

Stability and Complexity of Minimising Probabilistic Automata

Stefan Kiefer and Björn Wachter

University of Oxford, UK

Abstract. We consider the state-minimisation problem for weighted and probabilistic automata. We provide a numerically stable polynomial-time minimisation algorithm for weighted automata, with guaranteed bounds on the numerical error when run with floating-point arithmetic. Our algorithm can also be used for “lossy” minimisation with bounded error. We show an application in image compression. In the second part of the paper we study the complexity of the minimisation problem for probabilistic automata. We prove that the problem is NP-hard and in PSPACE, improving a recent EXPTIME-result.

1 Introduction

Probabilistic and weighted automata were introduced in the 1960s, with many fundamental results established by Schützenberger [24] and Rabin [22]. Nowadays probabilistic automata are widely used in automated verification, natural-language processing, and machine learning.

Probabilistic automata (PAs) generalise deterministic finite automata (DFAs): The transition relation specifies, for each state q and each input letter a , a probability distribution on the successor state. Instead of a single initial state, a PA has a probability distribution over states; and instead of accepting states, a PA has an acceptance probability for each state. As a consequence, the language induced by a PA is a *probabilistic language*, i.e., a mapping $L : \Sigma^* \rightarrow [0, 1]$, which assigns each word an acceptance probability. Weighted automata (WAs), in turn, generalise PAs: the numbers appearing in the specification of a WA may be arbitrary real numbers. As a consequence, a WA induces a *weighted language*, i.e., a mapping $L : \Sigma^* \rightarrow \mathbb{R}$. Loosely speaking, the weight of a word w is the sum of the weights of all accepting w -labelled paths through the WA.

Given an automaton, it is natural to ask for a small automaton that accepts the same weighted language. A small automaton is particularly desirable when further algorithms are run on the automaton, and the runtime of those algorithms depends crucially on the size of the automaton [17]. In this paper we consider the problem of minimising the number of states of a given WA or PA, while preserving its (weighted or probabilistic) language.

WAs can be minimised in polynomial time, using, e.g., the standardisation procedure of [24]. When implemented efficiently (for instance using triangular matrices), one obtains an $O(|\Sigma|n^3)$ minimisation algorithm, where n is the number of states. As PAs are special WAs, the same holds in principle for PAs.

There are two problems with these algorithms: (1) numerical instability, i.e., round-off errors can lead to an automaton that is not minimal and/or induces a different probabilistic language; and (2) minimising a PA using WA minimisation algorithms does not necessarily result in a PA: transition weights may, e.g., become negative. This paper deals with those two issues.

Concerning problem (1), numerical stability is crucial under two scenarios: (a) when the automaton size makes the use of exact rational arithmetic prohibitive, and thus necessitates floating-point arithmetic [17]; or (b) when exact minimisation yields an automaton that is still too large and a “lossy compression” is called for, as in image compression [15]. Besides finding a numerically stable algorithm, we aim at two further goals: First, a stable algorithm should also be efficient; i.e., it should be as fast as classical (efficient, but possibly unstable) algorithms. Second, stability should be provable, and ideally there should be easily computable error bounds. In Section 3 we provide a numerically stable $O(|\Sigma|n^3)$ algorithm for minimising WAs. The algorithm generalises the *Arnoldi iteration* [2] which is used for locating eigenvalues in numerical linear algebra. The key ingredient, leading to numerical stability and allowing us to give error bounds, is the use of special orthonormal matrices, called *Householder reflectors* [14]. To the best of the authors’ knowledge, these techniques have not been previously utilised for computations on weighted automata.

Problem (2) suggests a study of the computational complexity of the *PA minimisation problem*: given a PA and $m \in \mathbb{N}$, is there an equivalent PA with m states? In the 1960s and 70s, PAs were studied extensively, see the survey [7] for references and Paz’s influential textbook [21]. PAs appear in various flavours and under different names. For instance, in *stochastic sequential machines* [21] there is no fixed initial state distribution, so the semantics of a stochastic sequential machine is not a probabilistic language, but a mapping from initial distributions to probabilistic languages. This gives rise to several notions of minimality in this model [21]. In this paper we consider only PAs with an initial state distribution; equivalence means equality of probabilistic languages.

One may be tempted to think that PA minimisation is trivially in NP, by guessing the minimal PA and verifying equivalence. However, it is not clear that the minimal PA has rational transition probabilities, even if this holds for the original PA.

For DFAs, which are special PAs, an automaton is minimal (i.e., has the least number of states) if and only if all states are reachable and no two states are equivalent. However, this equivalence does in general not hold for PAs. In fact, even if a PA has the property that no state behaves like a convex combination of other states, the PA may nevertheless not be minimal. As an example, consider the PA in the middle of Figure 2 on page 9. State 3 behaves like a convex combination of states 2 and 4: state 3 can be removed by splitting its incoming arc with weight 1 in two arcs with weight 1/2 each and redirecting the new arcs to states 2 and 4. The resulting PA is equivalent and no state can be replaced by a convex combination of other states. But the PA on the right of the figure is equivalent and has even fewer states.

In Section 4 we show that the PA minimisation problem is NP-hard by a reduction from 3SAT. A step in our reduction is to show that the following problem, the *hypercube problem*, is NP-hard: given a convex polytope P within the d -dimensional unit hypercube and $m \in \mathbb{N}$, is there a convex polytope with m vertices that is nested between P and the hypercube? We then reduce the hypercube problem to PA minimisation. To the best of the authors' knowledge, no lower complexity bound for PA minimisation has been previously obtained, and there was no reduction from the hypercube problem to PA minimisation. However, towards the converse direction, the textbook [21] suggests that an algorithm for the hypercube problem could serve as a “subroutine” for a PA minimisation algorithm, leaving the decidability of both problems open. In fact, problems similar to the hypercube problem were subsequently studied in the field of computational geometry, citing PA minimisation as a motivation [25,20,11,10].

The PA minimisation problem was shown to be decidable in [19], where the authors provided an exponential reduction to the existential theory of the reals, which, in turn, is decidable in PSPACE [8,23], but not known to be PSPACE-hard. In Section 4.2 we give a polynomial-time reduction from the PA minimisation problem to the existential theory of the reals. It follows that the PA minimisation problem is in PSPACE, improving the EXPTIME result of [19].

2 Preliminaries

In the technical development that follows it is more convenient to talk about vectors and transition matrices than about states, edges, alphabet labels and weights. However, a PA “of size n ” can be easily viewed as a PA with states $1, 2, \dots, n$. We use this equivalence in pictures.

Let $\mathbb{N} = \{0, 1, 2, \dots\}$. For $n \in \mathbb{N}$ we write \mathbb{N}_n for the set $\{1, 2, \dots, n\}$. For $m, n \in \mathbb{N}$, elements of \mathbb{R}^m and $\mathbb{R}^{m \times n}$ are viewed as vectors and matrices, respectively. Vectors are row vectors by default. Let $\alpha \in \mathbb{R}^m$ and $M \in \mathbb{R}^{m \times n}$. We denote the entries by $\alpha[i]$ and $M[i, j]$ for $i \in \mathbb{N}_m$ and $j \in \mathbb{N}_n$. By $M[i, \cdot]$ we refer to the i th row of M . By $\alpha[i..j]$ for $i \leq j$ we refer to the sub-vector $(\alpha[i], \alpha[i+1], \dots, \alpha[j])$, and similarly for matrices. We denote the transpose by α^T (a column vector) and $M^T \in \mathbb{R}^{n \times m}$. We write I_n for the $n \times n$ identity matrix. When the dimension is clear from the context, we write $e(i)$ for the vector with $e(i)[i] = 1$ and $e(i)[j] = 0$ for $j \neq i$. A vector $\alpha \in \mathbb{R}^m$ is *stochastic* if $\alpha[i] \geq 0$ for all $i \in \mathbb{N}_m$ and $\sum_{i=1}^m \alpha[i] \leq 1$. A matrix is *stochastic* if all its rows are stochastic. By $\|\cdot\| = \|\cdot\|_2$, we mean the 2-norm for vectors and matrices throughout the paper unless specified otherwise. If a matrix M is stochastic, then $\|M\| \leq \|M\|_1 \leq 1$. For a set $V \subseteq \mathbb{R}^n$, we write $\langle V \rangle$ to denote the vector space spanned by V , where we often omit the braces when denoting V . For instance, if $\alpha, \beta \in \mathbb{R}^n$, then $\langle \{\alpha, \beta\} \rangle = \langle \alpha, \beta \rangle = \{r\alpha + s\beta \mid r, s \in \mathbb{R}\}$.

An *\mathbb{R} -weighted automaton (WA)* $\mathcal{A} = (n, \Sigma, M, \alpha, \eta)$ consists of a size $n \in \mathbb{N}$, a finite alphabet Σ , a map $M : \Sigma \rightarrow \mathbb{R}^{n \times n}$, an initial (row) vector $\alpha \in \mathbb{R}^n$, and a final (column) vector $\eta \in \mathbb{R}^n$. Extend M to Σ^* by setting $M(a_1 \cdots a_k) := M(a_1) \cdots M(a_k)$. The *language* $L_{\mathcal{A}}$ of a WA \mathcal{A} is the mapping

$L_{\mathcal{A}} : \Sigma^* \rightarrow \mathbb{R}$ with $L_{\mathcal{A}}(w) = \alpha M(w)\eta$. WAs \mathcal{A}, \mathcal{B} over the same alphabet Σ are said to be *equivalent* if $L_{\mathcal{A}} = L_{\mathcal{B}}$. A WA \mathcal{A} is *minimal* if there is no equivalent WA \mathcal{B} of smaller size.

A *probabilistic automaton (PA)* $\mathcal{A} = (n, \Sigma, M, \alpha, \eta)$ is a WA, where α is stochastic, $M(a)$ is stochastic for all $a \in \Sigma$, and $\eta \in [0, 1]^n$. A PA is a *DFA* if all numbers in M, α, η are 0 or 1.

3 Stable WA Minimisation

In this section we discuss WA minimisation. In Section 3.1 we describe a WA minimisation algorithm in terms of elementary linear algebra. The presentation reminds of Brzozowski's algorithm for NFA minimisation [6].¹ WA minimisation techniques are well known, originating in [24], cf. also [4, Chapter II] and [3]. Our algorithm and its correctness proof may be of independent interest, as they appear to be particularly succinct. In Sections 3.2 and 3.3 we take further advantage of the linear algebra setting and develop a numerically stable WA minimisation algorithm.

3.1 Brzozowski-like WA Minimisation

Let $\mathcal{A} = (n, \Sigma, M, \alpha, \eta)$ be a WA. Define the *forward space* of \mathcal{A} as the (row) vector space $\mathbf{F} := \langle \alpha M(w) \mid w \in \Sigma^* \rangle$. Similarly, let the *backward space* of \mathcal{A} be the (column) vector space $\mathbf{B} := \langle M(w)\eta \mid w \in \Sigma^* \rangle$. Let $\vec{n} \in \mathbb{N}$ and $F \in \mathbb{R}^{\vec{n} \times n}$ such that the rows of F form a basis of \mathbf{F} . Similarly, let $\overleftarrow{n} \in \mathbb{N}$ and $B \in \mathbb{R}^{n \times \overleftarrow{n}}$ such that the columns of B form a basis of \mathbf{B} . Since $\mathbf{F}M(a) \subseteq \mathbf{F}$ and $M(a)\mathbf{B} \subseteq \mathbf{B}$ for all $a \in \Sigma$, there exist maps $\vec{M} : \Sigma \rightarrow \mathbb{R}^{\vec{n} \times \vec{n}}$ and $\overleftarrow{M} : \Sigma \rightarrow \mathbb{R}^{\overleftarrow{n} \times \overleftarrow{n}}$ such that

$$FM(a) = \vec{M}(a)F \quad \text{and} \quad M(a)B = B\overleftarrow{M}(a) \quad \text{for all } a \in \Sigma. \quad (1)$$

We call (F, \vec{M}) a *forward reduction* and (B, \overleftarrow{M}) a *backward reduction*. We will show that minimisation reduces to computing such reductions. By symmetry we can focus on forward reductions. We call a forward reduction (F, \vec{M}) *canonical* if $F[1, \cdot]$ (i.e., the first row of F) is a multiple of α , and the rows of F are orthonormal, i.e., $FF^T = I_{\vec{n}}$.

Let $\mathcal{A} = (n, \Sigma, M, \alpha, \eta)$ be a WA with forward and backward reductions (F, \vec{M}) and (B, \overleftarrow{M}) , respectively. Let $\vec{\alpha} \in \mathbb{R}^{\vec{n}}$ be a row vector such that $\alpha = \vec{\alpha}F$; let $\overleftarrow{\eta} \in \mathbb{R}^{\overleftarrow{n}}$ be a column vector such that $\eta = B\overleftarrow{\eta}$. (If (F, \vec{M}) is canonical, we have $\vec{\alpha} = (\pm\|\alpha\|, 0, \dots, 0)$.) Call $\vec{\mathcal{A}} := (\vec{n}, \Sigma, \vec{M}, \vec{\alpha}, F\eta)$ a *forward WA* of \mathcal{A} with base F and $\overleftarrow{\mathcal{A}} := (\overleftarrow{n}, \Sigma, \overleftarrow{M}, \alpha B, \overleftarrow{\eta})$ a *backward WA* of \mathcal{A} with base B . By extending (1) one can see that these automata are equivalent to \mathcal{A} :

¹ In [5] a very general Brzozowski-like minimization algorithm is presented in terms of universal algebra. One can show that it specialises to ours in the WA setting.

Proposition 1. *Let \mathcal{A} be a WA. Then $L_{\mathcal{A}} = L_{\overleftarrow{\mathcal{A}}} = L_{\overrightarrow{\mathcal{A}}}$.*

Further, applying both constructions consecutively yields a minimal WA:

Theorem 2. *Let \mathcal{A} be a WA. Let $\mathcal{A}' = \overleftarrow{\overrightarrow{\mathcal{A}}}$ or $\mathcal{A}' = \overrightarrow{\overleftarrow{\mathcal{A}}}$. Then \mathcal{A}' is minimal and equivalent to \mathcal{A} .*

Theorem 2 mirrors Brzozowski’s NFA minimisation algorithm. We give a short proof in Appendix A.2.

3.2 Numerically Stable WA Minimisation

Theorem 2 reduces the problem of minimising a WA to the problem of computing a forward and a backward reduction. In the following we focus on computing a *canonical* (see above for the definition) forward reduction (F, \overrightarrow{M}) . Figure 1 shows a generalisation of Arnoldi’s iteration [2] to multiple matrices. Arnoldi’s iteration is typically used for locating eigenvalues [12]. Its generalisation to multiple matrices is novel, to the best of the authors’s knowledge. Using (1) one can see that it computes a canonical forward reduction by iteratively extending a partial orthonormal basis $\{f_1, \dots, f_j\}$ for the forward space F .

```

function ArnoldiReduction
input:  $\alpha \in \mathbb{R}^n$ ;  $M : \Sigma \rightarrow \mathbb{R}^{n \times n}$ 
output: canonical forward reduction  $(F, \overrightarrow{M})$  with  $F \in \mathbb{R}^{\vec{n} \times n}$  and  $\overrightarrow{M} : \Sigma \rightarrow \mathbb{R}^{\vec{n} \times \vec{n}}$ 
   $\ell := 0$ ;  $j := 1$ ;  $f_1 := \alpha / \|\alpha\|$  (or  $f_1 := -\alpha / \|\alpha\|$ )
  while  $\ell < j$  do
     $\ell := \ell + 1$ 
    for  $a \in \Sigma$  do
      if  $f_\ell M(a) \notin \langle f_1, \dots, f_j \rangle$ 
         $j := j + 1$ 
        define  $f_j$  orthonormal to  $f_1, \dots, f_{j-1}$  such that
           $\langle f_1, \dots, f_{j-1}, f_\ell M(a) \rangle = \langle f_1, \dots, f_j \rangle$ 
        define  $\overrightarrow{M}(a)[\ell, \cdot]$  such that  $f_\ell M(a) = \sum_{i=1}^j \overrightarrow{M}(a)[\ell, i] f_i$ 
          and  $\overrightarrow{M}(a)[\ell, j+1..n] = (0, \dots, 0)$ 
     $\vec{n} := j$ ; form  $F \in \mathbb{R}^{\vec{n} \times \vec{n}}$  with rows  $f_1, \dots, f_{\vec{n}}$ 
  return  $F$  and  $\overrightarrow{M}(a)[1..\vec{n}, 1..\vec{n}]$  for all  $a \in \Sigma$ 

```

Fig. 1: Generalised Arnoldi iteration.

For efficiency, one would like to run generalised Arnoldi iteration (Figure 1) using floating-point arithmetic. This leads to *round-off errors*. The check “if $f_\ell M(a) \notin \langle f_1, \dots, f_j \rangle$ ” is particularly problematic: since the vectors f_1, \dots, f_j are computed with floating-point arithmetic, we cannot expect that $f_\ell M(a)$ lies *exactly* in the vector space spanned by those vectors, even if that would be the case without round-off errors. As a consequence, we need to introduce an *error tolerance parameter* $\tau > 0$, so that the check “ $f_\ell M(a) \notin \langle f_1, \dots, f_j \rangle$ ” returns *true* only if $f_\ell M(a)$ has a “distance” of more than τ to the vector space

$\langle f_1, \dots, f_j \rangle$.² Without such a “fuzzy” comparison the resulting automaton could even have more states than the original one. The error tolerance parameter τ causes further errors.

To assess the impact of those errors, we use the *standard model of floating-point arithmetic*, which assumes that the elementary operations $+, -, \cdot, /$ are computed exactly, up to a relative error of at most the *machine epsilon* $\varepsilon_{\text{mach}} \geq 0$. It is stated in [13, Chapter 2]: “This model is valid for most computers, and, in particular, holds for IEEE standard arithmetic.” The bit length of numbers arising in a numerical computation is bounded by hardware, using suitable roundoff. So we adopt the convention of numerical linear algebra to take the number of arithmetic operations as a measure of time complexity.

The algorithm `ArnoldiReduction` (Figure 1) leaves open how to implement the conditional “if $f_\ell M(a) \notin \langle f_1, \dots, f_j \rangle$ ”, and how to compute the new basis element f_j . In Appendix A.3 we propose an instantiation *HouseholderReduction* of `ArnoldiReduction` based on so-called *Householder reflectors* [14], which are special orthonormal matrices. We prove the following stability property:

Proposition 3. *Consider the algorithm `HouseholderReduction` in Appendix A.3, which has the following interface:*

function `HouseholderReduction`

input: $\alpha \in \mathbb{R}^n$; $M : \Sigma \rightarrow \mathbb{R}^{n \times n}$; error tolerance parameter $\tau \geq 0$

output: canonical forward reduction (F, \vec{M}) with $F \in \mathbb{R}^{\vec{n} \times n}$ and $\vec{M} : \Sigma \rightarrow \mathbb{R}^{\vec{n} \times \vec{n}}$

We have:

1. *The number of arithmetic operations is $O(|\Sigma|n^3)$.*
2. *HouseholderReduction instantiates ArnoldiReduction.*
3. *The computed matrices satisfy the following error bound: For each $a \in \Sigma$, the matrix $\mathcal{E}(a) \in \mathbb{R}^{\vec{n} \times n}$ with $\mathcal{E}(a) := FM(a) - \vec{M}(a)F$ satisfies*

$$\|\mathcal{E}(a)\| \leq 2\sqrt{n}\tau + cmn^3\varepsilon_{\text{mach}},$$

where $m > 0$ is such that $\|M(a)\| \leq m$ holds for all $a \in \Sigma$, and $c > 0$ is an input-independent constant.

The proof follows classical error-analysis techniques for QR factorisations with Householder reflectors [13, Chapter 19], but is substantially complicated by the presence of the “if” conditional and the resulting need for the τ parameter. By Proposition 3.2. `HouseholderReduction` computes a precise canonical forward reduction for $\varepsilon_{\text{mach}} = \tau = 0$. For positive $\varepsilon_{\text{mach}}$ and τ the error bound grows linearly in $\varepsilon_{\text{mach}}$ and τ , and with modest polynomials in the WA size n . In practice $\varepsilon_{\text{mach}}$ is very small³, so that the term $cmn^3\varepsilon_{\text{mach}}$ can virtually be ignored.

The use of Householder reflectors is crucial to obtain the bound of Proposition 3. Let us mention a few alternative techniques, which have been used for computing certain matrix factorisations. Such factorisations (QR or LU) are

² This will be made formal in our algorithm.

³ With IEEE double precision, e.g., it holds $\varepsilon_{\text{mach}} = 2^{-53}$ [13].

related to our algorithm. *Gaussian elimination* can also be used for WA minimisation in time $O(|\Sigma|n^3)$, but its stability is governed by the *growth factor*, which can be exponential even with *pivoting* [13, Chapter 9], so the bound on $\|\mathcal{E}(a)\|$ in Proposition 3 would include a term of the form $2^n \varepsilon_{\text{mach}}$. The most straightforward implementation of ArnoldiReduction would use the *Classical Gram-Schmidt* process, which is highly unstable [13, Chapter 19.8]. A variant, the *Modified Gram-Schmidt* process is stable, but the error analysis is complicated by a possibly loss of orthogonality of the computed matrix F . The extent of that loss depends on certain condition numbers (cf. [13, Equation (19.30)]), which are hard to estimate or control in our case. In contrast, our error bound is independent of condition numbers.

Using Theorem 2 we can prove:

Theorem 4. *Consider the following algorithm:*

function HouseholderMinimisation

input: WA $\mathcal{A} = (n, \Sigma, M, \alpha, \eta)$; error tolerance parameter $\tau \geq 0$

output: minimised WA $\mathcal{A}' = (n', \Sigma, M', \alpha', \eta')$.

 compute forward reduction (F, \vec{M}) of \mathcal{A} using HouseholderReduction

 form $\vec{\mathcal{A}} := (\vec{n}, \Sigma, \vec{M}, \vec{\alpha}, \vec{\eta})$ as the forward WA of \mathcal{A} with base F

 compute backward reduction (B, M') of $\vec{\mathcal{A}}$ using HouseholderReduction

 form $\mathcal{A}' := (n', \Sigma, M', \alpha', \eta')$ as the backward WA of $\vec{\mathcal{A}}$ with base B

 return \mathcal{A}'

We have:

1. The number of arithmetic operations is $O(|\Sigma|n^3)$.
2. For $\varepsilon_{\text{mach}} = \tau = 0$, the computed WA \mathcal{A}' is minimal and equivalent to \mathcal{A} .
3. Let $\tau > 0$. Let $m > 0$ such that $\|A\| \leq m$ holds for all $A \in \{M(a), \vec{M}(a), M'(a) \mid a \in \Sigma\}$. Then for all $w \in \Sigma^*$ we have

$$|L_{\mathcal{A}}(w) - L_{\mathcal{A}'}(w)| \leq 4|w|\|\alpha\|m^{|w|-1}\|\eta\|\sqrt{n}\tau \\ + c \max\{|w|, 1\}\|\alpha\|m^{|w|}\|\eta\|n^3\varepsilon_{\text{mach}},$$

where $c > 0$ is an input-independent constant.

The algorithm computes a backward reduction by running the straightforward backward variant of HouseholderReduction. We remark that for PAs one can take $m = 1$ for the norm bound m from part 3. of the theorem (or $m = 1 + \varepsilon$ for a small ε if unfortunate roundoff errors occur). It is hard to avoid an error bound exponential in the word length $|w|$, as $|L_{\mathcal{A}}(w)|$ itself may be exponential in $|w|$ (consider a WA of size 1 with $M(a) = 2$). Theorem 4 is proved in Appendix A.5.

The error bounds in Proposition 3 and Theorem 4 suggest to choose a small value for the error tolerance parameter τ . But as we have discussed, the computed WA may be non-minimal if τ is set too small or even to 0, intuitively because round-off errors may cause the algorithm to overlook minimisation opportunities. So it seems advisable to choose τ smaller (by a few orders of magnitude) than the desired bound on $\|\mathcal{E}(a)\|$, but larger (by a few orders of magnitude) than $\varepsilon_{\text{mach}}$.

Note that for $\varepsilon_{\text{mach}} > 0$ Theorem 4 does not provide a bound on the number of states of \mathcal{A}' .

To illustrate the stability issue we have experimented with minimising a PA \mathcal{A} derived from Herman’s protocol as in [17]. The PA has 190 states and $\Sigma = \{a\}$. When minimising with the (unstable) Classical Gram-Schmidt process, we have measured a huge error of $|L_{\mathcal{A}}(a^{190}) - L_{\mathcal{A}'}(a^{190})| \approx 10^{36}$. With the Modified Gram-Schmidt process and the method from Theorem 4 the corresponding errors were about 10^{-7} , which is in the same order as the error tolerance parameter τ .

3.3 Lossy WA Minimisation

A larger error tolerance parameter τ leads to more “aggressive” minimisation of a possibly already minimal WA. The price to pay is a shift in the language: one would expect only $L'_{\mathcal{A}}(w) \approx L_{\mathcal{A}}(w)$. Theorem 4 provides a bound on this imprecision. In this section we illustrate the trade-off between size and precision using an application in image compression.

Weighted automata can be used for image compression, as suggested by Culik et al. [15]. An image, represented as a two-dimensional matrix of grey-scale values, can be encoded as a weighted automaton where each pixel is addressed by a unique word. To obtain this automaton, the image is recursively subdivided into quadrants. There is a state for each quadrant and transitions from a quadrant to its sub-quadrants. At the level of the pixels, the automaton accepts with the correct grey-scale value.

Following this idea, we have implemented a prototype tool for image compression based on the algorithm of Theorem 4. We give details and show example pictures in Appendix A.6. This application illustrates lossy minimisation. The point is that Theorem 4 guarantees bounds on the loss.

4 The Complexity of PA Minimisation

Given a PA $\mathcal{A} = (n, \Sigma, M, \alpha, \eta)$ and $n' \in \mathbb{N}$, the *PA minimisation problem* asks whether there exists a PA $\mathcal{A}' = (n', \Sigma, M', \alpha', \eta')$ so that \mathcal{A} and \mathcal{A}' are equivalent. For the complexity results in this section we assume that the numbers in the description of the given PA are fractions of natural numbers represented in binary, so they are rational. In Section 4.1 we show that the minimisation problem is NP-hard. In Section 4.2 we show that the problem is in PSPACE by providing a polynomial-time reduction to the existential theory of the reals.

4.1 NP-Hardness

We will show:

Theorem 5. *The PA minimisation problem is NP-hard.*

For the proof we reduce from a geometrical problem, the *hypercube problem*, which we show to be NP-hard. Given $d \in \mathbb{N}$, a finite set $P = \{p_1, \dots, p_k\} \subseteq [0, 1]^d$ of vectors (“points”) within the d -dimensional unit hypercube, and $\ell \in \mathbb{N}$, the *hypercube problem* asks whether there is a set $Q = \{q_1, \dots, q_\ell\} \subseteq [0, 1]^d$ of at most ℓ points within the hypercube such that $\text{conv}(Q) \supseteq P$, where

$$\text{conv}(Q) := \{\lambda_1 q_1 + \dots + \lambda_\ell q_\ell \mid \lambda_1, \dots, \lambda_\ell \geq 0, \lambda_1 + \dots + \lambda_\ell = 1\}$$

denotes the convex hull of Q . Geometrically, the convex hull of P can be viewed as a convex polytope, nested inside the hypercube, which is another convex polytope. The hypercube problem asks whether a convex polytope with at most ℓ vertices can be nested in between those polytopes. The answer is trivially yes, if $\ell \geq k$ (take $Q = P$) or if $\ell \geq 2^d$ (take $Q = \{0, 1\}^d$). We speak of the *restricted hypercube problem* if P contains the origin $(0, \dots, 0)$. We prove the following:

Proposition 6. *The restricted hypercube problem can in polynomial time be reduced to the PA minimisation problem.*

Proof (sketch). Let $d \in \mathbb{N}$ and $P = \{p_1, \dots, p_k\} \subseteq [0, 1]^d$ and $\ell \in \mathbb{N}$ be an instance of the restricted hypercube problem, where $p_1 = (0, \dots, 0)$ and $\ell \geq 1$. We construct in polynomial time a PA $\mathcal{A} = (k+1, \Sigma, M, \alpha, \eta)$ such that there is a set $Q = \{q_1, \dots, q_\ell\} \subseteq [0, 1]^d$ with $\text{conv}(Q) \supseteq P$ if and only if there is a PA $\mathcal{A}' = (\ell+1, \Sigma, M', \alpha', \eta')$ equivalent to \mathcal{A} . Take $\Sigma := \{a_2, \dots, a_k\} \cup \{b_1, \dots, b_d\}$. Set $M(a_i)[1, i] := 1$ and $M(b_s)[i, k+1] := p_i[s]$ for all $i \in \{2, \dots, k\}$ and all $s \in \mathbb{N}_d$, and set all other entries of M to 0. Set $\alpha := e(1)$ and $\eta := e(k+1)^T$. Figure 2 shows an example of this reduction. We prove the correctness of this reduction in Appendix B.1. \square




Fig. 2: Reduction from the hypercube problem to the minimisation problem. The left figure shows an instance of the hypercube problem with $d = 2$ and $P = \{p_1, \dots, p_5\} = \{(0, 0), (0, \frac{3}{4}), (\frac{1}{4}, \frac{1}{2}), (\frac{1}{2}, \frac{1}{4}), (\frac{1}{2}, \frac{3}{4})\}$. It also suggests a set $Q = \{(0, 0), (0, 1), (1, \frac{1}{2})\}$ with $\text{conv}(Q) \supseteq P$. The middle figure depicts the PA \mathcal{A} obtained from P . The right figure depicts a minimal equivalent PA \mathcal{A}' , corresponding to the set Q suggested in the left figure.

Next we show that the hypercube problem is NP-hard, which together with Proposition 6 implies Theorem 5. A related problem is known⁴ to be NP-hard:

Theorem 7 (Theorem 4.2 of [10]). *Given two nested convex polyhedra in three dimensions, the problem of nesting a convex polyhedron with minimum faces between the two polyhedra is NP-hard.*

Note that this NP-hardness result holds even in $d = 3$ dimensions. However, the outer polyhedron is not required to be a cube, and the problem is about minimising the number of faces rather than the number of vertices. Using a completely different technique we show:

Proposition 8. *The hypercube problem is NP-hard. This holds even for the restricted hypercube problem.*

The proof is by a reduction from 3SAT, see Appendix B.2.

Remark 9. The hypercube problem is in PSPACE, by appealing to decision algorithms for $ExTh(\mathbb{R})$, the existential fragment of the first-order theory of the reals. For every fixed d the hypercube problem is⁵ in P , exploiting the fact that $ExTh(\mathbb{R})$ can be decided in polynomial time, if the number of variables is fixed. (For $d = 2$ an efficient algorithm is provided in [1].) It is an open question whether the hypercube problem is in NP. It is also open whether the search for a minimum Q can be restricted to sets of points with rational coordinates (this holds for $d = 2$).

Propositions 6 and 8 together imply Theorem 5.

4.2 Reduction to the Existential Theory of the Reals

In this section we reduce the PA minimisation problem to $ExTh(\mathbb{R})$, the existential fragment of the first-order theory of the reals. A formula of $ExTh(\mathbb{R})$ is of the form $\exists x_1 \dots \exists x_m R(x_1, \dots, x_m)$, where $R(x_1, \dots, x_m)$ is a boolean combination of comparisons of the form $p(x_1, \dots, x_m) \sim 0$, where $p(x_1, \dots, x_m)$ is a multivariate polynomial and $\sim \in \{<, >, \leq, \geq, =, \neq\}$. The validity of closed formulas ($m = n$) is decidable in PSPACE [8,23], and is not known to be PSPACE-hard.

Proposition 10. *Let $\mathcal{A}_1 = (n_1, \Sigma, M_1, \alpha_1, \eta_1)$ be a PA. A PA $\mathcal{A}_2 = (n_2, \Sigma, M_2, \alpha_2, \eta_2)$ is equivalent to \mathcal{A}_1 if and only if there exist matrices $\vec{M}(a) \in \mathbb{R}^{(n_1+n_2) \times (n_1+n_2)}$ for $a \in \Sigma$ and a matrix $F \in \mathbb{R}^{(n_1+n_2) \times (n_1+n_2)}$ such that $F[1, \cdot] = (\alpha_1, \alpha_2)$, and $F(\eta_1^T, -\eta_2^T)^T = (0, \dots, 0)^T$, and*

$$F \begin{pmatrix} M_1(a) & 0 \\ 0 & M_2(a) \end{pmatrix} = \vec{M}(a)F \quad \text{for all } a \in \Sigma.$$

The proof is in Appendix B.3. The conditions of Proposition 10 on \mathcal{A}_2 , including that it be a PA, can be phrased in $ExTh(\mathbb{R})$. Thus it follows:

⁴ The authors thank Joseph O'Rourke for pointing out [10].

⁵ This observation is in part due to Radu Grigore.

Theorem 11. *The PA minimisation problem can be reduced in polynomial time to $ExTh(\mathbb{R})$. Hence, PA minimisation is in PSPACE.*

Theorem 11 improves on a result in [19] where the minimisation problem was shown to be in EXPTIME. (More precisely, Theorem 4 of [19] states that a minimal PA can be computed in EXPSPACE, but the proof reveals that the decision problem can be solved in EXPTIME.)

5 Conclusions and Open Questions

We have developed a numerically stable and efficient algorithm for minimising WAs, based on linear algebra and Brzozowski-like automata minimisation. We have given bounds on the minimisation error in terms of both the machine epsilon and the error tolerance parameter τ .

We have shown NP-hardness for PA minimisation, and have given a polynomial-time reduction to $ExTh(\mathbb{R})$. Our work leaves open the precise complexity of the PA minimisation problem. The authors do not know whether the search for a minimal PA can be restricted to PAs with rational numbers. As stated in the Remark after Proposition 8, the corresponding question is open even for the hypercube problem. If rational numbers indeed suffice, then an NP algorithm might exist that guesses the (rational numbers of the) minimal PA and checks for equivalence with the given PA. Proving PSPACE-hardness would imply PSPACE-hardness of $ExTh(\mathbb{R})$, thus solving a longstanding open problem.

For comparison, the corresponding minimisation problems involving WAs (a generalisation of PAs) and DFAs (a special case of PAs) lie in P . More precisely, minimisation of WAs (with rational numbers) is in randomised NC [18], and DFA minimisation is NL-complete [9]. NFA minimisation is PSPACE-complete [16].

Acknowledgements. The authors would like to thank James Worrell, Radu Grigore, and Joseph O’Rourke for valuable discussions, and the anonymous referees for their helpful comments. Stefan Kiefer is supported by a Royal Society University Research Fellowship.

References

1. A. Aggarwal, H. Booth, J. O’Rourke, S. Suri, and C. K. Yap. Finding minimal convex nested polygons. *Information and Computation*, 83(1):98–110, 1989.
2. W.E. Arnoldi. The principle of minimized iteration in the solution of the matrix eigenvalue problem. *Quarterly of Applied Mathematics*, 9:17–29, 1951.
3. A. Beimel, F. Bergadano, N.H. Bshouty, E. Kushilevitz, and S. Varricchio. Learning functions represented as multiplicity automata. *Journal of the ACM*, 47(3):506–530, 2000.
4. J. Berstel and C. Reutenauer. *Rational Series and Their Languages*. Springer, 1988.
5. F. Bonchi, M.M. Bonsangue, H.H. Hansen, P. Panangaden, J.J.M.M. Rutten, and A. Silva. Algebra-coalgebra duality in Brzozowski’s minimization algorithm. *ACM Transactions on Computational Logic*, to appear.

6. J.A. Brzozowski. Canonical regular expressions and minimal state graphs for definite events. In *Symposium on Mathematical Theory of Automata*, volume 12 of *MRI Symposia Series*, pages 529–561. Polytechnic Press, Polytechnic Institute of Brooklyn, 1962.
7. R.G. Bukharaev. Probabilistic automata. *Journal of Soviet Mathematics*, 13(3):359–386, 1980.
8. J. Canny. Some algebraic and geometric computations in PSPACE. In *Proceedings of STOC’88*, pages 460–467, 1988.
9. S. Cho and D.T. Huynh. The parallel complexity of finite-state automata problems. *Information and Computation*, 97(1):1–22, 1992.
10. G. Das and M.T. Goodrich. On the complexity of approximating and illuminating three-dimensional convex polyhedra. In *Proceedings of Workshop on Algorithms and Data Structures*, volume 955 of *LNCS*, pages 74–85. Springer, 1995.
11. G. Das and D. Joseph. Minimum vertex hulls for polyhedral domains. *Theoretical Computer Science*, 103(1):107–135, 1992.
12. G.H. Golub and C.F. van Loan. *Matrix Computations*. John Hopkins University Press, 1989.
13. N.J. Higham. *Accuracy and Stability of Numerical Algorithms*. SIAM, second edition, 2002.
14. A.S. Householder. Unitary triangularization of a nonsymmetric matrix. *Journal of the ACM*, 5(4):339–342, 1958.
15. K. Culik II and J. Kari. Image compression using weighted finite automata. *Computers & Graphics*, 17(3):305–313, 1993.
16. T. Jiang and B. Ravikumar. Minimal NFA problems are hard. *SIAM Journal on Computing*, 22(6):1117–1141, 1993.
17. S. Kiefer, A.S. Murawski, J. Ouaknine, B. Wachter, and J. Worrell. Language equivalence for probabilistic automata. In *Proceedings of CAV*, volume 6806 of *LNCS*, pages 526–540, 2011.
18. S. Kiefer, A.S. Murawski, J. Ouaknine, B. Wachter, and J. Worrell. On the complexity of equivalence and minimisation for Q-weighted automata. *Logical Methods in Computer Science*, 9(1:8):1–22, 2013.
19. P. Mateus, D. Qiu, and L. Li. On the complexity of minimizing probabilistic and quantum automata. *Information and Computation*, 218:36–53, 2012.
20. J.S.B. Mitchell and S. Suri. Separation and approximation of polyhedral objects. In *Proceedings of SODA*, pages 296–306, 1992.
21. A. Paz. *Introduction to probabilistic automata*. Academic Press, 1971.
22. M.O. Rabin. Probabilistic automata. *Information and Control*, 6 (3):230–245, 1963.
23. J. Renegar. On the computational complexity and geometry of the first-order theory of the reals. Parts I–III. *Journal of Symbolic Computation*, 13(3):255–352, 1992.
24. M.-P. Schützenberger. On the definition of a family of automata. *Information and Control*, 4:245–270, 1961.
25. C.B. Silio. An efficient simplex coverability algorithm in E^2 with application to stochastic sequential machines. *IEEE Transactions on Computers*, C-28(2):109–120, 1979.
26. W. Tzeng. A polynomial-time algorithm for the equivalence of probabilistic automata. *SIAM Journal on Computing*, 21(2):216–227, 1992.

A Proofs of Section 3

A.1 Proof of Proposition 1

Proposition 1. *Let \mathcal{A} be a WA. Then $L_{\mathcal{A}} = L_{\overleftarrow{\mathcal{A}}} = L_{\overrightarrow{\mathcal{A}}}$.*

Proof. Observe that the equalities (1) extend inductively to words:

$$FM(w) = \overrightarrow{M}(w)F \quad \text{and} \quad M(w)B = B\overleftarrow{M}(w) \quad \text{for all } w \in \Sigma^*. \quad (2)$$

Using (2) and the definition of $\overrightarrow{\alpha}$ we have for all $w \in \Sigma^*$:

$$L_{\mathcal{A}}(w) = \alpha M(w)\eta = \overrightarrow{\alpha} FM(w)\eta = \overrightarrow{\alpha} \overrightarrow{M}(w)F\eta = L_{\overrightarrow{\mathcal{A}}}(w).$$

Symmetrically one can show $L_{\mathcal{A}} = L_{\overleftarrow{\mathcal{A}}}$. \square

A.2 Proof of Theorem 2

We will use the notion of a *Hankel matrix* [4,3]:

Definition 12. *Let $L : \Sigma^* \rightarrow \mathbb{R}$. The Hankel matrix of L is the matrix $H^L \in \mathbb{R}^{\Sigma^* \times \Sigma^*}$ with $H^L[x, y] = L(xy)$ for all $x, y \in \Sigma^*$. We define $\text{rank}(L) := \text{rank}(H^L)$.*

We have the following proposition:

Proposition 13. *Let \mathcal{A} be an automaton of size n . Then $\text{rank}(L_{\mathcal{A}}) \leq n$.*

Proof. Consider the matrices $\widehat{F} : \mathbb{R}^{\Sigma^* \times n}$ and $\widehat{B} : \mathbb{R}^{n \times \Sigma^*}$ with $\widehat{F}[w, \cdot] := \alpha M(w)$ and $\widehat{B}[\cdot, w] := M(w)\eta$ for all $w \in \Sigma^*$. Note that $\text{rank}(\widehat{F}) \leq n$ and $\text{rank}(\widehat{B}) \leq n$. Let $x, y \in \Sigma^*$. Then $(\widehat{F}\widehat{B})[x, y] = \alpha M(x)M(y)\eta = L_{\mathcal{A}}(xy)$, so $\widehat{F}\widehat{B}$ is the Hankel matrix of $L_{\mathcal{A}}$. Hence $\text{rank}(L_{\mathcal{A}}) = \text{rank}(\widehat{F}\widehat{B}) \leq \min\{\text{rank}(\widehat{F}), \text{rank}(\widehat{B})\} \leq n$. \square

Now we can prove the theorem:

Theorem 2. *Let \mathcal{A} be a WA. Let $\mathcal{A}' = \overleftarrow{\overleftarrow{\mathcal{A}}}$ or $\mathcal{A}' = \overrightarrow{\overrightarrow{\mathcal{A}}}$. Then \mathcal{A}' is minimal and equivalent to \mathcal{A} .*

Proof. W.l.o.g. we assume $\mathcal{A}' = \overrightarrow{\overrightarrow{\mathcal{A}}}$. Let $\mathcal{A} = (n, \Sigma, M, \alpha, \eta)$. Let $\overleftarrow{\mathcal{A}} = (\overleftarrow{n}, \Sigma, \overleftarrow{M}, \alpha B, \overleftarrow{\eta})$ be a backward automaton of \mathcal{A} with base B . Let $\overrightarrow{\overrightarrow{\mathcal{A}}}$ be a forward automaton of $\overleftarrow{\mathcal{A}}$ with base \widetilde{F} . Equivalence of \mathcal{A} , $\overleftarrow{\mathcal{A}}$ and $\overrightarrow{\overrightarrow{\mathcal{A}}}$ follows from Proposition 1. Assume that $\overrightarrow{\overrightarrow{\mathcal{A}}}$ has $\overrightarrow{\overrightarrow{n}}$ states. For minimality, by Proposition 13,

it suffices to show $\overrightarrow{\text{rank}}_n = \text{rank}(H)$, where H is the Hankel matrix of $L_{\mathcal{A}}$. Let \hat{F} and \hat{B} be the matrices from the proof of Proposition 13. We have:

$$\begin{aligned}
\overrightarrow{\text{rank}}_n &= \text{rank}(\tilde{F}) && \text{(definition of } \overrightarrow{\text{rank}}_n) \\
&= \dim \langle \alpha B \overleftarrow{M}(w) \mid w \in \Sigma^* \rangle && \text{(definition of } \tilde{F}) \\
&= \dim \langle \alpha M(w) B \mid w \in \Sigma^* \rangle && \text{(by (2))} \\
&= \text{rank}(\hat{F}B) && \text{(definition of } \hat{F}) \\
&= \dim \langle \hat{F}M(w)\eta \mid w \in \Sigma^* \rangle && \text{(definition of } B) \\
&= \text{rank}(\hat{F}\hat{B}) && \text{(definition of } \hat{B}) \\
&= \text{rank}(H) && \text{(proof of Proposition 13).}
\end{aligned}$$

□

A.3 Instantiation of ArnoldiReduction with Householder Reflectors

Fix $n \geq 1$. For a row vector $x \in \mathbb{R}^k$ with $k \in \mathbb{N}_n$ and $\|x\| \neq 0$, the *Householder reflector* P for x is defined as the matrix

$$P = \begin{pmatrix} I_{n-k} & 0 \\ 0 & R \end{pmatrix} \in \mathbb{R}^{n \times n},$$

where $R = I_k - 2v^T v \in \mathbb{R}^{k \times k}$, and

$$v = \frac{(x[1] + \text{sign}(x[1])\|x\|, x[2], \dots, x[k])}{\|(x[1] + \text{sign}(x[1])\|x\|, x[2], \dots, x[k])\|} \in \mathbb{R}^k,$$

where $\text{sign}(r) = +1$ if $r \geq 0$ and -1 otherwise. (The careful choice of the sign here is to ensure numerical stability.) To understand this definition better, first observe that v is a row vector with $\|v\| = 1$. It is easy to verify that R and thus P are orthonormal and symmetric. Moreover, we have $RR = I_k$ and thus $PP = I_n$, i.e., $P = P^T = P^{-1}$. Geometrically, R describes a reflection about the hyperplane through the origin and orthogonal to $v \in \mathbb{R}^k$ [13, Chapter 19.1]. Crucially, the vector v is designed so that R reflects x onto the first axis, i.e., $xR = (\pm\|x\|, 0, \dots, 0)$.

Figure 3 shows an instantiation of the algorithm from Figure 1 using Householder reflectors. Below we prove correctness by showing that HouseholderReduction indeed refines ArnoldiReduction, assuming $\tau = 0$ for the error tolerance parameter. For efficiency it is important not to form the reflectors P_1, P_2, \dots explicitly. If P is the Householder reflector for $x \in \mathbb{R}^k$, it suffices to keep the vector $v \in \mathbb{R}^k$ from the definition of Householder reflectors. A multiplication $y := yP$ (for a row vector $y \in \mathbb{R}^n$) can then be implemented in $O(n)$ with $y[(n-k+1)..n] := y[(n-k+1)..n] - 2(y[(n-k+1)..n] \cdot v^T)v$. This gives an $O(|\Sigma|n^3)$ number of arithmetic operations of HouseholderReduction.

```

function HouseholderReduction
input:  $\alpha \in \mathbb{R}^n$ ;  $M : \Sigma \rightarrow \mathbb{R}^{n \times n}$ ; error tolerance parameter  $\tau \geq 0$ 
output: canonical forward reduction  $(F, \vec{M})$  with  $F \in \mathbb{R}^{\vec{n} \times n}$  and  $\vec{M} : \Sigma \rightarrow \mathbb{R}^{\vec{n} \times \vec{n}}$ 
   $P_1 :=$  Householder reflector for  $\alpha/\|\alpha\|$ 
   $\ell := 0$ ;  $j := 1$ ;  $f_1 := e(1)P_1$ 
  while  $\ell < j$  do
     $\ell := \ell + 1$ 
    for  $a \in \Sigma$  do
       $\vec{M}(a)[\ell, \cdot] := f_\ell M(a)P_1 \cdots P_j$ 
      if  $j + 1 \leq n$  and  $\|\vec{M}(a)[\ell, j+1..n]\| > \tau$ 
         $j := j + 1$ 
         $P_j :=$  Householder reflector for  $\vec{M}(a)[\ell, j..n]$ 
         $\vec{M}(a)[\ell, \cdot] := \vec{M}(a)[\ell, \cdot]P_j$ 
         $f_j := e(j)P_j \cdots P_1$ 
       $\vec{n} := j$ ; form  $F \in \mathbb{R}^{\vec{n} \times n}$  with rows  $f_1, \dots, f_{\vec{n}}$ 
    return  $F$  and  $\vec{M}(a)[1..\vec{n}, 1..\vec{n}]$  for all  $a \in \Sigma$ 

```

Fig. 3: Instantiation of ArnoldiReduction (Fig. 1) using Householder reflectors.

HouseholderReduction **Instantiates** **ArnoldiReduction.** Let $P_1, \dots, P_{\vec{n}} \in \mathbb{R}^{n \times n}$ be the reflectors computed in HouseholderReduction. Recall that we have $P_j^T = P_j^{-1} = P_j$. For $0 \leq j \leq \vec{n}$ define

$$F_j := P_j P_{j-1} \cdots P_1.$$

The matrices $F_j \in \mathbb{R}^{n \times n}$ are orthonormal. Since for all $j < \vec{n}$ the reflector P_{j+1} leaves the first j rows of F_j unchanged, the first j rows of $F_j, \dots, F_{\vec{n}}$ coincide with the vectors f_1, \dots, f_j computed in HouseholderReduction.

First we consider the initialisation part before the while loop. Since P_1 is the Householder reflector for $\alpha/\|\alpha\|$, we have $(\alpha/\|\alpha\|)P_1 = (\pm 1, 0, \dots, 0)$. It follows that we have $f_1 = e(1)P_1 = \pm \alpha/\|\alpha\|$, as in ArnoldiReduction.

Now we consider the while loop. Consider an arbitrary iteration of the “for $a \in \Sigma$ do” loop. Directly after the loop head we have $\vec{M}(a)[\ell, \cdot] = f_\ell M(a)P_1 \cdots P_j$, hence $f_\ell M(a) = \vec{M}(a)[\ell, \cdot]F_j$. As F_j is orthonormal and its first j rows coincide with f_1, \dots, f_j , we have $f_\ell M(a) \in \langle f_1, \dots, f_j \rangle$ if and only if $\vec{M}(a)[\ell, j+1..n] = (0, \dots, 0)$. This corresponds to the “if” conditional in ArnoldiReduction. It remains to be shown that $\vec{M}(a)[\ell, \cdot]$ is defined as in ArnoldiReduction:

- Let $\vec{M}(a)[\ell, j+1..n] = (0, \dots, 0)$. Then at the end of the loop we have $f_\ell M(a) = \sum_{i=1}^j \vec{M}(a)[\ell, i]f_i$, as required in ArnoldiReduction.
- Let $\vec{M}(a)[\ell, j+1..n] \neq (0, \dots, 0)$. Then we have:

$$\begin{aligned}
f_\ell M(a) &= \vec{M}(a)[\ell, \cdot]F_j && \text{(as argued above)} \\
&= \vec{M}(a)[\ell, \cdot]P_{j+1}P_{j+1}F_j && \text{(as } P_{j+1}P_{j+1} = I_n) \\
&= \vec{M}(a)[\ell, \cdot]P_{j+1}F_{j+1} && \text{(by definition of } F_{j+1})
\end{aligned}$$

Further, the reflector P_{j+1} is designed such that $\vec{M}(a)[\ell, j+1..n] = (r, 0, \dots, 0)$ for some $r \neq 0$. So after increasing j , and updating $\vec{M}(a)[\ell, \cdot]$, and defining the new vector f_j , at the end of the loop we have $f_\ell M(a) = \sum_{i=1}^j \vec{M}(a)[\ell, i] f_i$, as required in `ArnoldiReduction`.

A.4 Proof of Proposition 3

Proposition 3. *Consider `HouseholderReduction` (Figure 3). We have:*

1. *The number of arithmetic operations is $O(|\Sigma|n^3)$.*
2. *`HouseholderReduction` instantiates `ArnoldiReduction`.*
3. *The computed matrices satisfy the following error bound: For each $a \in \Sigma$, the matrix $\mathcal{E}(a) \in \mathbb{R}^{\vec{n} \times n}$ with $\mathcal{E}(a) := FM(a) - \vec{M}(a)F$ satisfies*

$$\|\mathcal{E}(a)\| \leq 2\sqrt{n}\tau + cmn^3\varepsilon_{\text{mach}},$$

where $m > 0$ is such that $\|M(a)\| \leq m$ holds for all $a \in \Sigma$, and $c > 0$ is an input-independent constant.

Proof. The fact that `HouseholderReduction` is an instance of `ArnoldiReduction` was proved in the previous subsection. There we also showed the bound on the number of arithmetic operations. It remains to prove the error bound. In order to highlight the gist of the argument we consider first the case $\varepsilon_{\text{mach}} = 0$. For $0 \leq j \leq n$ we write $F_j := P_j P_{j-1} \cdots P_1$. Recall that f_1, \dots, f_j , i.e., the first j rows of F , coincide with the first j rows of F_j . We have at the end of the for loop:

$$\begin{aligned} f_\ell M(a) &= \vec{M}(a)[\ell, 1..n] F_j \\ &= \vec{M}(a)[\ell, 1..\vec{n}] F - \vec{M}(a)[\ell, j+1..\vec{n}] F[j+1..\vec{n}, \cdot] \\ &\quad + \vec{M}(a)[\ell, j+1..n] F_j[j+1..n, \cdot] \end{aligned} \quad (3)$$

So we have, for $\ell \in \mathbb{N}_{\vec{n}}$:

$$\mathcal{E}(a)[\ell, \cdot] = -\vec{M}(a)[\ell, j+1..\vec{n}] F[j+1..\vec{n}, \cdot] + \vec{M}(a)[\ell, j+1..n] F_j[j+1..n, \cdot]$$

Thus:

$$\begin{aligned} \|\mathcal{E}(a)[\ell, \cdot]\| &= \|-\vec{M}(a)[\ell, j+1..\vec{n}] F[j+1..\vec{n}, \cdot] + \vec{M}(a)[\ell, j+1..n] F_j[j+1..n, \cdot]\| \\ &\leq \|\vec{M}(a)[\ell, j+1..\vec{n}]\| + \|\vec{M}(a)[\ell, j+1..n]\| \\ &\leq 2\|\vec{M}(a)[\ell, j+1..n]\| \\ &\leq 2\tau, \end{aligned} \quad (4)$$

where the first inequality is by the fact that the rows of F and F_j are orthonormal, and the last inequality by the “if” conditional in `HouseholderReduction`.

From the row-wise bound (4) we get the following bound on the matrix norm, see [13, Lemma 6.6.a]:

$$\|\mathcal{E}(a)\|_2 \leq 2\sqrt{n}\tau$$

We now consider $\varepsilon_{\text{mach}} > 0$. We perform an error analysis similar to the one in [13, Chapter 19] for QR factorisations with Householder reflectors. Our situation is complicated by the error tolerance parameter τ ; i.e., we have to handle a combination of the errors caused by τ (as analysed above) and numerical errors. In the following we distinguish between *computed* quantities (with numerical errors and indicated with a “hat” accent) and *ideal* quantities (without numerical errors, no “hat” accent). It is important to note that when we speak of ideal quantities, we assume that we perform the *same* arithmetic operations (multiplications, additions, etc.) as in the computed case; i.e., the only difference between computed and ideal quantities is that the computed quantities come with numerical errors. In particular, we assume that the boolean values that the “if” conditional in HouseholderReduction evaluates to are the *same* for the ideal computation. The error caused by “wrong” evaluations of the conditional are already captured in the analysis above. For the remaining analysis it suffices to add the (purely) numerical error.

Concretely, we write $\hat{F} \in \mathbb{R}^{\vec{n} \times n}$ for the computed version of F , and \hat{f}_ℓ for its rows. For $a \in \Sigma$ we write $\vec{M}(a) \in \mathbb{R}^{\vec{n} \times n}$ for the *computed* quantity, as we do not consider an ideal version (and to avoid clutter). So we wish to bound the norm of $\mathcal{E}(a) := \hat{F}M(a) - \vec{M}(a)[\cdot, 1..\vec{n}]\hat{F}$. By observing that \hat{F} arises by (numerically) multiplying (at most n) Householder reflectors, and by invoking the analysis from [13, Chapter 19.3] (specifically the computation leading to Equation (19.13)) we have

$$\hat{F} = F + \tilde{O}(n^{2.5}\varepsilon_{\text{mach}}), \quad (5)$$

where by $\tilde{O}(p(n)\varepsilon_{\text{mach}})$ for a polynomial p we mean a matrix or vector A of appropriate dimension with $\|A\| \leq cp(n)\varepsilon_{\text{mach}}$ for a constant $c > 0$.⁶ (In (5) we would have $A \in \mathbb{R}^{\vec{n} \times n}$ with $\|A\| \leq cn^{2.5}\varepsilon_{\text{mach}}$.) For $a \in \Sigma$ and $\ell \in \mathbb{N}_{\vec{n}}$ define:

$$g_{\ell,a} := \hat{f}_\ell M(a) \quad (6)$$

Let $\hat{g}_{\ell,a}$ be the corresponding computed quantity. Then we have, see [13, Chapter 3.5]:

$$g_{\ell,a} = \hat{g}_{\ell,a} + \tilde{O}(mn^{1.5}\varepsilon_{\text{mach}}) \quad (7)$$

Observe from the algorithm that $\vec{M}(a)[\ell, \cdot]$ is computed by applying at most n Householder reflectors to $\hat{g}_{\ell,a}$. It follows with [13, Lemma 19.3]:

$$\vec{M}(a)[\ell, \cdot] = \left(\hat{g}_{\ell,a} + \tilde{O}(mn^2\varepsilon_{\text{mach}}) \right) P_1 P_2 \dots P_j, \quad (8)$$

⁶ We do not “hunt down” the constant c . The analysis in [13, Chapter 19.3] is similar. There the analogous constants are not computed either but are called “small”.

where the P_i are ideal Householder reflectors for subvectors of the computed \vec{M} , as specified in the algorithm. As before, define $F_j := P_j P_{j-1} \cdots P_1$. Then it follows directly from (8):

$$\vec{M}(a)[\ell, \cdot] F_j = \hat{g}_{\ell, a} + \tilde{O}(mn^2 \varepsilon_{\text{mach}}) \quad (9)$$

By combining (6), (7) and (9) we obtain

$$\hat{f}_\ell M(a) = \vec{M}(a)[\ell, \cdot] F_j + \tilde{O}(mn^2 \varepsilon_{\text{mach}}) \quad (10)$$

Using again the fact that the first j rows of F coincide with the first j rows of F_j , we have as in (3):

$$\begin{aligned} \vec{M}(a)[\ell, \cdot] F_j &= \vec{M}(a)[\ell, 1.. \vec{n}] F - \vec{M}(a)[\ell, j+1.. \vec{n}] F[j+1.. \vec{n}, \cdot] \\ &\quad + \vec{M}(a)[\ell, j+1.. n] F_j[j+1.. n, \cdot] \end{aligned} \quad (11)$$

Using (5) we have:

$$\vec{M}(a)[\ell, 1.. \vec{n}] F = \vec{M}(a)[\ell, 1.. \vec{n}] \hat{F} + \tilde{O}(mn^{2.5} \varepsilon_{\text{mach}}) \quad (12)$$

By combining (10), (11) and (12) we get:

$$\begin{aligned} &\hat{f}_\ell M(a) - \vec{M}(a)[\ell, 1.. \vec{n}] \hat{F} \\ &= -\vec{M}(a)[\ell, j+1.. \vec{n}] F[j+1.. \vec{n}, \cdot] \\ &\quad + \vec{M}(a)[\ell, j+1.. n] F_j[j+1.. n, \cdot] + \tilde{O}(mn^{2.5} \varepsilon_{\text{mach}}) \end{aligned} \quad (13)$$

We have:

$$\begin{aligned} &\|\mathcal{E}(a)[\ell, \cdot]\| \\ &= \|\hat{f}_\ell M(a) - \vec{M}(a)[\ell, 1.. \vec{n}] \hat{F}\| && \text{def. of } \mathcal{E}(a) \\ &\leq \|\vec{M}(a)[\ell, j+1.. \vec{n}] F[j+1.. \vec{n}, \cdot] + \vec{M}(a)[\ell, j+1.. n] F_j[j+1.. n, \cdot]\| && \text{by (13)} \\ &\quad + cmn^{2.5} \varepsilon_{\text{mach}} \\ &\leq 2\tau + cmn^{2.5} \varepsilon_{\text{mach}} && \text{as in (4)} \end{aligned}$$

From this row-wise bound we get the desired bound on the matrix norm, see [13, Lemma 6.6.a]:

$$\|\mathcal{E}(a)\|_2 \leq 2\sqrt{n}\tau + cmn^3 \varepsilon_{\text{mach}}$$

□

A.5 Proof of Theorem 4

Theorem 4. *Consider the following algorithm:*

function HouseholderMinimisation

input: WA $\mathcal{A} = (n, \Sigma, M, \alpha, \eta)$; error tolerance parameter $\tau \geq 0$

output: minimised WA $\mathcal{A}' = (n', \Sigma, M', \alpha', \eta')$.

compute forward reduction (F, \vec{M}) of \mathcal{A} using HouseholderReduction
form $\vec{\mathcal{A}} := (\vec{n}, \Sigma, \vec{M}, \vec{\alpha}, \vec{\eta})$ as the forward WA of \mathcal{A} with base F
compute backward reduction (B, M') of $\vec{\mathcal{A}}$ using HouseholderReduction
form $\mathcal{A}' := (n', \Sigma, M', \alpha', \eta')$ as the backward WA of $\vec{\mathcal{A}}$ with base B
return \mathcal{A}'

We have:

1. The number of arithmetic operations is $O(|\Sigma|n^3)$.
2. For $\varepsilon_{mach} = \tau = 0$, the computed WA \mathcal{A}' is minimal and equivalent to \mathcal{A} .
3. Let $\tau > 0$. Let $m > 0$ such that $\|A\| \leq m$ holds for all $A \in \{M(a), \vec{M}(a), M'(a) \mid a \in \Sigma\}$. Then for all $w \in \Sigma^*$ we have

$$|L_{\mathcal{A}}(w) - L_{\mathcal{A}'}(w)| \leq 4|w|\|\alpha\|m^{|w|-1}\|\eta\|\sqrt{n}\tau \\ + c \max\{|w|, 1\}\|\alpha\|m^{|w|}\|\eta\|n^3\varepsilon_{mach},$$

where $c > 0$ is an input-independent constant.

Part 1. follows from Proposition 3.1. Part 2. follows from Proposition 3.2. and Theorem 2. It remains to prove part 3. We use again the notation \hat{F} for the computed version of F , as in the proof of Proposition 3. We have the following lemma:

Lemma 14. Consider HouseholderReduction, see Figure 3. Let $m > 0$ such that $\|M(a)\| \leq m$ and $\|\vec{M}(a)\| \leq m$ hold for all $a \in \Sigma$. Let $b := 2\sqrt{n}\tau + cmn^3\varepsilon_{mach}$ be the bound from Proposition 3. For all $w \in \Sigma^*$ we have:

$$\|\hat{F}M(w) - \vec{M}(w)\hat{F}\| \leq b|w|m^{|w|-1}$$

Proof. We proceed by induction on $|w|$. The base case, $|w| = 0$, is trivial. Let $|w| \geq 0$ and $a \in \Sigma$. With the matrix $\mathcal{E}(a)$ from Proposition 3 we have:

$$\begin{aligned} \hat{F}M(aw) - \vec{M}(aw)\hat{F} &= \hat{F}M(a)M(w) - \vec{M}(a)\vec{M}(w)\hat{F} \\ &= (\vec{M}(a)\hat{F} + \mathcal{E}(a))M(w) - \vec{M}(a)\vec{M}(w)\hat{F} \\ &= \vec{M}(a)(\hat{F}M(w) - \vec{M}(w)\hat{F}) + \mathcal{E}(a)M(w) \end{aligned}$$

Using the induction hypothesis, Proposition 3, and the bounds on $\|\vec{M}(a)\| \leq m$ and $\|M(w)\| \leq m^{|w|}$ we obtain:

$$\|\hat{F}M(aw) - \vec{M}(aw)\hat{F}\| \leq mb|w|m^{|w|-1} + bm^{|w|} = b(|w| + 1)m^{|w|}$$

□

Now we can prove part 3. of Theorem 4:

Proof (of Theorem 4, part 3.). We use again the \tilde{O} -notation from Proposition 3. The vector \hat{f}_1 (the first row of \hat{F}) is computed by applying one Householder reflector to $e(1)$. So we have by [13, Lemma 19.3]:

$$\hat{f}_1 = \frac{\pm\alpha}{\|\alpha\|} + \tilde{O}(n\varepsilon_{\text{mach}})$$

Hence it follows:

$$\vec{\alpha}\hat{F} = \vec{\alpha}[1]\hat{f}_1 = \alpha + \tilde{O}(\|\alpha\|n\varepsilon_{\text{mach}}) \quad (14)$$

The vector $\vec{\eta}$ is computed by multiplying \hat{F} with η . So we have by [13, Chapter 3.5]:

$$\vec{\eta} = \hat{F}\eta + \tilde{O}(\|\eta\|n^{1.5}\varepsilon_{\text{mach}}) \quad (15)$$

Let $w \in \Sigma^*$. We have:

$$\begin{aligned} & |L_{\mathcal{A}}(w) - L_{\vec{\mathcal{A}}}(w)| \\ &= |\alpha M(w)\eta - \vec{\alpha}\vec{M}(w)\vec{\eta}| \\ &= |\vec{\alpha}\hat{F}M(w)\eta - \vec{\alpha}\vec{M}(w)\hat{F}\eta| + \tilde{O}(\|\alpha\|m^w\|\eta\|n^{1.5}\varepsilon_{\text{mach}}) \quad \text{by (14), (15)} \\ &= |w|\|\alpha\| (2\sqrt{n}\tau + cmn^3\varepsilon_{\text{mach}}) m^{|w|-1}\|\eta\| + \tilde{O}(\|\alpha\|m^w\|\eta\|n^{1.5}\varepsilon_{\text{mach}}) \quad \text{Lemma 14} \\ &= 2|w|\|\alpha\|m^{|w|-1}\|\eta\|\sqrt{n}\tau + \tilde{O}(\max\{|w|, 1\}\|\alpha\|m^{|w|}\|\eta\|n^3\varepsilon_{\text{mach}}) \end{aligned}$$


One can show the same bound on $|L_{\vec{\mathcal{A}}}(w) - L_{\mathcal{A}'}(w)|$ in the same way. The statement follows. \square

A.6 Image Compression


We have implemented the algorithm of Theorem 4 in a prototype tool in C++ using the Boost uBLAS library, which provides basic matrix and vector data structures. Further, we have built a translator between images and automata: it loads compressed images, constructs an automaton based on recursive algorithm sketched in the main body of the paper, feeds this automaton to the minimiser, reads back the minimised automaton, and displays the resulting compressed image.

We have applied our image compression tool to some images. Figure 4a shows a picture with a resolution of 150×150 pixels. The weighted automaton that encodes the image exactly has 33110 states. Applying minimisation with error tolerance parameter $\tau = 10^{-6}$ yields an automaton with 229 states, which leads to the compressed picture in Figure 4b. Larger values of τ lead to smaller automata and blurrier pictures, see Figures 4c and 4d, where the pictures change perceptibly.


We remark that our tool is not meant to deliver state-of-the-art image compression, which would require many tweaks, as indicated in [15].




(a) original image: 33110 states



(b) $\tau = 10^{-6}$: 229 states



(c) $\tau = 1.5 \times 10^{-2}$: 175 states



(d) $\tau = 2 \times 10^{-2}$: 121 states

Fig. 4: Image compression via WA minimisation with different values of the error tolerance parameter τ .

B Proofs of Section 4

B.1 Continuation of the proof of Proposition 6

Proposition 6. *The restricted hypercube problem can in polynomial time be reduced to the PA minimisation problem.*

Proof. Consider the reduction given in the main body of the paper. Observe that we have for all $i \in \{2, \dots, k\}$ and all $s \in \mathbb{N}_d$ that

$$L_{\mathcal{A}}(a_i b_s) = M(a_i)[1, i] M(b_s)[i, k+1] = p_i[s]. \quad (16)$$

It remains to show the correctness of the reduction, i.e., we need to show that there is a set $Q = \{q_1, \dots, q_\ell\} \subseteq [0, 1]^d$ with $\text{conv}(Q) \supseteq P$ if and only if there is a PA $\mathcal{A}' = (\ell + 1, \Sigma, M', \alpha', \eta')$ equivalent to \mathcal{A} .

For the “only if” direction, suppose that $Q = \{q_1, \dots, q_\ell\}$ such that $\text{conv}(Q) \supseteq P$. As $p_1 = (0, \dots, 0) \in P$ is a vertex of the hypercube, we have $p_1 \in Q$, say $q_1 = (0, \dots, 0)$. As $\text{conv}(Q) \supseteq P$, for each $i \in \{2, \dots, k\}$ there are $\lambda_1^{(i)}, \dots, \lambda_\ell^{(i)} \geq 0$ with $\sum_{j=1}^\ell \lambda_j^{(i)} = 1$ and

$$p_i = \sum_{j=1}^\ell \lambda_j^{(i)} q_j = \sum_{j=2}^\ell \lambda_j^{(i)} q_j. \quad (17)$$

Build the PA $\mathcal{A}' = (\ell + 1, \Sigma, M', \alpha', \eta')$ as follows. Set $M'(a_i)[1, j] := \lambda_j^{(i)}$ and $M'(b_s)[j, \ell + 1] := q_j[s]$ for all $i \in \{2, \dots, k\}$ and all $j \in \{2, \dots, \ell\}$ and all $s \in \mathbb{N}_d$, and set all other entries of M' to 0. Set $\alpha' := e(1)$ and $\eta' := e(\ell + 1)^T$. Then $\mathcal{A}, \mathcal{A}'$ are equivalent, as

$$L_{\mathcal{A}}(a_i b_s) \stackrel{(16)}{=} p_i[s] \stackrel{(17)}{=} \sum_{j=2}^\ell \lambda_j^{(i)} q_j[s] = \sum_{j=2}^\ell M'(a_i)[1, j] M'(b_s)[j, \ell + 1] = L_{\mathcal{A}'}(a_i b_s).$$

For the “if direction”, suppose that $\mathcal{A}' = (\ell + 1, \Sigma, M', \alpha', \eta')$ is a PA with $L_{\mathcal{A}'} = L_{\mathcal{A}}$. For any vector $\beta \in [0, 1]^{\ell+1}$ we define $\text{supp}(\beta) := \{j \in \mathbb{N}_{\ell+1} \mid \beta[j] > 0\}$. Define the following subsets of $\mathbb{N}_{\ell+1}$:

$$\begin{aligned} J_1 &:= \text{supp}(\alpha') \cap \bigcup_{i \in \{2, \dots, k\}, s \in \mathbb{N}_d} \text{supp}(M'(a_i)M'(b_s)\eta') \\ J_2 &:= \bigcup_{i \in \{2, \dots, k\}} \text{supp}(\alpha' M'(a_i)) \cap \bigcup_{s \in \mathbb{N}_d} \text{supp}(M'(b_s)\eta') \\ J_3 &:= \bigcup_{i \in \{2, \dots, k\}, s \in \mathbb{N}_d} \text{supp}(\alpha' M'(a_i)M'(b_s)) \cap \text{supp}(\eta') \end{aligned}$$

Recall that $L_{\mathcal{A}'} = L_{\mathcal{A}}$. Since $L_{\mathcal{A}}(b_s) = 0$ for all s , we have $J_1 \cap J_2 = \emptyset$. Since $L_{\mathcal{A}}(\tau) = 0$, we have $J_1 \cap J_3 = \emptyset$. Since $L_{\mathcal{A}}(a_i) = 0$ for all i , we have $J_2 \cap J_3 = \emptyset$. If one of J_1, J_2, J_3 is the empty set, then $J_1 = J_2 = J_3 = \emptyset$ and we have

$L_{\mathcal{A}}(w) = 0$ for all $w \in \Sigma^*$, so then by (16) we have $P = \{(0, \dots, 0)\}$, and one can take $Q = P$. So we can assume for the rest of the proof that J_1, J_2, J_3 are all non-empty. But they are pairwise disjoint, so it follows $|J_2| \leq \ell - 1$. Without loss of generality, assume $1 \notin J_2$. For $i \in \{2, \dots, k\}$ and $j \in J_2$, define $\lambda_j^{(i)} \geq 0$ and $q_j \in [0, 1]^d$ with

$$\lambda_j^{(i)} = (\alpha' M'(a_i)) [j] \quad \text{and} \quad q_j [s] = (M'(b_s) \eta') [j] \quad \text{for } s \in \mathbb{N}_d.$$

Let $q_1 := (0, \dots, 0)$. Since $\alpha' M'(a_i)$ is stochastic, one can choose $\lambda_1^{(i)} \geq 0$ so that $\sum_{j \in \{1\} \cup J_2} \lambda_j^{(i)} = 1$. We have:

$$\begin{aligned} p_i[s] &\stackrel{(16)}{=} L_{\mathcal{A}}(a_i b_s) = L_{\mathcal{A}'}(a_i b_s) = \alpha' M'(a_i) M'(b_s) \eta' \\ &= \sum_{j \in J_2} (\alpha' M'(a_i)) [j] (M'(b_s) \eta') [j] = \sum_{j \in \{1\} \cup J_2} \lambda_j^{(i)} q_j [s] \end{aligned}$$

It follows that $P \subseteq \text{conv}(Q)$ holds for $Q := \{q_j \mid j \in \{1\} \cup J_2\}$, with $|Q| \leq \ell$. \square

B.2 Proof of Proposition 8

Proposition 8. *The hypercube problem is NP-hard. This holds even for the restricted hypercube problem.*

Proof. We reduce 3SAT to the hypercube problem. Let x_1, \dots, x_N be the variables and let $\varphi = c_1 \wedge \dots \wedge c_M$ be a 3SAT formula. Each clause c_j is a disjunction of three literals $c_j = l_{j,1} \vee l_{j,2} \vee l_{j,3}$, where $l_{j,k} = x_{j,k}^-$ or $l_{j,k} = x_{j,k}^+$ and $x_{j,k} \in \{x_1, \dots, x_N\}$. (It is convenient in the following to distinguish between a variable and a positive literal, so we prefer the notation x_i^- and x_i^+ over the more familiar $\neg x_i$ and x_i for literals.) We can assume that no clause appears twice in φ and that no variable appears twice in the same clause. Define a set D of coordinates:

$$D := \{x_i^*, y_i, z_i \mid i \in \mathbb{N}_N\} \cup \{c_j^* \mid j \in \mathbb{N}_M\}$$

We take $d := |D| = 3N + M$. For $u \in D$ denote by $e(u) \in \{0, 1\}^D$ the vector with $e(u)[u] = 1$ and $e(u)[u'] = 0$ for $u' \in D \setminus \{u\}$. For $i \in \mathbb{N}_N$, define shorthands $f(x_i^-) := e(y_i)$ and $f(x_i^+) := e(y_i) + e(z_i)$. Observe that those points are vertices of the hypercube.

Define:

$$\begin{aligned}
P_{var} &:= \{e(x_i^*) + f(x_i^-), e(x_i^*) + f(x_i^+) \mid i \in \mathbb{N}_N\} \\
P_{cla} &:= \{e(c_j^*) + f(l_{j,1}), e(c_j^*) + f(l_{j,2}), e(c_j^*) + f(l_{j,3}) \mid j \in \mathbb{N}_M\} \\
p(x_i) &:= \frac{1}{2}e(x_i^*) + e(y_i) + \frac{1}{2}e(z_i) \\
&= \frac{1}{2}e(x_i^*) + \frac{1}{2}f(x_i^-) + \frac{1}{2}f(x_i^+) && \text{for } i \in \mathbb{N}_N \\
p(c_j) &:= \frac{2}{3}e(c_j^*) + \frac{1}{3}f(l_{j,1}) + \frac{1}{3}f(l_{j,2}) + \frac{1}{3}f(l_{j,3}) && \text{for } j \in \mathbb{N}_M \\
P &:= P_{var} \cup P_{cla} \cup \{p(x_1), \dots, p(x_N), p(c_1), \dots, p(c_M)\}
\end{aligned}$$

Figure 5 visualizes the points in P .




Fig. 5: Reduction from 3SAT to the hypercube problem. The left figure visualizes the $\{z_i, x_i^*\}$ -face of the hypercube with the y_i -coordinate = 1 and all other coordinates = 0. The black points are in P . Observe that $p(x_i) \in \text{conv}(\{e(x_i^*) + f(x_i^-), f(x_i^+)\})$ and $p(x_i) \in \text{conv}(\{e(x_i^*) + f(x_i^+), f(x_i^-)\})$. The right figure visualizes six hypercube vertices and a point $p(c_j) \in P$. The black points are in P . Observe that $p(c_j) \in \text{conv}(\{e(c_j^*) + f(l_{j,k(1)}), e(c_j^*) + f(l_{j,k(2)}), f(l_{j,k(3)})\})$ for all $k(1), k(2), k(3)$ with $\{k(1), k(2), k(3)\} = \{1, 2, 3\}$.

Observe that $|P| = 3N + 4M$. Take $\ell := 3N + 3M$.

First we show that if φ is satisfiable, then there is a set $Q \subseteq [0, 1]^d$ with $|Q| \leq \ell$ and $\text{conv}(Q) \supseteq P$. Let $\sigma : \{x_1, \dots, x_N\} \rightarrow \{\text{true}, \text{false}\}$ be an assignment that satisfies φ . Define:

$$\begin{aligned}
s_i &:= \begin{cases} f(x_i^-) & \text{if } \sigma(x_i) = \text{false} \\ f(x_i^+) & \text{if } \sigma(x_i) = \text{true} \end{cases} && \text{for } i \in \mathbb{N}_N \\
Q &:= P_{var} \cup P_{cla} \cup \{s_1, \dots, s_N\}
\end{aligned}$$

We have $|Q| = 3N + 3M$. Clearly, $\text{conv}(Q) \supseteq P_{var} \cup P_{cla}$. Moreover:

- Let $i \in \mathbb{N}_N$. If $\sigma(x_i) = \text{false}$, then $p(x_i) = \frac{1}{2}(e(x_i^*) + f(x_i^+)) + \frac{1}{2}s_i$; if $\sigma(x_i) = \text{true}$, then $p(x_i) = \frac{1}{2}(e(x_i^*) + f(x_i^-)) + \frac{1}{2}s_i$.
- Let $j \in \mathbb{N}_M$. As σ satisfies φ , there are $k(1), k(2), k(3)$ such that $\{k(1), k(2), k(3)\} \in \{1, 2, 3\}$ and $\sigma(l_{j,k(1)}) = \text{true}$. Let $i \in \mathbb{N}_N$ such that $l_{j,k(1)} \in \{x_i^-, x_i^+\}$. Then $p(c_j) = \frac{1}{3}(e(c_j^*) + f(l_{j,k(2)})) + \frac{1}{3}(e(c_j^*) + f(l_{j,k(3)})) + \frac{1}{3}s_i$.

Hence we have that $\text{conv}(Q) \supseteq P$.

For the converse, let $Q \subseteq [0, 1]^d$ with $|Q| \leq 3N + 3M$ and $\text{conv}(Q) \supseteq P$. Then $Q \supseteq P_{\text{var}} \cup P_{\text{cla}}$, as P_{var} and P_{cla} consist of hypercube vertices.

Let $i \in \mathbb{N}_N$. Let $Q_i^{\text{var}} \subseteq Q$ be a minimal subset of Q with $p(x_i) \in \text{conv}(Q_i^{\text{var}})$, i.e., if Q'_i is a proper subset of Q_i^{var} , then $p(x_i) \notin \text{conv}(Q'_i)$. As $p(x_i)[x_{i'}^*] = p(x_i)[c_j^*] = 0$ holds for all $i' \in \mathbb{N}_N \setminus \{i\}$ and all $j \in \mathbb{N}_M$, we have

$$Q_i^{\text{var}} \cap (P_{\text{var}} \cup P_{\text{cla}}) \subseteq \{e(x_i^*) + f(x_i^-), e(x_i^*) + f(x_i^+)\}.$$

As $p(x_i)[x_i^*] = \frac{1}{2}$ and $p(x_i)[y_i] = 1$, there is a point $s_i \in Q_i^{\text{var}}$ with $s_i[x_i^*] \leq \frac{1}{2}$ and $s_i[y_i] = 1$ and $s_i[z_i] \in [0, 1]$ and $s_i[u] = 0$ for all other coordinates $u \in D$.

It follows that $Q = P_{\text{var}} \cup P_{\text{cla}} \cup \{s_1, \dots, s_N\}$ and $|Q| = 3N + 3M$. Let σ be any assignment with

$$\sigma(x_i) = \begin{cases} \text{false} & \text{if } s_i[z_i] = 0 \\ \text{true} & \text{if } s_i[z_i] = 1 \end{cases}.$$

We show that σ satisfies φ . Let $j \in \mathbb{N}_M$. Let $Q_j^{\text{cla}} \subseteq Q$ be a minimal subset of Q with $p(c_j) \in \text{conv}(Q_j^{\text{cla}})$, i.e., if Q'_j is a proper subset of Q_j^{cla} , then $p(c_j) \notin \text{conv}(Q'_j)$. As $p(c_j)[c_{j'}^*] = p(c_j)[x_i^*] = 0$ holds for all $j' \in \mathbb{N}_M \setminus \{j\}$ and all $i \in \mathbb{N}_N$, we have

$$Q_j^{\text{cla}} \subseteq \{e(c_j^*) + f(l_{j,1}), e(c_j^*) + f(l_{j,2}), e(c_j^*) + f(l_{j,3})\} \cup \{s_1, \dots, s_N\}.$$

As $p(c_j)[c_j^*] = \frac{2}{3} < 1$, there exists an i such that $s_i \in Q_j^{\text{cla}}$. As $s_i[y_i] = 1 > 0$, we have $p(c_j)[y_i] > 0$. Hence the variable x_i appears in c_j , so one of the following two cases holds:

- The literal x_i^- appears in c_j . As we have $p(c_j)[z_i] = 0$, it follows that $q[z_i] = 0$ holds for all $q \in Q_j^{\text{cla}}$. In particular, we have $s_i[z_i] = 0$, so $\sigma(x_i) = \text{false}$.
- The literal x_i^+ appears in c_j . Note that for all points $q \in Q_j^{\text{cla}}$ we have $q[y_i] \geq q[z_i]$. As we have $p(c_j)[y_i] = \frac{1}{3} = p(c_j)[z_i]$, it follows that $q[y_i] = q[z_i]$ holds for all $q \in Q_j^{\text{cla}}$. In particular, we have $s_i[z_i] = 1$, so $\sigma(x_i) = \text{true}$.

For both cases it follows that σ satisfies c_j . As j was chosen arbitrarily, we conclude that σ satisfies φ . This completes the reduction to the hypercube problem.

The given reduction does not put the origin in P . However, $P_{\text{var}} \cup P_{\text{cla}} \subseteq P$ consist of corners of the hypercube. One can pick one of the corners in P and apply a simple linear coordinate transformation to all points in P such that the picked corner becomes the origin. Hence the restricted hypercube problem is NP-hard as well. \square

B.3 Proof of Proposition 10

Proposition 10. *Let $\mathcal{A}_1 = (n_1, \Sigma, M_1, \alpha_1, \eta_1)$ be a PA. A PA $\mathcal{A}_2 = (n_2, \Sigma, M_2, \alpha_2, \eta_2)$ is equivalent to \mathcal{A}_1 if and only if there exist matrices $\vec{M}(a) \in \mathbb{R}^{(n_1+n_2) \times (n_1+n_2)}$ for $a \in \Sigma$ and a matrix $F \in \mathbb{R}^{(n_1+n_2) \times (n_1+n_2)}$ such that $F[1, \cdot] = (\alpha_1, \alpha_2)$, and $F(\eta_1^T, -\eta_2^T)^T = (0, \dots, 0)^T$, and*

$$F \begin{pmatrix} M_1(a) & 0 \\ 0 & M_2(a) \end{pmatrix} = \vec{M}(a)F \quad \text{for all } a \in \Sigma.$$

Proof. We say, a WA \mathcal{A} is *zero* if $L_{\mathcal{A}}(w) = 0$ holds for all $w \in \Sigma^*$. For two WAs $\mathcal{A}_i = (n_i, \Sigma, M_i, \alpha_i, \eta_i)$ (with $i = 1, 2$), define their *difference* WA $\mathcal{A} = (n, \Sigma, M, \alpha, \eta)$, where $n = n_1 + n_2$, and $M(a) = \begin{pmatrix} M_1(a) & 0 \\ 0 & M_2(a) \end{pmatrix}$ for $a \in \Sigma$, and $\alpha = (\alpha_1, \alpha_2)$, and $\eta = (\eta_1^T, -\eta_2^T)^T$. Clearly, $L_{\mathcal{A}}(w) = L_{\mathcal{A}_1}(w) - L_{\mathcal{A}_2}(w)$ holds for all $w \in \Sigma^*$. So WAs \mathcal{A}_1 and \mathcal{A}_2 are equivalent if and only if their difference WA is zero.

A WA $\mathcal{A} = (n, \Sigma, M, \alpha, \eta)$ is zero if and only if all vectors of the forward space $\langle \alpha M(w) \mid w \in \Sigma^* \rangle$ are orthogonal to η (see, e.g., [26]). It follows that a WA $\mathcal{A} = (n, \Sigma, M, \alpha, \eta)$ is zero if and only if there is a vector space $F \subseteq \mathbb{R}^n$ with $\alpha \in F$, and $F\eta = \{0\}$, and $FM(a) \subseteq F$ holds for all $a \in \Sigma$. (Here, the actual forward space is a subset of F .)

The proposition follows from those observations. \square