

Fast gradient-free activation maximization for neurons in spiking neural networks

Nikita Pospelov^{a,*}, Andrei Chertkov^{b,c}, Maxim Beketov^d, Ivan Oseledets^{b,c}, Konstantin Anokhin^a

^a Laboratory of Neuronal Intelligence, Institute for Advanced Brain Studies, Moscow, Russia

^b Artificial Intelligence Research Institute (AIRI), Moscow, Russia

^c Skolkovo Institute of Science and Technology, Moscow, Russia

^d HSE University, Moscow, Russia

ARTICLE INFO

Communicated by J. Liu

Keywords:

Neuronal selectivity
Activation maximization
Effective stimuli
Explainable AI
Neural representation
Derivative-free optimization
Spiking neural networks
Tensor Train decomposition
Generative models

ABSTRACT

Elements of neural networks, both biological and artificial, can be described by their selectivity for specific cognitive features. Understanding these features is important for understanding the inner workings of neural networks. For a living system, such as a neuron, whose response to a stimulus is unknown and not differentiable, the only way to reveal these features is through a feedback loop that exposes it to a large set of different stimuli. The properties of these stimuli should be varied iteratively in order to maximize the neuronal response. To utilize this feedback loop for a biological neural network, it is important to run it quickly and efficiently in order to reach the stimuli that maximizes certain neurons' activation with the least number of iterations possible. Here we present a framework with an efficient design for such a loop. We successfully tested it on an artificial spiking neural network (SNN), which is a model that simulates the asynchronous spiking activity of neurons in the brain. Our optimization method for activation maximization is based on the low-rank Tensor Train decomposition of the discrete activation function. The optimization space is the latent parameter space of images generated by SN-GAN or VQ-VAE generative models. To our knowledge, this is the first time that gradient-free activation maximization has been applied to SNNs. We track changes in the optimal stimuli for artificial neurons during training and show that highly selective neurons can form already in the early epochs of training and in the early layers of a convolutional spiking network. This formation of refined optimal stimuli is associated with an increase in classification accuracy. Some neurons, especially in the deeper layers, may gradually change the concepts they are selective for during learning, potentially explaining their importance for model performance. The source code of our framework, MANGO (for Maximization of neuronal Activation via Non-Gradient Optimization) is available on GitHub (<https://github.com/iabs-neuro/mango>)

1. Introduction

Brain neurons respond selectively to certain properties of stimuli. For instance, neurons in the primary visual cortex are specialized on simple properties of visible images, such as the direction of motion [1] or color [2]. Meanwhile, neurons from higher level visual areas specialize in more complex stimuli, such as the presence of a face within the visual field [3]. Identifying the specialization of biological neurons involves varying the properties of a stimulus (e.g., the image presented to the neuron) iteratively to find those that elicit the most intense response [4,5] - these approaches are known as activation maximization (AM). However, to effectively solve this problem *in vivo*, one needs to:

- Find an optimal stimulus in as few iterations and as accurately as possible.

- Investigate the stimulus space for multiple activation optima, including local and global optima.

Significance of feedback loops in studies of intelligence can hardly be overstated – these were extensively studied in the field of Cybernetics [6] – both in biological and artificial systems, and can be argued essential for a system to be intelligent. Feedback loops are at the heart of many Reinforcement Learning (RL) [7,8] approaches: there exist so-called reinforcement learning from human feedback (RLHF) techniques. Present work can be seen in this light, only with a SNN as an intelligent agent providing feedback. SNNs were chosen as a testbed for the development of our framework to the point of it being fast and efficient enough before any *in vivo* studies can be performed on living organisms.

* Corresponding author.

E-mail address: pospelov.na14@physics.msu.ru (N. Pospelov).

This is especially important considering the ethical implications of such studies in neuroscience [9,10]. The faster and better one can obtain results, the less stress is imposed onto the studied animals, and the fewer animals are needed. To address this challenge, it is necessary to develop and incorporate new mathematical optimization methods into experimental setups. This will allow to obtain the best approximations of the global activation optimum with a minimal number of iterations, meaning that the animal will be exposed to stimuli (such as images or sounds) for a shorter period of time.

Nowadays, it is widely believed that the learning mechanisms of deep neural networks are significantly different from those in the brain [11,12]. Therefore, all direct comparisons between the two should be made with caution. However, modern artificial neural networks, despite their diverse architectures and training methods, exhibit similar selectivity towards complex stimuli [13]. Interestingly, just like in the brain, ANN neurons can exhibit selectivity for multiple stimuli [14].

Since ANNs are essentially computational graphs of functions parameterized by the weights of layers, and modern computational hardware such as GPUs and TPUs are capable of efficiently computing both the value of the function at a point for inference and its derivatives with respect to parameters for backpropagation, it is relatively easy to “dissect” ANNs [15] and analyze the “responsibility areas” of each artificial neuron [16].

This makes ANNs ideal for applying activation maximization techniques based on gradients computation (see, for example, [17] for a review). Gradient-free approaches are also applicable to this problem, such as those described in [18], where 15 different methods were compared for maximizing the response of a specific neuron in a deep ANN using a pre-trained generative adversarial network (GAN) [19]. A comprehensive comparative analysis of various optimization techniques, including those based on genetic and evolutionary optimization, was conducted in [20]. An original approach, presented in the latter, based on the Tensor Train decomposition (TT decomposition, Oseledets [21]) was found to have significant advantages in terms of speed, accuracy, and stability (given the same number of iterations).

TT-decomposition allows one to represent a tensor (multidimensional array) in a compact and descriptive low-parameter form. In addition to reducing memory consumption, TT-format makes it possible to perform many algebraic operations on tensors: element-wise addition and multiplication, convolution and integration, solving systems of linear equations, calculating statistical moments, etc. — with complexity being linear in dimension. TT-format provides a wide range of methods for constructing surrogate models and effectively performing linear-algebraic operations. Recently it was shown in [20,22,23] that TT-format can also be successfully applied for black-box (gradient-free) optimization, and the quality of solutions in many cases turns out to be higher than of methods based on alternative gradient-free approaches (e.g. genetic/evolutionary algorithms). Since in the problem considered in present work one does not have explicit access to the gradients of neurons’ activation, we have chosen this promising tensor-based gradient-free approach as the basis of our method.

A more biologically realistic [24] (compared to classic ANNs) model of living neural networks is the family of artificial spiking neural networks (SNNs) [25], the activity of which unfolds in time — in contrast to conventional feedforward neural networks, which are mathematically just functions without any inherent temporal dynamics. SNNs represent an important step towards bridging the gap between the inner workings of artificial and biological neurons, as spike timing has been shown to play a crucial role in brain functioning [26,27]. With this biological relevance in mind, in present work we focus on the activation maximization of neurons in SNNs. Recently, the inner neuronal representations of data in SNNs have been studied and compared to those in ordinary artificial neural networks [28].

In our work we attempt to solve a related but different problem — finding effective stimuli (MEIs for Most Exciting Images, termed

in [29]) specifically for SNN neurons. We aim to do this efficiently, with as few iterations as possible, and with a view towards future applications to biological neural networks. There have been recent efforts to reveal what excites neurons in SNNs [30] — but, to the best of our knowledge, these efforts have used optimization methods with explicit access to gradients (or surrogate approximations thereof) of neurons’ activation functions. In present work we propose to specifically focus on gradient-free activation maximization for revealing MEIs of SNN neurons — to the best of our knowledge, this is being done for the first time.

This work aims to contribute to the developing field of interpretable AI and feature visualization [31–36], as well as potentially provide useful tools for optimal stimulus search in biological experiments. By doing so, it aims to help establish connections between natural and artificial intelligence through a unified computational framework. Our contributions in this paper can be summarized as follows:

- We propose a new activation maximization method based on the low-rank Tensor Train decomposition.
- We compare it to existing optimizers in the Nevergrad library [37] and show its good convergence to optimal stimuli and superior performance. We apply our method to learning the MEIs of a spiking ResNet convolutional neural network, which is, to the best of our knowledge, the first time that gradient-free stimuli optimization methods have been applied to SNNs.
- Furthermore, we studied the computed MEIs layer-wise throughout the network training and correlated model performance with the formation of highly selective neurons.

2. Background

2.1. Activation maximization in artificial and biological neural networks

Our goal in present work is to optimize the activation of a specific neuron in an artificial (ANN) or spiking (SNN) neural network by finding its most exciting input (MEI) stimulus. This problem lies at the intersection of contemporary neuroscience and deep learning (DL). In DL, well-known studies such as XDream [38] and Inception Loops [29], as well as Nguyen et al. [35] — have demonstrated how generative models can effectively map the latent space of stimuli (images) to perform activation maximization in ANNs. In neuroscience, recent research has shown that it is possible to study the function of neurons in living organisms by incorporating a generative model such as XDream into a feedback loop that exposes a living subject (a mammal with a well-developed visual system) to generated images. In [4], XDream was used to identify the MEIs in the visual cortex of macaque monkeys. In [5] further exploration found that MEIs do not necessarily encode exact faces, but can also encode abstract objects (many of which look face-like).

In present work, we aim to contribute to the rapidly developing field of integrating deep learning techniques into neuroscience to address the long-standing question of neural coding principles. We introduce a novel approach that efficiently solves the activation maximization problem, and does not require any information about the underlying system’s structure.

With the activation intensity/frequency of a certain neuron viewed, mathematically, as a function of the input “fed” to the neural network, whether it is a biological or artificial neuron, the MEI is defined as

$$\text{MEI}_i \stackrel{\text{def}}{=} \arg \max_{x \in \mathcal{S}} A_i(x) \quad (1)$$

where \mathcal{S} denotes the stimulus space, $A_i(x)$ is the activation function of i -th neuron in the system. In real-world applications, $A_i(x)$ is always measured with noise which complicates MEI determination. In practice, the MEI of a neuron is calculated iteratively. Additionally, due to constraints imposed by real-world experiments, we want to calculate MEI,

as efficiently as possible using the least number of iterations possible. This is essential to save time and maintain the focus of the animal being studied in an *in vivo* experiment, and to save computational resources in an *in silico* experiment.

When studying a biological system, such as an animal’s brain, we do not have access to its internal structure and do not fully understand how information about stimuli is transmitted to specific neurons. This is informally known as a black box function optimization problem — without explicit access to the gradient (derivatives) of the function.

Optimization methods in general can be roughly divided into:

1. Gradient-based methods, dating back to Newton and his method. Algorithms of this class have evolved dramatically, with variations of Stochastic Gradient Descent (SGD) [39] powering the training of millions of neural network-based machine learning models around the world today.
2. Gradient-free methods that do not explicitly use information about the system’s structure. There are several families of these methods, including deterministic and stochastic ones. Evolutionary algorithms, which have gained particular attention in the context of activation maximization in recent years [18,38,40], fall into this category.

Gradient-free methods are naturally more appropriate for the problem of AM in real-world systems due to activation being a black-box-like function. Of course, even if not explicitly accessible, the gradient can be approximated by finite differences. However, if the function’s landscape is very “wobbly” – quickly changing – the error of gradient estimates will be effectively unbounded; and noise that is present in measurements of function’s values will also contribute to the error of gradient estimates.

In the context of AM, quite recently such gradient-free methods have been explored [40], and extensively compared to each other [18] — both in silico and in vivo, with Covariance Matrix Adaptation (CMA) [41] outperforming most in both settings.

In present work, we introduce (for the first time) another family of gradient-free optimization algorithms for the problem of SNN’s neuron activation maximization — ones based on low-rank tensor decompositions (Tensor Train, Oseledets [21]) of the optimized function. Some of these methods are essentially grid search for a maximum in an array, but a “smart” one, utilizing the low-rank tensor nature of the array, thus achieving exponential speedup over straightforward search. Others belong to the Monte-Carlo (MC) family of methods, which are less prone to being stuck in local optima by stochastically “jumping” around the array, not necessarily towards the optimum. See the corresponding section of Methods for a self-contained introduction of these methods.

2.2. Spiking Neural Networks (SNNs)

From a mathematical perspective, ordinary ANNs are simply parameterized functions that transform input data to produce output — such ANNs are called feedforward. Of course, the field of contemporary deep learning [42,43] is rich in further complications of this idea, e.g. NN-based generative models can be said to be non-deterministic functions, but for the purposes of this discussion we limit ourselves to the basics. Unlike in ANNs, the activity of biological NNs unfolds in time. Biological neurons are constantly accumulating signals from neighbor neurons they share synaptic connections with, and, if the total voltage amplitude of incoming signals (membrane potential, MP) exceeds some threshold, the neuron “fires”. That is — it emits what is called a spike, a signal the form of which is very special (and well-studied [44]) due to the bio-chemical dynamics that governs its emission and transmission. The spike then travels along the axon to downstream neurons — to provoke them to spike in their turn. This inherent temporal nature of biological NNs was omitted when the first artificial models of neural networks, namely the 1943 neuron model of

McCulloch & Pitts [45], further developed and implemented in hardware and termed Perceptron by Rosenblatt [46,47] in 1957–58, were introduced. The invention of the error back-propagation algorithm [48] in 1986 as a useful method for training such ANNs (e.g. to perform classification tasks after learning on exemplar data) further cemented the dominance of feedforward ANNs as machine models of intelligence. It was Maass’s work [49] in 1997 that introduced networks of artificial spiking neurons as a model of biological NNs incorporating temporal dynamics.

An artificial spiking neural network (SNN), just like a biological one, is a dynamical system, with each neuron being a subsystem with its own internal dynamics. Neurons are connected with each other, forming layers with connections being directional. Over time, a neuron is receiving signals from its sourcing neighbors. These signals are being summed with weights designating the importance of connections — to give the total value of “membrane potential”. If this value exceeds a certain threshold (one of the hyperparameters of this model), the neuron “fires” — emits a spike into its output. For this to be effectively modeled digitally, this dynamics is discretized: the time-dependence of signal is encoded by an array of numbers (signal amplitudes). These arrays are basically a discrete model of what is called spike trains (the term used for signals in biological NNs as well). This discretization is not too far-fetched as a model of biological reality: spikes in biological neurons are almost “binary”.

Nowadays there are several software implementations of artificial SNNs, with at least two outstanding frameworks — `snnTorch` [50] & `SpikingJelly` [51], allowing to do complete training and inference procedures — these two were used in present work.

There are several important questions to answer when working with an SNN. First, how is input data encoded for the SNN to process it? In our work we “expose” an SNN to static images. For that there are at least two approaches [50]:

1. simply repeating the image (unaltered) a certain number of frames (kind of a static video)
2. showing the image on a certain number of frames, but altering it with some sort of noise.

We have chosen the first approach for simplicity, so in our work the first spiking layer was exposed to input image repeated for 20 to 100 frames. The model we were working with is basically a spiking analogue of ResNet [52], see Section 2.2 for further details.

The second basic question is: what exact model of a spiking neuron does one use? The above description is very schematic, one needs to specify a certain mechanistic model (an electric circuit, to be precise) of the biological neuron, so that the dynamics of such a model accurately approximates reality. For practical applications, Hodgkin–Huxley model [53], while being most biologically and physically accurate, is too complicated to implement and work with. So a practical choice is the one of so-called Leaky Integrate-and-Fire (LIF) neuron models, that date back to Louis Lapicque’s 1907 work [54].

In LIF models, a spiking neuron is represented as an RC circuit — instead of directly summing the incoming voltages, such a neuron integrates input signal over time with leakage (due to R for resistance present in the circuit). A LIF neuron abstracts away the shape and profile of the output spike: it is simply treated as a discrete event. As a result, information is not stored within the spike, but rather the timing (or frequency) of spikes. Discrete-time equations of the spiking dynamics for such a neuron thus are:

$$S[t] = \Theta(U[t] - U_{\text{thr}}) = \begin{cases} 1, & \text{if } U[t] > U_{\text{thr}} \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

with $S[t]$ being the signal (spike) intensity of a neuron at a discrete moment of time t , $U[t]$ being the MP. With $\Theta(x)$ being Heaviside’s step function, a neuron discretely fires if its MP $U[t]$ exceeds a threshold

value. The RC-circuit nature of a LIF neuron is described by the following (discretized) temporal dynamics equation on $U[t]$:

$$U[t+1] = \underbrace{\beta U[t]}_{\text{decay}} + \underbrace{W X[t+1]}_{\text{input}} - \underbrace{S[t] U_{\text{thr}}}_{\text{reset}} \quad (3)$$

— at next moment of time, $t + 1$, the MP is given by a linear combination (with learnable weights W) of inputs $X[t]$ from sourcing neighbor neurons. There is also some decay (leakage) and a reset term. With U_{thr} set to, say, 1, the only hyperparameter left to set is the decay rate, β .

The third basic question is: how does one learn the weights of neurons’ connections W ? For that to be solvable with common machine learning approaches, the total loss function of SNN’s activation on certain input data, $\mathcal{L}_W(\text{input})$, should be differentiable w.r.t. parameters of the network, W . The problem is that Heaviside’s theta-function in Eq. (2) is not differentiable at zero, and has zero derivative elsewhere. A traditional [50] way to overcome this is to smooth Θ , replacing $S[t]$ with some surrogate, $\tilde{S}[t]$, which could be any sigmoidal function, typically logistic function or arctangent. With \tilde{S} being a smooth function, the NN’s weights can be updated with gradient optimization methods. The choice of this surrogate function is another hyperparameter to be set when defining an SNN.

In this work, we view artificial SNNs as biologically-realistic models [24] of NNs of the brain. In a typical biological experiment, the internal state parameters of each neuron in a biological neural network are not completely observable — rather, the neuron’s activity, resulting from this state and the state of the environment, is measurable. This activity is thus known as a “black box” function of response to the stimuli the NN is exposed to: only its values at some points of the stimuli space are known, not the gradients. For this reason, we constrain ourselves to consider this as a problem to be solved with gradient-free optimization methods. This is in contrast with the problem of AM for non-spiking ANNs, the activity of which is explicitly differentiable, and hence gradient-based optimization methods are available. Present work was heavily inspired, in particular, by Xiao and Kreiman [40] — where, although ANNs in consideration are differentiable, exactly this gradient-free paradigm is advocated. Such approach can be viewed as a testing ground for future *in vivo* experiments, with the ANNs replaced by biological NNs, the activity of which is not differentiable.

3. Methods

3.1. Tensor Train (TT) decomposition and TT-based optimization techniques

As described above, our goal is to maximize the activation (spiking frequency) of certain neurons in an SNN trained to classify CIFAR-10 images. So we are optimizing a black-box function, the domain of which is the $\sim 10^{2-3}$ -dimensional latent space of image features (the latent representations of images are available with a generative model). If one picks a certain region of this domain and discretizes it (in our experiments, the latent space of images was discretized into a 128-dimensional cube with 64 points on each side), this high-dimensional tensor might be of quite low rank. The reasoning behind this is that both ANNs (by design) and biological NNs (due to restrictions of the physical world) do not have exponential resources to encode all the places of interest in this space. This hypothesis of low rank nature of biological NNs’ activity is an open question, currently being researched [55,56]. Under such assumptions, optimization methods that are based on low-rank tensor decompositions might prove very effective.

One kind of such decompositions that turned into a whole fruitful research field is the Tensor Train (TT) decomposition, introduced in 2011 by Oseledets [21]. It allows to encode a low-rank high-dimensional tensor in a compact and convenient format, only using

a polynomial (in dimension of the tensor) number of variables, instead of exponential required in general. The format resembles, if depicted graphically as what is called a tensor network, linked train cars, hence the name. In this format, tensors can be operated with: added, multiplied, convoluted, etc — very effectively (in the amount of operations and memory). All that provides a basis for applying TT-decomposition to solve various linear and nonlinear equations, PDEs, etc. Of many possible applications of TT-decomposition, other than to optimization problems, one could especially note “tensorizing NNs” [57,58] — effectively building a compact tensor approximation to a nonlinear ANN. Apart from speeding up inference tasks, this seems very related to the subject of present work: the whole landscape of activation values of a certain neuron in a neural network could be compressed and investigated in detail. If the conjecture of low rank holds for activation functions of biological neurons, this would mean that only a polynomial (in latent space dimension) number of “requests” (expositions to stimuli) is needed to effectively approximate the response of the biological NN. This seems to be a promising direction of future research.

There are thus several new optimization algorithms based on the TT-decomposition, where the discretized optimized function $f(x)$ is better and better approximated $\tilde{f}_{TT}(x)$ with a Tensor Train at each iteration of computing its values at different points. At the same time, the optimum candidate (optimum of $\tilde{f}_{TT}(x)$) is quickly found exploiting the structure of TT. This is the basic idea behind TTOpt optimization method [20]. Another related method is Optima-TT [59]. Also see [60,61] for an overview of tensor approximation methods for optimization.

For our problem we have tried several TT-based optimization methods, and found PROTES [22] to be the most effective. PROTES stands for “Probabilistic Optimization with Tensor Sampling” — it is a probabilistic optimization method (see Fig. 2). The main idea is similar to that of Simulated Annealing algorithms (see [62]) of the Monte-Carlo family. While gradient-based methods may be stuck in local optima, such methods have some probability of “jumping” out of these, if the temperature (yet another hyperparameter) of the “optimum-seeking particles” is high enough. The intuition behind PROTES is as follows: given a target function $f(x)$ to minimize (maximization is done by minimizing $-f(x)$), apply the following monotonic transformation to it:

$$F[f](x) = \frac{1}{\exp((f(x) - y_{\min} - E)/T) + 1}. \quad (4)$$

In physics this expression is known as Fermi-Dirac statistic. With y_{\min} being an exact or approximate minimum of f , $T > 0$ - the “temperature” parameter, and E — some “energy” threshold. With F being a CDF-like function, one can stably find the maximum of its expectation: $\max_{\theta} \mathbb{E}_{\xi_{\theta}} F[f](\xi_{\theta})$ where a family of random variables ξ_{θ} has a parametric distribution with density $p_{\theta}(x)$. Using the REINFORCE trick [63], one can estimate the expectation gradient as

$$\nabla_{\theta} \mathbb{E}_{\xi_{\theta}} F[f](\xi_{\theta}) \approx \frac{1}{M} \sum_{i=1}^M F[f](x_i) \nabla_{\theta} \log p_{\theta}(x_i) \quad (5)$$

with Monte-Carlo — $\{x_i\}_1^M$ being i.i.d. realizations of the r.v. ξ_{θ} . If one manages to find optimal parameter values $\tilde{\theta}$ of p_{θ} , then $p_{\tilde{\theta}}$ is expected to have a peak at the maximum of $F[f]$. At “low temperature” (small values of T), only a few terms contribute to the sum (5) — namely, with those x_i for which $f(x_i) - y_{\min} < E$ (“low energy particles”) — for those, $F[f] \approx 1$, while for others it is close to zero. Thus one keeps only a few best values of the sample. With p_{θ} having a low-rank TT-representation, the above procedure (sampling and finding top-n values of the array) can be done quickly and effectively.

3.2. Generative modeling

The task of a generative model in the pipeline of our experiment is to effectively “map” the space of stimuli, providing lower-dimensional latent coordinates for the subspace of “natural” images in the huge

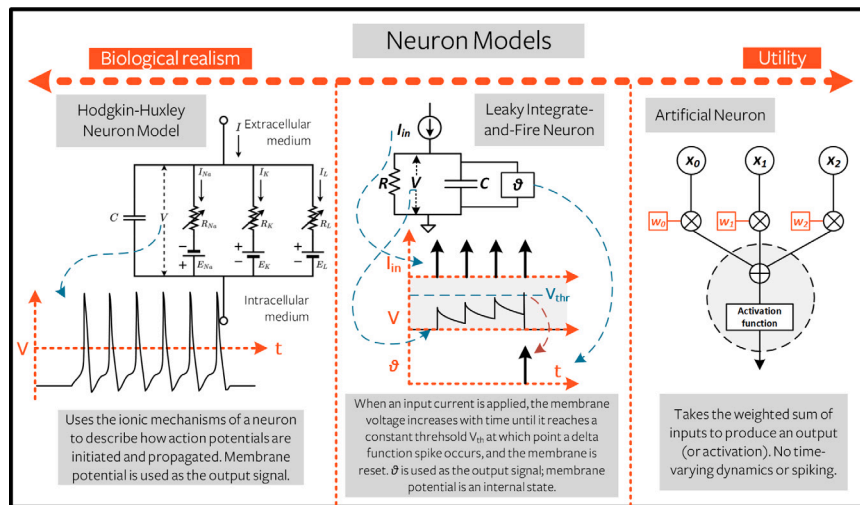


Fig. 1. Overview of spiking neuron models
Source: Image from [snnTorch online tutorial](#) [50].

space of all possible images of a given size. The models were trained on CIFAR-10 [64] — a dataset of 60,000 color images of size 32×32 pixels, split into 10 classes (6,000 images per class): (0) airplane, (1) automobile, (2) bird, (3) cat, (4) deer, (5) dog, (6) frog, (7) horse, (8) ship, (9) truck — see Fig. B.16 in Appendix B for examples.

We have tried two types of generative models that are known to perform well on this task:

1. Generative Adversarial Networks (GANs) [19], specifically their spectral-normalized version (SN-GAN) [65]
2. Variational AutoEncoders (VAEs) [66], specifically their “discretized” (vector-quantized) version – VQ-VAEs [67]

We have found SN-GANs more suitable in our setting — see Appendix B for a review on VQ-VAEs and comparison of both models performance on our task. Thus in this section we only provide an overview of GANs.

A Generative Adversarial Network (GAN) [19] is comprised of two NNs, a generator G and a discriminator D , playing a minimax game: G has to learn a distribution close to that of natural examples (images), fooling D that they are realistic; D has to discriminate images produced by G from real ones — by learning the probability distribution of the latter (supported on the space of latent features of the images).

Due to the unstable nature of this minimax game G and D are playing, original GANs do not perform exceptionally well on data such as CIFAR-10, so quite soon after their introduction, many regularization techniques were proposed. Quite popular one being the Spectral Normalization (SN-GAN) [65] — with the layers’ weight matrices being penalized so that their Lipschitz constant is bounded from above. This simple yet elegant solution allowed the authors of SN-GAN reach then-state-of-the-art results on CIFAR-10 in terms of metrics of inception score [68] and Fréchet inception distance (FID) [69]. Subjectively, the quality of images generated by an SN-GAN trained on CIFAR-10, is very impressive.

After running several experiments, we have settled on using an SN-GAN with the latent space of dimension 128, with each dimension discretized into 64 parts.

3.3. SNN architectures

Before moving on to experiments with maximizing the neuronal responses *in vivo*, it is necessary to fully prepare and test the pipeline of such an experiment, where biological neurons are replaced by some adequate *in silico* model. The most studied class of such models are

artificial spiking neural networks (SNN). Among all the software implementations of SNN, two most well-known and developed one were chosen — SNNtorch [50] and SpikingJelly [51] frameworks.

What is special about spiking neural networks, both artificial and living ones — is that the spikes are discrete in time. So it is not immediately obvious how to backpropagate error through the output of an artificial SNN, as noted in Section 2.2. In present work, we use the solution introduced in [50], based on the so-called surrogate gradient method, see Fig. 1. That is, local learning rules common in SNN research were not applied, instead, we trained the network from scratch using surrogate gradients to backpropagate through spiking layers.

First, a simple convolutional NN based on SNNtorch framework was implemented and trained to solve the problem of image classification from the CIFAR-10 dataset (10 classes of color images of 32×32 pixels). Model architecture consisted of “ordinary” (feedforward, not spiking) convolutional layers, as well as fully-connected layers. The spiking nature of the model was achieved by projecting the input of ordinary layers to the Leaky Integrate-and-Fire neurons with a membrane decay constant of 0.9. It should be noted that many possible approaches to this so-called ANN-to-SNN (ANN2SNN) conversion exist [70,71], as well as many interesting hybrid training [72,73] methods.

The resulting SNN was trained to achieve classification accuracy of 72% on the CIFAR-10 dataset. This number is lower than that of state-of-the-art feedforward (not spiking) architectures, but we considered it quite sufficient for our preliminary studies.

Next, to speed up inference and increase network depth, we integrated SpikingJelly [51] into our framework. We used a spiking analogue of a celebrated ResNet18 model [52] for the main results of present work, since this architecture provided an optimal trade-off between depth and inference speed. Again, we used LIF neurons with membrane decay constant equal to 0.9. The trained network achieved 86% accuracy on CIFAR10 dataset.

Single neuron activity was determined by the number of spikes that it produced during “exposure” time (time was discretized into $T = 100$ counts for our SNNtorch-based network and $T = 50$ for SpikingJelly-based one). To address the spiking nature of the network, each image fed as the input to the model was copied T times along a separate tensor dimension. For each training epoch, we scanned neural activity from all spiking layers of the model (19 layers in total). Activity was averaged over all feature channels and normalized to the range $[0,1]$.

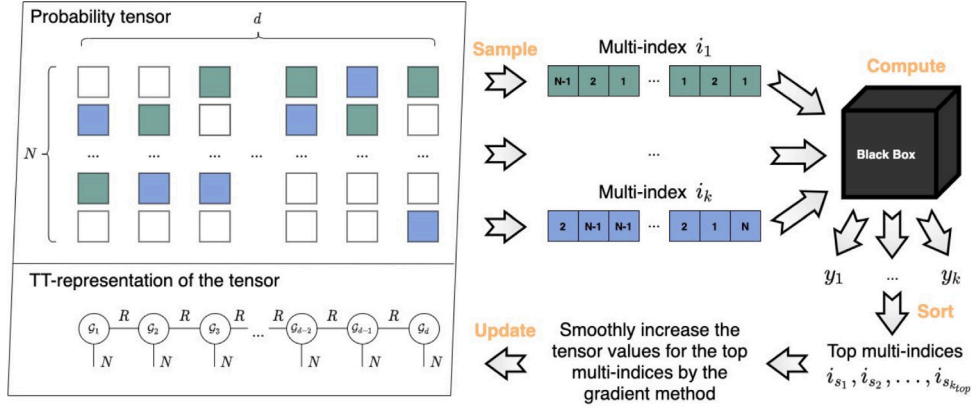


Fig. 2. Schematic of the PROTES method.

3.4. Optimization algorithms

One of the most popular and effective software packages for gradient-free optimization is the Nevergrad framework [37]. Methods implemented there were used as benchmarks for our research. One distinctive feature of Nevergrad is that it implements techniques that combine several optimization approaches into one, called “Portfolios”.

Based on our testing, we found that the basic Portfolio method, which combines methods of (1) covariance matrix adaptation evolution strategy (CMA-ES), (2) differential evolution (DE), and (3) SCR-Hammersley, performed the best in our problem setting. We also compared it to other popular gradient-free methods implemented in Nevergrad, such as (1) OnePlusOne, (2) NoisyBandit, and (3) simultaneous perturbation stochastic approximation (SPSA), which all showed poorer results compared to the default Portfolio. It should be noted that in [18], it was demonstrated that the CMA-based optimization method outperformed the maximum activation of genetic algorithms (by 66% in silico and 44% in vivo), which are popular choices for non-gradient-based AM tasks.

As an alternative to Nevergrad algorithms, we used PROTES optimization method [22] (described in previous sections), based on the low-rank Tensor Train decomposition [21]. In addition to baseline PROTES, we applied quantization of the tensor modes (transforming each tensor dimension into a new set of dimensions with smaller modes), which allowed us to achieve better results compared to the baseline PROTES method.

To summarize, we compared the following optimization methods with the following hyperparameters:

1. (baseline) Random search
2. (baseline) Portfolio from the Nevergrad package
3. TT = PROTES with $K = 10$; $k_{top} = 1$
4. TT-s = PROTES with $K = 5$; $k_{top} = 1 + \text{mode quantization}$
5. TT-b = PROTES with $K = 25$; $k_{top} = 5 + \text{mode quantization}$

For PROTES, K stands for number of optimum candidates, sampled from the model probability density; and k_{top} is the number best candidates out of K , based on which the parameters of this model density are updated on the last step of each iteration.

The tensor rank hyperparameter for all PROTES instances was set $r = 5$ — that is, each “car” of the Tensor Train approximation was a tensor of rank 5. We have run experiments with different r values in the nearby range, and observed effects were similar, so we report results with $r = 5$.

The outline of our TT-based optimization algorithm can be found in Algorithm 1.

Algorithm 1: MEI construction using PROTES

Data: function that computes the activation of the target neuron $a = f(x)$ for any input image $x \in \mathbb{R}^{d_{inp}}$ having d_{inp} pixels; generative model $x = g(z)$ that produces artificial input (image) x for the provided latent vector $z \in \mathbb{R}^d$, where d is a dimension of the latent space; maximum number of queries m (i.e., the computational budget).

Result: input image $x^{(opt)}$ that maximizes the activation.

Function $\text{loss}(z)$:

- Generate the artificial input (image): $x = g(z)$
- Compute the activation of the neuron: $a = f(x)$
- return** a

Optimize until the budget is exhausted: $z^{(opt)} = \text{protes}(\text{loss}, \text{budget} = m)$

The optimizer sequentially generates candidates for the optimum, performs corresponding queries to the objective function, and updates the internal expectation tensor for the optimum based on the returned values. Please see Algorithm 1 with details of the PROTES optimization method in the original work [22].

Generate the artificial input (image): $x^{(opt)} = g(z^{(opt)})$

return $x^{(opt)}$

3.5. MANGO framework

We have developed a software framework for fast and accurate calculation of MEIs in artificial neural networks — called MANGO (Maximization of neuronal Activation via Non-Gradient Optimization). We were particularly interested in considering models that are more closely related to the functioning mechanisms of biological neurons, specifically spiking neural networks (SNNs), as one of our main future goals is to apply these methods to living systems. However, the gradient-free optimization techniques we propose in this paper are also well suited for classical neural networks (see Fig. 3).

Within the framework, one can select a dataset, a generator model, a target neural network model, and an optimization method. One can also generate MEI search using various hardware backends and save and analyze the results.

The following list shows the options that have been added to the framework:

1. Datasets: MNIST [74], Fashion-MNIST [75], CIFAR10 [64], Imagenet [76]
2. Generative models: VQ-VAE [67], SN-GAN [65]
3. Classic neural networks: AlexNet [77], Densenet [78], VGG [79], ResNet [52]
4. Spiking neural networks: SNN-Torch [50]-supported model convolutional networks and SpikingJelly-supported spiking ResNet18 [51]

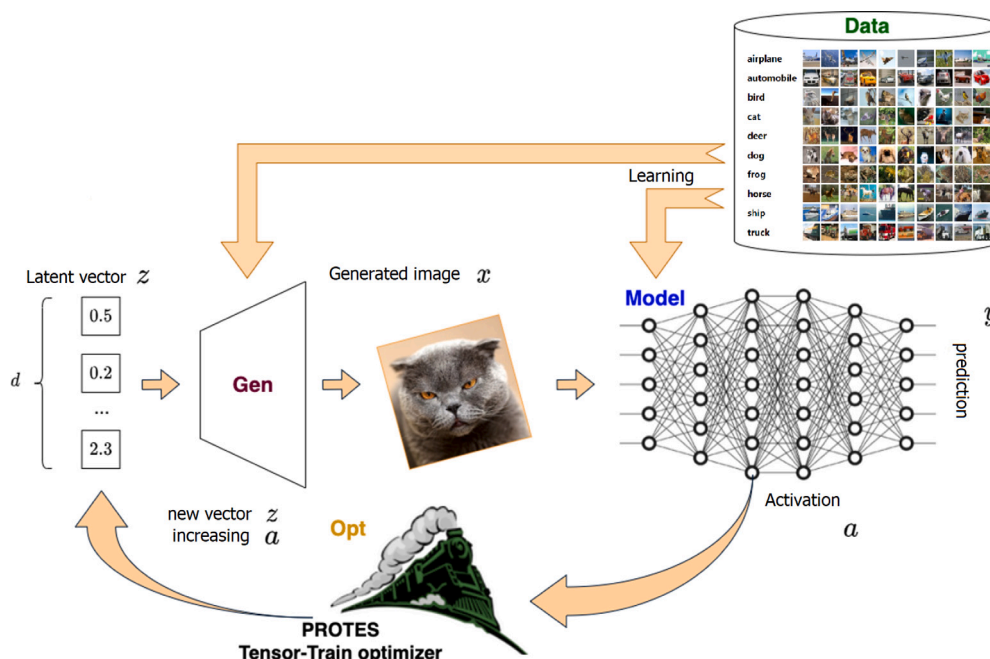


Fig. 3. Schematic of the MANGO framework.

5. Optimization methods: Random search, Nevergrad-based benchmarks [37] and Tensor Train decomposition-based methods (TTOpt, PROTES, etc.) [20,22]
6. backends: CPU (with multithreading), GPU, CuPy for SNN-based models (support provided by provided SpikingJelly authors [51])

The selected dataset is used to train a convolutional (in our case, also spiking) network for the image classification task, as well as to train a generator network for the task of compressing input data into an effective latent representation. After training, the generator creates random latent representations, which are transformed into images and presented to the convolutional network. The activation of the neuron of interest is measured and fed as input to the optimizer, which produces an improved latent vector that maximizes the likely response of the neuron studied. This process is repeated multiple times until convergence or until the optimizer query budget is exhausted.

MANGO code is available on GitHub <https://github.com/iabs-neuro/mango>, including training and analysis scripts. Full sets of MEIs will be provided to interested readers upon request.

4. Results

4.1. Performance of tensor train-based optimization methods

Optimization method performance was measured in terms of target neuron activation in response to generated MEI. Given a large number of iterations, approaches based on Tensor Train decomposition demonstrated performance slightly lower than the baseline Nevergrad Portfolio method, which was the most effective among the non-gradient methods considered, see Fig. 4 and Table 1. However, on low or intermediate number of iterations, which are the most relevant in real-world experiments, TT-based methods demonstrated comparable or better performance. All methods performed significantly better than the random search baseline.

At the same time the PROTES method in various modifications, has been shown to be 3–15 times faster than the Nevergrad Portfolio benchmark, depending on the number of iterations, hyperparameters and properties of the image. This substantial difference makes TT-based methods very competitive in a real experimental setting.

Three different (although related) tensor methods and a non-gradient benchmark from the Nevergrad library often converged on the same images. This is surprising, given the vast number of image variations sampled from the generator’s latent space. Despite the fact that the differences between MEIs increased in deeper layers (see Section 4.3), they were very similar in the early layers of the network (see Figs. A.13, B.17, B.18).

This suggests that all discrete optimization algorithms indeed converge to a good optimum in the latent space of the generator. Considering the results from Appendix B, we can extend these findings to the general image space and conclude that discrete optimization techniques, especially those based on Tensor Train decomposition, are effective in finding MEIs.

4.2. Emergence and dynamics of neuronal specializations

To study the dynamics of MEI emergence and their content we analyzed optimal stimuli for the Resnet18 SNN spiking layers. We were interested in the maximal neuronal activations achieved during the optimization process and their correspondence with objects in the training images. The training was conducted for 1.000 epochs and the network states were recorded after each epoch. Nine spiking layers were chosen, evenly spaced between the input and output layers, with 64 neurons in each layer selected for MEI construction. Neuronal activation was maximized for all 576 considered neurons sequentially by searching for the most exciting visual input using three variants of the PROTES method based on TT decomposition. SN-GAN was used as the generative model (we also used VQ-VAE image generator, the comparison between these two models is given in Appendix B).

An important aspect when analyzing MEIs is the relationship between identified neuronal specializations and the patterns the network has learned from the dataset. To study this, the resulting MEIs were then fed into the network to determine if the target neuron was specific to a particular image class. To quantify neuronal selectivity for specific classes, we have run the network on the MEIs and obtained the model class probabilities. Arguably, a more uniform class probability distribution corresponds to more abstract and general images. Conversely, a high probability for a specific class indicates the presence of that respective pattern in the MEI. For each MEI, we calculated the entropy

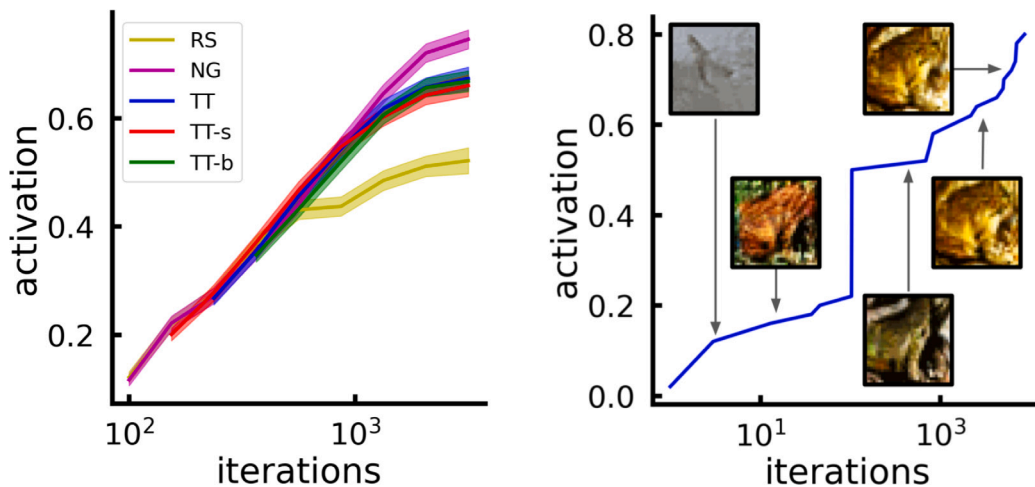


Fig. 4. Left: Neuron activations for different optimization methods, averaged for all units and SNN layers. Shaded regions represent standard errors. Right: process of a single MEI formation at different iterations with corresponding activation values.

Table 1

Average MEI-related activations and MEI computation time for various non-gradient optimization methods. Methods: RS — Random search; NG — Nevergrad baseline Portfolio; TT, TT-b, TT-s — Tensor Train based optimization variants.

Optimization methods performance						
	Act. (10^2 it.)	Act. (10^3 it.)	Act. (10^4 it.)	Time (10^2 it.)	Time (10^3 it.)	Time (10^4 it.)
RS	0.384 ± 0.001	0.468 ± 0.001	0.537 ± 0.001	10.202 ± 0.033	91.123 ± 0.334	861.266 ± 3.176
NG	0.393 ± 0.001	0.574 ± 0.001	0.755 ± 0.001	12.947 ± 0.036	163.45 ± 0.294	1685.464 ± 3.04
TT	0.411 ± 0.001	0.581 ± 0.001	0.675 ± 0.001	3.267 ± 0.006	14.456 ± 0.007	219.277 ± 0.289
TT-b	0.418 ± 0.001	0.579 ± 0.001	0.662 ± 0.001	3.321 ± 0.005	14.283 ± 0.026	383.852 ± 0.489
TT-s	0.392 ± 0.001	0.564 ± 0.001	0.679 ± 0.001	3.724 ± 0.005	15.236 ± 0.007	127.406 ± 0.173




Fig. 5. “Pink horse” neuron 52 from LIF layer 1.1 of spiking ResNet18. Images were generated using Tensor Train based methods. Numbers show activations of the target neuron on MEI.

of the final class probability distribution and normalized it against the maximum possible entropy (for a uniform distribution over 10 classes). Some other notable measures of neuronal selectivity in SNNs [80] exist as well.

The results are shown in Fig. 6. One can see a sharp peak of low-entropy MEIs in the last spiking layers of the trained model, which corresponds to high model confidence in assigning specific classes to MEIs. Even in the middle of the network, most neurons are already specifically associated with a certain class. In contrast, at the early stage of training, the model does not show any MEI specificity for a particular class.

Following the MEI formation process, we observed that it often consisted of two distinct stages. At first, the optimization was searching the latent image space almost randomly, with the network’s classification confidence (negative entropy of the output class probability distribution) being low. Then, the AM process converged to some MEI with further optimization only refining it. Transition from the first stage to the second one was marked by a sudden increase in network classification confidence especially on MEIs that resemble training images of clearly distinguishable classes (see Fig. 4).

Notably, the three TT-based methods used often converged to similar images. Given the high complexity of the underlying landscape of potential MEIs, we used the following composite criterion of neuron selectivity with respect to a specific class:

- “Stability”: All three maximization techniques converged on images of the same class (the MEI class was defined by the maximum probability at the network’s output).
- “Confidence”: The probability of the best class was at least 75%.
- “Activation”: The activity of the targeted neuron was at least 75% of the maximum possible. Here, activity was calculated by simply adding up the number of spikes across all channels for a given neuron.

As expected, the vast majority of specialized neurons were found in the final layers of the model, where MEIs contained the most complex patterns (Fig. 7, see Appendix A for characteristic examples of MEIs). Surprisingly, however, we also found a significant number of neurons related to a certain class in the intermediate layers, accounting for up to 40% of the total. Unexpectedly, highly selective neurons were found as early as the third spiking layer (Fig. 5).

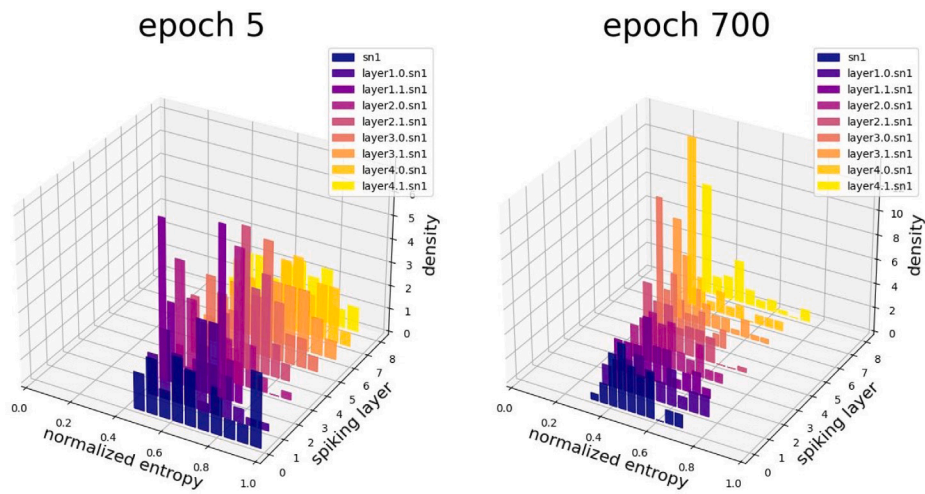


Fig. 6. Distribution of neurons by entropy of class prediction probabilities (lower entropy = more class-specialized). **Left:** early stage of learning (epoch #5), **right:** late stage of learning (epoch #700).




Fig. 7. Fraction of selective neurons. Horizontal axis — epoch number, vertical axis — layers (lower = deeper). **Left:** activation threshold — spikes $\geq 75\%$ of exposure time; **Right:** additionally, max class probability $\geq 75\%$, with max class stable over different maximization methods.

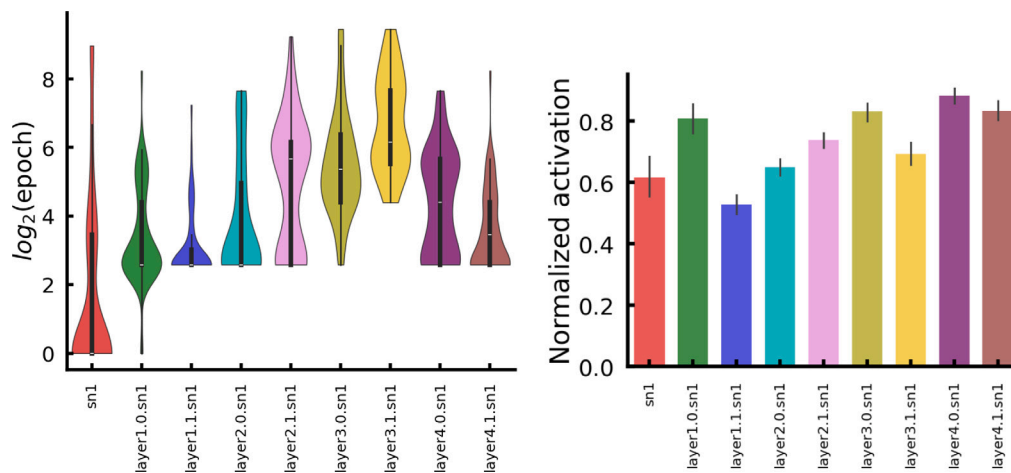


Fig. 8. Emergence of neuronal specializations. Horizontal axis — layer depth (increasing left to right). **Left:** Distribution (over several runs) of the first epoch numbers when a strong MEI appears (neuron spikes $\geq 75\%$ of exposure time). **Right:** Normalized activation by MEI (errorbars — 95% confidence interval).

As the model learns, highly specialized neurons increase in number. To quantify the “goodness” of MEIs, we calculated averaged neuronal response caused by MEIs across layers and learning epochs. The results are shown in Figs. 8 and 9. In general, MEI-related activity in the last layers increased with learning, correlating well with model performance (Fig. 9 right). For intermediate layers, MEI-related activity reaches plateau and may even show a slight decrease at late epochs. We discuss these effects and compare them to the ones taking place in a non-spiking ANN model in Appendix C. However, we observed an abnormally late development of MEIs before the final layers of the neural network (layer 3.1). These MEIs led to weaker activation

of targeted neurons (Fig. 8 right). We believe this decrease may be due to ongoing restructuring of MEIs in these layers, as they act as a “gateway” to more complex specialized layers. It is also possible that insufficient optimization budget prevented obtaining the “true” MEIs of these neurons.

We then addressed the issue of “labile” neurons, which change their selectivity as training progresses. To do this, we tracked neurons that met the criteria for specialization on more than one class during training. These neurons are interesting because they help understand the mechanisms of “narrow” specializations formation. These often specialize on abstract concepts that can be used as building blocks

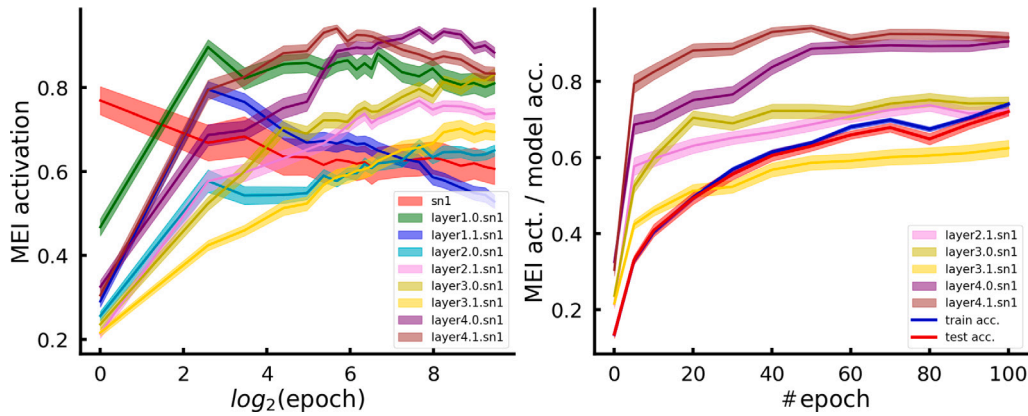


Fig. 9. Left: Dynamics of neurons' activations in response to their MEIs. Right: the same for the first 100 epochs with the model classification accuracy in the same axes.

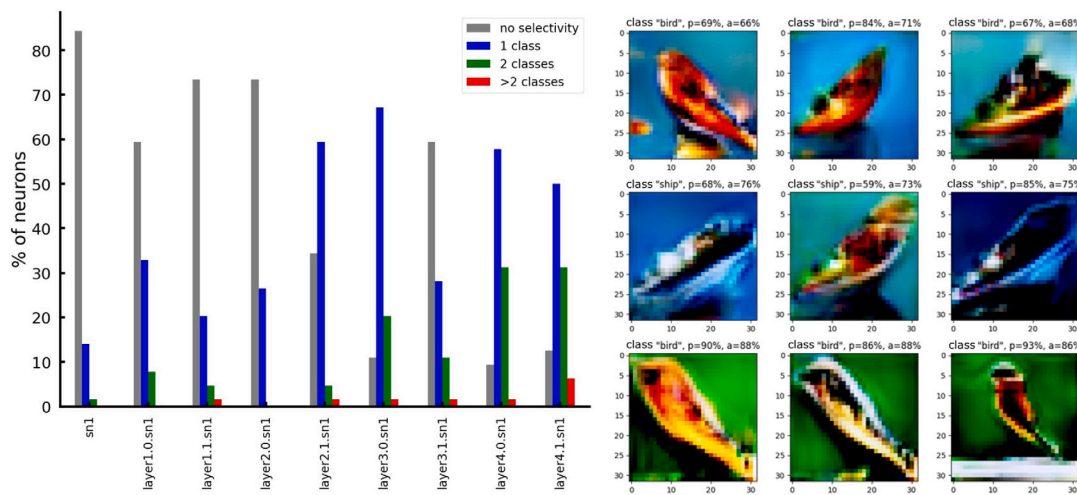


Fig. 10. Left: fractions of non-specialized, stable (selective to only 1 class) and labile (selective to 2 or more) neurons vs. layer depth. Right: oscillating MEIs of labile “bird-ship” neuron, rows: epoch # 40, 100, 300 of training; columns: results obtained with TT, TT-S, TT-B optimization methods. The captions above each MEI provide information about the predominant class according to classification results, the probability, and the activation of the neuron as a percentage of the maximum possible.

for classifying more complex patterns. An example of MEI of such a neuron can be seen on Fig. 7, right. It selectively activates on elongated, diagonally arranged structures on a blue/green background. In this area of the latent space, the embeddings of birds and ships are similar, which is reflected in the neuron’s classification at different stages of training. Such labile neurons were virtually non-existent in the initial layers of the network (see Fig. 7, left), but their fraction rose to approximately 40% in the final layer. This could indicate a high degree of variability in the specialization of the final layers and their capacity for rapid reconfiguration when receiving new training data, drawing on information from the neurons in earlier layers. Specifically, this supports the empirical principle behind reservoir computing, which is to train only the output layers while allowing the internal layers to identify complex features in the data set [81].

4.3. Complexity and diversity of specializations

Based on our observations, the entropy of class probability predictions by the model on MEIs decreases with layer depth and training epochs number, indicating that the neurons are gradually forming more refined class-related specializations.

In general, as expected, neurons of early layers are maximally activated by simple geometric features or color patterns, while neurons in deeper layers show selective activation for one class or another (see Figs. A.13, A.14, A.15 for examples).

MEIs generated using different optimization methods contain the same patterns of varying degrees of complexity, but are not identical to each other. In order to evaluate the variability of the resulting images, we calculated the average Euclidean and cosine distances between the latent vectors corresponding to MEIs (SN-GAN generator was used).

The variability of possible images in the last layers of spiking ResNet18 turned out to be greater than in the first ones (see Fig. 11). Average distances between the generated MEIs in the last spiking layer turned out to be about 20% higher than in the first one, revealing significant transformations of neurons’ specializations landscape. Possible explanation of this effect is that there are many ways in which a complex concept, such as a horse, can be realized in an image. NNs learn to nonlinearly project the data space during training, simultaneously identifying some of its areas with each other [82]. Thus, the number of local maxima in the overall image space grows with increasing specialization of neurons. At the same time, this effect measured via cosine distance became weaker during model training (see Fig. 11, right).

The growing diversity of MEIs is accompanied by an increase in the activations they produce on corresponding neurons (Fig. 8). This may be explained by finer specialization of neurons in deeper layers, in accordance with the results from Section 4.2. At the same time, the dips in maximal activation values observed in layer 3.1 are reproduced by both generator types (GAN and VAE) and show that the overall picture may be more complex.

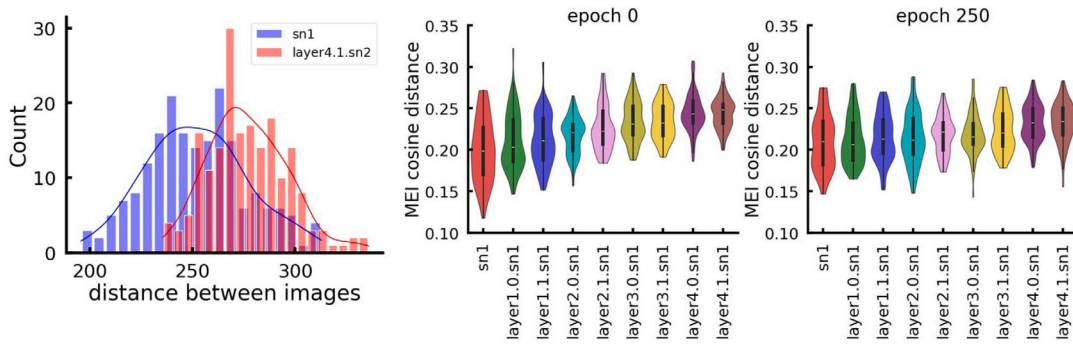


Fig. 11. Diversity of specializations in terms of distances between MEIs in the latent space. **Left:** Euclidean distances. **Right:** cosine similarities for different layers (layer depth increases to the right), at different training stages.




Fig. 12. Complexity of MEIs, measured as compressed file size (JPEG and ZIP compressions).

As the network learns, the content of MEIs becomes more complex. Calculating an exact measure of image complexity is a challenging task, for which several complementary approaches have been developed [83]. However, it is possible to gain a general understanding of this measure by using simple heuristics based on information compression algorithms. This approach has proven successful in assessing entropy and mutual information [84] as well as the quantitative assessment of consciousness level [85]. To quantify visual complexity of generated MEIs, we calculated the compression ratios of the optimized images in the .jpg and .zip formats and compared them to the size of the original images (see Fig. 12). Both compression methods yielded similar results: as the network learns, the complexity of MEIs in the first layers increases, gradually reaching a steady state. We believe that this is due to rapid formation of more complex specializations in the early layers of the network (Fig. 8 see Section 5.1 in Discussion).

5. Discussion

5.1. Neuronal selectivity in spiking convolutional networks

We track neuronal selectivity properties throughout the training process of the model and observe MEIs of varying complexity with more refined features corresponding to deeper layers, in accordance with the principles shown in both ANNs [31,77,86] and animal visual cortex [3,87].

Neurons from intermediate layers are often selective for ambiguous images potentially useful for multiple class discrimination (see Fig. A.14 (left) for examples). This effect can provide a potential basis

for complex image specializations. For instance, it has been shown that abstract feature representations emerge naturally in neural networks trained to perform multiple tasks [88].

The influence of different sensory modalities on the visual cortex in animals is not negligible, potentially leading to complex neuron specializations [89]. For example, neurons in the V2 visual area can be selective for “naturalistic” texture stimuli [90]. Neurons selective for complex vocal stimuli have also been found in the primary auditory cortex [91]. These findings coincide with the formation of complex specializations in early layers, as demonstrated in this work (see Fig. 7).

An important research question is the patterns and mechanisms of neuronal specializations formation over time. It is notable that neuronal specializations in our setting form early during the training process, as shown in Fig. 7 — this phenomenon has also been observed in non-spiking ANNs [92]. The model’s neurons are able to form complex representations, expressed in the activity of neurons, after several exposures to the dataset. This occurs both in deep and intermediate layers, even though the recognition quality has not yet reached a plateau at this stage of learning. This is consistent with data on the early formation of specialized neurons in biological neural networks, as obtained in experiments [93]. This behavior may indicate that complex neural representations are rapidly formed, and the network then “rebuilds” the weights over time to optimally convey information from these highly specialized neurons to the output layers.

The results obtained suggest that the reconfiguration of neural specialization is a fundamental process in both biological and artificial neural networks. Additionally, heterogeneity in the selectivity of

neurons is observed in these systems, with some neurons retaining their specialization as they learn, while others' specializations are more labile. In biological networks, it was shown that neurons with stronger selectivity tend to be more stable, supporting the hypothesis of the log-dynamic brain [94]. In future work, we plan to investigate how this behavior relates to the strength of specialization, as well as investigate the weights of connections between artificial neurons with specific specializations in order to gain better understanding of their influence on one another and trace the development of complex specializations.

5.2. Limitations of the current work and directions of future research

As our results show, the quality of the calculated MEIs depends significantly on the image generator used. In the future, we plan to extend MANGO by adding support for new generators that are well suited for the MEI search task (i.e. combining the discrete structure of the latent space with the high plausibility of the generated images). It is of interest to consider the popular direction of discrete GANs [95].

In parallel with single-neuron selectivity, neuronal computations can be described by low-dimensional activity manifolds, which arise from the collective, coordinated spiking of numerous neurons. This type of information encoding in the brain is known as population coding and has a long and rich history in neuroscience [96,97]. Recently, population behavior of artificial neurons in large language models (LLMs) has received particular attention, potentially leading to the emergence of the new field of representation engineering [98]. To our knowledge, the relationships between optimal stimuli for single neurons and their contribution to population coding have not yet been analyzed, and we consider this to be one of the important directions for future work.

Recently, it has become apparent that the activity of individual components can significantly influence the operation of a large neural network. Researchers working with LLMs have recently demonstrated that manipulating the activity of a single specialized artificial neuron can lead to significant changes in the behavior of an LLM [99]. At the same time, it is known that stimulating a small number of neurons (or even a single neuron in some cases) in a biological neural network can cause the entire ensemble to become highly active, potentially influencing animal behavior [100]. Given this evidence of the importance of individual computational units, it would be interesting to investigate the effect of inhibiting or exciting a single neuron on model performance and correlate this with the MEI properties of that neuron. We plan to conduct such analysis in future work.

5.3. Activation maximization in ANNs & explainable AI

The problem of maximizing neuronal activation in artificial neural networks (ANNs) is being actively researched within the field of ANN visualization & interpretation [31–34], a sub-field of the more general area of explainable AI. In order to successfully integrate ANN-based solutions into critical systems, such as medical and legal practices, it is necessary to have a human-understandable interpretation of the algorithm's decision-making process. In recent years, methods based on visualizing the computation graph, loss function, parameter space of certain layers, or even individual neurons [101] have gained popularity in creating interpretable deep learning models. These methods allow researchers to better understand the inner workings of NNs.

A recent review [102] of visualization techniques for unraveling neuronal stimuli covers some of the most successful approaches: the method of activation maximization [31], the Grad-CAM method [103], layer-wise relevance propagation (LRP) [34] technique, well as the Integrated Gradients method [104]. Another outstanding recent work is Goh et al. [14], which demonstrates a rich range of stimuli that maximize the activation of different neurons in an ANN. This work also demonstrates the effect of multimodal selectivity in artificial neurons.

However, all of the aforementioned methods are mostly based on automatic differentiation: the ANN's computation graph is formed allowing to compute gradients with minimal computational complexity. This allows the use of gradient descent (ascent), and its stochastic modifications. Gradient-based methods are clearly not applicable to living systems, so one must either use gradient-free methods, such as genetic algorithms (as in [38]), or other methods that use local approximations of gradients, such as evolutionary algorithms, or alternative approaches, including optimization techniques based on TT-decomposition.

Our results presented in Section 4.1 indicate that TT-based optimization methods outperform Nevergrad's [37] Portfolio in our problem setting (for small and moderate optimization budget). Portfolio is a balanced mixture of optimization methods that includes CMA-ES [41]. It was shown [18] that CMA-ES outperforms other genetic algorithms on the problem of activation maximization. This allows one to carefully conclude that TT-based optimization methods are more suitable for AM problems, although this claim requires further accurate examination.

5.4. Possible applications for TT-based activation maximization in biological neurons

Modern optical imaging techniques allow us to monitor the activity of hundreds of neurons in their temporal dynamics *in vivo*. Activation maximization of multiple neurons, both individually and in groups, could be a powerful tool for analyzing the function of the nervous system at the cellular level, shedding light on the formation and spread of cognitive specializations in biological neurons. However, the design of biological experiments poses serious limitations on the applicability of optimization methods. These methods need to be quick, accurate, and able to tune themselves to new data. We expect our TT-decomposition-based methods to possess all these properties.

Thus, one of the most interesting and potentially fruitful applications of our tools is maximizing neuronal activation *in vivo*. Since biological neurons have an irreducible minimum response time to a stimulus, research into algorithms aimed at reducing optimizer budget and lowering computation time required for MEI generation seems particularly relevant. In particular, we plan to extend MANGO with more optimization methods based on Tensor Train decomposition, specifically optimized for this task.

A separate challenge is finding optimal stimuli in non-visual domains, such as the space of sounds or smells, which may be far more relevant to the animal under study than to humans [105]. We also hope to gain some insight into the applicability of TT-decomposition based methods in these latent stimulus spaces.

Studies of neuronal selectivity in spiking neural networks are of great interest, as they reveal the mechanisms of the functional units' operation in the neural network using a biologically plausible model. The use of spiking artificial neurons makes the analysis in this study significantly closer to real electrophysiological experiments [4] and allows us to test hypotheses about the temporal patterns of neurons' responses to MEIs *in silico*. We hope that our algorithms will be useful in the future study of optimal stimuli for biological neural networks.

Our framework aims to bridge the gap between neurobiological and computational experiments, enabling the possibility to explore the principles of information encoding in deep neural networks. Additionally, we hope that this research will bring activation maximization algorithms closer to the requirements of real biological experiments, providing a new tool for analyzing cognitive functions of individual neurons.

CRedit authorship contribution statement

Nikita Pospelov: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Conceptualization. **Andrei Chertkov:** Visualization, Validation, Software, Methodology, Investigation, Data curation. **Maxim Beketov:** Writing – review & editing, Writing – original draft, Software, Resources, Methodology, Investigation. **Ivan Oseledets:** Writing – review & editing, Supervision, Resources, Methodology, Conceptualization. **Konstantin Anokhin:** Writing – review & editing, Validation, Supervision, Resources, Methodology, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the Non-Commercial Foundation for the Support of Science and Education “INTELLECT” and Lomonosov Moscow State University. N. Pospelov acknowledges the support of the Brain Program of the IDEAS Research Center. The work of M. Beketov has been supported by HSE University Basic Research Program. The main computational experiments were conducted in the Institute for Advanced Brain Studies. Some experiments used the computational resources of the HPC facilities at HSE University [106]. The authors are grateful to all members of the Laboratory of Neuronal Intelligence and especially to Ksenia Toropova and Olga Ivashkina for fruitful discussions.

Appendix A. Gallery of notable MEIs

In this section we present some notable MEIs for neurons of different (spiking) layers and attempt to interpret their specializations.

Fig. A.13 shows MEIs of neurons of the first spiking layer (sn1). Despite the spiking nature of neurons, these clearly specialize on visual primitives – patterns/textures, uniform monochrome “blobs” – just like non-spiking neurons in feedforward ANNs [77,86]. See Figs. B.17, B.18 in Appendix B for more examples of first layer neurons’ MEIs. MEIs get more interesting with layer depth.

Fig. A.14 shows MEIs of different neurons of an moderately deep layer. In the left 3×3 block, first row neuron specializes on elongated shapes on blue background: MEIs are close to typical images of classes 8 – ship and 0 – airplane. Second row neuron specializes on green bottom and blue top background (grass and skies) with objects in between (automobiles/trucks/airplanes on runway strips). Third row neuron is possibly the most interesting, as it seems to specialize on images sharing features of birds and ships (green freshwater) – we call it the bird-ship neuron. Also see Section 4.2 and Fig. 10.

In the right 3×3 block, we demonstrate 3 different neurons with very similar specializations. Clearly the MEIs share the features of dataset classes 7 – horse, 4 – deer (possibly also 5 – dog), but were often found to be almost mirror-symmetric, so it seems these neurons specialize on such “brown tetrapods” on green (grass) background without much knowledge of head vs. tail placement. Overall it is expected that moderately deep neurons specialize on such simple image types – more complicated than visual primitives, but less complicated than detailed images clearly distinguishable as dataset classes. What is notable however is that such smooth but not yet very class-specific MEIs are stably obtained with a gradient-free optimization procedure, not by simply back-propagating as in feedforward ANNs [16].

Fig. A.15 shows MEIs of the last (deepest) spiking layer (layer4.1.sn1). Its neurons clearly specialize on dataset-like images, with classes 4 – deer, 6 – frog, 7 – horse which are clearly distinguishable.

Appendix B. Generative models comparison

B.1. CIFAR-10 dataset

The models were trained on CIFAR-10 [64] – a dataset of 60.000 color images of size 32×32 pixels, split into 10 classes (6.000 images per class): (0) airplane, (1) automobile, (2) bird, (3) cat, (4) deer, (5) dog, (6) frog, (7) horse, (8) ship, (9) truck. Fig. B.16 provides examples of dataset images.

B.2. Overview of VAEs

Apart from SN-GANs, we have also tested VQ-VAEs for image generation. A Vector-Quantized Variational AutoEncoder (VQ-VAEs) [67] is a generative model of the VAE family [66]. VAEs were chosen for our task mainly due being quick to generate data – only one forward pass is needed to obtain a sample; but also due to simplicity of architecture – encoder–decoder, where the endoder gives the latent coordinates of a sample.

VAEs take a sample of data, x (e.g. a batch of images) and transform it, layer by layer of the encoder (if data is images, some convolutional layers are applied first), to the latent space. On that space, a family of parameterized probability distributions (posteriors) $p_{\theta}(z|x)$ is defined (one can say that the encoder just outputs a vector of parameters θ of this distribution, thus specifying a representative of this family). One then generates a sample $z \sim p_{\theta}(z|x)$ from that distribution – this is a non-deterministic action and so at first it seems non-differentiable for one to design a backward pass (error backpropagation), but a clever idea of the reparameterization trick [66] overcomes this problem (by separating parameterization of a distribution from sampling from it). This “latent code” z that describes the input sample x is then passed to the decoder (if the VAE has to generate images – last layers of the decoder will be deconvolutional, etc.) to produce an output – a new sample of data. There is also what is called a prior distribution $p(z)$ of latent codes, which, for the case of latent space being \mathbb{R}^d , is uninformatively chosen to be standard normal (uniform distribution cannot be supported on \mathbb{R}^d). There exists a problem of “posterior collapse” – when the posterior distribution gets too close to the uninformative prior on some latent codes. Various modifications of VAEs aim to fix this problem, with VQ-VAE doing so quite successfully under certain conditions.

Vector-Quantized VAEs [67] take this idea to a “discrete” setting – the family of posterior distributions $p_{\theta}(z|x)$ is now discrete rather than continuous (multidimensional normal in original VAEs), supported on a fixed-size dictionary of latent codes (codewords). The motivation behind this is that, for any finite sample of natural images, there will only be a discrete set of classes in it – dogs, cats, cars, etc. So VQ-VAEs tend to produce far sharper images than their original VAE counterparts (which tend to blur the image, with the uncertainty of the posterior distribution translating into Gaussian-like, albeit nonlinearly transformed, blur on the resulting images), which at first seemed better for our needs. The encoder of VQ-VAE thus learns to output these discrete posteriors (the latent codes and their posterior probabilities) so that samples from them (output of the decoder) were alike the training data. Since the prior $p(z)$ is also supported on this discrete codebase (which can support a uniform distribution), choosing it to be uniform helps prevent posterior collapse.

In our experiments, we have tested VQ-VAEs with latent dimension 64 with dictionary size of 512 – while CIFAR-10 only has 10 classes, codes are assigned not to singular images, but to their batches (samples). Both encoder and decoder were built with ResidualStacks [52] of 3–4 convolutional layers.




Fig. A.13. MEIs of first spiking (sn1) layer neurons. Each row is a certain neuron, several activation optima.



Fig. A.14. MEIs of intermediary (layer2.1.sn1) layer neurons. Each rows of 3 images: one neuron, several activation optima. Left 3 × 3 block: 3 different neurons with different specializations. Right 3 × 3 block: 3 different neurons with similar specializations.

B.3. Comparison of GAN- and VAE- based generators

For MEIs created for each neuron of all spiking layers from the spiking ResNet18, analysis has shown that the maximal activations of neurons on images generated by VQ-VAE were smaller than on images generated by SN-GAN. The magnitude of the gap depended on the layer and ranged from 50% in early layers to 5% in deep layers.

In general, despite neuronal activations were significantly above chance, the MEIs generated by VAE often contain no visible structure and do not provide any insight about neuronal specializations.

The exception is the early layers, in which the MEIs obtained using both generators are quite similar. This allows one to expect that discrete optimization methods actually find generator-agnostic maxima in the global image space, and difficulties in obtaining interpretable images can be solved by further improving the structure of generators' latent space.

Thus, VQ-VAE-generated images are closest to richer SN-GAN-generated ones where there are uniform colors or simple geometric patterns in the image (see Figs. B.17, B.18). We suppose that this



Fig. A.15. MEIs of last spiking layer (layer4.1.sn1). Each row is a certain neuron, several activation optima.

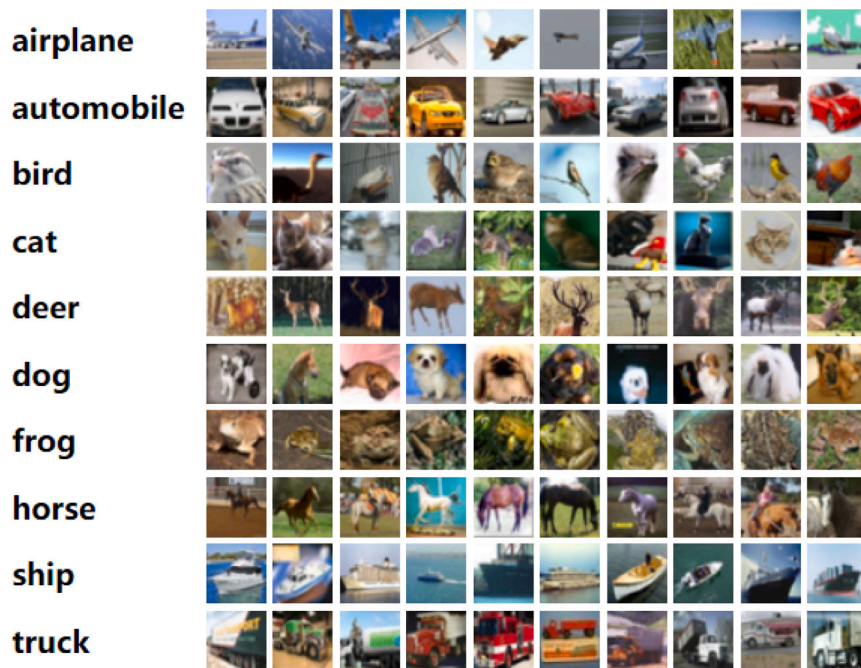


Fig. B.16. 10 classes of images in CIFAR10 dataset [64].

is a sign of excessive information compression in the latent space of VQ-VAE generator, which leads to the loss of significant image details.

Our hypothesis as for why VQ-VAE is performing poorer in our setup is: as described in the above sections, VQ-VAEs learn discretized, rather than continuous posterior distributions of data samples — supported on a finite set (dictionary) of (learned) vectors in the latent space, the

posterior distribution then being multinomial on that discrete pointset. For that reason, VQ-VAEs seem to be not so good at interpolating between probable datapoints (images) — designed to mitigate posterior collapse and Gaussian-like blur of generated images in original VAEs, it seems that VQ-VAEs loose in the ability to generate interesting out-of-distribution samples. Of course, some interpolation is possible, and it




Fig. B.17. MEIs for unit 16 of the first spiking layer from spiking ResNet18. Top: images generated with SN-GAN. Bottom: images generated with VQ-VAE. Columns correspond to different TT-based optimization methods.



Fig. B.18. MEIs for unit 35 of the first spiking layer from spiking ResNet18. Top: images generated with SN-GAN. Bottom: images generated with VQ-VAE. Columns correspond to different TT-based optimization methods.

seems we did not find good balance between discretization parameters of VQ-VAE and TT — that is matter of future investigation.

SN-GANs performed far better in our experiment, providing seemingly smoother (due to benefits of Spectral Normalization) coordinates on the latent space of images. However, the above results on neuron specialization distributions seem to be universal — no matter what generative model was used.

Appendix C. Comparison of optimal stimuli dynamics in ANN and SNN

To investigate to what extent the results obtained are determined by the spiking nature of the network under investigation, we repeated our analysis on a standard (non-spiking) ANN of ResNet18 architecture. Recall that the SNN we analyzed was essentially a spiking layer-augmented version of this model. We applied TT-based optimization

methods to reveal optimal stimuli for ResNet18 block neurons (8 layers total) during learning, and then applied the same analysis as outlined in the Results section, including MEI tracking throughout learning, to recorded data. For this analysis, we removed the “sufficient activity” selectivity criterion since the output magnitude of non-spiking neurons is not constrained from above. Our results have shown significant differences between the MEI formation processes in non-spiking convolutional networks and their spiking counterparts.

- In an SNN, the proportion of class-selective neurons grows or remains intact during learning in all layers, while in an ANN, this process is clearly two-staged. Initial growth of the proportion of selective units is followed by a decrease in all layers except for the last one, where it monotonically increases during training (see Fig. C.19). It seems that the selectivity “flows” into the last layer,

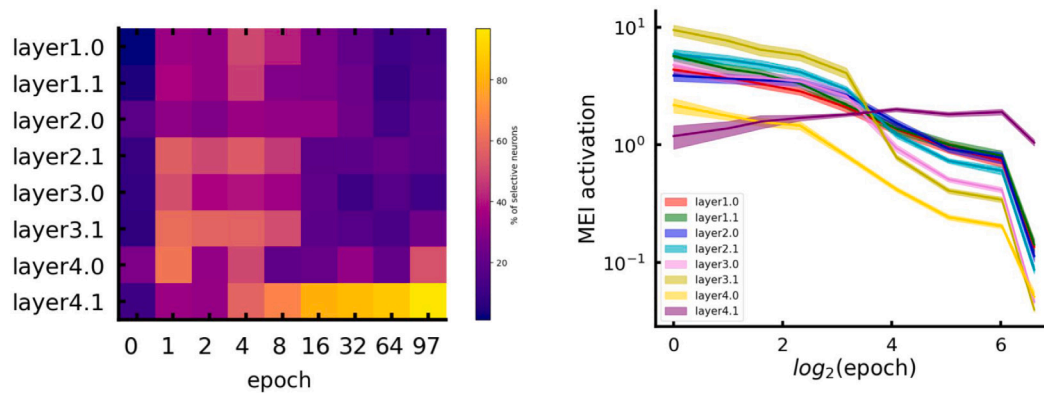


Fig. C.19. **Left:** fraction of class-selective neurons in an ANN model. Horizontal axis — epoch number, vertical axis — layers (lower = deeper). **Right:** Unbounded MEI-related activations in an ANN model per layer during training. Shaded regions indicate standard errors.

and the optimal stimuli in intermediate layers become less refined and class-specific.

- In general, the model's confidence (measured by the neg-entropy of the outputs) for generated MEIs was significantly lower in the ANN version of ResNet18 than in the spiking one. The only layer with stable and reliable class-related MEIs was the last one, especially during the late epochs of training.

Given low network confidence during inference on generated MEIs for a non-spiking ResNet18, it is possible that the TT-based optimization method may not perform well in this task. Further analysis using gradient-based optimization methods, which are commonly used for ANNs, could help one better understand this issue. However, these results can have an intriguing explanation in the context of general learning dynamics of neural network learning. It has been shown that forgetting (decay of information contained in the weights) is critical to achieving invariance and disentanglement in representation learning [92]. This can explain the decrease in selectivity and activation amplitudes in intermediate layers. At the same time, the computational role of the last layers increases, making them disproportionately important for the model's decision making while the rest of the network serves as a high-dimensional “feature extractor” - the idea lying at the heart of reservoir computing systems [107].

These processes also seem to occur in a spiking network. Note that the maximum MEI-related activations of the first SNN layers occur at intermediate stages of learning (see Fig. 9). However, the spiking nature of neurons makes it more challenging for the network to learn, extending the critical phases of this process. This suggests that one may observe effects similar to those taking place in ANNs in SNNs if one trains models for a long enough time.

Data availability

Data will be made available on request.

References

- [1] D.H. Hubel, T.N. Wiesel, *Brain and Visual Perception: The Story of a 25-Year Collaboration*, Oxford University Press, 2004.
- [2] J. Vierock, S. Rodriguez-Rozada, A. Dieter, F. Pieper, R. Sims, F. Tenedini, A.C. Bergs, I. Bendifallah, F. Zhou, N. Zeitzschel, et al., BiPOLES is an optogenetic tool developed for bidirectional dual-color control of neurons, *Nature Commun.* 12 (1) (2021) 4527.
- [3] R. Quian Quiroga, M. Boscaglia, J. Jonas, H.G. Rey, X. Yan, L. Maillard, S. Colnat-Coulbois, L. Koessler, B. Rossion, Single neuron responses underlying face recognition in the human midfusiform face-selective cortex, *Nature Commun.* 14 (1) (2023) <http://dx.doi.org/10.1038/s41467-023-41323-5>.
- [4] C.R. Ponce, W. Xiao, P.F. Schade, T.S. Hartmann, G. Kreiman, M.S. Livingstone, Evolving images for visual neurons using a deep generative network reveals coding principles and neuronal preferences, *Cell* 177 (4) (2019) 999–1009.
- [5] A. Bardou, W. Xiao, C.R. Ponce, M.S. Livingstone, G. Kreiman, Face neurons encode nonsemantic features, *Proc. Natl. Acad. Sci.* 119 (16) (2022) e2118705119.
- [6] W.R. Ashby, *An introduction to cybernetics*, 1956.
- [7] S.J. Russell, P. Norvig, *Artificial Intelligence: A Modern Approach*, Pearson, 2016.
- [8] R.S. Sutton, *Reinforcement learning: An introduction*, Bradford Book (2018).
- [9] A.J. Shriver, T.M. John, *Neuroethics and animals: report and recommendations from the University of Pennsylvania Animal Research Neuroethics Workshop*, *ILAR J.* 60 (3) (2019) 424–433.
- [10] P. Singer, Y.F. Tse, AI ethics: the case for including animals, *AI Ethics* 3 (2) (2023) 539–551.
- [11] A.M. Zador, A critique of pure learning and what artificial neural networks can learn from animal brains, *Nat. Commun.* 10 (1) (2019) 3770.
- [12] R. Schaeffer, M. Khona, I. Fiete, No free lunch from deep learning in neuroscience: A case study through models of the entorhinal-hippocampal circuit, *Adv. Neural Inf. Process. Syst.* 35 (2022) 16052–16067.
- [13] K. Dobs, J. Martinez, A.J. Kell, N. Kanwisher, Brain-like functional specialization emerges spontaneously in deep neural networks, *Sci. Adv.* 8 (11) (2022) eab18913.
- [14] G. Goh, N. Cammarata, C. Voss, S. Carter, M. Petrov, L. Schubert, A. Radford, C. Olah, Multimodal neurons in artificial neural networks, *Distill* 6 (3) (2021) e30.
- [15] D. Bau, B. Zhou, A. Khosla, A. Oliva, A. Torralba, Network dissection: Quantifying interpretability of deep visual representations, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 6541–6549.
- [16] C. Olah, A. Mordvintsev, L. Schubert, Feature visualization, *Distill* (2017) <http://dx.doi.org/10.23915/distill.00007>, <https://distill.pub/2017/feature-visualization>.
- [17] A. Nguyen, J. Yosinski, J. Clune, Understanding neural networks via feature visualization: A survey, in: *Lecture Notes in Computer Science*, Springer International Publishing, 2019, pp. 55–76, http://dx.doi.org/10.1007/978-3-030-28954-6_4.
- [18] B. Wang, C.R. Ponce, High-performance evolutionary algorithms for online neuron control, in: *Proceedings of the Genetic and Evolutionary Computation Conference*, 2022, pp. 1308–1316.
- [19] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, *Adv. Neural Inf. Process. Syst.* 27 (2014).
- [20] K. Sozykin, A. Chertkov, R. Schutski, A.-H. Phan, A.S. Cichocki, I. Oseledets, TTOpt: A maximum volume quantized tensor train-based optimization and its application to reinforcement learning, *Adv. Neural Inf. Process. Syst.* 35 (2022) 26052–26065.
- [21] I.V. Oseledets, Tensor-train decomposition, *SIAM J. Sci. Comput.* 33 (5) (2011) 2295–2317.
- [22] A. Batsheva, A. Chertkov, G. Ryzhakov, I. Oseledets, PROTES: Probabilistic optimization with tensor sampling, *Adv. Neural Inf. Process. Syst.* (2023).
- [23] A. Chertkov, G. Ryzhakov, G. Novikov, I. Oseledets, Tensor extrema estimation via sampling: A new approach for determining min/max elements, *Comput. Sci. Eng.* (2023).
- [24] A. Taherkhani, A. Belatreche, Y. Li, G. Cosma, L.P. Maguire, T.M. McGinnity, A review of learning in biologically plausible spiking neural networks, *Neural Netw.* 122 (2020) 253–272.
- [25] K. Yamazaki, V.-K. Vo-Ho, D. Bulsara, N. Le, Spiking neural networks and their applications: A review, *Brain Sci.* 12 (7) (2022) 863.
- [26] S. Panzeri, R.S. Petersen, S.R. Schultz, M. Lebedev, M.E. Diamond, The role of spike timing in the coding of stimulus location in rat somatosensory cortex, *Neuron* 29 (3) (2001) 769–777, [http://dx.doi.org/10.1016/s0896-6273\(01\)00251-3](http://dx.doi.org/10.1016/s0896-6273(01)00251-3).

- [27] Y. Andrade-Talavera, A. Fisahn, A. Rodríguez-Moreno, Timing to be precise? An overview of spike timing-dependent plasticity, brain rhythmicity, and glial cells interplay within neuronal circuits, *Mol. Psychiatry* 28 (6) (2023) 2177–2188, <http://dx.doi.org/10.1038/s41380-023-02027-w>.
- [28] Y. Li, Y. Kim, H. Park, P. Panda, Uncovering the representation of spiking neural networks trained with surrogate gradient, 2023, arXiv preprint arXiv:2304.13098.
- [29] E.Y. Walker, F.H. Sinz, E. Cobos, T. Muhammad, E. Froudarakis, P.G. Fahey, A.S. Ecker, J. Reimer, X. Pitkow, A.S. Tolias, Inception loops discover what excites neurons most using deep predictive models, *Nature Neurosci.* 22 (12) (2019) 2060–2065.
- [30] Y. Kim, P. Panda, Visual explanations from spiking neural networks using inter-spike intervals, *Sci. Rep.* 11 (1) (2021) 19037.
- [31] D. Erhan, Y. Bengio, A. Courville, P. Vincent, Visualizing higher-layer features of a deep network, *Univ. Montr.* 1341 (3) (2009) 1.
- [32] J. Yosinski, J. Clune, A. Nguyen, T. Fuchs, N. Lipson, Understanding neural networks through deep visualization, 2015, arXiv preprint arXiv:1506.06579.
- [33] W. Samek, A. Binder, G. Montavon, S. Lapuschkin, K.-R. Müller, Evaluating the visualization of what a deep neural network has learned, *IEEE Trans. Neural Netw. Learn. Syst.* 28 (11) (2016) 2660–2673.
- [34] G. Montavon, W. Samek, K.-R. Müller, Methods for interpreting and understanding deep neural networks, *Digit. Signal Process.* 73 (2018) 1–15.
- [35] A. Nguyen, A. Dosovitskiy, J. Yosinski, T. Brox, J. Clune, Synthesizing the preferred inputs for neurons in neural networks via deep generator networks, *Adv. Neural Inf. Process. Syst.* 29 (2016).
- [36] C. Olah, A. Satyanarayan, I. Johnson, S. Carter, L. Schubert, K. Ye, A. Mordvintsev, The building blocks of interpretability, *Distill* 3 (3) (2018) e10.
- [37] J. Rapin, O. Teytaud, Nevergrad - A gradient-free optimization platform, 2018, <https://GitHub.com/FacebookResearch/Nevergrad>.
- [38] W. Xiao, G. Kreiman, XDream: Finding preferred stimuli for visual neurons using generative networks and gradient-free optimization, *PLoS Comput. Biol.* 16 (6) (2020) e1007973.
- [39] S. Ruder, An overview of gradient descent optimization algorithms, 2016, <http://dx.doi.org/10.48550/ARXIV.1609.04747>, URL: <https://arxiv.org/abs/1609.04747>.
- [40] W. Xiao, G. Kreiman, Gradient-free activation maximization for identifying effective stimuli, 2019, arXiv preprint arXiv:1905.00378.
- [41] N. Hansen, A. Ostermeier, Adapting arbitrary normal mutation distributions in evolution strategies: The covariance matrix adaptation, in: *Proceedings of IEEE International Conference on Evolutionary Computation, IEEE*, 1996, pp. 312–317.
- [42] Y. LeCun, Y. Bengio, G. Hinton, Deep learning, *Nature* 521 (7553) (2015) 436–444.
- [43] I. Goodfellow, Y. Bengio, A. Courville, *Deep Learning*, MIT Press, 2016.
- [44] E.M. Izhikevich, Simple model of spiking neurons, *IEEE Trans. Neural Netw.* 14 (6) (2003) 1569–1572.
- [45] W.S. McCulloch, W. Pitts, A logical calculus of the ideas immanent in nervous activity, *Bull. Math. Biophys.* 5 (1943) 115–133.
- [46] F. Rosenblatt, *The Perceptron, A Perceiving and Recognizing Automaton Project Para*, Cornell Aeronautical Laboratory, 1957.
- [47] F. Rosenblatt, The perceptron: a probabilistic model for information storage and organization in the brain, *Psychol. Rev.* 65 (6) (1958) 386.
- [48] D.E. Rumelhart, G.E. Hinton, R.J. Williams, Learning representations by back-propagating errors, *Nature* 323 (6088) (1986) 533–536.
- [49] W. Maass, Networks of spiking neurons: the third generation of neural network models, *Neural Net.* 10 (9) (1997) 1659–1671.
- [50] J.K. Eshraghian, M. Ward, E. Neftci, X. Wang, G. Lenz, G. Dwivedi, M. Bannamoun, D.S. Jeong, W.D. Lu, Training spiking neural networks using lessons from deep learning, *Proc. IEEE* 111 (9) (2023) 1016–1054.
- [51] W. Fang, Y. Chen, J. Ding, Z. Yu, T. Masquelier, D. Chen, L. Huang, H. Zhou, G. Li, Y. Tian, SpikingJelly: An open-source machine learning infrastructure platform for spike-based intelligence, *Sci. Adv.* 9 (40) (2023) eadi1480, <http://dx.doi.org/10.1126/sciadv.adi1480>, arXiv:<https://www.science.org/doi/pdf/10.1126/sciadv.adi1480>, URL: <https://www.science.org/doi/abs/10.1126/sciadv.adi1480>.
- [52] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778.
- [53] A.L. Hodgkin, A.F. Huxley, A quantitative description of membrane current and its application to conduction and excitation in nerve, *J. Physiol.* 117 (4) (1952) 500.
- [54] N. Brunel, M.C. Van Rossum, Lapicque's 1907 paper: from frogs to integrate-and-fire, *Biol. Cybern.* 97 (5–6) (2007) 337–339.
- [55] J.T. Kaardal, F.E. Theunissen, T.O. Sharpee, A low-rank method for characterizing high-level neural computations, *Front. Comput. Neurosci.* 11 (2017) 68.
- [56] D.A. Ruff, C. Xue, L.E. Kramer, F. Baqai, M.R. Cohen, Low rank mechanisms underlying flexible visual representations, *Proc. Natl. Acad. Sci.* 117 (47) (2020) 29321–29329.
- [57] A. Novikov, D. Podoprikin, A. Osokin, D.P. Vetrov, Tensorizing neural networks, *Adv. Neural Inf. Process. Syst.* 28 (2015).
- [58] D. Liu, L.T. Yang, P. Wang, R. Zhao, Q. Zhang, Tt-tsvd: A multi-modal tensor train decomposition with its application in convolutional neural networks for smart healthcare, *ACM Trans. Multimed. Comput. Commun. Appl. (TOMM)* 18 (1s) (2022) 1–17.
- [59] A. Chertkov, G. Ryzhakov, G. Novikov, I. Oseledets, Optimization of functions given in the tensor train format, 2022, arXiv preprint arXiv:2209.14808.
- [60] A. Cichocki, N. Lee, I. Oseledets, A.-H. Phan, Q. Zhao, D.P. Mandic, et al., Tensor networks for dimensionality reduction and large-scale optimization: Part 1 low-rank tensor decompositions, *Found. Trends® Mach. Learn.* 9 (4–5) (2016) 249–429.
- [61] A. Cichocki, A.-H. Phan, Q. Zhao, N. Lee, I. Oseledets, M. Sugiyama, D.P. Mandic, et al., Tensor networks for dimensionality reduction and large-scale optimization: Part 2 applications and future perspectives, *Found. Trends® Mach. Learn.* 9 (6) (2017) 431–673.
- [62] M.W. Trosset, What is simulated annealing? *Optim. Eng.* 2 (2001) 201–213.
- [63] R.J. Williams, Simple statistical gradient-following algorithms for connectionist reinforcement learning, *Mach. Learn.* 8 (1992) 229–256.
- [64] A. Krizhevsky, G. Hinton, et al., Learning multiple layers of features from tiny images, 2009.
- [65] T. Miyato, T. Kataoka, M. Koyama, Y. Yoshida, Spectral normalization for generative adversarial networks, 2018, arXiv preprint arXiv:1802.05957.
- [66] D.P. Kingma, M. Welling, Auto-encoding variational bayes, 2013, arXiv preprint arXiv:1312.6114.
- [67] A. Van Den Oord, O. Vinyals, et al., Neural discrete representation learning, *Adv. Neural Inf. Process. Syst.* 30 (2017).
- [68] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen, Improved techniques for training gans, *Adv. Neural Inf. Process. Syst.* 29 (2016).
- [69] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, S. Hochreiter, Gans trained by a two time-scale update rule converge to a local nash equilibrium, *Adv. Neural Inf. Process. Syst.* 30 (2017).
- [70] T. Bu, W. Fang, J. Ding, P. Dai, Z. Yu, T. Huang, Optimal ANN-SNN conversion for high-accuracy and ultra-low-latency spiking neural networks, 2023, arXiv preprint arXiv:2303.04347.
- [71] J. Ding, Z. Yu, Y. Tian, T. Huang, Optimal ann-snn conversion for fast and accurate inference in deep spiking neural networks, 2021, arXiv preprint arXiv:2105.11654.
- [72] N. Rathi, G. Srinivasan, P. Panda, K. Roy, Enabling deep spiking neural networks with hybrid conversion and spike timing dependent backpropagation, 2020, arXiv preprint arXiv:2005.01807.
- [73] Q. Xu, Y. Li, J. Shen, J.K. Liu, H. Tang, G. Pan, Constructing deep spiking neural networks from artificial neural networks with knowledge distillation, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 7886–7895.
- [74] Y. LeCun, The MNIST database of handwritten digits, 1998, <http://yann.lecun.com/exdb/mnist/>.
- [75] H. Xiao, K. Rasul, R. Vollgraf, Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms, 2017, arXiv preprint arXiv:1708.07747.
- [76] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: *2009 IEEE Conference on Computer Vision and Pattern Recognition*, Ieee, 2009, pp. 248–255.
- [77] A. Krizhevsky, I. Sutskever, G.E. Hinton, Imagenet classification with deep convolutional neural networks, *Adv. Neural Inf. Process. Syst.* 25 (2012).
- [78] G. Huang, Z. Liu, L. Van Der Maaten, K.Q. Weinberger, Densely connected convolutional networks, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 4700–4708.
- [79] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, 2014, arXiv preprint arXiv:1409.1556.
- [80] J. Chevallier, Stimulus sensitivity of a spiking neural network model, *J. Stat. Phys.* 170 (4) (2018) 800–808.
- [81] G. Tanaka, T. Yamane, J.B. Héroux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, A. Hirose, Recent advances in physical reservoir computing: A review, *Neural Netw.* 115 (2019) 100–123, <http://dx.doi.org/10.1016/j.neunet.2019.03.005>.
- [82] G. Montúfar, R. Pascanu, K. Cho, Y. Bengio, On the number of linear regions of deep neural networks, in: *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS '14*, MIT Press, Cambridge, MA, USA, 2014, pp. 2924–2932.
- [83] V. Chikhanov, V. Bondarko, M. Danilova, A. Goluzina, Y. Shelepin, Complexity of images: Experimental and computational estimates compared, *Perception* 41 (6) (2012) 631–647, <http://dx.doi.org/10.1068/p6987>.
- [84] A. Baroncchelli, E. Caglioti, V. Loreto, Measuring complexity with zippers, *Eur. J. Phys.* 26 (5) (2005) S69–S77, <http://dx.doi.org/10.1088/0143-0807/26/5/s08>.
- [85] A.G. Casali, O. Gosseries, M. Rosanova, M. Boly, S. Sarasso, K.R. Casali, S. Casarotto, M.-A. Bruno, S. Laureys, G. Tononi, M. Massimini, A theoretically based index of consciousness independent of sensory processing and behavior, *Sci. Transl. Med.* 5 (198) (2013) <http://dx.doi.org/10.1126/scitranslmed.3006294>.


- [86] J. Yosinski, J. Clune, Y. Bengio, H. Lipson, How transferable are features in deep neural networks? *Adv. Neural Inf. Process. Syst.* 27 (2014).
- [87] J.J. DiCarlo, D. Zoccolan, N.C. Rust, How does the brain solve visual object recognition? *Neuron* 73 (3) (2012) 415–434, <http://dx.doi.org/10.1016/j.neuron.2012.01.010>.
- [88] W.J. Johnston, S. Fusi, Abstract representations emerge naturally in neural networks trained to perform multiple tasks, *Nature Commun.* 14 (1) (2023) <http://dx.doi.org/10.1038/s41467-023-36583-0>.
- [89] C.M.A. Pennartz, M.N. Oude Lohuis, U. Olcese, How ‘visual’ is the visual cortex? The interactions between the visual cortex and other sensory, motivational and motor systems as enabling factors for visual perception, *Phil. Trans. R. Soc. B* 378 (1886) (2023) <http://dx.doi.org/10.1098/rstb.2022.0336>.
- [90] J.A. Movshon, E.P. Simoncelli, Representation of naturalistic image structure in the primate visual cortex, *Cold Spring Harb. Symp. Quant. Biol.* 79 (2014) 115–122, <http://dx.doi.org/10.1101/sqb.2014.79.024844>.
- [91] P. Montes-Lourido, M. Kar, S.V. David, S. Sadagopan, Neuronal selectivity to complex vocalization features emerges in the superficial layers of primary auditory cortex, *PLoS Biol.* 19 (6) (2021) e3001299, <http://dx.doi.org/10.1371/journal.pbio.3001299>.
- [92] A. Achille, M. Rovere, S. Soatto, Critical learning periods in deep neural networks, 2017, arXiv preprint [arXiv:1711.08856](https://arxiv.org/abs/1711.08856).
- [93] V.P. Sotskov, N.A. Pospelov, V.V. Plusnin, K.V. Anokhin, Calcium imaging reveals fast tuning dynamics of hippocampal place cells and CA1 population activity during free exploration task in mice, *Int. J. Mol. Sci.* 23 (2) (2022) 638, <http://dx.doi.org/10.3390/ijms23020638>.
- [94] G. Buzsáki, K. Mizuseki, The log-dynamic brain: how skewed distributions affect network operations, *Nature Rev. Neurosci.* 15 (4) (2014) 264–278, <http://dx.doi.org/10.1038/nrn3687>.
- [95] P. Esser, R. Rombach, B. Ommer, Taming transformers for high-resolution image synthesis, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021*, pp. 12873–12883.
- [96] S. Panzeri, J.H. Macke, J. Gross, C. Kayser, Neural population coding: combining insights from microscopic and mass signals, *Trends in Cognitive Sciences* 19 (3) (2015) 162–172, <http://dx.doi.org/10.1016/j.tics.2015.01.002>.
- [97] J.A. Gallego, M.G. Perich, L.E. Miller, S.A. Solla, Neural manifolds for the control of movement, *Neuron* 94 (5) (2017) 978–984, <http://dx.doi.org/10.1016/j.neuron.2017.05.025>.
- [98] A. Zou, L. Phan, S. Chen, J. Campbell, P. Guo, R. Ren, A. Pan, X. Yin, M. Mazeika, A.-K. Dombrowski, S. Goel, N. Li, M.J. Byun, Z. Wang, A. Mallen, S. Basart, S. Koyejo, D. Song, M. Fredrikson, J.Z. Kolter, D. Hendrycks, Representation engineering: A top-down approach to AI transparency, 2023, <http://dx.doi.org/10.48550/ARXIV.2310.01405>, URL: <https://arxiv.org/abs/2310.01405>.
- [99] A. Templeton, T. Conerly, J. Marcus, J. Lindsey, T. Bricken, B. Chen, A. Pearce, C. Citro, E. Ameisen, A. Jones, H. Cunningham, N.L. Turner, C. McDougall, M. MacDiarmid, C.D. Freeman, T.R. Sumers, E. Rees, J. Batson, A. Jermyn, S. Carter, C. Olah, T. Henighan, Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet, *Transform. Circuits Thread* (2024) URL: <https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html>.
- [100] T. Alejandro-García, S. Kim, J. Pérez-Ortega, R. Yuste, Intrinsic excitability mechanisms of neuronal ensemble formation, *eLife* 11 (2022) <http://dx.doi.org/10.7554/eLife.77470>.
- [101] D. Bau, J.-Y. Zhu, H. Strobelt, A. Lapedriza, B. Zhou, A. Torralba, Understanding the role of individual units in a deep neural network, *Proc. Natl. Acad. Sci.* 117 (48) (2020) 30071–30078.
- [102] S.A. Matveev, I.V. Oseledets, E.S. Ponomarev, A.V. Chertkov, Overview of visualization methods for artificial neural networks, *Comput. Math. Math. Phys.* 61 (5) (2021) 887–899.
- [103] R.R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, D. Batra, GradCam: Visual explanations from deep networks via gradient-based localization, in: *Proceedings of the IEEE International Conference on Computer Vision, 2017*, pp. 618–626.
- [104] M. Sundararajan, A. Taly, Q. Yan, Axiomatic attribution for deep networks, in: *International Conference on Machine Learning, PMLR, 2017*, pp. 3319–3328.
- [105] H. Arakawa, D.C. Blanchard, K. Arakawa, C. Dunlap, R.J. Blanchard, Scent marking behavior as an odorant communication in mice, *Neurosci. Biobehav. Rev.* 32 (7) (2008) 1236–1248, <http://dx.doi.org/10.1016/j.neubiorev.2008.05.012>.
- [106] P. Kostenetskiy, R. Chulkevich, V. Kozyrev, HPC resources of the higher school of economics, in: *Journal of Physics: Conference Series*, IOP Publishing, 2021, 012050.
- [107] D.J. Gauthier, E. Bollt, A. Griffith, W.A.S. Barbosa, Next generation reservoir computing, *Nature Commun.* 12 (1) (2021) <http://dx.doi.org/10.1038/s41467-021-25801-2>.




Nikita Pospelov is a Ph.D. student at the MSU Institute for Advanced Brain Studies working on extracting biologically relevant variables from complex brain activity and particularly interested in bridging the gap between theoretical neuroscience and current AI research. He holds an MSc degree in Biophysics from Moscow State University. N.P. applies and develops methods of complex systems analysis to networks of the brain and to neural signals in order to unravel their inner structure and shed light on information processing mechanisms in the brain.



Andrei Chertkov, Ph.D. from Skoltech, is a Research Scientist at Artificial Intelligence Research Institute (AIRI), working on efficient high-dimensional computation methods based on the Tensor Train decomposition, applied to optimization, machine learning and artificial intelligence, physics.



Maxim Beketov is a Ph.D. student at the CS Department of Higher School of Economics (HSE), with research interests in statistics, geometry and topology and their applications in machine learning as well as neuroscience. He holds an M.Sc. degree from Moscow Institute of Physics and Technology (MIPT) and Skoltech.



Ivan Oseledets, D.Sc., is Professor and Head of the Laboratory of Computational Intelligence at Skoltech, as well as Director of Artificial Intelligence Research Institute (AIRI). Besides other accomplishments, he is widely known as author of the Tensor Train decomposition – an effective representation of low-rank high-dimensional tensors – a method widely used now in numerical computations, simulations, machine learning; applied in physics, chemistry, biology, data analysis. His current research interests include the development of new algorithms in machine learning and artificial intelligence, such as the construction of adversarial examples, theory of generative adversarial networks, and compression of neural networks.



Konstantin Anokhin, D.Sc., is a full member of the Russian Academy of Sciences, Professor and the head of MSU Institute for Advanced Brain Studies. He has extensively investigated the immediate early genes that are activated in the brain during learning and memory formation, playing a crucial role in constructing various types of long-term memory and stabilizing pathological conditions of the nervous system. He proposed and experimentally justified a scheme for molecular signaling pathways for immediate early — late genes underlying the consolidation of all known forms of long-lasting memory, and also developed new methods for imaging and controlling functions of the central nervous system. Prof. Anokhin’s research interests include molecular mechanisms of memory, theoretical neuroscience, principles of learning in biological and artificial neural networks and studies of consciousness.