

On defining PPT-search problems

Oded Goldreich

Department of Computer Science
Weizmann Institute of Science, Rehovot, ISRAEL.

January 17, 2025

Abstract

We propose a new definition of the class of search problems that correspond to BPP . Specifically, a problem in this class is specified by a polynomial-time approximable function $q : \{0, 1\}^* \times \{0, 1\}^* \rightarrow [0, 1]$ that associates, with each possible solution y to an instance x , a quality $q(x, y)$. Intuitively, quality 1 corresponds to perfectly valid solutions, quality 0 corresponds to perfectly invalid solutions, but the quality of other solutions can be anywhere in-between. The class of PPT-search problems contains q if there exists a PPT algorithm that, on input x , finds a solution with value close to $\max_y \{q(x, y)\}$.

We relate this definition to previously studied definitions of “BPP-search problems” and articulate our preference for it. More importantly, we show that any PPT-search problem can be reduced in deterministic polynomial-time to a promise problem in prBPP (i.e., promise-BPP).

A preliminary version of this paper has been posted as TR24-161 of *ECCC*. The current revision is more detailed and more opinionated.

On the companion paper. This paper represents my perspective on a joint project conducted together with Roei Tell, whereas Roei’s perspective is presented in [6]. The technical contents of both papers is similar and both papers present the foregoing definition of PPT-search problems (i.e., Definition 2.4), but the perspectives presented in the two papers differ.

... it is nevertheless a sacred duty to prefer the truth to one’s friends.
[Aristotle, *Nicomachean Ethics*, Book I, Chapter 6, 1st paragraph]

Contents

1	Introduction	1
2	Definitions	1
2.1	Towards Defining PPT-search: Background	2
2.2	Towards Defining PPT-search: Alternatives	4
2.3	In Favor of Definition 2.4	5
2.4	Relating the Definitions	5
3	Discussion	6
3.1	Specifying Sets and Decisional Problems	7
3.2	Specifying Search Problems	8
3.3	Defining Natural Complexity Classes of Search Problems	9
3.4	Relating Complexity Classes of Different Types	10
4	Proofs	11
4.1	Proof of Theorem 2.5	11
4.2	On the Failure of Other Versions of Theorem 2.5	13
4.3	Emulations	15
5	Conclusion	16

1 Introduction

While search problems are of great intuitive importance, most research in complexity theory is focused on decision problems. This phenomenon is justified by the fact that the study of search problems seems more cumbersome and that in many cases this study can be reduced to the study of decision problems. Unfortunately, the latter phenomenon does not always hold, and in such cases one is forced to deal directly with search problems.

One such case, which is the actual focus of the current text, is the case of search problems that are solvable by probabilistic polynomial-time (PPT) algorithms. Such problems were the focus of [4], which provided a definition of a class called “BPP-search” (see [4, Def. 3.2]). We view this definition as a tentative proposal, and will propose an alternative that we consider much better. We also propose to change the terminology, replacing the term BPP-search by PPT-search, which we consider more appropriate.

Organization: In Section 2, we present four definitions of the notion of PPT-search problems: In Section 2.1 we present two definitions that appeared in [4, Sec. 3], whereas in Section 2.2 we present two new alternatives. In Sections 2.3 and 3 we explain why we prefer Definition 2.4 over the other alternatives: Section 2.3 focuses on the specific issues at hand, whereas Section 3 provides a wider perspective. Some results about the various definitions are stated in Sections 2.3 and 2.4 and proved in Section 4. We conclude by restating our perspective on the various definitions (Section 5).

2 Definitions

We shall first review the definitions proposed in [4], and then make our own proposals and explore their consequences. A more extensive discussion of the principles underlying these definitions appears in Section 3.

2.1 Towards Defining PPT-search: Background

Recall that search problems are typically associated with binary relations. Specifically, the search problem associated with the binary relation $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$ consists of finding a solution y to a given instance x , where y is considered a valid solution to x if $(x, y) \in R$. In other words, letting $R(x) \stackrel{\text{def}}{=} \{y : (x, y) \in R\}$, given an instance x such that $R(x) \neq \emptyset$, the task is find an element of $R(x)$.

Turning to PPT-search, the naive proposal is to say that a binary relation $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$ constitutes a PPT-search problem if there exists a probabilistic polynomial-time (PPT) algorithm A such that $\Pr[A(x) \in R(x)] \geq 2/3$ for every x such that $R(x) \neq \emptyset$. One drawback of this naive definition is that it does not support *error reduction*. This is a consequence of the fact that this definition does *not* require that valid solutions can be easily recognized as valid. We consider the latter fact to be a fundamental drawback of the naive definition.

There is an easy fix for the foregoing drawbacks: Augmenting the naive definition so that the missing feature is made an explicit condition. This leads to [4, Def. 3.1], which is reproduced next.

Definition 2.1 (PPT-search problem, Type 1, reproducing [4, Def. 3.1]):¹ *The binary relation $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$ constitutes a PPT-search problem (of Type 1) if the following two conditions hold.*

1. *R satisfies the naive definition (i.e., efficiently solving R): There exists a PPT algorithm A such that for every $x \in \{0, 1\}^*$ if $R(x) \neq \emptyset$, then $\Pr[A(x) \in R(x)] \geq 2/3$.*
2. *Efficient recognition of valid solutions (i.e., R as a set is in \mathcal{BPP}): There exists a PPT algorithm D such that for every $x, y \in \{0, 1\}^*$ if $y \in R(x)$ then $\Pr[D(x, y) = 1] \geq 2/3$ and otherwise $\Pr[D(x, y) = 1] \leq 1/3$.*

We say that a probabilistic mechanism, M , solves (T1-solves) the search problem R if M satisfies the first condition; that is, $\Pr[M(x) \in R(x)] \geq 2/3$ for every $x \in \{0, 1\}^$ such that $R(x) \neq \emptyset$.*

The fact that the distinguishing algorithm D is allowed to be probabilistic is natural within the current mind-frame in which PPT is viewed as the yardstick for efficient computation. (Indeed, a more restricted class of PPT-search problems would arise if we were to require that D is deterministic; but, as will become clear shortly, we wish to enlarge the class of PPT-search problems rather than make it smaller.)

Before proceeding, let us note some advantages offered by Definition 2.1. First, it allows for error reduction (in Condition 1), since one can repeatedly invoke algorithm A on input x (till obtaining a valid solution (per Condition 2)). Second, the current formulation, which is actually a “candid search problem” (see [3, Def. 2.30]), is computationally equivalent to one that also requires that if $R(x) = \emptyset$, then $\Pr[A(x) = \perp] \geq 2/3$, where $\perp \notin \{0, 1\}^*$ is a special failure symbol. Furthermore, these PPT-search problems can be solved by a deterministic reduction to $\text{pr}\mathcal{BPP}$ (see [4, Thm. 3.5]).² Hence, the study of this class of search problems is reduced to the study of corresponding decision problems.

¹The notion of T1-solving is not stated explicitly in [4, Def. 3.1].

²Recall that, while \mathcal{BPP} denotes the class of (pure) decision problems that can be decided by a PPT algorithm, $\text{pr}\mathcal{BPP}$ denotes the corresponding class of decisional *promise* problems. The foregoing reduction suffices for relating the question of whether every PPT-search problem has a pseudodeterministic algorithm [2] to the question of whether every problem in $\text{pr}\mathcal{BPP}$ can be extended to a problem in \mathcal{BPP} (see the first two pages of [5, Sec. 4]).

Unfortunately, many natural approximation problems that can be solved by PPT algorithms, including the archetypical problem of estimating the acceptance probability of a Boolean circuit, can not be captured by Definition 2.1. Typically, in these cases, a perfect solution is hard to find and recognize, whereas one can easily find and recognize “good enough” solutions.³ The hardness of evaluating the exact quality of solutions leads to relaxing the requirement such that one is only required to distinguish good solutions from bad ones, whereas there is a gap between being good and being bad. This is captured by the formalism of promise problems.

While promise problems have appeared in the penultimate paragraph (as candid search problems (and also in the class prBPP)), in the case of (candid) search problems the promise was on the instance x (in Condition 1) only (i.e., $R(x) \neq \emptyset$). This formulation falls within the framework of [3, Def. 2.29], and its conceptual meaning is clear (i.e., we totally discard instances that violate the promise and do not care what the algorithm does with them). In contrast, allowing (as in [4, Def. 3.2] (reproduced below)) a promise on the set of instance-solution pairs leads to a tri-partition of the set of possible solutions into valid, invalid, and in-between solutions. This relaxes the distinguishing task (i.e., Condition 2), which only refers to valid and invalid solutions, but raises the question of how to treat the in-between solutions in the solving task (i.e., Condition 1). The conservative approach towards the solving task, which was taken by [4, Def. 3.2], is to ignore the in-between solutions in the solving task (i.e., insists that solving means providing a valid solution).

Definition 2.2 (PPT-search problem, Type 2, reproducing [4, Def. 3.2]):⁴ *Let $R_{\text{yes}}, R_{\text{no}} \subseteq \{0, 1\}^* \times \{0, 1\}^*$ be two disjoint binary relations. The pair $(R_{\text{yes}}, R_{\text{no}})$ constitutes a PPT-search problem (of Type 2) if the following two conditions hold.*

1. *Efficiently solving R_{yes} : There exists a PPT algorithm A such that for every $x \in \{0, 1\}^*$ if $R_{\text{yes}}(x) \neq \emptyset$, then $\Pr[A(x) \in R_{\text{yes}}(x)] \geq 2/3$.*
2. *Efficient distinguishing between R_{yes} and R_{no} (i.e., $(R_{\text{yes}}, R_{\text{no}})$ as a decisional promise problem is in prBPP): There exists a PPT algorithm D such that for every $x, y \in \{0, 1\}^*$ if $y \in R_{\text{yes}}(x)$ then $\Pr[D(x, y) = 1] \geq 2/3$, whereas if $y \in R_{\text{no}}(x)$ then $\Pr[D(x, y) = 1] \leq 1/3$.*

We say that a probabilistic mechanism, M , solves (T2-solves) the search problem $(R_{\text{yes}}, R_{\text{no}})$ if M satisfies the first condition; that is, $\Pr[M(x) \in R_{\text{yes}}(x)] \geq 2/3$ for every $x \in \{0, 1\}^$ such that $R_{\text{yes}}(x) \neq \emptyset$.*

The elements of $R_{\text{yes}}(x)$ are viewed as **valid solutions** (for x), the elements of $R_{\text{no}}(x)$ are viewed as **invalid solutions** (for x), whereas the rest are in *limbo* (and are hereafter referred to as **limbo solutions**). Note that Condition 1 treats the limbo solutions as invalid, whereas Condition 2 allows to treat them as “wild cards” (i.e., allows viewing each of them as either valid or invalid, at whim).⁵ This discrepancy does not allow for error reduction in Condition 1, nor does it provide for the recognition of instances that have no valid solution. Furthermore, only an (unspecified) subset of these PPT-search problems can be solved by a deterministic reduction to prBPP (see [4, Thm. 3.5], which refers to search problems having a “companion” that is a BPP-search problem).

³In retrospect, this perception of the quality of solutions fits Definition 2.4 more than Definition 2.2.

⁴The notion of T2-solving is not stated explicitly in [4, Def. 3.2]. Evidently, PPT-search problems of Type 1 are a special case of PPT-search problems of Type 2, since the binary relation R may be viewed as the disjoint pair $(R, U \setminus R)$, where $U = \{0, 1\}^* \times \{0, 1\}^*$.

⁵Actually, Condition 2 allows to treat each such solution as neither valid nor invalid, and accept it with arbitrary probability.

2.2 Towards Defining PPT-search: Alternatives

The first alternative we consider is streamlining the two conditions in Definition 2.2. Specifically, given that Condition 2 treats limbo solutions as “wild cards” (i.e., allows viewing them either valid or invalid, at whim), we allow Condition 1 to do the same. Hence, the alternative definition is as follows.

Definition 2.3 (PPT-search problem, Type 3): *A pair $(R_{\text{yes}}, R_{\text{no}})$ as in Definition 2.2 constitutes a PPT-search problem (of Type 3) if the following two conditions hold.*

1. *A relaxed solving condition: There exists a PPT algorithm A such that for every $x \in \{0, 1\}^*$ if $R_{\text{yes}}(x) \neq \emptyset$, then $\Pr[A(x) \in (\{0, 1\}^* \setminus R_{\text{no}}(x))] \geq 2/3$.*
2. *Efficient distinguishing between R_{yes} and R_{no} (as in Definition 2.2). There exists a PPT algorithm D such that for every $x, y \in \{0, 1\}^*$ if $y \in R_{\text{yes}}(x)$ then $\Pr[D(x, y) = 1] \geq 2/3$, whereas if $y \in R_{\text{no}}(x)$ then $\Pr[D(x, y) = 1] \leq 1/3$.*

We say that a probabilistic mechanism, M , solves (T3-solves) the search problem $(R_{\text{yes}}, R_{\text{no}})$ if M satisfies the first condition; that is, $\Pr[M(x) \in (\{0, 1\}^ \setminus R_{\text{no}}(x))] \geq 2/3$ for every $x \in \{0, 1\}^*$ such that $R_{\text{yes}}(x) \neq \emptyset$.*

Note that PPT-search problems of Type 2 are a special case of PPT-search problems of Type 3, since satisfying Definition 2.2 implies satisfying Definition 2.3; that is, if an algorithm T2-solves $(R_{\text{yes}}, R_{\text{no}})$, then it T3-solves $(R_{\text{yes}}, R_{\text{no}})$.

Definition 2.3 enjoys greater internal consistency than Definition 2.2, yet it still leaves us wondering why is it OK to “accept” solutions that cannot be certified as valid (but can be certified as not being invalid). Furthermore (like Definition 2.2), Definition 2.3 does not support error reduction in Condition 1.

One attempt to justify the foregoing convention (i.e., the “acceptability” of limbo solutions) is saying that limbo solutions have a quality that is in-between valid and invalid. This qualitative perspective, which reflects the aforementioned trichotomy, leads to asking why not allow several levels of intermediate quality. Actually, why not view quality in a quantitative manner; that is, consider solutions of arbitrary quality ranging, say, in $[0, 1]$. In fact, such a quantitative interpretation of the quality of potential solutions is more natural than the trichotomy underlying Definitions 2.2 and 2.3, and it leads us to the second alternative.

The second alternative consists of replacing the two disjoint binary relations $R_{\text{yes}}, R_{\text{no}} \subseteq \{0, 1\}^* \times \{0, 1\}^*$ by a quality function $q : \{0, 1\}^* \times \{0, 1\}^* \rightarrow [0, 1]$ and defining solving the instance x as finding y that almost maximizes $q(x, y)$, where the function q is easy to approximate. Specifically, we propose the following definition.

Definition 2.4 (PPT-search problem, Type 4): *The function $q : \{0, 1\}^* \times \{0, 1\}^* \rightarrow [0, 1]$ constitutes a PPT-search problem (of Type 4) if the following two conditions hold.*

1. *A relaxed solving condition: There exists a PPT algorithm A such that, for every $x \in \{0, 1\}^*$ and $t \in \mathbb{N}$, it holds that $\Pr[q(x, A(x, 1^t)) \geq q'(x) - (1/t)] \geq 2/3$, where $q'(x) \stackrel{\text{def}}{=} \max_y \{q(x, y)\}$.*
2. *Efficient approximation of q : There exists a PPT algorithm Q such that, for every $x, y \in \{0, 1\}^*$ and $t \in \mathbb{N}$, it holds that $\Pr[|Q(x, y, 1^t) - q(x, y)| \leq 1/t] \geq 2/3$.*

We say that a probabilistic mechanism, M , solves (T4-solves) the search problem q if M satisfies the first condition; that is, $\Pr[q(x, M(x, 1^t)) \geq q'(x) - (1/t)] \geq 2/3$ for every $x \in \{0, 1\}^*$ and $t \in \mathbb{N}$.

Note that Definition 2.4 allows for error reduction (see the first part of Section 4.1). Furthermore, Definition 2.1 can be cast as a special case of Definition 2.4 by considering $q(x, y) \stackrel{\text{def}}{=} 1$ if $(x, y) \in R$ and $q(x, y) \stackrel{\text{def}}{=} 0$ otherwise.

2.3 In Favor of Definition 2.4

As hinted above, we prefer Definition 2.4 over Definitions 2.2 and 2.3, because we prefer the smooth qualitative treatment of the quality of potential solutions (underlying Definition 2.4) over the rigid trichotomy underlying Definitions 2.2 and 2.3. We further articulate the preference in Section 3.3. However, we have no decisive argument in favor of this preference nor can we rule out the possibility of a better alternative.

One technical advantage of Definition 2.4 over Definitions 2.2 and 2.3 is that it allows for error reduction in Condition 1. Indeed, error reduction in Condition 2 is straightforward for all definitions, whereas approximating the quality of a potential solution allows for selecting the best solution among several solutions that are found by repeated invocation of the T4-solver. (For details, see the first part of Section 4.1.)

Another technical advantage of Definition 2.4 is that its solving task is deterministically reducible to a (decisional promise) problem in $\text{pr}\mathcal{BPP}$. Hence, the study of PPT-search problems of Type 4 is reduced to the study of decisional problems in $\text{pr}\mathcal{BPP}$.⁶

Theorem 2.5 (solving PPT-search of Type 4 is deterministically reducible to $\text{pr}\mathcal{BPP}$):⁷ *Let $q : \{0, 1\}^* \times \{0, 1\}^* \rightarrow [0, 1]$ be a PPT-search problem of Type 4. Then, T4-solving q is reducible in deterministic polynomial-time to a promise problem in $\text{pr}\mathcal{BPP}$; that is, there exists $\Pi \in \text{pr}\mathcal{BPP}$ and a deterministic polynomial-time oracle machine M such that $q(x, M^\Pi(x, 1^t)) \geq q'(x) - (1/t)$ for every $x \in \{0, 1\}^*$ and $t \in \mathbb{N}$.*

The proof of Theorem 2.5, which uses the ideas that underlie the proof of [4, Thm. 3.5], appears in Section 4.1. An analogous result regarding PPT-search problems of Type 2 is impossible and for Type 3 it is unlikely. As shown in Section 4.2, these facts are rooted in the impossibility (resp., unlikelihood) of deterministic solvers for some problems in these classes.

2.4 Relating the Definitions

In contrast to Theorem 2.5, the following result, which is stated for PPT-search problems of Type 4, holds also for PPT-search problems of Type 2 and 3.

⁶This fact allows for extending the aforementioned result of [5, Sec. 4] from PPT-search of Type 1 to PPT-search of Type 4. Recall that this result relates the question of whether every PPT-search problem has a pseudodeterministic algorithm [2] to the question of whether every problem in $\text{pr}\mathcal{BPP}$ can be extended to a problem in \mathcal{BPP} .

⁷This result is analogous to [4, Thm. 3.5], which shows that solving an (indirectly specified) subclass of the problems captured by [4, Def. 3.2] is (deterministically) reducible to $\text{pr}\mathcal{BPP}$. Specifically, this subclass consists of search problems that have a “companion” that is a BPP-search problem, where the notion of companion is defined in [4, Obs. 3.3].

Theorem 2.6 (probabilistic FPTAS is a special case of PPT-search):⁸ *Consider a fully PPT approximation scheme for a function $f : \{0, 1\}^* \rightarrow [0, 1]$; that is, a PPT algorithm that, on input $x \in \{0, 1\}^*$ and $k \in \mathbb{N}$, returns a value in $[f(x) \pm (1/k)]$. Then, $q(x, v) \stackrel{\text{def}}{=} \max(1 - |f(x) - v|, 0)$ is a PPT-search problem of Type 4 and approximating f reduces to solving the search problem of q .*

The proof amounts to observing that $q(x, v) \approx 1$ if and only if $f(x) \approx v$. Analogous statements for PPT-search problems of Type 2 and 3 follow from the fact that PPT-search problems of Type 4 can be expressed as (or emulated by) PPT-search problems of Type 2 (resp., Type 3). The latter assertion is captured by Proposition 4.6.

Actually, it is in place to ask what do we mean by saying that one type of problems can be expressed as (or emulated by) another type of problems, let alone when their syntax differs. Loosely speaking, our answer is that this means that for each problem Π of the first type (resp., class) there exists a problem Π' of the second type (resp., class) such that solving Π is reducible to solving Π' . This answer raises the question of what does *solving a problem* mean, and the default and natural answer is that it means the notion of solving that is part of the definition of the class (of problems). (In fact, we made this interpretation explicit in the last part of Definitions 2.1–2.4). Nevertheless, other notions of solvability may be meaningful too. Specifically, the notion of T3-solvability, which is part of Definition 2.3, can be used for the problems of Type 2 (defined in Definition 2.2). We may view T3-solvability of problems of Type 2 as a weak notion of solvability of problems of Type 2.

The usefulness of the foregoing proposal is demonstrated in Proposition 4.4. Loosely speaking, Proposition 4.4 asserts that *for every problem $\bar{R} = (R_{\text{yes}}, R_{\text{no}})$ of Type 2 there exists a problem q of Type 4 such that T3-solving \bar{R} is reducible in deterministic polynomial-time to T4-solving q* . This may be interpreted as saying that problems of Type 2 can be weakly emulated by problems of Type 4. Combining Proposition 4.4 with Theorem 2.5 implies that every problem \bar{R} of Type 2 can be T3-solved by a deterministic polynomial-time reduction to $\text{pr}\mathcal{BPP}$.

An alternative way of interpreting Proposition 4.4 is to say that it refers to a subclass of the problems of Type 3 (i.e., the class of problems of Type 2). Hence, we can say that Proposition 4.4 identifies a subclass of problems of Type 3 that can be emulated by problems of Type 4 (in the very sense stated above), except that viewing \bar{R} as a problem of Type 3 justifies considering T3-solving it. Of course, this view ignores the fact that the problem \bar{R} is of Type 2 and so a stronger notion of solving is relevant to it.

3 Discussion

In this section we try to identify the principles that underlie the definition of complexity classes and apply them towards evaluating the various definitions presented in Section 2. Hence, we shall start with notions that are well-known and agreed upon by all (e.g., decision problems), and proceed to the more problematic and less explored notions, while trying to identify the underlying principles. Indeed, our real focus is on PPT-search problems; this focus comes to the front in Sections 3.2 and 3.3, which contain the discussions that are most important towards evaluating the various definitions of PPT-search problems.

Complexity classes are classes of sets. Such a class is typically defined by imposing conditions on a broader class of sets, which in turn are specified in some manner. So we start by recalling how sets are typically specified in the context of complexity theory.

⁸Approximating the acceptance probability of a given Boolean circuit is a special case; that is $f(C) = \Pr_x[C(x) = 1]$.

3.1 Specifying Sets and Decisional Problems

Decisional problems are commonly associated with sets of binary strings. Two methods are commonly used to specify such sets.⁹

Specification by recognition: The set is specified by an algorithm (i.e., an operational rule) that distinguishes members of the set from non-members.

We stress that this algorithm needs not be the most efficient algorithm for the corresponding decision problem. It is often the conceptually simplest (or most intuitive) algorithm. For example, the set of primes is defined as the set of natural numbers that have no non-trivial divisor.

Specification by verification: The set is specified by a verification procedure that accepts members of the set, when coupled with adequate proofs, but never accepts non-members. The verification procedure is typically efficient, but the length of possible proofs may vary arbitrarily.

Although this specification can be replaced by one of the former type (i.e., the decision procedure may try all possible proofs), it is often more appealing and intuitive to use the current type. In particular, even sets in \mathcal{P} are often specified in this manner (which only places them in \mathcal{NP}). Examples include the set of satisfiable 2CNF formulae, the set of bipartite graphs, and the set of graphs having perfect matchings.

We stress that whereas \mathcal{P} is defined as the class of sets that can be recognized in polynomial-time, specific sets in \mathcal{P} may be specified in other means (see the foregoing examples). In general, we distinguish between the way that specific sets are specified and the way that a complexity class is defined (and specific sets are shown to reside in it).

Decisional problems of the *pure type* (a.k.a. “decision problems”) consist of determining, for every given instance, whether or not it is a member of a predetermined set. In contrast, *decisional problems of the promise type* (a.k.a. “promise problems” [1]) only require such distinguishing for instances that are in a (predetermined) promise set. Hence, **decisional problems of the promise type** are defined by the specification of two sets, a main set S and a promise set P , and consist of determining membership in S for instances in P only; equivalently, distinguishing between $P \cap S$ and $P \setminus S$. Indeed, these two sets determine a three-way partition of $\{0, 1\}^*$ (i.e., the partition $(P \cap S, P \setminus S, \{0, 1\}^* \setminus P)$), where instances in $\{0, 1\}^* \setminus P$ are said to violate the promise. We stress that when considering decision procedures for such promise problems, their behavior on instances that violate the promise is immaterial (and is being ignored).

Defining complexity classes. Usually, classes of decisional promise problems are defined by considering the *complexity of the task of distinguishing between $P \cap S$ and $P \setminus S$* . For example, prBPP (promise-BPP) is defined as the collection of pairs of sets (S, P) such that there exists a PPT algorithm A such that $\Pr[A(x) = 1] \geq 2/3$ if $x \in P \cap S$ and $\Pr[A(x) = 1] \leq 1/3$ if $x \in P \setminus S$. Equivalently, one may require that for every $x \in P$, it holds that $\Pr[A(x) = \chi_S(x)] \geq 2/3$, where $\chi_S(x) = 1$ if $x \in S$ and $\chi_S(x) = 0$ otherwise.

⁹In addition, sets are also specified by set-operations that involve previously specified sets.

3.2 Specifying Search Problems

Search problems are commonly associated with binary relations $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$, which are typically *specified by recognition algorithms*. In this case, the search problem consists of finding, when given an instance x , a string y such that $(x, y) \in R$. Such a string y is called a **valid solution** for the instance x , and $R(x) \stackrel{\text{def}}{=} \{y : (x, y) \in R\}$ is the set of valid solutions for x . The bottom-line is that search problems are specified by specifying the binary relation that is associated with them.

A **promise problem of the search type** [3, Def. 2.29] is specified by a promise set $P \subseteq \{0, 1\}^*$ and a binary relation R (as above), and consists of finding valid solutions for instances in P (only). We stress that the promise refers to the instance only. Hence, just as in the case of decisional promise problems, the meaning of the promise in such search problems is clear: Instances that violate the promise do not matter and the performance of search algorithms on them is ignored.

In contrast, considering the task of recognizing members of R leads to a relaxation of the promise type, but here the promise is a subset P of instance-solution pairs. In this case, the semantics of violating the promise P is less clear (i.e., what does $(x, y) \notin P$ mean?). The problem arises when trying to formulate the searching task. Specifically, given an instance x , *are we required to return an element of $R(x) \cap P(x)$* (i.e., a solution y such that $(x, y) \in P \cap R$) or *are we allowed to return any element of $R(x)$* ?

Indeed, the source of trouble is our choice to specify a search problem via a trichotomy of instance-solution pairs, which yields three types of solutions (i.e., valid, invalid, and in-between (a.k.a limbo)). The motivation for this choice is extending natural classes of search problems by capitalizing on the possibility that distinguishing between $P \cap R$ and $P \setminus R$ may be easier than distinguishing between R and $\{0, 1\}^* \setminus R$ (e.g., in the context of probabilistic polynomial-time). *But is this the best way of achieving the aforementioned goal?*

On the actual specification of PPT-search problems in Section 2. Before addressing the foregoing question, we note that underlying Definitions 2.1–2.4 are general mechanisms for defining search problems. Such a **mechanism** consists of associating a search problem with a specific formal structure and defining what is considered as solving the problem. Here we explicitly state the four mechanisms that are used in these four definitions.

In Definition 2.1 a problem is specified by a binary relation $R \subseteq \{0, 1\}^* \times \{0, 1\}^*$, and solving the problem means providing an element of $R(x)$ when given an instance x such that $R(x) \neq \emptyset$.

In Definition 2.2 a problem is specified by a pair of disjoint binary relation $R_{\text{yes}}, R_{\text{no}} \subseteq \{0, 1\}^* \times \{0, 1\}^*$, and solving the problem means providing an element of $R_{\text{yes}}(x)$ when given an instance x such that $R_{\text{yes}}(x) \neq \emptyset$.

In Definition 2.3 a problem is specified as in Definition 2.2 (i.e., by a pair of disjoint binary relation $R_{\text{yes}}, R_{\text{no}} \subseteq \{0, 1\}^* \times \{0, 1\}^*$), and solving the problem means providing an element of $\{0, 1\}^* \setminus R_{\text{no}}(x)$ when given an instance x such that $R_{\text{yes}}(x) \neq \emptyset$.

In Definition 2.4 a problem is specified by a function $q : \{0, 1\}^* \times \{0, 1\}^* \rightarrow [0, 1]$, and solving the problem means providing an string y such that $q(x, y) \geq q'(x) - (1/t)$ when given an instance x and a parameter $t \in \mathbb{N}$.

The definitions presented in Section 2 were obtained by imposing complexity bounds on the relevant structures as well as on the relevant solvers.

We stress that all decisional (promise) problems are specified by one type of structure (i.e., a pair of sets) and the same notion of solving (the corresponding problem). In contrast, Definitions 2.2–2.4 exhibit two different types of structures and three different notion of solving. While the structure-type used in Definitions 2.2–2.3 is aligned with the traditional structure-type used in Definition 2.1, we see no reason to prefer it over the structure-type used in Definition 2.4. That is, the trichotomy of potential solutions (i.e., valid, invalid, limbo) does not seem more appealing than the qualitative view of the quality of potential solutions.

3.3 Defining Natural Complexity Classes of Search Problems

Let us take our cue from the class of search problems that arises, implicitly, in the context of the P-vs-NP question (cf. [3, Sec. 2.1.1]). First, we define the class of search problems that consists of all binary relations for which the recognition problem can be solved in polynomial-time (cf. [3, Def. 2.3]). When considering only polynomially bounded relations (i.e., R such that $(x, y) \in R$ implies $|y| \leq \text{poly}(|x|)$), this yields a class of search problems, denoted \mathcal{PC} , which is analogous to \mathcal{NP} . The P-vs-NP question is equivalent to asking whether solutions for every problem in \mathcal{PC} can be found in polynomial-time, and the implicit class that emerges is the class of all search problems in \mathcal{PC} for which solutions can be found in polynomial-time.

We seek to follow the same path towards defining a corresponding class for probabilistic polynomial-time (PPT) algorithms. Doing so leads to [4, Def. 3.1], which is reproduced as Definition 2.1. Unfortunately, the resulting class feels too restrictive; in particular, it does not contain the natural problem of approximating the acceptance probability of a given Boolean circuit (i.e., given a Boolean circuit C and $t \in \mathbb{N}$, find v such that $|v - \Pr_r[C(r)=1]| \leq 1/t$).

The path taken by [4, Def. 3.2] is the one outlined in (the middle part of) Section 3.2; indeed, this is the path that was followed in Definitions 2.2 and 2.3. The starting point of this path is a focus on the complexity of distinguishing between valid solutions and invalid solutions. The distinguishing task is made easier by introducing limbo solutions (on which the distinguisher may answer arbitrarily), but *the treatment of these limbo solutions in the solving task is conceptually discomforting*. This discomfort is amplified by our opinion that the study of search problems should give priority to the solving task. Hence, it seems more natural to take an alternative path, one that avoids the question of how to treat limbo solutions. Such an alternative path was taken by Definition 2.4. Let us outline the principles underlying this alternative path.

The first principle is that, when considering search problems that can be solve within some resource bounds, it is natural to confine the attention to problems for which solutions can be evaluated for validity within the same resource bounds. Intuitively, searching for a solution makes less sense if one cannot evaluate the validity (or quality) of the solution once found. In other words, the solution is meaningful for the finder, because it can recognize the validity (or quality) of the solution. We call this principle the **meaningfulness principle**. Formally, the search version of the P-vs-NP question is obtained by applying the meaningfulness principle (i.e., restricting attention to the class \mathcal{PC}). The meaningfulness principle also led us to Definitions 2.1–2.4.

The second principle calls for trying to provide the most general definition that corresponds to a natural notion, and yet do so in an intuitive manner (i.e., the generalized definition should be appealing). Hence, we tried to extend Definition 2.1, which is natural *per se* but was deemed too restrictive. Unfortunately, the extensions proposed by Definitions 2.2 and 2.3 were problematic: They replaced the clear valid-vs-invalid dichotomy (regarding potential solutions) by a trichotomy that included in-between (or limbo) solutions. Definitions 2.2 and 2.3 differ in the way they handle

these limbo solutions, but in both cases this treatment does not abide the meaningfulness principle: Definition 2.2 does not allow the searcher to return a limbo solution although this solution may be indistinguishable from a valid solution, whereas Definition 2.3 allows the searcher to return a limbo solution although this solution may be indistinguishable from an invalid solution.

In contrast, Definition 2.4 softens the valid-vs-invalid (qualitative) dichotomy and replaces it by a quantitative spectrum. Viewed from this perspective, perfectly valid solutions are viewed as having quality 1, evidently invalid solutions are viewed as having quality 0, whereas, in general, potential solutions have a value in $[0, 1]$. The task of the searcher is to find a solution of maximal quality (or rather almost maximal quality), whereas the quality of solutions can be well-approximated (by a PPT algorithm). Hence, in our opinion, Definition 2.4 satisfies both principles stated above.

In light of the foregoing, we prefer Definition 2.4 over Definitions 2.2 and 2.3. Specifically, we consider the quantitative view of the quality of potential solutions far more appealing than the valid vs limbo vs invalid trichotomy, let alone the problematic treatment of limbo solutions. Furthermore, we consider the fact that Definition 2.4 supports error reduction, whereas Definitions 2.2 and 2.3 do not seem to do so, a fundamental issue. Supporting error reduction allows for having a class that is independent of the specific error bound used in its definition (as long as this bound is reasonable).

We note that, although not intended for that purpose, Definition 2.4 captures the natural notion of a PPT-solvable optimization problem: Specifically, for a polynomial-time computable function $f : \{0, 1\}^* \times \{0, 1\}^* \rightarrow [0, 1]$, consider the task in which, on input x and $\epsilon > 0$, one is required to find y such that $f(x, y) \geq \max_z \{f(x, z)\} - \epsilon$ and do so in probabilistic $\text{poly}(|x|/\epsilon)$ -time. Indeed, this is a special case of Definition 2.4, where $q = f$ is *computable in deterministic polynomial-time* rather than only *approximable in probabilistic polynomial-time*.

3.4 Relating Complexity Classes of Different Types

A typical class of computational problems incorporates, in its definition, a notion of solving problems in the class. In the case of decisional problems solving refers to deciding membership (possibly in a relaxed manner as captured by promise problems). In the case of search problems solving refers to finding solutions. As illustrated by these cases, different types of problems call for different notions of solving these problems. Furthermore, as suggested by Definitions 2.1–2.4, even different types of PPT-search problems refer to different notions of solving these problems. This fact makes relating different types of problems less straightforward than it may seem.

Indeed, the natural choice is to say that a problem Π of one class is emulated by problem Π' of another class if solving Π in the sense of the first class reduces to solving Π' in the sense of the second class. Furthermore, this interpretation of emulation is called for when referring to classes of problems rather than to individual problems.

Still, saying that solving Π in some other sense reduces to solving Π' in yet another sense is also informative (albeit there is a gap between being informative and being appealing). In fact, Proposition 4.4 asserts that T3-solving (rather than T2-solving) a PPT-search problem of Type 2 is reducible to T4-solving a problem of Type 4. That is, the first problem is solved in a sense that is weaker than the sense associated with the corresponding class. We stress that, using the foregoing flexibility, Proposition 4.4 asserts a result that is weaker than the natural formulation. Yet, this is an informative result.

4 Proofs

This section contains proofs of the results stated in Section 2 and some additional observations.

4.1 Proof of Theorem 2.5

The proof of Theorem 2.5 is analogous to the proof of [4, Thm. 3.5], which refers to an indirectly specified subclass of PPT-search problems of Type 2. Towards proving Theorem 2.5, we first strengthen the two conditions of Definition 2.4 (when referring to the function $q : \{0, 1\}^* \times \{0, 1\}^* \rightarrow [0, 1]$). Starting with Condition 2 and employing error-reduction (while taking the median value), we note that this condition implies the existence of a PPT algorithm Q such that for every $x, y \in \{0, 1\}^*$ and $t \in \mathbb{N}$ it holds that

$$\Pr \left[|Q(x, y, 1^t) - q(x, y)| \leq \frac{1}{t} \right] \geq 1 - 2^{-t}. \quad (1)$$

Next, turning to Condition 1 and employing error-reduction (while using the solution of the maximum approximate quality¹⁰), we note that this condition implies the existence of a PPT algorithm F such that for every $x \in \{0, 1\}^*$ and $t \in \mathbb{N}$ it holds that

$$\Pr \left[q(x, F(x, 1^t)) \geq q'(x) - \frac{1}{t} \right] \geq 1 - 2^{-t}, \quad (2)$$

where $q'(x) \stackrel{\text{def}}{=} \max_y \{q(x, y)\}$. We shall also consider arbitrary probabilistic procedures F that satisfy Eq. (2). At this point, we restate Theorem 2.5 as follows.

Theorem 4.1 (PPT solvers vs reduction to promise-BPP): *Let $q : \{0, 1\}^* \times \{0, 1\}^* \rightarrow [0, 1]$ be a function such that there exists a PPT algorithm Q that satisfies Eq. (1). Then, the following two statements are equivalent:*

1. *There exists a PPT algorithm F that solves the search problem of q ; that is, F satisfies Eq. (2).*
2. *There exists a deterministic polynomial-time reduction of solving the search problem of q to a decisional problem in promise-BPP; that is, there exists a deterministic polynomial-time oracle machine M and a decisional promise problem Π in promise-BPP such that M^Π satisfies Eq. (2). (Actually, since M is deterministic, $q(x, M^\Pi(x, 1^t)) \geq q'(x) - (1/t)$ always holds).*

Recall that, by definition (of a reduction to a promise problem), the reduction in Statement 2 yields the correct answer regardless of how queries that violate the promise are answered. The proof of Theorem 4.1 uses the ideas that underlie the proof of [4, Thm. 3.5].

Proof: Clearly, if Statement 2 holds, then we can derive a PPT algorithm that solves q (i.e., establish Statement 1) by emulating the oracle in a straightforward manner (after adequate error reduction).

The proof of the opposite direction (i.e., Statement 1 implies Statement 2) relies on the fact that a promise-BPP oracle can be utilized for finding a sequence of choices (for F 's randomness) that leads F to a solution of almost maximal quality. This is done by an iterative process that

¹⁰Specifically, we invoke the original solver A as well as the evaluator Q with deviation bound $1/3t$, while capitalizing on the fact that each invocation of Q errs with probability at most 2^{-3t} .

extends the prefix of such a sequence, one bit at a time, while relying on the ability to approximate the expected quality of the solution found by the residual solver (that arises by fixing a prefix of such a sequence). Indeed, this approximation is obtained from the promise-BPP oracle. Details follow.

Let F be a PPT algorithm that solves the search problem of q (i.e., satisfies Eq. (2)), and let F_r denote the residual deterministic algorithm obtained from F when fixing its randomness to $r \in \{0, 1\}^m$, where $m = \text{poly}(|x| + t)$. Consider the function $p : \{0, 1\}^* \times \{1\}^* \times \{0, 1\}^* \rightarrow [0, 1]$ defined by

$$p(x, 1^t, r') \stackrel{\text{def}}{=} \mathbb{E}_{r'' \in \{0, 1\}^{m-|r'|}} [q(x, F_{r'r''}(x, 1^t))]. \quad (3)$$

By the hypothesis that F satisfies Eq. (2), it follows that

$$p(x, 1^t, \lambda) \geq (1 - 2^{-t}) \cdot (q'(x) - (1/t)) > q'(x) - (2/t).$$

On the other hand, for every r' , it holds that $p(x, 1^t, r') = (p(x, 1^t, r'0) + p(x, 1^t, r'1))/2$. Hence, approximating the function p up to an additive term of ϵ allows to find a bit b such that $p(x, 1^t, r'b) > p(x, 1^t, r') - \epsilon$, which in turn yields an iterative process that on input x and 1^t yields a string $r \in \{0, 1\}^m$ such that $q(x, 1^t) > q'(x) - m \cdot \epsilon$.

Now, observe that p can be approximated up-to an additive term of $2/mt$ by making mt oracle calls to the decisional promise problem $\Pi = (\Pi_{\text{yes}}, \Pi_{\text{no}})$ such that

$$\begin{aligned} \Pi_{\text{yes}} &\stackrel{\text{def}}{=} \{(x, 1^t, r', 1^k, v) : |p(x, 1^t, r') - v| \leq 1/k\} \\ \Pi_{\text{no}} &\stackrel{\text{def}}{=} \{(x, 1^t, r', 1^k, v) : |p(x, 1^t, r') - v| > 2/k\}. \end{aligned}$$

Specifically, we approximate $p(x, 1^t, r')$ (up to an additive term of $2/mt$) by making the queries $(x, 1^t, r', 1^{mt}, (i-0.5)/mt)$ for all $i \in [mt]$, and setting the estimate as an arbitrary value $(i-0.5)/mt$ such that the oracle answered the corresponding query with a **yes**.¹¹

Next, we shall show (in Claim 4.1.1) that solving the search problem of q is reducible (in deterministic polynomial-time) to approximating p up to an additive term of $2/mt$. Recalling that such an approximation can be obtained by making mt queries to the decisional promise problem $\Pi = (\Pi_{\text{yes}}, \Pi_{\text{no}})$, we shall complete the proof by showing that Π is in promise-BPP (see Claim 4.1.2).

Claim 4.1.1 (solving q reduces to approximating p): *Solving the search problem of q reduces in deterministic polynomial-time to approximating $p(x, 1^t, \cdot)$ up to an additive term of $1/\text{poly}(|x| + t)$.*

Recall that approximating p in the foregoing sense reduces to solving the decisional promise problem Π .

Proof: Recall that on input $(x, 1^t)$, we seek y such that $q(x, y) \geq q'(x) - (1/t)$. Equivalently, we seek $r \in \{0, 1\}^m$ such that $q(x, F_r(x, 1^t)) \geq q'(x) - (1/t)$. We find such an r in m iterations by gradually extending a candidate prefix of r by one bit in each iteration. Specifically, in the i^{th} iteration the current $(i-1)$ -bit long prefix is extended by a single bit towards the i -bit long prefix of the higher value; that is, $r' \in \{0, 1\}^{i-1}$ is extended to $r'b$ if the estimate of $p(x, 1^t, r'b)$ is at least as large as the estimate of $p(x, 1^t, r'\bar{b})$. Consequently, it holds that $p(x, 1^t, r'b) \geq p(x, 1^t, r') - (2/mt)$, where

¹¹We stress that this reduction also issue queries that violate the promise of the decisional promise problem, but the result of the reduction is correct regardless of how these (“promise violating”) queries are answered. Actually, in general, such (“promise violating”) queries are crucial (cf. [3, Exer. 2.36]).

we assume that p is approximated upto an additive term of $2/mt$. Thus, after m iterations, we obtain $r \in \{0, 1\}^m$ such that $p(x, 1^t, r) \geq p(x, 1^t, \lambda) - (2m/mt) > q'(x) - (4/t)$, and output the value $F_r(x, 1^t)$ (such that $q(x, F_r(x, 1^t)) > q'(x) - (4/t)$). Lastly, by substituting t with $4t$, we can output a value that exceeds $q'(x) - (1/t)$. \square

Claim 4.1.2 (Π is in promise-BPP): *The promise problem Π is in prBPP .*

Proof: The key observation is that $p(x, 1^k, r')$ can be approximated up-to an additive term of $1/k$ (with error probability at most $\exp(-\Omega(k))$) by approximating $q(x, F_{r'r''}(x, 1^t))$ on k^3 random choices of $r'' \in \{0, 1\}^{m-|r'|}$. Specifically, we invoke $Q(x, F_{r'r''}(x, 1^t), 1^{2k})$ on k^3 random choices of $r'' \in \{0, 1\}^{m-|r'|}$, and use the average of the obtained values. Indeed, each of these k^3 invocations satisfies $\Pr[|Q(x, F_{r'r''}(x, 1^t), 1^{2k}) - q(x, F_{r'r''}(x, 1^t))| \leq (1/2k)] \geq 1 - 2^{-2k}$, whereas (with probability at least $1 - \exp(-\Omega(k))$) the average of the k^3 sampled $q(x, F_{r'r''}(x, 1^t))$'s is within $1/2k$ of the expected value. \square

Combining Claims 4.1.1 and 4.1.2, the theorem follows. \blacksquare

4.2 On the Failure of Other Versions of Theorem 2.5

Here “ T_i -search” stands for *PPT-search problem of Type i* . Recall that Theorem 2.5 asserts that solving (i.e., T_4 -solving) T_4 -search problems is reducible in deterministic polynomial-time to prBPP , and that an analogous result holds for T_1 -search problems (see [4, Thm. 3.5]). In contrast, we show that an analogous result does not hold for T_2 -search problems, and is unlikely also for T_3 -search problems. We start with the first claim, which follows from the fact that there exist T_2 -search problems that cannot be solved by deterministic algorithms (regardless of their running time).

Proposition 4.2 (T_2 -solving some T_2 -search problems cannot be deterministically reduced to prBPP): *There exists a T_2 -search problem $\bar{R} = (R_{\text{yes}}, R_{\text{no}})$ that cannot be T_2 -solved by a deterministic algorithm. In particular, T_2 -solving \bar{R} cannot be reduced in deterministic polynomial-time to a promise problem in prBPP .*

Proof: Denoting the Kolmogorov complexity of a string x by $K(x)$, consider $R_{\text{yes}} = \{(1^{|x|}, x) : K(x) > |x|/2\}$ and $R_{\text{no}} = \emptyset$. Then, $\bar{R} = (R_{\text{yes}}, R_{\text{no}})$ satisfies the conditions of Definition 2.2 (e.g., Condition 1 is satisfied by an algorithm that, on input 1^n , generates a uniformly distributed n -bit string, whereas Condition 2 holds trivially). On the other hand, no deterministic algorithm can T_2 -solve \bar{R} , because $A(1^n) = x \in \{0, 1\}^n$ implies $K(x) \leq |A| + O(1) + \log_2 n$ (which is smaller than $|x|/2$ for fixed A and sufficiently large n). Lastly, observe that a deterministic polynomial-time reduction of T_2 -solving \bar{R} to $\Pi \in \text{prBPP}$ implies a deterministic algorithm for T_2 -solving \bar{R} , because Π can be decided by a deterministic (exponential-time) algorithm. \blacksquare

Proposition 4.3 (on T_3 -solving some T_3 -search problems either in deterministic polynomial-time or by a reduction to prBPP):

1. *If $\mathcal{P} \neq \mathcal{NP}$, then there exists a T_3 -search problem $\bar{R} = (R_{\text{yes}}, R_{\text{no}})$ that cannot be T_3 -solved by a deterministic polynomial-time algorithm.*

2. If $\mathcal{NP} \not\subseteq \mathcal{BPP}$, then there exists a T3-search problem $\bar{R} = (R_{\text{yes}}, R_{\text{no}})$ that cannot be T3-solved by a deterministic polynomial-time reduction to a promise problem in $\text{pr}\mathcal{BPP}$.

Proof: Starting with Part 1, let R be the NP-witness relation of some set S in $\mathcal{NP} \setminus \mathcal{P}$ (i.e., $x \in S$ iff $R(x) \neq \emptyset$), and assume (w.l.o.g.) that $(x, w) \in R$ implies $|w| = 3 \cdot |x|$. Consider $R_{\text{yes}} = R$ and

$$R_{\text{no}} = \{(x, w) : |w| \neq 3 \cdot |x|\} \cup \{(x, w) : w \in (\{0, 1\}^{3|x|} \setminus R(x)) \ \& \ K(w) < |w|/2\}$$

where $K(w)$ denotes the Kolmogorov complexity of w . Hence, w is a limbo solution for x if and only if $w \in (\{0, 1\}^{3|x|} \setminus R(x))$ and $K(w) \geq |w|/2 = 3|x|/2$

We first note that $\bar{R} = (R_{\text{yes}}, R_{\text{no}})$ satisfies the conditions of Definition 2.3: Condition 1 is satisfied by an algorithm that, on input x , generates a uniformly distributed $3|x|$ -bit string (which is likely to be a limbo solution), whereas Condition 2 holds by virtue of $R_{\text{yes}} = R \in \mathcal{P}$. (Indeed, we satisfy Condition 1 by outputting a limbo solution, but satisfy Condition 2 by rejecting limbo solutions.)

To show that no deterministic polynomial-time algorithm can T3-solve \bar{R} , we consider two cases regarding the possible behavior of a hypothetical deterministic polynomial-time algorithm A that T3-solves \bar{R} (i.e., for each $x \in S$, it holds that $A(x) \in (\{0, 1\}^* \setminus R_{\text{no}}(x)) = \{0, 1\}^{3|x|} \setminus R_{\text{no}}(x)$).

Case 1: A solves R ; that is, for all but finitely many x 's in S , it holds that $A(x) \in R(x)$. In this case, we obtain a polynomial-time decision procedure for S : On input x , this procedure invokes A on x , and accepts if and only if $(x, A(x)) \in R$. This contradicts the hypothesis $S \notin \mathcal{P}$.

Case 2: A does not solve R ; that is, for infinitely many x 's in S , it holds that $A(x) \in \{0, 1\}^{3|x|}$ and $K(A(x)) \geq 3|x|/2$. But this is impossible, because $K(A(x)) \leq |x| + |A| + O(1)$.

This completes the proof of the first part of the proposition. To prove the second part, we follow a similar strategy, except that we use an NP-witness relation R of a set S in $\mathcal{NP} \setminus \mathcal{BPP}$. Letting M be a hypothetical deterministic polynomial-time reduction of T3-solving \bar{R} to $\Pi \in \text{pr}\mathcal{BPP}$, we consider the analogous two cases.

Case 1: M^Π solves R ; that is, for all but finitely many x 's in S , it holds that $M^\Pi(x) \in R(x)$. In this case, we obtain a PPT decision procedure for S : On input x , this procedure invokes M on x , while emulating the oracle calls to Π (and accepting if and only if M 's output is in $R(x)$). This contradicts the hypothesis $S \notin \mathcal{BPP}$.

Case 2: M^Π does not solve R ; that is, for infinitely many x 's in S , it holds that $M^\Pi(x) \in \{0, 1\}^{3|x|}$ and $K(M^\Pi(x)) \geq 3|x|/2$. But this is impossible, because $K(M^\Pi(x)) \leq |x| + |M| + |P| + O(1)$, where P is a deterministic algorithm that solves Π .

This completes the proof of the second part of the proposition. ■

Digest. The foregoing proofs capitalize on a concern that was raised in Section 3.3. Specifically, Definition 2.2 does not allow the searcher to return a limbo solution although this solution may be indistinguishable from a valid solution, whereas Definition 2.3 allows the searcher to return a limbo solution although this solution may be indistinguishable from an invalid solution. Indeed, this is exactly what happens in the proofs of Proposition 4.2 and Proposition 4.3, respectively.

4.3 Emulations

Again, “ Ti -search” stands for *PPT-search problem of Type i* . Recall that both T2- and T3-search refer to $\bar{R} = (R_{\text{yes}}, R_{\text{no}})$ that satisfies the same Condition 2 (i.e., \bar{R} as a decisional promise problem is in prBPP). The difference between these two types is captured by Condition 1: In T2-search solving requires outputting a valid solution, whereas in T3-search solving allows outputting a limbo solution. In the sequel, we refer to the first notion of solving (i.e., T2-solving) as to **strongly-solving**, and to the second (relaxed) notion of solving (i.e., T3-solving) as **weakly-solving**. In addition, we generalize Definition 2.4 so as to allow also functions of the form $q : P \times \{0, 1\}^* \rightarrow [0, 1]$, where $P \subseteq \{0, 1\}^*$ is a promise set. In this generalization (of Definition 2.4), we require:

1. A relaxed solving condition: There exists a PPT algorithm A such that, for every $x \in P$ and $t \in \mathbb{N}$, it holds that $\Pr[q(x, A(x, 1^t)) \geq q'(x) - (1/t)] \geq 2/3$, where $q'(x) \stackrel{\text{def}}{=} \max_y \{q(x, y)\}$.
2. Efficient approximation of q : There exists a PPT algorithm Q such that, for every $x \in P$, $y \in \{0, 1\}^*$ and $t \in \mathbb{N}$, it holds that $\Pr[|Q(x, y, 1^t) - q(x, y)| \leq 1/t] \geq 2/3$.

Lastly, for a binary relation R , we use the notation $S_R \stackrel{\text{def}}{=} \{x : R(x) \neq \emptyset\}$.

Proposition 4.4 (weak emulation of T2-search by T4-search): *For every T2-search problem $\bar{R} = (R_{\text{yes}}, R_{\text{no}})$ there exists a T4-search problem $q : S_{R_{\text{yes}}} \times \{0, 1\}^* \rightarrow [0, 1]$ such that weakly-solving \bar{R} is reducible in deterministic polynomial-time to solving q .*

Note that the hypothesis (that \bar{R} is a T2-search problem) implies that \bar{R} is strongly-solvable (in PPT), but this does not mean that strongly-solving \bar{R} is reducible in deterministic polynomial-time to solving q . We also note that the T4-search problem q depends on the algorithm that distinguishes between R_{yes} and R_{no} (rather than only depending on $(R_{\text{yes}}, R_{\text{no}})$ itself).

Proof: Letting D be an algorithm distinguishing between R_{yes} and R_{no} (as provided in Definition 2.2), we define the function $q : S_{R_{\text{yes}}} \times \{0, 1\}^* \rightarrow [0, 1]$ such that

$$q(x, y) \stackrel{\text{def}}{=} \min(\Pr[D(x, y) = 1], 2/3)$$

and observe that q can be approximated as required by Condition 2 (of Definition 2.4). Letting A be the algorithm that finds solutions w.r.t to R_{yes} (per Condition 1 of Definition 2.2), observe that when A is invoked on an input $x \in S_{R_{\text{yes}}}$, it outputs a valid solution, which has value $2/3 = q'(x)$. Hence, q constitutes a T4-search problem.

Next, we reduce weakly-solving \bar{R} to solving q . On input $x \in S_{R_{\text{yes}}}$, we invoke the T4-solver, denoted F , with parameter $t > 3$ (e.g., $t = 4$). In this case (w.p. at least $2/3$), F finds y such that $q(x, y) \geq q'(x) - (1/t) > 2/3 - 1/3 = 1/3$, which means that y is *not* invalid. Hence, $A(x) \stackrel{\text{def}}{=} F(x, 1^t)$ weakly-solves \bar{R} . ■

Remark 4.5 (an alternative version of Proposition 4.4): *In the proof of Proposition 4.4, the function q was restricted to instances in $S_{R_{\text{yes}}}$ in order to cope with the fact that Definition 2.2 makes no postulation about the output of the search algorithm in case the instance is not in $S_{R_{\text{yes}}}$. An alternative that allows using Definition 2.4 itself, rather than its generalization, is augmenting Definition 2.2 with the postulation that $S_{R_{\text{yes}}}$ is in BPP . The proof of Proposition 4.4 can then be adapted by redefining q such that $q(x, y) \stackrel{\text{def}}{=} 0$ for every $x \notin S_{R_{\text{yes}}}$ and $y \in \{0, 1\}^*$.*

Proposition 4.6 (emulation of T4-search by T2-search): *For every T4-search problem $q : P \times \{0, 1\}^* \rightarrow [0, 1]$ there exists a T2-search problem $\overline{R} = (R_{\text{yes}}, R_{\text{no}})$ such that solving q is reducible in deterministic polynomial-time to weakly-solving \overline{R} .*

Needless to say, this implies that solving q is reducible in deterministic polynomial-time to strongly-solving \overline{R} . Hence, Definitions 2.2 and 2.3 are each at least as expressive as Definition 2.4.

Proof: Given q as in the hypothesis, define

$$\begin{aligned} R_{\text{yes}} &\stackrel{\text{def}}{=} \{(\langle x, 1^t \rangle, y) : q(x, y) \geq q'(x) - (1/t)\} \\ R_{\text{no}} &\stackrel{\text{def}}{=} \{(\langle x, 1^t \rangle, y) : q(x, y) < q'(x) - (2/t)\}. \end{aligned}$$

Using both algorithms associated with q , we can distinguish between R_{yes} and R_{no} : This is done by approximating $q(x, y)$ as well as $q'(x)$, where q' is approximating by finding a solution of almost maximal quality and approximating its quality. Hence, Condition 2 of Definition 2.2 is satisfied. Turning to Condition 1 (of Definition 2.2), we find a valid solution for $x' = \langle x, 1^t \rangle$ by invoking $A(x, 1^t)$, while observing that (w.p. at least $2/3$) the output y satisfies $q(x, y) \geq q'(x) - (1/t)$, which means that $y \in R_{\text{yes}}(x')$. Hence, \overline{R} is a T2-search problem.

Next, we reduce solving q to weakly-solving \overline{R} . Let A denote the algorithm that weakly-solves \overline{R} , and note that $S_{R_{\text{yes}}} = \{0, 1\}^* \times \{1\}^*$. Then, on input $x \in \{0, 1\}^*$ and 1^t , we invoke A on input $\langle x, 1^{2t} \rangle$, and obtain a (non-invalid) solution y , which means that $q(x, y) \geq q'(x) - (2/2t)$. Hence, $F(x, 1^t) \stackrel{\text{def}}{=} A(x, 1^{2t})$ solves q . ■

Digest. Combining Propositions 4.4 and 4.6 we infer that Definitions 2.2 and 2.4 can weakly emulate each other in the sense that for every problem Π in one class there is a problem Π' in the other class such that X-solving Π is reducible to X'-solving Π' , where X (resp., X') equals T3 when the PPT-search problem is of Type 2 and equals T4 otherwise (i.e., when the PPT-search problem is of Type 4). Indeed, the fact that we used T3-solving rather than T2-solving for PPT-search problems of Type 2 is a drawback. Furthermore, unlike Tell [6], we are unwilling to view “T3-solving a problem of Type 2” as a natural notion.¹²

5 Conclusion

As articulated in Section 3.3, we prefer Definition 2.4 over Definitions 2.2 and 2.3, and tentatively suggest using Definition 2.4 as the definition of PPT-search problems.

Our main conceptual consideration is that the quantitative view of the quality of potential solutions is more appealing than the qualitative trichotomy of valid vs limbo vs invalid solutions, let alone the controversial treatment of limbo solutions. In contrast to the qualitative distinction between valid and invalid solutions, a sharp distinction between valid (resp., invalid) and limbo solutions seems artificial; a quantitative spectrum (of the quality of solutions) with valid and invalid (solutions) at its extremes is far more justified when considering solutions that are in-between valid and invalid.

Another important conceptual consideration is having probabilistic complexity classes that are independent of the specific error bound used in their definition (as long as this bound is reasonable).

¹²We stress that seeing the usefulness of a notion falls short of viewing it as a natural notion.

This feature is provided by Definitions 2.1 and 2.4 (but not by Definitions 2.2 and 2.3), because the former classes support error reduction. (In retrospect, it may be nicer to set the quite arbitrary bound in Condition 1 of these definitions to $1/2$ (rather than to $2/3$).)

Yet another advantage of Definitions 2.1 and 2.4 is that the study of the search problems captured by them can be reduced to the study of corresponding decision problems. Specifically, by Theorem 2.5, solving any PPT-search problem of Type 4 is reducible in deterministic polynomial-time to a decisional problem in prBPP . Recall that Definition 2.1 is a special case of Definition 2.4, and that the latter was preferred due to its larger expressive power (which is essential for capturing many natural problems).

Lastly, in our opinion, the fact that Definition 2.4 associates search problems with bivariate functions rather than with binary relations (or pairs of disjoint binary relations) does not diminish its claim for capturing the intuitive notion of PPT-search problems. In general, the formal structure that is used to capture an intuitive notion is unimportant per se; what matters is how well does the entire mechanism (which includes also the notion of solving) refer to the intuitive notion.

References

- [1] Shimon Even, Alan L. Selman, and Yacov Yacobi. The Complexity of Promise Problems with Applications to Public-Key Cryptography. *Information and Control*, Vol. 61 (2), pages 159–173, 1984.
- [2] Eran Gat and Shafi Goldwasser. Probabilistic Search Algorithms with Unique Answers and Their Cryptographic Applications. *ECCC*, TR11-136, 2011.
- [3] Oded Goldreich. *Computational Complexity: A Conceptual Perspective*. Cambridge University Press, 2008.
- [4] Oded Goldreich. In a World of $P=BPP$. In *Studies in Complexity and Cryptography*, LNCS Vol. 6650, Springer, pages 191–232, 2011.
- [5] Oded Goldreich. Multi-pseudodeterministic algorithms. *ECCC*, TR19-012, 2019.
- [6] Roei Tell. On defining PPT-search problems and PPT-optimization problems. *ECCC*, TR24-163, 2024.