

# Strategies for Basing the CS Theory Course on Non-decision Problems

John MacCormick\*  
Dickinson College  
jmac@dickinson.edu

## ABSTRACT

Computational and complexity theory are core components of the computer science curriculum, and in the vast majority of cases are taught using *decision problems* as the main paradigm. For experienced practitioners, decision problems are the best tool. But for undergraduates encountering the material for the first time, we present evidence that *non-decision problems* (such as optimization problems and search problems) are preferable. In addition, we describe technical definitions and pedagogical strategies that have been used successfully for teaching the theory course using non-decision problems as the central concept.

## CCS CONCEPTS

• **Theory of computation** → **Complexity classes; Problems, reductions and completeness**; • **Applied computing** → **Education**;

## KEYWORDS

theory course; decision problems; computability theory; complexity theory

### ACM Reference Format:

John MacCormick. 2018. Strategies for Basing the CS Theory Course on Non-decision Problems. In *SIGCSE '18: The 49th ACM Technical Symposium on Computer Science Education, Feb. 21–24, 2018, Baltimore, MD, USA*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3159450.3159557>

## 1 INTRODUCTION

Many undergraduate computer science curricula include a “theory” course, covering some aspects of automata theory, computability, and complexity. Central questions addressed by a theory course include: (a) Given a computational problem  $C$ , can we solve  $C$  using a computer program? (b) If so, can we do so in polynomial time? A strong majority of popular theory textbooks focus on computational problems that are in fact *decision problems*—problems that have a yes/no answer. An alternative approach, advocated by this paper, is to consider *general* computational problems—including non-decision problems such as *optimization problems* and *search*

*problems*—while narrowing in on decision problems when appropriate.

There are two principal objectives of this paper. First, we explain why non-decision problems are pedagogically advantageous for introductory theory courses, supporting the explanation with empirical evidence based on a survey of computer science majors. Second, we describe technical details and pedagogical strategies for using non-decision problems as the central concept in the CS theory course, based on experience with four years of course offerings and a recent textbook [16]. The use of non-decision problems is the single most important idea for the suggested approach, but additional strategies for making the CS theory course accessible are also described. This includes the use of real computer programs processing ASCII strings, instead of the more common model of Turing machines processing symbols from an arbitrary alphabet.

The nature of the audience is crucial to our considerations. In this paper, the primary target is undergraduate students who are seeing computability and complexity theory for the first, and quite possibly the last, time. We will refer to such students as the *novice audience*. We focus only on the novice audience in this paper, ignoring other theory course scenarios such as graduate courses or hybrid courses taught to a mixture of graduate students and advanced undergraduates. We also ignore the possibility of theory courses that focus mostly on automata theory, assuming instead that computability and complexity theory comprise at least a majority of the course.

For this audience, the key advantage of non-decision problems is that they are more realistic: they match the previous programming and algorithms experience of undergraduates more closely. As a concrete example (details are given later), a decision problem may ask the yes/no question, “Does this graph have a Hamilton cycle?” The corresponding non-decision problem is, “Please give me a Hamilton cycle of this graph if it has one.” A program that solves the non-decision version produces a useful result that could conceivably be used in a real-world application, whereas a program solving the decision version yields only a single bit: it tells us whether or not a Hamilton cycle exists, but gives no additional information. For the novice audience, focusing on these single-bit decision problems potentially positions the course as abstruse and irrelevant. Focusing on non-decision problems with meaningful solutions, on the other hand, provides more direct connections to earlier courses. In section 4, we provide a combination of educational theory and empirical evidence to support these claims.

However, it is important to note at the outset that we do not suggest decision problems should be jettisoned from the theory course. On the contrary, there are some compelling reasons to incorporate decision problems into any theory course, including: (i) some theorems and proofs are more concise and elegant when phrased in terms of decision problems; (ii) the vast majority of

\* Also affiliated with School of Computing Sciences, University of East Anglia, 2017–19

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*SIGCSE '18, Feb. 21–24, 2018, Baltimore, MD, USA*  
© 2018 Association for Computing Machinery.  
ACM ISBN 978-1-4503-5103-4/18/02...\$15.00  
<https://doi.org/10.1145/3159450.3159557>

existing CS theory literature considers only decision problems, so it is essential that any student planning to take a subsequent graduate-level theory course has been exposed to the decision problem framework. Therefore, in this paper we advocate a hybrid approach, whereby the earlier topics in the theory course are taught with a focus on non-decision problems. When NP-completeness is introduced—typically sometime in the second half of a one-semester course—we suggest transitioning to the classical approach and focusing only on decision problems (with some minor exceptions described later). This hybrid approach allows the novice audience to establish an initial appreciation of the relevance and importance of CS theory via non-decision problems, while still being well-versed in the classical approach by the end of the course.

As a basis for comparing and contrasting various approaches to the theory course, we will refer to a selection of ten textbooks covering the relevant material [2, 5, 10–13, 18, 19, 21, 24]; please see the technical report [15] for details about the selection of comparison textbooks.

## 2 RELATED WORK

There are several strands of related work aiming to make CS theory courses more practical, accessible, or meaningful. We might classify these strands as: (i) interactive automata software tools such as JFLAP and DEM [3, 22, 23]; (ii) lab assignments and visualizations for NP-completeness, an approach sometimes described as “NP-completeness for all” [4, 6, 14]; (iii) recasting the theoretical ideas themselves, for example by emphasizing non-decision problems and including holistic discussions about the implications of NP-completeness and P-versus-NP [8–10, 17, 20]. The present paper falls firmly in category (iii). The ideas described here are orthogonal to categories (i) and (ii). Indeed, the ideas of this paper have been employed over a four-year period in a course that also uses JFLAP and practical programming assignments, thus benefiting from all three strands of literature on the CS theory course.

To the best of our knowledge, Goldreich’s position paper [9] and subsequent textbook [10] comprise the previous work most similar to the present paper. Goldreich makes several important, insightful, and useful suggestions for improving the theory course, including strong advocacy for the use of non-decision problems. The novel contribution of the present paper is to recast the use of non-decision problems in a manner that is more accessible to the novice audience, thus enabling instructors to focus on non-decision problems and deliver the consequent educational benefits to students.

## 3 A PRACTICAL DEFINITION OF COMPUTATIONAL PROBLEMS

Before the advantages of non-decision problems can be explained, we review some elementary background material to establish notation and give a formal definition of a “computational problem,” which can of course be either a decision problem or a non-decision problem. The background review in this section also includes several ancillary recommendations for how to present this material to a novice audience.

Computational theory is often developed in terms of an arbitrary *alphabet* (i.e. a finite set of symbols) denoted  $\Sigma$ . The set of all possible *strings* (i.e. finite sequences of symbols) on  $\Sigma$  is denoted  $\Sigma^*$ .

Experienced practitioners understand that the choice of alphabet is irrelevant for most purposes. But for students being introduced to complexity theory for the first time, it may be preferable to employ a more familiar and obviously relevant alphabet—one that is used by programmers to describe the inputs and outputs of computer programs in practice. This motivates us here to take the ASCII alphabet as our primary example of  $\Sigma$ .

For similar reasons, we use real computer programs as our main computational model, rather than the Turing machines which are more common in the theory literature. A program  $P$  receives a single ASCII string  $w$  as input and the output is either undefined (e.g. if  $P$  crashes or enters an infinite loop) or is an ASCII string denoted  $P(w)$ . When  $P(w)$  is defined, we say  $P$  *rejects*  $w$  if  $P(w) = \text{“no”}$  and  $P$  *accepts*  $w$  if  $P(w) \neq \text{“no”}$ .

A central concept in complexity theory is a *language* (sometimes called a *formal language*), defined as a subset of  $\Sigma^*$ . The intuitive notion of “solving a problem” is usually formalized as “deciding a language.” Thus, a central concept in many treatments is that a program  $P$  *decides* a language  $L$  if  $P$  accepts all strings  $s \in L$  and rejects all  $s \notin L$ .

What is the connection between “deciding a language” and “solving a problem”? For so-called *decision problems*, there is a direct and simple connection. Informally, decision problems are computational questions that have a yes/no answer, such as “Given an integer  $m > 0$ , is  $m$  prime?” or “Given a graph  $G$ , does  $G$  contain a Hamilton cycle?”. Formally, one could define a decision problem as a function from  $\Sigma^*$  to {“yes”, “no”}. When taking  $\Sigma$  as the ASCII alphabet, we first agree an encoding of the relevant mathematical objects into ASCII. For example: the integer  $m = 43552$  might be encoded as the ASCII string “43552”; a graph that is a three-cycle with three vertices labeled  $a, b, c$  might be encoded as the ASCII string “a, b b, c c, a”.

For a given decision problem  $D$ , strings for which the answer is “yes” are called *positive instances* of  $D$  and all other strings are *negative instances*. This leads to a simple and obvious mapping between languages and decision problems. Any language  $L$  corresponds to a decision problem  $D_L$  defined as follows: for  $s \in \Sigma^*$ ,  $D_L(s) = \text{“yes”}$  if and only if  $s \in L$ . The reverse mapping is also simple and obvious: a decision problem  $D$  corresponds to a language  $L_D$  with the property that  $s \in L$  if and only if  $D_L(s) = \text{“yes”}$ .

As a result of the exact correspondence between decision problems and languages, the theory of computability and complexity is usually described in terms of languages. As already mentioned in the introduction, the advantages of doing so include simple notation and compact statements of certain theorems, but disadvantages for the novice audience include a high level of abstraction and the restriction to decision problems, which may appear unfamiliar and irrelevant.

To illustrate the distinction between decision and non-decision problems, we will use two running examples throughout the paper:

- **FACTOR:** The input is a string representing a positive integer  $m$  in decimal notation (e.g. “35”), and a solution is any factor of the input other than 1 and  $m$  (e.g. “5” or “7” for  $m = 35$ ), or “no” if no such factor exists. FACTORD is the corresponding decision problem, with solution “yes” if  $m$  has a non-trivial factor and “no” otherwise.

- **HAMCYCLE**: The input is a string representing a graph  $G$  (e.g. “a, b c c, a”), and a solution is any Hamilton cycle of  $G$  (e.g. “a, b, c” for the given example of  $G$ ), or “no” if no such cycle exists. **HAMCYCLED** is the corresponding decision problem, with solution “yes” if  $G$  has a Hamilton cycle and “no” otherwise.

We also assume the reader is familiar with **SAT** (which asks for a satisfying assignment to a Boolean formula) and **SATD** (the decision variant, which asks whether a satisfying assignment exists).

Several important differences between the decision and non-decision variants should be immediately obvious from the **FACTOR** and **HAMCYCLE** examples:

- Non-decision problems can have multiple solutions (e.g. an integer  $m$  can have multiple non-trivial factors; a graph  $G$  can have multiple distinct Hamilton cycles), whereas decision problems always have a unique solution (“yes” or “no”).
- It is possible for decision and non-decision variants of the same problem to have apparently different complexity properties. **FACTOR** is a classic example here, since **FACTORD** can be solved in polynomial time by the AKS algorithm [1], whereas **FACTOR** has no known polynomial-time method of solution.<sup>1</sup>
- Non-decision variants appear to be more “natural.” For example, it is hard for the novice audience to imagine an application where it is useful to determine the mere existence of a Hamilton cycle, rather than determining the sequence of vertices in the cycle.<sup>2</sup>


A significant fraction of theory textbooks give no formal definition of a non-decision problem. In our sample of ten books, seven give no formal definition [2, 5, 11–13, 18, 24]. One [21] gives a brief definition; Papadimitriou [19] and Goldreich [10] give detailed definitions and analysis. These last two define non-decision problems in terms of binary relations with certain technical properties: they are *polynomially decidable* and *polynomially bounded* (or *balanced*). This approach is correct and concise, but perhaps more abstract than necessary for a novice undergraduate audience. Instead, the following concrete definition is recommended for the novice audience:

A *computational problem* (which may or may not be a decision problem) is a function  $F$ , mapping ASCII strings to sets of ASCII strings. If  $F(x) = \{s_1, s_2, \dots\}$ , we call  $\{s_1, s_2, \dots\}$  the *solution set* for  $x$ , and each  $s_i$  is a *solution* for  $x$ . If  $F(x) = \{\text{“no”}\}$ , then  $x$  is a *negative instance* of  $F$ ; otherwise  $x$  is a *positive instance*. If  $F$  has unique solutions,  $F$  is a *function problem*. For function problems, we can drop the set notation, for example writing  $F(x) = y$  instead of  $F(x) = \{y\}$ .

For example: “35” is a positive instance of **FACTOR**, and we have  $\text{FACTOR}(\text{“35”}) = \{\text{“5”}, \text{“7”}\}$ ; “29” is a negative instance, and we have  $\text{FACTOR}(\text{“29”}) = \{\text{“no”}\}$ . Note that neither **FACTOR** nor **HAMCYCLE** is a function problem, since both problems can have multiple solutions. The definition of computational problem can obviously be adapted to non-ASCII alphabets, but we omit those details here.

<sup>1</sup>For NP-complete problems like **HAMCYCLED**, however, this never happens: given a non-decision problem  $C$  whose decision variant is an NP-complete problem  $C'$ , it is known that  $C$  always has a polynomial time reduction to  $C'$ . This property is sometimes known as *self-reducibility* [10].

<sup>2</sup>In contrast, expert practitioners know that sometimes the single-bit decision is useful. For example, determining whether an integer is prime or composite is important in cryptographic applications.



**Figure 1: Students perceive programs solving non-decision problems as considerably more useful than programs solving decision problems. The red bars represent 99% confidence intervals for the mean.**

The above definition has an intuitive connection to computer programs that produce meaningful output (as opposed to a single accept/reject bit). Formally, we say a computer program  $P$  solves the computational problem  $F$  if  $P(x) \in F(x)$  for all  $x$ . That is, the program always terminates and outputs a correct solution.

#### 4 EMPIRICAL EVIDENCE FOR PEDAGOGICAL BENEFIT OF NON-DECISION PROBLEMS

The core motivation for our approach is the well-established pedagogical principle that students learn new concepts more effectively when those concepts are placed in a familiar context. The effectiveness of learning is further enhanced if the new concept is perceived as *useful* or *applicable* [7]. Non-decision problems conform to these two criteria—familiarity and applicability—much more closely than decision problems. The familiarity criterion is indisputable, since almost every computer program written by students, beginning with the first programming course and continuing throughout their careers, computes meaningful answers to general problems rather than producing a single-bit accept/reject decision. The second criterion of perceived usefulness or applicability also seems plausible, but requires supporting evidence. To investigate this, a sample of 41 computer science majors were given descriptions of four computer programs which solved decision and non-decision variants of two different problems. Participants rated the usefulness of each program on a Likert scale from 1 (extremely useful) to 5 (not at all useful). The order of presentation (decision versus non-decision variants) was varied to eliminate ordering effects. We control for the intrinsic perceived usefulness of any given problem by presenting decision and non-decision variants of the *same* problems. A total of  $n = 81$  paired (i.e. decision versus non-decision) responses were received from the 41 participants. Further details of the survey design, demographics, and results are available in an accompanying technical report [15].

The results show that programs solving non-decision problems are perceived as significantly more useful than programs solving decision problems (mean 1.8 vs 3.1 on the Likert scale, with standard deviations of about 0.7 and 0.8 respectively, leading to standard errors of less than 0.1 in both cases). A 99% confidence interval for the population mean, formed from plus-or-minus three standard errors, is shown on figure 1; the intervals suggest a large and significant difference. Two more rigorous tests confirm this: (i) a Wilcoxon signed-rank test for asymmetry of the paired differences has a negligible  $p$ -value ( $p < 10^{-11}$ ); (ii) the same test run on shifted data,

with the non-decision responses shifted away from “useful” by an entire Likert gradation, also has a negligible  $p$ -value ( $p < 10^{-6}$ ). We conclude that the difference in student perceptions of usefulness is rather large—certainly more than one Likert gradation and hence exceeding the distinction between “very useful” and “mildly useful,” as shown on figure 1.

It is worth emphasizing here that this difference in perceived usefulness exists in a *novice audience* (as defined in section 1), and this underpins the key pedagogical point of the paper. Experienced practitioners know that decision programs can often be converted to equivalent non-decision programs with only a logarithmic increase in running time. But because the novice audience does not share this intuition, we should exploit the novices’ perception that non-decision programs are useful. We conclude that elementary concepts in an introductory computer science theory course should be taught using non-decision problems whenever possible.

## 5 COMPUTABILITY FOR NON-DECISION PROBLEMS

The previous section established the key pedagogical advantages of non-decision problems: familiarity and applicability. The remainder of the paper describes technical details of how to teach the standard material of a CS theory course using non-decision problems. This brief section tackles *computability theory*, and the following section tackles *complexity theory*.

We say a computational problem is *computable* if there exists a program that solves it. Note that computability is a generalization of decidability, which applies only to decision problems. Most undergraduate theory treatments use *undecidability* as the central concept to convey the profound idea that “there are some problems that computers can’t solve.” In the approach advocated by this paper, the notion of *uncomputability* replaces undecidability as the central concept. We can still use classical decision problems as examples of uncomputable problems (e.g. the halting problem, or the question of whether a given program computes a given function). But in addition, we can discuss other interesting uncomputable problems that are not decision problems (see [15, 16]). In all cases, we retain the pedagogical advantage of working in a framework that is perceived by the novice audience as familiar and applicable.

## 6 ELEMENTARY COMPLEXITY THEORY FOR NON-DECISION PROBLEMS

We now move from computability theory to complexity theory. Here, the use of non-decision problems as the central focus requires some new notation and terminology, compared to the traditional approach. The most fundamental classical complexity classes (P, NP, Exp) contain only decision problems, so we need new notation for analogous classes that contain both decision and non-decision problems. Here, we will denote these new classes by Poly, NPoly, and Expo respectively. Formally, then, Poly is the set of computational problems for which there exists a program (or Turing machine) that solves the problem in polynomial time. NPoly is the same but we

allow nondeterministic programs<sup>3</sup>; and Expo is the same but we allow exponential time.

The literature already has classes FP and FNP, which are similar in spirit to Poly and NPoly respectively—so we need good justification for introducing new terminology into such well-trodden territory. One reason is that FP is sometimes defined in terms of function problems only (which excludes problems with multiple solutions such as FACTOR and HAMCYCLE). And even when FP is defined in terms of general problems, it is usually done via the polynomially bounded relations mentioned above, which is unnecessarily complicated for the novice audience. Therefore, we prefer to introduce the new classes Poly, NPoly, and Expo.

A comparison with our ten sample books is interesting here: five do not mention FP, FNP, or any analogous complexity class [5, 11–13, 24]; three briefly cover one or both of FP and FNP [2, 18, 21]; and two give more thorough analyses of FP and FNP or analogous classes (Papadimitriou [19], Goldreich [10]). Papadimitriou’s definitions of FP and FNP are closest to the Poly and NPoly defined here. However, they are not equivalent, because of the restriction to polynomially balanced relations mentioned above. Goldreich is the only one of the ten sample books to adopt non-decision problems as a central concept. Goldreich defines the interesting non-decision complexity classes PF (for “polynomial time find”) and PC (for “polynomial time checkable”), but they are again not equivalent to Poly and NPoly. Moreover, PF is not a subset of PC (in contrast to the intuitive relationships  $P \subseteq NP$ ,  $\text{Poly} \subseteq \text{NPoly}$ ,  $\text{FP} \subseteq \text{FNP}$ ), making PF and PC less than ideal for the novice audience. The overall conclusion from analyzing the selection of textbooks is twofold: (i) few textbooks discuss non-decision problems in any detail, even via the well-established complexity classes FP and FNP; and (ii) for technical reasons, the various non-decision classes FP, FNP, PF, and PC are unsuitable for the novice audience and we instead recommend Poly, NPoly, and Expo as defined above.

Introducing novice audiences to complexity theory via the non-decision classes Poly, NPoly, and Expo has a striking advantage that is worth discussing further: the concrete impact of polynomial time algorithms on cryptography is more obvious. To see this, note that one way to crack the popular RSA cryptosystem is to factor a large integer. So, if it turned out to be true that  $\text{FACTOR} \in \text{Poly}$ , the extraordinary consequences (namely, that RSA would be vulnerable to attack) are immediately obvious to the novice audience. In contrast, the consequences are unclear when the same concepts are taught in terms of decision problems: we already know that  $\text{FACTOR} \in P$ , but that doesn’t help to crack RSA since we get only the existence and not the value of the factors. There are ways of rephrasing factorization in terms of decision problems,<sup>4</sup> but the extra technical complexity of rephrasing the problem obscures the key point for novice audiences.

What about more abstract problems such as HAMCYCLE or SAT? In these cases, the pedagogical advantages of familiarity and applicability are compelling. To see why, first note that both HAMCYCLE and SAT can be reduced fairly easily to their decision variants.

<sup>3</sup>The details of defining the outputs of nondeterministic programs that can produce more than a yes/no solution are interesting and important, but are omitted here for space reasons. See [16] for details.

<sup>4</sup>Specifically, we can ask if there exists a factor within a given range. The search problem of finding a factor reduces to this decision problem, via binary search.

Therefore, to an experienced practitioner, it makes absolutely no difference whether we discuss algorithms for `HAMCYCLE` or `HAMCYCLED` (and similarly for `SAT` vs `SATD`). But to the novice audience, the act of finding a Hamilton cycle or satisfying assignment is much more compelling than determining their existence. This distinction becomes especially important when discussing nondeterminism, since nondeterministic programs can easily and efficiently compute factors, Hamilton cycles, and satisfying assignments. In fact, a beneficial and enlightening homework assignment is for students to write multithreaded programs that compute solutions to these problems in nondeterministic polynomial time. As emphasized in section 4, the fact that the outputs of such programs are perceived as “useful” is an important factor in achieving positive learning outcomes.

## 7 THE VERIFIER-BASED DEFINITION OF NPOLY

It is well-known that NP has two equivalent definitions: (i) decision problems that can be decided by a polynomial time nondeterministic program; and (ii) decision problems whose positive instances can be verified in polynomial time by a deterministic program, when provided with a suitable *certificate* (also known as a *witness*, or *hint*). Most modern textbooks cover both definitions, usually introducing (ii) first and later proving the equivalence to (i). Can we generalize these definitions to non-decision problems in a way that appeals to the novice audience? Yes we can, but with some caveats.

Definition (i), based on nondeterminism, was discussed in the previous section. Here we focus on definition (ii), based on verification. According to this definition, NPoly is the class of computational problems which have polynomial time verifiers. And NP is defined in a precisely analogous way: it is the class of *decision* problems which have polynomial time verifiers. When we move to defining *verifier*, however, there are key differences in the decision (NP) and non-decision (NPoly) scenarios. Details of these differences can be found in the accompanying technical report [15]. Here, we mention just one key aspect. Recall that every positive instance of an NP problem must be verifiable in polynomial time when provided with some certificate  $c$ . Typically, the certificate  $c$  encodes a solution to the underlying non-decision problem. For example, a `HAMCYCLED` instance can be verified by providing a legitimate Hamilton cycle as the certificate. But unfortunately, it turns out that certificates for NP problems can also work in less intuitive ways. For example, it is possible to define a `HAMCYCLED` verifier that uses only partial Hamilton cycles as certificates. This discrepancy between certificates and underlying solutions is a source of confusion for the novice audience.

It turns out that if we work with non-decision problems, this potentially confusing “looseness” in the definition of certificates is eliminated. In essence (and please see the technical report [15] or textbook [16] for additional details), the certificate  $c$  is replaced by two separate strings: a solution  $s$  and hint  $h$ . Both strings are given as input to the verifier. For successful verifications of positive instances, the solution  $s$  really is a solution for the given instance; the hint  $h$  provides any additional information needed for verification. The explicit roles of  $s$  and  $h$  provide additional clarity to the novice audience.

To the best of our knowledge, no textbook other than [16] employs this formulation, which ensures that verified solutions are meaningful while still permitting hints for problems that need them. This is one key difference between the approach of this paper and that of Goldreich [10].

## 8 REDUCTIONS AND NP-COMPLETENESS

We noted in the introduction that some of the theorems and proofs in computational complexity theory are more elegant when expressed in terms of decision problems, and the vast majority of literature takes this approach. Therefore, even if we choose for pedagogical reasons to introduce elementary concepts in terms of non-decision problems, it makes sense to eventually transition to the decision problem viewpoint. Over a four-year period, we experimented with making this transition at various points in the course. For reasons discussed below, we found the best choice is to make the transition when students first encounter polynomial time reductions, and hence also before introducing NP-completeness.

For the novice audience it seems preferable to focus on the simplest type of reduction, which is variously known as a *Karp reduction*, *polynomial time mapping reduction*, or *many-one reduction*. In focusing on Karp reductions, we are now in line with eight of our ten sample textbooks (with the same two outliers [10, 19] as previously). However, we still break slightly from the eight standard treatments, because instead of defining reductions between *decision* problems, we define reductions from decision computational problems to *general* problems (i.e. *to*, but not *from*, general problems). This difference again calls for new terminology, and we use the term “polyreduction” for this.

Formally, a *polyreduction* from the decision problem  $D$  to the general computational problem  $G$  is a map  $r : \Sigma^* \rightarrow \Sigma^*$ , such that  $r(w)$  is a positive instance of  $G$  if and only if  $w$  is a positive instance of  $D$ . We also require that  $r$  is computable in polynomial time. With this definition, students can first examine trivial polyreductions such as `HAMCYCLED`→`HAMCYCLE`, then move on to more interesting ones such as `DIRECTEDHAMCYCLED`→`UNDIRECTEDHAMCYCLED`.

With a little more effort, we can define polyreductions between general problems  $F$  and  $G$ . To do so, we need polynomial time maps in both directions. First,  $r : \Sigma^* \rightarrow \Sigma^*$  maps instances of  $F$  to instances of  $G$ . Then  $r' : \Sigma^* \rightarrow \Sigma^*$  maps solutions of  $G$  back to solutions of  $F$ . The correctness condition is that, given any program  $G$  that solves the problem  $G$ , we must have  $r'(G(r(w))) \in F(w)$  for all  $w$ . (In words, the solution of  $G$  for  $r(w)$  must map back to a solution of  $F$  for  $w$ .) We have experimented with teaching this concept to novice audiences and met with a certain amount of success. Goldreich also recommends this approach [9]. However, as already stated, on balance we recommend avoiding this extra level of generality, instead restricting polyreductions to be maps from decision problems to general problems via the simpler definition in the previous paragraph.

In the same spirit, we might attempt to teach “NPoly-completeness” for general computational problems. But in this case, our experience has been that the benefits of using non-decision problems are clearly outweighed by the disadvantages of dealing with arbitrary solutions when stating and proving theorems about NPoly-completeness. Therefore, it seems preferable to stay solidly in the traditional

realm of decision problems when teaching NP-completeness. As a small bonus, the fact that our definition of polyreduction allows decision problems to be reduced to *general* problems leads to an elegant definition of NP-hardness: a problem  $G$  is NP-hard if some NP-complete problem  $D$  polyreduces to  $G$ . Holistic discussions of P versus NP (although defined formally in terms of decision problems) also take on a more practical tone when the majority of concepts earlier in the course have been taught in terms of non-decision problems.

## 9 CONCLUSION

The paper advocated the use of non-decision problems in CS theory courses. The paper consisted of two separate strands: (i) an explanation of *why* non-decision problems are preferable, based on a combination of empirical survey results and educational theory; and (ii) an explanation of *how* to use non-decision problems in the theory course, using reformulations of classical concepts. In the first strand (section 4), we demonstrated via an empirical survey of CS majors that programs solving non-decision problems are perceived as much more useful and applicable than programs solving decision problems. Invoking the well-known educational principle that learning outcomes are likely to be better for a course that uses materials perceived as “useful,” we concluded that non-decision problems should be used if possible. In the second strand (sections 3 and 5–8), we suggested new definitions, terminology, and notation that are designed to employ non-decision problems and to maximize the accessibility of CS theory concepts for the novice undergraduate audience. The approach has been tested and refined over four years of teaching, culminating in a recent textbook [16]. Looking to the future, we hope this approach will not only improve learning outcomes for undergraduates receiving their first taste of computer science theory, but will also lead to theory courses that are taught earlier in the curriculum, to a wider range of undergraduates, at a wider range of institutions.

## REFERENCES

- [1] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena. 2004. PRIMES is in P. *Annals of Mathematics* 160, 2 (2004), 781–793.
- [2] S. Arora and B. Barak. 2009. *Computational Complexity: A Modern Approach*. Cambridge University Press. <https://books.google.com/books?id=8Wjqvsoo48MC>
- [3] Carlos I. Chesñevar, María L. Cobo, and William Yurcik. 2003. Using Theoretical Computer Simulators for Formal Languages and Automata Theory. *SIGCSE Bulletin* 35, 2 (June 2003), 33–37. <https://doi.org/10.1145/782941.782975>
- [4] Pierluigi Crescenzi, Emma Enström, and Viggo Kann. 2013. From Theory to Practice: NP-completeness for Every CS Student. In *Proc. ITiCSE*. 16–21. <https://doi.org/10.1145/2462476.2465582>
- [5] Martin Davis, Ron Sigal, and Elaine J Weyuker. 1994. *Computability, Complexity, and Languages: Fundamentals of Theoretical Computer Science* (2nd ed.). Morgan Kaufmann.
- [6] Emma Enström and Viggo Kann. 2010. Computer lab work on theory. In *Proc. ITiCSE*. 93–97.
- [7] L Dee Fink. 2013. *Creating significant learning experiences: An integrated approach to designing college courses* (2nd ed.). John Wiley & Sons.
- [8] Lance Fortnow. 2013. *The golden ticket: P, NP, and the search for the impossible*. Princeton University Press.
- [9] Oded Goldreich. 2006. *On Teaching the Basics of Complexity Theory*. Springer, 348–374. [https://doi.org/10.1007/11685654\\_15](https://doi.org/10.1007/11685654_15) In *Theoretical Computer Science: Essays in Memory of Shimon Even*, ed. Goldreich, Rosenberg, Selman.
- [10] Oded Goldreich. 2010. *P, NP, and NP-Completeness: The Basics of Computational Complexity*. Cambridge University Press.
- [11] John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. 2006. *Introduction to Automata Theory, Languages, and Computation* (3rd ed.). Pearson.
- [12] Harry R Lewis and Christos H Papadimitriou. 1997. *Elements of the Theory of Computation* (2nd ed.). Prentice Hall.
- [13] Peter Linz. 2011. *An Introduction to Formal Languages and Automata* (5th ed.). Jones & Bartlett.
- [14] Andrea F. Lobo and Ganesh R. Baliga. 2006. NP-completeness for All Computer Science Undergraduates: A Novel Project-based Curriculum. *J. Comput. Sci. Coll.* 21, 6 (June 2006), 53–63. <http://dl.acm.org/citation.cfm?id=1127442.1127450>
- [15] John MacCormick. 2017. *Strategies for basing the CS theory course on non-decision problems*. Technical Report. Dickinson College.
- [16] John MacCormick. 2018. *What Can Be Computed?: A Practical Guide to the Theory of Computation*. Princeton University Press.
- [17] Dino Mandrioli. 1982. On teaching theoretical foundations of Computer Science. *ACM SIGACT News* 14, 3 (1982), 36–53. Part 2 appears in 14(4):58–69.
- [18] Christopher Moore and Stephan Mertens. 2011. *The Nature of Computation*. Oxford University Press.
- [19] Christos H. Papadimitriou. 1994. *Computational Complexity*. Addison Wesley, Massachusetts.
- [20] Christos H Papadimitriou. 1997. NP-completeness: A retrospective. In *International Colloquium on Automata, Languages, and Programming*. Springer, 2–6.
- [21] Elaine Rich. 2007. *Automata, Computability and Complexity: Theory and Applications*. Pearson.
- [22] Susan H Rodger and Thomas W Finley. 2006. *JFLAP: an interactive formal languages and automata package*. Jones & Bartlett Learning.
- [23] Susan H Rodger, Eric Wiebe, Kyung Min Lee, Chris Morgan, Kareem Omar, and Jonathan Su. 2009. Increasing engagement in automata theory with JFLAP. In *ACM SIGCSE Bulletin*, Vol. 41. ACM, 403–407.
- [24] Michael Sipser. 2013. *Introduction to the Theory of Computation* (3rd ed.). Cengage.