\mathcal{A}} = 1 | g(s) = y_0] \equiv \Pr [\mathcal{G}_{y_0}^{\mathcal{A}} = 1]$$

but this is a trivial exercise that is left to the reader.

CHOOSING BETWEEN SEVERAL SIMPLE HYPOTHESIS. Already a superficial inspection of the game \mathcal{G}_* reveals that an adversary \mathcal{A} must choose between multiple simple hypotheses $\mathcal{H}_{y_0} = [y \stackrel{?}{=} y_0]$. As a next step, we can express

$$\begin{aligned} \Pr [\mathcal{G}_*^{\mathcal{A}} = 1] &= \Pr [y \leftarrow \mathcal{Y}_0 : y = y_1] \Pr [\mathcal{G}_{y_1}^{\mathcal{A}} = 1] \\ &\quad + \sum_{y_0 \in \mathcal{Y} \setminus y_1} \Pr [y \leftarrow \mathcal{Y}_0 : y = y_0] \Pr [\mathcal{G}_{y_0}^{\mathcal{A}} = 1] \end{aligned}$$

where y_1 is the most probable index element. Now note that

$$\begin{aligned} \Pr [y_1] \Pr [\mathcal{G}_{y_1}^{\mathcal{A}} = 1] &= \Pr [y_1] - \Pr [y_1] \cdot \sum_{y_0 \in \mathcal{Y} \setminus y_1} \Pr [s \leftarrow \mathcal{S}_{y_1} : \mathcal{A}(f(s)) = y_0] \\ &\leq \Pr [y_1] - \sum_{y_0 \in \mathcal{Y} \setminus y_1} \Pr [y_0] \Pr [s \leftarrow \mathcal{S}_{y_1} : \mathcal{A}(f(s)) = y_0] \end{aligned}$$

since $\Pr [y_0] \leq \Pr [y_1]$ by the choice of y_1 . Consequently,

$$\begin{aligned} \Pr [\mathcal{G}_*^{\mathcal{A}} = 1] &\leq \Pr [y_1] + \sum_{y_0 \in \mathcal{Y} \setminus y_1} \Pr [y_0] \left| \Pr \left[\begin{array}{c} s \leftarrow \mathcal{S}_{y_0} : \\ \mathcal{A}(\cdot) = y_0 \end{array} \right] - \Pr \left[\begin{array}{c} s \leftarrow \mathcal{S}_{y_1} : \\ \mathcal{A}(\cdot) = y_0 \end{array} \right] \right| \\ &\leq \Pr [y_1] + \max_{y_0 \in \mathcal{Y}} \left| \Pr \left[\begin{array}{c} s \leftarrow \mathcal{S}_{y_0} : \\ \mathcal{A}(\cdot) = y_0 \end{array} \right] - \Pr \left[\begin{array}{c} s \leftarrow \mathcal{S}_{y_1} : \\ \mathcal{A}(\cdot) = y_0 \end{array} \right] \right|. \end{aligned}$$

FROM COMPLEX HYPOTHESES TO SIMPLE HYPOTHESES. Note that in terms of \mathcal{S}_0 the last term in the upper bound obtained above is very close to a computational distance between complex hypotheses $[s \leftarrow \mathcal{S}_{y_0}]$ and $[s \leftarrow \mathcal{S}_{y_1}]$.

In fact, if we know the maximising value y_0 , then we can convert a t -time algorithm \mathcal{A} that maximises the term into $2t$ -time distinguisher $\mathcal{B} : \mathcal{X} \rightarrow \{0, 1\}$. As y_0 can be at most t -bits long, we can test $[\mathcal{A}(f(s)) \stackrel{?}{=} y_0]$ in t -time and consequently we can build $2t$ -time algorithm \mathcal{B} such that

$$\begin{aligned} p_0 &\doteq \Pr [s \leftarrow \mathcal{S}_{y_0} : \mathcal{B}(f(s)) = 1] = \Pr [s \leftarrow \mathcal{S}_{y_0} : \mathcal{A}(f(s)) = y_0] \quad , \\ p_1 &\doteq \Pr [s \leftarrow \mathcal{S}_{y_1} : \mathcal{B}(f(s)) = 1] = \Pr [s \leftarrow \mathcal{S}_{y_1} : \mathcal{A}(f(s)) = y_0] \quad . \end{aligned}$$

Since

$$\begin{aligned} p_0 &= \sum_{\substack{s_0 \in \mathcal{S}_{y_0} \\ s_1 \in \mathcal{S}_{y_1}}} \Pr [s \leftarrow \mathcal{S}_{y_0} : s = s_0] \Pr [s \leftarrow \mathcal{S}_{y_1} : s = s_1] \Pr [\mathcal{B}(f(s_0)) = 1] \\ p_1 &= \sum_{\substack{s_0 \in \mathcal{S}_{y_0} \\ s_1 \in \mathcal{S}_{y_1}}} \Pr [s \leftarrow \mathcal{S}_{y_0} : s = s_0] \Pr [s \leftarrow \mathcal{S}_{y_1} : s = s_1] \Pr [\mathcal{B}(f(s_1)) = 1] \end{aligned}$$

we can estimate

$$\begin{aligned}
|p_0 - p_1| &= \left| \sum_{\substack{s_0 \in \mathcal{S}_{y_0} \\ s_1 \in \mathcal{S}_{y_1}}} \Pr[s_0] \Pr[s_1] (\Pr[\mathcal{B}(f(s_0)) = 1] - \Pr[\mathcal{B}(f(s_1)) = 1]) \right| \\
&\leq \sum_{\substack{s_0 \in \mathcal{S}_{y_0} \\ s_1 \in \mathcal{S}_{y_1}}} \Pr[s_0] \Pr[s_1] |\Pr[\mathcal{B}(f(s_0)) = 1] - \Pr[\mathcal{B}(f(s_1)) = 1]| \\
&\leq \max_{\substack{s_0 \in \mathcal{S}_{y_0} \\ s_1 \in \mathcal{S}_{y_1}}} |\Pr[\mathcal{B}(f(s_0)) = 1] - \Pr[\mathcal{B}(f(s_1)) = 1]| \\
&\leq \max_{s_0, s_1 \in \mathcal{S}_0} \text{cd}_x^{2t}([s \stackrel{?}{=} s_0], [s \stackrel{?}{=} s_1]) \leq \varepsilon .
\end{aligned}$$

THE FINAL STEP. To summarise, we have obtained

$$\begin{aligned}
\Pr[\mathcal{G}^{\mathcal{A}} = 1] &= \Pr[\mathcal{G}_*^{\mathcal{A}} = 1] \leq \max_{y_0 \in \mathcal{Y}} \Pr[y \leftarrow \mathcal{Y}_0 : y = y_0] + \varepsilon \\
&\leq \max_{y_0 \in \mathcal{Y}} \Pr[s \leftarrow \mathcal{S}_0 : g(s) = y_0] + \varepsilon .
\end{aligned}$$

□

3 Final Remarks

Note that the proof given above is strictly non-constructive and does not show how to convert a good predictor of g -values into a good distinguisher of hidden states $s_0, s_1 \in \mathcal{S}_0$. Non-constructivity of the proof is simultaneously the main strength and weakness of this approach. Theorem 1 is extremely powerful as a mathematical claim, since it assumes nothing from the sample distribution \mathcal{S}_0 and holds for all functions g . As a result, non-constructivity is essential in the proof, since we cannot assume that elements of \mathcal{S}_0 can be *efficiently* sampled neither we can assume that the function g is *efficiently* computable.

To be precise, non-constructivity enters into the proof in three places. First, we fix random coins ω so that $g(\cdot, \omega)$ behaves better than $g(\cdot)$. Second, we fix a most probable output y_1 and sub-distributions \mathcal{S}_y for $y \in \mathcal{Y}$. Third, we use two proof steps for finding an output $y_0 \in \mathcal{Y}$ and two states $s_0 \in \mathcal{S}_{y_0}, s_1 \in \mathcal{S}_{y_1}$ that maximises the difference $|\Pr[\mathcal{A}(f(s_0)) = y_0] - \Pr[\mathcal{A}(f(s_1)) = y_0]|$.

For a fixed distribution \mathcal{S}_0 and a fixed function $g(\cdot)$, the complexity of these non-constructive steps is irrelevant. For a hypothetical t -time algorithm \mathcal{A} that achieves $\text{Adv}_{f,g}^{\text{sem}}(\mathcal{A}) > \varepsilon$, there exist states $s_0, s_1 \in \mathcal{S}_0$ and a corresponding $2t$ -time algorithm \mathcal{B} such that $\text{Adv}_{s_0, s_1}^{\text{ind}}(\mathcal{B}) \geq \varepsilon$ and consequently the existence of such a t -time algorithm \mathcal{A} directly contradicts the security premise.

However, the efficiency is important if the distribution \mathcal{S}_0 is not fixed beforehand and a contradiction triple (\mathcal{B}, s_0, s_1) must be discovered on the fly. For instance, Theorem 1 does not cover semantic security of ciphertexts in the settings, where an adversary can influence which messages are encrypted. The latter

is not so far-fetched assumption. For instance, communication in a war is mostly about the adversarial behaviour and thus clearly controllable by adversary.

To prove security in such settings, we have to weaken the claim so that the construction in the proof would become efficiently constructable. In particular, the distribution \mathcal{S}_0 must be efficiently samplable and g efficiently computable. Then we can convert the original non-constructive proof into a constructive reduction. The details of this approach are thoroughly discussed in [BDJR97].

References

- [BDJR97] Mihir Bellare, Anand Desai, E. Jorjani, and Phillip Rogaway. A concrete security treatment of symmetric encryption. Available from <http://www-cse.ucsd.edu/users/mihir/papers/sym-enc.html>, 1997. Full version of the FOCS publication.