Indistinguishability and Unpredictability Hardcore Lemmas: New Proofs with Applications to Pseudoentropy

Maciej Skorski

Cryptology and Data Security Group, University of Warsaw
maciej.skorski@gmail.com

Abstract. Hardcore lemmas are results in complexity theory which state that average-case hardness must have a very hard "kernel", that is a subset of instances where the problem is extremely hard. Such results find important applications in hardness amplification. In this paper we revisit two classical results:

(a) The hardcore lemma for unpredictability, proved first by Impagliazzo. It states that if a boolean function f is "moderately" hard to predict on average, then there must be a set of noticeable size on which f is "extremely" hard to predict.

(b) The hardcore lemma for indistnguishability, proved by Maurer and Tesaro, states that for two random variables X and Y which are ϵ-computationally close, there exist events A and B of probability $1-\epsilon$ such that the distributions of $X|A$ and $Y|B$ are "almost" identical.

We provide alternative proofs and some generalizations of these result in the nonuniform setting. As an interesting application, we show a strengthening of the transformation between two most popular pseudoentropy variants: HILL and Metric Entropy, and apply it to show how to extract pseudorandomness from a sequence of metric-entropy sources of poor quality. Comparing to the best known techniques we significantly improve security parameters.

1 Introduction

1.1 Hardcore Lemmas and Their Applications

Unpredictability Hardcore Lemma. Suppose that we have a predicate f that is mildly hard to predict by a class of circuits; for every circuit D from this class, $\mathrm{D}(x)$ and $f(x)$ agree on at most, lets say, a 0.99 fraction of inputs x. One of the reasons for that, which could intuitively explain this behavior, is the existence of a "kernel" for this hardness- a set of noticeable size on which f is extremely hard to predict. Quite surprisingly, this intuitive characterization is true. The first such result was proved by Impagliazzo [8]. Below we present the tight improvement due to Holenstein.

Theorem 1 (Unpredictability Hardcore Lemma [7]). *Let $f : \{0,1\}^n$ be a predicate such that f is ϵ-unpredictable by circuits of size s, that is*

$$\Pr_{x \leftarrow \{0,1\}^n}[\mathrm{D}(x) = f(x)] \leqslant 1 - \frac{\epsilon}{2}$$

holds for all D *of size at most s. Then for any $\delta \in (0,1)$ there exists a "hardcore" set S of size $\epsilon 2^n$ such that f on S is $1-\delta$ unpredictable by circuits of size $s' = \mathcal{O}\left(s\delta^2/n\right)$, that is*

$$\Pr_{x \leftarrow S}[\mathrm{D}(x) = f(x)] \leqslant \frac{1+\delta}{2}, \quad \textit{for all } \mathrm{D} \textit{ of size at most } s\delta^2/32n.$$

Remark 1. Some authors use different conventions for ϵ-unpredictability. We follow the approach of [7]. The definition above is intuitive since 1-unpredictability would means that f is totally unpredictable.

Note that the size of the hardcore set, guaranteed to be at least $2^n\epsilon$, is tight. Indeed, if the second part of the theorem is satisfied, i.e. f is almost unpredictable on a set of size ϵ, it implies that f, on average over the whole domain, cannot be predicted better than $1 - \frac{\epsilon+\delta}{2} \approx 1 - \frac{\epsilon}{2}$. An uniform version, with the tight hardcore density, is given also in [7]. Constructive versions of the hardcore lemma can be obtained by actually an arbitrary boosting algorithm [1,9], however such results are typically not tight, without additional optimization.

Indistinguishability Hardcore Lemma. It is know that if two distributions X_1, X_2 have the statistical distance at most ϵ, then there exists events A_1, A_2 of probability at least $1-\epsilon$ such that the distributions $X_1|A_1$ and $X_2|A_2$ are *identical*. Based on the reduction to the unpredictability hardcore lemma, Maurer and Tessaro proved the following computational generalization of this fact

Theorem 2 (Indistinguishability Hardcore Lemma [11]). *Let X_1 and X_2 be distributions on $\{0,1\}^n$, with the computational distance ϵ against circuits of size s, that is*

$$|\,\mathbf{E}\,\mathrm{D}(X_1) - \mathbf{E}\,\mathrm{D}(X_2)| < \epsilon \quad \textit{for all } \mathrm{D} \textit{ of size } s.$$

Then there exist events A_1 and A_2 of probability $1-\epsilon$ such that A_1 and A_2 are computationally indistinguishable, that is

$$|\,\mathbf{E}\,\mathrm{D}(X_1|A_1) - \mathbf{E}\,\mathrm{D}(X_2|A_2)| \leqslant \delta \quad \textit{for all } \mathrm{D} \textit{ of size } s = s\delta^2/128n.$$

which states that if two distributions are (computationally) not too far away from each other then after conditioning on an event of noticeable probability they are almost indistinguishable. Since the lower bound $1-\epsilon$ on the probabilities of hardcore events is tight[1], this theorem can be viewed as a characterization of computational indistinguishability.

[1] By the similar reasoning as in the unpredictability case.

Applications of hardcore lemmas. Hardcore lemmas are fundamental result in complexity theory and find applications in cryptography and learning theory. They are particularly important in the context of hardness amplification, i.e. transforming somewhat hard problems into hard problems. See for instance [5,7,8,10,11].

1.2 Our Results

An unpredictability hardcore lemma under arbitrary distributions. We prove a nonuniform version of a hardcore lemma that is true when inputs for functions are sampled from *arbitrary* distribution, not necessarily uniform. Due to connections of machine learning theory and hardcore lemma, well explained in [9], it is clear that there is nothing special in the uniform distribution and *qualitatively* similar statements indeed could be derived for any distribution. However, our approach has the following advantages:

(a) The proof strategy is very simple and natural: we observe that it is straightforward to construct a hardcore for any fixed distinguisher and then we use the min-max theorem to "reverse" the quantifiers. We believe that the proof of this form can be useful to derive some complexity lower bounds for hardcore lemma, which is (besides boosting proofs) an open problem.
(b) Our hardcore lemma is *quantitatively* tight, that is the weight of the hardcore event for ϵ-unpredictability is guaranteed to be ϵ and the loss in the complexity is $\mathcal{O}\left(\delta^2/n\right)$ for δ-closeness, which matches the best known result for the case of the uniform distribution. Actually we slightly improve security bounds with better explicit constants and replacing the dimension n by a smaller factor.
(c) The only technical difficulty in the proof, the proof that if a hardcore can be constructed for any fixed boolean circuit then the same is true for real valued circuits, is overcome by the technique of explicitly characterizing the "worst case" measures. This technique, which might be of independent interest, allows us to give a relatively short and direct (without reducing to the unpredictability version) proof of the indistinguishability hardcore lemma and, more interestingly, a variant of the indistinguishability hardcore lemma dedicated for computational entropy.

A simplified reduction from indistinguishability hardcores to unpredictability hardcores. Basing on our generalized unpredictability hardcore lemma we show an alternative proof for the indistinguishability hardcore lemma of Maurer and Tessaro. In [11] the reduction goes from the indistinuishablity hardcore lemma to the "standard" unpredictability hardcore lemma, that is where inputs are sampled from the uniform distribution. In contrary, we find it much easier and natural to reduce it to unpredictability of some predicate which explicitly depends on the distributions X_1, X_2- is simply equal to the sign of the difference between probability mass functions. In our reduction, besides better constants and decreasing the factor depending on the dimension, we also gain

an additional factor in security if the statistical distance of X_1 and X_2 is small. This slightly improves the security bounds, however this improvement seems to be of limited applicability for cryptographic applications.

A direct proof of the Indistinguishability Hardcore Lemma. By adapting the proof given for the unpredictability case, we can derive the (nonuniform) Indistinguishability Hardcore Lemma of Maurer and Tessaro *directly*, that is *without reducing* it to unpredictability hardcore lemmas. This can be interesting in the context of lower bounds. Indeed, no lower bounds on unpredictability hardcore lemmas, if they were discovered, would imply lower bounds for the indistinguishability version based on this reduction.

An Indistinguishability Hardcore Lemma for Pseudoentropy. In some situations, for instance in extracting entropy, we do not really need our distribution X to be indistinugishable from a *particular* Y but rather from *a class* of distributions Y (which is a weaker requirement). To illustrate this, consider the following alternatives to formalize the statement "X almost has property P".

(i) X is (s, δ)-close to having property P, if there exists a distribution Y with property P such that for every circuit D of size s, we have $\Delta^{\mathrm{D}}(X;Y) \leqslant \delta$
(ii) X is (s, δ)-close to having property P, if for every D of size s there exists a distribution Y with property P such that we have $\Delta^{\mathrm{D}}(X;Y) \leqslant \delta$.

where $\Delta^{\mathrm{D}}(X;Y) = \mathbf{E}\,\mathrm{D}(X) - \mathbf{E}\,\mathrm{D}(Y)$ is the advantage of the attacker D. In condition (i) we have the standard computational indistinugishability of X and Y. In turn, the second condition can be thought of as indistinguishability between X and the set of all distributions Y having property P, since it means that there is no D that separates X and all Y. Clearly condition (ii) is strictly weaker, thought it is easy to see that for convex properties P (i.e. closed under taking a convex combinations) both are equivalent up to the loss of a factor $\mathcal{O}\left(\delta^2\right)$ in circuit size [2]. Surprisingly it turns out that for many applications the weaker definition is good enough. If, for instance, P means "distribution has min-entropy at least k", then condition (ii), provided that it holds for probabilistic distinugishers, is strong enough to ensure that any (k, ϵ)-extractor applied to X yields an ϵ-pseudorandom distribution. The concept of "weak" indistinguishability, i.e. indistinugishability in the sense of (ii), is very useful in studying computational generalizations of entropy.

Set the property P to be "having min-entropy at least k". For case (i), we obtain the notion of the HILL entropy [6]: X has k bits of pseudoentropy, with quality (s, ϵ), if there is a distribution Y with k-bits of min-entropy such that no circuit of size s can distinguish X and Y with the advantage better than ϵ. In case (ii) we obtain a relaxed notion called Metric Pseudoentropy [2]: no adversary of size s can distinguish between X and *all* distributions of min-entropy at least k. As mentioned, metric pseudoentropy is very useful and widely used as a convenient substitute of HILL and find many application in studying pseudorandomness [2-4,13]. It is known [2] that metric entropy with parameters (s, ϵ) can be converted into HILL entropy with no loss in the amount and the parameters $(s', \epsilon') = \left(\mathcal{O}\left(s \cdot \delta^2/n\right), \epsilon + \delta\right)$ for any δ. Applying our techniques we

obtain a nice and much stronger version of this transformation: if X has metric entropy of quality (s,ϵ) (even against weakest deterministic circuits) then after conditioning on an event of probability $1-\epsilon$, it the same amount of HILL entropy of quality $(\delta, \mathcal{O}\left(s\cdot\delta^2/n\right))$.

Application: extracting pseudorandomness from pseudoentropy of poor quality. Using our generalized indistinguishability hardcore lemma, we prove that for a sequence of independent distributions $X_1,\ldots,X_\ell$, each having metric-entropy k with parameters (s,ϵ) for some *large* ϵ and against *deterministic* circuits of size s, the concatenated string $X=X_1,X_2,\ldots,X_\ell$ has HILL entropy roughly $(1-\epsilon)\ell k$ with parameters $(s',\delta')=(\delta, s\delta^2\ell^{-2}/n)$. In other words, for a metric pseudoentropy source of quality (s,ϵ) we achieve, sampling many times, the entropy extraction rate $\alpha=1-\epsilon$ with good security. Comparing to the state of art we save a quite large factor δ^2 in security[2].

1.3 Outline of the Paper

Section 2 provides necessary definitions for hardness of unpredictability, computational indistinguishabilty and computational entropy. In Section 3 we present a generalization of the unpredictability hardcore lemma and a slightly simplified proof of the indistinguishability hardcore lemma. A hardcore lemma dedicated for pseudoentropy is given in Section 4. An application to the problem of extracting from a pseudoentropy source of very bad quality is discussed in Section 5.

2 Preliminaries

Computational and Statistical Indistinguishability. Let X and Y be two random variables taking values in the same space. The advantage of D in distinguishing between X and Y is defined to be $\Delta^{\mathrm{D}}(X;Y)=\mathbf{E}\,\mathrm{D}(X)-\mathbf{E}\,\mathrm{D}(Y)$. The statistical distance between two random variables X and Y, is defined as $\Delta(X;Y)=\frac{1}{2}\sum_x|\Pr[X=x]-\Pr[Y=x]|$ and is equal to the maximum of $\Delta^{\mathrm{D}}(X;Y)$ over all $[0,1]$-valued functions D. The computational distance between X and Y is defined as $\max_{\mathrm{D}\in\mathcal{D}}\Delta^{\mathrm{D}}(X;Y)$ where $\mathcal{D}$ is a fixed class of boolean functions. We say that X and Y are (s,ϵ)-close or (s,ϵ)-indistinguishable if $\Delta^{\mathrm{D}}(X;Y)\leqslant\epsilon$ for all D of size at most s.

Hardness of Unpredictability. A boolean funciton $f:\{0,1\}^n\to\{0,1\}$ is said to be (s,δ)-unpredictable if $\Pr_{x\leftarrow}[\mathrm{D}(x)=f(x)]\leqslant 1-\delta/2$ for all D of size at most s. We also say that f is δ-hard against circuits of size s. We say that f is (s,δ)-unpredictable under the distribution V if $\Pr_{x\leftarrow V}[\mathrm{D}(x)=f(x)]\leqslant 1-\delta/2$ for all D of size at most s.

[2] We note that the following issues makes this problem challenging: (a) since ϵ is large, no hybrid technique can be applied and (b) pseudoentropy is *only* against deterministic adversaries so no extractor can be directly applied.

Computational Entropy. There are many ways to define computational analogues of entropy. We follow the most popular approach, which is based on the concept of computational indistinguishability.

Definition 1 (HILL Pseudoentropy [6]). *Let X be a distribution with the following property: there exists Y of min-entropy at least k such that for all circuits D of size at most s we have $|\Delta^{\mathrm{D}}(X;Y)| \leqslant \epsilon$. Then we say that X has k bits of HILL entropy of quality (s,ϵ) and denote by $\mathbf{H}^{\mathrm{HILL}}(X)_{s,\epsilon} \geqslant k$.*

It is known that for HILL Entropy all kind of circuits: deterministic boolean, deterministic real valued and randomized boolean, are equivalent (with the same size s). The following definition differs in the order of quantifiers

Definition 2 (Metric Pseudoentropy [2]). *Let X be a distribution with the following property: for every deterministic boolean (respectively: deterministic real valued or boolean randomized) circuit D of size at most s there exists Y of min-entropy at least k such that $|\Delta^{\mathrm{D}}(X;Y)| \leqslant \epsilon$. Then we say that X has k bits of deterministic (respectively: deterministic real valued or boolean randomized) metric entropy of quality (s,ϵ) and denote by $\mathbf{H}^{\mathrm{Metric,det}\{0,1\}}(X)_{s,\epsilon} \geqslant k$ (respectively: $\mathbf{H}^{\mathrm{Metric,det}[0,1]}(X)$ and $\mathbf{H}^{\mathrm{Metric,rand}}(X)$).*

3 Hardcore Lemmas

3.1 Approximating Convex Hulls

The following standard facts, derived by the Hoeffding-Chernoff Inequality, are useful when we want to approximate possibly long convex combinations of functions by a combination of few functions; for instance, when we use the min-max theorem and need to approximate any mixed strategy by an efficient strategy.

Lemma 1 ([12], Lemma 2.1). *Let $\mathcal{X}$ be a finite domain, $\mathcal{G}$ be any set of functions $g : \mathcal{X} \to [-1,1]$ and let $\overline{g}$ be a convex combinations of functions from $\mathcal{G}$. Then for any $\epsilon \in (0,1)$ and for some $k \leqslant \frac{4}{\epsilon^2}\log\left(\frac{2}{\epsilon}\right)$, there exist functions $g_1,\ldots,g_k$ such that*

$$\mathbf{E}_{x\leftarrow\nu}\left|\overline{g}(x) - \left(\frac{1}{k}\sum_{i=1}^{k} g_i(x)\right)\right| \leqslant \epsilon$$

Lemma 2 ([2]). *Let $\mathcal{X}$ be a finite domain, ν be a distribution on $\mathcal{X}$ and let $\mathcal{G}$ be any set of functions $g : \mathcal{X} \to [-1,1]$ and let $\overline{g}$ be a convex combinations of functions from $\mathcal{G}$. Then for any $\epsilon \in (0,1)$ and for some $k \leqslant \frac{\log|\mathcal{X}|}{2\epsilon^2}$, there exist functions $g_1,\ldots,g_k$ such that*

$$\max_{x\in\mathcal{X}}\left|\overline{g}(x) - \left(\frac{1}{k}\sum_{i=1}^{k} g_i(x)\right)\right| \leqslant \epsilon$$

3.2 Hardcore Lemma for Unpredictability under Arbitrary Distributions

Below we prove a hardcore lemma, with the optimal weight of the hardcore, valid for an *arbitrary* distribution. We also obtain better constants than in Theorem 1 and an improvement for the case when ϵ is bounded away from 0: we replace then the dimension n by $\log(1/\delta)$ which is typically much smaller.

Theorem 3 (Unpredictability Hardcore Lemma for arbitrary distributions). *Let V be an arbitrary distribution on $\{0,1\}^n$ and suppose that an n-bit boolean function f is (s,ϵ)-unpredictable under sampling from a distribution V. Then for any δ there exists an event A of probability at least 2ϵ such that f is $(s',1-\delta)$-unpredictable under $V|A$, where $s'=\max\big(s\cdot 2\delta^2/n, s\cdot 4\epsilon^2\delta^2/\log(4/\epsilon\delta)\big)$.*

Note that f is essentially *almost* unbiased under $V|A$: by applying trivial distinguishers $\mathrm{D}\equiv 1$ and $\mathrm{D}\equiv 0$ we get $\frac{1}{2}-\delta \leqslant \Pr[f(V|A)=1] \leqslant \frac{1}{2}+\delta$. For some technical reasons we need the following observation, which states that the hardcore event "preserves" unbiased predicates.

Corollary 1 (Unpredictability Hardcore Lemma for unbiased predicates). *Suppose that Theorem 3 holds for f and V such that $\mathbf{P}\left(f(V)=-1\right)=\frac{1}{2}=\mathbf{P}\left(f(V)=1\right)$. Then the hardcore event A can be chosen in such a way that $\mathbf{P}\left(f(V|A)=-1\right)=\mathbf{P}\left(f(V|A)=1\right)=\frac{1}{2}$, with the additional loss of the factor 3 in circuit size.*

The proof of Corollary 1 appears in the full version of the paper. It is relatively simple and uses the idea of "mass-shifting". The proof of Theorem 3 consists of the three important steps: (a) the trivial observation that for any fixed boolean function D of size s we can find a hardcore event, (b) the observation that we can find a hardcore for any $[0,1]$-valued function of approximately the same size s and (c) using the min-max theorem and approximation lemmas to switch the order of quantifiers to find a one hardcore event that works with all functions of size s. Only step (b) is non-trivial and novel (and allows us to slightly improve Hollentein's bounds). We prove (b) by assuming that there exists a function D for which one cannot find a suitable hardcore measure. Then we argue that if this is true, we cannot find a hardcore measure for some threshold transformation of D. We characterize the measure which is "closest" to be a hardcore measure and use it to show that the same measure is optimal for all threshold transformations of D. By taking an appropriate threshold we conclude that we cannot find a hardcore measure for some threshold versions of D, which is impossible since it is now boolean and this contradicts the assumptions. The proof appears in the full version of the paper.

3.3 Hardcore Lemma for Indistinguishability - Reduction to Unpredictability Case

The following lemma shows that indistinguishability of two distributions is equivalent to the hardness of predicting some boolean function, which explicitly depends on these distributions. This function is quite natural: it equals the sign of the difference between the probability mass functions.

Lemma 3. *Let $\mathcal{D}$ be a class of boolean functions, $X, Y \in \{0,1\}^n$ be random variables, and let $\Delta = \Delta(X,Y)$ be different than 0. Then the following are equivalent:*

(a) X and Y are $(\mathcal{D}, \epsilon)$-indistinguishable
(b) $f(x)$ is $(\mathcal{D}, 1-\epsilon/\Delta)$-unpredictable under V, where $f(x)$ is the indicator of the set $\{x : \mathbf{P}_X(x) > \mathbf{P}_Y(x)\}$ and the distribution of V is given by $\mathbf{P}_V(x) = |\mathbf{P}_X(x) - \mathbf{P}_Y(x)|/2\Delta$.

Proof. For any boolean D we obtain

$$\begin{aligned}\mathbf{E}\,\mathrm{D}(X) - \mathbf{E}\,\mathrm{D}(Y) &= \sum_x (\mathbf{P}_X(x) - \mathbf{P}_Y(x))\mathrm{D}(x) \\ &= 2\Delta\,(\Pr[f(V)=1]\,\mathbf{E}[\mathrm{D}(V)|f(V)=1] \\ &\quad - \Pr[f(V)=0]\,\mathbf{E}[\mathrm{D}(V)|f(V)=0])\end{aligned}$$

Observe that $\Pr[f(V)=1] = \Pr[f(V)=0] = \frac{1}{2}$. Therefore

$$\mathbf{E}\,\mathrm{D}(X) - \mathbf{E}\,\mathrm{D}(Y) = 2\Delta\left(-\frac{1}{2} + \frac{1}{2}\,\mathbf{E}[\mathrm{D}(V)|f(V)=1] + \frac{1}{2}\,\mathbf{E}[(1-\mathrm{D}(V))|f(V)=0]\right).$$

Since D is boolean, the last equation is equivalent to

$$\mathbf{E}\,\mathrm{D}(X) - \mathbf{E}\,\mathrm{D}(Y) = 2\Delta\left(\Pr[\mathrm{D}(V) = f(V)] - \frac{1}{2}\right),$$

which finishes the proof. □

Based on Lemma 3 we prove the following result

Theorem 4 (Indistinguishability Hardcore Lemma). *Suppose that X and Y are arbitrary (s,ϵ) indistinguishable by boolean circuits Then there exists an events $A(X), A(Y)$ both of equal probability at least $1-\epsilon$ such that $X|A(X)$ and $Y|A(Y)$ are $(\mathcal{O}\left(s\cdot\delta^2/n\right), \Delta(X;Y)\cdot\delta)$ indistinguishable.*

Proof. From the construction of V, we know that f is $1-\epsilon/\Delta(X,Y)$-unpredictable under V. From Theorem 3 we obtain that there exists a hardcore A with probability at least $1-\epsilon/\Delta(X,Y)$ such that f is *extremely* unpredictable under $V|A$. This hardcore event can be described as follows: there exists a measure $M = M_A$ that satisfies $M(x) \leqslant \mathbf{P}_V(x)$ and $\mathbf{P}(A) = \mu(M) \geqslant 1-\epsilon/\Delta(X,Y)$ and such that $f(x)$ is unpredictable for sampling according to M, i.e. $\mathbf{P}_{x\leftarrow M}\left(D(x) = f(x)\right) < 1/2+\delta$. The distribution $V|A$ is then defined by $\mathbf{P}_{V|A} = \mathbf{P}_M$. Consider the events $S^- = \{x :\ f(x) = 0\}$ and $S^+ = \{x :\ f(x) = 1\}$. From the definition of V and f it follows that $\mathbf{P}_V\left(S^-\right) = \mathbf{P}_V\left(S^+\right) = \frac{1}{2}$. As shown in Corollary 1, the sets

S^+, S^- can be assumed to be *perfectly unbiased* also under $V|A$. Define now two measures $M_0 = M_X$ and $M_1 = M_Y$ as follows:

$$M_0(x) = \begin{cases} \mathbf{P}_X(x) - 2\Delta(X,Y)\left(\mathbf{P}_V(x) - M'(x)\right) & \text{if } \mathbf{P}_X(x) > \mathbf{P}_Y(x) \\ \mathbf{P}_X(x) & \text{otherwise} \end{cases} \tag{1}$$

and similarly,

$$M_1(x) = \begin{cases} \mathbf{P}_Y(x) - 2\Delta(X,Y)\left(\mathbf{P}_V(x) - M'(x)\right) & \text{if } \mathbf{P}_X(x) < \mathbf{P}_Y(x) \\ \mathbf{P}_Y(x) & \text{otherwise} \end{cases} \tag{2}$$

Note that both measures are well defined since $\mathbf{P}_V(x) = |\mathbf{P}_X(x) - \mathbf{P}_Y(x)| / 2\Delta(X,Y)$ and $M'(x) \leqslant \mathbf{P}_V(x)$. Then from the definition of (V, A) and the definition of f it follows that

$$\begin{aligned} \mu(M_0) &= 1 - 2\Delta(X,Y) \sum_{x:\ f(x)=1} \mathbf{P}_V(x) + 2\Delta(X,Y) \sum_{x:\ f(x)=1} M'(x) \\ &= 1 - \Delta(X,Y) + 2\Delta(X,Y)\mathbf{P}(A) \cdot \mathbf{P}_{V|A}\left(S^+\right) \\ &= 1 - \Delta(X,Y)\mathbf{P}(A^c) \\ &\geqslant 1 - \epsilon \end{aligned} \tag{3}$$

and similarly that the same estimate holds for $\mu(M_1)$. Observe also that since S^+ and S^- are perfectly unbiased with respect to M', and since the same holds for V, we have $\mu(M_0) = \mu(M_1)$. These measures give rise to the joint distributions $X, A(X)$ and $Y, A(Y)$ for some events $A(X), A(Y)$ with probabilities at least $\mu(M_0) = \mu(M_1)$. It remains to calculate the advantage in distinguishing. Let V' and f' be a distrubtion and a predicate corresponding to $X|A(X)$ and $Y|A(Y)$ according to the statement of Lemma 3. Observe that $M_0(x) > M_1(x)$ if and only if $f(x) = 1$, hence $f'(x) = f(x)$. Since $|M_0(x) - M_1(x)| = 2\Delta(X,Y)M'(x)$ for every x, we get $\mathbf{P}_V(x) = M'(x)/\mu(M') = \mathbf{P}_{V|A}(x)$ and $\Delta(X|A(X), Y|A(Y)) = \Delta(X,Y)$. Therefore

$$\begin{aligned} \Delta^D(X|A(X), Y|A(Y)) &= \Delta(X,Y) \cdot \left(2\mathbf{P}_{x \leftarrow V'}\left(D(x) = f'(x)\right) - 1\right) \\ &= \Delta(X,Y) \cdot \left(2\mathbf{P}_{x \leftarrow V|A'}\left(D(x) = f(x)\right) - 1\right) \\ &< \Delta(X,Y) \cdot \delta, \end{aligned} \tag{4}$$

and we have finished the proof. □

Remark 2. We note that without Corollary 1 we would obtain a slightly weaker version of the indistinguishability hardcore lemma where the probability of the hardcore events is guaranteed to be at least $1 - \epsilon - \delta$, which is very close to the optimal $1 - \epsilon$ and equally good in applications.

4 Indistinguishability Hardcore Lemma for Pseudoentropy

In this section we prove the following theorem, discussed in the introduction, which gives the existence of a "HILL-entropy-hardcore" for metric pseudoentropy.

Theorem 5 (Indistinguishability Hardcore Lemma for pseudoentropy). *Suppose that* $\mathbf{H}^{\mathrm{Metric,det}\{0,1\}}_{s,\epsilon}(X) \geqslant k$. *Then for any* δ *and* $s' = \mathcal{O}\left(s \cdot \delta^2/n\right)$ *there exists an event* A *of probability* $1-\epsilon$ *such that* $\mathbf{H}^{\mathrm{HILL}}_{s',\delta}(X|A) \geqslant k - \log(1/(1-\epsilon))$.

This theorem shows that metric entropy not only can be converted to HILL entropy with the loss of factor δ in advantage and δ^2 in circuit size; It has a hardcore of HILL entropy with the same quality parameters.

Remark 3. The constant hidden in the big "O" term is at most 2/3.

Before we give the proof, let us observe that this result implies the transformation between metric and HILL entropy (up to the lose of at most one bit)

Corollary 2 (Metric entropy - HILL entropy transformation [2]). *Suppose that* $\mathbf{H}^{\mathrm{Metric,det}\{0,1\}}_{s,\epsilon}(X) \geqslant k$. *Then* $\mathbf{H}^{\mathrm{HILL}}_{s',\epsilon'}(X) \geqslant k$ *where* $s' = \mathcal{O}\left(s \cdot \delta^2/n\right)$ *and* $\epsilon' = \epsilon + \delta$.

Proof (Proof of Corollarz 2). We apply Theorem 5 obtaining a distribution $Y|A$ which is (s', δ)-indistinguishable from $X|A$, and then we define $\Pr[Y' = x] = \Pr[A] \cdot \Pr[Y = x|A] + 2^{-n} \Pr[A^c]$. Note that $\mathbf{H}_\infty(Y') \geqslant k-1$ and Y' is $(s', \epsilon+\delta)$-indistinguishable from X. We remark that one can actually show without the loss of 1 bit, because Theorem 5 actually is slightly stronger that stated, namely $\mathbf{H}^{\mathrm{HILL}}_{s',\delta}(X|A) \geqslant k - \log(1/(1-\epsilon))$ can be replaced by the following: $X|A$ is (s', δ)-indistinguishable from $Y|A$ where Y has k bits of min-entropy. □

The proof strategy for Theorem 5 is exactly the same as in the case of Theorem 3; the proof appears in the full version of this paper. Note that the result in Theorem 5 with much worse parameters follows by converting metric-entropy into HILL entropy using Corollary 2 and then applying Theorem 2. This way we lose δ^4 in circuit size.

Corollary 3 (Direct proof of the Indistinguishability Hardcore Lemma). *The proof of Theorem 5 can be easily adapted to give a direct proof of Theorem 4 without reducing it to Theorem 3. Namely, in the proof we replace the condition* $M_2 \leqslant 2^{-k}$ *by* $M_2 \leqslant \mathbf{P}_Y$.

5 Applications: Extracting from Metric Pseudoentropy of Poor Quality

Suppose that we have a source of metric pseudoentropy that produces samples secure against deterministic adversaries of high complexity but only with a very big advantage ϵ (for instance, $\epsilon = 0.25$). Since the metric entropy is only against deterministic adversaries, for which it is not known if we can extract pseudorandomness directly[3], one needs to convert in into the HILL entropy. However, it

[3] The problem of randomized vs deterministic adversaries is the matter of metric entropy only; as already mentioned, for the HILL entropy all kind of circuits are equivalent.

still does not solve the problem of large ϵ. In the next step one can use Theorem 2 to prove that a concatenated sequence of many samples has large HILL entropy[4], with the rate of roughly $1-\epsilon$. This approach loses $\mathcal{O}\left(\delta^4\right)$ in security. Below we show that these two steps can be done *at the same time* which allows us to save a factor of $\mathcal{O}\left(\delta^2\right)$ in security.

Theorem 6. *Suppose that X_i, for $i=1,\ldots,\ell$, are independent n-bit random variables such that $\mathbf{H}^{\mathrm{Metric,det}\{0,1\}}_{s,\epsilon}(X) \geqslant k$. Then for any $\gamma>0$ we have*

$$\mathbf{H}^{\mathrm{HILL}}_{s',\delta'}(X) \geqslant (1-\epsilon-\gamma)\ell\left(k-\log(1/(1-\epsilon))\right),$$

where $s'=\mathcal{O}\left(s\cdot\delta^2/n\ell^2\right)$ and $\delta'=\delta+2\exp(-2\ell\gamma^2)$

Proof. Fix δ and let $s'=\mathcal{O}\left(s\cdot\delta^2/n\right)$. We apply Theorem 5 to X_i, for $i=1,,\ldots,\ell$, obtaining hardcore events A_i of probability at least $1-\epsilon$ such that $\mathbf{H}^{\mathrm{HILL}}_{s',\delta}(X_i|A_i) \geqslant k-\log(1/(1-\epsilon))$. By the Chernoff Bound we know that the probability that $m=\ell(1-\epsilon-\gamma)$ of them happen simultaneously, is at least $1-2\exp(-2\ell\gamma^2)$. The result follows now by the observation that concatenating ℓ random variables $Y_1,\ldots,Y_\ell$ of HILL entropy $k_1,\ldots,k_\ell$ with parameters (s',δ) yields a distribution of HILL entropy $k_1+k_2+\ldots+k_\ell$ with parameters $(s',\ell\delta)$ (the proof if by standard hybrid technique). □

6 Conclusion

An interesting open problem is to check if the indistinguishability hardcore lemma can be derived from the unpredictability hardcore lemma, that is show the reduction in other direction than in [11] and this paper. Another problem worth of mentioning is the question about the lower bounds on the necessary loss in security for hardcore lemmas. To the best of our knowledge, nothing is known about negative results so far, except the lower bounds for proofs based on boosting, which follow from general machine learning theory [9].

References

1. Barak, B., Hardt, M., Kale, S.: The uniform hardcore lemma via approximate bregman projections. In Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '09, pp. 1193–1200, Philadelphia, PA, USA, Society for Industrial and Applied Mathematics (2009)
2. Barak, B., Shaltiel, R., Wigderson, A.: Computational analogues of entropy. In Arora, S., Jansen, K., Rolim, J.D.P., Sahai, A.(eds.), RANDOM-APPROX, vol. 2764 of Lecture Notes in Computer Science, pp. 200–215. Springer (2003)
3. Benjamin, F., Leonid, R.: A unified approach to deterministic encryption: New constructions and a connection to computational entropy. In TCC 2012, vol. 7194 of LNCS, pp. 582–599. Springer (2012)

[4] Maurer and Tessaro construct in the same way a PRG from a weak PRG.

4. Dziembowski, S., Pietrzak, K.: Leakage-resilient cryptography in the standard model. IACR Cryptology ePrint Archive, 2008:240 (2008)
5. Goldreich, O., Nisan, N., Wigderson, A.: Studies in complexity and cryptography. chapter On Yao's XOR-lemma, pp. 273–301. Springer-Verlag, Berlin, Heidelberg (2011)
6. Hastad, J., Impagliazzo, R., Levin, L.A., Luby, M.: A pseudorandom generator from any one-way function. SIAM J. Comput., 28(4):1364–1396 (1999)
7. Holenstein, T.: Key agreement from weak bit agreement. In Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing, STOC '05, pp. 664–673, New York, NY, USA (2005)
8. Impagliazzo, R.: Hard-core distributions for somewhat hard problems. In Proceedings of the 36th Annual Symposium on Foundations of Computer Science, FOCS '95, pp. 538–, Washington, DC, USA, IEEE Computer Society (1995)
9. Klivans, A.R., Servedio, R.A.: Boosting and hard-core sets. In Proceedings of the 40th Annual Symposium on Foundations of Computer Science, FOCS '99, pp. 624–, Washington, DC, USA IEEE Computer Society (1999)
10. Lin, H., Tessaro, S.: Amplification of chosen-ciphertext security. In Thomas Johansson and PhongQ. Nguyen, editors, Advances in Cryptology EUROCRYPT 2013, vol. 7881 of Lecture Notes in Computer Science, pp. 503–519. Springer Berlin Heidelberg (2013)
11. Maurer, U., Tessaro, S.: A hardcore lemma for computational indistinguishability: Security amplification for arbitrarily weak prgs with optimal stretch. In Proceedings of the 7th International Conference on Theory of Cryptography, TCC'10, pp. 237–254, Berlin, Heidelberg (2010) Springer-Verlag.
12. Trevisan, L., Tulsiani, M., Vadhan, S.: Regularity, boosting, and efficiently simulating every high-entropy distribution. In Proceedings of the 2009 24th Annual IEEE Conference on Computational Complexity, CCC '09, pp. 126–136, Washington, DC, USA, IEEE Computer Society (2009)
13. Vadhan, S., Zheng, C.J.: Characterizing pseudoentropy and simplifying pseudorandom generator constructions. In Proceedings of the 44th symposium on Theory of Computing, STOC '12, pp. 817–836, New York, NY, USA (2012)