

The Computational Complexity of Understanding Network Decisions

Stephan Wäldchen¹, Jan Macdonald¹, Sascha Hauch¹, and Gitta Kutyniok^{1,2}

¹Institut für Mathematik, Technische Universität Berlin

²Fakultät Elektrotechnik und Informatik, Technische Universität Berlin
{stephanw, macdonald, hauch, kutyniok}@math.tu-berlin.de

May 22, 2019

Abstract

For a Boolean function $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$ and an assignment to its variables $\mathbf{x} = (x_1, x_2, \dots, x_d)$ we consider the problem of finding the subsets of the variables that are sufficient to determine the function value with a given probability δ . This is motivated by the task of interpreting predictions of binary classifiers described as Boolean circuits (which can be seen as special cases of neural networks).

We show that the problem of deciding whether such subsets of relevant variables of limited size $k \leq d$ exist is complete for the complexity class NP^{PP} and thus generally unfeasible to solve. We introduce a variant where it suffices to check whether a subset determines the function value with probability at least δ or at most $\delta - \gamma$ for $0 < \gamma < \delta$. This reduces the complexity to the class NP^{BPP} .

Finally, we show that finding the minimal set of relevant variables can not be reasonably approximated, i.e. with an approximation factor $d^{1-\alpha}$ for $\alpha > 0$, by a polynomial time algorithm unless $\text{P} = \text{NP}$ (this holds even with the probability gap).

1. Introduction

Algorithmic problem solving in real-world scenarios often requires reasoning in an uncertain environment. This necessity lead to the investigation of probabilistic satisfiability problems and probabilistic computational complexity classes such as PP and NP^{PP} . One prototypical example, the E-MAJ-SAT problem [14, 15], is an extension of the classical satisfiability problem that includes an element of a model counting formulation. The class of NP^{PP} -complete problems contains many relevant artificial intelligence (AI) problems such as probabilistic conformant planning [4, 13], calculating maximum expected utility (MEU) solutions [2], and maximum a posteriori (MAP) hypotheses [17].

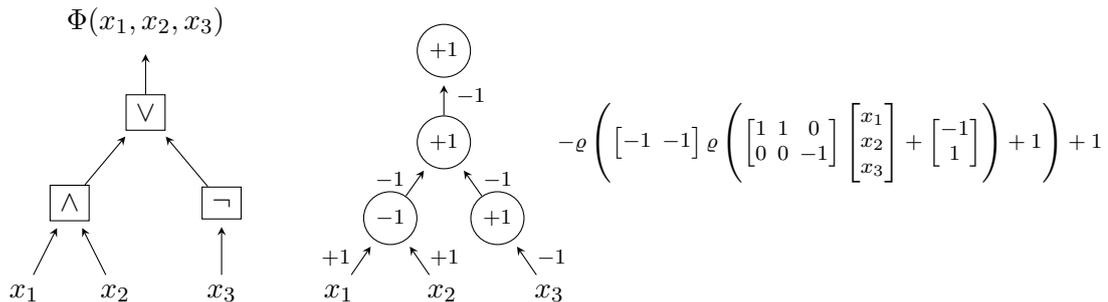


Figure 1: The Boolean function $\Phi(x_1, x_2, x_3) = (x_1 \wedge x_2) \vee (\neg x_3)$ viewed as a Boolean circuit (left) and a rectified linear unit (ReLU) neural network in its graphical (middle) and algebraic representation (right). Network weights and biases are denoted at the edges and nodes respectively. The ReLU activation $\varrho(x) = \max\{x, 0\}$ is applied components-wise.

We connect these probabilistic reasoning tasks to an important problem in machine learning, namely the problem of interpreting the decisions of neural network classifiers. Neural networks are parameter-rich and highly non-linear models and can be seen as continuous generalisations of Boolean circuits. This is briefly visualised in Figure 1. They have achieved impressive success in classification [7, 12, 22] and regression tasks [21, 24, 23] and are increasingly also used to solve various inverse problems [10, 25].

However, the same expressiveness that allows for hierarchical reasoning and universal approximation makes understanding and interpreting these models more challenging compared to traditional machine learning methods like linear regression or decision trees. Treating neural networks as “black box” solvers without accessible reasoning is not feasible in many circumstances, for example in critical applications such as medical imaging and diagnosis [16, 19].

A significant first step towards understanding network decisions is to distinguish the relevant input parameters from the less relevant ones for a specific prediction. This goal has been pursued predominantly for image classification problems in the form of visual maps that assign importance values to the inputs variables, for example in [1, 5, 20, 26].

We formalise this notion as a probabilistic decision problem in the following sense. Given a Boolean function and an assignment to its variables, is there a small set of input variables that if held constant determines the function value for almost all possible assignments to the rest of the variables?

This problem should be contrasted with the feature selection problem, where the goal is to find a subset of globally important variables, independent from a concrete input. Instead, here we are interested in finding the important variables only for one specific input assignment.

We formulate the decision and minimisation version of the problem at hand and compare it to related problems in Section 2. We then analyse its computational complexity in Section 3 and prove that it is NP^{PP} -complete. We propose a more practically relevant

variant of the problem with a gap promise that allows for efficient computations of expectation values in Section 4. Analysing its complexity will finally show that any reasonable approximation algorithm to the minimisation version of the problem cannot be polynomial time unless $P = NP$.

Notation Throughout the paper $d \in \mathbb{N}$ denotes the arity of the considered Boolean function $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$ and $\mathbf{x} = (x_1, \dots, x_d) \in \{0, 1\}^d$ is an arbitrary fixed assignment to its variables for which we are interested in finding the subsets of important variables. We denote the d -dimensional vectors of all zeros or ones by $\mathbf{0}_d$ and $\mathbf{1}_d$ respectively. The function Φ is assumed to be described in terms of standard logical operations like AND, OR, and NOT. The description length of Φ is thus equal to the description length of the logical expression describing it. We denote $[d] = \{1, \dots, d\}$ and for a subset $S \subseteq [d]$ denote by $\mathbf{x}_S = (x_i)_{i \in S}$ the restriction of \mathbf{x} to components indexed by S . Further, we will use Boolean functions also interchangeably as logical propositions, in the sense that $\Phi(\mathbf{x})$ is shorthand for the logical proposition $\Phi(\mathbf{x}) = 1$. Whenever we talk about statements concerning probabilities of logical propositions to hold we assume independent uniform distributions for all involved variables. Thus, we have for example

$$P_{\mathbf{y}}(\Phi(\mathbf{y})) = \frac{|\{ \mathbf{y} \in \{0, 1\}^d : \Phi(\mathbf{y}) = 1 \}|}{|\{ \mathbf{y} \in \{0, 1\}^d \}|}.$$

We omit the subscript whenever it is clear from the context over which variables the probability is taken. If the probability is taken over all variables of a Boolean function we simply write $P(\Phi)$ instead of $P_{\mathbf{y}}(\Phi(\mathbf{y}))$.

2. Problem Formulation

Intuitively, a subset $S \subseteq [d]$ of variables is *relevant* for the function value $\Phi(\mathbf{x})$ if fixing \mathbf{x} on S and randomising it on the complement S^c does not change the value of Φ with high probability. The complement then consists of the *non-relevant* variables.

Definition 2.1. Let $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$, $\mathbf{x} \in \{0, 1\}^d$, and $\delta \in [0, 1]$. We call $S \subseteq [d]$ a δ -*relevant* set for Φ and \mathbf{x} , if

$$P_{\mathbf{y}}(\Phi(\mathbf{y}) = \Phi(\mathbf{x}) \mid \mathbf{y}_S = \mathbf{x}_S) \geq \delta.$$

For δ close to one this means that the input \mathbf{x} supported on S already determines the output $\Phi(\mathbf{x})$ with high probability. It is clear that $S = [d]$ is always 1-relevant and any subset $S \subseteq [d]$ is 0-relevant. Now the interesting question arises if for a given δ there exists a δ -relevant set of a certain maximal size. Similarly, one could ask to find the smallest δ -relevant set. This set would then be composed of the most important variables for the function value $\Phi(\mathbf{x})$. There is an obvious “rate-distortion” trade-off in the sense that generally a larger δ will require a larger set S .

Definition 2.2. For $\delta \in (0, 1]$ we define the RELEVANT-INPUT problem as follows.

GIVEN: $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$, $\mathbf{x} \in \{0, 1\}^d$, and $k \in \mathbb{N}$, $1 \leq k \leq d$.

DECIDE: Does there exist $S \subseteq [d]$ with $|S| \leq k$ such that S is δ -relevant for Φ and \mathbf{x} ?

The minimisation formulation of the above decision problem can be defined in the obvious way.

Definition 2.3. For $\delta \in (0, 1]$ we define the MIN-RELEVANT-INPUT problem as follows.

GIVEN: $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$ and $\mathbf{x} \in \{0, 1\}^d$.

MINIMIZE: $k \in \mathbb{N}$ such that there exists $S \subseteq [d]$ with $|S| \leq k$ and S is δ -relevant for Φ and \mathbf{x} .

The majority of the remainder of the paper will deal with analysing the computational complexity of the RELEVANT-INPUT and MIN-RELEVANT-INPUT problems and related variants thereof. Before, we want to point out some similarities to a problem from cooperative game theory.

2.1. Related Works

Another concept for measuring the relevance or the *contribution* of individual variables to a collective are the Shapley values [18] in cooperative game theory. Here the variables are seen as players of a coalitional game and the Shapley values describe a method to distribute the value achieved by a coalition of players to the individual players. This distribution fulfils a set of game theoretic properties that make it “fair”.

Let $\nu: 2^{[d]} \rightarrow \mathbb{R}$ be a function that assigns a value to each subset of variables (coalition of players). It is called the *characteristic function* of the game. Then the Shapley value of the i -th variable (i -th player) is defined as

$$\varphi_{i,\nu} = \sum_{S \subseteq [d] \setminus \{i\}} \frac{|S|!(d - |S| - 1)!}{d!} (\nu(S \cup \{i\}) - \nu(S)),$$

which can be interpreted as the marginal contribution of the i -th variable to the value ν averaged over all possible coalitions. In general it is #P-hard to compute Shapley values [3]. However, in some cases efficient approximation algorithms exist [6].

In our scenario the value of a subset of variables S can be measured by the expected difference in Φ when fixing variables in S and randomising the remaining variables. In [11] it was proposed to use

$$\nu(S) = \frac{1}{2^{d-|S|}} \sum_{\substack{\mathbf{y} \in \{0,1\}^d \\ \mathbf{y}_S = \mathbf{x}_S}} \Phi(\mathbf{y}) - \mathbb{E}_{\mathbf{y}}(\Phi(\mathbf{y}))$$

for the analysis of classifier decisions, which uses the expectation of the completely randomised classifier score as a reference value to determine the coalition value. We observe that

$$\begin{aligned} P_{\mathbf{y}}(\Phi(\mathbf{y}) = \Phi(\mathbf{x}) \mid \mathbf{y}_S = \mathbf{x}_S) &= 1 - \frac{1}{2^{d-|S|}} \sum_{\substack{\mathbf{y} \in \{0,1\}^d \\ \mathbf{y}_S = \mathbf{x}_S}} |\Phi(\mathbf{y}) - \Phi(\mathbf{x})| \\ &= 1 - |\nu(S) + \mathbb{E}_{\mathbf{y}}(\Phi(\mathbf{y})) - \Phi(\mathbf{x})|, \end{aligned}$$

hence $S \subseteq [d]$ is δ -relevant for Φ and \mathbf{x} exactly if $|\nu(S) + \mathbb{E}_{\mathbf{y}}(\Phi(\mathbf{y})) - \Phi(\mathbf{x})| \leq 1 - \delta$.

Despite this relation between δ -relevant sets and the characteristic function ν our problem formulation is considerably different from the Shapley value approach. The task considered in this paper is not to distribute the value of coalitions amongst the variables but to find (small) coalitions that are guaranteed to have a certain value. We will see that unlike for Shapley values no efficient approximation algorithms can exist for this problem unless $\text{P} = \text{NP}$.

3. Computational Complexity Analysis

The first main theorem shows that the RELEVANT-INPUT problem is generally hard to solve for $\delta \in [\frac{1}{2}, 1)$.

Theorem 3.1. *For $\delta \in [\frac{1}{2}, 1)$ the RELEVANT-INPUT problem is NP^{PP} -complete.*

The proof of Theorem 3.1 will be split into two parts. We will show that RELEVANT-INPUT is NP^{PP} -hard in Section 3.1 and that it is contained in NP^{PP} in Section 3.2. Before we continue to give the proof we want to make a few more observations and remarks.

Intuitively, the NP -part of the problem complexity comes from the necessity to check all subsets $S \subseteq [d]$ as possible candidates for being δ -relevant. The PP -part of the complexity comes from the fact that for any given set S checking if it is δ -relevant is by itself a hard problem (in fact PP -hard)¹. The problem class NP^{PP} is beyond the scope of conventional computing. In particular, MIN-RELEVANT-INPUT is at least as hard to solve as the corresponding decision problem, which makes it unfeasible to solve exactly. However, in applications it is rarely required to exactly find the smallest relevant set. It would be desirable to obtain good approximate solutions within feasible computational complexity.

There are two potential ways for simplifying the problem by allowing approximations: Firstly, approximating the size of the minimal relevant set itself, and secondly relaxing the requirement that the set has to be exactly δ -relevant. The former would address the NP part whereas the latter would address the PP aspect.

Calculating probabilities or expectation values may be hard in theory, yet it is often easy to calculate them (approximately) in practice, for example by sampling. Checking

¹Checking if a subset is 1-relevant is in coNP instead of PP . Thus we excluded $\delta = 1$ in Theorem 3.1.

whether a logical proposition is satisfied with probability more than δ this way only fails if the true probability can be arbitrarily close to δ . These edge cases cause the hardness of the problem but in our scenario we do not necessarily care about their resolution. We discuss in Section 4.1 that this can be relaxed formally by stating a promise problem. This reduces the problem complexity from PP to BPP. Unfortunately, even in this simplified case, it remains NP-hard to approximate the size of the optimal set S to within any reasonable approximation factor. This is discussed in Section 4.2.

3.1. Relevant-Input is NP^{PP} -hard

We now give the first part of the proof of Theorem 3.1. This is done by constructing a polynomial-time reduction of a NP^{PP} -complete problem to RELEVANT-INPUT. The canonical complete problem for NP^{PP} is E-MAJ-SAT [14].

Definition 3.2. The E-MAJ-SAT problem is defined as follows.

GIVEN: $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$ and $k \in \mathbb{N}$, $1 \leq k \leq d$.

DECIDE: Does there exist $\mathbf{x} \in \{0, 1\}^k$ such that $P_{\mathbf{y}}(\Phi(\mathbf{y}) \mid \mathbf{y}_{[k]} = \mathbf{x}) > \frac{1}{2}$?

In other words E-MAJ-SAT asks if there is an assignment to the first k variables of Φ such that a majority of assignments to the remaining $d - k$ variables satisfies Φ . There are three hurdles to take if we want to reduce this to RELEVANT-INPUT.

1. Instead of assigning values to a given set of k variables we can freely choose the set S of size k .
2. Instead of freely assigning values to a subset of variables we are given an assignment to all variables and can only choose to fix a subset of them and randomise the rest.
3. Instead of checking whether the majority of assignments satisfies Φ we check if the fraction of satisfying assignments is larger or equal to some δ .

We address each of these hurdles and give a chain of polynomial-time reductions $\text{E-MAJ-SAT} \preceq_p \text{IP1} \preceq_p \text{IP2} \preceq_p \text{RELEVANT-INPUT}$ in three steps with intermediate auxiliary problems IP1 and IP2. The following observations will turn out to be useful.

Remark 3.3. Let Φ and Ψ be Boolean functions, not necessarily of different variables. Then

$$\begin{aligned} P(\Psi) = 0 &\Rightarrow P(\Phi \oplus \Psi) = P(\Phi), \\ P(\Psi) = 1 &\Rightarrow P(\Phi \oplus \Psi) = 1 - P(\Phi), \end{aligned}$$

and if Φ and Ψ are independent, i.e. $P(\Phi \wedge \Psi) = P(\Phi)P(\Psi)$, also

$$P(\Psi) = \frac{1}{2} \Rightarrow P(\Phi \oplus \Psi) = \frac{1}{2}.$$

Lemma 3.4. *Let $B: \{0, 1\}^k \times \{0, 1\}^k \rightarrow \{0, 1\}$ and $\Psi: \{0, 1\}^k \times \{0, 1\}^k \times \{0, 1\} \rightarrow \{0, 1\}$ be defined as*

$$B(\mathbf{u}, \mathbf{v}) = \bigwedge_{i=1}^k \neg(u_i \oplus v_i)$$

$$\Psi(\mathbf{u}, \mathbf{v}, t) = \left(\bigvee_{i=1}^k (u_i \oplus v_i) \right) \wedge t.$$

Then, for any $\Phi: \{0, 1\}^k \times \{0, 1\}^{d-k} \rightarrow \{0, 1\}$ and $A: \{0, 1\}^k \times \{0, 1\}^k \rightarrow \{0, 1\}$, we have

$$P(\Phi(\mathbf{u}, \mathbf{r}) \oplus \Psi(\mathbf{u}, \mathbf{v}, t) \mid A(\mathbf{u}, \mathbf{v})) > \frac{1}{2} \iff P(\Phi(\mathbf{u}, \mathbf{r}) \mid A(\mathbf{u}, \mathbf{v}), B(\mathbf{u}, \mathbf{v})) > \frac{1}{2}.$$

Proof. We can rewrite $\Psi(\mathbf{u}, \mathbf{v}, t) = (\neg B(\mathbf{u}, \mathbf{v})) \wedge t$ and therefore

$$P(\Psi \mid B) = 0$$

$$P(\Psi \mid \neg B) = \frac{1}{2}.$$

Since $\Phi \mid A$ and $\Psi \mid A$ are conditionally independent given $\neg B$ (in this case Ψ depends on t only), we obtain from Remark 3.3 that

$$P(\Phi \oplus \Psi \mid A, B) = P(\Phi \mid A, B)$$

and

$$P(\Phi \oplus \Psi \mid A, \neg B) = \frac{1}{2}.$$

Therefore

$$\begin{aligned} P(\Phi \oplus \Psi \mid A) &= P(\Phi \oplus \Psi \mid A, B) P(B) + P(\Phi \oplus \Psi \mid A, \neg B) P(\neg B) \\ &= P(\Phi \mid A, B) P(B) + \frac{1}{2} (1 - P(B)) \\ &= \frac{1}{2} + \left(P(\Phi \mid A, B) - \frac{1}{2} \right) P(B). \end{aligned}$$

This directly implies $P(\Phi \oplus \Psi \mid A) > \frac{1}{2}$ if and only if $P(\Phi \mid A, B) > \frac{1}{2}$. □

Let us now come to the first step of the reductive chain. In this, we translate the option to freely assign the first k variables into the choice of fixing a set of variables from a given assignment or randomising them. This choice is however still restricted to the first k variables.

Definition 3.5. We define the INTERMEDIATE PROBLEM 1 (IP1) as follows.

GIVEN: $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$, $\mathbf{x} \in \{0, 1\}^d$ and $k \in \mathbb{N}$, $1 \leq k \leq d$.

DECIDE: Does there exist $S \subseteq [k]$ such that $P_{\mathbf{y}}(\Phi(\mathbf{y}) \mid \mathbf{y}_S = \mathbf{x}_S) > \frac{1}{2}$?

In other words IP1 asks the questions if there is a subset of the first k variables of Φ , such that when fixing these to the values given by \mathbf{x} , a majority of assignments to the remaining variables satisfies Φ .

Lemma 3.6. E-MAJ-SAT \preceq_p IP1, in particular IP1 is NPP^{P} -hard.

Proof. Let $\{\Phi, k\}$ be an E-MAJ-SAT instance. We will construct $\{\Phi', \mathbf{x}', k'\}$ that is a *Yes*-instance for IP1 if and only if $\{\Phi, k\}$ is a *Yes*-instance for E-MAJ-SAT. For convenience we split the d variables of Φ into the first k variables and the remaining $d - k$ variables and denote this $\Phi(\mathbf{x}) = \Phi(\mathbf{u}, \mathbf{r})$. The main idea is to duplicate the first k variables and choose \mathbf{x}' in such a way that fixing the original variables or their duplicates corresponds to assigning zeros or ones in the E-MAJ-SAT instance respectively. More precisely, we define

- $\Phi': \{0, 1\}^k \times \{0, 1\}^k \times \{0, 1\}^{d-k} \times \{0, 1\} \rightarrow \{0, 1\}$ as

$$\Phi'(\mathbf{u}, \mathbf{v}, \mathbf{r}, t) = \Phi(\mathbf{u}, \mathbf{r}) \oplus \left(\bigvee_{i=1}^k (u_i \oplus v_i) \wedge t \right),$$

- $\mathbf{x}' = (\mathbf{0}_k, \mathbf{1}_k, \mathbf{0}_{d-k}, 0) \in \{0, 1\}^k \times \{0, 1\}^k \times \{0, 1\}^{d-k} \times \{0, 1\}$,
- $k' = 2k$.

This is a polynomial time construction. With Ψ as in Lemma 3.4 we can rewrite $\Phi'(\mathbf{u}, \mathbf{v}, \mathbf{r}, t) = \Phi(\mathbf{u}, \mathbf{r}) \oplus \Psi(\mathbf{u}, \mathbf{v}, t)$.

Necessity: Assume that $\{\Phi, k\}$ is a *Yes*-instance for E-MAJ-SAT. Then there exists an assignment $\mathbf{u}^* \in \{0, 1\}^k$ to the first k variables of Φ such that $P_{\mathbf{r}}(\Phi(\mathbf{u}^*, \mathbf{r})) > \frac{1}{2}$. Now choose $S' = \{i \in \{1, \dots, k\} : u_i^* = 0\} \cup \{i \in \{k+1, \dots, 2k\} : u_{i-k}^* = 1\} \subseteq [k'] = [2k]$. Let $A: \{0, 1\}^k \times \{0, 1\}^k \rightarrow \{0, 1\}$ be given by

$$A(\mathbf{u}, \mathbf{v}) = \left(\bigwedge_{i \in S' \cap \{1, \dots, k\}} \neg u_i \right) \wedge \left(\bigwedge_{i \in S' \cap \{k+1, \dots, 2k\}} v_{i-k} \right)$$

and B as in Lemma 3.4. Note that A depends on S' and thus implicitly on \mathbf{u}^* . In fact, $A(\mathbf{u}, \mathbf{v}) = 1$ holds if and only if $u_i^* = 0$ implies $u_i = 0$ and $u_i^* = 1$ implies $v_i = 1$ for all $i \in [k]$. In particular, we have $A(\mathbf{u}, \mathbf{u}) = 1$ if and only if $\mathbf{u} = \mathbf{u}^*$. Also $B(\mathbf{u}, \mathbf{v}) = 1$ if and only if $\mathbf{u} = \mathbf{v}$. Thus, by the choice of A , B , and S' we have

$$\begin{aligned} P_{\mathbf{r}}(\Phi(\mathbf{u}^*, \mathbf{r})) &= P(\Phi(\mathbf{u}, \mathbf{r}) \mid \mathbf{u} = \mathbf{u}^*) \\ &= P(\Phi(\mathbf{u}, \mathbf{r}) \mid A(\mathbf{u}, \mathbf{u})) \\ &= P(\Phi(\mathbf{u}, \mathbf{r}) \mid A(\mathbf{u}, \mathbf{v}), B(\mathbf{u}, \mathbf{v})), \end{aligned}$$

and by the choice of \mathbf{x}' , A , and S' we get

$$\begin{aligned} P_{\mathbf{y}'}(\Phi'(\mathbf{y}') | \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) &= P(\Phi'(\mathbf{u}, \mathbf{v}, \mathbf{r}, t) | A(\mathbf{u}, \mathbf{v})) \\ &= P(\Phi(\mathbf{u}, \mathbf{r}) \oplus \Psi(\mathbf{u}, \mathbf{v}, t) | A(\mathbf{u}, \mathbf{v})). \end{aligned}$$

Since $P_{\mathbf{r}}(\Phi(\mathbf{u}^*, \mathbf{r})) > \frac{1}{2}$, we can use Lemma 3.4 to conclude $P_{\mathbf{y}'}(\Phi'(\mathbf{y}') | \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) > \frac{1}{2}$ which shows that $\{\Phi', \mathbf{x}', k'\}$ is a *Yes*-instance for IP1.

Sufficiency: Now, conversely, assume that $\{\Phi', \mathbf{x}', k'\}$ is a *Yes*-instance for IP1. Then there exists $S' \subseteq [k'] = [2k]$ such that $P_{\mathbf{y}'}(\Phi'(\mathbf{y}') | \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) > \frac{1}{2}$. Following the same grouping of variables as before, we write $\mathbf{x}' = (\mathbf{u}', \mathbf{v}', \mathbf{r}', t')$ and observe that $u'_i \neq v'_i$ for all $i \in [k]$. Towards a contradiction, assume there exists an $i \in [k]$ with $i \in S'$ and $i + k \in S'$. Then $u_i \neq v_i$ if $(\mathbf{u}, \mathbf{v})_{S'} = (\mathbf{u}', \mathbf{v}')_{S'}$ and hence

$$P(\Psi(\mathbf{u}, \mathbf{v}, t) | (\mathbf{u}, \mathbf{v})_{S'} = (\mathbf{u}', \mathbf{v}')_{S'}) = P(\Psi(\mathbf{u}, \mathbf{v}, t) | \neg B(\mathbf{u}, \mathbf{v})) = \frac{1}{2}.$$

Thus, Remark 3.3 would imply $P_{\mathbf{y}'}(\Phi'(\mathbf{y}') | \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) = \frac{1}{2}$, which is a contradiction. Therefore for all $i \in [k]$ we can have $i \in S'$ or $i + k \in S'$ but not both. Similarly, if we assume there exists an $i \in [k]$ with neither $i \in S'$ nor $i + k \in S'$ we have

$$P_{\mathbf{y}'}(\Phi'(\mathbf{y}') | \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) = P_{\mathbf{y}'}(\Phi'(\mathbf{y}') | \mathbf{y}'_{S' \cup \{i+k\}} = \mathbf{x}'_{S' \cup \{i+k\}}).$$

This holds because the only difference is that the clause $u_i \oplus v_i$ in Ψ is replaced by $u_i \oplus 1 = \neg u_i$ and this does not change the overall probability. So without loss of generality we can assume that for each $i \in [k]$ exactly one of the cases $i \in S'$ or $i + k \in S'$ occurs. Then we can define $\mathbf{u}^* \in \{0, 1\}^k$ as $u_i^* = 0$ if $i \in S'$ and $u_i^* = 1$ otherwise. We observe that S' and \mathbf{u}^* are exactly as in the previous step and the rest of the proof follows analogously. Again we use Lemma 3.4 and conclude from $P_{\mathbf{y}'}(\Phi'(\mathbf{y}') | \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) > \frac{1}{2}$ that $P_{\mathbf{r}}(\Phi(\mathbf{u}^*, \mathbf{r})) > \frac{1}{2}$. This shows that $\{\Phi, k\}$ is a *Yes*-instance for E-MAJ-SAT. \square

We continue with the second step of the reductive chain. We translate the question whether a majority of assignments to the non-fixed variables satisfies a Boolean formula into the question whether the fraction of assignments leaving the function value unchanged is larger or equal to some δ .

Definition 3.7. For $\delta \in (0, 1]$ we define the INTERMEDIATE PROBLEM 2 (IP2) as follows.

GIVEN: $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$, $\mathbf{x} \in \{0, 1\}^d$ and $k \in \mathbb{N}$, $1 \leq k \leq d$.

DECIDE: Does there exist $S \subseteq [k]$ such that $P_{\mathbf{y}}(\Phi(\mathbf{y}) | \mathbf{y}_S = \mathbf{x}_S) \geq \delta$?

In other words IP2 asks the question if there exists a subset of the first k variables of Φ such that when fixing these to the values given by \mathbf{x} a fraction δ of assignments to the remaining variables leaves Φ unchanged. In order to relate this to IP1, we have to change the acceptance probability from $\frac{1}{2}$ to δ . This can be done using the following Lemma.

Lemma 3.8. *Given $d \in \mathbb{N}$ and $0 < \delta_1 < \delta_2 < 1$, there exists a monotone function $\Pi: \{0, 1\}^n \rightarrow \{0, 1\}$ such that $\Pi(\mathbf{0}_n) = 0$, $\Pi(\mathbf{1}_n) = 1$, and for all $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$ we have*

$$P_{\mathbf{y}}(\Phi(\mathbf{y})) > \delta_1 \iff P_{(\mathbf{y}, \mathbf{r})}(\Phi(\mathbf{y}) \vee \Pi(\mathbf{r})) \geq \delta_2$$

with $n \in \mathcal{O}\left(\left(d + \log_2\left(\frac{1-\delta_1}{1-\delta_2}\right)\right)^2 + \log_2\left(\frac{1}{\delta_2-\delta_1}\right)\right)$. It can be constructed in $\mathcal{O}(n)$ time.

The constructive proof of Lemma 3.8 can be found in Appendix A.

Lemma 3.9. *For $\delta \in \left[\frac{1}{2}, 1\right)$ we have $\text{IP1} \preceq_p \text{IP2}$, in particular in this case IP2 is NP^{PP} -hard.*

Proof. Let $\{\Phi, \mathbf{x}, k\}$ be an IP1 instance. We will construct $\{\Phi', \mathbf{x}', k'\}$ that is a *Yes*-instance for IP2 if and only if $\{\Phi, \mathbf{x}, k\}$ is a *Yes*-instance for IP1. Let $\Pi: \{0, 1\}^n \rightarrow \{0, 1\}$ be as in Lemma 3.8 for $\delta_1 = \frac{1}{4}$ and $\delta_2 = \delta$. We define

- $\Phi': \{0, 1\}^d \times \{0, 1\} \times \{0, 1\}^n \rightarrow \{0, 1\} : (\mathbf{y}, t, \mathbf{r}) \mapsto (\Phi(\mathbf{y}) \wedge t) \vee \Pi(\mathbf{r})$,
- $\mathbf{x}' = (\mathbf{x}, 1, \mathbf{1}_n) \in \{0, 1\}^d \times \{0, 1\} \times \{0, 1\}^n$,
- $k' = k$.

This is a polynomial time construction. By the choice of Φ' and \mathbf{x}' we guarantee $\Phi'(\mathbf{x}') = 1$ regardless of the value of $\Phi(\mathbf{x})$ since $\Pi(\mathbf{1}_n) = 1$.

Then, for all $S' = S \subseteq [k] = [k']$, we have

$$\begin{aligned} P_{\mathbf{y}}(\Phi(\mathbf{y}) \mid \mathbf{y}_S = \mathbf{x}_S) &> \frac{1}{2} \\ \iff P_{(\mathbf{y}, t)}(\Phi(\mathbf{y}) \wedge t \mid \mathbf{y}_S = \mathbf{x}_S) &> \frac{1}{4} \\ \iff P_{(\mathbf{y}, t, \mathbf{r})}((\Phi(\mathbf{y}) \wedge t) \vee \Pi(\mathbf{r}) \mid \mathbf{y}_S = \mathbf{x}_S) &\geq \delta \\ \iff P_{\mathbf{y}'}(\Phi'(\mathbf{y}') = \Phi'(\mathbf{x}') \mid \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) &\geq \delta. \end{aligned}$$

Thus $\{\Phi, \mathbf{x}, k\}$ is a *Yes*-instance for IP1 if and only if $\{\Phi', \mathbf{x}', k'\}$ is a *Yes*-instance for IP2. \square

We continue with the third and final step of the reductive chain. We translate the option to choose which of the first k variables to fix into the choice of fixing any set of k variables.

Lemma 3.10. *For $\delta \in \left[\frac{1}{2}, 1\right)$ we have $\text{IP2} \preceq_p \text{RELEVANT-INPUT}$, in particular in this case RELEVANT-INPUT is NP^{PP} -hard.*

Proof. Let $\{\Phi, \mathbf{x}, k\}$ be an IP2 instance. We will construct $\{\Phi', \mathbf{x}', k'\}$ that is a *Yes*-instance for RELEVANT-INPUT if and only if $\{\Phi, \mathbf{x}, k\}$ is a *Yes*-instance for IP2. For convenience, we split the d variables of Φ into the first k variables and the remaining $d - k$ variables and denote this $\Phi(\mathbf{x}) = \Phi(\mathbf{u}, \mathbf{r})$. The main idea is to extend Φ with clauses that force the set S to be chosen from the first k variables. More precisely, we define

- $\Phi' : \{0, 1\}^k \times \{0, 1\}^k \times \{0, 1\}^{d-k} \times \{0, 1\}^{d-k} \times \{0, 1\}^{d-k} \rightarrow \{0, 1\}$ with

$$\Phi'(\mathbf{u}, \mathbf{v}, \mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3) = (\Phi(\mathbf{u}, \mathbf{r}_1 \oplus \mathbf{r}_2 \oplus \mathbf{r}_3) \oplus (\neg\Phi(\mathbf{x}))) \wedge \left(\bigwedge_{i=1}^k ((u_i \oplus \neg x_i) \vee v_i) \right),$$

where $\mathbf{r}_1 \oplus \mathbf{r}_2 \oplus \mathbf{r}_3$ is understood component-wise,

- $\mathbf{x}' = (\mathbf{x}_{[k]}, \mathbf{1}_k, \mathbf{x}_{[k]^c}, \mathbf{x}_{[k]^c}, \mathbf{x}_{[k]^c}) \in \{0, 1\}^k \times \{0, 1\}^k \times \{0, 1\}^{d-k} \times \{0, 1\}^{d-k} \times \{0, 1\}^{d-k}$,
- $k' = k$.

This is a polynomial time construction. By the choice of Φ' and \mathbf{x}' we guarantee $\Phi'(\mathbf{x}') = 1$ regardless of the value of $\Phi(\mathbf{x})$.

Necessity: Assume that $\{\Phi, \mathbf{x}, k\}$ is a *Yes*-instance for IP2. Then there exists $S \subseteq [k]$ such that $P_{\mathbf{y}}(\Phi(\mathbf{y}) = \Phi(\mathbf{x}) \mid \mathbf{y}_S = \mathbf{x}_S) \geq \delta$. Now choose

$$S' = S \cup \{i \in \{k+1, \dots, 2k\} : i - k \notin S\}.$$

Then $|S'| = |S| + (k - |S|) = k = k'$ and for each $i \in [k]$ exactly one of the cases $i \in S'$ or $i + k \in S'$ can occur. The former corresponds to fixing $u_i = x'_i = x_i$ and the latter to fixing $v_i = 1$. Therefore

$$P_{(\mathbf{u}, \mathbf{v})} \left(\bigwedge_{i=1}^k ((u_i \oplus \neg x_i) \vee v_i) \mid (\mathbf{u}, \mathbf{v})_{S'} = \mathbf{x}'_{S'} \right) = 1,$$

which means, conditioned on $(\mathbf{u}, \mathbf{v})_{S'} = \mathbf{x}'_{S'}$, the probability of satisfying Φ' only depends on $\Phi(\mathbf{u}, \mathbf{r}_1 \oplus \mathbf{r}_2 \oplus \mathbf{r}_3)$. Now, since the random vector $\mathbf{r}_1 \oplus \mathbf{r}_2 \oplus \mathbf{r}_3$ is independent of this condition, it has the exact same distribution as the random vector \mathbf{r} and we obtain

$$\begin{aligned} P_{\mathbf{y}'}(\Phi'(\mathbf{y}') = \Phi'(\mathbf{x}') \mid \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) &= P_{\mathbf{y}'}(\Phi'(\mathbf{y}) \mid \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) \\ &= P(\Phi'(\mathbf{u}, \mathbf{v}, \mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3) \mid (\mathbf{u}, \mathbf{v})_{S'} = \mathbf{x}'_{S'}) \\ &= P(\Phi'(\mathbf{u}, \mathbf{v}, \mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3) \mid \mathbf{u}_S = \mathbf{x}_S, \mathbf{v}_{S^c} = \mathbf{1}) \\ &= P(\Phi(\mathbf{u}, \mathbf{r}) \oplus (\neg\Phi(\mathbf{x})) \mid \mathbf{u}_S = \mathbf{x}_S) \\ &= P(\Phi(\mathbf{u}, \mathbf{r}) = \Phi(\mathbf{x}) \mid \mathbf{u}_S = \mathbf{x}_S) \\ &= P_{\mathbf{y}}(\Phi(\mathbf{y}) = \Phi(\mathbf{x}) \mid \mathbf{y}_S = \mathbf{x}_S) \geq \delta. \end{aligned}$$

This shows that S' is a δ -relevant set for Φ' and \mathbf{x}' . Therefore $\{\Phi', \mathbf{x}', k'\}$ is a *Yes*-instance for RELEVANT-INPUT.

Sufficiency: Now, conversely, assume that $\{\Phi', \mathbf{x}', k'\}$ is a *Yes*-instance for RELEVANT-INPUT. Then there exists a δ -relevant set $S' \subseteq [2k + 3(d - k)]$ for Φ' and \mathbf{x}' with $|S'| \leq k' = k$. Hence, we have

$$P_{\mathbf{y}'}(\Phi'(\mathbf{y}') \mid \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) = P_{\mathbf{y}'}(\Phi'(\mathbf{y}') = \Phi'(\mathbf{x}') \mid \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) \geq \delta.$$

Without any conditions on \mathbf{u} and \mathbf{v} we have $P((u_i \oplus \neg x_i) \vee v_i) = \frac{3}{4}$ for all $i \in [k]$. Since $(\frac{3}{4})^3 < \frac{1}{2} \leq \delta$, we have $i \in S'$ or $i+k \in S'$ for all $i \in [k]$ except for at most two cases. Hence, $|S'| \leq k$ implies $|S' \cap [2k]^c| \leq 2$. Therefore, $\mathbf{r}_1 \oplus \mathbf{r}_2 \oplus \mathbf{r}_3$ conditioned on $\mathbf{y}'_{S'} = \mathbf{x}'_{S'}$ has the same distribution as $\mathbf{r}_1 \oplus \mathbf{r}_2 \oplus \mathbf{r}_3$ without the condition. So, without loss of generality, we can even assume $S' \cap [2k]^c = \emptyset$.

Similarly, if $i \in S'$ we have $P((u_i \oplus \neg x_i) \vee v_i | (\mathbf{u}, \mathbf{v})_{S'} = \mathbf{x}'_{S'}) = 1$ and additionally having $i+k \in S'$ could not increase the probability of satisfying Φ' . Contrary, if $i \notin S'$, we have

$$P((u_i \oplus \neg x_i) \vee v_i | (\mathbf{u}, \mathbf{v})_{S'} = \mathbf{x}'_{S'}) = \frac{1}{2} + \frac{1}{2}P(v_i | (\mathbf{u}, \mathbf{v})_{S'} = \mathbf{x}'_{S'}),$$

which is one if $i+k \in S'$ and $\frac{3}{4}$ otherwise. So including $i+k$ in S' does not decrease the probability. Altogether, without loss of generality we can assume $S' \subseteq [2k]$, $|S'| = k$ and for each $i \in [k]$ exactly one of the cases $i \in S'$ or $i+k \in S'$ occurs.

We now choose $S = S' \cap [k]$. Then the rest of the proof proceeds exactly as above and we conclude

$$\begin{aligned} P_{\mathbf{y}}(\Phi(\mathbf{y}) = \Phi(\mathbf{x}) | \mathbf{y}_S = \mathbf{x}_S) &= P(\Phi(\mathbf{u}, \mathbf{r}) \oplus (\neg \Phi(\mathbf{x})) | \mathbf{u}_S = \mathbf{x}_S) \\ &= P(\Phi'(\mathbf{u}, \mathbf{v}, \mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3) | \mathbf{u}_S = \mathbf{x}_S, \mathbf{v}_{S^c} = \mathbf{1}) \\ &= P(\Phi'(\mathbf{u}, \mathbf{v}, \mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3) | (\mathbf{u}, \mathbf{v})_{S'} = \mathbf{x}'_{S'}) \\ &= P_{\mathbf{y}'}(\Phi'(\mathbf{y}'_{S'}) | \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) \\ &= P_{\mathbf{y}'}(\Phi'(\mathbf{y}') = \Phi'(\mathbf{x}') | \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) \geq \delta, \end{aligned}$$

which shows that $\{\Phi, \mathbf{x}, k\}$ is a *Yes*-instance for IP2. □

Altogether, this section finishes the first part of the proof of Theorem 3.1.

3.2. Relevant-Input is contained in NP^{PP}

We now come to the second part of the proof of Theorem 3.1. We will show that RELEVANT-INPUT is in fact contained in NP^{PP} meaning that it can be solved in polynomial time by a non-deterministic Turing machine with access to a PP-oracle. The following Lemma will be useful.

Lemma 3.11. *Given $d \in \mathbb{N}$ and $0 < \delta_1 < \delta_2 < 1$, there exists a monotone function $\Pi: \{0, 1\}^n \rightarrow \{0, 1\}$ such that, for all $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$, we have*

$$P_{\mathbf{y}}(\Phi(\mathbf{y})) \geq \delta_2 \iff P_{(\mathbf{y}, \mathbf{r})}(\Phi(\mathbf{y}) \wedge \Pi(\mathbf{r})) > \delta_1$$

with $n \in \mathcal{O}\left(\left(d + \log_2\left(\frac{\delta_2}{\delta_1}\right)\right)^2\right)$. It can be constructed in $\mathcal{O}(n)$ time.

The constructive proof of Lemma 3.11 can be found in Appendix B.

Lemma 3.12. *For $\delta \in (0, 1)$ the RELEVANT-INPUT problem is contained in NP^{PP}.*

We will prove this for $\delta \in \left(\frac{1}{2}, 1\right)$ by lowering the probability threshold from δ to $\frac{1}{2}$. The case $\delta \in \left(0, \frac{1}{2}\right]$ can be treated analogously by raising the threshold.

Proof. Let $\{\Phi, \mathbf{x}, k\}$ be an instance of RELEVANT-INPUT. It suffices to show that the decision problem whether a given set $S \subseteq [d]$ is δ -relevant for Φ and \mathbf{x} is in PP. Without loss of generality we can assume $\Phi(\mathbf{x}) = 1$. Otherwise we could consider $\neg\Phi$ instead. Now, choose $\Pi: \{0, 1\}^n \rightarrow \{0, 1\}$ as in Lemma 3.11 for $\delta_1 = \frac{1}{2}$ and $\delta_2 = \delta$. Then

$$P_{\mathbf{y}}(\Phi(\mathbf{y}) \mid \mathbf{y}_S = \mathbf{x}_S) \geq \delta \iff P_{(\mathbf{y}, \mathbf{r})}(\Phi(\mathbf{y}) \wedge \Pi(\mathbf{r}) \mid \mathbf{y}_S = \mathbf{x}_S) > \frac{1}{2}.$$

A probabilistic Turing machine can now draw a random assignment (\mathbf{y}, \mathbf{r}) conditioned on $\mathbf{y}_S = \mathbf{x}_S$ and evaluate $\Phi(\mathbf{y}) \wedge \Pi(\mathbf{r})$. Thus, the machine will answer *Yes* with probability strictly greater than $\frac{1}{2}$ if and only if S is δ -relevant. This means the subproblem of checking a set for δ -relevance is contained in PP.

A non-deterministic Turing-machine with a PP-oracle can thus search over all possible sets $S \subseteq [d]$ with $|S| \leq k$ and check whether any of them is δ -relevant using the oracle. \square

4. Variations of the Problem Formulation

We want to consider two variations of the RELEVANT-INPUT problem. The first variation relaxes the requirement to check if a candidate set S is exactly δ -relevant or not by introducing a probability gap γ . In short, we then ask if a δ -relevant set of size k exists or if all sets of size k are not even $(\delta - \gamma)$ -relevant.

The second variation concerns the optimisation version of the problem. Here we introduce a set size gap and relax the requirement to find the smallest δ -relevant set. Instead for $k < m$ we ask if a δ -relevant set of size k exists or if all relevant sets must be of size at least m .

We show that these problems remain hard to solve (even in combination, that is with both a gap in probability and set size). This can be used to show that no polynomial time approximation algorithm for MIN-RELEVANT-INPUT with approximation factor better than the trivial factor d can exist unless $P = NP$. Due to the connection between Boolean functions and neural networks, as described in Section 1, this inapproximability result shows theoretical limitations of interpretation methods for neural network decision.

4.1. The Probability Gap

Probabilities and expectation values may be hard to calculate in theory, yet are often easy to approximate in practice. Checking whether a Boolean function is satisfied by more than a fraction δ of the possible assignments to its variables can often be done probabilistically. This only fails if the true probability is close to δ . These edge cases cause the hardness of the problem. Yet, in our scenario we do not necessarily care about their resolution. In fact, the choice of δ in the definition of δ -relevant sets is somewhat

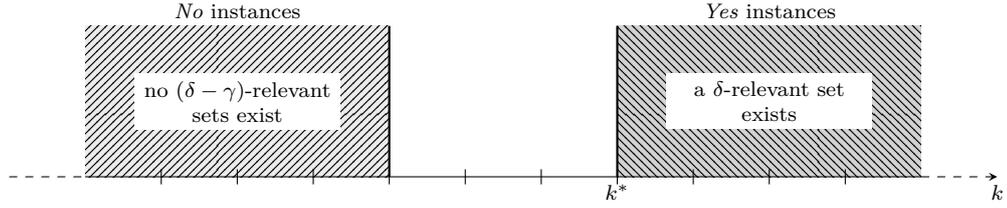


Figure 2: Visualisation of the GAPPED-RELEVANT-INPUT problem for some fixed Φ and \mathbf{x} and for various k . In the unmarked region in the centre no δ -relevant set exists but $\tilde{\delta}$ -relevant sets could exist for any $\tilde{\delta} < \delta$, in particular also for $\tilde{\delta} = \delta - \gamma$. In this region we do not expect an answer for the gapped problem. The solution k^* of the ungapped optimisation problem MIN-RELEVANT-INPUT is the left boundary of the *Yes*-instance region.

arbitrary. Thus, it seems impractical to defend the hardness of the RELEVANT-INPUT problem with the exact evaluation of probabilities.

Therefore, we introduce a variant of the problem including a probability gap. This can be seen as a promise problem with the promise that all sets S are either δ -relevant or not even $(\delta - \gamma)$ -relevant. Alternatively, this can be seen as the RELEVANT-INPUT problem where we want to answer *Yes* if a δ -relevant set of size k exists but only want to answer *No* if all sets of size k are not even $(\delta - \gamma)$ -relevant. For cases inbetween we do not expect an answer at all or do not care about the exact answer. This is illustrated in Figure 2.

Definition 4.1. For $\delta \in (0, 1]$ and $\gamma \in [0, \delta)$ we define the GAPPED-RELEVANT-INPUT problem as follows.

GIVEN: $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$, $\mathbf{x} \in \{0, 1\}^d$, and $k \in \mathbb{N}$, $1 \leq k \leq d$.

DECIDE:

Yes: There exists $S \subseteq [d]$ with $|S| \leq k$ and S is δ -relevant for Φ and \mathbf{x} .

No: All $S \subseteq [d]$ with $|S| \leq k$ are not $(\delta - \gamma)$ -relevant for Φ and \mathbf{x} .

For $\gamma = 0$ we exactly retrieve the original RELEVANT-INPUT problem but for $\gamma > 0$ this is an easier question.

Lemma 4.2. For $\delta \in (0, 1)$ and $\gamma \in (0, \delta)$ the GAPPED-RELEVANT-INPUT problem is contained in NP^{BPP} .

Proof. Let $\{\Phi, \mathbf{x}, k\}$ be an instance of GAPPED-RELEVANT-INPUT. It suffices to show that the decision problem whether a given set $S \subseteq [d]$ is either δ -relevant (*Yes*) or not $(\delta - \gamma)$ -relevant (*No*) for Φ and \mathbf{x} is in BPP. To see this, we describe an explicit algorithm with bounded error probability.

Draw $n = \lceil \frac{2 \ln(3)}{\gamma^2} \rceil$ independent random binary vectors $\mathbf{b}^{(i)} \in \{0, 1\}^{d-|S|}$ for $i \in [n]$ from the uniform distribution on $\{0, 1\}^{d-|S|}$ and define $\mathbf{y}^{(i)} \in \{0, 1\}^d$ as $\mathbf{y}_S^{(i)} = \mathbf{x}_S$ and

$\mathbf{y}_{S^c}^{(i)} = \mathbf{b}^{(i)}$. Set

$$\xi = \frac{1}{n} \sum_{i=1}^n \xi_i \quad \text{where} \quad \xi_i = \begin{cases} 1, & \text{if } \Phi(\mathbf{x}) = \Phi(\mathbf{y}^{(i)}) \\ 0, & \text{if } \Phi(\mathbf{x}) \neq \Phi(\mathbf{y}^{(i)}) \end{cases} \quad \text{for } i = 1, \dots, n.$$

Then answer *No* if $\xi < \delta - \frac{\gamma}{2}$ and answer *Yes* if $\xi \geq \delta - \frac{\gamma}{2}$.

The random variables ξ_i are independently and identically Bernoulli distributed variables with

$$p = \mathbb{E}[\xi_i] = \mathbb{E}[\xi] = P_{\mathbf{y}}(\Phi(\mathbf{y}_S) = \Phi(\mathbf{x}) \mid \mathbf{y}_S = \mathbf{x}_S).$$

Therefore S is δ -relevant if $p \geq \delta$ and not $(\delta - \gamma)$ -relevant if $p < \delta - \gamma$. We use Hoeffding's inequality [9] to bound the error probability of the algorithm. Firstly, assume $p \geq \delta$. Then we make an error if $\xi < \delta - \frac{\gamma}{2}$, which implies $p - \xi > \frac{\gamma}{2}$. The probability for this event can be bounded by

$$P\left(p - \xi > \frac{\gamma}{2}\right) \leq e^{-\frac{n\gamma^2}{2}} \leq \frac{1}{3}.$$

Secondly, assume $p < \delta - \gamma$. Then we can bound the probability that $\xi \geq \delta - \frac{\gamma}{2}$, and thus $\xi - p > \frac{\gamma}{2}$, by

$$P\left(\xi - p > \frac{\gamma}{2}\right) \leq e^{-\frac{n\gamma^2}{2}} \leq \frac{1}{3}.$$

Altogether the algorithm answers correctly with probability $\frac{2}{3}$ showing that the problem lies in BPP.

A non-deterministic Turing machine with BPP-oracle can thus search over all possible sets $S \subseteq [d]$ with $|S| \leq k$ and check if any of them is δ -relevant or all of them are not $(\delta - \gamma)$ -relevant. \square

Similar to the original problem formulation we can also state an optimisation version of the gapped problem. In this case, we relax the optimality condition on the set size k by allowing also sizes in the region between *Yes*- and *No*-instances of GAPPED-RELEVANT-INPUT (cf. Figure 2). In other words, we want to find any k that is large enough so that it is not a *No*-instance for the gapped problem but not larger than the optimal solution of the ungapped minimisation problem. Strictly speaking, this results in a search problem and not an optimisation problem. However, problems of this type can be referred to as weak optimisation problems [8].

Definition 4.3. For $\delta \in (0, 1]$ and $\gamma \in [0, \delta)$ we define the MIN-GAPPED-RELEVANT-INPUT problem as follows.

GIVEN: $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$ and $\mathbf{x} \in \{0, 1\}^d$.

FIND: $k \in \mathbb{N}$, $1 \leq k \leq d$ such that

- (i) There exists $S \subseteq [d]$ with $|S| = k$ and S is $(\delta - \gamma)$ -relevant for Φ and \mathbf{x} .
- (ii) All $S \subseteq [d]$ with $|S| < k$ are not δ -relevant for Φ and \mathbf{x} .

Note that both for GAPPED-RELEVANT-INPUT and MIN-GAPPED-RELEVANT-INPUT a solution for γ_1 will always also be a solution for $\gamma_2 > \gamma_1$. Specifically, being able to solve the ungapped problems introduced in Section 2 provides a solution to the gapped problems for any $\gamma > 0$.

4.2. The Set Size Gap (Approximability)

Even the gapped version of the minimisation problem is hard to approximate. We prove this by introducing another intermediate problem which we show to be NP-hard but which would be in P if there exists a “good” polynomial time approximation algorithm for MIN-GAPPED-RELEVANT-INPUT. As already mentioned before, strictly speaking MIN-GAPPED-RELEVANT-INPUT is not an optimisation but a search problem. In order to give a meaning to the concept of approximation factors we use the following convention.

Definition 4.4. An algorithm for MIN-GAPPED-RELEVANT-INPUT has an approximation factor $c \geq 1$ if, for any instance $\{\Phi, \mathbf{x}\}$, it produces an approximate solution k such that there exists a true solution \tilde{k} (satisfying both conditions in Definition 4.3) with $\tilde{k} \leq k \leq c\tilde{k}$.

An algorithm that always produces the trivial approximate solution $k = d$ achieves an approximation factor d . We will show that it is generally hard to obtain better factors. More precisely, for any $\alpha > 0$ an algorithm achieving an approximation factor $d^{1-\alpha}$ can not be in polynomial time unless $P = NP$.

Definition 4.5. For $\delta \in (0, 1]$ and $\gamma \in [0, \delta)$ we define the INTERMEDIATE PROBLEM 3 (IP3) as follows.

GIVEN: $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$, $\mathbf{x} \in \{0, 1\}^d$, and $k, m \in \mathbb{N}$, $1 \leq k \leq m \leq d$.

DECIDE:

Yes: There exists $S \subseteq [d]$ with $|S| \leq k$ and S is δ -relevant for Φ and \mathbf{x} .

No: All $S \subseteq [d]$ with $|S| \leq m$ are not $(\delta - \gamma)$ -relevant for Φ and \mathbf{x} .

The restriction to the case $k = m$ is exactly the GAPPED-RELEVANT-INPUT problem. However here we also allow the case $k < m$ with a gap in the set sizes. This is illustrated in Figure 3.

Lemma 4.6. For $\delta \in (0, 1)$ and $\gamma \in [0, \delta)$ we have $\text{SAT} \preceq_p \text{IP3}$, in particular, in this case IP3 is NP-hard.

Proof. Let $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$ be a SAT instance. We will construct $\{\Phi', \mathbf{x}', k', m'\}$ that is *Yes*-instance for IP3 if and only if Φ is a *Yes*-instance for SAT. Let

$$q = \left\lceil \log_2 \left(\frac{d}{1-\delta} \right) \right\rceil \quad \text{and} \quad p = \left\lceil \log_2 \left(\frac{1}{\delta-\gamma} \right) \right\rceil + 1.$$

We set

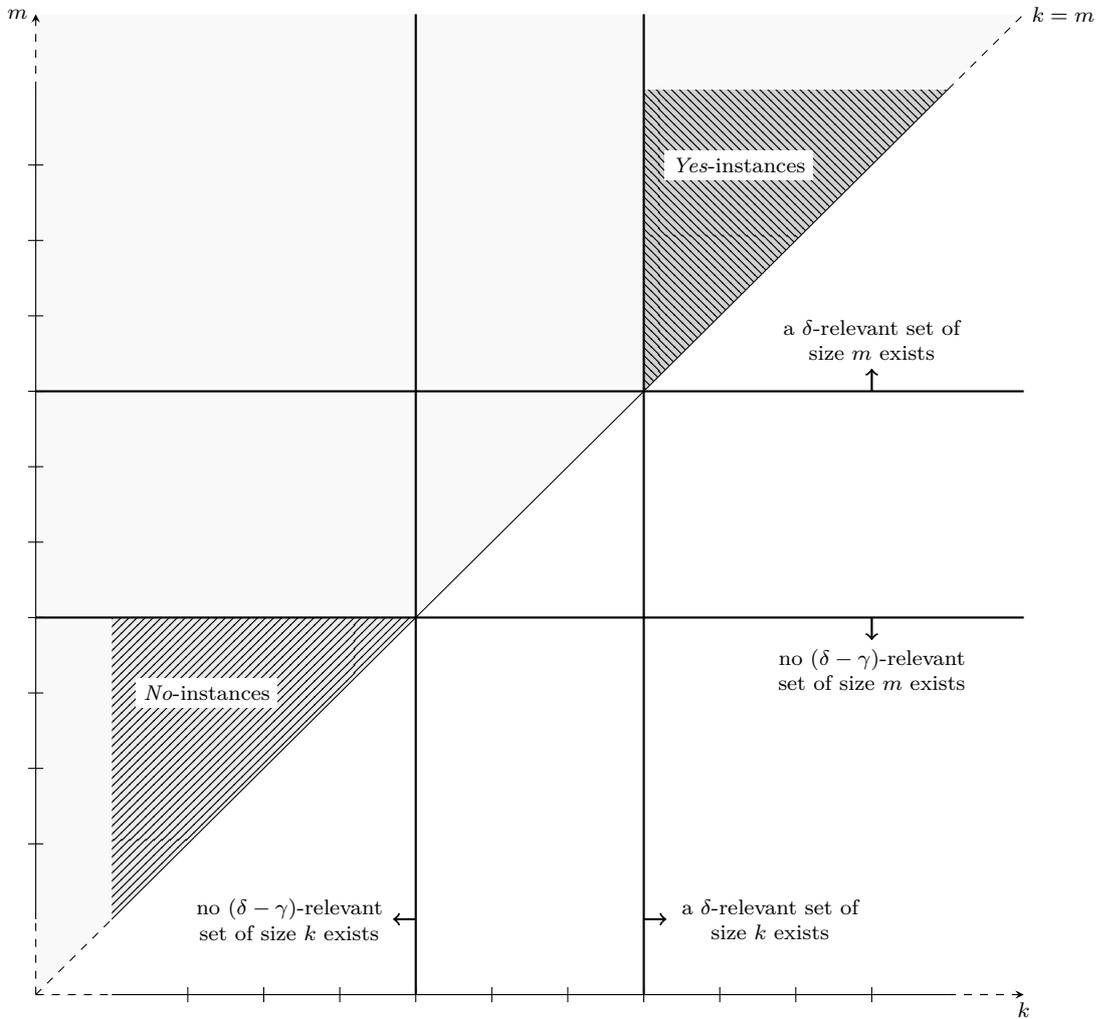


Figure 3: Visualisation of the INTERMEDIATE PROBLEM 3 for some fixed Φ and \mathbf{x} and for various k and m . As before we do not expect an answer for this problem in the unmarked regions. The restriction to the diagonal $k = m$ corresponds to the GAPPED-RELEVANT-INPUT problem (cf. Figure 2).

- $k' = dq$,
- $m' \geq k'$ arbitrary but at most polynomial in d ,
- $\Phi': \{0, 1\}^{d \times q} \times \{0, 1\}^{m'+p} \rightarrow \{0, 1\}$ with

$$\Phi'(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(q)}, \mathbf{v}) = \Phi \left(\bigwedge_{j=1}^q \mathbf{u}^{(j)} \right) \vee \left(\bigwedge_{i=1}^{m'+p} v_i \right),$$

where each $\mathbf{u}^{(j)} \in \{0, 1\}^d$ and the conjunction within Φ is understood component-wise,

- $\mathbf{x}' = \mathbf{1}_{dq+m'+p}$.

This is a polynomial time construction. By the choice of Φ' and \mathbf{x}' we guarantee $\Phi'(\mathbf{x}') = 1$ regardless of the satisfiability of Φ .

Necessity: Let Φ be a *Yes*-instance for SAT. This means that there exists $\mathbf{x} \in \{0, 1\}^d$ with $\Phi(\mathbf{x}) = 1$. Let $S = \{i \in [d] : x_i = 1\}$ and $S' = S \times [q]$. Then $|S'| \leq k'$. Denote

$$A(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(q)}) = \bigwedge_{(i,j) \in S'} u_i^{(j)}.$$

Hence, S' is δ -relevant for Φ' and \mathbf{x}' if $P(\Phi'(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(q)}, \mathbf{v}) \mid A(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(q)})) \geq \delta$. We have

$$\begin{aligned} P(\Phi'(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(q)}, \mathbf{v}) \mid A(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(q)})) &\geq P\left(\Phi\left(\bigwedge_{j=1}^q \mathbf{u}^{(j)}\right) \mid A(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(q)})\right) \\ &\geq P\left(\bigwedge_{j=1}^q \mathbf{u}^{(j)} = \mathbf{x} \mid A(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(q)})\right). \end{aligned}$$

From this, with a union bound, we obtain

$$\begin{aligned} P\left(\bigwedge_{j=1}^q \mathbf{u}^{(j)} = \mathbf{x} \mid A(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(q)})\right) &= 1 - P\left(\neg \bigwedge_{j=1}^q \mathbf{u}^{(j)} = \mathbf{x} \mid A(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(q)})\right) \\ &= 1 - P\left(\exists i \in S^c : \bigwedge_{j=1}^q u_i^{(j)}\right) \\ &\geq 1 - |S^c|2^{-q} \\ &\geq \delta, \end{aligned}$$

which shows that $\{\Phi', \mathbf{x}', k', m'\}$ is a *Yes*-instance for IP3.

Sufficiency: Now, conversely, let Φ be a *No*-instance for SAT. Then for any subset $S' \subseteq [dq + m' + p]$ with $|S'| \leq m'$ we have

$$\begin{aligned} P_{\mathbf{y}'}(\Phi'(\mathbf{y}') = \Phi'(\mathbf{x}') \mid \mathbf{y}'_{S'} = \mathbf{x}'_{S'}) &= P_{\mathbf{y}'}(\Phi'(\mathbf{y}') \mid \mathbf{y}'_{S'} = \mathbf{1}) \\ &= P_{(\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(q)}, \mathbf{v})} \left(\bigwedge_{i=1}^{m'+p} v_i \mid (\mathbf{u}^{(1)}, \dots, \mathbf{u}^{(q)}, \mathbf{v})_{S'} = \mathbf{1} \right) \\ &\leq 2^{-(m'+p-|S'|)} \\ &\leq 2^{-p} \\ &< \delta - \gamma. \end{aligned}$$

This shows that S' is not $(\delta - \gamma)$ -relevant for Φ' and \mathbf{x}' , hence $\{\Phi', \mathbf{x}', k', m'\}$ is a *No*-instance for IP3. \square

Finally, we come to the second main theorem of the paper which shows the inapproximability of the MIN-GAPPED-RELEVANT-INPUT problem.

Theorem 4.7. *Let $\delta \in (0, 1)$ and $\gamma \in [0, \delta)$. Then for any $\alpha \in (0, 1)$ there is no polynomial time approximation algorithm for MIN-GAPPED-RELEVANT-INPUT with an approximation factor of $d^{1-\alpha}$ unless $\mathsf{P} = \mathsf{NP}$.*

Proof. We prove this by showing that the existence of such an approximation algorithm would allow us to decide IP3 in polynomial time for certain instances. These can be chosen as in the proof of Lemma 4.6, which in turn implies that we could decide SAT in polynomial time. This is only possible if $\mathsf{P} = \mathsf{NP}$.

Let $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$ be a SAT instance and $\{\Phi', \mathbf{x}', k', m'\}$ an equivalent IP3 instance as in the proof of Lemma 4.6. We have seen that there is some freedom in the choice of m' as long as it satisfies $k' \leq m'$ and is at most polynomial in d . We choose $m' = \lceil \max(2k'(k'^{1-\alpha} + p^{1-\alpha}), (2k')^{\frac{1}{\alpha}} + 1) \rceil$ with $p = \lceil \log_2 \left(\frac{1}{\delta - \gamma} \right) \rceil + 1$ as before. Recall that $k' = dq$ with $q = \lceil \log_2 \left(\frac{d}{1-\delta} \right) \rceil$, so clearly m' is polynomial in d and $k' \leq m'$. Further, we have $m' > (2k')^{\frac{1}{\alpha}}$ so $1 - k'm'^{-\alpha} > \frac{1}{2}$ and therefore

$$m'(1 - k'm'^{-\alpha}) > \frac{m'}{2} \geq k'(k'^{1-\alpha} + p^{1-\alpha}).$$

Now let $d' = k' + m' + p$ denote the number of variables of Φ' . By the subadditivity of the map $z \mapsto z^{1-\alpha}$ we finally obtain

$$k'd'^{1-\alpha} = k'(k' + m' + p)^{1-\alpha} \leq k'(k'^{1-\alpha} + m'^{1-\alpha} + p^{1-\alpha}) < m'.$$

It remains to show that an IP3 instance with $m' > k'd'^{1-\alpha}$ can be decided by an approximation algorithm for MIN-GAPPED-RELEVANT-INPUT with approximation factor $d'^{1-\alpha}$. Assume such an algorithm exists and let k be an approximate solution. Then there exists a true solution \tilde{k} with $\tilde{k} \leq k \leq d'^{1-\alpha}\tilde{k}$.

Firstly, assume that $\{\Phi', \mathbf{x}', k', m'\}$ is a *Yes*-instance for IP3. Then there is a δ -relevant set of size k' . But no set smaller than \tilde{k} can be δ -relevant. This implies $\tilde{k} \leq k'$ and therefore $k \leq d^{1-\alpha} k' < m'$.

Secondly, assume that $\{\Phi', \mathbf{x}', k', m'\}$ is a *No*-instance for IP3. Then all sets of size at most m' are not $(\delta - \gamma)$ -relevant. But there exists a $(\delta - \gamma)$ -relevant set of size \tilde{k} . This implies $k \geq \tilde{k} > m'$.

Altogether, checking whether $k < m'$ or $k > m'$ decides $\{\Phi', \mathbf{x}', k', m'\}$. \square

5. Conclusion

We showed that the task of identifying the relevant components of an input assignment to the variables of a Boolean function is complete for the complexity class NP^{PP} and thus related in difficulty to the problem of planing under uncertainty. We showed furthermore that even relaxing the problem by introducing a probability gap promise that allows for efficient estimation of the fraction of assignments satisfying the Boolean function does not significantly simplify the problem. It remains hard to only approximately solve this gapped problem in any reasonable sense (beyond the trivial solution) unless $\text{P} = \text{NP}$.

We want to stress that this is a worst-case analysis and does not imply that the task is impossible in practical applications. Many non-linear optimisation problems are NP -hard in general and yet performed successfully on a regular basis. But it strongly suggests that anything beyond heuristic solution strategies requires further restrictions on the problem.

A. Raising the Probability

We give a constructive proof of Lemma 3.8. Let $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$ be arbitrary and $0 < \delta_1 < \delta_2 < 1$. We will construct a monotone function $\Pi: \{0, 1\}^n \rightarrow \{0, 1\}$ such that $\Pi(\mathbf{0}_n) = 0$, $\Pi(\mathbf{1}_n) = 1$, and

$$P_{\mathbf{y}}(\Phi(\mathbf{y})) > \delta_1 \iff P_{(\mathbf{y}, \mathbf{r})}(\Phi(\mathbf{y}) \vee \Pi(\mathbf{r})) \geq \delta_2, \quad (1)$$

with

$$n \in \mathcal{O} \left(\left(d + \log_2 \left(\frac{1 - \delta_1}{1 - \delta_2} \right) \right)^2 + \log_2 \left(\frac{1}{\delta_2 - \delta_1} \right) \right).$$

Denote $\Phi': \{0, 1\}^d \times \{0, 1\}^n \rightarrow \{0, 1\} : (\mathbf{y}, \mathbf{r}) \mapsto \Phi(\mathbf{y}) \vee \Pi(\mathbf{r})$, then

$$P(\Phi') = P(\Phi) + (1 - P(\Phi))P(\Pi), \quad (2)$$

which is monotonously increasing in both $P(\Phi)$ and $P(\Pi)$. Thus it suffices to consider the edge case when $P(\Phi)$ is close to δ_1 . Since $P(\Phi)$ can only take values in $\left\{ \frac{0}{2^d}, \frac{1}{2^d}, \dots, \frac{2^d}{2^d} \right\}$

we see that (1) is equivalent to the two conditions

$$P(\Phi) = \frac{\lfloor \delta_1 2^d \rfloor}{2^d} \implies P(\Phi') < \delta_2,$$

$$P(\Phi) = \frac{\lfloor \delta_1 2^d \rfloor + 1}{2^d} \implies P(\Phi') \geq \delta_2,$$

which together with (2) is equivalent to

$$\frac{\lfloor \delta_1 2^d \rfloor}{2^d} + \frac{2^d - \lfloor \delta_1 2^d \rfloor}{2^d} P(\Pi) < \delta_2 \tag{3}$$

$$\frac{\lfloor \delta_1 2^d \rfloor + 1}{2^d} + \frac{2^d - \lfloor \delta_1 2^d \rfloor - 1}{2^d} P(\Pi) \geq \delta_2. \tag{4}$$

In the case $\delta_1 < \delta_2 \leq \frac{\lfloor \delta_1 2^d \rfloor + 1}{2^d}$ the condition (4) is always fulfilled. Further, we have $\frac{\lfloor \delta_1 2^d \rfloor}{2^d} + \frac{2^d - \lfloor \delta_1 2^d \rfloor}{2^d} P(\Pi) < \delta_1 + P(\Pi)$ so that if $P(\Pi) < \delta_2 - \delta_1$ also (3) holds. Therefore, we choose

$$\Pi(\mathbf{r}) := \bigwedge_{i=1}^n r_i \quad \text{and} \quad n := \left\lceil \log_2 \left(\frac{1}{\delta_2 - \delta_1} \right) \right\rceil + 1.$$

If otherwise $\delta_2 > \frac{\lfloor \delta_1 2^d \rfloor + 1}{2^d}$ rearranging (3) and (4) finally gives us the bounds

$$a \leq P(\Pi) < b \tag{5}$$

on $P(\Pi)$ where

$$a = \frac{\delta_2 2^d - \lfloor \delta_1 2^d \rfloor - 1}{2^d - \lfloor \delta_1 2^d \rfloor - 1},$$

$$b = \frac{\delta_2 2^d - \lfloor \delta_1 2^d \rfloor}{2^d - \lfloor \delta_1 2^d \rfloor}.$$

It is not hard to check that indeed we have $0 \leq a < b \leq 1$. In Appendix C we show for $\eta \in [0, 1]$ and $\ell \in \mathbb{N}$ the existence of a monotone Boolean function $\Pi_{\eta, \ell}: \{0, 1\}^n \rightarrow \{0, 1\}$ such that $\Pi_{\eta, \ell}(\mathbf{0}_n) = 0$, $\Pi_{\eta, \ell}(\mathbf{1}_n) = 1$, and

$$|P(\Pi_{\eta, \ell}) - \eta| \leq 2^{-\ell}$$

with $n \leq \frac{\ell(\ell+3)}{2} \in \mathcal{O}(\ell^2)$. We conclude by choosing

$$\eta = \frac{b+a}{2},$$

$$\ell = \left\lceil -\log_2 \left(\frac{b-a}{2} \right) \right\rceil + 1 = \left\lceil -\log_2 \left(\frac{2^d(1-\delta_2)}{(2^d - \lfloor \delta_1 2^d \rfloor)(2^d - \lfloor \delta_1 2^d \rfloor - 1)} \right) \right\rceil + 1$$

and setting $\Pi = \Pi_{\eta, \ell}$. Then Π satisfies (5) and thus also (1) and we get

$$\ell \in \mathcal{O} \left(d + \log_2 \left(\frac{1 - \delta_1}{1 - \delta_2} \right) \right)$$

and therefore $n \in \mathcal{O}(\ell^2) = \mathcal{O} \left(\left(d + \log_2 \left(\frac{1 - \delta_1}{1 - \delta_2} \right) \right)^2 \right)$.

B. Lowering the Probability

We give a constructive proof of Lemma 3.11. Let $\Phi: \{0, 1\}^d \rightarrow \{0, 1\}$ be arbitrary and $0 < \delta_1 < \delta_2 < 1$. We will construct a monotone function $\Pi: \{0, 1\}^n \rightarrow \{0, 1\}$ such that

$$P_{\mathbf{y}}(\Phi(\mathbf{y})) \geq \delta_2 \iff P_{(\mathbf{y}, \mathbf{r})}(\Phi(\mathbf{y}) \wedge \Pi(\mathbf{r})) > \delta_1, \quad (6)$$

with

$$n \in \mathcal{O} \left(\left(d + \log_2 \left(\frac{\delta_2}{\delta_1} \right) \right)^2 \right).$$

Denote $\Phi': \{0, 1\}^d \times \{0, 1\}^n \rightarrow \{0, 1\} : (\mathbf{y}, \mathbf{r}) \mapsto \Phi(\mathbf{y}) \wedge \Pi(\mathbf{r})$, then

$$P(\Phi') = P(\Phi)P(\Pi), \quad (7)$$

which is monotonously increasing in both $P(\Phi)$ and $P(\Pi)$. Thus it suffices to consider the edge case when $P(\Phi)$ is close to δ_2 . Since $P(\Phi)$ can only take values in $\left\{ \frac{0}{2^d}, \frac{1}{2^d}, \dots, \frac{2^d}{2^d} \right\}$ we see that (6) is equivalent to the two conditions

$$P(\Phi) = \frac{\lceil \delta_2 2^d \rceil}{2^d} \implies P(\Phi') > \delta_1,$$

$$P(\Phi) = \frac{\lceil \delta_2 2^d \rceil - 1}{2^d} \implies P(\Phi') \leq \delta_1,$$

which together with (7) is equivalent to

$$\frac{\lceil \delta_2 2^d \rceil}{2^d} P(\Pi) > \delta_1 \quad (8)$$

$$\frac{\lceil \delta_2 2^d \rceil - 1}{2^d} P(\Pi) \leq \delta_1. \quad (9)$$

In the case $\frac{\lceil \delta_2 2^d \rceil - 1}{2^d} \leq \delta_1 < \delta_2$, notice that we can simply choose $\Pi \equiv 1$. Thus, from now on we assume that this is not the case. Then, rearranging (8) and (9) finally yields the bounds

$$a < P(\Pi) \leq b \quad (10)$$

on $P(\Pi)$, where

$$a = \frac{\delta_1 2^d}{\lceil \delta_2 2^d \rceil},$$

$$b = \frac{\delta_1 2^d}{\lceil \delta_2 2^d \rceil - 1}.$$

It is not hard to check that indeed we have $0 \leq a < b \leq 1$. In Appendix C, we show for $\eta \in [0, 1]$ and $\ell \in \mathbb{N}$ the existence of a monotone Boolean function $\Pi_{\eta, \ell}: \{0, 1\}^n \rightarrow \{0, 1\}$ such that $\Pi_{\eta, \ell}(\mathbf{0}_n) = 0$, $\Pi_{\eta, \ell}(\mathbf{1}_n) = 1$, and

$$|P(\Pi_{\eta, \ell}) - \eta| \leq 2^{-\ell}$$

with $n \leq \frac{\ell(\ell+3)}{2} \in \mathcal{O}(\ell^2)$. We conclude by choosing

$$\eta = \frac{b+a}{2},$$

$$\ell = \left\lceil -\log_2 \left(\frac{b-a}{2} \right) \right\rceil + 1 = \left\lceil -\log_2 \left(\frac{\delta_1 2^d}{2^{\lceil \delta_2 2^d \rceil} (\lceil \delta_2 2^d \rceil - 1)} \right) \right\rceil + 1$$

and setting $\Pi = \Pi_{\eta, \ell}$. Then Π satisfies (10) and thus also (6). Hence, we finally conclude that

$$\ell \in \mathcal{O} \left(d + \log_2 \left(\frac{\delta_2}{\delta_1} \right) \right)$$

and therefore $n \in \mathcal{O}(\ell^2) = \mathcal{O} \left(\left(d + \log_2 \left(\frac{\delta_2}{\delta_1} \right) \right)^2 \right)$.

C. Construction of the Functions $\Pi_{\eta, \ell}$

For $\eta \in [0, 1]$ (the target probability) and $\ell \in \mathbb{N}$ (the accuracy) we construct a Boolean function $\Pi_{\eta, \ell}: \{0, 1\}^n \rightarrow \{0, 1\}$ in disjunctive normal form with $n \in \mathcal{O}(\ell^2)$, $\Pi_{\eta, \ell}(\mathbf{0}_n) = 0$, $\Pi_{\eta, \ell}(\mathbf{1}_n) = 1$, and

$$|\eta - P(\Pi_{\eta, \ell})| \leq 2^{-\ell}.$$

If $\eta \leq 2^{-\ell}$ we can simply choose $\Pi_{\eta, \ell}(x_1, \dots, x_\ell) = \bigwedge_{k=1}^{\ell} x_k$. So from now on assume $2^{-\ell} < \eta \leq 1$. In this case we construct a sequence of functions $\Pi_i: \{0, 1\}^{n_i} \rightarrow \{0, 1\}$ such that $p_i = P(\Pi_i)$ is monotonely increasing and converges to η from below. We proceed according to the following iterative procedure: Start with the constant function $\Pi_0 \equiv 0$. Given Π_i and p_i we can stop and set $\Pi_{\eta, \ell} = \Pi_i$ if $|\eta - p_i| \leq 2^{-\ell}$. Otherwise we set $n_{i+1} = n_i + \Delta n_i$ with

$$\Delta n_i = \operatorname{argmin} \{ n \in \mathbb{N} : p_i + (1 - p_i)2^{-n} \leq \eta \}, \quad (11)$$

and

$$\Pi_{i+1}(x_1, \dots, x_{n_{i+1}}) = \Pi_i(x_1, \dots, x_{n_i}) \vee \left(\bigwedge_{k=n_i+1}^{n_{i+1}} x_k \right).$$

Clearly, we obtain $p_{i+1} = p_i + (1 - p_i)2^{-\Delta n_i}$. We will see below that Δn_i can not be too large and thus (11) can be efficiently computed by sequential search.

Lemma C.1. *The sequence $(p_i)_{i \in \mathbb{N}}$ is monotonely increasing and we have*

$$|\eta - p_{i+1}| \leq \frac{1}{2} |\eta - p_i|$$

for all $i \in \mathbb{N}$. In particular $|\eta - p_i| \leq 2^{-i}$ and $p_i \rightarrow \eta$ as $i \rightarrow \infty$.

Proof. Since $0 = p_0 \leq \eta$ and by choice of Δn_i , we have $p_i \leq \eta$ for all $i \in \mathbb{N}$. Also from (11) we know that $p_i + (1 - p_i)2^{-(\Delta n_i - 1)} > \eta$ since otherwise Δn_i would be chosen smaller. Therefore

$$\begin{aligned} \eta - p_{i+1} &= \eta - p_i - (1 - p_i)2^{-\Delta n_i} \\ &= \eta - \frac{1}{2}p_i - \frac{1}{2}\left(p_i + (1 - p_i)2^{-(\Delta n_i - 1)}\right) \\ &\leq \frac{1}{2}(\eta - p_i). \end{aligned}$$

The second part simply follows by repeatedly applying the above recursion i times and from the fact that $\eta - p_0 = \eta \leq 1$. \square

We conclude that the desired accuracy is reached after at most ℓ iterations in which case we stop and set $\Pi_{\eta, \ell} = \Pi_\ell$. It remains to determine how many variables need to be used in total. We first bound how many variables are added in each step.

Lemma C.2. *For any $i \in \mathbb{N}$, we have $\Delta n_i < -\log_2(\eta - p_i) + 1$.*

Proof. As before we know $p_i + (1 - p_i)2^{-(\Delta n_i - 1)} > \eta$ since otherwise Δn_i would be chosen smaller. This implies

$$2^{-(\Delta n_i - 1)} > \frac{\eta - p_i}{1 - p_i} \geq \eta - p_i,$$

and therefore $\Delta n_i < -\log_2(\eta - p_i) + 1$. \square

This can finally be used to bound how many variables are used in total.

Lemma C.3. *The total number of variables for $\Pi_{\eta, \ell} = \Pi_\ell$ is*

$$n = n_\ell = \sum_{i=1}^{\ell} \Delta n_{i-1} \in \mathcal{O}(\ell^2).$$

Proof. From Lemma C.1 we get $\eta - p_i \geq 2(\eta - p_{i+1})$ and thus $\eta - p_i \geq 2^{\ell-1-i}(\eta - p_{\ell-1})$. Without loss of generality we can assume $\eta - p_{\ell-1} \geq 2^{-\ell}$ since otherwise we can stop the iterative construction of $\Pi_{\eta, \ell}$ at $\ell - 1$. Using Lemma C.2, this immediately results in

$$\begin{aligned} n &= \sum_{i=1}^{\ell} \Delta n_{i-1} \\ &\leq \sum_{i=1}^{\ell} -\log_2(\eta - p_{i-1}) + 1 \\ &\leq \sum_{i=1}^{\ell} -\log_2\left(2^{\ell-i}(\eta - p_{\ell-1})\right) + 1 \\ &\leq \sum_{i=1}^{\ell} -\log_2\left(2^{-i}\right) + 1 \\ &= \frac{\ell(\ell+1)}{2} + \ell \in \mathcal{O}(\ell^2). \end{aligned} \quad \square$$

References

- [1] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLOS ONE*, 10(7):1–46, 07 2015.
- [2] R. Dechter. Bucket elimination: A unifying framework for probabilistic inference. In *Learning in graphical models*, pages 75–104. Springer, 1998.
- [3] X. Deng and C. H. Papadimitriou. On the complexity of cooperative solution concepts. *Mathematics of Operations Research*, 19(2):257–266, 1994.
- [4] M. Drummond and J. Bresina. *Anytime synthetic projection: Maximizing the probability of goal satisfaction*. NASA, Ames Research Center, Artificial Intelligence Research Branch, 1990.
- [5] D. Erhan, Y. Bengio, A. Courville, and P. Vincent. Visualizing higher-layer features of a deep network. Technical Report 1341, University of Montreal, 06 2009. Also presented at the ICML 2009 Workshop on Learning Feature Hierarchies, Montréal, Canada.
- [6] S. S. Fatima, M. Wooldridge, and N. R. Jennings. A linear approximation method for the shapley value. *Artificial Intelligence*, 172(14):1673–1699, 2008.
- [7] A. Graves, A.-R. Mohamed, and G. Hinton. Speech recognition with deep recurrent neural networks. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6645–6649, 05 2013.
- [8] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer, 1988.
- [9] W. Hoeffding. Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding*, pages 409–426. Springer, 1994.
- [10] E. Kang, J. Min, and J. C. Ye. A deep convolutional neural network using directional wavelets for low-dose x-ray ct reconstruction. *Medical Physics*, 44(10):e360–e375, 2017.
- [11] I. Kononenko et al. An efficient explanation of individual classifications using game theory. *Journal of Machine Learning Research*, 11(Jan):1–18, 2010.
- [12] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.
- [13] N. Kushmerick, S. Hanks, and D. S. Weld. An algorithm for probabilistic planning. *Artificial Intelligence*, 76(1-2):239–286, 1995.
- [14] M. L. Littman, J. Goldsmith, and M. Mundhenk. The computational complexity of probabilistic planning. *Journal of Artificial Intelligence Research*, 9:1–36, 1998.
- [15] M. L. Littman, S. M. Majercik, and T. Pitassi. Stochastic boolean satisfiability. *Journal of Automated Reasoning*, 27(3):251–296, 2001.
- [16] M. P. McBee, O. A. Awan, A. T. Colucci, C. W. Ghobadi, N. Kadom, A. P. Kansagra, S. Tridandapani, and W. F. Auffermann. Deep learning in radiology. *Academic Radiology*, 25(11):1472–1480, 2018.

- [17] J. D. Park. Map complexity results and approximation methods. In *Proceedings of the Eighteenth conference on Uncertainty in artificial intelligence*, pages 388–396. Morgan Kaufmann Publishers Inc., 2002.
- [18] L. S. Shapley. A value for n-person games. In H. W. Kuhn and A. W. Tucker, editors, *Contributions to the Theory of Games II*, pages 307–317. Princeton University Press, Princeton, 1953.
- [19] D. Shen, G. Wu, and H.-I. Suk. Deep learning in medical image analysis. *Annual Review of Biomedical Engineering*, 19(1):221–248, 2017.
- [20] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.
- [21] Y. Sun, X. Wang, and X. Tang. Deep convolutional network cascade for facial point detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3476–3483, 2013.
- [22] C. Szegedy, A. Toshev, and D. Erhan. Deep neural networks for object detection. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 2553–2561. Curran Associates, Inc., 2013.
- [23] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf. Deepface: Closing the gap to human-level performance in face verification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1701–1708, 2014.
- [24] A. Toshev and C. Szegedy. Deeppose: Human pose estimation via deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1653–1660, 2014.
- [25] J. Xie, L. Xu, and E. Chen. Image denoising and inpainting with deep neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 341–349. Curran Associates, Inc., 2012.
- [26] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 818–833, Cham, 2014. Springer International Publishing.