

The Computational Complexity of Satisfiability in State Space Models

Eric Alsmann^{a,*} and Martin Lange^{a,**}

^aUniversity of Kassel, Germany

Abstract. We analyse the complexity of the satisfiability problem *ssmSAT* for State Space Models (SSM), which asks whether an input sequence can lead the model to an accepting configuration. We find that *ssmSAT* is undecidable in general, reflecting the computational power of SSM. Motivated by practical settings, we identify two natural restrictions under which *ssmSAT* becomes decidable and establish corresponding complexity bounds. First, for SSM with bounded context length, *ssmSAT* is NP-complete when the input length is given in unary and in NEXPTIME (and PSPACE-hard) when the input length is given in binary. Second, for quantised SSM operating over fixed-width arithmetic, *ssmSAT* is PSPACE-complete resp. in EXPSpace depending on the bit-width encoding. While these results hold for diagonal gated SSM we also establish complexity bounds for time-invariant SSM. Our results establish a first complexity landscape for formal reasoning in SSM and highlight fundamental limits and opportunities for the verification of SSM-based language models.

1 Introduction

State Space Models (SSM) [5, 6, 2, 18, 9, 16, 12] have recently emerged as a promising alternative to transformer-based architectures for sequence modelling tasks. By relying on linear recurrences and pointwise transformations, SSM are able to capture long-range dependencies while maintaining computational simplicity. This structural difference compared to transformers raises new questions about computational properties of SSM, in particular with regards to formal verification. While the empirical performance of SSM has been studied extensively, formal verification remains largely unexplored. In particular, understanding the complexity of basic reasoning tasks such as is there an input the SSM accepts (satisfiability) or does the SSM produce a specific output if the input satisfies some property (reachability) [7], is essential for assessing the reliability and robustness of models deployed in safety-critical applications. Satisfiability serves as a baseline for investigating formal reasoning and formal verification due to its foundational role in capturing the core computational challenges inherent in these tasks. As detailed by Sälzer et al. [13], satisfiability abstracts a wide class of formal reasoning problems—including verification of safety properties and formal interpretability—by focusing on the existence of inputs that lead a model to exhibit a specific behavior. This abstraction removes dependency on domain-specific properties while preserving the essential computational characteristics. Take for instance

a robustness property of an SSM like: given some set of suspicious symbols $E \subseteq \Sigma$, will every input containing some symbols of E be rejected? This can be reduced to checking the satisfiability of the negated condition, i.e., whether there exists an input containing symbols from E that is still accepted. A more general introduction to the different verification problems, in particular for language models, can be found in [7]. In this way, the satisfiability problem provides a unifying and theoretically robust framework for analysing the decidability and complexity of formal reasoning across a variety of verification and interpretability scenarios.

In this paper, we initiate a systematic investigation of the satisfiability problem for SSM. The satisfiability problem asks whether there exists an input sequence that causes the model to reach an accepting configuration. We first show that, in full generality, the problem is undecidable by a reduction from the halting problem for Minsky machines. Motivated by practical settings, we then consider two natural restrictions that make the problem decidable. First, we study SSM under bounded context length, a common assumption in deployed language models. Second, we investigate SSM operating over fixed-width arithmetic, reflecting the use of quantised computation in practical implementations. We investigate these problems for different practically motivated classes of SSM.

These results establish a first complexity landscape for formal reasoning in SSM. They provide a foundation for future work on verification techniques and a better theoretical understanding of the strengths and limitations of SSM.

Related Work. There is only limited work on formal properties of State Space Models. Recent theoretical work [17, 14, 10] has exclusively dealt with the expressive power of SSM. Sarrof et al. [14] showed that SSM can express all star-free languages and that a specific class of SSM exactly characterises the star-free languages. Furthermore, Merrill et al. [10] proved that quantised SSM can recognise only languages contained in the class TC^0 .

Complexity questions regarding the satisfiability problem have also been investigated for transformer architectures: Sälzer et al. [13] studied the satisfiability problem for transformer encoders, establishing complexity results similar to ours.

2 Preliminaries

Notation. We denote the set of natural numbers as \mathbb{N} and the set of real numbers as \mathbb{R} . For a vector \mathbf{v} , we denote its i -th component as v_i . We use $\mathbf{0}$ to denote the zero vector, and $\mathbf{1}$ to denote the vector with all components set to 1. We write $\mathbf{0}_d$ and $\mathbf{1}_d$ to denote the zero and one vectors of dimension d , respectively. For a finite set M , we fix

* Corresponding Author. Email: eric.alsmann@uni-kassel.de.

** Corresponding Author. Email: martin.lange@uni-kassel.de.

an arbitrary but fixed ordering of its elements. We define the one-hot encoding $e_M : 2^M \rightarrow \{0, 1\}^{|M|}$ such that for any subset $S \subseteq M$, the vector $e_M(S)$ has 1 in each position corresponding to an element of S , and 0 elsewhere. In particular, for any single element $a \in M$, we write $e_M(a) := e_M(\{a\})$. We denote the i -th standard basis vector by e_i , with its dimension being clear from the context in which it is used. When convenient, we define large vectors by writing them as tuples of smaller vectors. For example, (\mathbf{u}, \mathbf{v}) denotes the vector obtained by concatenating the entries of \mathbf{u} and \mathbf{v} .

Feedforward Neural Networks. An (FNN-)node is a function $v : \mathbb{R}^k \rightarrow \mathbb{R}$ with $v(\mathbf{x}) = \text{relu}(\sum_{i=1}^k c_i x_i + b)$, where k is the *input dimension*, the $c_i \in \mathbb{R}$ are called *weights*, $b \in \mathbb{R}$ is the *bias* and $\text{relu} : \mathbb{R} \rightarrow \mathbb{R}$ with $\text{relu}(x) = \max(0, x)$ is the *activation function* of v . An (FNN-)layer l is a tuple of some n nodes (v_1, \dots, v_n) where each node has the same input dimension m . It computes the function $l : \mathbb{R}^m \rightarrow \mathbb{R}^n$ via $l(\mathbf{x}) = (v_1(\mathbf{x}), \dots, v_n(\mathbf{x}))$. We call m the *input* and n the *output dimension* of l . A *Feedforward Neural Network (FNN)* N consists of k layers l_1, \dots, l_k , where l_1 has input dimension m , the output dimension of l_i is equal to the input dimension of l_{i+1} for $i < k$ and the output dimension of l_k is n . The FNN N computes a function from \mathbb{R}^m to \mathbb{R}^n by $N(\mathbf{x}) = l_k(l_{k-1}(\dots l_1(\mathbf{x}) \dots))$.

For our lower bound constructions we sometimes need to make use of FNN checking specific properties, for instance: is some vector entry equal to a specific value? For easier notation we will define FNN *gadgets* for specific properties we will need later.

Lemma 1. *Let $n, n_1, \dots, n_k, m, b \in \mathbb{N}$. There are FNN*

- $N_{=b}$ s.t. $N_{=b}(n) = 1$ if $n = b$ and $N_{=b}(n) = 0$ otherwise.
- $N_{\leq b}$ s.t. $N_{\leq b}(n) = 1$ if $n \leq b$ and $N_{\leq b}(n) = 0$ otherwise.
- N_{\wedge} s.t. $N_{\wedge}(n_1, \dots, n_k) = 1$ if $n_i = 1$ for all $i \leq k$ and $N_{\wedge}(n_1, \dots, n_k) = 0$ otherwise.

Proof. Let $N_{=b}$ be the FNN computing the function $N_{=b}(x) = \text{relu}(\text{relu}(x - (b - 1)) - 2 \cdot \text{relu}(x - b))$, $N_{\leq b}$ be the FNN computing $N_{\leq b}(x) = \text{relu}(\text{relu}(b + 1 - x) - \text{relu}(b - x))$ and N_{\wedge} be the FNN computing $N_{\wedge}(x_1, \dots, x_k) = N_{=k}(x_1 + \dots + x_k)$. It is straightforward to see that the given FNN compute the required functions. \square

We will sometimes need to combine FNN. We define two operations composition (\circ) and concatenation (\parallel). Let N_1 and N_2 be FNN computing functions $\mathbb{R}^{m_i} \rightarrow \mathbb{R}^{n_i}$ for $i \in \{1, 2\}$. When $n_2 = m_1$, the FNN $N_1 \circ N_2$ computes the function $\mathbb{R}^{m_1} \rightarrow \mathbb{R}^{n_2}$ with $N_1 \circ N_2(\mathbf{x}) = N_1(N_2(\mathbf{x}))$. Syntactically, composition simply concatenates the output layer of N_2 and the input layer of N_1 setting all weights to 1.

The FNN $N_1 \parallel N_2$ computes the function $\mathbb{R}^{m_1+m_2} \rightarrow \mathbb{R}^{n_1+n_2}$ with $N_1 \parallel N_2(\mathbf{x}_1, \mathbf{x}_2) = (N_1(\mathbf{x}_1), N_2(\mathbf{x}_2))$ for all $\mathbf{x}_1 \in \mathbb{R}^{m_1}, \mathbf{x}_2 \in \mathbb{R}^{m_2}$. Syntactically concatenation first extends N_1 to have the same number of layers as N_2 by adding identity nodes for additional layers. Then each layer of N_1 is extended by the respective layer of N_2 while zeroing incoming connections from nodes of N_1 .

State Space Models. We explore State Space Models as presented in [6]. For a structured analysis of this model, we will formalize its architecture following the approach in [14].

An SSM layer l is defined as a tuple $(\mathbf{h}_0, \text{gate}, \text{inc}, \phi)$, where $\mathbf{h}_0 \in \mathbb{R}^d$. The function gate is a mapping $\mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$, inc is a mapping $\mathbb{R}^d \rightarrow \mathbb{R}^d$, and ϕ is a mapping $\mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^d$. An SSM layer transforms an input sequence of vectors $\mathbf{x}_1 \dots \mathbf{x}_k \in (\mathbb{R}^d)^+$ into an output sequence $\mathbf{z}_1 \dots \mathbf{z}_k \in (\mathbb{R}^d)^+$ through the following

process: it computes an intermediate sequence $\mathbf{h}_1 \dots \mathbf{h}_k \in (\mathbb{R}^d)^+$ via a linear recurrence

$$\mathbf{h}_t = \text{gate}(\mathbf{x}_t) \cdot \mathbf{h}_{t-1} + \text{inc}(\mathbf{x}_t) \quad \text{for } 1 \leq t \leq k$$

and subsequently generates the output via $\mathbf{z}_t = \phi(\mathbf{h}_t, \mathbf{x}_t)$.

An SSM comprising L layers is expressed as a tuple $(\text{emb}, l_1, \dots, l_L, \text{out})$. Each l_i refers to a layer as defined above, while emb is a function $\Sigma \rightarrow \mathbb{R}^d$, and out is a function $\mathbb{R}^d \rightarrow \mathbb{R}^d$, computed by an FNN. The SSM computes a function $\Sigma^+ \rightarrow \mathbb{R}$ as described: let $w = a_1 \dots a_k \in \Sigma^+$ be a word. Initially, the SSM computes the embedding $\mathbf{x}_1^0 \dots \mathbf{x}_k^0$ of the word by setting $\mathbf{x}_i^0 = \text{emb}(a_i)$. Subsequently, for each layer $1 \leq j \leq L$: compute $\mathbf{z}_1^j \dots \mathbf{z}_k^j = l_j(\mathbf{x}_1^{j-1} \dots \mathbf{x}_k^{j-1})$. Each layer's output serves as the input for the succeeding layer, meaning $\mathbf{x}_1^{j+1} \dots \mathbf{x}_k^{j+1} = \mathbf{z}_1^j \dots \mathbf{z}_k^j$. The SSM's final output, denoted as $\mathbf{y}_1 \dots \mathbf{y}_k$, is derived by applying out element-wise: $\mathbf{y}_i = \text{out}(\mathbf{z}_i^L)$. In the end, the output of $\mathcal{S}(w)$ is computed by \mathbf{y}_k . We say that \mathcal{S} accepts a word w if $\mathcal{S}(w) = 1$. Otherwise it is rejected. Let \mathfrak{S} be some class of SSM. We define the decision problem $\text{SSMSAT}[\mathfrak{S}]$ as the problem, given an SSM $\mathcal{S} \in \mathfrak{S}$ over an alphabet Σ , decide whether there is $w \in \Sigma^*$ such that $\mathcal{S}(w) = 1$.

To establish upper complexity bounds, we define the representation size of an SSM. Given an SSM over an alphabet Σ with dimension d and number of layers L , we set $|\mathcal{S}| = |\Sigma| + L + d$. We assume that the syntactic representation of \mathcal{S} is polynomial in $|\mathcal{S}|$. For any input word, the output of an SSM is computed layer-wise, with each layer requiring only a linear number of operations. This gives rise to the *polynomial-evaluation property*: for any word $w \in \Sigma^*$, the value $\mathcal{S}(w)$ can be computed in time polynomial in $|\mathcal{S}| + |w|$. Note that we assume that each SSM can be finitely represented. This implies that all weights and constants used in the SSM have a finite representation.

Fixed-Width Arithmetic. Our notion of fixed-width arithmetics are representations of numbers using a fixed amount of bits, like floating- or fixed-point arithmetic. Our results are independent of the specific choice of an implementation. We assume that all values represented in a fixed-width arithmetic use b bits for representing numbers. We say that an SSM works over fixed-width arithmetics if all computations and values occurring in the computation of the SSM are carried out using only b bits.

Minsky-Machines. A Minsky machine $M = (Q, q_0, q_f, \delta)$ consists of a finite set of states Q , a start state $q_0 \in Q$, a final state $q_f \in Q$ and a finite set of transitions $\delta \subseteq Q \times \text{Act} \times Q$ where $\text{Act} = \{\text{inc}_1, \text{inc}_2, \text{dec}_1, \text{dec}_2, \text{ztest}_1, \text{ztest}_2\}$. We additionally assume that for every state $q \in Q$ there is either only one $i \in \{1, 2\}$ and one transition with action inc_i or two transitions with actions dec_i and ztest_i . A configuration of M is a tuple $(q, c_1, c_2) \in Q \times \mathbb{N}^2$, consisting of the current state and two counter values. A step of M $(q, c_1, c_2) \vdash (q', c'_1, c'_2)$ is valid if (1) $(q, \text{inc}_i, q') \in \delta$ and $c'_i = c_i + 1$, (2) $(q, \text{dec}_i, q') \in \delta$, $c_i > 0$ and $c'_i = c_i - 1$ or (3) $(q, \text{ztest}_i, q') \in \delta$, $c_i = 0$ and $c'_i = c_i$ for $i \in \{1, 2\}$. Note that the other counter in each step remains unchanged. A run of M is a sequence valid steps starting from $(q_0, 0, 0)$. A run is accepting if it ends in a configuration with state q_f .

The problem to decide, given a Minsky machine M , whether it has an accepting run, is called MHALT. Note that due to our definition of a Minsky machine M , given a specific state of M there is always only one possible successor state. Thus, we can describe an accepting run of M as a sequence of state-action pairs rather than full configurations.

Theorem 2 ([11]). *MHALT is undecidable.*

Linear Temporal Logic LTL_f . Let \mathcal{P} be a finite set of atomic propositions. The syntax of LTL_f is defined as follows:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \mathbf{X}\varphi \mid \varphi \cup \varphi$$

Formulas of LTL_f are interpreted over finite words over the alphabet $\Sigma = 2^{\mathcal{P}}$. Given a word $w = a_1 \cdots a_n \in \Sigma^*$ and $i \in [n]$, the semantics of LTL_f is inductively defined as follows:

$$\begin{aligned} w, i \models p &\iff p \in a_i \\ w, i \models \neg\varphi &\iff w, i \not\models \varphi \\ w, i \models \varphi \wedge \psi &\iff w, i \models \varphi \text{ and } w, i \models \psi \\ w, i \models \mathbf{X}\varphi &\iff i < n \text{ and } w, i+1 \models \varphi \\ w, i \models \varphi \cup \psi &\iff \text{ex. } i \leq k \leq n : w, k \models \varphi \text{ and} \\ &\quad \text{f.a. } i \leq j < k : w, j \models \psi \end{aligned}$$

We say that w is a model of φ iff $w, 1 \models \varphi$ (or simply $w \models \varphi$). For easier notation we define $\llbracket \varphi \rrbracket^w = 1$ if $w \models \varphi$ and 0 otherwise. We may omit the superscript if w is clear from context. Furthermore, formula φ is satisfiable iff there exists $w \in \Sigma^*$ such that w is a model of φ . The size of φ is denoted by $|\varphi|$.

Theorem 3 ([3]). *Satisfiability for LTL_f is PSPACE-complete.*

3 Overview

Recent work on the expressiveness of SSM [10, 14] distinguishes between two main classes of SSM which differ in the allowed *gate* functions. The class of *time-invariant* SSM [9, 16, 12] (which we will denote \mathfrak{S}_{ti}) only allow *gate* functions such that there is a matrix $A \in \mathbb{R}^{d \times d}$ with $\text{gate}(\mathbf{x}) = A$ for all $\mathbf{x} \in \mathbb{R}^d$. The class of *diagonal* SSM [5, 2, 18] (which we denote $\mathfrak{S}_{\text{diag}}$) only use *gate* functions such that $\text{gate}(\mathbf{x})$ is a diagonal matrix but can depend on \mathbf{x} . The output of each layer is computed by a non-linear function ϕ which also gets the initial input modelling a residual connection. In practice, different kinds of non-linear functions are used. In this paper we assume ϕ to be represented by an FNN. Despite this assumption, all upper complexity bounds established in this paper also hold for any other non-linear activation function used in practice.

An overview of the results established in this paper is depicted in Table 1. In Section 4, we establish the undecidability of the satisfiability problem for SSM (Thm. 6). Our proof is based on a reduction from the halting problem for Minsky machines (MHALT), and it proceeds by constructing, for each instance of a Minsky machine M , an SSM that accepts precisely those sequences that encode accepting runs of M . Finally, we show in Theorem 7 that this undecidability result persists even under finite-width arithmetic, provided the precision grows logarithmically with the input length.

In Section 5, we turn our attention to decidable fragments of the problem. We consider two natural restrictions motivated by practical applications: bounding the input length and restricting arithmetic to fixed-width representations. In both cases we additionally make a distinction between unary and binary encodings of the word length or bit-width. While this distinction may seem subtle, it has significant implications for complexity. Intuitively, unary encoding corresponds to a setting where the available computational resources—such as the length of input sequences or the precision of arithmetic—are given explicitly and can be directly used during computation. In contrast, binary encoding specifies the size of these resources succinctly: we

are told how much space we are allowed to use, but we must first construct or simulate that space ourselves.

For the bounded context length variant $\text{SSMSAT}_{\{\text{un}, \text{bin}\}}^{\leq}[\mathfrak{S}]$, we established that the unary-encoded version is NP-complete (Corollary 14) for time-invariant and diagonal gated SSM. When considering binary encoding, we establish membership in NEXPTIME (Theorem 9) and PSPACE-hardness (Theorem 17) for diagonal SSM as well as NP-hardness for time-invariant SSM (Theorem 13)

In the second setting, we study SSM operating over fixed-width arithmetic. We establish an exponential-model property (Lemma 10), which ensures that if a word is accepted under b -bit precision, there exists an accepting word of length at most exponential in $|\mathcal{S}| + b$. We use this to prove that when the bit-width b is constant or given in unary, the problem lies in PSPACE (Theorem 11), while for binary-encoded bit-width the complexity increases to EXPSPACE (Theorem 12).

To complement these upper bounds, we develop a polynomial-time reduction from the satisfiability problem for linear-time temporal logic over finite traces (LTL_f) to SSM satisfiability (Lemma 15). This construction relies on diagonal gated SSM with constant precision and shows that the corresponding fixed-width satisfiability problems are all PSPACE-hard (Theorem 18). This yields PSPACE-completeness for constant and unary bit-width (Corollary 19), and PSPACE-hardness for binary bit-width, leaving a gap to the EXPSPACE upper bound.

For time-invariant SSM we establish NP-hardness (Theorem 20) for the cases with fixed-width arithmetic. Possible reasons for this complexity gap between time-invariant and diagonal gated SSM are discussed at the end of Section 5. Our results provide a comprehensive classification of decidable fragments of the SSM satisfiability problem and reveal fine-grained complexity distinctions based on model structure and input encoding.

4 Undecidability

In this section, we will show undecidability for time-invariant as well as diagonal gated SSM working over arbitrary precision, by reducing from the halting problem for Minsky machines. Furthermore, it is shown that the undecidability still holds if we allow the precision to grow logarithmically with the input length. The reduction constructs an SSM which recognises accepting runs of a Minsky machine. In a first step, we will show how runs of a Minsky machine can be encoded as a word over a finite alphabet.

Lemma 4. *Let $M = (Q, q_0, q_f, \delta)$ be a Minsky-machine. There is an embedding function emb and an SSM layer $l_1 \in \mathfrak{S}_{\text{ti}} \cap \mathfrak{S}_{\text{diag}}$ such that for each word $w = (q_1, a_1) \cdots (q_n, a_n) \in \Sigma^*$ with $\Sigma = \{(q', a) \mid \exists q \in Q : (q, a, q') \in \delta\}$ we have $l_1(\text{emb}(w)) = z_1 \cdots z_{|w|}$ with $z_i = (e_Q(q_i), e_Q(q_i), e_{\text{Act}}(a), c_i^1, c_i^2, 0)$.*

Proof. Let $M = (Q, q_0, q_f, \delta)$ be a Minsky machine. We define $\Sigma = \{(q', a) \mid \exists q \in Q : (q, a, q') \in \delta\}$. Therefore, each word $w \in \Sigma^*$ describes a potential run of M , skipping the starting state. We now define an SSM \mathcal{S} with one layer and $d = 2|Q| + 9$ dimensions which computes the counter states for the given run encoded in w . Let emb be the function $\Sigma \rightarrow \mathbb{R}^d$ with:

$$\text{emb}((q, a)) = (e_Q(q), e_Q(q), e_{\text{Act}}(a), u_1(a), u_2(a), 0)$$

where $u_i : \text{Act} \rightarrow \{-1, 0, 1\}$ with $u_i(\text{inc}_i) = 1$, $u_i(\text{dec}_i) = -1$, and $u_i(\text{ztest}_i) = 0$. Other inputs are mapped to zero. Let $l_1 \in \mathfrak{S}_{\text{ti}} \cap \mathfrak{S}_{\text{diag}}$ computing

$$\mathbf{h}_t = E_{[2|Q|+7, 2|Q|+8]} \cdot \mathbf{h}_{t-1} + I_d \cdot \mathbf{x}_t \quad \phi(\mathbf{x}_t, \mathbf{h}_t) = \mathbf{h}_t$$

Table 1. Overview of the results established in this paper.

	time-invariant SSM (\mathfrak{S}_{ti})	diagonal SSM (\mathfrak{S}_{diag})
arbitrary precision	undecidable (Thm. 6)	undecidable (Thm. 6)
log-precision	undecidable (Thm. 7)	undecidable (Thm. 7)
bounded context length (unary)	NP-complete (Cor. 14)	NP-complete (Cor. 14)
bounded context length (binary)	NP-hard (Thm. 13), \in NEXPTIME (Thm. 9)	PSPACE-hard (Thm. 17), \in NEXPTIME (Thm. 9)
constant bit-width	\in PSPACE (Thm. 11)	PSPACE-complete (Cor. 19)
fixed bit-width (unary)	NP-hard (Thm. 20), \in PSPACE (Thm. 11)	PSPACE-complete (Cor. 19)
fixed bit-width (binary)	NP-hard (Thm. 20), \in EXPSPACE (Thm. 12)	PSPACE-hard, \in EXPSPACE (Thm. 12)

where $E_{[i,j]}$ is the modified identity matrix with $E_{k,k} = 1$ if $i \leq k \leq j$ and 0 otherwise. During this computation only the counter dimension get accumulated, resulting in

$$h_t = \left(e_Q(q_t), e_Q(q_t), e_{Act}(a), \sum_{1 \leq i \leq t} u_1(a_i), \sum_{1 \leq i \leq t} u_2(a_i), 0 \right)$$

This concludes that l_1 together with the embedding function, maps a potential run of M in form of a state-action pair sequence to a sequence of vectors each containing the current state, the action and the intermediate counter values. \square

Once we have encoded the current machine state and the corresponding counter values, we have to check whether the state-action pairs from the input are valid transitions with respect to δ . Moreover, it has to be checked if the transitions taken were allowed with respect to the counter values.

To verify these properties, we must transfer information about the previous state into the vector at the next position. This can be achieved using a trick, which involves encoding the information about the previous state into the binary expansion of a value. This information can then be decoded back into a state using an FNN.

Lemma 5. *There is an SSM layer $l_{prev} \in \mathfrak{S}_{ti} \cap \mathfrak{S}_{diag}$ that maps a sequence $x_1 \cdots x_k \in \{0, 1\}^*$ to a sequence $z_1 \cdots z_k \in \{0, 1\}^*$ such that $z_i = x_{i-1}$ for $0 < i \leq k$ and $z_0 = 0$.*

Proof Sketch. To achieve this, we adopt an encoding technique used by [14], which encodes the previous bit in the binary expansion of a value. This encoding allows the retrieval of historical information using a simple FNN. This also works in a setting when finite-width arithmetic is used by ensuring that each bit of the history is separated by a zero in the binary expansion, which avoids false results due to rounding errors and enables precise recovery via thresholding. The SSM layer computes the recurrence:

$$h_t = \frac{1}{4} \cdot h_{t-1} + x_t$$

This recurrence corresponds to a right shift of the binary expansion of h_t by two positions. This encodes the full history of $x_1 \cdots x_{t-1}$ into the binary expansion of h_t . Subsequently the previous bit can be read out by an FNN even when using fixed-width arithmetic with a bit-width of at least six. More details can be found in the supplementary material [1]. \square

We now combine the previous lemmas to construct an SSM that verifies whether a given sequence of state-action pairs encodes an accepting run of a Minsky machine.

Theorem 6. *SSMSAT[$\mathfrak{S}_{ti} \cap \mathfrak{S}_{diag}$] is undecidable.*

Proof Sketch. We establish a reduction from MHALT to SSMSAT[$\mathfrak{S}_{ti} \cap \mathfrak{S}_{diag}$]. For each instance $M = (Q, q_0, q_f, \delta)$ of MHALT we construct an SSM $\mathcal{S}_M \in \mathfrak{S}_{ti} \cap \mathfrak{S}_{diag}$, which accepts only valid encodings of accepting runs of M .

\mathcal{S}_M has $2|Q| + 9$ dimensions and three layers. The embedding function emb and layer l_1 are given by Lemma 4 and are used to encode the state and the current counter values in each vector. Layer l_2 , based on Lemma 5, retrieves the previous state for each configuration. By setting h_0^2 to encode the starting state, we ensure that the first state-action pair encoded a valid transition from the starting state. Additionally l_2 uses the output FNN to check the validity of each transition. Thus after l_2 the SSM computes the sequence $z_1^2 \cdots z_k^2$ with

$$z_i^2 = (e_Q(q_i), e_Q(q_{i-1}), e_{Act}(a), c_i^1, c_i^2, \text{check}(q_{t-1}, a, q_t, c_i^1, c_i^2)).$$

$\text{check}(q_{t-1}, a, q_t, c_i^1, c_i^2)$ is the output of a pointwise applied FNN checking (1) if $(q_{t-1}, a, q_t) \in \delta$ and (2) if depending on the action, the counter states have a valid value. This verification is particularly crucial for *dec* and *ztest* actions, where it must be ensured that the counter state is either not negative in case of *dec* or is zero in case of *ztest*. If the check is valid, the function will output zero and in case of an invalid transition it will be one. Consequently, layer l_3 applies a linear recurrence to accumulate values in the *check* dimension. The output FNN then only has to check two things (1) whether *check* has value zero, in which case each transition was valid or has a nonzero value, which leads to rejection of the input sequence and (2) if the last vector represents the accepting state q_f . This reduction establishes the undecidability of SSMSAT[$\mathfrak{S}_{ti} \cap \mathfrak{S}_{diag}$]. The detailed proof can be found in the supplementary material [1]. \square

Next, we consider the case when the SSM works over finite precision, but the precision depends logarithmically on the input length. This setting has recently been studied in terms of expressiveness [10]. I.e. when evaluating an SSM on a word $w \in \Sigma^*$ with length $|w|$, the SSM works with $\mathcal{O}(\log(\max(|\Sigma|, |w|)))$ precision.

Theorem 7. *SSMSAT $_{\log}^{\text{fix}}[\mathfrak{S}_{ti} \cap \mathfrak{S}_{diag}]$ is undecidable.*

Proof. This proof follows the same lines as the proof of Theorem 6. Since we are working with binary representations, the number of bits required to represent a number that grows linearly with the input length increases only logarithmically. Therefore, a fixed-precision SSM that operates with $\mathcal{O}(\log(\max(|\Sigma|, |w|)))$ precision suffices to simulate the computation accurately. This ensures that all values required to verify a run can still be represented precisely enough to preserve correctness, which establishes that the undecidability result still holds under logarithmically bounded fixed-width arithmetic. \square

5 Decidable Cases

In this section, we investigate restrictions to the satisfiability problem for SSM that yield decidability. A natural restriction, frequently encountered in practical language models, is bounded context length. Let \mathfrak{S} be a class of SSM. The bounded context length satisfiability problem for SSM, denoted by $\text{SSMSAT}_{\{\text{un}, \text{bin}\}}^{\leq}[\mathfrak{S}]$, is as follows: given $\mathcal{S} \in \mathfrak{S}$ and $n \in \mathbb{N}$ (encoded in unary or binary), decide whether there exists a word w with $|w| \leq n$ such that $\mathcal{S}(w) = 1$.

The second restriction, also frequently encountered in practice, is the use of SSM in quantised environments, where computations are performed using fixed-width arithmetic, such as fixed- or floating-point representations. We consider two different settings for SSM over fixed-width arithmetic. The first setting assumes the bit-width b to be constant. For $b \in \mathbb{N}$, define the decision problem $b\text{-SSMSAT}^{\text{fix}}[\mathfrak{S}]$: given $\mathcal{S} \in \mathfrak{S}$, does there exist a word w such that $\mathcal{S}(w) = 1$ when \mathcal{S} is evaluated over fixed-width arithmetic with bit-width b ? The second setting considers the bit-width as part of the input. We distinguish between the cases where the bit-width is given in *unary* or *binary*. The decision problem $\text{SSMSAT}_{\{\text{un}, \text{bin}\}}^{\text{fix}}[\mathfrak{S}]$ is: given $\mathcal{S} \in \mathfrak{S}$ and $b \in \mathbb{N}$ (encoded in unary or binary), decide whether there exists a word w such that $\mathcal{S}(w) = 1$ when \mathcal{S} is evaluated over fixed-width arithmetic with bit-width b .

5.1 Upper Bounds

We start the complexity analysis by establishing upper bounds for the mentioned problems. All of our upper bounds will hold for time-invariant as well as diagonal gated SSM. Therefore we will use \mathfrak{S} to refer to the class of $\mathfrak{S}_{\text{ti}} \cup \mathfrak{S}_{\text{diag}}$.

Theorem 8. $\text{SSMSAT}_{\text{un}}^{\leq}[\mathfrak{S}]$ is in NP.

Proof. We first show membership in NP using a standard guess-and-check approach. Given an SSM $\mathcal{S} \in \mathfrak{S}$ and $n \in \mathbb{N}$ (encoded in unary), we nondeterministically guess a word $w \in \Sigma^*$ of length at most n and check whether $\mathcal{S}(w) = 1$. By the polynomial-evaluation property, this verification can be performed in polynomial time. \square

Theorem 9. $\text{SSMSAT}_{\text{bin}}^{\leq}[\mathfrak{S}]$ is in NEXPTIME.

Proof. Let $\mathcal{S} \in \mathfrak{S}$ be an SSM and n be the binary encoded word length. Hence, the maximum word length to check is exponential in n . This naturally gives us in NEXPTIME algorithm to decide $\text{SSMSAT}_{\text{bin}}^{\leq}[\mathfrak{S}]$ by just guessing a word w with $|w| \leq 2^n$ and then checking if $\mathcal{S}(w) = 1$. Using the polynomial-evaluation property $\mathcal{S}(w)$ can be computed in time polynomial in $|\mathcal{S}| + |w| = |\mathcal{S}| + 2^n$. \square

Turning our attention to the fixed-width arithmetic restriction, we first show that SSM operating over fixed-width arithmetic possess an *exponential-model property*: if an SSM \mathcal{S} accepts some word w when using b -bit fixed-width arithmetic, then \mathcal{S} also accepts a word w' with $|w'| \leq 2^{\text{poly}(|\mathcal{S}|+b)}$. Consequently, restricting SSM to fixed-width arithmetic has a similar effect to bounding the context length.

Lemma 10. Let $\mathcal{S} \in \mathfrak{S}$ be an SSM with L layers and d dimensions, operating over fixed-point arithmetic with bit-width b . If \mathcal{S} accepts a word w , then there exists a word w' of length at most 2^{2Ldb} such that \mathcal{S} accepts w' .

Proof. Consider \mathcal{S} with L layers and d dimensions, using b bits of fixed-point precision. For each position $1 \leq i \leq n$ and each layer $0 \leq l \leq L$, \mathcal{S} computes vectors \mathbf{h}_i^l . Each \mathbf{h}_i^l depends only on \mathbf{h}_{i-1}^l and the input \mathbf{x}_i^{l-1} , and since each entry is represented with b bits, there are at most 2^{db} possible distinct vectors per layer. Across all L layers, the total number of possible combined layer states is at most 2^{Ldb} .

Now, along the computation on a word of length $n > 2^{2Ldb}$, by the pigeonhole principle, there must exist indices $1 \leq i < j \leq n$ such that, for all $1 \leq l \leq L$, $\mathbf{h}_i^l = \mathbf{h}_j^l$ and $\mathbf{h}_i^{l-1} = \mathbf{h}_j^{l-1}$. That is, the complete vector state across all layers at positions i and j coincides. Therefore, the computation from position $j + 1$ onward is identical to the computation from $i + 1$ onward. Thus, the word obtained by removing the segment $a_{i+1} \cdots a_j$ from w is also accepted by \mathcal{S} . Iterating this argument, we obtain an accepted word w' of length at most 2^{2Ldb} . \square

We will use this exponential-model property in order to prove membership in PSPACE for the cases where the bit-width is constant or part of the input and encoded in unary.

Theorem 11. $b\text{-SSMSAT}^{\text{fix}}[\mathfrak{S}]$ and $\text{SSMSAT}_{\text{un}}^{\text{fix}}[\mathfrak{S}]$ are in PSPACE.

Proof. The computation of each next vector \mathbf{h}_i^l in an SSM layer depends only on the previous state \mathbf{h}_{i-1}^l and the embedding of the current symbol. By Lemma 10, it suffices to consider words of length up to 2^{2Ldb} , where L is the number of layers, d the dimension, and b the bit-width.

We describe a nondeterministic algorithm using polynomial space. According to Savitch's Theorem [15] this suffices to prove membership in PSPACE. Maintain a counter for the current word length (requiring $2Ldb$ bits), and nondeterministically generate a word $a_1 \cdots a_k \in \Sigma^*$ with $k \leq 2^{2Ldb}$, symbol by symbol. After each symbol a_i is chosen, compute \mathbf{h}_i^l from \mathbf{h}_{i-1}^l and $\text{emb}(a_i)$ for all $1 \leq l \leq L$, discarding \mathbf{h}_{i-1}^l as it is no longer needed. Each step requires only polynomial time and space in $|\mathcal{S}|$. Storing all \mathbf{h}_{i-1}^l for all layers requires Ldb bits. When b is constant, the total space is $\mathcal{O}(|\mathcal{S}|^2)$. Thus, $b\text{-SSMSAT}^{\text{fix}}[\mathfrak{S}]$ is in PSPACE. In the case of b being part of the input and encoded in binary the total space is $\mathcal{O}((|\mathcal{S}| + b)^3)$. Thus, $\text{SSMSAT}_{\text{un}}^{\text{fix}}[\mathfrak{S}]$ is in PSPACE. \square

For binary-encoded bit-width, the problem lies in EXPSpace, due to the succinctness of the encoding.

Theorem 12. $\text{SSMSAT}_{\text{bin}}^{\text{fix}}[\mathfrak{S}]$ is in EXPSpace.

Proof. The proof is analogous to Theorem 11, except that with binary encoding of the bit-width, the space required to store an intermediate SSM computation is $Ld2^{|b|}$. This yields an EXPSpace algorithm for the problem. \square

5.2 Lower Bounds

Having established upper bounds for all of the problems, we now aim to show lower bounds. While the upper bounds did not depend on the class of gates the SSM used, we will see that this is not the case when it comes to lower bounds.

Theorem 13. $\text{SSMSAT}_{\text{un}}^{\leq}[\mathfrak{S}_{\text{ti}} \cap \mathfrak{S}_{\text{diag}}]$ is NP-hard.

Proof. We reduce from the Zero-One Integer Programming Problem, which is known to be NP-complete [8]. Given an integer matrix $A \in \mathbb{N}^{d \times d}$ and a vector $\mathbf{b} \in \mathbb{N}^d$ (with numbers encoded in binary), the problem asks whether there exists $\mathbf{v} \in \{0, 1\}^d$ such that $A\mathbf{v} = \mathbf{b}$.

For a vector $\mathbf{v} \in \{0, 1\}^d$, let $\text{supp}(\mathbf{v}) = \{i \in \{1, \dots, d\} \mid v_i \neq 0\}$ denote its support. Then \mathbf{v} can be written as a sum of standard base vectors: $\mathbf{v} = \sum_{i \in \text{supp}(\mathbf{v})} \mathbf{e}_i$, so $A\mathbf{v} = \sum_{i \in \text{supp}(\mathbf{v})} A\mathbf{e}_i$.

We construct an SSM $S_{A,b}$ over the alphabet $\Sigma = \{1, \dots, d\}$, with dimension $2d$ and embedding function $\text{emb}(i) = (\mathbf{e}_i, \mathbf{0}_d)$. The input word encodes the support of \mathbf{v} as a sequence of indices which are mapped to the corresponding basis vectors. To ensure that each dimension appears at most once, we duplicate the vector: the first d components accumulate $A\mathbf{v}$, while the second d components count occurrences of each index. We define a single layer with recurrence $\mathbf{h}_0 = \mathbf{0}$ and

$$\mathbf{h}_t = I \cdot \mathbf{h}_{t-1} + \begin{pmatrix} A & 0 \\ I_d & 0 \end{pmatrix} \mathbf{x}_t,$$

so that

$$\mathbf{h}_t = \begin{pmatrix} \sum_{1 \leq i \leq t} A\mathbf{x}_i, \\ \sum_{1 \leq i \leq t} \mathbf{x}_i \end{pmatrix}.$$

The output FNN N_b then verifies that the first d entries equal \mathbf{b} and each entry in the second d components is at most 1, ensuring a valid encoding of $\mathbf{v} \in \{0, 1\}^d$. Using the gadgets from Lemma 1, N_b can be defined as

$$N_\wedge \circ (N_{=b_1} \parallel \dots \parallel N_{=b_d} \parallel N_{\leq 1} \parallel \dots \parallel N_{\leq 1})$$

The resulting SSM $S_{A,b}$ has only polynomial size A, b yielding a polynomial time reduction. The constructed SSM is both time-invariant and diagonal. Thus, $\text{SSMSAT}_{\text{im}}^{\leq}$ is NP-complete for both time-invariant and diagonal SSM. \square

As a consequence of Theorem 13 and Theorem 8 we get the following result.

Corollary 14. $\text{SSMSAT}_{\text{im}}^{\leq}[\mathfrak{S}_{ii}]$, $\text{SSMSAT}_{\text{im}}^{\leq}[\mathfrak{S}_{\text{diag}}]$ are NP-complete.

In order to prove PSPACE lower bounds for some problems, the next lemma shows that we can reduce satisfiability of LTL_f to SSM-SAT by constructing an SSM which is polynomial in the size of the formula and additionally only needs a constant bit-width to operate.

Lemma 15. For any LTL_f-formula φ over Σ , there exists an SSM $S_\varphi \in \mathfrak{S}_{\text{diag}}$ such that for all $w = a_1 \dots a_n \in \Sigma^*$ holds: $w \models \varphi$ iff S_φ accepts $\overleftarrow{w} = a_n \dots a_1$.

Proof. Given a formula φ over Σ , we construct an SSM S_φ over the same alphabet with $\mathcal{O}(|\varphi|)$ layers and dimensions, such that for all $w = a_1 \dots a_n \in \Sigma^*$, w is a model of φ if and only if S_φ accepts the word $\overleftarrow{w} = a_n \dots a_1$. The reversal is due to the fact that while in LTL_f one can specify properties about future positions, SSM can only transfer information about previous positions to the current state. By reversing the models the SSM can check a formula $\mathbb{X}\varphi$ by looking one step back and a formula $\varphi \cup \psi$ by remembering that ψ held in the past and φ ever since. Moreover, S_φ requires only 6 bits of precision, regardless of φ .

We introduce the *copy matrix* $C^{(i \rightarrow j)} \in \mathbb{R}^{d \times d}$, defined by $C^{(i \rightarrow j)} = \mathbf{e}_j \mathbf{e}_i^\top$. Multiplying $C^{(i \rightarrow j)}$ with a vector $\mathbf{x} \in \mathbb{R}^d$ yields a vector whose j -th entry is x_i and all other entries are zero. The copy matrix is used to transfer values between specific dimensions in the SSM state vector, enabling the encoding of logical dependencies.

Let φ be an LTL_f-formula over the alphabet $\Sigma = 2^P$, with k subformulas and k_x X-formulas. We construct an SSM S_φ with $k + k_x$ layers and $|\mathcal{P}| + k + 1$ dimensions. The embedding function $\text{emb} : \Sigma \rightarrow \mathbb{R}^{k+1}$ is defined as $\text{emb}(a) = (\mathbf{e}_P(a), \mathbf{0}_k, 1)$. Since

SSM layers are linear while the subformulas of φ form a tree, we fix a total order on φ 's subformulas $\varphi_1 < \dots < \varphi_k = \varphi$ such that $i < j$ if φ_i is a subformula of φ_j . This order is used to construct the SSM layer-wise, ensuring that when evaluating a subformula, all its dependencies have already been computed in lower layers. Each subformula, proposition $a \in \mathcal{P}$, and the constant 1 correspond to a particular dimension in the vectors during the computation. For easier notation we identify subformulas, propositions and the constant 1 with the corresponding dimension. The layers l_1, \dots, l_k , with $l_i \in \mathfrak{S}_{\text{diag}}$, are defined as follows:

- *Case 1:* $\varphi_i = a$ for $a \in \mathcal{P}$. The value is read from dimension a and copied to dimension φ_i :

$$\mathbf{h}_t = 0 \cdot \mathbf{h}_{t-1} + (I + C^{(a \rightarrow \varphi_i)}) \cdot \mathbf{x}_t, \quad \phi(\mathbf{x}, \mathbf{h}) = \mathbf{h}.$$

- *Case 2:* $\varphi_i = \neg\varphi_j$ with $j < i$. Compute 1 minus the value at φ_j :

$$\mathbf{h}_t = 0 \cdot \mathbf{h}_{t-1} + (I + C^{(1 \rightarrow \varphi_i)} - C^{(\varphi_j \rightarrow \varphi_i)}) \cdot \mathbf{x}_t, \quad \phi(\mathbf{x}, \mathbf{h}) = \mathbf{h}.$$

- *Case 3:* $\varphi_i = \varphi_j \wedge \varphi_k$ with $j, k < i$. The conjunction is computed as $\max(0, \llbracket \varphi_j \rrbracket + \llbracket \varphi_k \rrbracket - 1)$:

$$\mathbf{h}_t = 0 \cdot \mathbf{h}_{t-1} + (I + C^{(\varphi_j \rightarrow \varphi_i)} + C^{(\varphi_k \rightarrow \varphi_i)} - C^{(1 \rightarrow \varphi_i)}) \cdot \mathbf{x}_t,$$

and $\phi(\mathbf{x}, \mathbf{h}) = \mathbf{h}'$, where $h'_j = \text{relu}(h_j)$ if $j = \varphi_i$, and $h'_j = h_j$ otherwise.

- *Case 4:* $\varphi_i = \mathbb{X}\varphi_j$ with $j < i$. In the reversed setting of the SSM, in order to evaluate $\mathbb{X}\varphi_j$ at some position we have to check the previous state vector. This is done in two consecutive layers. This first layer simply copies the evaluation of φ_j to φ_i at each position.

$$\mathbf{h}_t = 0 \cdot \mathbf{h}_{t-1} + (I + C^{\varphi_j \rightarrow \varphi_i} \cdot \mathbf{x}_t) \quad \phi(\mathbf{x}, \mathbf{h}) = \mathbf{h}$$

The second layer uses Lemma 5, to construct a layer that overrides these values with the previous value. As stated in Lemma 5 this requires only six bits of precision.

- *Case 5:* $\varphi_i = \varphi_j \cup \varphi_k$ with $j, k < i$. Unrolling the until operator we get:

$$\varphi_j \cup \varphi_k = (\varphi_j \wedge \mathbb{X}(\varphi_j \cup \varphi_k)) \vee \varphi_k$$

Because the SSM checks the formulas reversed, we check if since the last position where φ_k was satisfied φ_j held continuously. This check actually needs a non time-invariant diagonal gate, because the previous state of φ_i has to be multiplied with the current state of φ_j in order to check the recurrence. Additionally the current state of φ_k is added to φ_i to restart the recurrence as soon as φ_k holds again. This can be encoded as

$$\mathbf{h}_t = \text{diag}(C^{(\varphi_j \rightarrow \varphi_i)} \mathbf{x}_t) \cdot \mathbf{h}_{t-1} + (I + C^{(\varphi_k \rightarrow \varphi_i)}) \cdot \mathbf{x}_t,$$

and $\phi(\mathbf{x}, \mathbf{h}) = \mathbf{h}'$, where $h'_j = \min(1, h_j)$ if $j = \varphi_i$, and $h'_j = h_j$ otherwise. The function $\min(1, x)$ is FNN-computable by $\text{relu}(x) - \text{relu}(-x) - \text{relu}(x-1)$. The min-operation is needed to scale down the resulting value to be either one or zero.

Let $w = a_1 \dots a_n \in \Sigma^*$ and $\mathbf{z}_1 \dots \mathbf{z}_n$ the sequence of vectors after all layers defined above. Correctness of the reduction follows now from the fact that $w, i \models \varphi_j$ iff $(\mathbf{z}_{n-(i-1)})_{\varphi_j} = 1$. The output FNN simply has to check whether \mathbf{z}_n at dimension φ equals 1. It remains to note that all matrix coefficients and intermediate values used in this construction can be represented using a constant number of bits, as well as intermediate numbers during the computation do not need more than six bits ensuring that the SSM operates even with constant bit-width. As stated above S_φ has polynomial size in $|\varphi|$. \square

We note that in the previous lemma, evaluating the *until*-operator required a time-dependent diagonal gate, due to the need to multiply the previous and next hidden state. Thus, the lower bounds based on the construction will only hold for the class $\mathfrak{S}_{\text{diag}}$. First, we look at the case of bounded context length SSM, where the length is given in binary. The previous lemma showed that we can find a polynomial time reduction from the satisfiability problem for LTL_f -formula to SSMSAT . In order to make this reduction also work in the case of bounded context length we need an argument, that a satisfiable LTL_f -formula has a model of length at most exponential in the size of the formula.

Lemma 16 ([4, 3]). *For every satisfiable LTL_f formula φ , there is a word w with $|w| \leq |\varphi| \cdot 2^{|\varphi|}$ such that $w \models \varphi$.*

Using this exponential-model property for LTL_f -formula we show a lower bound for diagonal gated SSM with bounded context length and with the length encoded in binary.

Theorem 17. $\text{SSMSAT}_{\text{bin}}^{\leq}[\mathfrak{S}_{\text{diag}}]$ is PSPACE-hard.

Proof. We use the reduction from Lemma 15. Given an LTL_f -formula φ , we can construct an SSM $\mathcal{S}_\varphi \in \mathfrak{S}_{\text{diag}}$ which accepts exactly the reversed models of φ . Lemma 16 showed that if φ is satisfiable, it has a model $w \models \varphi$ with $|w| \leq |\varphi| \cdot 2^{|\varphi|}$. Due to Lemma 15 the same holds for \mathcal{S}_φ . Therefore, we can specify the context length as $\log(|\varphi| \cdot 2^{|\varphi|}) = \log(|\varphi|) + |\varphi|$ yielding a polynomial time reduction from the satisfiability problem for LTL_f -formula to $\text{SSMSAT}_{\text{bin}}^{\leq}[\mathfrak{S}_{\text{diag}}]$. \square

Next, we consider the case of SSM working over fixed-width arithmetic. First, we show a tight lower bound for the case of the bit-width being constant.

Theorem 18. $b\text{-SSMSAT}_{\text{bin}}^{\text{fix}}[\mathfrak{S}_{\text{diag}}]$ (for $b \geq 6$), $\text{SSMSAT}_{\text{un}}^{\text{fix}}[\mathfrak{S}_{\text{diag}}]$ and $\text{SSMSAT}_{\text{bin}}^{\text{fix}}[\mathfrak{S}_{\text{diag}}]$ are PSPACE-hard.

Proof. PSPACE-hardness follows from the polynomial reduction from LTL_f to SSMSAT in Lemma 15. The resulting SSM only needed a constant bit-width of six, independently from the given formula. This of course also shows PSPACE-hardness for $\text{SSMSAT}_{\text{un}}^{\text{fix}}[\mathfrak{S}_{\text{diag}}]$ and $\text{SSMSAT}_{\text{bin}}^{\text{fix}}[\mathfrak{S}_{\text{diag}}]$, because of the independence of the bit-width from the given formula. \square

While this leaves an exponential complexity gap for $\text{SSMSAT}_{\text{bin}}^{\text{fix}}[\mathfrak{S}_{\text{diag}}]$ as it is contained in EXPSPACE and PSPACE-hard, it is a tight bound for the other two problems.

Corollary 19. $b\text{-SSMSAT}_{\text{un}}^{\text{fix}}[\mathfrak{S}_{\text{diag}}]$ (for $b \geq 6$) and $\text{SSMSAT}_{\text{bin}}^{\text{fix}}[\mathfrak{S}_{\text{diag}}]$ are PSPACE-complete.

Establishing an EXPSPACE lower bound for $\text{SSMSAT}_{\text{bin}}^{\text{fix}}[\mathfrak{S}_{\text{diag}}]$ does not seem to be an easy task. Normally succinct encodings of problems imply an exponential increase in complexity. In this case though, only a portion of the input—the bit-width—is succinctly encoded, whereas the syntactical description of the SSM is given in standard (non-succinct) representation. It remains an open problem if this is enough to establish EXPSPACE-hardness or if the upper bound can be lowered by exploiting some inherent properties of the SSM. We have seen that the construction in Lemma 15 for establishing the PSPACE-hardness depended on the use of non time-invariant diagonal gates, due to the recursive nature of the *until*-operator. For time-invariant SSM we will show NP-hardness as a lower bound.

Theorem 20. $\text{SSMSAT}_{\text{un}}^{\text{fix}}[\mathfrak{S}_{\text{ti}}]$ and $\text{SSMSAT}_{\text{bin}}^{\text{fix}}[\mathfrak{S}_{\text{ti}}]$ are NP-hard.

Proof. This follows by the same reduction as in Theorem 13 for the context length restriction. Given a matrix $A \in \mathbb{N}^{d \times d}$ and vector $\mathbf{b} \in \mathbb{N}^d$, we must show that the bit-width required for all intermediate computations of $A\mathbf{v}$ for some $\mathbf{v} \in \{0, 1\}^d$ is polynomial in the size of A and \mathbf{b} . Let m be the largest entry in A . In the worst case, $\mathbf{v} = \mathbf{1}_d$, so the largest entry in $A\mathbf{v}$ is $d \cdot m$. Thus, all intermediate values can be represented with bit-width $\log(d) + \log(m)$, which is polynomial in the size of A and \mathbf{b} . Thus, $\text{SSMSAT}_{\text{un}}^{\text{fix}}[\mathfrak{S}_{\text{ti}}]$ is NP-hard. The same obviously holds for $\text{SSMSAT}_{\text{bin}}^{\text{fix}}[\mathfrak{S}_{\text{ti}}]$ as $\log(\log(d) + \log(m))$ is even smaller. \square

This leaves the question why we could not close the complexity gap between time-invariant and diagonal SSM. Intuitively in time-invariant SSM, the calculation of a hidden state $\mathbf{h}_t = M\mathbf{h}_{t-1} + B(\mathbf{x}_t)$ is an affine linear transformation of \mathbf{h}_{t-1} , while in the case of diagonal SSM the transformation is non-linear, due to the multiplication of the input-dependent gate with the previous state.

We have seen that this is necessary for our PSPACE lower bound. In the NP lower bound however, the hidden state computation only uses a linear transformation. It remains an open question if this non-linearity actually has an impact on the complexity and the PSPACE, resp. EXPSPACE upper bounds can be improved or if the lower bounds can be tightened despite the restriction to linear transformations.

6 Outlook

The satisfiability problem studied in this paper provides a natural foundation for the study on formal verification of State Space Models. Although abstract, it captures the core challenge underlying many verification tasks, such as proving or refuting safety properties. Its general formulation, detached from specific property types, ensures that undecidability and complexity results immediately extend to broader reasoning problems. Thus, satisfiability serves as a baseline for understanding the fundamental limits of verifying SSM and motivates future research into efficient and sound verification techniques. Establishing an EXPSPACE lower bound for $\text{SSMSAT}_{\text{bin}}^{\text{fix}}[\mathfrak{S}_{\text{diag}}]$ remains an open problem. While the binary encoding of the bit-width introduces a form of succinctness, the SSM itself is not succinctly represented, leaving it unclear whether this suffices for EXPSPACE-hardness or allows for a lower upper bound by exploiting structural properties of the model.

Additionally, the complexity gap between time-invariant and diagonal SSMs has not been closed. A key distinction lies in the state update: time-invariant SSMs apply affine linear updates, whereas diagonal SSMs involve non-linear gating mechanisms. While such non-linearity is essential for our PSPACE lower bound, the NP lower bound relies solely on linear updates. It is an open question whether this non-linearity influences the overall complexity or whether tighter lower bounds can be achieved despite the linearity constraint. Future work should further investigate the role of non-linearity in complexity bounds and explore potential refinements of both upper and lower bounds for these SSM variants.

This paper demonstrates that sound and complete formal reasoning techniques for SSM face significant computational complexity challenges, even in settings with bounded context length and fixed-width arithmetic. Nevertheless, practical approaches such as SMT encodings of the verification problem or the development of incomplete or unsound verification procedures could offer promising pathways for ensuring desirable properties in SSM-based language models.

References

- [1] E. Alsmann and M. Lange. The computational complexity of satisfiability in state space models, 2025. arXiv, Full version of this paper.
- [2] S. De, S. L. Smith, A. Fernando, A. Botev, G. Cristian-Muraru, A. Gu, R. Haroun, L. Berrada, Y. Chen, S. Srinivasan, G. Desjardins, A. Doucet, D. Budden, Y. W. Teh, R. Pascanu, N. D. Freitas, and C. Gulcehre. Griffin: Mixing Gated Linear Recurrences with Local Attention for Efficient Language Models, Feb. 2024.
- [3] G. De Giacomo and M. Y. Vardi. Linear temporal logic and linear dynamic logic on finite traces. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI '13*, pages 854–860, Beijing, China, Aug. 2013. AAAI Press.
- [4] S. Demri, V. Goranko, and M. Lange. *Temporal Logics in Computer Science: Finite-State Systems*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge, 2016. ISBN 978-1-107-02836-4. doi: 10.1017/CBO9781139236119.
- [5] A. Gu and T. Dao. Mamba: Linear-Time Sequence Modeling with Selective State Spaces. In *First Conference on Language Modeling*, Aug. 2024.
- [6] A. Gu, K. Goel, and C. Re. Efficiently Modeling Long Sequences with Structured State Spaces. In *International Conference on Learning Representations*, 2022.
- [7] X. Huang, W. Ruan, W. Huang, G. Jin, Y. Dong, C. Wu, S. Bensalem, R. Mu, Y. Qi, X. Zhao, K. Cai, Y. Zhang, S. Wu, P. Xu, D. Wu, A. Freitas, and M. A. Mustafa. A survey of safety and trustworthiness of large language models through the lens of verification and validation. *Artificial Intelligence Review*, 57(7):175, June 2024. ISSN 1573-7462. doi: 10.1007/s10462-024-10824-0.
- [8] R. M. Karp. Reducibility among Combinatorial Problems. In R. E. Miller, J. W. Thatcher, and J. D. Bohlinger, editors, *Complexity of Computer Computations: Proceedings of a Symposium on the Complexity of Computer Computations, Held March 20–22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, and Sponsored by the Office of Naval Research, Mathematics Program, IBM World Trade Corporation, and the IBM Research Mathematical Sciences Department*, pages 85–103. Springer US, Boston, MA, 1972. ISBN 978-1-4684-2001-2. doi: 10.1007/978-1-4684-2001-2_9.
- [9] H. Mehta, A. Gupta, A. Cutkosky, and B. Neyshabur. Long range language modeling via gated state spaces. In *The Eleventh International Conference on Learning Representations*, 2023.
- [10] W. Merrill, J. Petty, and A. Sabharwal. The illusion of state in state-space models. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *ICML'24*, pages 35492–35506, Vienna, Austria, July 2024. JMLR.org.
- [11] M. L. Minsky. *Computation: Finite and Infinite Machines*. Prentice-Hall, Inc., USA, 1967. ISBN 978-0-13-165563-8.
- [12] A. Orvieto, S. L. Smith, A. Gu, A. Fernando, C. Gulcehre, R. Pascanu, and S. De. Resurrecting recurrent neural networks for long sequences. In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*, Honolulu, Hawaii, USA, 2023. JMLR.org.
- [13] M. Sälzer, E. Alsmann, and M. Lange. Transformer Encoder Satisfiability: Complexity and Impact on Formal Reasoning. In *The Thirteenth International Conference on Learning Representations*, Oct. 2024.
- [14] Y. Sarrof, Y. Veitsman, and M. Hahn. The Expressive Capacity of State Space Models: A Formal Language Perspective. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, Nov. 2024.
- [15] W. J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4(2):177–192, 1970. ISSN 0022-0000. doi: 10.1016/S0022-0000(70)80006-X.
- [16] Y. Sun, L. Dong, S. Huang, S. Ma, Y. Xia, J. Xue, J. Wang, and F. Wei. Retentive Network: A Successor to Transformer for Large Language Models, Aug. 2023.
- [17] A. Terzic, M. Hersche, G. Camposampiero, T. Hofmann, A. Sebastian, and A. Rahimi. On the expressiveness and length generalization of selective state space models on regular languages. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(19):20876–20884, Apr. 2025. doi: 10.1609/aaai.v39i19.34301.
- [18] S. Yang, B. Wang, Y. Shen, R. Panda, and Y. Kim. Gated linear attention transformers with hardware-efficient training. In *Proceedings of the 41st International Conference on Machine Learning, ICML'24*, Vienna, Austria, 2024. JMLR.org.