
Turing Completeness of Bounded-Precision Recurrent Neural Networks

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Previous works have proved that recurrent neural networks (RNNs) are Turing-
2 complete. However, in the proofs, it is assumed that the RNNs can have neurons
3 with unbounded precision, which is neither practical in implementation nor biolog-
4 ically plausible. To remove this assumption, we propose a dynamically growing
5 memory module made of neurons of fixed precision. The memory module dynam-
6 ically recruits new neurons when more memories are needed, and releases them
7 when memories become irrelevant. To illustrate the memory module’s capacity, we
8 prove that a 54-neuron bounded-precision RNN with growing memory modules can
9 simulate any Turing Machines. Furthermore, we analyze the Turing completeness
10 of both unbounded-precision RNNs and bounded-precision RNNs, revisiting and
11 extending the theoretical foundation of RNNs.

12 1 Introduction

13 Symbolic (such as Turing machines) and sub-symbolic processing (such as adaptive neural networks)
14 are two competing methods for representing and processing information, each with its pros and cons.
15 An ultimate way to combine symbolic and sub-symbolic capabilities is by enabling the running of
16 algorithms on a neural substrate, which means a neural network that can simulate a Universal Turing
17 Machine. Previous works [1, 2, 3] have shown that this is possible – there exists a recurrent neural
18 network (RNN) that can simulate a Universal Turing Machine. However, these proofs assumed
19 a couple of neurons with unbounded precision that hold the same number of symbols used when
20 processing the Turing machine’s tape. Here we provide another simulation of Turing machines by an
21 RNN with bounded-precision neurons only.

22 The general idea works as follows: the Turing machine’s tape is stored in a growing memory module
23 - a stack of neurons with pushing and popping operations controlled by neurons in the RNN. The
24 size of the growing memory module is determined by the usage of the tape – the RNN dynamically
25 recruits new neurons when more memory is needed (i.e. for a task that is more memory-intensive) and
26 releases them when memories become irrelevant. The neurons in the stack (except the top neuron)
27 are not updated in the steps of RNN, saving computational cost for memories that are not in the focus
28 of computing and thus do not require changing. The capability of the growing memory module is
29 illustrated by the existence of a 54-neuron bounded-precision RNN with growing memory modules
30 that can simulate any Turing machines.

31 [4, 5] proposed an RNN with a memory bank, called Neural Turing Machine, that is fully differentiable
32 and thus can be trained by gradient descent, sharing similarities with previous work [6]. Though
33 inspired by Turing Machines, bounded-precision Neural Turing Machines are not Turing-complete
34 since it has a fixed-size memory, but can likely be used for training the one we propose here.

35 Our proposed growing memory module is inspired by biological memory systems. The process of
36 dynamically recruiting new neurons when more memory is necessary is also observed in biological

37 memory systems. Neurogenesis is the process by which new neurons are produced in the central
38 nervous system. It is most active during early development, but continuous throughout life. In
39 adult vertebrates, neurogenesis is known to occur in the dentate gyrus (DG) of the hippocampal
40 formation [7] and the subventricular zone (SVZ) of the lateral ventricles [8]. Since DG is well-known
41 in neuroscience for its role in pattern separation for memory encoding [9, 10], this suggests that
42 biological memory systems also dynamically recruit new neurons. The rate of neurogenesis in adult
43 mice has been shown to be higher if they are exposed to a wider variety of experiences [11]. This
44 further suggests a role for self-regulated neurogenesis in scaling up the number of new memories
45 that can be encoded and stored during one’s lifetime without catastrophic forgetting of previously
46 consolidated memories. Besides the mechanism of recruiting new neurons, the process of storing
47 neurons in the growing memory module also shares some similarities with biological memory
48 consolidation, a process by which short-term memory is transformed into long-term memory [12, 13].
49 Compared to short-term memory, long-term memory is more long-lasting and robust to interference.
50 This is similar to the neurons stored in the growing memory module – the values of these neurons
51 (except the top neuron in the stack) remain unchanged and cannot be interfered by the RNN, providing
52 a mechanism to store information stably.

53 Simulation of a Turing machine by an RNN with growing memory modules poses a practical and
54 biologically inspired way to combine symbolic and sub-symbolic capabilities. All neurons in the RNN
55 and growing memory modules have fixed precision, and the enhanced RNN can thus be implemented
56 easily. The number of neurons required in the RNN is constant in the tape’s length, and the time
57 complexity of the simulation is linear in the number of steps required for the Turing machine to
58 compute the output. And most interestingly, the number of computations required for the simulation
59 does not grow with the memory size or the tape’s length since neurons in the growing memory module
60 (except the top neuron in the stack) are not updated.

61 Besides the design of the growing memory module, this paper also analyzes the Turing complete-
62 ness of both unbounded-precision RNN and bounded-precision RNN, revisiting and extending the
63 theoretical foundation of RNNs. For unbounded-precision RNNs, we prove that there exists a 40-
64 neuron unbounded-precision RNN that is Turing-complete, which is the smallest Turing-complete
65 RNN to date. For bounded-precision RNNs, we analyze the relationship between the number of
66 neurons and the precision of an RNN when simulating a Turing machine, showing that the number of
67 bounded-precision neurons required to simulate a Turing machine is linear in the tape’s length.

68 The remainder of the paper is structured as follows. Section 2 describes the preliminary of the paper,
69 including the definition of Turing machines and RNNs. Section 3 revisits and extends theories relating
70 to simulating a Turing machine with unbounded-precision RNNs, proving the existence of a 40-
71 neuron unbounded-precision RNN that is Turing-complete. Section 4 presents the growing memory
72 module and proves the existence of a 54-neuron bounded-precision RNN with two growing memory
73 modules that is Turing-complete. Section 5 establishes new theories relating to the relationship
74 between the number of neurons and the precision of an RNN when simulating a Turing machine.
75 Section 6 concludes the paper.

76 2 Background and Notation

77 A Turing machine is a 7-tuple $\mathcal{M} = (Q, \Sigma, \Gamma, \delta, q_0, B, F)$, where Q is a finite set of states, Σ is a finite
78 set of input symbols, Γ is a finite set of tape symbols (note that $\Sigma \subset \Gamma$), $\delta : Q \times \Gamma \rightarrow Q \times \Gamma \times \{L, R\}$
79 is the transition rule, $q_0 \in Q$ is the initial starting state, B is the blank symbol (note that $B \in \Gamma$
80 but $B \notin \Sigma$), and $F \subset Q$ is the set of final state. We only consider deterministic Turing machines
81 in this paper. The *instantaneous description* of a Turing machine can be represented by a 3-tuple
82 $(q, t_l, t_r) \in (Q, \Gamma^*, \Gamma^*)$ where q denotes the state, t_l and t_r denotes the string of symbols in the left
83 and right tape respectively. We assume the leftmost symbol is the closest symbol to the read/write
84 head for both t_l and t_r (that is, the left tape is reversed in the representation) and the infinite blank
85 symbols are omitted. The set of all possible instantaneous description is denoted as $\mathcal{X} := (Q, \Gamma^*, \Gamma^*)$.
86 The *complete dynamic map* of \mathcal{M} , denoted as $\mathcal{P}_{\mathcal{M}} : \mathcal{X} \rightarrow \mathcal{X}$, is defined by: 1. determinate the next
87 transition by $q', r', d = \delta(q, r)$ where r is the leftmost symbol in t_l , denoted by $t_{l,1}$, and q' is the
88 next state; 2. replace r by r' ; 3. move $t_{l,1}$ to the leftmost of t_r if $d = L$ and $t_{r,1}$ to the leftmost of
89 t_l if $d = R$ (If there are no symbols left in t_l or t_r , append a blank symbol B to it). The *partial*
90 *input-output function* of \mathcal{M} , denoted as $\mathcal{P}_{\mathcal{M}}^* : \mathcal{X} \rightarrow \mathcal{X}$, is defined by applying $\mathcal{P}_{\mathcal{M}}$ repeatedly until
91 $q \in F$, and is undefined if it is not possible to have $q \in F$ by applying $\mathcal{P}_{\mathcal{M}}$ repeatedly.

92 A recurrent neural network (RNN) is a neural network consisting of n neurons. The value of neuron
 93 i at time $t \in \{1, 2, \dots\}$, denoted as $x_i(t) \in \mathbb{Q}$ (\mathbb{Q} is the set of rational numbers), is computed by
 94 an affine transformation of the values of all neurons in the previous state followed by an activation
 95 function σ , i.e. $x_i(t) = \sigma(\sum_{j=1}^n w_{ij}x_j(t-1) + b_i)$, where w_{ij} is the weight and b_i is the bias; or in
 96 vector form:

$$\mathbf{x}(t) = \sigma(W\mathbf{x}(t-1) + \mathbf{b}), \quad (1)$$

97 where $\mathbf{x}(t) \in \mathbb{Q}^n$, $W \in \mathbb{R}^{n \times n}$ and $\mathbf{b} \in \mathbb{R}^n$. This defines a mapping $\mathcal{T}_{W,\mathbf{b}} : \mathbb{Q}^n \rightarrow \mathbb{Q}^n$ which
 98 characterizes an RNN. Also, we only consider the saturated-linear function in this paper; that is:

$$\sigma(x) := \begin{cases} 0, & \text{if } x < 0 \\ x, & \text{if } 0 \leq x \leq 1 \\ 1, & \text{if } x > 1. \end{cases} \quad (2)$$

99 We call a neuron $x_i(t)$ to have precision p in base b if for all $t > 0$, $x_i(t)$ can be expressed as
 100 $\sum_{k=1}^p \frac{a_k(t)}{\prod_{j=1}^k b_j(t)}$ for some strings $a(t) \in \{0, 1, \dots, b\}^p$ and $b(t) \in \{1, \dots, b\}^p$.

101 For a string a , we use a_k to denote the k^{th} symbol in a and $a_{j:k}$ to denote the string $a_j a_{j+1} \dots a_k$. For
 102 a function f that maps from a set \mathbb{Y} to a subset of \mathbb{Y} , we denote f^n as the n^{th} iterate of f where
 103 $1 \leq n < \infty$. For any two vectors $\mathbf{x} \in \mathbb{R}^m$, $\mathbf{y} \in \mathbb{R}^n$, we denote $\mathbf{x} \oplus \mathbf{y} \in \mathbb{R}^{m+n}$ as the concatenation
 104 of the two vectors.

105 3 Turing Completeness of Unbounded-Precision RNNs

106 To simulate a Turing machine \mathcal{M} by an RNN, we first consider how to encode the instantaneous
 107 description $(q, t_l, t_r) \in \mathcal{X}$ by a vector of rational numbers, with which an RNN can be initialized.
 108 For the state $q \in Q$, we encode it with $\lceil \log_2 |Q| \rceil$ binary values, denoted as $\rho^q : Q \rightarrow \{0, 1\}^{\lceil \log_2 |Q| \rceil}$,
 109 with each possible combination of binary values representing a specific state.

110 For the left tape $t_l \in \Gamma^*$ and the right tape $t_r \in \Gamma^*$, we use fractal encoding to encode them into
 111 two rational numbers. The fractal encoding we use is similar to [1, 2], but we generalize here
 112 to an arbitrary number of symbols. Without loss of generality, assume that the tape symbols are
 113 $\Gamma = \{1, 2, \dots, |\Gamma|\}$. Then, define the fractal encoding $\rho^t : \Gamma^* \rightarrow \mathbb{Q}$ by:

$$\rho^t(t_m) := \left(\sum_{i=1}^{|t_m|} \frac{2t_{m,i} - 1}{(2|\Gamma|)^i} \right) + \frac{B}{(2|\Gamma|)^{|t_m|} \cdot (2|\Gamma| - 1)}. \quad (3)$$

114 *Example.* Let $|\Gamma| = 4$, $t_l = (2, 3, 4, 2, 3)$, and $B = 1$. Then $\rho^t(t_l) = \frac{3}{8} + \frac{5}{8^2} + \frac{7}{8^3} + \frac{3}{8^4} + \frac{5}{8^5} + \frac{1}{8^6} +$
 115 $\frac{1}{8^7} + \frac{1}{8^8} + \dots = \frac{3}{8} + \frac{5}{8^2} + \frac{7}{8^3} + \frac{3}{8^4} + \frac{5}{8^5} + \frac{1}{8^{5.7}}$. Note that the last term in (3) represents the infinite
 116 blank symbols of a tape.

117 The particular choice of this fractal encoding is due to the easy manipulation of the top symbols, and
 118 the details can be found in Appendix A. The possible values of fractal encoding also have a close
 119 relationship with the Cantor set [1, 2].

120 However, this fractal encoding assumes that neurons can have the same precision as $|t_m|$, the tape's
 121 size, in base $2|\Gamma|$. As the tape in a Turing machine has an unbounded length, it means that we
 122 require neurons with unbounded precision. Note that this is different from infinite precision - at
 123 any steps of the Turing machine, the precision of $\rho^t(t_m)$ is finite. Still, the tape's size may be large,
 124 and it is difficult to compute with the high-precision encoding. We will discuss how to remove this
 125 unbounded-precision assumption in Section 4.

126 Finally, we need to have a specific encoding for the top symbol in each tape: $\rho^r : \Gamma \rightarrow \{0, 1\}^{|\Gamma|-1}$,
 127 defined by:

$$\rho_i^r(s) := 1\{s > i\}, \quad (4)$$

128 where $1 \leq i \leq |\Gamma| - 1$. That is, the i coordinate of $\rho^r(s)$ is one if and only if the symbol s has a
 129 value larger than i . This fully encodes the symbol s and is used to encode both $t_{l,1}$ and $t_{r,1}$.

130 Together, define the encoding function $\rho : \mathcal{X} \rightarrow \mathbb{Q}^{2|\Gamma| + \lceil \log_2 |Q| \rceil + |Q||\Gamma| + 5}$ by:

$$\rho(q, t_l, t_r) = \rho^q(q) \oplus \rho^t(t_l) \oplus \rho^t(t_r) \oplus \rho^r(t_{l,1}) \oplus \rho^r(t_{r,1}) \oplus \mathbf{0}, \quad (5)$$

131 where $\mathbf{0}$ is a zero vector of size $|Q||\Gamma| + 5$. Also, let $\rho^{-1} : \rho(\mathcal{X}) \rightarrow \mathcal{X}$ be the function such that
 132 $\rho^{-1}(\rho(x)) = x$ for all $x \in \mathcal{X}$ (note that ρ is injective), and we call ρ^{-1} the decoder function.

133 It should be noted that $\rho^q(q) \oplus \rho^t(t_l) \oplus \rho^t(t_r)$ is sufficient to decode the instantaneous description. We
 134 include $\rho^r(t_{l,1}) \oplus \rho^r(t_{r,1}) \oplus \mathbf{0}$ only because it facilitates the operation of the RNN. For example, some
 135 neurons that are initialized with the zero values will compute the values required in the intermediate
 136 steps of updates.

137 Given an instantaneous description $x \in \mathcal{X}$, we can then initialize neurons in an RNN with values
 138 $\rho(x)$. Then, it is possible to construct the parameters of RNN such that the update given by the RNN
 139 on these neurons is the same as the update given by the Turing machine:

140 **Theorem 1.** *Given a Turing machine \mathcal{M} , there exists an injective function $\rho : \mathcal{X} \rightarrow \mathbb{Q}^N$ and an*
 141 *n -neuron unbounded-precision RNN $\mathcal{T}_{W,b} : \mathbb{Q}^n \rightarrow \mathbb{Q}^n$, where $n = 2|\Gamma| + \lceil \log_2 |Q| \rceil + |Q||\Gamma| + 5$,*
 142 *such that for all instantaneous descriptions $x \in \mathcal{X}$,*

$$143 \quad \rho^{-1}(\mathcal{T}_{W,b}^3(\rho(x))) = \mathcal{P}_{\mathcal{M}}(x). \quad (6)$$

144 The proof can be found in Appendix A. In other words, for any Turing machine \mathcal{M} , there exists an
 145 RNN such that every three steps of the RNN yield the same result as one step of the Turing Machine.
 146 This means that we can simulate any Turing machines using an RNN with a linear time. To be
 147 specific, the *partial input-output function* of an RNN, denoted as $\mathcal{T}_{W,b}^* : \mathbb{Q}^N \rightarrow \mathbb{Q}^N$, is defined by
 148 applying $\mathcal{T}_{W,b}^3$ repeatedly until $q \in F$ (where q is the state that the RNN simulates), and is undefined
 149 if it is not possible to have $q \in F$ by applying $\mathcal{T}_{W,b}^3$ repeatedly. This is similar to the definition of
 150 $\mathcal{P}_{\mathcal{M}}^*$. Based on this definition and Theorem 1, it follows that:

151 **Corollary 1.1.** *Given a Turing machine \mathcal{M} , there exists an injective function $\rho : \mathcal{X} \rightarrow \mathbb{Q}^n$ and an*
 152 *n -neuron unbounded-precision RNN $\mathcal{T}_{W,b} : \mathbb{Q}^n \rightarrow \mathbb{Q}^n$, where $n = 2|\Gamma| + \lceil \log_2 |Q| \rceil + |Q||\Gamma| + 5$,*
 153 *such that for all instantaneous descriptions $x \in \mathcal{X}$, the following holds: If $\mathcal{P}_{\mathcal{M}}^*(x)$ is defined, then*

$$154 \quad \rho^{-1}(\mathcal{T}_{W,b}^*(\rho(x))) = \mathcal{P}_{\mathcal{M}}^*(x), \quad (7)$$

155 *and if $\mathcal{P}_{\mathcal{M}}^*(x)$ is not defined, then $\mathcal{T}_{W,b}^*(\rho(x))$ is also not defined.*

157 *Moreover, if $\mathcal{P}_{\mathcal{M}}^*$ is defined and computed in T steps by \mathcal{M} , then $\mathcal{T}_{W,b}^*(\rho(x))$ is computed in $3T$ steps*
 158 *by the RNN.*

159 Corollary 1.1 shares some similarities with the Theorem 1 proposed in [1]. However, our theorem
 160 states that $3T$, instead of $4T$ plus a linear order of length of the tape, is sufficient to simulate a Turing
 161 machine. Our theorem also gives a relationship between the number of neurons required for the RNN
 162 and the size of Q and Γ in the Turing machine. The RNN constructed is also a direct simulation of a
 163 Turing Machine instead of a two-stack machine with a binary tape.

164 [14] proposed a Universal Turing Machine (UTM), a Turing machine with 6 states and 4 symbols
 165 denoted by $U_{6,4}$, that can simulate any Turing machines in time $\mathcal{O}(T^6)$, where T is the number of
 166 steps required for the Turing machine (the one to be simulated) to compute the result. As $U_{6,4}$ is also
 167 a Turing machine, we can thus apply Corollary 1.1 to simulate $U_{6,4}$, leading to a Turing-complete
 168 RNN. Plugging in $|Q| = 6$ and $|\Gamma| = 4$, we obtain the following result:

169 **Theorem 2.** *There exists a 40-neuron unbounded-precision RNN that can simulate any Turing*
 170 *machines in $\mathcal{O}(T^6)$, where T is the number of steps required for the Turing machine to compute the*
 171 *result.*

172 Note that [1] proved that there exists a Turing-complete RNN with 1058 neurons, while [15] proposed
 173 a Turing-complete RNN with 52 neurons. However, [15] does not have a rigorous proof, and some
 174 key formulas seem to be missing. As far as we are aware, our proposed 40-neuron RNN is the
 175 smallest Turing-complete RNN to date.

176 4 Turing Completeness of Bounded-Precision RNNs with Growing 177 Memories

178 In the following, we consider how to remove the assumption of unbounded precision that is required
 179 in all the results in Section 3. If we assume all neurons to have precision bounded by p in base $2|\Gamma|$,

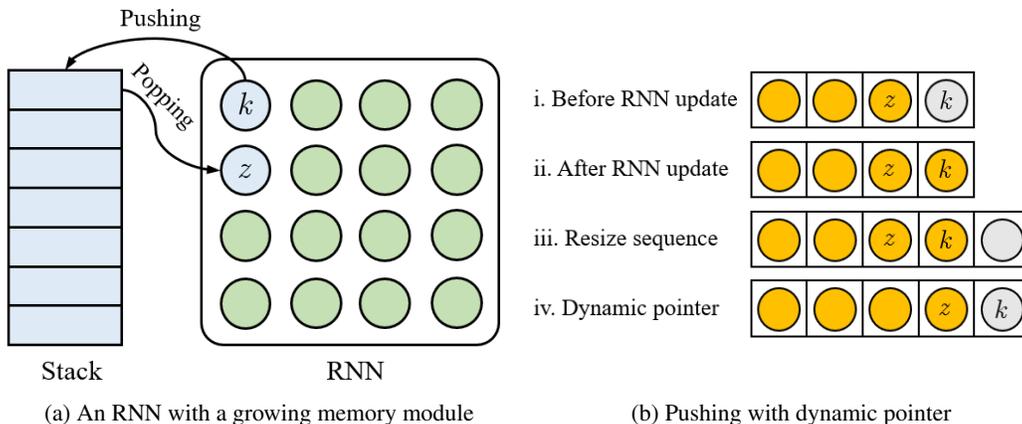


Figure 1: (a) An RNN with a growing memory module. The RNN controls the pushing and popping of the stack by two neurons $k(t)$ and $z(t)$. (b) Illustration of pushing in growing memory module as a dynamic pointer. Orange (dark) and gray (light) circles represent non-zero neurons and zero neurons respectively. i. Before RNN update, $k(t-1) = 0$. ii. RNN set $k(t) = c$ where $c > 0$ is the value to be pushed. iii. One zero neuron is appended to the sequence since there are no zero neurons left. iv. Both z and k point to different neurons in the sequence, ready to be read by other neurons in the next update of RNN.

180 then the tape can be encoded by $\lceil |t_m|/p \rceil$ neurons using fractal encoding, where $m \in \{l, r\}$. We
 181 do this by encoding every p symbols in the tape with a single neuron. However, notice that most of
 182 these neurons do not require updating when simulating a single step of the Turing machine. This
 183 is similar to how most parts of the tape are not updated during one step of a Turing machine. It is
 184 thus not efficient to include all these neurons (which may be large in number) in the RNN. Instead,
 185 we propose a *growing memory module* that can be equipped to an RNN to manipulate neurons in a
 186 stack-like mechanism:

187 **Definition 3.** A growing memory module is a stack of non-zero neurons with push and pop operations
 188 controlled by two neurons in an RNN, denoted as push neuron $k(t)$ and pop neuron $z(t)$, in the
 189 following way: for every step t after the RNN finished updating the values of neurons, (i) if $k(t) > 0$,
 190 then a new neuron of value $k(t)$ is pushed to the stack and $k(t)$ is set to 0; (ii) if $z(t) = 0$ and the
 191 stack is not empty, then the top neuron is popped from the stack and $z(t)$ is set to the value of the top
 192 neuron in the updated stack; (iii) if $z(t) = 0$ and the stack is empty, then $z(t)$ is set to a default value
 193 c .

194 With this definition, an RNN with a growing memory module is a mapping $\mathcal{T}_{W,b} : (\mathbb{Q}^N, \mathbb{Q}^*) \rightarrow$
 195 $(\mathbb{Q}^N, \mathbb{Q}^*)$ where the first element of the tuple corresponds to values of neurons in the RNN and the
 196 second element of the tuple corresponds to the stack of the growing memory module. We can also
 197 equip an RNN with multiple growing memory modules, with each module having its own push and
 198 pop neurons controlled by the RNN. For example, an RNN with two growing memory modules
 199 defines a mapping $\mathcal{T}_{W,b} : (\mathbb{Q}^N, \mathbb{Q}^*, \mathbb{Q}^*) \rightarrow (\mathbb{Q}^N, \mathbb{Q}^*, \mathbb{Q}^*)$.

200 We can also view the growing memory module as a way of dynamically pointing to a sequence of
 201 neurons - consider a sequence of non-zero neurons appended by a zero neuron. $z(t)$ can be viewed
 202 as the pointer for the last non-zero neuron in the sequence, and $k(t)$ can be viewed as the pointer
 203 for the zero neuron in the sequence. Also, let m be the number of zero neurons in the sequence. If
 204 $m = 0$, then we add one zero neuron at the end of the sequence. If $m > 1$, then we remove $m - 1$
 205 zero neurons from the end of sequence. In this way, $k(t)$ is always defined. An illustration of this is
 206 shown in Figure 1.

207 The advantage of the proposed growing memory module is that it can dynamically recruit new
 208 neurons when more memories are needed, and releases them when memories become irrelevant. We
 209 therefore do not need to estimate the number of neurons required before simulating a Turing machine.
 210 It also saves computational cost significantly since neurons in the stack (except the top neuron) are
 211 not updated.

212 Based on the growing memory module, we can construct an RNN with bounded-precision neurons
 213 that can simulate any Turing Machines. We require two growing memory modules in the RNN: one
 214 for the left tape and one for the right tape. Similar to the previous section, we first describe how we
 215 can encode the instantaneous description $(q, t_l, t_r) \in \mathcal{X}$ by a vector of rational numbers and two
 216 stacks, with which an RNN and its growing memory modules can be initialized. In the following
 217 discussion, we assume that all neurons have precision bounded by $p \geq 2$ in base $2|\Gamma|$; that is, each
 218 neuron can encode p symbols at most.

219 Both the state q and the top symbols $t_{l,1}, t_{r,1}$ are encoded with binary values using the same method
 220 as in Section 3. For the tape t_m ($m \in \{l, r\}$), we define the fractal encoding $\rho^t : \Gamma^* \rightarrow \mathbb{Q}$ by:

$$\rho^t(t) := \sum_{i=1}^{|t|} \frac{2t_i}{(2|\Gamma|)^i}, \quad (8)$$

221 which is the same as (3) but with the encoding for infinite blank symbols removed. Then, we encode
 222 the tape t_m into a stack of neurons as follows: First, encode the rightmost p symbols of it with ρ^t and
 223 push it into an empty stack, denoted as M_m . Then, encode the next rightmost p symbols and push it
 224 to M_m again, and repeat until there are no more than p symbols left in the tape. Denote this encoding
 225 function for the tape as $\rho^M : \Gamma^* \rightarrow \mathbb{Q}^*$. The remaining symbols in the tape, denoted as $t_{m,1:h(|t_m|)}$,
 226 where $h(y) := ((y-1) \bmod p) + 1$, will be encoded with fractal encoding ρ^t as well but appears in
 227 the initialization values of the RNN's neurons instead of the stacks.

228 The general idea is that the symbols closest to the read/write head ($t_{m,1:h(|t_m|)}$) resides in the RNN
 229 since they are the symbols that require updating, and the remaining symbols will be stored in a stack.
 230 If the number of symbols residing in the RNN reaches 0 or p , then we pop or push from the stack
 231 respectively to ensure that at least 1 and at most p symbols reside in the RNN. It is also interesting to
 232 note that the k^{th} neuron in the stack (from the top) will require at least kp steps of the Turing machine
 233 before it is changed, so the values of neurons near the bottom of the stack will not be changed for
 234 many steps.

235 Finally, we need to encode $h(|t_m|)$, which keeps track of the number of symbols residing in the RNN,
 236 such that the RNN can know when to push or pop from the stack. We use the following function:
 237 $\rho^h : \{1, 2, \dots, p\} \rightarrow \mathbb{Q}$ by:

$$\rho^h(y) = \frac{y}{p+1}. \quad (9)$$

238 Together, define the encoding function $\rho : \mathcal{X} \rightarrow (\mathbb{Q}^{2|\Gamma| + \lceil \log_2 |Q| \rceil + |Q||\Gamma| + 19}, \mathbb{Q}^*, \mathbb{Q}^*)$ by:

$$\begin{aligned} \rho(q, t_l, t_r) = & (\rho^q(q) \oplus \rho^t(t_{l,1:h(|t_l|)}) \oplus \rho^t(t_{r,1:h(|t_r|)}) \oplus \rho^t(t_{l,h(|t_l|)+1:h(|t_l|)+p}) \oplus \\ & \rho^t(t_{r,h(|t_r|)+1:h(|t_r|)+p}) \oplus \rho^r(t_{l,1}) \oplus \rho^r(t_{r,1}) \oplus \rho^h(h(|t_l|)) \oplus \rho^h(h(|t_r|))) \oplus \\ & \mathbf{0}, \rho^M(t_l), \rho^M(t_r)), \end{aligned} \quad (10)$$

239 where $\mathbf{0}$ is a zero vector of size $|Q||\Gamma| + 15$. The first element of the tuple is for initializing the
 240 neurons in the RNN, while the second and third element of the tuple is for initializing stacks of
 241 the two growing memory modules. Note that all encoded values have precision of p in the above
 242 definition. Similar to the previous section, ρ is injective and so we can define the decoder function
 243 $\rho^{-1} : \rho(\mathcal{X}) \rightarrow \mathcal{X}$.

244 With the new encoder function ρ and the growing memory module, we can prove an alternative
 245 version of Theorem 1 that only requires an RNN with bounded precision:

246 **Theorem 4.** *Given a Turing machine \mathcal{M} , there exists an injective function $\rho : \mathcal{X} \rightarrow (\mathbb{Q}^n, \mathbb{Q}^*, \mathbb{Q}^*)$
 247 and an n -neuron p -precision (in base $2|\Gamma|$) RNN with two growing memory modules $\mathcal{T}_{W,\mathbf{b}} :$
 248 $(\mathbb{Q}^n, \mathbb{Q}^*, \mathbb{Q}^*) \rightarrow (\mathbb{Q}^n, \mathbb{Q}^*, \mathbb{Q}^*)$, where $n = 2|\Gamma| + \lceil \log_2 |Q| \rceil + |Q||\Gamma| + 19$ and $p \geq 2$, such
 249 that for all instantaneous descriptions $x \in \mathcal{X}$,*

$$\rho^{-1}(\mathcal{T}_{W,\mathbf{b}}^3(\rho(x))) = \mathcal{P}_{\mathcal{M}}(x). \quad (11)$$

251 *Proof.* The detailed proof is in Appendix B. To illustrate the construction of the RNN, the parameters
 252 for the neuron initialized with $\rho^h(h(|t_l|))$, called the left-selector neuron $s_l(t)$, will be described
 253 here.

254 We assume all neurons in the RNN are initialized with $\rho(x)$ at time $t = 1$. The selector neuron $s_l(t)$
 255 encodes the number of left-tape symbols residing in the RNN. In three steps of an RNN, we need
 256 to update its value from $s_l(1) = h(|t_l|)/(p + 1)$ to $s_l(4) = h(|t'_l|)/(p + 1)$, where t'_l is the left tape
 257 after one step of the Turing machine. First, notice that $h(|t'_l|)$ can be expressed as:

$$h(|t'_l|) = \begin{cases} h(|t_l|) - 1, & \text{if } d = L \text{ and } h(|t_l|) \geq 2 \\ p, & \text{if } d = L \text{ and } h(|t_l|) = 1 \\ h(|t_l|) + 1, & \text{if } d = R \text{ and } h(|t_l|) \leq p - 1 \\ 1, & \text{if } d = R \text{ and } h(|t_l|) = p, \end{cases} \quad (12)$$

258 where d is the direction that the Turing machine is moving. For example, if the Turing machine is
 259 moving left, then there will be one less left-tape symbol residing in the RNN after one step of the
 260 Turing machine. However, if there are no left-tape symbols left in the RNN, the top neuron will be
 261 popped from the left-tape stack to the RNN, so there will be p left-tape symbols residing in the RNN
 262 instead. A similar process holds when the Turing machine is moving right.

263 Then we consider how to implement (12) with an RNN: Define stage neurons as:

$$c_1(t + 1) = \sigma(1 - c_1(t) - c_2(t)), \quad (13)$$

$$c_2(t + 1) = \sigma(c_1(t)), \quad (14)$$

264 with both neurons initialized to be zero. Define $\mathbf{c}(t) := [c_1(t), c_2(t), c_3(t)]$ where $c_3(t) := 1 -$
 265 $c_1(t) - c_2(t)$, then $\mathbf{c}(1) = [0, 0, 1]$; $\mathbf{c}(2) = [1, 0, 0]$; $\mathbf{c}(3) = [0, 1, 0]$. These stage neurons signal
 266 which one of the three steps that the RNN is in.

267 Also, in the construction of our RNN, there exists a linear sum of neurons, denoted as, $\mathbf{d}_m(t) =$
 268 $[d_l(t), d_r(t)]$, such that if the Turing machine is moving left, $\mathbf{d}(1) = [0, 0]$; $\mathbf{d}(2) = [1, 0]$; $\mathbf{d}(3) =$
 269 $[0, 0]$; and if the Turing machine is moving right, $\mathbf{d}(1) = [0, 0]$; $\mathbf{d}(2) = [0, 1]$; $\mathbf{d}(3) = [0, 0]$; this
 270 signals which direction the Turing machine is moving to (formulas of $\mathbf{d}_m(t)$ appears in Appendix B).

271 Then consider the following update rule for the left-selector neurons:

$$s_l(t + 1) = \sigma(s_l(t) + d_r(t) - d_l(t) - ps'_l(t) + ps''_l(t)), \quad (15)$$

$$s'_l(t + 1) = \sigma((p + 1)s_l(t) + d_r(t) - p - c_2(t) - c_3(t)), \quad (16)$$

$$s''_l(t + 1) = \sigma(2 - (p + 1)s_l(t) - d_r(t) - c_2(t) - c_3(t)), \quad (17)$$

272 where $s_l(1) = h(|t_l|)/(p + 1)$, $s'_l(1) = s''_l(1) = 0$. It can be verified that $s_l(4) = h(|t'_l|)/(p + 1)$ as
 273 defined by (12), completing the proof for $s_l(t)$. The proof for the remaining neurons can be found in
 274 Appendix B. \square

275 Similar to Corollary 1.1, it follows that:

276 **Corollary 4.1.** *Given a Turing machine \mathcal{M} , there exists an injective function $\rho : \mathcal{X} \rightarrow (\mathbb{Q}^n, \mathbb{Q}^*, \mathbb{Q}^*)$
 277 and an n -neuron p -precision (in base $|\Gamma|$) RNN with two growing memory modules $\mathcal{T}_{W, \mathbf{b}}^* :$
 278 $(\mathbb{Q}^n, \mathbb{Q}^*, \mathbb{Q}^*) \rightarrow (\mathbb{Q}^n, \mathbb{Q}^*, \mathbb{Q}^*)$, where $n = 2|\Gamma| + \lceil \log_2 |Q| \rceil + |Q||\Gamma| + 19$ and $p \geq 2$, such
 279 that for all instantaneous descriptions $x \in \mathcal{X}$, the following holds: If $\mathcal{P}_{\mathcal{M}}^*(x)$ is defined, then*

$$\rho^{-1}(\mathcal{T}_{W, \mathbf{b}}^*(\rho(x))) = \mathcal{P}_{\mathcal{M}}^*(x), \quad (18)$$

280

281 and if $\mathcal{P}_{\mathcal{M}}^*(x)$ is not defined, then $\mathcal{T}_{W, \mathbf{b}}^*(\rho(x))$ is also not defined.

282 Moreover, if $\mathcal{P}_{\mathcal{M}}^*(x)$ is defined and computed in T steps by \mathcal{M} , then $\mathcal{T}_{W, \mathbf{b}}^*(\rho(x))$ is computed in $3T$
 284 steps by the RNN.

285 Finally, applying Corollary 4.1 to $U_{6,4}$, we obtain:

286 **Theorem 5.** *There exists a 54-neuron p -precision (in base 8) RNN with two growing memory modules
 287 that can simulate any Turing machines in $\mathcal{O}(T^6)$, where T is the number of steps required for the
 288 Turing machine to compute the result and $p \geq 2$.*

289 The architecture of the Turing-complete 54-neuron RNN follows the one used in the proof of Theorem
 290 4, which can be illustrated as a 4-layer model as shown in Figure 2.

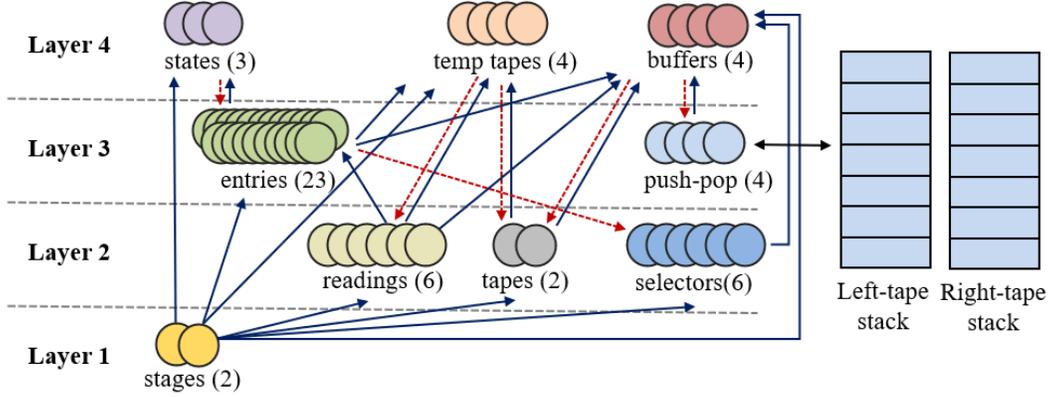


Figure 2: The architecture of the Turing-complete 54-neuron RNN with two growing memory modules. Blue (solid) lines denote bottom-up connections, and red (dotted) lines denote top-down connections. Notable neurons include: state neurons - represent the current state; tape neurons - represent the tape symbols near the read/write head; selector neurons - represent the number of symbols residing in the RNN; push and pop neurons - control the pushing and popping of left-tape and right-tape stacks; stage neurons - represent the current stage and inhibits computation of other neurons if the neurons do not require updating in a stage. A detailed description of these neurons can be found in Appendix B.

291 5 Turing Completeness of Bounded-Precision RNNs

292 As discussed previously, it is inefficient to update all neurons that encode the tape in an RNN, since
 293 most of the neurons do not require updating. However, we do not need the growing memory modules
 294 if we update all neurons at the same time. To extend the theoretical foundation of RNNs, we also
 295 construct a bounded-precision RNN to simulate a Turing machine without any growing memory
 296 modules.

297 We define a Turing machine with a bounded tape as a Turing machine that halts if the read/write head
 298 reaches the end of the tape, which is assumed to have a finite size. Then it is possible to construct an
 299 RNN with $\mathcal{O}(\lceil F/p \rceil)$ p -precision neurons to simulate a Turing machine with a bounded tape of size
 300 F :

301 **Theorem 6.** *Given a Turing machine \mathcal{M} with a bounded tape of size F , there exists an injective*
 302 *function $\rho : \mathcal{X} \rightarrow \mathbb{Q}^N$ and an n -neuron p -precision (in base $|2\Gamma|$) RNN $\mathcal{T}_{W,b} : \mathbb{Q}^n \rightarrow \mathbb{Q}^n$, where*
 303 *$n = \mathcal{O}(\lceil F/p \rceil)$ and $p \geq 2$, such that for all instantaneous descriptions $x \in \mathcal{X}$,*

$$304 \quad \rho^{-1}(\mathcal{T}_{W,b}^3(\rho(x))) = \mathcal{P}_{\mathcal{M}}(x). \quad (19)$$

305 The proof can be found in Appendix C. The general idea of the proof is to implement the growing
 306 memory module in Section 4 by an RNN as well and place all neurons inside the RNN. A corollary
 307 similar to Corollary 1.1 and 4.1 follows and is omitted here.

308 This theorem shows that there is a trade-off between the number of neurons and precision when
 309 trying to simulate a Turing machine with an RNN. It also establishes that with the same precision
 310 p , an RNN with more neurons has a larger capability since it can simulate a Turing machine with
 311 a larger tape. Taking it to the extreme, we can use an infinite-neuron bounded-precision RNN to
 312 simulate $\bar{U}_{6,4}$, which has an unbounded tape. This leads to the following theorem:

313 **Theorem 7.** *There exists an infinite-neuron bounded-precision RNN that can simulate any Turing*
 314 *machines in $\mathcal{O}(T^6)$, where T is the number of steps required for the Turing machine to compute the*
 315 *result.*

316 6 Discussion and Conclusion

317 It is inevitable that either unbounded precision (Theorem 2) or an infinite number of neurons
 318 (Theorem 7) is required to construct a Turing-complete RNN, since the tape in a Turing machine

319 has an unbounded length, and an RNN can only encode information by neurons. However, both
320 unbounded precision or infinite neurons are not practical in implementation. This highlights a major
321 limitation of RNNs. To provide a practical solution, we propose a growing memory module that
322 allows an RNN to control a stack with two neurons and prove the Turing completeness of the enhanced
323 RNN (Theorem 5).

324 All the RNNs consider in this paper have a fixed parameter. Given the input and target output, we may
325 also consider learning the parameters of RNNs with error backpropagation. However, the operation of
326 the growing memory module is not fully differentiable. Nonetheless, it may be possible to construct
327 a differentiable version of the growing memory module and learn the parameters of RNNs directly,
328 which is left for future work. The growing memory module may facilitate RNNs to learn tasks that
329 are more memory-intensive.

330 In conclusion, we prove the Turing completeness of a 40-neuron unbounded-precision RNN, which
331 is the smallest Turing-complete RNN to date. We propose a novel growing memory module that
332 can be equipped to an RNN to remove the unbounded-precision requirement. With the growing
333 memory module, a 54-neuron bounded-precision RNN can be Turing-complete. We also analyze the
334 relationship between the number of neurons and the precision of an RNN when simulating a Turing
335 machine. Finally, we prove the Turing completeness of an infinite-neuron bounded-precision RNN.

336 References

- 337 [1] Hava T Siegelmann and Eduardo D Sontag. On the computational power of neural nets. *Journal*
338 *of computer and system sciences*, 50(1):132–150, 1995.
- 339 [2] Hava T Siegelmann. Computation beyond the turing limit. *Science*, 268(5210):545–548, 1995.
- 340 [3] Hava T Siegelmann. *Neural networks and analog computation: beyond the Turing limit*.
341 Springer Science & Business Media, 2012.
- 342 [4] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. *arXiv preprint*
343 *arXiv:1410.5401*, 2014.
- 344 [5] Alex Graves, Greg Wayne, Malcolm Reynolds, Tim Harley, Ivo Danihelka, Agnieszka Grabska-
345 Barwińska, Sergio Gómez Colmenarejo, Edward Grefenstette, Tiago Ramalho, John Agapiou,
346 et al. Hybrid computing using a neural network with dynamic external memory. *Nature*,
347 538(7626):471–476, 2016.
- 348 [6] Joe Kilian and Hava T Siegelmann. The dynamic universality of sigmoidal neural networks.
349 *Information and computation*, 128(1):48–56, 1996.
- 350 [7] H Georg Kuhn, Heather Dickinson-Anson, and Fred H Gage. Neurogenesis in the dentate
351 gyrus of the adult rat: age-related decrease of neuronal progenitor proliferation. *Journal of*
352 *Neuroscience*, 16(6):2027–2033, 1996.
- 353 [8] Arturo Alvarez-Buylla and Jose Manuel Garcia-Verdugo. Neurogenesis in adult subventricular
354 zone. *Journal of Neuroscience*, 22(3):629–634, 2002.
- 355 [9] David Marr, David Willshaw, and Bruce McNaughton. Simple memory: a theory for archicortex.
356 In *From the Retina to the Neocortex*, pages 59–128. Springer, 1991.
- 357 [10] MW Jung and BL McNaughton. Spatial selectivity of unit activity in the hippocampal granular
358 layer. *Hippocampus*, 3(2):165–182, 1993.
- 359 [11] Gerd Kempermann, H Georg Kuhn, and Fred H Gage. Experience-induced neurogenesis in the
360 senescent dentate gyrus. *Journal of Neuroscience*, 18(9):3206–3212, 1998.
- 361 [12] Georg Elias Müller and Alfons Pilzecker. *Experimentelle beiträge zur lehre vom gedächtniss*,
362 volume 1. JA Barth, 1900.
- 363 [13] James L McGaugh. Memory—a century of consolidation. *Science*, 287(5451):248–251, 2000.
- 364 [14] Turlough Neary and Damien Woods. Four small universal turing machines. *Fundamenta*
365 *Informaticae*, 91(1):123–144, 2009.

366 [15] J Nicholas Hobbs and Hava Siegelmann. Implementation of universal computation via small
367 recurrent finite precision neural networks. In *2015 International Joint Conference on Neural*
368 *Networks (IJCNN)*, pages 1–5. IEEE, 2015.

369 Checklist

370 The checklist follows the references. Please read the checklist guidelines carefully for information on
371 how to answer these questions. For each question, change the default **[TODO]** to **[Yes]**, **[No]**, or
372 **[N/A]**. You are strongly encouraged to include a **justification to your answer**, either by referencing
373 the appropriate section of your paper or providing a brief inline description. For example:

- 374 • Did you include the license to the code and datasets? **[Yes]** See Section ??.
- 375 • Did you include the license to the code and datasets? **[No]** The code and the data are
376 proprietary.
- 377 • Did you include the license to the code and datasets? **[N/A]**

378 Please do not modify the questions and only use the provided macros for your answers. Note that the
379 Checklist section does not count towards the page limit. In your paper, please delete this instructions
380 block and only keep the Checklist section heading above along with the questions/answers below.

- 381 1. For all authors...
 - 382 (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s
383 contributions and scope? **[Yes]**
 - 384 (b) Did you describe the limitations of your work? **[Yes]**
 - 385 (c) Did you discuss any potential negative societal impacts of your work? **[Yes]** This is a
386 theoretical work and does not present any foreseeable societal consequences.
 - 387 (d) Have you read the ethics review guidelines and ensured that your paper conforms to
388 them? **[Yes]**
- 389 2. If you are including theoretical results...
 - 390 (a) Did you state the full set of assumptions of all theoretical results? **[Yes]**
 - 391 (b) Did you include complete proofs of all theoretical results? **[Yes]**
- 392 3. If you ran experiments...
 - 393 (a) Did you include the code, data, and instructions needed to reproduce the main experi-
394 mental results (either in the supplemental material or as a URL)? **[N/A]**
 - 395 (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they
396 were chosen)? **[N/A]**
 - 397 (c) Did you report error bars (e.g., with respect to the random seed after running experi-
398 ments multiple times)? **[N/A]**
 - 399 (d) Did you include the total amount of compute and the type of resources used (e.g., type
400 of GPUs, internal cluster, or cloud provider)? **[N/A]**
- 401 4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - 402 (a) If your work uses existing assets, did you cite the creators? **[N/A]**
 - 403 (b) Did you mention the license of the assets? **[N/A]**
 - 404 (c) Did you include any new assets either in the supplemental material or as a URL? **[N/A]**
405
 - 406 (d) Did you discuss whether and how consent was obtained from people whose data you’re
407 using/curating? **[N/A]**
 - 408 (e) Did you discuss whether the data you are using/curating contains personally identifiable
409 information or offensive content? **[N/A]**
- 410 5. If you used crowdsourcing or conducted research with human subjects...
 - 411 (a) Did you include the full text of instructions given to participants and screenshots, if
412 applicable? **[N/A]**

413
414
415
416

- (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
- (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]