Impagliazzo's Worlds Through the Lens of Conditional Kolmogorov Complexity

Zhenjian Lu
University of Warwick, UK

Rahul Santhanam
University of Oxford, UK

Abstract

We develop new characterizations of Impagliazzo's worlds Algorithmica, Heuristica and Pessiland by the intractability of conditional Kolmogorov complexity K and conditional probabilistic time-bounded Kolmogorov complexity pK^t.

In our first set of results, we show that $\mathsf{NP} \subseteq \mathsf{BPP}$ iff $\mathsf{pK}^t(x \mid y)$ can be computed efficiently in the worst case when t is sublinear in $|x| + |y|$; $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$ iff $\mathsf{pK}^t(x \mid y)$ can be computed efficiently over all polynomial-time samplable distributions when t is sublinear in $|x| + |y|$; and infinitely-often one-way functions fail to exist iff $\mathsf{pK}^t(x \mid y)$ can be computed efficiently over all polynomial-time samplable distributions for t a sufficiently large polynomial in $|x| + |y|$. These results characterize Impagliazzo's worlds Algorithmica, Heuristica and Pessiland purely in terms of the tractability of conditional pK^t. Notably, the results imply that Pessiland fails to exist iff the average-case intractability of conditional pK^t is insensitive to the difference between sublinear and polynomially bounded t. As a corollary, while we prove conditional pK^t to be NP-hard for sublinear t, showing NP-hardness for large enough polynomially bounded t would eliminate Pessiland as a possible world of average-case complexity.

In our second set of results, we characterize Impagliazzo's worlds Algorithmica, Heuristica and Pessiland by the distributional tractability of a natural problem, i.e., approximating the conditional Kolmogorov complexity, that is provably intractable in the worst case. We show that $\mathsf{NP} \subseteq \mathsf{BPP}$ iff conditional Kolmogorov complexity can be approximated in the *semi-worst case*; and $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$ iff conditional Kolmogorov complexity can be approximated on average over all *independent polynomial-time samplable distributions*. It follows from a result by Ilango, Ren, and Santhanam (STOC 2022) that infinitely-often one-way functions fail to exist iff conditional Kolmogorov complexity can be approximated on average over all *polynomial-time samplable distributions*. Together, these results yield the claimed characterizations. Our techniques, combined with previous work, also yield a characterization of auxiliary-input one-way functions and equivalences between different average-case tractability assumptions for conditional Kolmogorov complexity and its variants. Our results suggest that novel average-case tractability assumptions such as tractability in the semi-worst case and over independent polynomial-time samplable distributions might be worthy of further study.

2012 ACM Subject Classification Theory of computation → Computational complexity and cryptography

Keywords and phrases meta-complexity, Kolmogorov complexity, one-way functions, average-case complexity

Digital Object Identifier 10.4230/LIPIcs.ICALP.2024.110

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: `https://eccc.weizmann.ac.il/report/2024/085/` [20]

Acknowledgements We thank Shuichi Hirahara, Yanyi Liu, Igor C. Oliveira, and Hanlin Ren for useful discussions.

51st International Colloquium on Automata, Languages, and Programming (ICALP 2024).
Editors: Karl Bringmann, Martin Grohe, Gabriele Puppis, and Ola Svensson;
Article No. 110; pp. 110:1–110:17

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

In his influential survey on average-case complexity [12], Impagliazzo described five possible computational worlds: Algorithmica, Heuristica, Pessiland, Minicrypt and Cryptomania. Algorithmica is a world where NP is easy in the worst case; Heuristica a world where NP is hard in the worst case but easy on average; Pessiland a world where NP is hard on average but one-way functions do not exist; Minicrypt a world where one-way functions exist but public-key cryptography does not; and Cryptomania a world where public-key cryptography exists. The general belief among complexity theorists and cryptographers is that we live in Cryptomania, but we are very far from a proof, as even ruling out Algorithmica would involve showing $\mathsf{NP} \neq \mathsf{P}$.

There is the possibility, however, that we might be able to unconditionally rule out some of the intermediate worlds, such as Heuristica, Pessiland and Minicrypt. Until recently, there was little progress on ruling out these intermediate worlds. All that was known was that there are various black-box and relativization barriers to ruling out these worlds.

The study of *meta-complexity*, i.e., the complexity of computational problems that are themselves about complexity, has enabled new attacks on these questions. Examples of meta-complexity problems are the Minimum Circuit Size Problem (MCSP), which asks whether a Boolean function represented by its truth table has circuits of a given size, and the problem of computing Kolmogorov complexity and its resource-bounded variants such as Levin's time-bounded Kolmogorov complexity. The average-case complexity of meta-complexity problems is of particular interest [9]. Hirahara [5] gave an approach via meta-complexity to ruling out the analogue of Heuristica for the Polynomial Hierarchy. More recently, the Polynomial Hierarchy analogue of Pessiland has been ruled out [10], again using meta-complexity techniques.

There have been several successful efforts to characterize the existence of one-way functions via meta-complexity. In [23], a conditional characterization was given, based on a believable but seemingly hard-to-establish conjecture. Liu and Pass [14] unconditionally characterized one-way functions by the average-case hardness of polynomial-time-bounded Kolmogorov complexity over the uniform distribution. This characterization was extended to other meta-complexity problems and notions of one-way function in [15, 21, 1]. A different characterization of one-way functions via the hardness of approximating Kolmogorov complexity over samplable distributions was given in [11]. More recently, Hirahara [7] introduced a meta-complexity problem whose NP-hardness and the worst-case hardness of NP characterize the existence of one-way functions.

These connections between meta-complexity, average-case complexity and one-way functions raise the following question: Can we characterize Impagliazzo's worlds Algorithmica, Heuristica and Pessiland by different notions of hardness for a single computational problem? A positive answer to this question is implicit in [16], who study the problem of conditional polynomial-time-bounded Kolmogorov complexity. They show that the worst-case hardness of conditional polynomial-time-bounded Kolmogorov complexity captures worst-case hardness of NP, and the average-case hardness of conditional polynomial-time-bounded Kolmogorov complexity over the uniform distribution captures the existence of one-way functions. Their result on worst-case hardness immediately implies that the average-case hardness of NP is equivalent to the hardness of conditional polynomial-time-bounded Kolmogorov complexity over some samplable distribution.

In this work, we give two new characterizations of Impagliazzo's worlds by different notions of hardness for a single problem - first for conditional probabilistic time-bounded Kolmogorov complexity pK^t [3], and second for the standard notion of conditional Kolmogorov complexity.

These new characterizations have some interesting features. The first characterization implies that ruling out Pessiland corresponds to *robustness* of the average-case tractability of conditional pK^t over time regimes t that vary from sublinear to polynomial. As a consequence, while we are able to prove (by building on [6]) that pK^t is NP-hard to compute exactly when t is sublinear, Pessiland would fail to exist if pK^t were NP-hard to compute for *arbitrary* polynomial t. This could be a promising route to ruling out Pessiland, since pK^t is a fairly powerful complexity measure with nice properties such as the coding theorem which could potentially be exploited when showing hardness, and the computational version is in (promise) AM but is not known to be in NP.

The second characterization is for a fundamental problem that is *provably intractable in the worst case*, i.e., the problem of approximating conditional Kolmogorov complexity. A somewhat surprising aspect of our results (which is also present in the main result of [11] on which we build) is that conditional Kolmogorov complexity is uncomputable, yet natural average-case hardness assumptions on conditional Kolmogorov complexity capture complexity worlds related to average-case hardness of NP. What this indicates is that the distinctions between Impagliazzo's worlds can be encoded in a natural way into the *distributional assumptions* that are made, while considering a single well-understood problem.

As a corollary of our second set of results together with those in [16], we get new equivalences between hardness assumptions for conditional Kolmogorov complexity and hardness assumptions for conditional time-bounded Kolmogorov complexity. The proofs of these equivalences crucially use the various characterizations of Impagliazzo's worlds, and it seems tricky to show such equivalences directly.

1.1 Results

We state our results formally in this subsection.

1.1.1 Characterizing Both $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$ and Non-Existence of One-Way Functions by Average-Case Easiness of Conditional pK^t

We present a meta-complexity problem whose average-case tractability over polynomial-time samplable distributions can be used to characterize both the non-existence of one-way functions and $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$, while considering different time regimes in the measure of time-bounded Kolmogorov complexity. Specifically, we consider the problem of computing conditional probabilistic t-time-bounded Kolmogorov complexity.

As defined in [3], we let $\mathsf{pK}^t_\lambda(x \mid y)$ be the smallest integer k such that, with probability at least λ over the choice of a random string $w \sim \{0,1\}^t$, there is a (deterministic) program of size k that, when running on w and given oracle access to y, prints x within t steps (see [20, Definition 16] for the formal definition).

For $\tau\colon \mathbb{N} \times \mathbb{N} \to \mathbb{N}$, let $\mathsf{Cond\text{-}pK}[\tau]$ be the following promise problem $(\mathsf{YES}, \mathsf{NO})$:

$$\mathsf{YES} := \left\{ (x, y, 1^s) \mid \mathsf{pK}^{\tau(|x|,|y|)}_{2/3}(x \mid y) \leq s \right\},$$
$$\mathsf{NO} := \left\{ (x, y, 1^s) \mid \mathsf{pK}^{\tau(|x|,|y|)}_{1/3}(x \mid y) > s \right\}.$$

We will refer to this problem as "computing conditional pK^t".

We will consider two specific settings for the time bound function τ. For the purpose of illustration, let us consider the following simplified problem. For $\tau\colon \mathbb{N} \times \mathbb{N} \to \mathbb{N}$, we are given x, y and s, and the task is to decide whether $\mathsf{K}^{\tau(|x|,|y|)}(x \mid y) \leq s$, i.e., whether there is a program of size at most s such that given *oracle access* to y, the program outputs x within time $\tau(|x|, |y|)$.

A typical setting of τ is $\tau(n,m) := n^c \cdot m^c$, where $c > 1$ is some constant. For this τ, we want to decide if there is a program of size at most s that, given *oracle access* to y, outputs x within time $\tau(|x|,|y|)$, and such a program has enough time to read the entire string y.

Now consider another setting of τ where $\tau(n,m) := n^c \cdot m^{1-1/c}$ for a constant $c > 1$. In this case, for a string $y \in \{0,1\}^m$, where $m := n^{2c^2}$, we have

$$\tau(n,m) = n^c \cdot m^{1-1/c} = n^{2c^2-c} \ll m.$$

Again, we want to decide if there is a program of size at most s that, given *oracle access* to y, outputs x within time $\tau(|x|,|y|)$. However, in this case any such program does not have time to read the entire string y

We will show that the non-existence of one-way functions corresponds to the average-case tractability of $\mathsf{Cond\text{-}pK}[\tau]$ over polynomial-time samplable distributions for the "polynomial-time regime" of τ, and that $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$ corresponds to that of the "sublinear-time regime".[1] We state our results formally next.

For an algorithm A, $x, y \in \{0,1\}^*$, and $s \in \mathbb{N}$, we say that A *decides* $\mathsf{Cond\text{-}pK}[\tau]$ *on* $(x, y, 1^s)$ if the following holds:

$$A(x,y,1^s) = \begin{cases} 1 & \text{if } \mathsf{pK}^{\tau(|x|,|y|)}_{2/3}(x \mid y) \le s, \\ 0 & \text{if } \mathsf{pK}^{\tau(|x|,|y|)}_{1/3}(x \mid y) > s, \\ \textbf{either } 0 \textbf{ or } 1 & \text{otherwise.} \end{cases}$$

▶ **Theorem 1.** *The following are equivalent.*

1. *Infinitely-often one-way functions do not exist.*
2. **(Computing conditional pK^t in the polynomial-time regime is easy-on-average over samplable distributions.)**
   For every polynomial-time samplable distribution family $\{\mathcal{D}_{\langle n,m\rangle}\}_{n,m}$ supported over $\{0,1\}^n \times \{0,1\}^m$, every polynomial q, and for all large enough constant c, there exists a probabilistic polynomial-time algorithm A such that for all $n, m, s \in \mathbb{N}$,

$$\Pr_{(x,y)\sim\mathcal{D}_{\langle n,m\rangle}}[A \text{ decides } \mathsf{Cond\text{-}pK}[\tau] \text{ on } (x,y,1^s)] \ge 1 - \frac{1}{q(n,m)},$$

   where $\tau(n,m) := n^c \cdot m^c$.

▶ **Theorem 2.** *The following are equivalent.*

1. $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$.
2. **(Computing conditional pK^t in the sublinear-time regime is easy-on-average over samplable distributions.)**
   For every polynomial-time samplable distribution family $\{\mathcal{D}_{\langle n,m\rangle}\}_{n,m}$ supported over $\{0,1\}^n \times \{0,1\}^m$, every polynomial q, and for all large enough constant c, there exists a probabilistic polynomial-time algorithm A such that for all $n, m, s \in \mathbb{N}$,

$$\Pr_{(x,y)\sim\mathcal{D}_{\langle n,m\rangle}}[A \text{ decides } \mathsf{Cond\text{-}pK}[\tau] \text{ on } (x,y,1^s)] \ge 1 - \frac{1}{q(n,m)},$$

   where $\tau(n,m) := n^c \cdot m^{1-1/c}$.

[1] Note that even in the "sublinear-time regime" of τ, the program can still run in polynomial time with respect to the the length of x; the word "sublinear-time" refers to the fact that the program runs in sublinear time with respect to the length of y.

In proving Theorem 2, we also show that it is NP-hard to compute conditional pK^t in the sublinear-time regime in the worst case.

▶ **Theorem 3** (Informal). *For any constant $c > 1$, $\mathsf{Cond\text{-}pK}[\tau]$ is NP-hard under randomized polynomial-time reductions, where $\tau(n, m) := n^c \cdot m^{1-1/c}$.*

In fact, Theorem 3 holds even if we consider the problem of approximating $\mathsf{pK}^t(x \mid y)$ in the sublinear-time regime within a multiplicative factor of $|x|^{1/\log\log|x|^{O(1)}}$. This also extends a result by Liu and Pass [16] and Hirahara [6], which showed that the problem of computing/approximating conditional K^t in the sublinear-time regime is NP-hard.

Theorem 3, Theorem 1 and Theorem 2 together give characterizations of Impagliazzo's worlds Algorithmica, Heuristica and Pessiland based on different hardness assumptions for the computation of conditional pK^t.

In particular, Theorem 1 and Theorem 2 imply that the task of ruling out Pessiland[2] is equivalent to showing that the problem of computing conditional pK^t on average over polynomial-time samplable distributions is robust with respect to the two different time regimes.

Also, we get that to rule out Pessiland, it suffices to show that it is NP-hard to compute conditional pK^t in the *polynomial-time regime* in the worst case.

▶ **Corollary 4** (Informal. See [20, Corollary 55] for the formal version). *If computing conditional pK^t in the polynomial-time regime is NP-hard, then Pessiland does not exist.*

A proof sketch of Corollary 4 can be found in [20, Section 4.4].

For comparison, it was observed in [7] that if one can show the NP-hardness of *approximating* a certain variant of time-bounded Kolmogorov complexity called q^t, then Pessiland does not exist. It is known that $\mathsf{q}^{\mathsf{poly}}$ and $\mathsf{pK}^{\mathsf{poly}}$ are equivalent to each other up to an additive logarithmic factor. This implies that showing the NP-hardness of *approximating* pK^t will allow us to rule out Pessiland.[3] It can also be shown that the problem of *approximating* pK^t is reducible to that of computing conditional pK^t.[4] On the other hand, Corollary 4 only requires showing the NP-hardness of computing conditional pK^t, which might be easier. Moreover, we note that the barrier of [22] to showing NP-hardness of approximating Kolmogorov complexity and its variants does not seem to apply directly to exact computation.

Equivalences between Average-Case Easiness of Approximating and Computing (Conditional) pK^t

By combining Theorem 1 with existing characterizations of one-way functions, we get that the average-case easiness of approximating and computing different variants of probabilistic (conditional) time-bounded Kolmogorov complexity are in fact equivalent. We state this result more formally below.

We say that "approximating pK^t is easy-on-average over samplable distributions" if the following holds.

---

[2] In this case, we mean basing infinitely-often one-way functions on $\mathsf{DistNP} \not\subseteq \mathsf{HeurBPP}$.

[3] Here, we refer to the problem called Gap-MINpKT. For a polynomial τ, $\mathsf{Gap\text{-}MINpKT}[\tau]$ is the (promise) problem of deciding, given as input $(x, 1^s, 1^t)$, whether $\mathsf{pK}^t(x) \le s$ or $\mathsf{pK}^{\tau(|x|,t)}(x) > s + \log\tau(|x|, t)$.

[4] More precisely, if we can solve $\mathsf{Cond\text{-}pK}[\tau]$ for some polynomial τ, then we can also solve $\mathsf{Gap\text{-}MINpKT}[\tau']$ for some polynomial τ'.

> For every polynomial-time samplable distribution family $\{\mathcal{D}_n\}_n$ supported over $\{0,1\}^n$, every polynomial q, and for all large enough polynomial τ, there is a probabilistic polynomial-time algorithm A that can decide, given as input $(x, 1^s, 1^t)$, whether $\mathsf{pK}^t(x) \leq s$ or $\mathsf{pK}^{\tau(|x|,t)}(x) > s + \log \tau(|x|, t)$,[5] with probability at least $1 - 1/q(n)$ over $x \sim \mathcal{D}_n$ and the internal randomness of A.

The above can be naturally generalized to the conditional setting, where we consider any samplable distribution family $\{\mathcal{D}_{\langle n,m\rangle}\}_{n,m}$ supported over $\{0,1\}^n \times \{0,1\}^m$, and for all large enough polynomial τ, we can decide whether $\mathsf{pK}^t(x \mid y) \leq s$ or $\mathsf{pK}^{\tau(|x|,|y|,t)}(x \mid y) > s + \log \tau(|x|, |y|, t)$ with high probability over (x, y) sampled from $\mathcal{D}_{\langle n,m\rangle}$. In this case, we say that "approximating conditional pK^t is easy-on-average over samplable distributions"

Also, we say that "computing pK^t is easy-on-average over samplable distributions" if the following holds.

> For every polynomial-time samplable distribution family $\{\mathcal{D}_n\}_n$ supported over $\{0,1\}^n$, every polynomial q, and for all large enough polynomial τ, there is a probabilistic polynomial-time algorithm A that can decide, given as input $(x, 1^s)$, whether $\mathsf{pK}^{\tau(|x|)}_{2/3}(x) \leq s$ or $\mathsf{pK}^{\tau(|x|)}_{1/3}(x) > s$,[6] with probability at least $1 - 1/q(n)$ over $x \sim \mathcal{D}_n$ and the internal randomness of A.

▶ **Theorem 5** (Informal). *The following are equivalent.*

1. *Infinitely-often one-way functions do not exist.*
2. *Approximating* pK^t *is easy-on-average over samplable distributions.*
3. *Approximating conditional* pK^t *is easy-on-average over samplable distributions.*
4. *Computing* pK^t *is easy-on-average over samplable distributions.*
5. *Computing conditional* pK^t *is easy-on-average over samplable distributions.*

A sketch of the proof of Theorem 5 can be found in [20, Section 3.3].

1.1.2 Characterizing Impagliazzo's Worlds by Tractability of Conditional Time-Unbounded Kolmogorov Complexity

We present a meta-complexity problem, namely approximating conditional Kolmogorov complexity up to an $O(\log n)$ additive term, that is unconditionally hard (even uncomputable) in the worst case, but such that its average-case intractability for different classes of distributions characterize Algorithmica, Heuristica and Pessiland.

Characterizing DistNP $\subseteq$ BPP and DistNP $\subseteq$ HeurBPP by Tractability of Time-Unbounded Kolmogorov Complexity

To begin, we recall a recent result by Ilango, Ren, and Santhanam [11] characterizing the non-existence of one-way functions by the tractability of approximating Kolmogorov complexity over polynomial-time samplable distributions. We consider the following conditional variant from [8].

---

[5] Note that this is the problem Gap-MINpKT mentioned in Footnote 3.
[6] This problem is referred to as $\mathsf{MpK}^\tau\mathsf{P}$ in [17].

▶ **Theorem 6** ([8, Lemma 27], cf. [11]). *The following are equivalent.*

1. *Infinitely-often one-way functions do not exist.*
2. **(Approximating conditional Kolmogorov complexity is easy-on-average over polynomial-time samplable distributions.)**
   For every polynomial-time samplable distribution family $\{\mathcal{D}_n\}_n$, where each $\mathcal{D}_n$ is over $\{0,1\}^n \times \{0,1\}^n$, and every polynomial q, there exist a probabilistic polynomial-time algorithm A and a polynomial p such that for all $n \in \mathbb{N}$,

$$\Pr_{(x,y)\sim\mathcal{D}_n}[\mathsf{K}(x \mid y) \leq A(x,y) \leq \mathsf{K}(x \mid y) + \log p(n)] \geq 1 - \frac{1}{q(n)}.$$

Note that a one-way function is a function that is efficiently computable but hard to invert on average; thus, this notion is based on *average-case* hardness. Theorem 6 characterizes the existence of one-way functions by the average-case hardness of approximating (conditional) Kolmogorov complexity. Then, for $\mathsf{NP} \not\subseteq \mathsf{BPP}$, which is a worst-case hardness notion, one might think that it can be characterized by the worst-case hardness of approximating (conditional) Kolmogorov complexity. However, it is well known that the task of approximating the conditional Kolmogorov complexity is provably intractable in the worst case, so such a characterization would imply $\mathsf{NP} \not\subseteq \mathsf{BPP}$ unconditionally.

Consider a polynomial-time samplable distribution $\mathcal{D}$ over $\{0,1\}^n \times \{0,1\}^n$. Also, let $\mathcal{D}^{(2)}$ be the marginal distribution of $\mathcal{D}$ on the second half, and let $\mathcal{D}(\cdot \mid y)$ denote the conditional distribution of $\mathcal{D}$ on the first half given that the second half is y. Now, observe the following equivalent way of sampling a pair of strings (x, y) from $\mathcal{D}$: We first sample y from $\mathcal{D}^{(2)}$ and then x from $\mathcal{D}(\cdot \mid y)$.

Note that Theorem 6 essentially says that one-way functions do not exist if and only if, for every polynomial-time samplable distribution $\mathcal{D}$, one can approximate $\mathsf{K}(x \mid y)$ on average over (x, y), where we sample y from $\mathcal{D}^{(2)}$ and then x from $\mathcal{D}(\cdot \mid y)$. In order to characterize $\mathsf{NP} \subseteq \mathsf{BPP}$, we consider the tractability of approximating conditional Kolmogorov complexity in the *semi-worst case*, meaning that we can approximate $\mathsf{K}(x \mid y)$ on average over x sampled from $\mathcal{D}(\cdot \mid y)$ for *all* $y \in \{0,1\}^n$ (instead of an average y from $\mathcal{D}^{(2)}$). Our first result is a characterization of $\mathsf{NP} \subseteq \mathsf{BPP}$ by the tractability of approximating conditional Kolmogorov complexity in this semi-worst case. Formally, we show the following.

▶ **Theorem 7.** *The following are equivalent.*

1. $\mathsf{NP} \subseteq \mathsf{BPP}$.
2. **(Approximating conditional Kolmogorov complexity is easy in the semi-worst case.)**
   For every polynomial-time samplable distribution family $\{\mathcal{D}_n\}_n$, where each $\mathcal{D}_n$ is over $\{0,1\}^n \times \{0,1\}^n$, and every polynomial q, there exist a probabilistic polynomial-time algorithm A and a polynomial p such that for all $n \in \mathbb{N}$ and $y \in \{0,1\}^n$,

$$\Pr_{x\sim\mathcal{D}_n(\cdot|y)}[\mathsf{K}(x \mid y) \leq A(x,y) \leq \mathsf{K}(x \mid y) + \log p(n)] \geq 1 - \frac{1}{q(n)}.$$

Theorem 7 shows that $\mathsf{NP} \subseteq \mathsf{BPP}$ if and only if for every polynomial-time samplable distribution $\mathcal{D}$, approximating $\mathsf{K}(x \mid y)$ is easy on average over x sampled from $\mathcal{D}(\cdot \mid y)$ for *every* $y \in \{0,1\}^n$. Now, instead of considering every $y \in \{0,1\}^n$ (a worst-case notion), it is also natural to consider an average y sampled from some polynomial-time samplable distribution $\mathcal{C}$ (an average-case notion). However, the distribution $\mathcal{C}$ here can be independent of $\mathcal{D}$. In particular, it does not necessarily have to be $\mathcal{D}^{(2)}$.

Next, we show that the average-case tractability of approximating conditional Kolmogorov complexity over such *independent polynomial-time samplable distributions*, in fact characterizes the *average-case* easiness of NP (i.e., DistNP $\subseteq$ HeurBPP). We first state formally the definition of independent polynomial-time samplable distributions.

▶ **Definition 8** (Independent Polynomial-Time Samplable [8]). *We say that a distribution family $\{\mathcal{D}_n\}_n$, where each $\mathcal{D}_n$ is over $\{0,1\}^n \times \{0,1\}^n$, is* independent polynomial-time samplable *if there exist two polynomial-time samplable distribution families $\{\mathcal{A}_n\}_n$ and $\{\mathcal{B}_n\}_n$, where each $\mathcal{A}_n$ is over $\{0,1\}^n$ and each $\mathcal{B}_n$ is over $\{0,1\}^n \times \{0,1\}^n$, such that $\mathcal{D}_n$ can be equivalently sampled as follows: sample $y \sim \mathcal{A}_n$, sample $x \sim \mathcal{B}_n(\cdot \mid y)$, and then output (x,y).*

It is easy to see that every polynomial-time samplable distribution is also independent polynomial-time samplable, by letting $\mathcal{A}$ be the marginal distribution of $\mathcal{D}$ on the second half and letting $\mathcal{B}$ be $\mathcal{D}$. However, the converse is not necessarily true. Nevertheless, Theorem 6 and Theorem 9 (which we state below) imply that the task of ruling out Pessiland is equivalent to showing that the hardness of approximating conditional Kolmogorov complexity remains unchanged over these two classes of distributions.

▶ **Theorem 9.** *The following are equivalent.*

1. DistNP $\subseteq$ HeurBPP.
2. **(Approximating conditional Kolmogorov complexity is easy-on-average over independent polynomial-time samplable distributions.)**
   For every independent polynomial-time samplable distribution family $\{\mathcal{D}_n\}_n$ and every polynomial q, there exist a probabilistic polynomial-time algorithm A and a polynomial p such that for all $n \in \mathbb{N}$,

$$\Pr_{(x,y)\sim\mathcal{D}_n}[\mathsf{K}(x \mid y) \le A(x,y) \le \mathsf{K}(x \mid y) + \log p(n)] \ge 1 - \frac{1}{q(n)}.$$

Finally, we extend Theorem 6 to characterize the non-existence of *auxiliary-input one-way functions* by the tractability of approximating conditional Kolmogorov complexity over P/poly-*samplable distributions*.

▶ **Theorem 10.** *The following are equivalent.*

1. *Auxiliary-input one-way functions do not exist.*
2. *For every sequence of strings $\{y_n\}_n$ where each $y_n \in \{0,1\}^n$, every distribution family $\{\mathcal{D}_n\}_n$ samplable in polynomial time using $\{y_n\}_n$ as advice, where each $\mathcal{D}_n$ is over $\{0,1\}^n$, and every polynomial q, there exist a probabilistic polynomial-time algorithm A and a polynomial p such that for all $n \in \mathbb{N}$,*

$$\Pr_{x\sim\mathcal{D}_n}[\mathsf{K}(x \mid y_n) \le A(x,y_n) \le \mathsf{K}(x \mid y_n) + \log p(n)] \ge 1 - \frac{1}{q(n)}.$$

The results above characterize the non-existence of one-way functions, DistNP $\subseteq$ HeurBPP, and NP $\subseteq$ BPP by the *distributional* tractability of approximating the conditional Kolmogorov complexity. They imply that the tasks of ruling out Impagliazzo's certain worlds are equivalent to showing that the hardness of this problem is the same with respect to different classes of distributions. For example, Theorem 6 and Theorem 9 imply that basing one-way functions on DistNP $\not\subseteq$ HeurBPP (a.k.a., ruling out Pessiland) is equivalent to showing that the hardness of approximating conditional Kolmogorov complexity over polynomial-time samplable distributions is the same as the hardness over independent polynomial-time samplable distributions.

Equivalences between Tractability of Time-Unbounded and Time-Bounded Kolmogorov Complexity

We first recall the definition of time-bounded Kolmogorov complexity. For $x, y \in \{0,1\}^*$ and $t \in \mathbb{N}$, we define $\mathsf{K}^t(x \mid y)$ to be the minimum length of a program $p \in \{0,1\}^*$ such that $U^y(p)$ outputs x within t steps. Here, U is a fixed time-optimal universal Turing machine and has oracle access to the string y.

For $\tau\colon \mathbb{N} \to \mathbb{N}$ and $\kappa\colon \mathbb{N} \to \mathbb{N}$, let $\mathsf{McK}^\tau\mathsf{P}[\kappa]$ be the problem where we are given input $x \in \{0,1\}^n$ and $y \in \{0,1\}^{\kappa(n)}$, and we are asked to compute $\mathsf{K}^{\tau(|x|)}(x \mid y)$. Given a polynomial τ and a polynomial κ, we say that:

- $\mathsf{McK}^\tau\mathsf{P}[\kappa]$ is *easy in the worst case* if $\mathsf{McK}^\tau\mathsf{P}[\kappa]$ can be solved in polynomial time.
- $\mathsf{McK}^\tau\mathsf{P}[\kappa]$ is *easy-on-average over polynomial-time samplable distributions* if $\mathsf{McK}^\tau\mathsf{P}[\kappa]$ admits a heuristic scheme. That is for any polynomial-time samplable distribution $\mathcal{D} = \{D_n\}_n$, where each $\mathcal{D}_n$ samples (x, z) with $x \in \{0,1\}^n$ and $y \in \{0,1\}^{\kappa(n)}$, there exists a probabilistic polynomial-time algorithm A such that for all $n, k \in \mathbb{N}$,

$$\mathop{\mathbf{Pr}}_{x,y\sim\mathcal{D}_n}\left[A(x, y; 1^n, 1^k) = \mathsf{K}^{\tau(|x|)}(x \mid y)\right] \geq 1 - \frac{1}{k}.$$

- $\mathsf{McK}^\tau\mathsf{P}[\kappa]$ is *easy-on-average over the uniform distribution* if for every polynomial p, there exists a probabilistic polynomial-time algorithm A such that for all $n \in \mathbb{N}$,

$$\mathop{\mathbf{Pr}}_{x\sim\{0,1\}^n, y\sim\{0,1\}^{\kappa(n)}}\left[A(x, y) = \mathsf{K}^{\tau(|x|)}(x \mid y)\right] \geq 1 - \frac{1}{p(n)}.$$

▶ **Theorem 11** (Implicit in [16]). *The following hold.*

- *For all polynomial $\tau(n) \geq n^2$, there exists a polynomial κ such that* $\mathsf{NP} \subseteq \mathsf{BPP}$ *if and only if* $\mathsf{McK}^\tau\mathsf{P}[\kappa]$ *is easy in the worst-case.*
- *For all polynomial $\tau(n) \geq n^2$, there exists a polynomial κ such that* $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$ *if and only if* $\mathsf{McK}^\tau\mathsf{P}[\kappa]$ *is easy-on-average over polynomial-time samplable distributions.*
- *For every polynomial $\tau(n) \geq 1.1n$ and polynomial κ, infinitely-often one-way functions do not exist if and only if* $\mathsf{McK}^\tau\mathsf{P}[\kappa]$ *is easy-on-average over the uniform distribution.*

As a corollary, we get the following equivalences between the tractability of conditional Kolmogorov complexity and that of conditional time-bounded Kolmogorov complexity.

▶ **Corollary 12** (Informal). *The following hold.*

- *For all polynomial $\tau(n) \geq n^2$, there exists a polynomial κ such that approximating conditional Kolmogorov complexity is easy in the semi-worst case if and only if* $\mathsf{McK}^\tau\mathsf{P}[\kappa]$ *is easy in the worst-case case.*
- *For all polynomial $\tau(n) \geq n^2$, there exists a polynomial κ such that approximating conditional Kolmogorov complexity is easy-on-average over independent polynomial-time samplable distributions if and only if* $\mathsf{McK}^\tau\mathsf{P}[\kappa]$ *is easy-on-average over polynomial-time samplable distributions.*
- *For every polynomial $\tau(n) \geq 1.1n$ and polynomial κ, approximating conditional Kolmogorov complexity is easy-on-average over polynomial-time samplable distributions if and only if* $\mathsf{McK}^\tau\mathsf{P}[\kappa]$ *is easy-on-average over the uniform distribution.*

Proof. This follows directly from Theorem 6, Theorem 7, Theorem 9, and Theorem 11. ◀

1.2 Techniques

In this section, we explain the main ideas behind our proofs.

Characterizing Non-Existence of One-Way Functions by Average-Case Easiness of Conditional pK^t

A recent result by Liu and Pass [17] characterized the non-existence of (infinitely-often) one-way functions by the average-case easiness of computing pK^t over polynomial-time samplable distributions. Here, we describe a proof of this result that is slightly different than the original one and show how to generalize it to *conditional* pK^t.

It will be convenient to think of the pK^t complexity of a string as its K^t complexity *conditioning on a random string* r (see [20, Proposition 17]).

First of all, by employing ideas from [14, 17], one can construct a function, which outputs the string x produced by a randomly selected (time-bounded) program (resp. conditioning on a random string r), and show that if this function can be inverted, then we can obtain a shortest program for x (resp. conditioning on r) "on average". In particular, it can be shown that if infinitely-often one-way functions do not exist, then for every time bound function $\tau(n) = n^{O(1)}$, there exists an efficient algorithm A (for simplicity, think of it as being deterministic) such that with high probability over a uniformly random string r, $A(x; r)$ computes $\mathsf{K}^\tau(x \mid r)$ for an average x sampled from some distribution $\mathcal{E}^\tau_r$, defined as $\mathcal{E}^\tau_r(x) := 2^{-\mathsf{K}^\tau(x|r)}$.

Next, we want to say that, for almost all r, the algorithm $A(-; r)$, which works for the distribution $\mathcal{E}^\tau_r$, also works for a given polynomial-time samplable distribution $\mathcal{D}$ (provided that τ is a sufficiently large polynomial). To get this, it suffices to show that $\mathcal{E}^\tau_r$ *dominates*[7] $\mathcal{D}$, i.e., $2^{-\mathsf{K}^\tau(x|r)} \gtrsim \mathcal{D}(x)$ for *every* x. The observation here is that this follows from the recently discovered coding theorem for $\mathsf{pK}^{\mathsf{poly}}$ [19], which asserts that for *every* string x, $\mathsf{pK}^\tau(x) \lesssim \log(1/\mathcal{D}(x))$ (again, provided that τ is a sufficiently large polynomial). To see this, note that by the definition of pK^t, we have for a uniform random r, $\mathsf{K}^\tau(x \mid r) \leq \mathsf{pK}^\tau(x)$.

Given the above, we have that with high probability over a uniformly random r, $A(x; r) = \mathsf{K}^\tau(x \mid r)$ for an average x sampled from $\mathcal{D}$. By an averaging argument, we get that with high probability over $x \sim \mathcal{D}$, $A(x; r) = \mathsf{K}^\tau(x \mid r)$ with high probability over a uniformly random r. For any such *good* x, if $\mathsf{pK}^\tau_{2/3}(x) \leq s$ (resp. $\mathsf{pK}^\tau_{1/3}(x) > s$), which means $\mathbf{Pr}_r[\mathsf{K}^\tau(x \mid r) \leq s] \geq 2/3$ (resp. $\mathbf{Pr}_r[\mathsf{K}^\tau(x \mid r) > s] \geq 2/3$), then $A(x, r) \leq s$ (resp. $A(x, r) > s$) with high probability over r. This allows us to solve the problem of computing pK^τ on average over the distribution $\mathcal{D}$.

Now we describe how to generalize the above to *conditional* pK^t.

Suppose we want to compute $\mathsf{pK}^\tau(x \mid y)$ over (x, y) sampled from some polynomial-time distribution $\mathcal{D}$. It will be convenient to consider the following equivalent way of sampling $\mathcal{D}$: We first sample $y \sim \mathcal{D}^{(2)}$, where $\mathcal{D}^{(2)}$ is the marginal distribution of $\mathcal{D}$ on the second half, and then sample $x \sim \mathcal{D}(\cdot \mid y)$, where $\mathcal{D}(\cdot \mid y)$ is the conditional distribution of $\mathcal{D}_n$ on the first half given that the second half is y. Finally, we output (x, y).

First of all, by modifying the construction of the candidate one-way function described above (e.g., by incorporating the distribution $\mathcal{D}^{(2)}$ into the construction), we can show that if infinitely-often one-way functions do not exist, then there exists an efficient algorithm A such that with high probability over a uniformly random string r and over y sampled from $\mathcal{D}^{(2)}$, $A(x; y, r)$ computes $\mathsf{K}^\tau(x \mid y, r)$ for an average x sampled from some distribution $\mathcal{E}^\tau_{y,r}$, where $\mathcal{E}^\tau_{y,r}(x) := 2^{-\mathsf{K}^\tau(x|y,r)}$.

[7] Recall that a distribution $\mathcal{D}$ dominates another distribution $\mathcal{D}'$ if $\mathcal{D}(x) \geq \mathcal{D}'(x)/\mathsf{poly}(n)$ for every x.

Now similar to the previous case, we want to say that, with high probability over r and $y \sim \mathcal{D}^{(2)}$, the algorithm $A(-; y, r)$, which works for the distribution $\mathcal{E}^{\tau}_{y,r}$, also works for the distribution $\mathcal{D}(\cdot \mid y)$. Again, it suffices to show that $\mathcal{E}^{\tau}_{y,r}(x) = 2^{-\mathsf{K}^{\tau}(x|y,r)} \gtrsim \mathcal{D}(x \mid y)$ for every x. However, this would require a conditional version of the coding theorem for $\mathsf{pK}^{\mathsf{poly}}$ applying to the distribution $\mathcal{D}(\cdot \mid y)$ (which is not necessarily efficiently samplable given y). Such a coding theorem is not known (in fact, is unlikely to hold).

The key observation is that in order to show that the algorithm $A(-; y, r)$, which works on average over the distribution $\mathcal{E}^{\tau}_{y,r}$, also works for $\mathcal{D}(\cdot \mid y)$, it suffices to have that $\mathcal{E}^{\tau}_{y,r}(x)$ dominates $\mathcal{D}(x \mid y)$ *on almost all* x, instead of every x. Then this weaker condition can be obtained from an *average-case* coding theorem for $\mathsf{pK}^{\mathsf{poly}}$, which has been shown under the assumption that infinitely-often one-way functions do not exist [8] (see [20, Theorem 29]).

More specifically, [8] showed that if infinitely-often one-way functions do not exist, then with high probability over $y \sim \mathcal{D}^{(2)}$ and $x \sim \mathcal{D}(\cdot \mid y)$, it holds that

$$\mathsf{pK}^{\tau}(x \mid y) \lesssim \log \frac{1}{\mathcal{D}(x \mid y)}.$$

Again, by the definition of pK^t and an averaging argument, this yields that with high probability over a uniformly random r and $y \sim \mathcal{D}^{(2)}$,

$$\mathsf{K}^{\tau}(x \mid y, r) \leq \log \frac{1}{\mathcal{D}(x \mid y)}$$

holds for almost all $x \sim \mathcal{D}(\cdot \mid y)$. This allows us to say that with high probability over r and $y \sim \mathcal{D}^{(2)}$, the distribution $\mathcal{E}^{\tau}_{y,r}$ dominates $\mathcal{D}(\cdot \mid y)$ *on average*, so the algorithm $A(-; y, r)$, which works for $\mathcal{E}^{\tau}_{y,r}$, also works for $\mathcal{D}(\cdot \mid y)$.

At this point, we get that with high probability over $(x, y) \sim \mathcal{D}$ and over a uniformly random r, $A(x; y, r) = \mathsf{K}^{\tau}(x \mid y, r)$. By the same argument as described above, this allows us to compute $\mathsf{pK}^{\tau}(x \mid y)$ on average over (x, y) sampled from $\mathcal{D}$.

The converse direction, i.e., that computing conditional pK^t on average allows us to break one-way functions, follows from the standard observation that computing pK^t on average over samplable distributions allows us to distinguish pseudo-random distributions (which are supported on strings of low pK^t complexity) from random strings (which have high pK^t complexity).

Characterizing DistNP $\subseteq$ HeurBPP by Average-Case Easiness of Conditional pK^t in Sublinear-Time Regime

To show that the average-case easiness of computing conditional pK^t (in the sublinear-time regime) implies the average-case easiness of NP (both with respect to polynomial-time samplable distributions), we first show that it is NP-hard to compute conditional pK^t (again, in the sublinear-time regime). Recently, Liu and Pass [16] and Hirahara [6] showed that the problem of computing the conditional K^t in the sublinear-time regime is NP-hard. We generalize this result to pK^t.

At a high level, our proof follows a similar approach but also requires some crucial observations to address the more complex notion of pK^t and to make it applicable to show Theorem 2. In particular, we adapt the proof in [6] which relies on the use of a *secret sharing scheme* (see [6, Section 2.3] for an exposition). More specifically, it reduces the problem of approximating the hamming weight of a *minimum satisfying assignment* of a given monotone formula, which is known to be NP-hard, to that of computing conditional K^t in the sublinear-time regime. That is, for every constant $c > 1$ and time bound function

$\tau(n,m) := n^c \cdot m^{1-1/c}$, there is a randomized reduction R such that if a given monotone formula ψ has a satisfying assignment of hamming weight at most ζ (resp. much larger than ζ), then with high probability, R produces a pair of strings (x,y) and ρ such that $\mathsf{K}^{\tau(|x|,|y|)}(x \mid y) \leq \rho$ (resp. $\mathsf{K}^{\tau(|x|,|y|)}(x \mid y) > \rho$).

Our key observation is that this reduction still works in the presence of any fixed string r. Roughly put, the reason for this is that a secret sharing scheme remains secure even if an adversary has access to some fixed string. More specifically, we can show that with respect to any string r, if a given monotone formula ψ has a satisfying assignment of hamming weight much larger than ζ, then with high probability the algorithm R produces a pair of strings (x,y) and ρ such that $\mathsf{K}^{\tau(|x|,|y|)}(x \mid y,r) > \rho$. This allows us to say that if the minimum weight of ψ is much larger than ζ, then with high probability over a random string r and over the internal randomness of R, $\mathsf{K}^{\tau(|x|,|y|)}(x \mid y,r) > \rho$. By an averaging argument, this gives that with high probability over the internal randomness of R, $\mathsf{K}^{\tau(|x|,|y|)}(x \mid y,r) > \rho$ for more than $2/3$ of the r's, which essentially means $\mathsf{pK}^{\tau(|x|,|y|)}_{1/3}(x \mid y) > \rho$.

Now we have showed that computing conditional pK^t (in the sublinear-time regime) is NP-hard. To solve an NP problem L over a given polynomial-time samplable distribution $\mathcal{D}$, we can compose $\mathcal{D}$ with the reduction R to obtain a new distribution $\mathcal{D}'$. Then we can show that computing conditional pK^t on average over $\mathcal{D}'$ will allow us to solve L on average over $\mathcal{D}$. However, there is an additional subtle issue here, the original reduction R depends on the time bound function (i.e., for every sublinear time bound τ, there is a reduction R that will work). On the other hand, to show Theorem 2 (Item 2 $\implies$ Item 1), it is required that the reduction works for all time bound functions τ of the form $\tau(n,m) = n^c \cdot m^{1-1/c}$. We will then need to further modify the reduction to achieve this. (See [20, Lemma 45] for the details.)

Now we need to show the other direction saying that the average-case easiness of NP implies the average-case easiness of computing conditional pK^t. Unlike the problem of computing (conditional) K^t, computing (conditional) pK^t is not known to be in NP, so we can not get the desired implication directly. However, it is not hard to see that the problem of computing conditional pK^t is in fact in (promise) AM.[8] If we can solve NP, then we can also solve AM (in the randomized setting), by a standard trick that combines the instance of an AM problem with a random string to produce an instance for an NP problem. (See [20, Lemma 53] for the details.)

Characterizing DistNP $\subseteq$ BPP and DistNP $\subseteq$ HeurBPP by Tractability of Time-Unbounded Kolmogorov Complexity.

First, we recap the proof of Theorem 6 as presented in [11]. We will ignore the issue of "infinitely often" in this subsection.

To show that the non-existence of one-way functions implies efficient algorithms for approximating conditional Kolmogorov complexity on average over polynomial-time samplable distributions, we use a powerful result from [13], which asserts that if one-way functions do not exist, then for any polynomial-time samplable distribution $\mathcal{D}$ over $\{0,1\}^n \times \{0,1\}^n$, one can efficiently estimate $\mathcal{D}(x \mid y)$ on average over $(x,y) \sim \mathcal{D}$. In addition, we combine two fundamental properties related to time-unbounded Kolmogorov complexity: The first is called the coding theorem, which roughly says that for every $(x,y) \in \mathsf{Support}(\mathcal{D})$,

[8] Here, we refer to the problem $\mathsf{Cond\text{-}pK}$ instead of the one that asks to decide whether $\mathsf{pK}^{\tau(|x|,|y|)}(x \mid y) \leq s$ for a given input $(x,y,1^s)$ and time bound τ.

$$\mathsf{K}(x \mid y) \lesssim \log \frac{1}{\mathcal{D}(x \mid y)},$$

and the second is the incompressibility property, which states that all $y \in \{0,1\}^n$ and for almost all $x \sim \mathcal{D}(\cdot \mid y)$,

$$\mathsf{K}(x \mid y) \gtrsim \log \frac{1}{\mathcal{D}(x \mid y)}.$$

It follows that for almost all $(x, y) \sim \mathcal{D}$,

$$\mathsf{K}(x \mid y) \approx \log \frac{1}{\mathcal{D}(x \mid y)}.$$

This allows us to approximate $\mathsf{K}(x \mid y)$ by estimating $\mathcal{D}(x \mid y)$, and the latter can be done efficiently if one-way functions do not exist.

For the other direction, the idea is that an efficient algorithm for approximating Kolmogorov complexity on average can be used to construct a function that distinguishes the output distribution of a cryptographic pseudorandom generator from the uniform distribution. Intuitively, this is because the outputs of such a generator have low $\mathsf{K}^{\mathsf{poly}}$ complexity while a random string has high Kolmogorov complexity. Then such an algorithm implies the non-existence of pseudorandom generators and hence of one-way functions [4].

Now, let us try to see if we can adapt the above proof paradigm to show Theorem 9, which characterizes $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$ by the tractability of approximating conditional Kolmogorov complexity on average over *independent polynomial-time samplable distributions*.

One direction is in fact easy by using tools developed in [8]. In particular, it is observed in [8] that if $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$, then every independent polynomial-time samplable distribution can be simulated by some polynomial-time samplable distribution (see [20, Lemma 26]). Consequently, if $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$ (which also implies that one-way functions do not exist), then we can reduce the task of approximating conditional Kolmogorov complexity over independent polynomial-time samplable distributions to that of approximating conditional Kolmogorov complexity over polynomial-time samplable distributions, which is tractable if one-way functions do not exist.

However, for the other direction, it is unclear how we can get $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$ from the tractability of approximating conditional Kolmogorov complexity over independent polynomial-time samplable distributions, by using ideas from the proof of the characterization for one-way functions. In that scenario, we use the algorithm for approximating conditional Kolmogorov complexity as a distinguisher to break the security of a cryptographic pseudorandom generator.

Here, we will use a different approach. Specifically, we rely on a recently discovered characterization of $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$ by the validity of a certain property called *conditional coding* for pK^t. More precisely, the authors of [8] showed that $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$ if and only if conditional coding property for $\mathsf{pK}^{\mathsf{poly}}$ holds on average over pairs of strings drawn from independent polynomial-time samplable distributions, i.e., for any independent polynomial-time samplable distribution $\mathcal{D}$ over $\{0,1\}^n \times \{0,1\}^n$ and for almost all $(x, y) \sim \mathcal{D}$,

$$\mathsf{pK}^{\mathsf{poly}(n)}(x \mid y) \lesssim \log \frac{1}{\mathcal{D}(x \mid y)}$$

(see [20, Theorem 30]).

Now given this characterization of $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$ using conditional coding, it suffices to show that conditional coding property for $\mathsf{pK}^{\mathsf{poly}}$ over independent polynomial-time samplable distributions follows from the tractability of approximating conditional Kolmogorov complexity over the same class of distributions.

How can we show this? First of all, note that by the coding theorem for *time-unbounded* Kolmogorov complexity, we have that for every $(x, y) \in \mathsf{Support}(\mathcal{D})$,

$$\mathsf{K}(x \mid y) \lesssim \log \frac{1}{\mathcal{D}(x \mid y)}.$$

Then to get the desired conditional coding property for $\mathsf{pK}^{\mathsf{poly}}$, it suffices to show that for almost all $(x, y) \sim \mathcal{D}$,

$$\mathsf{pK}^{\mathsf{poly}(n)}(x \mid y) \leq \mathsf{K}(x \mid y) + O(\log n). \tag{1}$$

Now, let us describe how to show the above, assuming efficient algorithms for approximating conditional Kolmogorov complexity over independent polynomial-time samplable distributions.

The key ingredient here is a pseudorandom generator construction with reconstruction property. Such a generator is instantiated with a target string, it then takes as input a random seed and outputs a string that is longer than the seed. The reconstruction property allows us to say that if there exists a function that can distinguish the output distribution of the generator from the uniform distribution, then it can be used to recover the target string, using an additional advice string. This enables us to say that given a distinguisher, the target string has poly-time-bounded Kolmogorov complexity bounded by the length of the advice string. An algorithm for approximating Kolmogorov complexity can naturally be used as such a distinguisher, since the outputs of the generator have low Kolmogorov complexity while a random string has high Kolmogorov complexity. By appropriately configuring the parameters of the generator, we can ensure that the length of the advice string is comparable to the Kolmogorov complexity of the target string. This allows us to upper bound the poly-time-bounded Kolmogorov complexity of the target string by its Kolmogorov complexity.

Using this approach, the authors of [8] showed that if efficient algorithms exist for approximating conditional Kolmogorov complexity over polynomial-time samplable distributions, then for every polynomial-time samplable distribution $\mathcal{D}$ over $\{0,1\}^n \times \{0,1\}^n$ and almost all $(x, y) \sim \mathcal{D}$,

$$\mathsf{rK}^{\mathsf{poly}(n)}(x \mid y) \leq \mathsf{K}(x \mid y) + O(\log^3 n). \tag{2}$$

Here, rK^t is a certain randomized variant of time-bounded Kolmogorov complexity measure [2, 18].

The $O(\log^3 n)$ additive term in Equation (2) results from the use of a specific pseudorandom generator construction with an rK^t-style reconstruction property (as they need to upper bound $\mathsf{rK}^{\mathsf{poly}}$ by K), and such a generator has sub-optimal "advice complexity" in its reconstruction. In our case, we need to upper bound $\mathsf{pK}^{\mathsf{poly}}$ by K, and we can use a different pseudorandom generator construction with a pK^t-style reconstruction property that is known to have optimal "advice complexity" (see [20, Section 2.7]). This results in only an $O(\log n)$ additive term instead of $O(\log^3 n)$ as in the previous case.

The description provided above does not address an important technical distinction between showing Equation (1) and showing Equation (2) in [8]. In our case, we need to show Equation (1) over *independent polynomial-time samplable distributions*, whereas the

other case involves the simpler class of *polynomial-time samplable samplable distributions*. In fact, in the proof of Equation (2), a crucial fact used is that the uniform mixture of two polynomial-time samplable distributions is also polynomial-time samplable. Intuitively, the reason why this is needed is that we need to obtain a function that can distinguish the output distribution of a pseudorandom generator (induced by a polynomial-time samplable distribution) and the uniform distribution (also combined with a polynomial-time samplable distribution), so we need to apply an algorithm to approximate Kolmogorov complexity over the mixture uniform of those two distributions.

However, in our case, we are dealing with independent polynomial-time samplable distributions, and the uniform mixture of two independent polynomial-time samplable distributions is not necessarily independent polynomial-time samplable. The key insight here is that we don't really need to be concerned with the uniform mixture of two *generic* independently polynomial-time samplable distributions. Instead, the two distributions have the property that they are identical when restricted to the second half. We then show that the uniform mixture of such two distributions remains independently polynomial-time samplable. (See the proofs of [20, Lemma 59] and [20, Lemma 63] for details.)

We now describe the proof of Theorem 7. Again, the direction of showing the tractability of approximating conditional Kolmogorov complexity in the semi-worst case from $\mathsf{NP} \subseteq \mathsf{BPP}$ can be done in a way similar to that of Theorem 6 (as described earlier in this subsection). This is because if $\mathsf{NP} \subseteq \mathsf{BPP}$, then one can estimate $\mathcal{D}(x \mid y)$ for every polynomial-time samplable distribution $\mathcal{D}$ and $(x, y) \in \mathsf{Support}(\mathcal{D})$, a result due to [24] (see also [20, Lemma 27]).

For the other direction, we will employ the same approach as used to to show Theorem 9. In this case, we will use a similar characterization of $\mathsf{NP} \subseteq \mathsf{BPP}$ through conditional coding. Specifically, it has been shown in [8] that $\mathsf{NP} \subseteq \mathsf{BPP}$ if and only if *worst-case* conditional coding for $\mathsf{pK}^{\mathsf{poly}}$ holds, i.e., for every polynomial-time samplable distribution $\mathcal{D}$ over $\{0,1\}^n \times \{0,1\}^n$ and *every* $(x, y) \in \mathsf{Support}(D)$,

$$\mathsf{pK}^{\mathsf{poly}(n)}(x \mid y) \lesssim \log \frac{1}{\mathcal{D}(x \mid y)}. \tag{3}$$

Unfortunately, it is unclear how we can obtain the above worst-case conditional coding property from the tractability of approximating conditional Kolmogorov complexity in the *semi-worst case* by following the same approach. To overcome this, we observe that we can modify the original proof in [8] to obtain a characterization of $\mathsf{NP} \subseteq \mathsf{BPP}$ by *semi-worst-case* conditional coding, which only requires Equation (3) to hold for almost all $x \sim \mathcal{D}(\cdot \mid y)$ and for all $y \in \{0,1\}^n$ (see [20, Lemma 64]).

By using this alternative characterization and addressing a similar issue that arises when transitioning from polynomial-time samplable distributions to semi-worst-case input distributions, as described above in the case of showing Theorem 9, we can now use efficient algorithms for approximating conditional Kolmogorov complexity in the semi-worst case to obtain the desired semi-worst-case conditional coding property, which then yields $\mathsf{NP} \subseteq \mathsf{BPP}$.

1.3 Open Problems

Can we show NP-hardness of computing conditional pK^t in the polynomial-time regime? By Corollary 4, this would imply that Pessiland does not exist. Are there any barriers to showing such an NP-hardness result?

Theorem 9 characterizes the *error-prone* average-case easiness of NP (i.e., $\mathsf{DistNP} \subseteq \mathsf{HeurBPP}$) by the tractability of approximating conditional Kolmogorov complexity over independent polynomial-time samplable distributions. Can we obtain a similar characterization for the *errorless* average-case easiness of NP (i.e., $\mathsf{DistNP} \subseteq \mathsf{AvgBPP}$)?

References

1 Eric Allender, Mahdi Cheraghchi, Dimitrios Myrisiotis, Harsha Tirumala, and Ilya Volkovich. One-way functions and a conditional variant of MKTP. In *Conference on Foundations of Software Technology and Theoretical Computer Science* (FSTTCS), pages 7:1–7:19, 2021. doi:10.4230/LIPIcs.FSTTCS.2021.7.

2 Harry Buhrman, Troy Lee, and Dieter van Melkebeek. Language compression and pseudorandom generators. *Comput. Complex.*, 14(3):228–255, 2005. doi:10.1007/s00037-005-0199-5.

3 Halley Goldberg, Valentine Kabanets, Zhenjian Lu, and Igor C. Oliveira. Probabilistic Kolmogorov complexity with applications to average-case complexity. In *Computational Complexity Conference* (CCC), pages 16:1–16:60, 2022. doi:10.4230/LIPIcs.CCC.2022.16.

4 Johan Håstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999. doi:10.1137/S0097539793244708.

5 Shuichi Hirahara. Characterizing average-case complexity of PH by worst-case meta-complexity. In *Symposium on Foundations of Computer Science* (FOCS), pages 50–60, 2020. doi:10.1109/FOCS46700.2020.00014.

6 Shuichi Hirahara. Symmetry of information from meta-complexity. In *Computational Complexity Conference* (CCC), pages 26:1–26:41, 2022. doi:10.4230/LIPIcs.CCC.2022.26.

7 Shuichi Hirahara. Capturing one-way functions via NP-hardness of meta-complexity. In *Symposium on Theory of Computing* (STOC), pages 1027–1038, 2023. doi:10.1145/3564246.3585130.

8 Shuichi Hirahara, Rahul Ilango, Zhenjian Lu, Mikito Nanashima, and Igor C. Oliveira. A duality between one-way functions and average-case symmetry of information. In *Symposium on Theory of Computing* (STOC), pages 1039–1050, 2023. doi:10.1145/3564246.3585138.

9 Shuichi Hirahara and Rahul Santhanam. On the average-case complexity of MCSP and its variants. In *Computational Complexity Conference* (CCC), pages 7:1–7:20, 2017. doi:10.4230/LIPIcs.CCC.2017.7.

10 Shuichi Hirahara and Rahul Santhanam. Excluding PH pessiland. In *Innovations in Theoretical Computer Science Conference* (ITCS), pages 85:1–85:25, 2022. doi:10.4230/LIPIcs.ITCS.2022.85.

11 Rahul Ilango, Hanlin Ren, and Rahul Santhanam. Robustness of average-case meta-complexity via pseudorandomness. In *Symposium on Theory of Computing* (STOC), pages 1575–1583, 2022. doi:10.1145/3519935.3520051.

12 Russell Impagliazzo. A personal view of average-case complexity. In *Proceedings of the Tenth Annual Structure in Complexity Theory Conference*, pages 134–147, 1995. doi:10.1109/SCT.1995.514853.

13 Russell Impagliazzo and Leonid A. Levin. No better ways to generate hard NP instances than picking uniformly at random. In *Symposium on Theory of Computing* (STOC), pages 812–821, 1990. doi:10.1109/FSCS.1990.89604.

14 Yanyi Liu and Rafael Pass. On one-way functions and Kolmogorov complexity. In *Symposium on Foundations of Computer Science* (FOCS), pages 1243–1254, 2020. doi:10.1109/FOCS46700.2020.00118.

15 Yanyi Liu and Rafael Pass. On the possibility of basing cryptography on EXP≠BPP. In *International Cryptology Conference* (CRYPTO), pages 11–40, 2021. doi:10.1007/978-3-030-84242-0_2.

16 Yanyi Liu and Rafael Pass. On one-way functions from NP-complete problems. In *Conference on Computational Complexity* (CCC), pages 36:1–36:24, 2022. doi:10.4230/LIPIcs.CCC.2022.36.

17 Yanyi Liu and Rafael Pass. One-way functions and the hardness of (probabilistic) time-bounded Kolmogorov complexity w.r.t. samplable distributions. In *Annual Cryptology Conference* (CRYPTO), pages 645–673, 2023. doi:10.1007/978-3-031-38545-2_21.

18 Zhenjian Lu, Igor C. Oliveira, and Rahul Santhanam. Pseudodeterministic algorithms and the structure of probabilistic time. In *Symposium on Theory of Computing* (STOC), pages 303–316, 2021. doi:10.1145/3406325.3451085.

19 Zhenjian Lu, Igor C. Oliveira, and Marius Zimand. Optimal coding theorems in time-bounded Kolmogorov complexity. In *International Colloquium on Automata, Languages, and Programming* (ICALP), pages 92:1–92:14, 2022. doi:10.4230/LIPICS.ICALP.2022.92.

20 Zhenjian Lu and Rahul Santhanam. Impagliazzo's worlds through the lens of conditional Kolmogorov complexity. *Electronic Colloquium on Computational Complexity* (ECCC), TR24-085, 2024. URL: https://eccc.weizmann.ac.il/report/2024/085.

21 Hanlin Ren and Rahul Santhanam. Hardness of KT characterizes parallel cryptography. In *Computational Complexity Conference* (CCC), pages 35:1–35:58, 2021. doi:10.4230/LIPIcs.CCC.2021.35.

22 Michael E. Saks and Rahul Santhanam. On randomized reductions to the random strings. In *Computational Complexity Conference* (CCC), pages 29:1–29:30, 2022. doi:10.4230/LIPIcs.CCC.2022.29.

23 Rahul Santhanam. Pseudorandomness and the minimum circuit size problem. In *Innovations in Theoretical Computer Science Conference* (ITCS), pages 68:1–68:26, 2020. doi:10.4230/LIPIcs.ITCS.2020.68.

24 Larry J. Stockmeyer. On approximation algorithms for #P. *SIAM J. Comput.*, 14(4):849–861, 1985. doi:10.1137/0214060.