

Unifying computational entropies via Kullback–Leibler divergence

Rohit Agrawal*, Yi-Hsiu Chen**, Thibaut Horel***, and Salil Vadhan†

John A. Paulson School of Engineering and Applied Sciences, Harvard University,
Cambridge, MA 02138, USA
{rohitagr,yhchen,thorel}@seas.harvard.edu
salil_vadhan@harvard.edu

Abstract. We introduce *hardness in relative entropy*, a new notion of hardness for search problems which on the one hand is satisfied by all one-way functions and on the other hand implies both *next-block pseudoentropy* and *inaccessible entropy*, two forms of computational entropy used in recent constructions of pseudorandom generators and statistically hiding commitment schemes, respectively. Thus, hardness in relative entropy unifies the latter two notions of computational entropy and sheds light on the apparent “duality” between them. Additionally, it yields a more modular and illuminating proof that one-way functions imply next-block inaccessible entropy, similar in structure to the proof that one-way functions imply next-block pseudoentropy (Vadhan and Zheng, STOC ‘12).

Keywords: one-way function · pseudorandom generator · pseudoentropy · computational entropy · inaccessible entropy · statistically hiding commitment · next-bit pseudoentropy

1 Introduction

1.1 One-way functions and computational entropy

One-way functions [3] are on one hand the minimal assumption for complexity-based cryptography [15], but on the other hand can be used to construct a remarkable array of cryptographic primitives, including such powerful objects as CCA-secure symmetric encryption, zero-knowledge proofs and statistical zero-knowledge arguments for all of **NP**, and secure multiparty computation

* Supported by the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program.

** Supported by NSF grant CCF-1763299.

*** Supported in part by the National Science Foundation under grants CAREER IIS-1149662, CNS-1237235 and CCF-1763299, by the Office of Naval Research under grants YIP N00014-14-1-0485 and N00014-17-1-2131, and by a Google Research Award.

† Supported by NSF grant CCF-1763299.

with an honest majority [5,7,6,14,20,16,9]. All of these constructions begin by converting the “raw hardness” of a one-way function (OWF) to one of the following more structured cryptographic primitives: a pseudorandom generator (PRG) [1,22], a universal one-way hash function (UOWHF) [18], or a statistically hiding commitment scheme (SHC) [2].

The original constructions of these three primitives from arbitrary one-way functions [14,20,9] were all very complicated and inefficient. Over the past decade, there has been a series of simplifications and efficiency improvements to these constructions [11,12,8,21], leading to a situation where the constructions of two of these primitives — PRGs and SHCs — share a very similar structure and seem “dual” to each other. Specifically, these constructions proceed as follows:

1. Show that every OWF $f : \{0,1\}^n \rightarrow \{0,1\}^n$ has a gap between its “real entropy” and an appropriate form of “computational entropy”. Specifically, for constructing PRGs, it is shown that the function $G(x) = (f(x), x_1, x_2, \dots, x_n)$ has “next-block pseudoentropy” at least $n + \omega(\log n)$ while its real entropy is $H(G(U_n)) = n$ [21] where $H(\cdot)$ denotes Shannon entropy. For constructing SHCs, it is shown that the function $G(x) = (f(x)_1, \dots, f(x)_n, x)$ has “next-block accessible entropy” at most $n - \omega(\log n)$ while its real entropy is again $H(G(U_n)) = n$ [11]. Note that the differences between the two cases are whether we break x or $f(x)$ into individual bits (which matters because the “next-block” notions of computational entropy depend on the block structure) and whether the form of computational entropy is larger or smaller than the real entropy.
2. An “entropy equalization” step that converts G into a similar generator where the real entropy in each block conditioned on the prefix before it is known. This step is exactly the same in both constructions.
3. A “flattening” step that converts the (real and computational) Shannon entropy guarantees of the generator into ones on (smoothed) min-entropy and max-entropy. This step is again exactly the same in both constructions.
4. A “hashing” step where high (real or computational) min-entropy is converted to uniform (pseudo)randomness and low (real or computational) max-entropy is converted to a small-support or disjointness property. For PRGs, this step only requires randomness extractors [14,19], while for SHCs it requires (information-theoretic) interactive hashing [17,4]. (Constructing full-fledged SHCs in this step also utilizes UOWHFs, which can be constructed from one-way functions [20]. Without UOWHFs, we obtain a weaker binding property, which nevertheless suffices for constructing statistical zero-knowledge arguments for all of **NP**.)

This common construction template came about through a back-and-forth exchange of ideas between the two lines of work. Indeed, the uses of computational entropy notions, flattening, and hashing originate with PRGs [14], whereas the ideas of using next-block notions, obtaining them from breaking $(f(x), x)$ into short blocks, and entropy equalization originate with SHCs [11]. All this leads to a feeling that the two constructions, and their underlying computational entropy notions, are “dual” to each other and should be connected at a formal level.

In this paper, we make progress on this project of unifying the notions of computational entropy, by introducing a new computational entropy notion that yields both next-block pseudoentropy and next-block accessible entropy in a clean and modular fashion. It is inspired by the proof of [21] that $(f(x), x_1, \dots, x_n)$ has next-block pseudoentropy $n + \omega(\log n)$, which we will describe now.

1.2 Next-block pseudoentropy via relative pseudoentropy

We recall the definition of next-block pseudoentropy, and the result of [21] relating it to one-wayness.

Definition 1.1 (next-block pseudoentropy, informal). *Let n be a security parameter, and $X = (X_1, \dots, X_m)$ be a random variable distributed on strings of length $\text{poly}(n)$. We say that X has next-block pseudoentropy at least k if there is a random variable $Z = (Z_1, \dots, Z_m)$, jointly distributed with X , such that:*

1. *For all $i = 1, \dots, m$, $(X_1, \dots, X_{i-1}, X_i)$ is computationally indistinguishable from $(X_1, \dots, X_{i-1}, Z_i)$.*
2. $\sum_{i=1}^m \mathbb{H}(Z_i | X_1, \dots, X_{i-1}) \geq k$.

Equivalently, for I uniformly distributed in $[m]$, X_I has conditional pseudoentropy at least k/m given (X_1, \dots, X_{i-1}) .

It was conjectured in [10] that next-block pseudoentropy could be obtained from any OWF by breaking its input into bits, and this conjecture was proven in [21]:

Theorem 1.2 ([21], informal). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a one-way function, let X be uniformly distributed in $\{0, 1\}^n$, and let $X = (X_1, \dots, X_m)$ be a partition of X into blocks of length $O(\log n)$. Then $(f(X), X_1, \dots, X_m)$ has next-block pseudoentropy at least $n + \omega(\log n)$.*

The intuition behind Theorem 1.2 is that since X is hard to sample given $f(X)$, then it should have some extra computational entropy given $f(X)$. This intuition is formalized using the following notion of “relative pseudoentropy,” which is a renaming of [21]’s notion of “KL-hard for sampling,” to better unify the terminology with the notions introduced in this work.

Definition 1.3 (relative pseudoentropy). *Let n be a security parameter, and (X, Y) be a pair of random variables, jointly distributed over strings of length $\text{poly}(n)$. We say that X has relative pseudoentropy at least Δ given Y if for all probabilistic polynomial-time S , we have*

$$\text{KL}(X, Y \parallel S(Y), Y) \geq \Delta,$$

where $\text{KL}(\cdot \parallel \cdot)$ denotes the relative entropy (a.k.a. Kullback–Leibler divergence).¹

¹ Recall that for random variables A and B with $\text{Supp}(A) \subseteq \text{Supp}(B)$, the relative entropy is defined by $\text{KL}(A \parallel B) = \mathbb{E}_{a \leftarrow A} [\log(\Pr[A = a] / \Pr[B = a])]$.

That is, it is hard for any efficient adversary S to sample the conditional distribution of X given Y , even approximately.

The first step of the proof of Theorem 1.2 is to show that one-wayness implies relative pseudoentropy (which can be done with a one-line calculation):

Lemma 1.4. *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a one-way function and let X be uniformly distributed in $\{0, 1\}^n$. Then X has relative pseudoentropy at least $\omega(\log n)$ given $f(X)$.*

Next, we break X into short blocks, and show that the relative pseudoentropy is preserved:

Lemma 1.5. *Let n be a security parameter, let (X, Y) be random variables distributed on strings of length $\text{poly}(n)$, let $X = (X_1, \dots, X_m)$ be a partition of X into blocks, and let I be uniformly distributed in $[m]$. If X has Δ relative pseudoentropy given Y , then X_I has relative pseudoentropy at least Δ/m given (Y, X_1, \dots, X_{I-1}) .*

Finally, the main part of the proof is to show that, once we have short blocks, relative pseudoentropy is *equivalent* to a gap between conditional pseudoentropy and real conditional entropy.

Lemma 1.6. *Let n be a security parameter, Y be a random variable distributed on strings of length $\text{poly}(n)$, and X a random variable distributed on strings of length $O(\log n)$. Then X has relative pseudoentropy at least Δ given Y iff X has conditional pseudoentropy at least $H(X|Y) + \Delta$ given Y .*

Putting these three lemmas together, we see that when f is a one-way function, and we break X into blocks of length $O(\log n)$ to obtain $(f(X), X_1, \dots, X_m)$, on average, the conditional pseudoentropy of X_I given $(f(X), X_1, \dots, X_{I-1})$ is larger than its real conditional entropy by $\omega(\log n)/m$. This tells us that the next-block pseudoentropy of $(f(X), X_1, \dots, X_m)$ is larger than its real entropy by $\omega(\log n)$, as claimed in Theorem 1.2.

We remark that Lemma 1.6 explains why we need to break the input of the one-way function into short blocks: it is false when X is long. Indeed, if f is a one-way function, then we have already seen that X has $\omega(\log n)$ relative pseudoentropy given $f(X)$ (Lemma 1.4), but it does not have conditional pseudoentropy noticeably larger than $H(X|f(X))$ given $f(X)$ (as correct preimages can be efficiently distinguished from incorrect ones using f).

1.3 Inaccessible entropy

As mentioned above, for constructing SHCs from one-way functions, the notion of next-block pseudoentropy is replaced with next-block accessible entropy:

Definition 1.7 (next-block accessible entropy, informal). *Let n be a security parameter, and $Y = (Y_1, \dots, Y_m)$ be a random variable distributed on strings of length $\text{poly}(n)$. We say that Y has next-block accessible entropy at most k if the following holds.*

Let \tilde{G} be any probabilistic $\text{poly}(n)$ -time algorithm that takes a sequence of uniformly random strings $\tilde{R} = (\tilde{R}_1, \dots, \tilde{R}_m)$ and outputs a sequence $\tilde{Y} = (\tilde{Y}_1, \dots, \tilde{Y}_m)$ in an “online fashion” by which we mean that $\tilde{Y}_i = \tilde{G}(\tilde{R}_1, \dots, \tilde{R}_i)$ depends on only the first i random strings of \tilde{G} for $i = 1, \dots, m$. Suppose further that $\text{Supp}(\tilde{Y}) \subseteq \text{Supp}(Y)$.

Then we require:

$$\sum_{i=1}^m H(\tilde{Y}_i | \tilde{R}_1, \dots, \tilde{R}_{i-1}) \leq k.$$

(Next-block) accessible entropy differs from (next-block) pseudoentropy in two ways:

1. Accessible entropy is useful as an *upper* bound on computational entropy, and is interesting when it is *smaller* than the real entropy $H(Y)$. We refer to the gap $H(Y) - k$ as the *next-block inaccessible entropy* of Y .
2. The accessible entropy adversary \tilde{G} is trying to *generate* the random variables Y_i conditioned on the history rather than recognize them. Note that we take the “history” to not only be the previous blocks $(\tilde{Y}_1, \dots, \tilde{Y}_{i-1})$, but the coin tosses $(\tilde{R}_1, \dots, \tilde{R}_{i-1})$ used to generate those blocks.

Note that one unsatisfactory aspect of the definition is that when the random variable Y is not *flat* (i.e. uniform on its support), then there can be an adversary \tilde{G} achieving accessible entropy even *larger* than $H(Y)$, for example by making \tilde{Y} uniform on $\text{Supp}(Y)$.

Similarly to (and predating) Theorem 1.2, it is known that one-wayness implies next-block inaccessible entropy.

Theorem 1.8 ([11]). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a one-way function, let X be uniformly distributed in $\{0, 1\}^n$, and let (Y_1, \dots, Y_m) be a partition of $Y = f(X)$ into blocks of length $O(\log n)$. Then (Y_1, \dots, Y_m, X) has next-block accessible entropy at most $n - \omega(\log n)$.*

Unfortunately, however, the existing proof of Theorem 1.8 is not modular like that of Theorem 1.2. In particular, it does not isolate the step of relating one-wayness to entropy-theoretic measures (like Lemma 1.4 does) or the significance of having short blocks (like Lemma 1.6 does).

1.4 Our results

We remedy the above state of affairs by providing a new, more general notion of hardness in relative entropy that allows us to obtain next-block inaccessible entropy in a modular way while also encompassing what is needed for next-block pseudoentropy.

Like in relative pseudoentropy, we will consider a pair of jointly distributed random variables (Y, X) . Following the spirit of accessible entropy, the adversary \tilde{G} for our new notion will try to *generate* Y together with X , rather than taking

Y as input. That is, $\tilde{\mathbf{G}}$ will take randomness \tilde{R} and output a pair $(\tilde{Y}, \tilde{X}) = \tilde{\mathbf{G}}(\tilde{R}) = (\tilde{\mathbf{G}}_1(\tilde{R}), \tilde{\mathbf{G}}_2(\tilde{R}))$, which we require to be always within the support of (Y, X) . Note that $\tilde{\mathbf{G}}$ need not be an online generator; it can generate both \tilde{Y} and \tilde{X} using the same randomness \tilde{R} . Of course, if (Y, X) is efficiently samplable (as it would be in most cryptographic applications), $\tilde{\mathbf{G}}$ could generate (\tilde{Y}, \tilde{X}) identically distributed to (Y, X) by just using the “honest” sampler \mathbf{G} for (Y, X) . So, in addition, we require that the adversary $\tilde{\mathbf{G}}$ also come with a *simulator* \mathbf{S} , that can simulate its coin tosses given only \tilde{Y} . The goal of the adversary is to minimize the relative entropy

$$\text{KL}(\tilde{R}, \tilde{Y} \parallel \mathbf{S}(Y), Y)$$

for a uniformly random \tilde{R} . This divergence measures both how well $\tilde{\mathbf{G}}_1$ approximates the distribution of Y as well as how well \mathbf{S} simulates the corresponding coin tosses of $\tilde{\mathbf{G}}_1$. Note that when $\tilde{\mathbf{G}}$ is the honest sampler \mathbf{G} , the task of \mathbf{S} is exactly to sample from the conditional distribution of \tilde{R} given $\mathbf{G}(\tilde{R}) = Y$. However, the adversary may reduce the divergence by instead designing the sampler $\tilde{\mathbf{G}}$ and simulator \mathbf{S} to work in concert, potentially trading off how well $\tilde{\mathbf{G}}(\tilde{R})$ approximates Y in exchange for easier simulation by \mathbf{S} . Explicitly, the definition is as follows.

Definition 1.9 (hardness in relative entropy, informal version of Definition 3.4). *Let n be a security parameter, and (Y, X) be a pair of random variables jointly distributed over strings of length $\text{poly}(n)$. We say that (Y, X) is Δ -hard in relative entropy if the following holds.*

Let $\tilde{\mathbf{G}} = (\tilde{\mathbf{G}}_1, \tilde{\mathbf{G}}_2)$ and \mathbf{S} be probabilistic $\text{poly}(n)$ -time algorithms such that $\text{Supp}(\tilde{\mathbf{G}}(\tilde{R})) \subseteq \text{Supp}((Y, X))$, where \tilde{R} is uniformly distributed. Then writing $\tilde{Y} = \tilde{\mathbf{G}}_1(\tilde{R})$, we require that

$$\text{KL}(\tilde{R}, \tilde{Y} \parallel \mathbf{S}(Y), Y) \geq \Delta.$$

Similarly to Lemma 1.4, we can show that one-way functions achieve this notion of hardness in relative entropy.

Lemma 1.10. *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a one-way function and let X be uniformly distributed in $\{0, 1\}^n$. Then $(f(X), X)$ is $\omega(\log n)$ -hard in relative entropy.*

Note that this lemma implies Lemma 1.4. If we take $\tilde{\mathbf{G}}$ to be the “honest” sampler $\tilde{\mathbf{G}}(x) = (f(x), x)$, then we have:

$$\text{KL}(X, f(X) \parallel \mathbf{S}(Y), Y) = \text{KL}(\tilde{R}, \tilde{Y} \parallel \mathbf{S}(Y), Y),$$

which is $\omega(\log n)$ by Lemma 1.10. That is, relative pseudoentropy (as in Definition 1.3 and Lemma 1.4) is obtained by fixing $\tilde{\mathbf{G}}$ and focusing on the hardness for the simulator \mathbf{S} , i.e. the divergence $\text{KL}(X|Y \parallel \mathbf{S}(Y)|Y)$. Furthermore,

the step of breaking into short blocks (Lemma 1.5) is equivalent to requiring the simulator be *online* and showing that relative pseudoentropy implies the following notion of *next-block relative pseudoentropy*:

Definition 1.11 (next-block relative pseudoentropy, informal). *Let n be a security parameter, (X, Y) be jointly distributed random variables over strings of length $\text{poly}(n)$, and let $X = (X_1, \dots, X_m)$ be a partition of X into blocks. We say that X has next-block relative pseudoentropy at least Δ given Y if for all probabilistic polynomial-time S , we have*

$$\sum_{i=1}^m \text{KL}(X_i | X_{<i}, Y \parallel S(X_{<i}, Y) | X_{<i}, Y) \geq \Delta,$$

where we use the notation $z_{<i} = (z_1, \dots, z_{i-1})$.

Here, the simulator S is required to be “online” in the sense that it cannot simulate (X_1, \dots, X_m) at once, but must simulate X_i only as a function of $X_{<i}$ and Y .

In particular, Lemma 1.6 is thus equivalent to the statement that having Δ next-block relative pseudoentropy for blocks of length $O(\log n)$ is equivalent to having next-block pseudoentropy at least $\Delta + \sum_{i=1}^m H(X_i | X_{<i}, Y)$ in the sense of Definition 1.1.

Conversely, we show that inaccessible entropy arises from hardness in relative entropy by first requiring the *generator* G to be online and breaking the relative entropy into blocks to obtain the following next-block hardness property.

Definition 1.12 (next-block hardness in relative entropy, informal). *Let n be a security parameter, and $Y = (Y_1, \dots, Y_m)$ be a random variable distributed on strings of length $\text{poly}(n)$. We say that Y is Δ next-block hard in relative entropy if the following holds.*

Let \tilde{G} be any probabilistic $\text{poly}(n)$ -time algorithm that takes a sequence of uniformly random strings $\tilde{R} = (\tilde{R}_1, \dots, \tilde{R}_m)$ and outputs a sequence $\tilde{Y} = (\tilde{Y}_1, \dots, \tilde{Y}_m)$ in an “online fashion” by which we mean that $\tilde{Y}_i = \tilde{G}(\tilde{R}_1, \dots, \tilde{R}_i)$ depends on only the first i random strings of \tilde{G} for $i = 1, \dots, m$. Suppose further that $\text{Supp}(\tilde{Y}) \subseteq \text{Supp}(Y)$. Additionally, let S be a probabilistic $\text{poly}(n)$ -time algorithm such for all $i = 1, \dots, m$, S takes as input $\hat{R}_1, \dots, \hat{R}_{i-1}$ and Y_i and outputs \hat{R}_i , where \hat{R}_j has the same length as \tilde{R}_j . Then we require that for all such (\tilde{G}, S) , we have:

$$\sum_{i=1}^m \text{KL}(\tilde{R}_i, \tilde{Y}_i | \tilde{R}_{<i}, \tilde{Y}_{<i} \parallel \hat{R}_i, Y_i | \hat{R}_{<i}, Y_{<i}) \geq \Delta.$$

Observe that hardness in relative entropy can be seen as the specific case of next-block hardness in relative entropy when there is only one block (*i.e.*, setting $m = 1$ in the previous definition).

Next, we fix the *simulator*, analogously to how relative pseudoentropy was obtained by fixing the generator, and obtain *next-block inaccessible relative entropy*:

Definition 1.13 (next-block inaccessible relative entropy, informal).

Let n be a security parameter, and $Y = (Y_1, \dots, Y_m)$ be a random variable distributed on strings of length $\text{poly}(n)$. We say that Y has next-block inaccessible relative entropy at least Δ if the following holds.

Let \tilde{G} be any probabilistic $\text{poly}(n)$ -time algorithm that takes a sequence of uniformly random strings $\tilde{R} = (\tilde{R}_1, \dots, \tilde{R}_m)$ and outputs a sequence $\tilde{Y} = (\tilde{Y}_1, \dots, \tilde{Y}_m)$ in an online fashion, and such that $\text{Supp}(\tilde{Y}) \subseteq \text{Supp}(Y)$. Then we require that for all such \tilde{G} , we have:

$$\sum_{i=1}^m \text{KL} \left(\tilde{Y}_i | \tilde{R}_{<i}, \tilde{Y}_{<i} \parallel Y_i | R_{<i}, Y_{<i} \right) \geq \Delta,$$

where $R = (R_1, \dots, R_m)$ is a dummy random variable independent of Y .

That is, the goal of the online generator \tilde{G} is to generate \tilde{Y}_i given the history of coin tosses $\tilde{R}_{<i}$ with the same conditional distribution as Y_i given $Y_{<i}$. As promised, there is no explicit simulator in the definition of next-block inaccessible relative entropy, as we essentially dropped all \tilde{R} variables from the definition of next-block hardness in relative entropy. Nevertheless we can obtain it from hardness in relative entropy by using sufficiently short blocks:

Lemma 1.14. *Let n be a security parameter, let Y be a random variable distributed on strings of length $\text{poly}(n)$, and let $Y = (Y_1, \dots, Y_m)$ be a partition of Y into blocks of length $O(\log n)$.*

If (Y_1, \dots, Y_m) is Δ next-block hard in relative entropy, then (Y_1, \dots, Y_m) has next-block inaccessible relative entropy at least $\Delta - \text{negl}(n)$.

An intuition for the proof is that since the blocks are of logarithmic length, given Y_i we can simulate the corresponding coin tosses of \tilde{R}_i of \tilde{G} by rejection sampling and succeed with high probability in $\text{poly}(n)$ tries.

A nice feature of the definition of next-block inaccessible relative entropy compared to inaccessible entropy is that it is meaningful even for non-flat random variables, as KL divergence is always nonnegative. Moreover, for flat random variables, it equals the inaccessible entropy:

Lemma 1.15. *Suppose $Y = (Y_1, \dots, Y_m)$ is a flat random variable. Then Y has next-block inaccessible relative entropy at least Δ if and only if Y has accessible entropy at most $H(Y) - \Delta$.*

Intuitively, this lemma comes from the identity that if Y is a flat random variable and $\text{Supp}(\tilde{Y}) \subseteq \text{Supp}(Y)$, then $H(\tilde{Y}) = H(Y) - \text{KL}(\tilde{Y} \parallel Y)$. We stress that we do not require the individual blocks Y_i have flat distributions, only that the random variable Y as a whole is flat. For example, if f is a function and X is uniform, then $(f(X), X)$ is flat even though $f(X)$ itself may be far from flat.

Putting together Lemmas 1.10, 1.14, and 1.15, we obtain a new, more modular (and slightly tighter) proof of Theorem 1.8. The reduction implicit in the combination of these lemmas is the same as the one in [11], but the analysis is

different. (In particular, [11] makes no use of KL divergence.) Like the existing proof of Theorem 1.2, this proof separates the move from one-wayness to a form of hardness involving relative entropies, the role of short blocks, and the move from hardness in relative entropy to computational entropy, as summarized in Figure 1. Moreover, this further illumination of and toolkit for notions of computational entropy may open the door to other applications in cryptography.

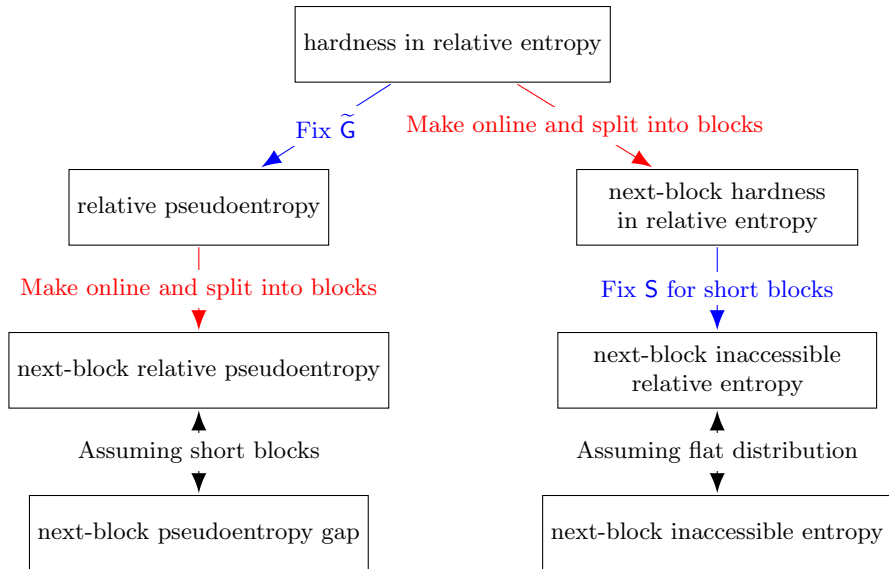


Fig. 1. Relationships between hardness notions.

We remark that another interesting direction for future work is to find a construction of universal one-way hash functions (UOWHFs) from one-way functions that follows a similar template to the above constructions of PRGs and SHCs. There is now a construction of UOWHFs based on a variant of inaccessible entropy [8], but it remains more complex and inefficient than those of PRGs and SHCs.

2 Preliminaries

Notations. For a tuple $x = (x_1, \dots, x_n)$, we write $x_{\leq i}$ for (x_1, \dots, x_i) , and $x_{< i}$ for (x_1, \dots, x_{i-1}) .

poly denotes the set of polynomial functions and negl the set of all negligible functions: $\varepsilon \in \text{negl}$ if for all $p \in \text{poly}$ and large enough $n \in \mathbb{N}$, $\varepsilon(n) \leq 1/p(n)$. We will sometimes abuse notations and write $\text{poly}(n)$ to mean $p(n)$ for some $p \in \text{poly}$ and similarly for $\text{negl}(n)$.

PPT stands for probabilistic polynomial time and can be either in the uniform or non-uniform model of computation. All our results are stated as uniform polynomial time oracle reductions and are thus meaningful in both models.

For a random variable X over \mathcal{X} , $\text{Supp}(X) \stackrel{\text{def}}{=} \{x \in \mathcal{X} : \Pr[X = x] > 0\}$ denotes the support of X . A random variable is *flat* if it is uniform over its support. Random variables will be written with uppercase letters and the associated lowercase letter represents a generic element from its support.

Information theory.

Definition 2.1 (Entropy). For a random variable X and $x \in \text{Supp}(X)$, the sample entropy (also called surprise) of x is $H_x^*(X) \stackrel{\text{def}}{=} \log(1/\Pr[X = x])$. The entropy $H(X)$ of X is the expected sample entropy: $H(X) \stackrel{\text{def}}{=} \mathbb{E}_{x \leftarrow X} [H_x^*(X)]$.

Definition 2.2 (Conditional entropy). Let (A, X) be a pair of random variables and consider $(a, x) \in \text{Supp}(A, X)$, the conditional sample entropy of (a, x) is $H_{a,x}^*(A|X) \stackrel{\text{def}}{=} \log(1/\Pr[A = a | X = x])$ and the conditional entropy of A given X is the expected conditional sample entropy:

$$H(A|X) \stackrel{\text{def}}{=} \mathbb{E}_{(a,x) \leftarrow (A,X)} \left[\log \frac{1}{\Pr[A = a | X = x]} \right].$$

Proposition 2.3 (Chain rule for entropy). Let (A, X) be a pair of random variables, then $H(A, X) = H(A|X) + H(X)$ and for $(a, x) \in \text{Supp}(A, X)$, $H_{a,x}^*(A, X) = H_{a,x}^*(A|X) + H_x^*(X)$.

Definition 2.4 (Relative entropy²). For a pair (A, B) of random variables and $(a, b) \in \text{Supp}(A, B)$ the sample relative entropy (log-probability ratio) is:

$$\text{KL}_a^*(A \parallel B) \stackrel{\text{def}}{=} \log \frac{\Pr[A = a]}{\Pr[B = a]},$$

and the relative entropy of A with respect to B is the expected sample relative entropy:

$$\text{KL}(A \parallel B) \stackrel{\text{def}}{=} \mathbb{E}_{a \leftarrow A} \left[\log \frac{\Pr[A = a]}{\Pr[B = a]} \right].$$

Definition 2.5 (Conditional relative entropy). For pairs of random variables (A, X) and (B, Y) , and $(a, x) \in \text{Supp}(A, X)$, the conditional sample relative entropy is:

$$\text{KL}_{a,x}^*(A|X \parallel B|Y) \stackrel{\text{def}}{=} \log \frac{\Pr[A = a | X = x]}{\Pr[B = a | Y = x]},$$

and the conditional relative entropy is:

$$\text{KL}(A|X \parallel B|Y) \stackrel{\text{def}}{=} \mathbb{E}_{(a,x) \leftarrow (A,X)} \left[\log \frac{\Pr[A = a | X = x]}{\Pr[B = a | Y = x]} \right].$$

² Relative entropy is also commonly referred to as *Kullback–Liebler divergence*, which explains the standard KL notation. We prefer to use relative entropy to have more uniformity across the notions discussed in this work.

Proposition 2.6 (Chain rule for relative entropy). For pairs of random variables (X, A) and (Y, B) :

$$\text{KL}(A, X \parallel B, Y) = \text{KL}(A|X \parallel B|Y) + \text{KL}(X \parallel Y) ,$$

and for $(a, x) \in \text{Supp}(A, X)$:

$$\text{KL}_{a,x}^*(A, X \parallel B, Y) = \text{KL}_{a,x}^*(A|X \parallel B|Y) + \text{KL}_x^*(X \parallel Y) .$$

Proposition 2.7 (Data-processing inequality). Let (X, Y) be a pair of random variables and let f be a function defined on $\text{Supp}(Y)$, then:

$$\text{KL}(X \parallel Y) \geq \text{KL}(f(X) \parallel f(Y)) .$$

Definition 2.8 (min relative entropy). Let (X, Y) be a pair of random variables and $\delta \in [0, 1]$. We define $\text{KL}_{\min}^\delta(X \parallel Y)$ to be the quantile of level δ of $\text{KL}_x^*(X \parallel Y)$, equivalently it is the smallest $\Delta \in \mathbb{R}$ satisfying:

$$\Pr_{x \leftarrow X} [\text{KL}_x^*(X \parallel Y) \leq \Delta] \geq \delta ,$$

and it is characterized by the following equivalence:

$$\text{KL}_{\min}^\delta(X \parallel Y) > \Delta \iff \Pr_{x \leftarrow X} [\text{KL}_x^*(X \parallel Y) \leq \Delta] < \delta .$$

Block generators

Definition 2.9 (Block generator). An m -block generator is a function $\mathbf{G} : \{0, 1\}^s \rightarrow \prod_{i=1}^m \{0, 1\}^{\ell_i}$. $\mathbf{G}_i(r)$ denotes the i -th block of \mathbf{G} on input r and $|\mathbf{G}_i| = \ell_i$ denotes the bit length of the i -th block.

Definition 2.10 (Online generator). An online m -block generator is a function $\tilde{\mathbf{G}} : \prod_{i=1}^m \{0, 1\}^{s_i} \rightarrow \prod_{i=1}^m \{0, 1\}^{\ell_i}$ such that for all $i \in [m]$ and $r \in \prod_{i=1}^m \{0, 1\}^{s_i}$, $\tilde{\mathbf{G}}_i(r)$ only depends on $r_{\leq i}$. We sometimes write $\tilde{\mathbf{G}}_i(r_{\leq i})$ when the input blocks $i + 1, \dots, m$ are unspecified.

Definition 2.11 (Support). The support of a generator \mathbf{G} is the support of the random variable $\text{Supp}(\mathbf{G}(R))$ for uniform input R . If \mathbf{G} is an $(m + 1)$ -block generator, and Π is a binary relation, we say that \mathbf{G} is supported on Π if $\text{Supp}(\mathbf{G}_{\leq m}(R), \mathbf{G}_{m+1}(R)) \subseteq \Pi$.

When \mathbf{G} is an $(m + 1)$ -block generator supported on a binary relation Π , we will often use the notation $\mathbf{G}_w \stackrel{\text{def}}{=} \mathbf{G}_{m+1}$ to emphasize that the last block corresponds to a witness for the first m blocks.

Cryptography.

Definition 2.12 (One-way Function). Let n be a security parameter, $t = t(n)$ and $\varepsilon = \varepsilon(n)$. A function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a (t, ε) -one-way function if:

1. For all time t randomized algorithm \mathbf{A} : $\Pr_{x \leftarrow U_n} [\mathbf{A}(f(x)) \in f^{-1}(f(x))] \leq \varepsilon$, where U_n is uniform over $\{0, 1\}^n$.
2. There exists a polynomial time algorithm \mathbf{B} such that $\mathbf{B}(x, 1^n) = f(x)$ for all $x \in \{0, 1\}^n$.

If f is $(n^c, 1/n^c)$ -one-way for every $c \in \mathbb{N}$, we say that f is (strongly) one-way.

3 Search Problems and Hardness in Relative Entropy

In this section, we first present the classical notion of hard-on-average search problems and introduce the new notion of hardness in relative entropy. We then relate the two notions by proving that average-case hardness implies hardness in relative entropy.

3.1 Search problems

For a binary relation $\Pi \subseteq \{0, 1\}^* \times \{0, 1\}^*$, we write $\Pi(y, w)$ for the predicate that is true iff $(y, w) \in \Pi$ and say that w is a *witness* for the *instance* y ³. To each relation Π , we naturally associate (1) a *search problem*: given y , find w such that $\Pi(y, w)$ or state that no such w exist and (2) the *decision problem* defined by the language $L_\Pi \stackrel{\text{def}}{=} \{y \in \{0, 1\}^* : \exists w \in \{0, 1\}^*, \Pi(y, w)\}$. **FNP** denotes the set of all relations Π computable by a polynomial time algorithm and such that there exists a polynomial p such that $\Pi(y, w) \Rightarrow |w| \leq p(|y|)$. Whenever $\Pi \in \mathbf{FNP}$, the associated decision problem L_Π is in **NP**. We now define average-case hardness.

Definition 3.1 (Ddistributional search problem). A distributional search problem is a pair (Π, Y) where $\Pi \subseteq \{0, 1\}^* \times \{0, 1\}^*$ is a binary relation and Y is a random variable supported on L_Π .

The problem (Π, Y) is (t, ε) -hard if $\Pr[\Pi(Y, A(Y))] \leq \varepsilon$ for all time t randomized algorithm A , where the probability is over the distribution of Y and the randomness of A .

Example 3.2. For $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$, the problem of inverting f is the search problem associated with the relation $\Pi^f \stackrel{\text{def}}{=} \{(f(x), x) : x \in \{0, 1\}^n\}$. If f is a (t, ε) -one-way function, then the distributional search problem $(\Pi^f, f(X))$ of inverting f on a uniform random input $X \in \{0, 1\}^n$ is (t, ε) -hard.

Remark 3.3. Consider a distributional search problem (Π, Y) . Without loss of generality, there exists a (possibly inefficient) two-block generator $G = (G_1, G_w)$ supported on Π such that $G_1(R) = Y$ for uniform input R . If G_w is polynomial-time computable, it is easy to see that the search problem $(\Pi^{G_1}, G_1(R))$ is at least as hard as (Π, Y) . The advantage of writing the problem in this “functional” form is that the distribution $(G_1(R), R)$ over (instance, witness) pairs is flat, which is a necessary condition to relate hardness to inaccessible entropy (see Theorem 4.10).

Furthermore, if G_1 is also polynomial-time computable and (Π, Y) is $(\text{poly}(n), \text{negl}(n))$ -hard, then $R \mapsto G_1(R)$ is a one-way function. Combined with the previous example, we see that the existence of one-way functions is equivalent to the existence of $(\text{poly}(n), \text{negl}(n))$ -hard search problems for which (instance, witness) pairs can be efficiently sampled.

³ We used the unconventional notation y for the instance (instead of x) because our relations will often be of the form Π^f for some function f ; in this case an instance is some y in the range of f and a witness for y is any preimage $x \in f^{-1}(y)$.

3.2 Hardness in relative entropy

Instead of considering an adversary directly attempting to solve a search problem (Π, Y) , the adversary in the definition of hardness in relative entropy comprises a pair of algorithm (\tilde{G}, S) where \tilde{G} is a two-block generator outputting valid (instance, witness) pairs for Π and S is a *simulator* for \tilde{G} : given an instance y , the goal of S is to output randomness r for \tilde{G} such that $\tilde{G}_1(r) = y$. Formally, the definition is as follows.

Definition 3.4 (hardness in relative entropy). *Let (Π, Y) be a distributional search problem. We say that (Π, Y) is (t, Δ) -hard in relative entropy if:*

$$\text{KL} \left(\tilde{R}, \tilde{G}_1(\tilde{R}) \parallel S(Y), Y \right) > \Delta,$$

for all pairs (\tilde{G}, S) of time t algorithms where \tilde{G} is a two-block generator supported on Π and \tilde{R} is uniform randomness for \tilde{G}_1 . Similarly, for $\delta \in [0, 1]$, (Π, Y) is (t, Δ) -hard in δ -min relative entropy if for all such pairs:

$$\text{KL}_{\min}^{\delta} \left(\tilde{R}, \tilde{G}_1(\tilde{R}) \parallel S(Y), Y \right) > \Delta.$$

Note that a pair (\tilde{G}, S) achieves a relative entropy of zero in Definition 3.4 if $\tilde{G}_1(\tilde{R})$ has the same distribution as Y and if $\tilde{G}_1(S(y)) = y$ for all $y \in \text{Supp}(Y)$. In this case, writing $\tilde{G}_w \stackrel{\text{def}}{=} \tilde{G}_2$, we have that $\tilde{G}_w(S(Y))$ is a valid witness for Y since \tilde{G} is supported on Π .

More generally, the composition $\tilde{G}_w \circ S$ solves the search problem (Π, Y) whenever $\tilde{G}_1(S(Y)) = Y$. When the relative entropies in Definition 3.4 are upper-bounded, we can lower bound the probability of the search problem being solved (Lemma 3.7) This immediately implies that hard search problems are also hard in relative entropy.

Theorem 3.5. *Let (Π, Y) be a distributional search problem. If (Π, Y) is (t, ε) -hard, then it is (t', Δ') -hard in relative entropy and (t', Δ'') -hard in δ -min relative entropy for every $\delta \in [0, 1]$ where $t' = \Omega(t)$,⁴ $\Delta' = \log(1/\varepsilon)$ and $\Delta'' = \log(1/\varepsilon) - \log(1/\delta)$.*

Remark 3.6. As we see, a “good” simulator S for a generator \tilde{G} is one for which $\tilde{G}_1(S(Y)) = Y$ holds often. It will be useful in Section 4 to consider simulators S which are allowed to fail by outputting a failure string $r \notin \text{Supp}(\tilde{R})$, (e.g. $r = \perp$) and adopt the convention that $\tilde{G}_1(r) = \perp$ whenever $r \notin \text{Supp}(\tilde{R})$. With this convention, we can without loss of generality add the requirement that $\tilde{G}_1(S(Y)) = Y$ whenever $S(Y) \in \text{Supp}(\tilde{R})$: indeed, S can always check that it

⁴ For the theorems in this paper that relate two notions of hardness, the notation $t' = \Omega(t)$ means that there exists a constant C depending *only* on the computational model such that $t' \geq C \cdot t$.

is the case and if not output a failure symbol. For such a simulator S , observe that for all $r \in \text{Supp}(\tilde{R})$, the second variable on both sides of the relative entropy in Definition 3.4 is obtained by applying \tilde{G}_1 on the first variable and can thus be dropped, leading to a simpler definition of hardness in relative entropy: $\text{KL}(\tilde{R} \parallel S(Y)) > \Delta$.

Theorem 3.5 is an immediate consequence of the following lemma.

Lemma 3.7. *Let (Π, Y) be a distributional search problem and (\tilde{G}, S) be a pair of algorithms with $\tilde{G} = (\tilde{G}_1, \tilde{G}_w)$ a two-block generator supported on Π . Define the linear-time oracle algorithm $A^{\tilde{G}_w, S}(y) \stackrel{\text{def}}{=} \tilde{G}_w(S(y))$. For $\Delta \in \mathbb{R}^+$ and $\delta \in [0, 1]$:*

1. *If $\text{KL}(\tilde{R}, \tilde{G}_1(\tilde{R}) \parallel S(Y), Y) \leq \Delta$ then $\Pr[\Pi(Y, A^{\tilde{G}_w, S}(Y))] \geq 1/2^\Delta$.*
2. *If $\text{KL}_{\min}^\delta(\tilde{R}, \tilde{G}_1(\tilde{R}) \parallel S(Y), Y) \leq \Delta$ then $\Pr[\Pi(Y, A^{\tilde{G}_w, S}(Y))] \geq \delta/2^\Delta$.*

Proof. We have:

$$\begin{aligned}
\Pr[\Pi(Y, A^{\tilde{G}_w, S}(Y))] &= \Pr[\Pi(Y, \tilde{G}_w(S(Y)))] \\
&\geq \Pr[\tilde{G}_1(S(Y)) = Y] \quad (\tilde{G} \text{ is supported on } \Pi) \\
&= \sum_{r \in \text{Supp}(\tilde{R})} \Pr[S(Y) = r \wedge Y = \tilde{G}_1(r)] \\
&= \mathbb{E}_{r \leftarrow \tilde{R}} \left[\frac{\Pr[S(Y) = r \wedge Y = \tilde{G}_1(r)]}{\Pr[\tilde{R} = r]} \right] \\
&= \mathbb{E}_{\substack{r \leftarrow \tilde{R} \\ y \leftarrow \tilde{G}_1(r)}} \left[2^{-\text{KL}_{r,y}^*(\tilde{R}, \tilde{G}_1(\tilde{R}) \parallel S(Y), Y)} \right].
\end{aligned}$$

Now, the first claim follows by Jensen's inequality (since $x \mapsto 2^{-x}$ is convex) and the second claim follows by Markov's inequality when considering the event that the sample relative entropy is smaller than Δ (which occurs with probability at least δ by assumption).

Relation to relative pseudoentropy. In [21], the authors introduced the notion of relative pseudoentropy⁵: for jointly distributed variables (Y, W) , W has relative pseudoentropy given Y if it is hard for a polynomial time adversary to approximate—measured in relative entropy—the conditional distribution W given Y . Formally:

⁵ As already mentioned in the introduction, this notion was in fact called “KL-hardness for sampling” in [21] but we rename it here to unify the terminology between the various notions discussed here.

Definition 3.8 (Relative pseudoentropy, Def. 3.4 in [21]). Let (Y, W) be a pair of random variables, we say that W has relative pseudoentropy (t, Δ) given Y if for all time t randomized algorithm S , we have:

$$\text{KL}(Y, W \parallel Y, S(Y)) > \Delta.$$

As discussed in Section 1.2, it was shown in [21] that if $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is a one-way function, then $(f(X), X_1, \dots, X_n)$ has next-bit pseudoentropy for uniform $X \in \{0, 1\}^n$ (see Theorem 1.2). The first step in proving this result was to prove that X has relative pseudoentropy given $f(X)$ (see Lemma 1.4).

We observe that when (Y, W) is of the form $(f(X), X)$ for some function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and variable X over $\{0, 1\}^n$, then relative pseudoentropy is implied by hardness in relative entropy by simply fixing \tilde{G} to be the “honest sampler” $\tilde{G}(X) = (f(X), X)$. Indeed, in this case we have:

$$\text{KL}(X, \tilde{G}_1(X) \parallel S(Y), Y) = \text{KL}(X, f(X) \parallel S(Y), Y).$$

We can thus recover Lemma 1.4 as a direct corollary of Theorem 3.5.

Corollary 3.9. Consider a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ and define $\Pi^f \stackrel{\text{def}}{=} \{(f(x), x) : x \in \{0, 1\}^n\}$ and $Y \stackrel{\text{def}}{=} f(X)$ for X uniform over $\{0, 1\}^n$. If f is (t, ε) -one-way, then (Π^f, Y) is $(t', \log(1/\varepsilon))$ -hard in relative entropy and X has $(t', \log(1/\varepsilon))$ relative pseudoentropy given Y with $t' = \Omega(t)$.

Witness hardness in relative entropy. We also introduce a relaxed notion of hardness in relative entropy called witness hardness in relative entropy. In this notion, we further require (\tilde{G}, S) to approximate the joint distribution of (instance, witness) pairs rather than only instances. For example, the problem of inverting a function f over a random input X is naturally associated with the distribution $(f(X), X)$. The relaxation in this case is analogous to the notion of *distributional one-way function* for which the adversary is required to approximate the uniform distribution over preimages.

Definition 3.10 (Witness hardness in relative entropy). Let Π be a binary relation and (Y, W) be a pair of random variables supported on Π . We say that (Π, Y, W) is (t, Δ) witness hard in relative entropy if for all pairs of time t algorithms (\tilde{G}, S) where \tilde{G} is a two-block generator supported on Π , for uniform \tilde{R} :

$$\text{KL}(\tilde{R}, \tilde{G}_1(\tilde{R}), \tilde{G}_w(\tilde{R}) \parallel S(Y), Y, W) > \Delta.$$

Similarly, for $\delta \in [0, 1]$, (Π, Y, W) is (t, Δ) -witness hard in δ -min relative entropy, if for all such pairs:

$$\text{KL}_{\min}^{\delta}(\tilde{R}, \tilde{G}_1(\tilde{R}), \tilde{G}_w(\tilde{R}) \parallel S(Y), Y, W) > \Delta.$$

We introduced hardness in relative entropy first, since it is the notion which is most directly obtained from the hardness of distribution search problems. Observe that by the data processing inequality for relative entropy (Proposition 2.7), dropping the third variable on both sides of the relative entropies in Definition 3.10 only decreases them. Hence, hardness in relative entropy implies witness hardness as stated in (Theorem 3.11). As we will see in Section 4 witness hardness in relative entropy is the “correct” notion to obtain inaccessible entropy from: it is in fact equal to inaccessible entropy up to 1/poly losses.

Theorem 3.11. *Let Π be a binary relation and (Y, W) be a pair of random variables supported on Π . If (Π, Y) is (t, ε) -hard, then (Π, Y, W) is (t', Δ') witness hard in relative entropy and (t', Δ'') witness hard in δ -min relative entropy for every $\delta \in [0, 1]$ where $t' = \Omega(t)$, $\Delta' = \log(1/\varepsilon)$ and $\Delta'' = \log(1/\varepsilon) - \log(1/\delta)$.*

Remark 3.12. The data processing inequality does not hold exactly for KL_{\min} , hence the statement about δ -min relative entropy in Theorem 3.11 does not follow with the claimed parameters in a black-box manner from Theorem 3.5. However, an essentially identical proof given in Appendix A yields the result.

4 Inaccessible Entropy and Hardness in Relative Entropy

In this section, we relate our notion of witness hardness in relative entropy to the inaccessible entropy definition of [13]. Roughly speaking, we “split” the relative entropy into blocks and obtain the intermediate notion of next-block inaccessible relative entropy (Section 4.1) that we then relate to inaccessible entropy (Section 4.2). Together, these results show that if f is a one-way function, the generator $G^f(X) = (f(X)_1, \dots, f(X)_n, X)$ has superlogarithmic inaccessible entropy.

4.1 Next-block hardness and rejection sampling

Consider a binary relation Π and a pair of random variables (Y, W) supported on Π . Let \tilde{G} be an online $(m + 1)$ -block generator supported on Π and write $\tilde{Y}_{\leq m} \stackrel{\text{def}}{=} \tilde{G}(\tilde{R}_{\leq m})$ for uniform $\tilde{R}_{\leq m}$. For such a generator \tilde{G} , it is natural to consider simulators operating in an online manner. Specifically, an online simulator in this context is a PPT algorithm S such that on input $(\tilde{R}_{< i}, Y_i)$, S outputs \tilde{R}_i of the same length as \tilde{R}_i . The goal of S is to output random coins such that $(\tilde{R}_i, \tilde{Y}_i)$ is “close” to (\tilde{R}_i, Y_i) conditioned on the past. This leads to the following natural blockwise notion of hardness in relative entropy for online generators and simulators.

Definition 4.1 (Next-block hardness in relative entropy). *The joint distribution (Y_1, \dots, Y_m) is (t, Δ) next-block hard in relative entropy if the following holds.*

Let \tilde{G} be any time t online m -block generator supported on $Y_{\leq m}$ and write $\tilde{Y}_{\leq m} \stackrel{\text{def}}{=} \tilde{G}(\tilde{R}_{\leq m})$ for uniform $\tilde{R}_{\leq m}$. Let S be an online simulator and define

inductively $\widehat{R}_i \stackrel{\text{def}}{=} S(\widehat{R}_{<i}, Y_i)$ with \widehat{R}_i having the same length as \widetilde{R}_i , where S is a probabilistic algorithm that uses time at most t to compute $\widehat{R}_{\leq m}$.

Then we require:

$$\sum_{i=1}^m \text{KL} \left(\widetilde{R}_i, \widetilde{Y}_i | \widetilde{R}_{<i}, \widetilde{Y}_{<i} \parallel \widehat{R}_i, Y_i | \widehat{R}_{<i}, Y_{<i} \right) > \Delta.$$

Similarly, for $\delta \in [0, 1]$, we say that (Y_1, \dots, Y_m) is (t, Δ) -next-block hard in δ -min relative entropy if, with the same notations as above:

$$\Pr_{\substack{r_{\leq m} \leftarrow \widetilde{R}_{\leq m} \\ y_{\leq m} \leftarrow \widetilde{G}(r_{\leq m})}} \left[\sum_{i=1}^m \text{KL}_{y_i, r_{<i}, y_{<i}}^* \left(\widetilde{R}_i, \widetilde{Y}_i | \widetilde{R}_{<i}, \widetilde{Y}_{<i} \parallel \widehat{R}_i, Y_i | \widehat{R}_{<i}, Y_{<i} \right) \leq \Delta \right] < \delta.$$

Observe that using the chain rule for relative entropy, the sum of relative entropies appearing in Definition 4.1 is exactly equal to the relative entropies appearing in Definition 3.4. Since, furthermore considering an online generator \widetilde{G} and online simulator S is only less general than arbitrary pairs (\widetilde{G}, S) , we immediately obtain the following theorem.

Theorem 4.2. *Let (Π, Y) be a distributional search problem. If (Π, Y) is (t, Δ) -hard in relative entropy then (Y_1, \dots, Y_m) is (t, Δ) next-block hard in relative entropy.*

Similarly, for any $\delta \in [0, 1]$, if (Π, Y) is (t, Δ) -hard in δ -min relative entropy then (Y_1, \dots, Y_m) is (t, Δ) next-block hard in δ -min relative entropy.

Proof. Immediate using the chain rule for relative (sample) entropy.

The next step is to obtain a notion of hardness that makes no reference to simulators by considering, for an online block generator \widetilde{G} , a specific simulator $\text{Sim}^{\widetilde{G}, T}$ which on input $(\widehat{R}_{<i}, Y_i)$, generates \widehat{R}_i using rejection sampling until $\widetilde{G}_i(\widehat{R}_{\leq i}) = Y_i$. The superscript T is the maximum number of attempts after which $\text{Sim}^{\widetilde{G}, T}$ gives up and outputs \perp . The formal definition of $\text{Sim}^{\widetilde{G}, T}$ is given in Algorithm 1.

For the rejection sampling simulator $\text{Sim}^{\widetilde{G}, T}$, we will show in Lemma 4.5 that the next-block hardness in relative entropy in Definition 4.1 decomposes as the sum of two terms:

1. A term measuring how well $\widetilde{G}_{\leq m}$ approximates the distribution Y in an online manner, without any reference to a simulator.
2. An error term measuring the failure probability of the rejection sampling procedure due to having a finite time bound T .

As we show in Lemma 4.6, the error term can be made arbitrarily small by setting the number of trials T in $\text{Sim}^{\widetilde{G}, T}$ to be a large enough multiple of $m \cdot 2^\ell$ where ℓ is the length of the blocks of $\widetilde{G}_{\leq m}$. This leads to a $\text{poly}(m)$ time algorithm

Algorithm 1 Rejection sampling simulator $\text{Sim}^{\tilde{\mathbf{G}}, T}$ for $1 \leq i \leq m$

Input: $y_i \in \{0, 1\}^*$, $\hat{r}_{<i} \in (\{0, 1\}^v \cup \{\perp\})^{i-1}$

Output: $\hat{r}_i \in \{0, 1\}^v \cup \{\perp\}$

```

if  $\hat{r}_{i-1} = \perp$  then
     $\hat{r}_i \leftarrow \perp$ ; return
end if
repeat
    sample  $\hat{r}_i \leftarrow \{0, 1\}^v$ 
until  $\tilde{\mathbf{G}}_i(\hat{r}_{\leq i}) = y_i$  or  $\geq T$  attempts
if  $\tilde{\mathbf{G}}_i(\hat{r}_{\leq i}) \neq y_i$  then
     $\hat{r}_i \leftarrow \perp$ 
end if

```

whenever ℓ is logarithmic in m . That is, given an online block generator $\tilde{\mathbf{G}}$ for which $\tilde{\mathbf{G}}_{\leq m}$ has short blocks, we obtain a corresponding simulator “for free”. Thus, considering only the first term leads to the following clean definition of next-block inaccessible relative entropy that makes no reference to simulators.

Definition 4.3 (Next-block inaccessible relative entropy). *The joint distribution (Y_1, \dots, Y_m) has (t, Δ) next-block inaccessible relative entropy, if for every time t online m -block generator $\tilde{\mathbf{G}}$ supported on $Y_{\leq m}$, writing $\tilde{Y}_{\leq m} \stackrel{\text{def}}{=} \tilde{\mathbf{G}}(\tilde{R}_{\leq m})$ for uniform $\tilde{R}_{\leq m}$, we have:*

$$\sum_{i=1}^m \text{KL} \left(\tilde{Y}_i | \tilde{R}_{<i}, \tilde{Y}_{<i} \parallel Y_i | R_{<i}, Y_{<i} \right) > \Delta,$$

where R_i is a “dummy” random variable over the domain of $\tilde{\mathbf{G}}_i$ and independent of $Y_{\leq m+1}$. Similarly, for $\delta \in [0, 1]$, we say that (Y_1, \dots, Y_{m+1}) has (t, Δ) -next-block inaccessible δ -min relative entropy if for every $\tilde{\mathbf{G}}$ as above:

$$\Pr_{\substack{r_{\leq m} \leftarrow \tilde{R}_{\leq m} \\ y_{\leq m} \leftarrow \tilde{\mathbf{G}}(r_{\leq m})}} \left[\sum_{i=1}^m \text{KL}_{y_i, r_{<i}, y_{<i}}^* \left(\tilde{Y}_i | \tilde{R}_{<i}, \tilde{Y}_{<i} \parallel Y_i | R_{<i}, Y_{<i} \right) \leq \Delta \right] < \delta,$$

where $(\tilde{Y}_{\leq m}, \tilde{R}_{\leq m})$ are defined as above.

Remark 4.4. Since $\tilde{Y}_{<i}$ is a function of $\tilde{R}_{<i}$, the first conditional distribution in the KL is effectively $\tilde{Y}_i | \tilde{R}_{<i}$. Similarly the second distribution is effectively $Y_i | Y_{<i}$. The extra random variables are there for syntactic consistency.

With this definition in hand, we can make formal the claim that, even as sample notions, the next-block hardness in relative entropy decomposes as next-block inaccessible relative entropy plus an error term.

Lemma 4.5. *For a joint distribution (Y_1, \dots, Y_m) , let $\tilde{\mathbf{G}}$ be an online m -block generator supported on $Y_{\leq m}$. Define $(\tilde{Y}_1, \dots, \tilde{Y}_m) \stackrel{\text{def}}{=} \tilde{\mathbf{G}}(\tilde{R})$ for uniform random*

variable $\tilde{R} = (\tilde{R}_1, \dots, \tilde{R}_m)$ and let R_i be a “dummy” random variable over the domain of \tilde{G}_i and independent of $Y_{\leq m+1}$. We also define $\hat{R}_i \stackrel{\text{def}}{=} \text{Sim}^{\tilde{G}, T}(\hat{R}_{< i}, Y_i)$ and $\hat{Y}_i = \tilde{G}(\hat{R}_{< i})$. Then, for all $r \in \text{Supp}(\tilde{R})$ and $y \stackrel{\text{def}}{=} \tilde{G}(r)$:

$$\begin{aligned} & \sum_{i=1}^m \text{KL}_{r,y}^* \left(\tilde{R}_i, \tilde{Y}_i | \tilde{R}_{< i}, \tilde{Y}_{< i} \parallel \hat{R}_i, Y_i | \hat{R}_{< i}, Y_{< i} \right) \\ &= \sum_{i=1}^m \text{KL}_{r,y}^* \left(\tilde{Y}_i | \tilde{R}_{< i}, \tilde{Y}_{< i} \parallel Y_i | R_{< i}, Y_{< i} \right) \\ & \quad + \sum_{i=1}^m \log \left(\frac{1}{\Pr \left[\hat{Y}_i = y_i | Y_i = y_i, \hat{R}_{< i} = r_{< i} \right]} \right). \end{aligned}$$

Moreover, the running time of $\text{Sim}^{\tilde{G}, T}$ on input $\hat{R}_{< i}, Y_i$ is $O(|r_i| \cdot T)$, with at most T oracle calls to \tilde{G} .

Proof. Consider $r \in \text{Supp}(\tilde{R})$ and $y \stackrel{\text{def}}{=} \tilde{G}(r)$. Then:

$$\begin{aligned} & \sum_{i=1}^m \text{KL}_{r,y}^* \left(\tilde{R}_i, \tilde{Y}_i | \tilde{R}_{< i}, \tilde{Y}_{< i} \parallel \hat{R}_i, Y_i | \hat{R}_{< i}, Y_{< i} \right) \\ &= \sum_{i=1}^m \text{KL}_{r,y}^* \left(\tilde{R}_i, \tilde{Y}_i | \tilde{R}_{< i}, \tilde{Y}_{< i} \parallel \hat{R}_i, \hat{Y}_i | \hat{R}_{< i}, \hat{Y}_{< i} \right) \\ &= \sum_{i=1}^m \left(\text{KL}_{r,y}^* \left(\tilde{R}_i | \tilde{R}_{< i}, \tilde{Y}_{\leq i} \parallel \hat{R}_i | \hat{R}_{< i}, \hat{Y}_{\leq i} \right) + \text{KL}_{r,y}^* \left(\tilde{Y}_i | \tilde{R}_{< i}, \tilde{Y}_{< i} \parallel \hat{Y}_i | \hat{R}_{< i}, \hat{Y}_{< i} \right) \right) \\ &= \sum_{i=1}^m \text{KL}_{r,y}^* \left(\tilde{Y}_i | \tilde{R}_{< i}, \tilde{Y}_{< i} \parallel \hat{Y}_i | \hat{R}_{< i}, \hat{Y}_{< i} \right) \\ &= \sum_{i=1}^m \text{KL}_{r,y}^* \left(\tilde{Y}_i | \tilde{R}_{< i} \parallel \hat{Y}_i | \hat{R}_{< i} \right), \end{aligned}$$

The first equality is because $Y_i = \hat{Y}_i$ since we are only considering non-failure cases ($r_i \neq \perp$). The second equality is the chain rule. The penultimate equality is by definition of rejection sampling: $\tilde{R}_i | \tilde{R}_{< i}, \tilde{Y}_{\leq i}$ and $\hat{R}_i | \hat{R}_{< i}, \hat{Y}_{\leq i}$ are identical on $\text{Supp}(\tilde{R}_i)$ since conditioning on $\hat{Y}_i = y$ implies that only non-failure cases ($r_i \neq \perp$) are considered. The last equality is because $\tilde{Y}_{< i}$ (resp. $\hat{Y}_{< i}$) is a deterministic function of $\tilde{R}_{< i}$ (resp. $\hat{R}_{< i}$).

We now relate $\widehat{Y}_i|\widehat{R}_{<i}$ to $Y_i|Y_{<i}$:

$$\begin{aligned}
& \Pr \left[\widehat{Y}_i = y_i | \widehat{R}_{<i} = r_{<i} \right] \\
&= \Pr \left[\widehat{Y}_i = y_i, Y_i = y_i | \widehat{R}_{<i} = r_{<i} \right] \quad (\widehat{Y}_i = y_i \Leftrightarrow \widehat{Y}_i = y_i \wedge Y_i = y_i) \\
&= \Pr \left[\widehat{Y}_i = y_i | Y_i = y_i, \widehat{R}_{<i} = r_{<i} \right] \cdot \Pr \left[Y_i = y_i | \widehat{R}_{<i} = r_{<i} \right] \quad (\text{Bayes' Rule}) \\
&= \Pr \left[\widehat{Y}_i = y_i | Y_i = y_i, \widehat{R}_{<i} = r_{<i} \right] \cdot \Pr [Y_i = y_i | Y_{<i} = y_{<i}],
\end{aligned}$$

where the last equality is because when $r \in \text{Supp}(\widetilde{R})$, $\widehat{R}_{<i} = r_{<i} \Rightarrow Y_{<i} = y_{<i}$ and because Y_i is independent of $\widehat{R}_{<i}$ given $Y_{<i}$ (as $\widehat{R}_{<i}$ is simply a randomized function of $Y_{<i}$). The conclusion of the lemma follows by combining the previous two derivations.

Observe that taking expectations with respect to a uniform \widetilde{R} on both sides in the conclusion of Lemma 4.5, we get that next-block hardness in relative entropy is equal to the sum of next-block inaccessible relative entropy and the expectation of the error term coming from the rejection sampling procedure. The following lemma upper bounds this expectation.

Lemma 4.6. *Let $\widetilde{\mathcal{G}}$ be an online m -block generator, and let $L_i \stackrel{\text{def}}{=} 2^{|\widetilde{\mathcal{G}}_i|}$ be the size of the codomain of $\widetilde{\mathcal{G}}_i$, $i \in [m]$. Then for all $i \in [m]$, $r_{<i} \in \text{Supp}(\widetilde{R}_{<i})$ and uniform \widetilde{R}_i :*

$$\mathbb{E}_{y_i \leftarrow \widetilde{\mathcal{G}}_i(r_{<i}, \widetilde{R}_i)} \left[\log \frac{1}{\Pr \left[\widehat{Y}_i = y_i | Y_i = y_i, \widehat{R}_{<i} = r_{<i} \right]} \right] \leq \log \left(1 + \frac{L_i - 1}{T} \right).$$

Proof (Proof of Lemma 4.6). By definition of $\text{Sim}^{\widetilde{\mathcal{G}}, T}$, we have:

$$\Pr \left[\widehat{Y}_i = y_i | Y_i = y_i, \widehat{R}_{<i} = r_{<i} \right] = 1 - \left(1 - \Pr \left[\widetilde{\mathcal{G}}_i(r_{<i}, \widetilde{R}_i) = y_i \right] \right)^T.$$

Applying Jensen's inequality, we have:

$$\begin{aligned}
& \mathbb{E}_{y_i \leftarrow \widetilde{\mathcal{G}}_i(r_{<i}, \widetilde{R}_i)} \left[\log \left(\frac{1}{\Pr \left[\widehat{Y}_i = y_i | Y_i = y_i, \widehat{R}_{<i} = r_{<i} \right]} \right) \right] \\
& \leq \log \mathbb{E}_{y_i \leftarrow \widetilde{\mathcal{G}}_i(r_{<i}, \widetilde{R}_i)} \left[\frac{1}{\Pr \left[\widehat{Y}_i = y_i | Y_i = y_i, \widehat{R}_{<i} = r_{<i} \right]} \right] \\
& = \log \left(\sum_{y \in \text{Im}(\widetilde{\mathcal{G}}_i(r_{<i}, \cdot))} \frac{p_y}{1 - (1 - p_y)^T} \right)
\end{aligned}$$

where $p_y = \Pr \left[\tilde{\mathbf{G}}_i(r_{<i}, \tilde{R}_i) = y \right]$. Since the function $x / (1 - (1 - x)^T)$ is convex (see Lemma A.1 in the appendix), the maximum of the expression inside the logarithm over probability distributions $\{p_y\}$ is achieved at the extremal points of the standard probability simplex. Namely, when all but one $p_y \rightarrow 0$ and the other one is 1. Since $\lim_{x \rightarrow 0} x / (1 - (1 - x)^T) = 1/T$:

$$\log \left(\sum_{y \in \text{Im}(\tilde{\mathbf{G}}_i)} \frac{p_y}{1 - (1 - p_y)^T} \right) \leq \log \left(1 + (L_i - 1) \cdot \frac{1}{T} \right).$$

By combining Lemmas 4.5 and 4.6, we are now ready to state the main result of this section, relating witness hardness in relative entropy to next-block inaccessible relative entropy.

Theorem 4.7. *Let Π be a binary relation and let (Y, W) be a pair of random variables supported on Π . Let $Y = (Y_1, \dots, Y_m)$ where the bit length of Y_i is at most ℓ . Then we have:*

1. *if (Π, Y, W) is (t, Δ) witness hard in relative entropy, then for every $0 < \Delta' \leq \Delta$, (Y_1, \dots, Y_m, W) has $(t', \Delta - \Delta')$ next-block inaccessible relative entropy where $t' = \Omega(t\Delta' / (m^2 2^\ell))$.*
2. *if (Π, Y, W) is (t, Δ) witness hard in δ -min relative entropy then for every $0 < \Delta' \leq \Delta$ and $0 \leq \delta' \leq 1 - \delta$, we have that (Y_1, \dots, Y_m, W) has $(t', \Delta - \Delta')$ next-block inaccessible $(\delta + \delta')$ -min relative entropy where $t' = \Omega(t\delta' \Delta' / (m^2 2^\ell))$.*

Proof. We consider an online generator $\tilde{\mathbf{G}}$ supported on (Y_1, \dots, Y_m, W) and the simulator $\text{Sim}^{\tilde{\mathbf{G}}, T}$. For convenience, we sometimes write Y_{m+1} for W . Define $\tilde{R} \stackrel{\text{def}}{=} \tilde{R}_{\leq m}$ where $\tilde{R}_{\leq m}$ is a sequence of independent and uniformly random variables, $\tilde{Y}_{\leq m+1} \stackrel{\text{def}}{=} \tilde{\mathbf{G}}(\tilde{R})$, $\tilde{\mathbf{G}}_1(\tilde{R}) \stackrel{\text{def}}{=} Y_{\leq m}$ and $\tilde{\mathbf{G}}_w(\tilde{R}) \stackrel{\text{def}}{=} \tilde{Y}_{m+1}$. We also write for $1 \leq i \leq m$, $\hat{R}_i \stackrel{\text{def}}{=} \text{Sim}^{\tilde{\mathbf{G}}, T}(\hat{R}_{<i}, Y_i)$, $\hat{Y}_i \stackrel{\text{def}}{=} \tilde{\mathbf{G}}(\hat{R}_{\leq i})_i$. Finally we define $\mathcal{S}^{\tilde{\mathbf{G}}, T}(Y) \stackrel{\text{def}}{=} \hat{R}_{\leq m}$.

Observe that $(\tilde{\mathbf{G}}_1, \tilde{\mathbf{G}}_w)$ is a two-block generator supported on Π , so the pair $(\tilde{\mathbf{G}}, \mathcal{S}^{\tilde{\mathbf{G}}, T})$ forms a pair of algorithms as in the definition of witness hardness in relative entropy (Definition 3.10). We focus on sample notions first, and consider $r \in \text{Supp}(\tilde{R})$, $y \in \text{Supp}(\tilde{Y}_{\leq m})$ and $w \in \text{Supp}(\tilde{Y}_{m+1})$. First we use the chain rule to isolate the witness block:

$$\begin{aligned} & \text{KL}_{r,y,w}^* \left(\tilde{R}, \tilde{\mathbf{G}}_1(\tilde{R}), \tilde{\mathbf{G}}_w(\tilde{R}) \parallel \mathcal{S}^{\tilde{\mathbf{G}}, T}(Y), Y, W \right) \\ &= \text{KL}_{r,y,w}^* \left(\tilde{\mathbf{G}}_w(\tilde{R}) \mid \tilde{R}, \tilde{\mathbf{G}}_1(\tilde{R}) \parallel W \mid \mathcal{S}^{\tilde{\mathbf{G}}, T}(Y), Y \right) \\ & \quad + \text{KL}_{r,y,w}^* \left(\tilde{R}, \tilde{\mathbf{G}}_1(\tilde{R}) \parallel \mathcal{S}^{\tilde{\mathbf{G}}, T}(Y), Y \right) \\ &= \text{KL}_{r,y,w}^* \left(\tilde{Y}_{m+1} \mid \tilde{R}_{\leq m}, \tilde{Y}_{\leq m} \parallel Y_{m+1} \mid R_{\leq m}, Y_{\leq m} \right) \\ & \quad + \text{KL}_{r,y,w}^* \left(\tilde{R}, \tilde{\mathbf{G}}_1(\tilde{R}) \parallel \mathcal{S}^{\tilde{\mathbf{G}}, T}(Y), Y \right). \end{aligned}$$

Next, as in Theorem 4.2 we apply the chain rule to decompose the second term on the right-hand side and obtain next-block hardness in relative entropy:

$$\begin{aligned} & \text{KL}_{r,y,w}^* \left(\tilde{R}, \tilde{G}_1(\tilde{R}) \parallel \mathcal{S}^{\tilde{G},T}(Y), Y \right) \\ &= \sum_{i=1}^m \text{KL}_{r,y,w}^* \left(\tilde{R}_i, \tilde{Y}_i | \tilde{R}_{<i}, \tilde{Y}_{<i} \parallel \hat{R}_i, Y_i | \hat{R}_{<i}, Y_{<i} \right). \end{aligned}$$

Finally, we use Lemma 4.5 to further decompose the right-hand side term into inaccessible relative entropy and the rejection sampling error:

$$\begin{aligned} & \sum_{i=1}^m \text{KL}_{r,y,w}^* \left(\tilde{R}_i, \tilde{Y}_i | \tilde{R}_{<i}, \tilde{Y}_{<i} \parallel \hat{R}_i, Y_i | \hat{R}_{<i}, Y_{<i} \right) \\ &= \sum_{i=1}^m \text{KL}_{r,y}^* \left(\tilde{Y}_i | \tilde{R}_{<i}, \tilde{Y}_{<i} \parallel Y_i | R_{<i}, Y_{<i} \right) \\ &+ \sum_{i=1}^m \log \left(\frac{1}{\Pr \left[\hat{Y}_i = y_i | Y_i = y_i, \hat{R}_{<i} = r_{<i} \right]} \right). \end{aligned}$$

Combining the previous derivations, we obtain:

$$\begin{aligned} & \sum_{i=1}^{m+1} \text{KL}_{r,y}^* \left(\tilde{Y}_i | \tilde{R}_{<i}, \tilde{Y}_{<i} \parallel Y_i | R_{<i}, Y_{<i} \right) \\ &= \text{KL}_{r,y,w}^* \left(\tilde{R}, \tilde{G}_1(\tilde{R}), \tilde{G}_w(\tilde{R}) \parallel \mathcal{S}^{\tilde{G},T}(Y), Y, W \right) \\ &- \sum_{i=1}^m \log \left(\frac{1}{\Pr \left[\hat{Y}_i = y_i | Y_i = y_i, \hat{R}_{<i} = r_{<i} \right]} \right). \end{aligned}$$

Now, the first claim of the theorem follows by taking expectations on both sides and observing that when $T = m \cdot 2^\ell / (\delta' \ln 2)$, Lemma 4.6 implies that the expected value of the rejection sampling error is smaller than δ' .

For the second claim, we first establish using Lemma 4.6 and Markov's inequality that:

$$\Pr_{\substack{y_{\leq m+1} \leftarrow \tilde{Y}_{\leq m+1} \\ r \leftarrow \tilde{R}}} \left[\sum_{i=1}^m \log \left(\frac{1}{\Pr \left[\hat{Y}_i = y_i | \hat{R}_{<i} = r_{<i}, \hat{Y}_{<i} = y_{<i} \right]} \right) \geq \frac{m \cdot 2^\ell}{T \delta' \ln 2} \right] \leq \delta'$$

and we reach a similar conclusion by setting $T = m \cdot 2^\ell / (\delta' \Delta' \ln 2)$.

Remark 4.8. For fixed distribution and generators, in the limit where T grows to infinity, the error term caused by the failure of rejection sampling in time T vanishes. In this case, hardness in relative entropy implies next-block inaccessible relative entropy without any loss in the hardness parameters.

4.2 Next-block inaccessible relative entropy and inaccessible entropy

We first recall the definition from [13], slightly adapted to our notations.

Definition 4.9 (Inaccessible Entropy). *Let (Y_1, \dots, Y_{m+1}) be a joint distribution.⁶ We say that (Y_1, \dots, Y_{m+1}) has t -inaccessible entropy Δ if for all $(m+1)$ -block online generators $\tilde{\mathbf{G}}$ running in time t and consistent with (Y_1, \dots, Y_{m+1}) :*

$$\sum_{i=1}^{m+1} \left(\mathbb{H}(Y_i | Y_{<i}) - \mathbb{H}(\tilde{Y}_i | \tilde{R}_{<i}) \right) > \Delta .$$

where $(\tilde{Y}_1, \dots, \tilde{Y}_{m+1}) = \tilde{\mathbf{G}}(\tilde{R}_1, \dots, \tilde{R}_{m+1})$ for a uniform $\tilde{R}_{\leq m+1}$. We say that (Y_1, \dots, Y_{m+1}) has (t, δ) -max-inaccessible entropy Δ if for all $(m+1)$ -block online generators $\tilde{\mathbf{G}}$ running in time t and consistent with (Y_1, \dots, Y_{m+1}) :

$$\Pr_{\substack{r_{\leq m+1} \leftarrow \tilde{R}_{\leq m+1} \\ y_{\leq m+1} \leftarrow \tilde{\mathbf{G}}(r_{\leq m+1})}} \left[\sum_{i=1}^{m+1} \left(\mathbb{H}_{y_i, y_{<i}}^* (Y_i | Y_{<i}) - \mathbb{H}_{y_i, r_{<i}}^* (\tilde{Y}_i | \tilde{R}_{<i}) \right) \leq \Delta \right] < \delta .$$

Unfortunately, one unsatisfactory aspect of Definition 4.9 is that inaccessible entropy can be negative since the generator $\tilde{\mathbf{G}}$ could have more entropy than (Y_1, \dots, Y_{m+1}) : if all the Y_i are independent biased random bits, then a generator $\tilde{\mathbf{G}}$ outputting unbiased random bits will have negative inaccessible entropy. On the other hand, next-block inaccessible relative entropy (Definition 4.3) does not suffer from this drawback.

Moreover, in the specific case where (Y_1, \dots, Y_{m+1}) is a flat distribution⁷, then no distribution with the same support can have higher entropy and in this case Definitions 4.3 and 4.9 coincide as stated in the following theorem.

Theorem 4.10. *Let (Y_1, \dots, Y_{m+1}) be a flat distribution and $\tilde{\mathbf{G}}$ be an $(m+1)$ -block generator consistent with $Y_{\leq m+1}$. Then for $\tilde{Y}_{\leq m+1} = \tilde{\mathbf{G}}(\tilde{R}_{\leq m+1})$ for uniform $\tilde{R}_{\leq m+1}$:*

1. For every $y_{\leq m+1}, r_{\leq m+1} \in \text{Supp}(\tilde{Y}_{\leq m+1}, \tilde{R}_{\leq m+1})$, it holds that

$$\begin{aligned} \sum_{i=1}^{m+1} \left(\mathbb{H}_{y_i, y_{<i}}^* (Y_i | Y_{<i}) - \mathbb{H}_{y_i, r_{<i}}^* (\tilde{Y}_i | \tilde{R}_{<i}) \right) \\ = \sum_{i=1}^{m+1} \text{KL}_{y_i, y_{<i}, r_{<i}}^* \left(\tilde{Y}_i | \tilde{R}_{<i}, \tilde{Y}_{<i} \parallel Y_i | R_{<i}, Y_{<i} \right) \end{aligned}$$

⁶ We write $m+1$ the total number of blocks, since in this section we will think of Y_{m+1} (also written as W) as the witness of distributional search problem and (Y_1, \dots, Y_m) are the blocks of the instance as in the previous section.

⁷ For example, the distribution $(Y_{\leq m}, Y_{m+1}) = (f(U), U)$ for a function f and uniform input U is always a flat distribution even if f itself is not regular.

In particular, (Y_1, \dots, Y_{m+1}) has (t, Δ) next-block inaccessible δ -min relative entropy if and only if it has (t, δ) -max-inaccessible entropy at least Δ .

2. Furthermore,

$$\sum_{i=1}^{m+1} \left(\mathbb{H}(Y_i|Y_{<i}) - \mathbb{H}(\tilde{Y}_i|\tilde{R}_{<i}) \right) = \sum_{i=1}^{m+1} \text{KL}(\tilde{Y}_i|\tilde{R}_{<i}, \tilde{Y}_{<i} \parallel Y_i|R_{<i}, Y_{<i}),$$

so in particular, (Y_1, \dots, Y_{m+1}) has (t, Δ) next-block inaccessible relative entropy if and only if it has t -inaccessible entropy at least Δ .

Proof. For the sample notions, the chain rule (Proposition 2.6) gives:

$$\sum_{i=1}^{m+1} \mathbb{H}_{y_i, y_{<i}}^*(Y_i|Y_{<i}) = \mathbb{H}_y^*(Y_{\leq m+1}) = \log |\text{Supp}(Y_{\leq m+1})|$$

for all y since Y is flat. Hence:

$$\begin{aligned} \log |\text{Supp}(Y_{\leq m+1})| &- \sum_{i=1}^{m+1} \mathbb{H}_{y_i, y_{<i}}^*(\tilde{Y}_i|\tilde{R}_{<i}) \\ &= \sum_{i=1}^{m+1} \left(\mathbb{H}_{y_i, y_{<i}}^*(Y_i|Y_{<i}) - \mathbb{H}_{y_i, r_{<i}}^*(\tilde{Y}_i|\tilde{R}_{<i}) \right) \\ &= \sum_{i=1}^{m+1} \text{KL}_{y_i, y_{<i}, r_{<i}}^*(\tilde{Y}_i|\tilde{R}_{<i}, \tilde{Y}_{<i} \parallel Y_i|R_{<i}, Y_{<i}), \end{aligned}$$

so the second claim follows by taking the expectation over $(\tilde{Y}_{\leq m+1}, \tilde{R}_{\leq m+1})$ on both sides.

By chaining the reductions between the different notions of hardness considered in this work (hardness in relative entropy, next-block inaccessible relative entropy and inaccessible entropy), we obtain a more modular proof of the theorem of Haitner *et al.* [13], obtaining inaccessible entropy from any one-way function.

Theorem 4.11. *Let n be a security parameter, $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a (t, ε) -one-way function, and X be uniform over $\{0, 1\}^n$. For $\ell \in \{1, \dots, n\}$, decompose $f(X) \stackrel{\text{def}}{=} (Y_1, \dots, Y_{n/\ell})$ into blocks of length ℓ . Then:*

1. *For every $0 \leq \Delta \leq \log(1/\varepsilon)$, $(Y_1, \dots, Y_{n/\ell}, X)$ has t' -inaccessible entropy at least $(\log(1/\varepsilon) - \Delta)$ for $t' = \Omega(t \cdot \Delta \cdot \ell^2 / (n^2 \cdot 2^\ell))$.*
2. *For every $0 < \delta \leq 1$ and $0 \leq \Delta \leq \log(1/\varepsilon) - \log(2/\delta)$, $(Y_1, \dots, Y_{n/\ell}, X)$ has (t', δ) -max-inaccessible entropy at least $(\log(1/\varepsilon) - \log(2/\delta) - \Delta)$ for $t' = \Omega(t \cdot \delta \cdot \Delta \cdot \ell^2 / (n^2 \cdot 2^\ell))$.*

Proof. Since f is (t, ε) -one-way, the distributional search problem $(\Pi^f, f(X))$ where $\Pi^f = \{(f(x), x) : x \in \{0, 1\}^n\}$ is (t, ε) -hard. Clearly, $(f(X), X)$ is supported on Π^f , so by applying Theorem 3.11, we have that $(\Pi^f, f(X), X)$ is

$(\Omega(t), \log(1/\varepsilon))$ witness hard in relative entropy and $(\Omega(t), \log(1/\varepsilon) - \log(2/\delta))$ witness hard in $\delta/2$ -min relative entropy. Thus, by Theorem 4.7 we have that $(Y_1, \dots, Y_{n/\ell}, X)$ has $(\Omega(t \cdot \Delta \cdot \ell^2 / (n^2 \cdot 2^\ell)), \log(1/\varepsilon) - \Delta)$ next-block inaccessible relative entropy and $(\Omega(t \cdot \delta \cdot \Delta \cdot \ell^2 / (n^2 \cdot 2^\ell)), \log(1/\varepsilon) - \log(2/\delta) - \Delta)$ next-block inaccessible δ -min relative entropy, and we conclude by Theorem 4.10.

Remark 4.12. For comparison, the original proof of [13] shows that for every $0 < \delta \leq 1$, $(Y_1, \dots, Y_{n/\ell}, X)$ has (t', δ) -max-inaccessible entropy at least $(\log(1/\varepsilon) - 2\log(1/\delta) - O(1))$ for $t' = \tilde{\Omega}(t \cdot \delta \cdot \ell^2 / (n^2 \cdot 2^\ell))$, which in particular for fixed t' has quadratically worse dependence on δ in terms of the achieved inaccessible entropy: $\log(1/\varepsilon) - 2 \cdot \log(1/\delta) - O(1)$ rather than our $\log(1/\varepsilon) - 1 \cdot \log(1/\delta) - O(1)$.

Corollary 4.13 (Theorem 4.2 in [13]). *Let n be a security parameter, $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ be a strong one-way function, and X be uniform over $\{0, 1\}^n$. Then for every $\ell = O(\log n)$, $(f(X)_{1.. \ell}, \dots, f(X)_{n-\ell+1..n}, X)$ has $n^{\omega(1)}$ -inaccessible entropy $\omega(\log n)$ and $(n^{\omega(1)}, \text{negl}(n))$ -max-inaccessible entropy $\omega(\log n)$.*

Acknowledgements

We thank Muthuramakrishnan Venkitasubramaniam for an inspiring conversation which sparked this work.

References

1. Blum, M., Micali, S.: How to generate cryptographically strong sequences of pseudo random bits. In: Proceedings of the 23th Annual Symposium on Foundations of Computer Science (FOCS). pp. 112–117 (1982)
2. Brassard, G., Chaum, D., Crépeau, C.: Minimum disclosure proofs of knowledge. *Journal of Computer and System Sciences* **37**(2), 156–189 (1988)
3. Diffie, W., Hellman, M.E.: New directions in cryptography. *IEEE Transactions on Information Theory* **22**(6), 644–654 (1976)
4. Ding, Y.Z., Harnik, D., Rosen, A., Shaltiel, R.: Constant-round oblivious transfer in the bounded storage model. In: *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004*. pp. 446–472. Springer (2004)
5. Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions. *Journal of the ACM* **33**(4), 792–807 (Oct 1986)
6. Goldreich, O., Micali, S., Wigderson, A.: How to play any mental game or a completeness theorem for protocols with honest majority. In: *Proceedings of the 19th Annual ACM Symposium on Theory of Computing (STOC)*. pp. 218–229. ACM Press (1987)
7. Goldreich, O., Micali, S., Wigderson, A.: Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM* **38**(1), 691–729 (1991)

8. Haitner, I., Holenstein, T., Reingold, O., Vadhan, S.P., Wee, H.: Universal one-way hash functions via inaccessible entropy. In: Advances in Cryptology - EUROCRYPT 2010, 29th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Monaco / French Riviera, May 30 - June 3, 2010. Proceedings. pp. 616–637 (2010), https://doi.org/10.1007/978-3-642-13190-5_31
9. Haitner, I., Nguyen, M., Ong, S.J., Reingold, O., Vadhan, S.: Statistically hiding commitments and statistical zero-knowledge arguments from any one-way function. *SIAM Journal on Computing* **39**(3), 1153–1218 (2009)
10. Haitner, I., Reingold, O., Vadhan, S.: Efficiency improvements in constructing pseudorandom generators from one-way functions. In: Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC). pp. 437–446 (2010)
11. Haitner, I., Reingold, O., Vadhan, S., Wee, H.: Inaccessible entropy. In: Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC '09). pp. 611–620 (31 May–2 June 2009)
12. Haitner, I., Reingold, O., Vadhan, S.P.: Efficiency improvements in constructing pseudorandom generators from one-way functions. *SIAM J. Comput.* **42**(3), 1405–1430 (2013). <https://doi.org/10.1137/100814421>, <https://doi.org/10.1137/100814421>
13. Haitner, I., Reingold, O., Vadhan, S.P., Wee, H.: Inaccessible entropy I: Inaccessible entropy generators and statistically hiding commitments from one-way functions. (2016), www.cs.tau.ac.il/~iftachh/papers/AccessibleEntropy/IE1.pdf, to appear. Preliminary version, named Inaccessible Entropy, appeared in STOC 09
14. Håstad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. *SIAM Journal on Computing* **28**(4), 1364–1396 (1999)
15. Impagliazzo, R., Luby, M.: One-way functions are essential for complexity based cryptography. In: Proceedings of the 30th Annual Symposium on Foundations of Computer Science (FOCS). pp. 230–235 (1989)
16. Naor, M.: Bit commitment using pseudorandomness. *Journal of Cryptology* **4**(2), 151–158 (1991)
17. Naor, M., Ostrovsky, R., Venkatesan, R., Yung, M.: Perfect zero-knowledge arguments for NP using any one-way permutation. *Journal of Cryptology* **11**(2), 87–108 (1998), preliminary version in *CRYPTO'92*.
18. Naor, M., Yung, M.: Universal one-way hash functions and their cryptographic applications. In: Proceedings of the 21st Annual ACM Symposium on Theory of Computing (STOC). pp. 33–43. ACM Press (1989)
19. Nisan, N., Zuckerman, D.: Randomness is linear in space. *Journal of Computer and System Sciences* **52**(1), 43–52 (Feb 1996)
20. Rompel, J.: One-way functions are necessary and sufficient for secure signatures. In: Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC). pp. 387–394 (1990)
21. Vadhan, S.P., Zheng, C.J.: Characterizing pseudoentropy and simplifying pseudorandom generator constructions. In: Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012. pp. 817–836 (2012), <http://doi.acm.org/10.1145/2213977.2214051>
22. Yao, A.C.: Theory and applications of trapdoor functions. In: Proceedings of the 23th Annual Symposium on Foundations of Computer Science (FOCS). pp. 80–91 (1982)

A Missing Proofs

Lemma A.1. For all $t \geq 1$, $f : x \mapsto \frac{x}{1-(1-x)^t}$ is convex over $[0, 1]$.

Proof. We instead show convexity of $\tilde{f} : x \mapsto f(1-x)$. A straightforward computation gives:

$$\tilde{f}''(x) = \frac{x^{t-2}t(t(1-x)(x^t+1) - (1+x)(1-x^t))}{(1-x^t)^3}$$

so that it suffices to show the non-negativity of $g(x) = t(1-x)(x^t+1) - (1+x)(1-x^t)$ over $[0, 1]$. The function g has second derivative $t(1-x)(t^2-1)x^{t-2}$, which is non-negative when $x \in [0, 1]$, and thus the first derivative g' is non-decreasing. Also, the first derivative at 1 is equal to zero, so that g' is non-positive over $[0, 1]$ and hence g is non-increasing over this interval. Since $g(1) = 0$, this implies that g is non-negative over $[0, 1]$ and f is convex as desired.

Theorem A.2 (Theorem 3.11 restated). Let Π be a binary relation and let (Y, W) be pair of random variables supported on Π . If (Π, Y) is (t, ε) -hard, then (Π, Y, W) is (t', Δ') witness hard in relative entropy and (t', Δ'') witness hard in δ -min relative entropy for every $\delta \in [0, 1]$ where $t' = \Omega(t)$, $\Delta' = \log(1/\varepsilon)$ and $\Delta'' = \log(\delta/\varepsilon)$.

Proof. We proceed similarly to the proof of Theorem 3.5. Let $(\tilde{\mathcal{G}}, \mathcal{S})$ be a pair of algorithms with $\tilde{\mathcal{G}} = (\tilde{\mathcal{G}}_1, \tilde{\mathcal{G}}_w)$ a two-block generator supported on Π . Define the linear-time oracle algorithm $A^{\tilde{\mathcal{G}}_w, \mathcal{S}}(y) \stackrel{\text{def}}{=} \tilde{\mathcal{G}}_w(\mathcal{S}(y))$. Then

$$\begin{aligned} \Pr \left[\Pi(Y, A^{\tilde{\mathcal{G}}_w, \mathcal{S}}(Y)) \right] &= \Pr \left[\Pi(Y, \tilde{\mathcal{G}}_w(\mathcal{S}(Y))) \right] \\ &\geq \Pr \left[\tilde{\mathcal{G}}_1(\mathcal{S}(Y)) = Y \right] \quad (\tilde{\mathcal{G}} \text{ is supported on } \Pi) \\ &= \sum_{r \in \text{Supp}(\tilde{R})} \Pr \left[\mathcal{S}(Y) = r \wedge Y = \tilde{\mathcal{G}}_1(r) \right] \\ &\geq \sum_{\substack{r \in \text{Supp}(\tilde{R}) \\ w \in \text{Supp}(\tilde{\mathcal{G}}_2(\tilde{R}))}} \Pr \left[\mathcal{S}(Y) = r \wedge Y = \tilde{\mathcal{G}}_1(r) \wedge W = w \right] \\ &= \mathbf{E}_{\substack{r \leftarrow \tilde{R} \\ w \leftarrow \tilde{\mathcal{G}}_2(r)}} \left[\frac{\Pr \left[\mathcal{S}(Y) = r \wedge Y = \tilde{\mathcal{G}}_1(r) \wedge W = w \right]}{\Pr \left[\tilde{R} = r \wedge \tilde{\mathcal{G}}_2(r) = w \right]} \right] \\ &= \mathbf{E}_{\substack{r \leftarrow \tilde{R} \\ y \leftarrow \tilde{\mathcal{G}}_1(r) \\ w \leftarrow \tilde{\mathcal{G}}_2(r)}} \left[2^{-\text{KL}_{r,y,w}^*(\tilde{R}, \tilde{\mathcal{G}}_1(\tilde{R}), \tilde{\mathcal{G}}_2(\tilde{R}) \| \mathcal{S}(Y), Y, W)} \right], \end{aligned}$$

The witness hardness in relative entropy then follows by applying Jensen's inequality (since 2^{-x} is convex) and the witness hardness in δ -min relative entropy follows by Markov's inequality by considering the event that the sample relative entropy is smaller than Δ (this event has density at least δ).