



**HAL**  
open science

# Generalization emerges from local optimization in a self-organized learning network

S. Barland, Lionel Gil

► **To cite this version:**

S. Barland, Lionel Gil. Generalization emerges from local optimization in a self-organized learning network. 2024. hal-04758774

**HAL Id: hal-04758774**

**<https://hal.science/hal-04758774v1>**

Preprint submitted on 29 Oct 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Generalization emerges from local optimization in a self-organized learning network

S. Barland, L. Gil

*Université Côte d'Azur, Institut de Physique de Nice (INPHYNI), France*

(Dated: October 11, 2024)

We design and analyze a new paradigm for building supervised learning networks, driven only by local optimization rules without relying on a global error function. Traditional neural networks with a fixed topology are made up of identical nodes and derive their expressiveness from an appropriate adjustment of connection weights. In contrast, our network stores new knowledge in the nodes accurately and instantaneously, in the form of a lookup table. Only then is some of this information structured and incorporated into the network geometry. The training error is initially zero by construction and remains so throughout the network topology transformation phase. The latter involves a small number of local topological transformations, such as splitting or merging of nodes and adding binary connections between them. The choice of operations to be carried out is only driven by optimization of expressivity at the local scale. What we're primarily looking for in a learning network is its ability to generalize, i.e. its capacity to correctly answer questions for which it has never learned the answers. We show on numerous examples of classification tasks that the networks generated by our algorithm systematically reach such a state of perfect generalization when the number of learned examples becomes sufficiently large. We report on the dynamics of the change of state and show that it is abrupt and has the distinctive characteristics of a first order phase transition, a phenomenon already observed for traditional learning networks and known as *grokking*. In addition to proposing a non-potential approach for the construction of learning networks, our algorithm makes it possible to rethink the grokking transition in a new light, under which acquisition of training data and topological structuring of data are completely decoupled phenomena.

PACS numbers: Self organized systems, Phase transitions, Physics of computation, Machine Learning Models

## I. INTRODUCTION

The computational properties of complex systems have long been a topic of fundamental interest for the physics community, from early examples of neural networks and cellular automata [1, 2] to very recent approaches of physics-based neuromorphic computing [3]. A common task for such systems is that of *supervised learning*, in which a neural network is trained to realize a particular function based on many examples of input-output couples. In this context, a surprising yet often observed property of neural networks is their capacity to provide statistically correct outputs even for unknown inputs. In spite of the ubiquity of this phenomenon known as *generalization* it is still far from being completely understood and the principles and tools of physics can certainly help in this quest [4–7]. In particular, early studies of shallow neural networks have shown that a phase transition to error-free generalization can take place for perfectly learnable rules [8–10]. This phenomenon has recently received the name of *grokking* as it was re-observed in deep neural networks [11] and subsequently analyzed as a phase transition in a shallow network [12]. Grokking, understood as a phase transition of a network towards a perfect generalization state, therefore emerges as a very general phenomenon in learning, neither limited to algorithmic data [13] nor to specific architectures or training processes [14–16]. It is however well established that the generalization property of neural networks comes at the cost of their propension to forget older inputs, a phenomenon known as catastrophic forgetting [17, 18]. Far from an implementation detail, this situation results from

a fundamental compromise to be found in the stability-sensitivity dilemma [19].

In the following, we demonstrate that a local optimization rule (as opposed to a global minimization procedure) can drive the error-free evolution of a network towards a perfect generalization state. Correspondingly, the learning dynamics shows distinctive features of a phase transition.

With this approach, we strike a new balance in the stability-sensitivity spectrum [20]: lookup tables (LUT) are perfect memories which lack any generalization capability and at the other end deep neural networks can generalize extremely complex problems but may catastrophically forget older data. Instead, our network evolves under a constraint of error free operation (thus ruling out forgetting) and dynamically transfers complexity from the nodes to the network topology, which provides generalization capability.

Our approach strongly differs from standard artificial neural networks in several ways. First, there is no global cost function such as the training losses of statistical learning. On the contrary, our network dynamics is ruled by a *local* minimization principle, as each node evolves based only on locally available information. Besides the long-questioned biological plausibility of a global minimization procedure (see *eg* [21, 22]), this is a fundamental difference since the dynamics we propose is not variational, as opposed to the usual training by gradient descent. Second, we do not update continuously tunable weights since our network evolves via the addition, merger or deletion of nodes and binary edges. Therefore, we do not fit an existing network to training data. In-

stead, the topology of our network evolves dynamically upon arrival of new data. Self-organized networks whose topology may evolve (see *eg* [23–26]) are often considered in the context of continual learning [27]. In this context, the focus is less on grokking than on adaptation to new data while avoiding catastrophic forgetting. Here instead, the network evolution involves first perfectly storing new data by adapting a node (which consist of a LUT) and then transferring the complexity from the individual nodes to the network, the evolution of each node being driven by greedy optimization.

The paper is organized as follows. In section II we describe the nature of our proposed network and the processes which drive its evolution during learning. In III we illustrate the resulting dynamics with the example of a network learning to solve the  $N$ -bit parity problem towards the final state of perfect generalization. In IV, we analyze the statistics of the generalization error with respect to the dimension of the problem and discuss it in relation with a phase transition. Finally, in section V we show that the proposed algorithm can generate learning networks capable of grokking many binary classification tasks, including the primality problem. We discuss the implications of our results in VI and conclude in VII.

## II. MODELISATION AND ALGORITHM

In deep learning algorithms, a distinction is made between the training phase, where the network is asked to reproduce the set of data provided as examples (training data samples or TDS), and the generalization phase, where the network is able to correctly predict the result on data it has not yet learned. The training threshold precedes the generalization threshold, but both are obtained by increasing the number of iterations of the same global error minimization procedure [11], wrongly suggesting that the generalization threshold is concomitant with better accuracy in the learning phase.

Here we design a learning network model in which the training phase is simplified to the extreme, as it simply disappears: Each new TDS is automatically, instantaneously and perfectly acquired on the fly. On the other hand, each new learning phase is followed by a network remodeling phase, during which the information learned is judiciously redistributed between the nodes and the information flow restructured by the topology of the connections. All the organization rules are local: local in time, as they act after each new TDS, and local in space, as they involve a single node or pair of nodes.

Our algorithm calls on multiple interdependent concepts, whose respective roles can only be understood through the input and output of the other concepts. A sequential presentation is not well adapted to this recurrent structure and we have therefore decided to present our algorithm in two complementary, if slightly redundant, stages. First, we give a synthetic description: we briefly introduce the various basic concepts (what the network

is made of, the dynamics of information on this network, the network’s own dynamics with the notions of splitting and merging, learning a new rule, conflict resolution and cleaning) and we describe the syntactic structure of the algorithm i.e. how these basic concepts relate to each other. The second stage takes up each of the concepts introduced in the first part and describes them in detail. Particular emphasis is placed on the node splitting mechanism and how it underpins the local generalization process at the heart of our algorithm. This second step may be omitted by the impatient reader on first reading.




FIG. 1. Network before (a) and after (b) topological evolution for the  $N$ -bits parity problem with  $N = 3$ . There are 4 interface nodes (3 input, 1 output), 1 hidden node for (a) and 2 for (b). There are 5 TDS, each associated with a line of the various tables involved. The notation  $i_A^p$  is used to designate the  $p$  th input connection to  $A$ ,  $o_A^q$  for its  $q$  th outgoing connection. As an example of how to read figure (a), we consider the case of the first (resp. 2nd) training sample: the top left node sends 0 (resp. 1), the middle left node 0 (resp. 0) and the bottom left node 1 (resp. 0). The central node then receives the code 0, 0, 1 (resp. 1, 0, 0). According to the node’s LUT, the corresponding output code is 1 (resp. 1), which is retrieved by the last node. Figure (b) reads the same way, from left to right, layer by layer. It is important to note that networks (a) and (b) are strictly equivalent for all 5 training samples, but that network (b) is able to provide a correct answer to all the  $2^N$  cases, including the three which are not part of the training samples (cf. text).

### A. Synthetic description

#### 1. What the network is made of

The network is a directed graph made up of nodes and links. The nodes are arranged on layers. The  $N_e$  nodes on the first layer and the single node on the last layer have a special role: The former (left side) have only one outgoing connection, the latter (right side) only one incoming connection. We’ll call them interface nodes. The

$N$  nodes between these two layers form the hidden nodes. During the network dynamics phase, the number, connections and behavior of the hidden nodes change. Interface nodes, on the other hand, are immutable in the sense that network dynamics preserve the uniqueness of outgoing and incoming connections respectively. We'll call last hidden node and note (LHN), the only hidden node connected to the output interface node. The similarity with neural systems ends here. The dynamics of a true numerical neuron (and, to a lesser extent, a biological neuron) are quite straightforward. A single action potential is calculated, and the same response is sent to all output links. Action potential and response are linked by a  $\mathbb{R} \rightarrow \mathbb{R}$  activation function. All the informations associated with the training data are stored in the network in the form of connection weights. In contrast, in our network, the informations associated with the TDS are stored partly in the network in the form of the topology of connections, but also partly in each node ( $i$ ) in the form of a multi-dimensional response function  $\{0, 1\}^{nil(i)} \rightarrow \{0, 1\}^{nol(i)}$  specific to each node, where  $nil(i)$  denotes its number of incoming links and  $nol(i)$  its number of outgoing links.

## 2. Dynamics on the network

Fig.1 deals with an elementary well known machine learning playground consisting in determining the odd or even nature of the sum of the  $Ne = 3$  components of a binary vector. The figure shows the situation after 5 TDS (hence the stack of 5 horizontal lines). The boxes associated with the input interface nodes correspond to the output codes of these nodes during each of the 5 TDS. The box for the output interface node lists the input codes for this node for all 5 TDS.

Each node in the network behaves like a LUT, matching any input code with a single output code. For example the node  $A$  in fig.1b has 2 incoming links  $i_A^1$  and  $i_A^2$  from nodes (1, 3) on its left as inputs, and 1 outgoing link  $o_A^1 = i_B^1$  to nodes ( $B$ ) on its right ( $nol$  can be greater than 1, but it is not the case here due to the extreme simplicity of the elementary problem under consideration). Any input code of  $A$  consists of an ordered list of bits (0 or 1), one bit for each input connection. The  $A$  node recognizes the current input code as one of the input codes in its LUT and transmits through its output connections the information bits of the output code corresponding to the identified input code. For example for the 1st TDS (resp. 2nd), the current input code is  $[0, 1]$  (resp.  $[1, 0]$ ), it is identified as the learning 1 (resp. 2) in the LUT and then  $A$  transmits the code  $[0]$  (resp.  $0$ ) to  $B$ . In fig.1b, when node  $A$  have transmitted its output code, then node  $B$  is in possession of the complete input code  $[0, 0]$  (resp.  $[0, 0]$ ) and can then transmit the output code  $[1]$  (resp.  $[1]$ ) corresponding to its own LUT. Information thus flows from left to right, and all the nodes form a feed-forward network.

A TDS for the network consists of a set of output codes

for the interface nodes. A learning for a hidden node consists of a pair of input and output codes for its connections. Despite the similarities between the two concepts, we'll use 2 distinct words to emphasize the difference in scale, but above all in data origin. The network receives unchanging data from outside, while the data accessible to a node has been processed by upstream nodes. The tricky part is mapping each network TDS to the multiple hidden nodes learning (going from (a) to (b) in fig.1), but later (in IIA 3) we'll discuss an algorithm that does this without loss of information. But now, the result is that each hidden node ends up storing information in the form [ learning number, input code, output code] . The length of the node LUT therefore increases by 1 unit with each new network TDS, and the same pair [entry code, output code] may be repeated several times. As a result, the learning process is straightforward: there are no connection weights to adjust, there is no non-linear thresholding function, there is no global minimization of error. All that is needed is to add a new item to the LUTs such that every new learning is exactly and immediately acquired.

## 3. Network's own dynamics

Once the new information has been rapidly acquired, the network topology may change through mechanisms, splitting and merging, that are both local ( i.e. they only involve processes associated with one node or a pair of nodes) and conservative (the training error is exactly zero throughout the processes).

Splitting and merging are illustrated in fig.1. Splitting is the conservative process whereby a hidden node  $C$  splits into 2 nodes  $A$  and  $B$ , sharing incoming and outgoing links and possibly creating links between  $A$  and  $B$ . Merging is the opposite conservative process, in which the two hidden nodes  $A$  and  $B$  merge to form a single node  $C$ . Both elementary processes can be chained together at will to profoundly modify the network topology, again with no effect on the training error. We discuss later on how these mechanisms can lead to a much more complex topology for more complex problems (see fig.5)

Splitting and merging are the basic antagonistic tools of topological transformations. When and how should they be used? Splitting is associated with data analysis and is responsible for proposing a coherent syntactic structure through inductive reasoning. It enables the network to propose a coherent (but not necessarily correct) interpretation of the local data. Merging occurs when a new TDS to be learned turns out to conflict with the previously coherent interpretation of the data. It allows us to backtrack.

The splitting action replaces the central node  $C$  in (a) with 2 hidden nodes  $A$  and  $B$  in (b), positioned on 2 distinct layers (fig.1). Local network dynamics are obviously modified, but the reader is strongly encouraged to check that (b) faithfully reproduce the 5 TDS described

in (a). Absolutely no training information was lost during the splitting operation. Now it's crucial to note that since the central node  $C$  has 3 incoming links, its LUT will only be complete after  $2^3$  independent learnings. Indeed, node  $C$  in fig.1a doesn't know any of the codes  $[0, 1, 0]$ ,  $[0, 1, 1]$  and  $[1, 0, 1]$ . Not only has it never encountered them before, but it's also unable of coming up with a plausible answer. On the contrary, we can check point by point that nodes  $A$  and  $B$  in the new topology (fig.1b) are really able to propose an answer. We emphasize that the proposed answer could have been wrong. Instead, the network provides here a correct answer for all the  $2^3$  possible codes, including the 3 unknown ones.

Finally, note that splitting can be chained. However, in this example, splitting one or both hidden nodes doesn't add anything.

#### 4. New training sample and conflict

Consider a network with  $N_e$  input interface nodes, which has already been taught  $m$  TDS. When we measure a non-zero generalization error, we deduce that we must continue to feed the network with new TDS. We therefore supply the left interface nodes with the output codes corresponding to the new training sample to be learned, and observe their evolution according to the dynamics described in II A 2. The input code seen by the nodes during this propagation is referred to as "current".

For values of  $m$  smaller than  $2^{N_e}$ , it often happens that one of the hidden nodes receives a current input code that is unknown to it, i.e. that doesn't correspond to any of the entries in its LUT. This current input code is then simply learned by integration into the LUT. In the case of LHN, the output code to be learned is the output code of the new network TDS. In the case of a hidden node other than LHN, an additional connection is created to each outgoing node to inform them of the presence of an unknown code. It transmits a bit 0 for old codes and 1 for the current input code, the other output bits remain unchanged in the first case, or are completed by analogy with the nearest known example in the second case. This procedure is then chained, from the first hidden node to receive an unknown code to all its downstream nodes.

For larger values of  $m$ , it may happen that the new TDS input propagates until the network output, with all current input codes known, but the current output code of LHN does not match the network TDS output to be learned. This causes a conflict. By construction, the conflict is not between the new network TDS and the previous  $m$  network TDS already learned, but rather between the new TDS and the coherent interpretation of the structure of the previous  $m$  achieved by successive splitting. This is where the 2-nodes merging procedure comes into play. Thanks to merging, we then progressively unravel the splitting procedure, starting with the ascending nodes closest to the conflicting one. The merging procedure continues until the new current entry code,

now fully fleshed out, becomes unknown.

#### 5. Cleaning

A node may have incoming connections that carry no relevant information. These redundant connections are useless, but they do add to the computational burden. Each node must therefore detect and eliminate them. Also this pruning may lead to the creation of hidden nodes with no outgoing connections. When these nodes are not connected to the interface nodes, they must self-destruct.

#### 6. Algorithm

Learning the network up to the transition to grokking is then a dynamic process involving the various concepts previously introduced in a recurrent loop :

1. We measure the generalization error  $E_g$  of a network that has learned  $m$  TDS. If this error is zero, nothing more needs to be done. Otherwise, an additional training sample  $m \rightarrow m + 1$  must be learned. Any conflicts are resolved as described in II A 4. The number of nodes and layers can be reduced during this phase.
2. Then, as long as this leads to a local improvement in generalization, a splitting procedure is applied to one of the network's hidden nodes taken at random. This phase increases the number of nodes and layers.
3. Clean up the network.
4. Return to point 1.

In section V we'll give a number of examples of successful applications of this algorithm. The problems addressed and their algorithmic complexity are very varied.

### B. Detailed Description

#### 1. 2-nodes merging

The mechanism for merging two nodes is fairly intuitive and easy to understand. It is illustrated in fig.2 for the topology and tab.I for the LUTs. In fig.2 top, before merging, node A has input connections  $i_A^1, i_A^2$  and  $i_A^3$  and output connections  $o_A^1, o_A^2$  and  $o_A^3$ . B has inputs connections  $i_B^p, p \in [1, 3]$  and outputs connections  $o_B^q, q \in [1, 2]$ . After merging (fig.2 bottom), the resulting node  $(A, B) \rightarrow C$  has inputs codes constructed simply by concatenating the input codes of the initial nodes A and B, as shown in tab.I. The same mechanism applies to output codes.

Two remarks are now in order. First in the example in fig.2 and tab.I, nodes  $A$  and  $B$  are not directly connected. But the merging procedure can easily be generalized to cases where  $A$  and  $B$  are connected, as in the example in fig.3 and tab.II (provided you start from fig.3b and go to fig.3a). In this case, the internal links between  $A$  and  $B$  disappear, and the information bits associated with them are not taken into account when the codes are concatenated. Second, it is important to realize that not all node pairs can be merged. This is the case, for example, with nodes  $B$  and  $E$  in fig.2. In this configuration, node  $D$  is both an output of  $B$  and an input of  $E$ , and therefore should appear as both an input and an output of their merger. This would be possible only in a recurrent network and we choose to restrict our analysis to a feedforward network.




FIG. 2. Network configuration before (top) and after (bottom) the merging  $(A, B) \rightarrow C$ . For node  $A$ , the left-hand  $p$  and right-hand  $q$  numbers overlapping the connections, label the incoming and outgoing links noted  $i_A^p$  and respectively  $o_A^q$  in the text. For the sake of clarity, we have used here a simplified notation. It is also important to note that all the pairs of nodes can not be grouped. This is the case, for example, of nodes  $B$  and  $E$  in the top figure.

## 2. 2-nodes splitting

The 2-nodes splitting procedure is illustrated in fig.3 for the topology and tab.II for the LUTs: it's an operation that allows you to switch from configuration (a) with a single node ( $C$ ) to configuration (b) with 2 nodes ( $A$  and  $B$ ), while maintaining the training error at exactly zero (tab.II). The transformation leaves the inputs and outputs (connections and nodes) invariant, such that it can be performed locally without any major changes to the surrounding nodes. During the procedure, every incoming connection to  $C$  becomes an incoming connection to  $A$  or (exclusive)  $B$ . The same applies to outgoing connections. Possible additional connections from  $A$  to  $B$  may appear, giving nodes  $A$  and  $B$  different roles. Finally, the operation can be completely cancelled by applying the 2-nodes merging procedure.

Learning: A	Input codes	Output codes
1	1 1 1	0 0 1
2	0 0 1	0 1 0
3	1 0 0	0 0 0
4	0 1 0	1 0 0
5	0 0 0	1 1 0
Learning: B	Input codes	Output codes
1	1 1 1	1 0
2	0 0 0	1 0
3	0 1 0	0 1
4	1 0 0	1 0
5	0 1 0	0 1
Learning: C	Input codes	Output codes
1	1 1 1 1 1 1	0 0 1 1 0
2	0 0 1 0 0 0	0 1 0 1 0
3	1 0 0 0 1 0	0 0 0 0 1
4	0 1 0 1 0 0	1 0 0 1 0
5	0 0 0 0 1 0	1 1 0 0 1

TABLE I. Input and output codes associated with nodes  $A$ ,  $B$  and  $C$  in fig.2.

Learning: C	Input codes	Output codes
1	1 0 0 1 0 0	1 1 0 0
2	0 1 1 0 1 0	1 0 0 1
3	0 1 0 1 1 1	0 0 0 1
4	1 1 1 1 0 0	1 1 1 0
5	0 0 1 0 1 0	0 1 0 0
Learning: A	Input codes	Output codes
1	1 0 0 1	1 1 0 <b>0</b>
2	0 1 1 0	1 0 1 <b>0</b>
3	0 1 0 1	0 0 1 <b>0</b>
4	1 1 1 1	1 1 0 <b>1</b>
5	0 0 1 0	0 1 0 <b>0</b>
Learning: B	Input codes	Output codes
1	<b>0</b> 0 0	0
2	<b>0</b> 1 0	0
3	<b>0</b> 1 1	0
4	<b>1</b> 0 0	1
5	<b>0</b> 1 0	0

TABLE II. Input and output codes associated with nodes  $A$ ,  $B$  and  $C$  in fig.3. The bits on a blue background are associated with the additional connection between  $A$  and  $B$ . They disappear in the  $C$  merger codes.

The advantage of 2-nodes splitting lies in the fact that there are multiple ways of dividing the incoming (resp. outgoing) connections of  $C$  into 2 disjoint sets, and that we can therefore make a "judicious" choice among




FIG. 3. Schematic representation of the splitting procedure  $C \rightarrow (A, B)$ . The diagram (a) shows the initial node  $C$  surrounded by its 6 inputs ( $i_C^p, p \in [1, 6]$ ) and 4 outputs ( $o_C^q, q \in [1, 4]$ ) connections. Plot (b) display the configuration after the splitting procedure. Note the extra link between  $A$  and  $B$  on the blue background, labeled  $i_B^1$  or  $o_A^4$  depending on whether it's considered as an ingoing or outcoming connection. This link carries the extra bits shown on the blue background in tab.II

a)			{Input codes: A}				
		$i_A^1$	1	0	0	1	0
		$i_A^2$	0	1	1	1	0
		$i_A^3$	0	1	0	1	1
		$i_A^4$	1	0	1	1	0
		$i_B^2$	$i_B^3$				
	{Input codes: B}	0	0	1			4
		1	0		2		5
		1	1			3	

b)				{Input codes: A}				
		$i_A^1$	1	0	0	1	0	
		$i_A^2$	0	1	1	1	0	
		$i_A^3$	0	1	0	1	1	
		$i_A^4$	1	0	1	1	0	
		$i_B^1$	$i_B^2$	$i_B^3$				
	{Input codes: B}	0	0	0	1	8	10	12
		1	0	0			4	
		0	1	0	6	2	11	? 5
		0	1	1	7	9	3	13

FIG. 4. Table of input codes for nodes  $A$  and  $B$  after  $m = 5$  training data learnings. (a): Numbers from 1 to 5 correspond to the projection on  $A$  and  $B$  of the  $C$  input codes for the 5 training data. Because of 1 and 4 (in red on (a)), the LUT of  $B$  does not define an application. We correct this by introducing an additional link  $o_A^4 = i_B^1$  (elements involved on a gray background in (b)). Numbers from 6 to 13 then correspond to consistent (although perhaps incorrect) responses from the pair  $(A, B)$  for input codes unknown to  $C$ .

the available configurations  $(A, B)$  without incurring any

change in the training error.

Let's start by identifying the possible configurations. Note that in the configuration fig.3a, the output codes of node  $A$  do not depend at all on the responses of node  $B$ . This means that  $A$ 's LUT is an application in the mathematical sense of the term, i.e. each input code of  $A$  has one and only one output image. This imposes a very strong constraint on the choice of input and output connections for  $A$ . For example, let's suppose we've selected

for  $A$  the input connection  $i_A^1$  and  $i_A^2$ , and the output connection  $o_A^1$ . The LUT for  $A$  would then be given by tab.III. Clearly, the same input code  $[0, 1]$  would then have been associated with two distinct output codes, depending on whether we had considered the 2 or 3 learnings. As this configuration was not an application, it had to be rejected.

Once we've identified a set of valid inputs and outputs

Learning: A	Input codes	Output codes
1	1 0	1
2	0 1	1
3	0 1	0
4	1 1	1
5	0 0	0

TABLE III. Example of an impossible splitting choice. LUT of  $A$  for the configuration in fig.3, when  $A$ 's entries are limited to  $i_A^1$  and  $i_A^2$ , while  $A$ 's exit to  $o_A^1$ .

connections for  $A$ , we can naturally consider a first definition of inputs and outputs for  $B$ : those for  $C$  minus those for  $A$ . At this point, the various LUTs are then described by tab.II, deprived of the 2 columns on a blue background and located respectively in the “Learning:  $A$ , Output codes, last column” and “Learning:  $B$ , input codes, first column” areas. But it turns out that the LUT of  $B$  thus constructed is not an application: the same input code  $[0,0]$  appears 2 times, for learning 1 and 4, but associated with distinct output codes. As the LUT of  $C$  is an application, it should be enough to retrieve some of the information stored in  $A$  to be able to distinguish the 2 input codes of  $B$ . This is the role assigned to the additional connection between  $A$  and  $B$  (blue background columns in tab.II). It transmits 0 for all learnings except 4, for which it transmits 1, thus lifting the degeneracy.

We’ve just seen how to perform the 2-nodes splitting procedure. Now we come to the fundamental question: what’s in it for us? As the gain is not to be found in the training error, since splitting does not modify it in any way, we focus on the generalization error. To do this, we’ll compare the responses of the pair  $(A, B)$  and those of  $C$  to the presentation of an unknown current input code. We begin by constructing the table in fig.4(a) of cross-referenced input codes for  $A$  and  $B$ . Until now, we’ve been working with lists of length  $m$  of input codes, where  $m$  is the number of training data. Here, we use the notation {input codes of  $A$ } to refer to the set of input codes of  $A$  (idem for  $B$ ). The difference is that in a set, unlike a list, an element is not allowed to repeat itself. The  $m = 5$  input codes of  $C$  break down into the input codes of  $A$  and  $B$  as described in tab.II and correspond to the intersections in fig.4a marked by the digits 1 to  $m$ . It is important to note that the  $[0, 0]$  and  $[1, 0]$  entry codes of  $B$  appear 2 times, respectively for learnings 1, 4 and 2, 5. For the 2 and 5 learnings, the  $B$  output codes associated with  $[1, 0]$  are identical (tab.II) and therefore present no difficulties. On the other hand, learnings 1 and 4 are problematic because the output codes associated with  $[0, 0]$  are not identical, and the LUT of  $B$  defined by fig.4a is not an application. The solution, as described in tab.II, is to create an additional link from  $A$  to  $B$  to remove the degeneracy. In fig.4b, this corresponds to the column  $i_B^1$ . Now consider the current code  $[1, 0, 0, 1, 1, 0]$ , unknown to  $C$ . Decomposed on the pair  $(A, B)$ , it corresponds to the

input code  $[1, 0, 0, 1]$  of  $A$  and  $[i_B^2, i_B^3] = [1, 0]$  identified by number 6 in fig.4b. The input code  $[1, 0, 0, 1]$  of  $A$  is associated with the output code  $[1, 1, 0, 0]$  (tab.II). Since  $o_A^4 = 0 = i_B^1$ , we know how to give a value to the output code of  $B$  associated with the input code  $[0, 1, 0]$ , the one associated with  $m = 2$  (i.e.  $[0]$ ). On the other hand, consider the current code  $[1, 1, 1, 1, 1, 0]$  unknown to  $C$ . Decomposed on the pair  $(A, B)$ , it corresponds to the input code  $[1, 1, 1, 1]$  of  $A$  and  $[i_B^2, i_B^3] = [1, 0]$  marked with “?” in table (b). For this input code of  $A$ , the answer is  $[1, 1, 0, 1]$  (tab.II) and so we have  $o_A^4 = 1 = i_B^1$ . But the current code  $[1, 1, 0]$  thus constructed is not an input code of  $B$ . So we don’t know how to assign it a value.

We’ve just shown that the pair  $(A, B)$  is sometimes able to provide coherent answers, whereas the initial node  $C$  is unable to do so. These answers are not necessarily accurate in the sense that a future network TDS may disavow them, but they have the merit of existing and proposing a coherent generalization.

In fact, in the particular case of fig.3, the pair  $(A, B)$  can provide a consistent answer in 13 input cases, 5 of which are exact (fig.4). As this number will play an important role in our algorithm, we’ll call it the number of consistent responses of the  $(A, B)$  pair and will denote it  $\text{NCR}(A, B) = 13$  (note that  $\text{NCR}(C) = 5$ ). For the same initial node  $C$ ,  $\text{NCR}(A, B)$  depends on the inputs and outputs repartition between  $A$  and  $B$ . With the idea of maximizing generalization, the algorithm tests several possible repartitions between  $A$  and  $B$  and selects the one that gives the highest value of  $\text{NCR}(A, B)$ . As the number of possible repartitions grows extremely quickly with the number of connections, we won’t perform an exhaustive exploration, but rather a random sampling.

Finally 2-nodes splitting operations can be chained together in a cascade process:

$$\begin{aligned}
 A_1 &\longrightarrow (A_{11}, A_{12}) \\
 A_{11} &\longrightarrow (A_{111}, A_{112}) \\
 A_{12} &\longrightarrow (A_{121}, A_{122}) \quad (1) \\
 \\ 
 A_{111} &\longrightarrow (A_{1111}, A_{1112}) \\
 \dots &
 \end{aligned}$$

which goes on as long as  $\text{NCR}(A, B) > \text{NCR}(C)$ , i.e. as long as splitting improves generalization.

### 3. New learning and unknown codes

Let’s suppose the network is in the configuration shown in fig.5 and has learned a number  $m = 5$  of TDS. We’re now trying to teach it a new TDS. The nodes of the first layer grab the bits of the new input code and pass them on to the nodes of the second layer. Each node of the second layer then receives an input code, its current






FIG. 5. A typical network configuration obtained after network topology evolution (a). The insert (b) corresponds to the initial configuration. Node A: The dotted red arrows represent the additional connections created by node A on receiving an unknown CIC. Nodes B, C and D: examples of configurations to be removed during the cleaning operation.

input code, which it may or may not already know. For the sake of argument, let's assume that all nodes in the second layer know their current input code. We then move on to the third layer. In this layer, let's assume that only node A receives a current input code (CIC(A)) that it doesn't know. The situation is summarized in tab.IV.

A little nomenclature is useful. We have already introduced  $noc(i)$  and  $nin(i)$  the number of outgoing and incoming connections of node  $i$ . Let  $non(i)$  and  $nin(i)$  be the number of outgoing and incoming nodes for node  $i$ . We have  $non(i) \leq noc(i)$  and  $nin(i) \leq nic(i)$ .

In order to deal with the CIC(A), node A and its outputs are then going to perform 3 operations:

1. Node A will add an extra connection for each of its outgoing nodes and inform them that they need to increment their  $nin$  by one. In our example (fig.5a):

$$\begin{cases} non(A) \longrightarrow non(A) \\ noc(A) \longrightarrow noc(A) + non(A) \end{cases}$$

$$\begin{cases} nin(4) \longrightarrow nin(4) \\ nic(4) \longrightarrow nic(4) + 1 \end{cases} \quad (2)$$

$$\begin{cases} nin(5) \longrightarrow nin(5) \\ nic(5) \longrightarrow nic(5) + 1 \end{cases}$$

2. Node A will modify its LUT: all its old output codes are completed by a length  $non(A)$  list of bits 0. A's output nodes modify their old input codes by adding 0 (tab.IV).
3. Node A constructs a new current output code (COC(A)): the first part of COC(A) is an approximation obtained by i) identifying the known input code of A closest to CIC(A) ii) then matching

its image through the LUT. The second part is a length  $non(A)$  list of bit 1 (tab.IV).

Tab.IV summarize the changes made to the LUTs, while the dotted red lines in fig.5 summarize the changes made to the network. We can check that the transformations made leave the old LUTs invariant, while taking into account the new information.

Learning: A	Input codes	Output codes
1	1 0 0	1 0 0 0
2	0 1 1	0 1 0 0
3	0 1 0	1 1 0 0
4	1 1 1	0 0 0 0
5	0 0 1	1 1 0 0
Current	Current input code	. . 1 1
Learning: 4	Input codes	Output codes
1	. 1 . 0	. . .
2	. 0 . 0	. . .
3	. 1 . 0	. . .
4	. 0 . 0	. . .
5	. 1 . 0	. . .
Current input code	. . . 1	? ? ?
Learning: 5	Input codes	Output codes
1	. 0 . 0	. . .
2	. 1 . 0	. . .
3	. 1 . 0	. . .
4	. 0 . 0	. . .
5	. 1 . 0	. . .
Current input code	. . . 1	? ? ?

TABLE IV. LUTs of nodes A,4 and 5 before and after adding the additional links described by dotted red lines in fig.5. The corresponding additional bits are on gray background, light for 4 and dark for 5. Dots (.) denote known values whose value does not need to be specified to understand the presentation. "???" stand for unknown ones.

Now it's the turn of the nodes in the fourth layer. By construction, nodes 4 and 5 of this layer receive a CIC that is unknown to them (uncertainty transmitted to them by A). We then repeat on these nodes the treatment just applied to A. We then pass from one layer to the next, propagating the lack of knowledge. If we reach LHN with an unknown CIC, then we just have to assign it the output code required by the current network learning operation.

Note that a node that has already been informed by one of its inputs that it is about to receive an unknown CIC, does not need to be informed again by another of its inputs. Hence node 7 on the fifth layer doesn't need to be informed by both 4 and 5 that its CIC is unknown.

#### 4. Conflicts

During the learning phase, and especially when the number of TDS starts to grow, it can happen that the LHN receives a CIC that it identifies as known, without the COC associated with it being equal to what the network should be learning. This is a conflict situation which clearly indicates that the up to now generalization choices made during the 2-nodes splittings and the consequent network topology are no longer compatible with the last TDS learned. This conflict is resolved by a cascade of 2-nodes merging propagating from the last to the first layer. The idea is to make as few changes as possible and therefore focus first on the latest generalizations. We start by merging LHN with one of the nodes of the previous layers. Then the current learning procedure is repeated until the CIC of the last node is computed. Either the conflict has disappeared, and the CICs and COCs can be saved in the LUT of each network node, or it hasn't, and the last node's merging procedure is repeated. Since the application of the 2-nodes merging procedure, repeated a sufficient number of times, transforms the topology of fig.5a into that of its insert fig.5b, and since in the latter configuration there is no possible conflict, since a new TDS process for the network boils down to simply increasing the LUT of the central node, we can be sure that the conflict resolution procedure described above converges.

#### 5. Cleaning

Learning, even when limited to a small number of input interface nodes ( $N_e$ ), is a computationally intensive process. This is particularly true of the current version of the algorithm, which is not parallelized. We therefore take great care to remove all local configurations which are clearly irrelevant and unnecessarily burden calculations. In addition to deleting redundant connections and hidden nodes (already discussed in II A 5), we can also delete gateway nodes. In fig.5a, node *C* only acts as a gateway between nodes 10 and 11, and node *D* between 1, 2 and *E*. They can be replaced by direct connections between 10 and 11 (respectively between 1, 2 and *E*) and an adequate redefinition of LUTs through composition.

#### 6. Generalisation test and no-fail mode

Recall that a network TDS is characterized by the output codes of the interface nodes of the first layer and the corresponding input codes of the last node. Let  $m$  be the number of independent TDS the network has already learned. If  $N_e$  denotes the number of nodes in the first layer, then the total number of possible network input codes is  $2^{N_e} \geq m$ . The generalization error is defined as the average error the network makes when responding to any of the  $2^{N_e}$  possible input codes.

New learning, unknown codes, conflicts and 2-nodes mergings and splittings, as described above, have been used to impress training information in the network. But when it comes to measure the generalization error, these learning mechanisms are no longer relevant. Faced with an unknown input code, the aim is no longer to learn it, but to always provide an answer, possibly approximate, to the outgoing nodes.

The no-fail mode corresponds to the node operating behavior during the generalization error measurement phase. A COC is injected at the output of each interface node in the first layer, and the information is propagated from layer to layer. In the presence of a known CIC, the answer is simply given by the LUT. On the other hand, in the presence of an unknown CIC, the node compares it with each entries in its LUT, selects the known input code closest to the CIC (Euclidean distance) and transmits the corresponding output code as a COC. Finally the network response in no-fail mode is then the CIC of the last-layer node. The error is 0 if the CIC is indeed the one expected, 1 otherwise. We emphasize that by construction, no-fail mode provides an exact response for any network input codes that have already been learned.

When  $m = 2^{N_e}$ , the generalisation error is trivially vanishing. But what we're really interested in, and what we'll call the "grokking transition", is getting the generalization error to cancel out for values of  $r = \frac{m}{2^{N_e}}$  well below 1.

### III. LEARNING DYNAMICS

We investigate the dynamics of the algorithm, taking as an example the special case of the parity problem: for any input code consisting of  $N_e$  bits 0 or 1, the network is expected to return 0 if the number of input bits equal to 1 is even, and 0 otherwise.

There are  $2^{N_e}$  possible network entry codes and therefore  $\frac{(2^{N_e})!}{(2^{N_e}-m)!}$  ways of choosing  $m$  codes from  $2^{N_e}$ , the order of choice being relevant. Such an ordered list of  $m$  input codes is called a training set. Fig.6 is associated with a single, randomly generated training set. The top figure shows the evolution of the number of hidden nodes as a function of  $r = \frac{m}{2^{N_e}}$ , while the bottom shows the evolution of the generalization error  $E_g(r)$ .

The first thing to notice is that the algorithm really works: we can clearly see (fig.6 bottom) that  $E_g$  eventually cancels out for a value of  $r$  greater than  $\simeq 0.13$ . In other words, although it has only been provided with 13% of possible TDS, the network is able to generalize to the remaining 87% of cases, those it has never learned.

The second observation is that this convergence is far from uniform: we observe fluctuations in the generalization error (fig.6 bottom) whose amplitude is large, being of the order of magnitude of  $E_g$  averaged over the duration of the measurement. Similarly, the evolution of the number  $N$  of hidden nodes with  $r$  (fig.6 top) shows




FIG. 6. Dynamics of the algorithm in the case of the parity problem with  $N_e = 10$ . The figures show the evolution of the hidden nodes number  $N$  (top) and the generalisation error  $E_g$  (bottom) versus the relative size of the training set  $r$ . The dots are the numerical measurements while the solid line connecting them is merely a visual aid.

the same kind of fluctuations, with a start at 1, several intermediate peaks including one at 34 and a stabilization at 4 after the grokking transition. A priori, the non-uniformity of convergence is not really a surprise for a greedy algorithm, since it is generally impossible to demonstrate that local optimization (of the generalization) at one node is not at the expense of optimization at another node. But the high amplitude of the fluctuations is indicative of an avalanche dynamic, characterized by the existence of pitfalls into which the entire network plunges, and which it abruptly leaves only when it learns a new lesson.

#### IV. STATISTICS OF GROKING

In the following we discuss (always on the parity problem) the statistical aspects of learning.

##### A. Grokking as a statistical phenomenon

As in fig.6 bottom, the top of fig.7 shows the evolution of the generalization error  $E_g$  as a function of  $r$  for the parity problem. The difference is that here we have

superimposed the trajectories of 20 independent training sets. Clearly,  $E_g$  always cancels out for a sufficiently large value of  $r > r_g$ , whatever the trajectory considered. On the other hand,  $r_g$  undoubtedly depends on the trajectory considered. This is because the learning process, and therefore the value of the grokking threshold  $r > r_g$ , are stochastic phenomena in more ways than one. First, there are  $2^{N_e}$  possible network entry codes and therefore  $\frac{(2^{N_e})!}{(2^{N_e}-m)!}$  ways of choosing  $m$  codes from  $2^{N_e}$ , the order of choice being relevant. Some of these training sets are particularly un-pedagogical. For example, a bad teacher trying to explain the parity problem might start by presenting all the entry codes that will give a 0 answer, then all the others that will give a 1. He will obtain a “eureka” with difficulty from his pupil only for  $r$  much greater than  $1/2$ . Fortunately, the number of training sets is gigantic ( $> 10^{300}$  for  $N_e = 10$  and  $m = 102$ ) and the probability of randomly drawing a wrong training set is very low. Second, even for a single training set, the algorithm includes many arbitrary choices: during the splitting cascade, the order of nodes splitting in two is stochastic. Also during conflict resolution, the choice of which pairs of nodes will merge is also random.

The threshold  $r_g$  is therefore a stochastic variable. The bottom of fig.7 is a numerical measure of its probability density function (PDF), performed on a (small) sample of 220 independent training sets. The sample size is limited by the length of the calculation times. Nevertheless, in view of the narrowness of the probability distribution, it seems reasonable to define a mean value. At the top of Fig.7, the continuous blue curve represents the mean generalization error as a function of  $r$ ,  $\langle E_g(r) \rangle$ , averaged over the 20 independent learning trajectories represented by the points.

##### B. Finite size effects

Simple statistical physics models, such as single-layer perceptrons, predict cancellation of the generalization error  $E_g$  with the number  $m$  of TDS, when  $m \geq \alpha N_p$  where  $\alpha \simeq 1.39$  and  $N_p$  is the number of independent adjustable parameters.  $E_g$  tends gently towards 0 when the adjustable parameters evolve continuously, and abruptly when they take only discrete values [8–10]. Since our algorithm involves several layers, with a constantly varying number of nodes and connections (and for which  $N_p$  remains to be defined), these theoretical results are not necessarily relevant. Nevertheless, we do observe a rapid fall in the generalization error with  $r$ , as expected for discrete models. Also, numerous recent numerical observations on a heterogeneous set of algorithmic data report that the grokking threshold value  $r_g$  decreases with the number of input interface nodes  $N_e$  [11]. In [15, 28], this same dependency is derived analytically for simple, solvable models.

Here, we’re interested in the effect of the network input size on the grokking transition, but for networks which




FIG. 7. Parity problem with  $N_e = 10$ . Top: Evolution of the generalization error  $E_g$  as a function of  $r$  for 20 independent training sets. The dots are the numerical measurements and the solid blue line represents the mean generalization error averaged over these 20 training sets. Bottom: PDF of  $r_g$  based on a sample of 220 independent training sets.

have been generated by our algorithm. Fig.8 top plots the evolution of the generalization error as a function of  $r$ , for several sizes  $N_e$  of the number of input interface nodes. Each curve is obtained by averaging over 20 independent training sets. The problem studied is always that of parity because it is a non-trivial example for which  $r_g$  is small and allows reasonably large values of  $N_e$  to be investigated, but the observations are identical for the other problems studied. We therefore observe not only that  $r_g$  decreases with  $N_e$ , but also that the size of the zone over which  $E_g$  falls seems to decrease with  $N_e$ . To quantify these observations, we fit the curves in fig.8 top with an ansatz of the form:

$$E_g(r) = \frac{1}{4} [1 - \text{erf}(\lambda(r - r_g))] \quad (3)$$

where  $\lambda^{-1}$  has the dimension of an interval of  $r$  and  $\text{erf}$  denotes the error function. Fig.8 bottom plots  $r_g$  and  $\lambda$  as a function of  $N_e$  and confirms the previous observation: larger  $N_e$  brings a sharper transition, which is reached for smaller values of  $r_g$ . Unfortunately, the asymptotic behavior at large  $N_e$ , potentially highly informative, is not accessible with our current computational resources.




FIG. 8. Top: Plots of  $E_g$  versus  $r$  for several values of  $N_e$ . From right to left, we have  $N_e = 5, 6, 7, 8, 9, 10, 11, 12$ . Each curve is obtained by averaging over 20 training sets; the points are the numerical measurements, the solid lines linking them are merely a visual aid. Note the logarithmic scale of the abscissa. Bottom: Plots of  $r_g$  and  $\lambda$  defined in eq.3 versus  $N_e$ . We were only able to estimate the error bars correctly for the case  $N_e = 10$ , for which the number of training sets is sufficiently large (because of the bottom of fig.7). We then measure a relative error of 3% for  $r_g$  and 16% for  $\lambda$ .

## V. EXAMPLES OF APPLICATION

Finally, we show that the proposed algorithm can indeed generate networks capable of performing a variety of simple or complex classification tasks.

These tasks have been chosen to be as heterogeneous as possible. Some are related to arithmetic operations. The  $N_e$  input interface nodes are divided into 2 groups of respective sizes  $N_{e1}$  and  $N_{e2}$ , with  $N_{e1} + N_{e2} = N_e$ . Each set of bits is then interpreted as the binary code of an integer  $n_1$  (resp  $n_2$ ). We then define 3 operations:

1. ADD: A priori, the binary encoding of the integer  $n_1 + n_2$  requires  $L = \max(N_{e1}, N_{e2}) + 1$  bits. The ADD function returns the most significant bit, i.e. the one associated with  $2^{L-1}$ .
2. MUL: Binary encoding of the integer  $n_1 * n_2$  requires at most  $L = N_{e1} + N_{e2}$  bits. MUL returns the most significant bit, i.e. the one associated with  $2^{L-1}$ .
3. SUP: returns 1 if  $n_1 \geq n_2$ , and 0 otherwise.




FIG. 9. Evolution of the generalization error  $E_g$  as a function of  $r$  for different problems with  $N_e = 9$ . The title of the figure refers to the problem considered (see text). For RAN, the continuous straight line is the theoretical prediction of eq.4.

Other tasks are associated with the presence of a particular pattern in the binary code entered:

1. PAR: returns 0 if the number of input interface bits equal to 1 is even, and 1 otherwise.
2. TRI: returns 1 if there are at least 3 equal successive input interface bits in the binary code, and 0 otherwise.

Finally, some tasks correspond to more complicated operations. Let  $n$  be the integer associated with the binary coding provided by the  $N_e$  input interface bits. Then:

1. FIB: returns 1 if  $n$  is part of the Fibonacci sequence 1, 1, 2, 3, 5, 8, 13..., and 0 otherwise.
2. PRI: returns 1 if  $n$  is a prime number, and 0 otherwise.

## VI. DISCUSSION

As we have shown above on several examples, the proposed learning approach can bring a network to grokking a variety of tasks. Beyond demonstrating that learning and generalizing can take place without a global cost function, these results may also have strong implications about generalization itself. Indeed, in order to answer the question of what distinguishes a neural network that generalizes well from one that generalizes poorly, a number of statistical learning theories have proposed measures of complexity capable of controlling generalization

Alongside these examples, we also tested our algorithm on a stochastic classifier, RAN, returning 0 or 1 randomly, with probability 1/2. For this problem, the training set can be learnt as perfectly as for the other problems but generalization is by construction impossible. We therefore expect a generalization error  $E_g$  of the form

$$E_g(r) = \frac{1}{2}(1-r) \quad (4)$$

Fig. 9 shows the results for all the examples studied. The first observation is that our algorithm provides a network capable of generalizing in all cases where there is a rule to grok and only in those cases. There are neither false negatives nor false positives. Another point to note is that the threshold of the grokking transition is highly dependent on the problem under consideration. So, alongside the theoretical classification of problem difficulty in terms of the size of the computing resources required (memory and number of instructions in a Turing machine), the numerical value of  $r_g$  could be considered as an “experimental” value.

error [29–32]. Nevertheless in [33], the authors report on numerical results that “rule out the above complexity measures as possible explanations for the generalization performance of state-of-the-art neural networks”. Our study doesn’t answer the question of the origin of the ability to generalize either, but it highlights 2 important points: 1) it provides a theoretical framework for designing learning networks in which learning and generalization are completely decoupled. Whatever complexity measure is used to control the generalization error, it can now be studied without concern for learning mechanisms. 2) The network’s ability to generalize occurs as an emergent property, i.e. as the collective result of a local

dynamic of expressiveness optimization at node level.

The key features of our approach can be summarized as follows:

1. Standard neural networks (and also KAN networks [34]) optimize the edges of a fixed-topology network to minimize global error function, while our network stores information in LUTs and achieves generalization via dynamically evolving topology, distributing complexity both in nodes and edges.
2. There is no global cost function to minimize, nor any gradient calculation. All operations are performed locally: (i) spatial locality, as the evolution rules involve only a single node or a pair of nodes, and (ii) temporal locality, as the temporal evolution takes place after each new learning. Our algorithm is therefore naturally parallelizable.
3. In a classic deep learning approach, the choice of meta-parameters (number of layers, number of nodes, number of connections) is more a matter of know-how than technique. Our algorithm solves this problem in an elegant way: all these choices are made naturally by the network itself as it evolves.
4. Being able to decouple the problem of training data learning from that of seeking generalization greatly

extends the range of possible ways of solving the later problems. The greedy algorithm is a simple solution, but more original (and certainly more efficient) procedures are conceivable.

## VII. CONCLUSION

We have proposed and analyzed a new framework for neural computation based on dynamically evolving binary connections between dynamically evolving lookup tables. The learning phase does not rely on global minimization but instead on storing examples in LUTs before dynamically reconfiguring the network topology via a greedy algorithm which maximizes local expressivity. We have shown that such an evolved network can generalize many binary classification example tasks and that grokking shows distinctive features of a phase transition. Beyond getting rid of a global cost function, our results also shed a different light on generalization, which can be sought independently from example learning, an observation which may contribute to rethinking generalization.

The present study is only a proof of concept, limited to very small input codes ( $N_e \leq 12$ ), several orders of magnitude smaller than those commonly used by current learning machines. We are currently working on increasing this size by parallelizing the algorithm.

- 
- [1] J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities., *Proceedings of the national academy of sciences* **79**, 2554 (1982).
  - [2] C. G. Langton, Computation at the edge of chaos: Phase transitions and emergent computation, *Physica D: non-linear phenomena* **42**, 12 (1990).
  - [3] D. Marković, A. Mizrahi, D. Querlioz, and J. Grollier, Physics for neuromorphic computing, *Nature Reviews Physics* **2**, 499 (2020).
  - [4] T. L. Watkin, A. Rau, and M. Biehl, The statistical mechanics of learning a rule, *Reviews of Modern Physics* **65**, 499 (1993).
  - [5] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Reviews of Modern Physics* **91**, 045002 (2019).
  - [6] L. Zdeborová, Understanding deep learning is also a job for physicists, *Nature Physics* **16**, 602 (2020).
  - [7] S. Goldt, M. Mézard, F. Krzakala, and L. Zdeborová, Modeling the influence of data structure on learning in neural networks: The hidden manifold model, *Physical Review X* **10**, 041044 (2020).
  - [8] G. Györgyi, First-order transition to perfect generalization in a neural network with binary synapses, *Physical Review A* **41**, 7097 (1990).
  - [9] H. Sompolinsky, N. Tishby, and H. S. Seung, Learning from examples in large neural networks, *Physical Review Letters* **65**, 1683 (1990).
  - [10] H. S. Seung, H. Sompolinsky, and N. Tishby, Statistical mechanics of learning from examples, *Physical review A* **45**, 6056 (1992).
  - [11] A. Power, Y. Burda, H. Edwards, I. Babuschkin, and V. Misra, Grokking: Generalization beyond overfitting on small algorithmic datasets, *arXiv preprint arXiv:2201.02177* (2022).
  - [12] N. Rubin, I. Seroussi, and Z. Ringel, Droplets of good representations: Grokking as a first order phase transition in two layer networks, *arXiv preprint arXiv:2310.03789* (2023).
  - [13] Z. Liu, E. J. Michaud, and M. Tegmark, Omnigrok: Grokking beyond algorithmic data, in *The Eleventh International Conference on Learning Representations* (2022).
  - [14] W. Kinzel, Phase transitions of neural networks, *Philosophical Magazine B* **77**, 1455 (1998).
  - [15] B. Žunković and E. Ilievski, Grokking phase transitions in learning local rules with gradient descent, *Journal of Machine Learning Research* **25**, 1 (2024).
  - [16] J. W. Miller, C. O’Neill, and T. D. Bui, Grokking beyond neural networks: An empirical exploration with model complexity, *Transactions on Machine Learning Research* (2024).
  - [17] R. M. French, Catastrophic forgetting in connectionist networks, *Trends in cognitive sciences* **3**, 128 (1999).
  - [18] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Rammalho, A. Grabska-Barwinska, et al., Overcoming catastrophic forgetting in neural networks, *Proceedings of the national academy of sciences* **114**, 3521 (2017).
  - [19] D. O. Hebb, *The organization of behavior: A neuropsychological theory*

- (Wiley, New York, 1949).
- [20] R. M. French, Pseudo-recurrent connectionist networks: An approach to the 'sensitivity-stability' dilemma, *Connection Science* **9**, 353 (1997).
- [21] S. Grossberg, Competitive learning: From interactive activation to adaptive resonance, *Cognitive science* **11**, 23 (1987).
- [22] T. P. Lillicrap, A. Santoro, L. Marris, C. J. Akerman, and G. Hinton, Backpropagation and the brain, *Nature Reviews Neuroscience* **21**, 335 (2020).
- [23] S. Marsland, J. Shapiro, and U. Nehmzow, A self-organising network that grows when required, *Neural networks* **15**, 1041 (2002).
- [24] A. Horzyk and R. Tadeusiewicz, Self-optimizing neural networks, in *Advances in Neural Networks—ISNN 2004: International Symposium on Intelligent Neural Networks, Dalian, China, August 2004*, Proceedings, Part 1 (Springer, 2004) pp. 150–155.
- [25] G. I. Parisi, J. Tani, C. Weber, and S. Wermter, Lifelong learning of human actions with deep neural network self-organization, *Neural Networks* **96**, 137 (2017).
- [26] C. Wiwatcharakoses and D. Berrar, Soinn+, a self-organizing incremental neural network for unsupervised learning from noisy data streams, *Expert Systems with Applications* **143**, 113069 (2020).
- [27] G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, and S. Wermter, Continual lifelong learning with neural networks: A review, *Neural networks* **113**, 54 (2019).
- [28] B. Žunkovič and E. Ilievski, Grokking phase transitions in learning local rules with gradient descent, arXiv preprint arXiv:2210.15435 (2022).
- [29] P. L. Bartlett and S. Mendelson, Rademacher and gaussian complexities: Risk bounds and structural results, *Journal of Machine Learning Research* **3**, 463 (2002).
- [30] S. Mukherjee, P. Niyogi, T. Poggio, and R. Rifkin, Statistical learning: Stability is sufficient for generalization and necessary and sufficient for consistency of empirical risk minimization, Technical Report AI Memo 2002-024, MIT (2002).
- [31] O. Bousquet and A. Elisseeff, Stability and generalization, *The Journal of Machine Learning Research* **2**, 499 (2002).
- [32] T. Poggio, Neural networks and the universal approximation conditions for predictivity in learning theory, *Nature* **428**, 419 (2004).
- [33] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, Understanding deep learning (still) requires rethinking generalization, *Communications of the ACM* **64**, 107 (2021).
- [34] Z. Liu, Y. Wang, S. Vaidya, F. Ruehle, J. Halverson, M. Soljačić, T. Y. Hou, and M. Tegmark, Kan: Kolmogorov-arnold networks, arXiv preprint arXiv:2404.19756 (2024).