



One-trial correction of legacy AI systems and stochastic separation theorems

Alexander N. Gorban^{a,e}, Richard Burton^{a,c}, Ilya Romanenko^d,
Ivan Yu. Tyukin^{a,b,e,1,*}

^a University of Leicester, Department of Mathematics, University Road, Leicester LE1 7RH, United Kingdom

^b Department of Automation and Control Processes, St. Petersburg State University of Electrical Engineering, Prof. Popova street 5, Saint-Petersburg 197376, Russian Federation

^c Imaging and Vision Group, ARM Ltd, 1 Summerpool Road, Loughborough LE11 5RD, United Kingdom

^d Spectral Edge Ltd., Bradfield Centre, 184 Science Park, Cambridge CB4 0GA, United Kingdom

^e Lobachevsky State University of Nizhny Novgorod, Prospekt Gagarina 23, Nizhny Novgorod 603950, Russian Federation



ARTICLE INFO

Article history:

Received 30 June 2017

Revised 29 January 2019

Accepted 1 February 2019

Available online 1 February 2019

Keywords:

Measure concentration

Separation theorems

Big data

Machine learning

ABSTRACT

We consider the problem of efficient “on the fly” tuning of existing, or *legacy*, Artificial Intelligence (AI) systems. The legacy AI systems are allowed to be of arbitrary class, albeit the data they are using for computing interim or final decision responses should possess an underlying structure of a high-dimensional topological real vector space. The tuning method that we propose enables dealing with errors without the need to re-train the system. Instead of re-training a simple cascade of perceptron nodes is added to the legacy system. The added cascade modulates the AI legacy system’s decisions. If applied repeatedly, the process results in a network of modulating rules “dressing up” and improving performance of existing AI systems. Mathematical rationale behind the method is based on the fundamental property of measure concentration in high dimensional spaces. The method is illustrated with an example of fine-tuning a deep convolutional network that has been pre-trained to detect pedestrians in images.

© 2019 Elsevier Inc. All rights reserved.

1. Introduction

Legacy information systems, i.e. information systems that have already been created and form crucial parts of existing solutions or service [3,4], are common in engineering practice and scientific computing [12]. They are becoming particularly wide-spread, as legacy Artificial Intelligence (AI) systems, in the areas of computer vision, image processing, data mining, and machine learning (see e.g. Caffè [27], MXNet [7] and Deeplearning4j [46] packages).

Despite their success and utility, legacy AI systems occasionally make mistakes. Their generalization errors may be caused by a number of issues, for example, by incomplete, insufficient or obsolete information used in their development and design, or by oversensitivity to some signals, etc. Adversarial images [36,45] are well-known examples of such issues for

* Corresponding author at: University of Leicester, Department of Mathematics, University Road, Leicester LE1 7RH, United Kingdom.

E-mail addresses: ag153@le.ac.uk (A.N. Gorban), Richard.Burton@arm.com (R. Burton), ilya@spectraledge.co.uk (I. Romanenko), I.Tyukin@le.ac.uk (I.Yu. Tyukin).

¹ The work was supported by the Ministry of Education and Science of Russia (Project no. 14.Y226.31.0022) and Innovate UK Knowledge Transfer Partnership grants KTP009890 and KTP010522.

classifiers based on Deep Learning Convolutional Neural Networks (CNNs) [30]. It has been shown in [45] that imperceptible changes in the images used in the training set may cause labeling errors, and at the same time completely unrecognizable to human perception objects can be classified as the ones that have already been learnt [36]. Retraining such legacy systems requires resources, e.g. computational power, time and/or access to training sets, that are not always available. Thus solutions that reliably correct mistakes of legacy AI systems without retraining are needed.

Significant progress has been made to date to understand and mitigate atypical and spurious mistakes. For example, in [21,23,24] using ensembles of classifiers was shown to improve performance of the overall system. With regards to adversarial images, augmenting data [31,35,37] and enforcing continuity of feature representations [51] help to increase reliability of classification. These approaches nevertheless do not warrant error-free behavior; AI and machine learning systems drawing conclusions on the basis of empirical data are expected to make mistakes, as human experts occasionally do too [11]. Hence having a method for fast and reliable rectification of these errors in legacy systems is crucial.

In this work we present a computationally efficient solution to the problem of correcting AI systems' mistakes. In contrast to conventional approaches [31,35,37,51] focusing on altering data and improving design procedures of legacy AI systems, we propose that the legacy AI system itself be augmented by miniature low-cost and low-risk additions. These additions are, in their essence, small neural network cascades that in principle can be easily incorporated into existing architectures already employing neural networks as their inherent components. Importantly, we show that for a large class of legacy AI systems in which decision criteria are calculated on the basis of high-dimensional data representation, such small neuronal cascades can be constructed via simple non-iterative procedures. We prove this by showing that *in an essentially high-dimensional finite random set with probability close to one all points are extreme*, i.e. *every point is linearly separable from the rest of the set*. Such a separability holds even for large random sets up to some upper limit, which grows *exponentially* with dimension. This is proven for finite samples i.i.d. drawn from an equidistribution in a ball or ellipsoid and demonstrated also for equidistributions in a cube and for normal distribution. We can hypothesize now that this is a very general property of all sufficiently regular essentially high dimensional distributions.

Hence, if a legacy AI system employs high-dimensional internal data representation then a single element, i.e. a mistake, in this representation can be separated from the rest by a simple perceptron, for example, by the Fisher discriminant. Several such mistakes can be collated into a common corrector implementable as a small neuronal cascade. The results are based on ideas of measure concentration [15,18,19] and can be related to the works of Gibbs [13] and Levi [32] and, on the other hand, to the Krein–Milman theorem [29].

The paper is organized as follows. Section 2 presents formal statement of the problem, Section 3 contains mathematical background for developing One-Trial correctors of Legacy AI systems, in Section 4 we state and discuss practically relevant generalizations, Section 5 provides experimental results showing viability of our method for correcting state-of-the-art Deep Learning CNN classifiers, and Section 6 concludes the paper. Main notational agreements are summarized in the Notation section below.

Notation

Throughout the paper the following notational agreements are used.

- \mathbb{R} denotes the field of real numbers;
- \mathbb{N} is the set of natural numbers;
- \mathbb{R}^n stands for the n -dimensional real space; unless stated otherwise symbol n is reserved to denote dimension of the underlying linear space;
- let $\mathbf{x} \in \mathbb{R}^n$, then $\|\mathbf{x}\|$ is the Euclidean norm of \mathbf{x} : $\|\mathbf{x}\| = \sqrt{x_1^2 + \dots + x_n^2}$;
- if \mathbf{x}, \mathbf{y} are two elements of \mathbb{R}^n then $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^n x_i y_i$;
- symbol $\mathbf{0}_n$ denotes an element of \mathbb{R}^n of which the values of all of its n coordinates are 0; if the dimension n is clear from the context, we will refer to such element as $\mathbf{0}$;
- $B_n(R)$ denotes a n -ball of radius R centered at $\mathbf{0}_n$: $B_n(R) = \{\mathbf{x} \in \mathbb{R}^n \mid \|\mathbf{x}\| \leq R\}$;
- $\mathcal{V}(\Xi)$ is the Lebesgue volume of $\Xi \subset \mathbb{R}^n$;
- \mathcal{M} is an i.i.d. sample equidistributed in $B_n(1)$;
- M is the number of points in \mathcal{M} , or simply the cardinality of the set \mathcal{M} .

Let E be a real vector space. Recall that a linear map $l: E \rightarrow \mathbb{R}$ is called a linear functional [20,40]. Similarly, an affine map $l: E \rightarrow \mathbb{R}$, i.e. a map defined as $l(\mathbf{x}) = l_0(\mathbf{x}) + b$ where l_0 is a linear functional and $b \in \mathbb{R}$, is referred to as an affine functional.

2. Problem formulation

Consider a legacy AI system that processes some *input* signals, produces *internal* representations of the input and returns some *outputs*. For the purposes of this work the exact nature and definition of the process and signals themselves are not important. Input signals may correspond to any physical measurements, outputs could be alarms or decisions, internal representations could include but not limited to extracted features, outputs of computational sub-routines etc. In the case

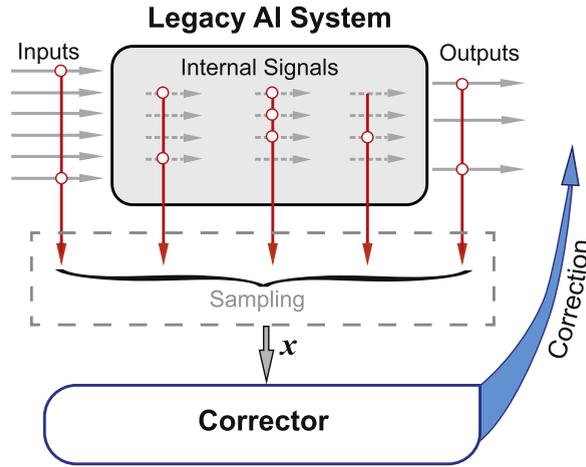


Fig. 1. Corrector of Legacy AI systems.

of a CNN classifier inputs are images, internal signals are outputs of convolutional and dense layers, and outputs are labels of objects.

We assume, however, that some relevant information about the input, internal signals, and outputs can be combined into a common object, \mathbf{x} , representing, but not necessarily defining, the *state* of the AI system. The objects \mathbf{x} are assumed to be elements of \mathbb{R}^n . Over relevant period of activity the AI system generates a set \mathcal{M} of representations \mathbf{x} . Since we do not wish to impose any prior knowledge of how \mathbf{x} are constructed and following standard assumptions in machine learning literature [47], it is natural to assume that \mathcal{M} is a random sample drawn from some distribution.

We suppose that some elements of the set \mathcal{M} are labeled as those corresponding to “errors” or mistakes which the original legacy AI system made. The task is to single out these errors and augment the AI systems’ response by an additional device, a *corrector*. A diagram illustrating this setting is shown in Fig. 1. The corrector system must in turn satisfy a list of properties ensuring that its deployment is practically feasible

1. The elements comprising the corrector must be re-usable in the original AI system;
2. The elements, as elementary learning machines, should be able to generalize;
3. The elements must be efficient computationally;
4. The corrector must allow for fast non-iterative learning;
5. The corrector should not affect functionality of the AI system, for reasonably large \mathcal{M} with $|\mathcal{M}| \gg n$.

For a broad range of AI systems, a candidate element satisfying requirements (1)–(3) is the parameterized affine functional

$$l(\mathbf{x}) = \langle \mathbf{x}, \mathbf{w} \rangle + b, \quad \mathbf{w} \in \mathbb{R}^n, \quad b \in \mathbb{R} \tag{1}$$

followed, if needed, by a suitable nonlinearity. It is a major building block of neural networks (convolutional, deep and shallow) as well as in decision trees and support vector machines. It’s computational efficiency and generalization capabilities [48] are well-known too. Whether remaining requirements could be guaranteed, however, is not clear. The problem therefore is to find an answer to this question.

In the next sections we show that, remarkably, in high dimensions the affine functionals do have these properties, with high probability.

3. Mathematical background

3.1. Separation of single points in finite random sets in high dimensions

Our basic example throughout this section is an equidistribution in the unit ball² $B_n(1)$ in \mathbb{R}^n . Generalizations to equidistribution in an ellipsoid will be formulated in Section 4.1, and other distributions will be discussed in Sections 4.2 and 5.

Definition 1. Let X and Y be subsets of \mathbb{R}^n . We say that a linear functional l on \mathbb{R}^n separates X and Y if there exists a $t \in \mathbb{R}$ such that

$$l(\mathbf{x}) > t > l(\mathbf{y}) \quad \forall \mathbf{x} \in X, \quad \mathbf{y} \in Y.$$

² Partially, the ideas concerning equidistributions in $B_n(1)$ have been presented in a conference talk [17]

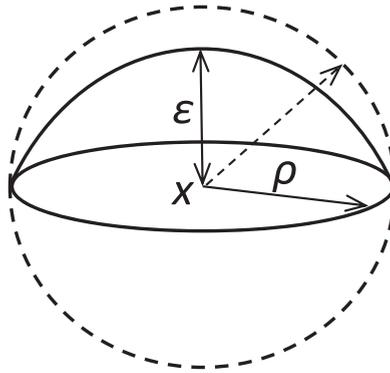


Fig. 2. A test point \mathbf{x} and a spherical cap on distance ε from the surface of unit sphere. Escribed sphere of radius ρ is showed by dashed line.

Let \mathcal{M} be an i.i.d. sample drawn from the equidistribution on the unit ball $B_n(1)$. We begin with evaluating the probability that a single element \mathbf{x} randomly and independently selected from the same equidistribution can be separated from \mathcal{M} by a linear functional. This probability, denoted as $P_1(\mathcal{M}, n)$, is estimated in the theorem below.

Theorem 1. Consider an equidistribution in a unit ball $B_n(1)$ in \mathbb{R}^n , and let \mathcal{M} be an i.i.d. sample from this distribution. Then

$$P_1(\mathcal{M}, n) \geq \max_{\varepsilon \in (0,1)} (1 - (1 - \varepsilon)^n) \left(1 - \frac{\rho(\varepsilon)^n}{2}\right)^M, \tag{2}$$

$$\rho(\varepsilon) = (1 - (1 - \varepsilon)^2)^{\frac{1}{2}}$$

Proof of Theorem 1. The proof of the theorem is contained mostly in the following lemma

Lemma 1. Let \mathbf{y} be random point from an equidistribution on a unit ball $B_n(1)$. Let $\mathbf{x} \in B_n(1)$ be a point inside the ball with $1 > \|\mathbf{x}\| > 1 - \varepsilon > 0$. Then

$$P\left(\left\langle \frac{\mathbf{x}}{\|\mathbf{x}\|}, \mathbf{y} \right\rangle < 1 - \varepsilon\right) \geq 1 - \frac{\rho(\varepsilon)^n}{2}. \tag{3}$$

Proof of Lemma 1. Recall that [32]: $\mathcal{V}(B_n(r)) = r^n \mathcal{V}(B_n(1))$ for all $n \in \mathbb{N}$ $r > 0$. The point \mathbf{x} is inside the spherical cap $C_n(\varepsilon)$:

$$C_n(\varepsilon) = B_n(1) \cap \left\{ \boldsymbol{\xi} \in \mathbb{R}^n \mid \left\langle \frac{\mathbf{x}}{\|\mathbf{x}\|}, \boldsymbol{\xi} \right\rangle > 1 - \varepsilon \right\} \tag{4}$$

The volume of this cap can be estimated from above [2] (see Fig. 2) as

$$\mathcal{V}(C_n(\varepsilon)) \leq \frac{1}{2} \mathcal{V}(B_n(1)) \rho(\varepsilon)^n. \tag{5}$$

The probability that the point $\mathbf{y} \in \mathcal{M}$ is outside of $C_n(\varepsilon)$ is equal to $1 - \mathcal{V}(C_n(\varepsilon))/\mathcal{V}(B_n(1))$. Estimate (3) now immediately follows from (5). \square

Let us now return to the proof of the theorem. If \mathbf{x} is selected independently from the equidistribution on $B_n(1)$ then the probabilities that $\mathbf{x} = \mathbf{0}_n$ or that it is on the boundary of the ball are 0. Let $\mathbf{x} \neq \mathbf{0}_n$ be in the interior of $B_n(1)$. According to Lemma 1, the probability that a linear functional l separates \mathbf{x} from a point $\mathbf{y} \in \mathcal{M}$ is larger than $1 - \frac{1}{2} \rho(\varepsilon)^n$. Given that points of the set \mathcal{M} are i.i.d. in accordance to the equidistribution on $B_n(1)$, the probability that l separates \mathbf{x} from \mathcal{M} is no smaller than $(1 - \frac{1}{2} \rho(\varepsilon)^n)^M$.

On the other hand

$$P(1 > \|\mathbf{x}\| > 1 - \varepsilon \mid \mathbf{x} \in B_n(1)) = (1 - (1 - \varepsilon)^n).$$

Given that \mathbf{x} and $\mathbf{y} \in \mathcal{M}$ are independently drawn from the same equidistribution and that the probabilities of randomly selecting the point \mathbf{x} exactly on the boundary of $B_n(1)$ or in its centre are zero, we can conclude that

$$P_1(\mathcal{M}, n) \geq (1 - (1 - \varepsilon)^n) \left(1 - \frac{1}{2} \rho(\varepsilon)^n\right)^M. \tag{6}$$

Finally, noticing that (6) holds for all $\varepsilon \in (0, 1)$ including the value of ε maximizing the rhs of (6), we can conclude that (2) holds true too. \square

Remark 1. For $\rho(\varepsilon)^n$ small (i.e. $\rho(\varepsilon)^n \ll 1$) the term $\left(1 - \frac{\rho(\varepsilon)^n}{2}\right)^M$ can be approximated by the exponential $e^{-M \frac{\rho(\varepsilon)^n}{2}}$. In this approximation, the rhs of estimate (2) becomes

$$\max_{\varepsilon \in (0,1)} (1 - (1 - \varepsilon)^n) e^{-M \frac{\rho(\varepsilon)^n}{2}}, \quad \rho(\varepsilon)^n \ll 1, \tag{7}$$

resulting in a convenient approximate lower bound for the probability $P_1(\mathcal{M}, n)$.

To see how large the probability $P_1(\mathcal{M}, n)$ could become for already rather modest dimensions, e.g. for $n = 50$, we estimate the value of $P_1(\mathcal{M}, 50)$ invoking (7) and letting $\varepsilon = 1/5$ and $\rho = 3/5$. The corresponding approximate lower bound for $P_1(\mathcal{M}, 50)$ is

$$0.999985727 \exp(-4 \times 10^{-12}M).$$

For $M \leq 10^9$ this figure is bounded from below by 0.995952534. We note that the original expression in the rhs of (2) for the same values of ε , n and M results in 0.995952506 which is different from the approximation only in the 8th digit. Thus in dimension 50 (and higher) a random point is linearly separable from a random set of 10^9 points (drawn independently from the same equidistribution in $B_n(1)$) with probability of approximately 0.996.

Remark 2. If \mathbf{x} is an element from the sample \mathcal{M} then the probability that the element \mathbf{x} is linearly separable from all other points in the sample may be bounded from below by the values of $P_1(\mathcal{M}, n)$, albeit with a possible replacement of M with $M - 1$ in (2) for a higher-accuracy estimate.

Remark 3. Let $\mathbf{x} \in \mathcal{M}$ be a given query point. This query point determines the value of $\varepsilon = 1 - \|\mathbf{x}\|$ as the least ε -thickening of the unit sphere containing \mathbf{x} . With probability 1 the values of ε belong to the open interval $(0,1)$. Let $p \in (0, 1)$ be the desired probability that \mathbf{x} is separated from the rest of the sample \mathcal{M} . It is clear that the estimate $P_1(\mathcal{M}, n) \geq p$ holds for M from some interval $[1, \bar{M}]$. Interestingly, for n large enough, the maximal number \bar{M} is exponentially large in dimension n . Indeed, let us fix the values of $\varepsilon \in (0, 1)$ and $p \in (0, 1)$. Then we find the estimate of the maximal possible sample size for which $P_1(\mathcal{M}, n) \geq p$ remains valid:

$$\max\{\bar{M}\} \geq \frac{\ln(p)}{\ln\left(1 - \frac{\rho(\varepsilon)^n}{2}\right)} - \frac{\ln(1 - (1 - \varepsilon)^n)}{\ln\left(1 - \frac{\rho(\varepsilon)^n}{2}\right)}$$

Using

$$\frac{x}{x-1} \leq \ln(1-x) \leq -x$$

we conclude that

$$\max\{\bar{M}\} \geq \left(\frac{1}{\rho(\varepsilon)}\right)^n C(n, \varepsilon)$$

where

$$C(n, \varepsilon) = 2 \left(|\ln(p)| \left(1 - \frac{\rho(\varepsilon)^n}{2}\right) - |\ln(1 - (1 - \varepsilon)^n)| \right).$$

Observe that for any fixed $\varepsilon \in (0, 1)$ there is an $N(\varepsilon)$ large enough such that $C(n, \varepsilon) \geq |\ln(p)|$ for all $n \geq N(\varepsilon)$. Hence, for n sufficiently large the following estimate holds:

$$\max\{\bar{M}\} \geq e^{n \ln(\rho(\varepsilon)^{-1})} |\ln(p)|. \tag{8}$$

Eq. (8) can be viewed as a *separation capacity* estimate of linear functionals. This estimate links the level of desired performance specified by p , maximal size of the sample, M , and parameters of the data, n and ε .

3.2. Extreme points of a random finite set

So far we have discussed the question of separability of a single random point \mathbf{x} , drawn from the equidistribution on $B_n(1)$, from a random i.i.d. sample \mathcal{M} drawn from the same distribution. In practice, however, the data or a training set are given or fixed. It is thus important to know if the “point” linear separability property formulated in Theorem 1 persists (in one form or another) when the test point \mathbf{x} belongs to the sample \mathcal{M} itself. In particular, the question is if the probability $P_M(\mathcal{M}, n)$ that each point $\mathbf{y} \in \mathcal{M}$ is linearly separable from $\mathcal{M} \setminus \{\mathbf{y}\}$ is close to 1 in high dimensions? If such a property does hold then one could conclude that in high dimensions with probability close to 1 all points of \mathcal{M} are the vertices (extreme points) of the convex hull of \mathcal{M} and none of $\mathbf{y} \in \mathcal{M}$ is a convex combination of other points. The fact that this is indeed the case follows from Theorem 2

Theorem 2. Consider an equidistribution in a unit ball $B_n(1)$ in \mathbb{R}^n , and let \mathcal{M} be an i.i.d. sample from this distribution. Then

$$P_M(\mathcal{M}, n) \geq \max_{\varepsilon \in (0,1)} \left[(1 - (1 - \varepsilon)^n) \left(1 - (M - 1) \frac{\rho(\varepsilon)^n}{2}\right) \right]^M. \tag{9}$$

Proof of Theorem 2. Let $P: \mathcal{F} \rightarrow [0, 1]$ be a probability measure and $A_i \in \mathcal{F}$, $i = 1, \dots, M$. It is well-known that

$$P(A_1 \vee A_2 \vee \dots \vee A_M) \leq \sum_{i=1}^M P(A_i) \tag{10}$$

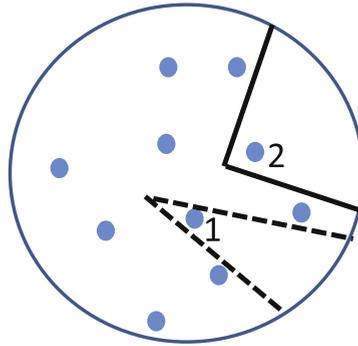


Fig. 3. Two-neuron separation in finite sets. Every point can be separated by a sufficiently acute angle or highly correlated neurons: Point 1 is separated from other points by an acute angle (dashed lines). Point 2 is separated by a right angle (non-correlated neurons), but cannot be separated by a linear functional (i.e. by a straight line).

The probability that a test point \mathbf{y} is in the ε -vicinity of the boundary of $B_n(1)$ is $1 - (1 - \varepsilon)^n$. Fix $\mathbf{y} \in \mathcal{M}$ and construct spherical caps $C_n(\varepsilon)$ for each element in $\mathcal{M} \setminus \{\mathbf{y}\}$ as specified by (4) but with \mathbf{x} replaced by the corresponding points from $\mathcal{M} \setminus \{\mathbf{y}\}$. According to (10) and Lemma 1, the probability that \mathbf{y} is in any of these caps is no larger than $(M - 1) \frac{\rho(\varepsilon)^n}{2}$. Hence the probability that a point $\mathbf{y} \in \mathcal{M}$ is separable from $\mathcal{M} \setminus \{\mathbf{y}\}$ is larger or equal to $(1 - (1 - \varepsilon)^n)(1 - (M - 1) \frac{\rho(\varepsilon)^n}{2})$. Given that points of \mathcal{M} are drawn independently and that there are exactly M points in \mathcal{M} , the probability that every single point is linearly separable from the rest satisfies (9). \square

Remark 4. Note that employing (10) one can obtain another estimate of P_M :

$$P_M(\mathcal{M}, n) \geq 1 - M(1 - P_1(\mathcal{M}, n)). \tag{11}$$

We can utilise this estimate together with (8) and estimate the maximal size of the sample from below. Indeed, if we require that $P_M(\mathcal{M}, n) \geq q$ for some probability q , $0 < q < 1$, then it is sufficient that $P_1(\mathcal{M}, n) > p$, where $1 - p = \frac{1}{M}(1 - q)$. Using in (8) $|\ln p| > 1 - p$, we get that $P_M(\mathcal{M}, n) > q$ if $M \leq \tilde{M}$ for some maximal value \tilde{M} , that satisfies the inequality

$$\tilde{M} \geq e^{n \ln(\rho(\varepsilon)^{-1})} \frac{1 - q}{\tilde{M}}.$$

Immediately from this inequality we get the explicit exponential estimate of the maximum of \tilde{M} from below:

$$\max\{\tilde{M}\} \geq e^{\frac{1}{2} n \ln(\rho(\varepsilon)^{-1})} \sqrt{1 - q} \tag{12}$$

Similarly to the example discussed in the end of the previous section, let us evaluate the right-hand side of (9) for some fixed values of n and M . If $n = 50$, $M = 1000$, and $\varepsilon = 1/5$, $\rho = 3/5$ then this estimate gives: $P_M(\mathcal{M}, 50) > 0.985$.

3.3. Two-functional (two-neuron) separation in finite sets

So far we have provided estimates of the probabilities that a single linear classifier or a learning machine can separate a given point from the rest of data and showed that two disjoint weakly compact subsets of a topological vector space can be separated by small networks of perceptrons. Let us now see how employing small networks may improve probabilities of separation of a point from the rest of the data in high dimensions. In particular, we will consider the case of a two-neuron separation in which the network is a simple cascade comprised of two perceptrons followed by a conjunction operation.

Before, however, going any further we need to somewhat clarify and adjust the notion of separability of a point from a finite data set by a network so that the question makes any practical and theoretical sense. Consider, for instance, the problem of separating a test point by just two perceptrons. If one projects the data onto a 2d plane so that projections of the test point and any other point from the rest of the data do not coincide then the problem always has a solution. This is illustrated with the diagram in Fig. 3. According to this diagram any given point from an arbitrary but finite data set could be cut out from the rest of the data by just two lines that intersect at a sufficiently acute angle. Thus two hyperplanes $\{\mathbf{x} \mid l_1(\mathbf{x}) = \theta_1\}$ and $\{\mathbf{x} \mid l_2(\mathbf{x}) = \theta_2\}$ whose projections onto the 2d plane are exactly these two lines already constitute the two-neuron separating cascade.

The problem with this solution is that if the acute angle determined by the inequalities $l_1(\mathbf{x}) > \theta_1$, $l_2(\mathbf{x}) > \theta_2$ is small then the coefficients of the corresponding linear functionals $l_1(\mathbf{x})$ and $l_2(\mathbf{x})$, i.e. synaptic weights of the neurons, are highly correlated. Their Pearson correlation coefficient is close to -1 . This implies that robustness of such a solution is low. Small changes in the coefficients of one perceptron could result in loss of separation. This motivates an alternative solution in which the coefficients are uncorrelated, i.e. the angles between two hyperplanes are right (or almost right).

Let us now analyse the problem of separation of a random i.i.d. finite sample \mathcal{M} drawn from an equidistribution in $B_n(1)$ from a point \mathbf{x} drawn independently from the same distribution by two non-correlated neurons. More formally, we are

interested in the probability $\mathcal{P}_1(\mathcal{M}, n)$ that a two-neuron cascade with uncorrelated synaptic weights separates \mathbf{x} from \mathcal{M} . An estimate of this probability is provided in the next theorem.

Theorem 3. Consider an equidistribution in a unit ball $B_n(1)$ in \mathbb{R}^n , and let \mathcal{M} be an i.i.d. sample from this distribution. Then

$$\begin{aligned} \mathcal{P}_1(\mathcal{M}, n) &\geq \max_{\varepsilon \in (0,1)} (1 - (1 - \varepsilon)^n) \\ &\quad \times \left(1 - \frac{\rho(\varepsilon)^n}{2}\right)^M e^{(M-n+1) \left[\frac{\rho(\varepsilon)^n}{1 - \frac{\rho(\varepsilon)^n}{2}}\right]} \\ &\quad \times \left(1 - \frac{1}{n!} \left((M - n + 1) \frac{\rho(\varepsilon)^n}{1 - \frac{\rho(\varepsilon)^n}{2}}\right)^n\right). \end{aligned} \tag{13}$$

Proof. Observe that in the case of general position, a single neuron (viz. linear functional) separates $n + 1$ points with probability 1. This means that if no more than $n - 1$ points from \mathcal{M} are in the spherical cap $C_n(\varepsilon)$ corresponding to the test point then the second perceptron whose weights are orthogonal to the first one will filter out these additional spurious $n - 1$ points with probability 1.

Let p_c be the probability that a point from \mathcal{M} falls within the spherical cap $C_n(\varepsilon)$. Then the probability that only up to $n - 1$ points of \mathcal{M} will be in the cap $C_n(\varepsilon)$ is

$$\mathcal{P}(M, n) = \sum_{k=0}^{n-1} \binom{M}{k} (1 - p_c)^{M-k} p_c^k$$

Observe that $\mathcal{P}(M, n)$, as a function of p_c , is monotone and non-increasing on the interval $[0,1]$, with $\mathcal{P}(M, n) = 0$ at $p_c = 1$ and $\mathcal{P}(M, n) = 1$ at $p_c = 0$. Hence taking estimate (5) into account one can conclude:

$$\mathcal{P}(M, n) \geq \sum_{k=0}^{n-1} \binom{M}{k} \left(1 - \frac{\rho(\varepsilon)^n}{2}\right)^{M-k} \left(\frac{\rho(\varepsilon)^n}{2}\right)^k.$$

Noticing that

$$\sum_{k=0}^{n-1} \binom{M}{k} \left(1 - \frac{\rho(\varepsilon)^n}{2}\right)^{M-k} \left(\frac{\rho(\varepsilon)^n}{2}\right)^k = \left(1 - \frac{\rho(\varepsilon)^n}{2}\right)^M \sum_{k=0}^{n-1} \binom{M}{k} \left(\frac{\frac{\rho(\varepsilon)^n}{2}}{1 - \frac{\rho(\varepsilon)^n}{2}}\right)^k,$$

and bounding $\binom{M}{k}$ from above and below as $\frac{(M-n+1)^k}{k!} \leq \binom{M}{k} \leq \frac{M^k}{k!}$ for $0 \leq k \leq n - 1$ we obtain

$$\mathcal{P}(\mathcal{M}, n) \geq \left(1 - \frac{\rho(\varepsilon)^n}{2}\right)^M \sum_{k=0}^{n-1} \frac{1}{k!} \left(\frac{(M - n + 1) \frac{\rho(\varepsilon)^n}{2}}{1 - \frac{\rho(\varepsilon)^n}{2}}\right)^k.$$

Invoking Taylor’s expansion of e^x at $x = 0$ with the Lagrange remainder term:

$$e^x = \sum_{k=0}^{n-1} \frac{x^k}{k!} + \frac{x^n}{n!} e^{\xi}, \quad \xi \in [0, x],$$

we can conclude that

$$\sum_{k=0}^{n-1} \frac{x^k}{k!} \geq e^x \left(1 - \frac{x^n}{n!}\right)$$

for all $x \geq 0$. Hence

$$\mathcal{P}(\mathcal{M}, n) \geq \left(1 - \frac{\rho(\varepsilon)^n}{2}\right)^M e^{(M-n+1) \left[\frac{\rho(\varepsilon)^n}{1 - \frac{\rho(\varepsilon)^n}{2}}\right]} \times \left(1 - \frac{1}{n!} \left((M - n + 1) \frac{\rho(\varepsilon)^n}{1 - \frac{\rho(\varepsilon)^n}{2}}\right)^n\right).$$

Finally, given that the probability that the test point \mathbf{x} is in the ε -vicinity of the boundary of $B_n(1)$ is at least $(1 - (1 - \varepsilon)^n)$ and that \mathbf{x} is independently selected from the same equidistribution as the set \mathcal{M} , we obtain $\mathcal{P}_1(\mathcal{M}, n) \geq (1 - (1 - \varepsilon)^n) \mathcal{P}(\mathcal{M}, n)$. This in turn implies that (13) holds. \square

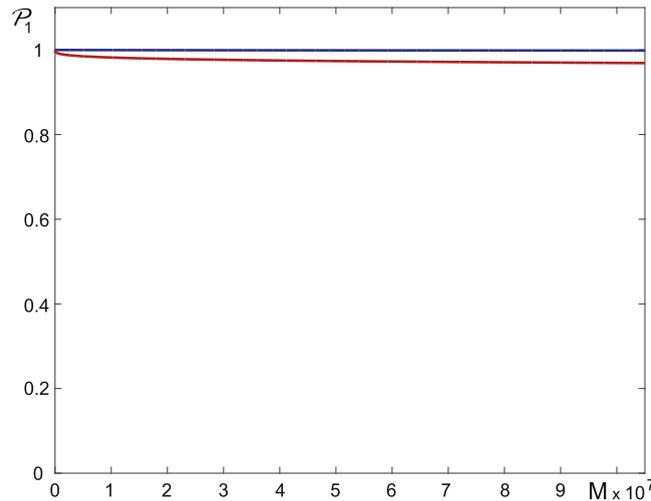


Fig. 4. Illustration to [Theorem 3](#). Blue line shows an estimate of the rhs of [\(13\)](#) as a function of M for $n = 30$. Red line depicts an estimate of the rhs of [\(2\)](#) as a function of M for the same values of n . (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

At the first glance estimate [\(13\)](#) looks more complicated as compared to e.g. [\(2\)](#). Yet, it differs from the latter by mere two factors. The first factor

$$\left(1 - \frac{1}{n!} \left((M - n + 1) \frac{\frac{\rho(\varepsilon)^n}{2}}{1 - \frac{\rho(\varepsilon)^n}{2}} \right)^n \right)$$

is close to 1 for $(M - n + 1) \frac{\rho(\varepsilon)^n}{2} < 1$ and n sufficiently large. The second factor:

$$e^{(M-n+1) \left[\frac{\frac{\rho(\varepsilon)^n}{2}}{1 - \frac{\rho(\varepsilon)^n}{2}} \right]},$$

is more important. It compensates for the decay of the probability of separation due to the term $(1 - \frac{\rho(\varepsilon)^n}{2})^M$ keeping the rhs of [\(13\)](#) close to 1 over large interval of values of M . The effect is illustrated with [Fig. 4](#). Observe that the probability corresponding to the two-neuron cascade remains in a close vicinity of 1, and the probability of one-neuron separation gradually decays with M .

Remark 5. Comparing performance of single vs two-neuron separability in terms of the probabilities $\mathcal{P}_1(\mathcal{M}, n)$ involves taking the maximum of

$$\mathcal{P}_1(\mathcal{M}, n, \varepsilon) = (1 - (1 - \varepsilon)^n) \left(1 - \frac{\rho(\varepsilon)^n}{2} \right)^M \tag{14}$$

and

$$\begin{aligned} \mathcal{P}_1(\mathcal{M}, n, \varepsilon) &= (1 - (1 - \varepsilon)^n) \\ &\times \left(1 - \frac{\rho(\varepsilon)^n}{2} \right)^M e^{(M-n+1) \left[\frac{\frac{\rho(\varepsilon)^n}{2}}{1 - \frac{\rho(\varepsilon)^n}{2}} \right]} \\ &\times \left(1 - \frac{1}{n!} \left((M - n + 1) \frac{\frac{\rho(\varepsilon)^n}{2}}{1 - \frac{\rho(\varepsilon)^n}{2}} \right)^n \right) \end{aligned} \tag{15}$$

with respect to ε over $(0,1)$. In some situations, when the testing point is already given, the probabilities $\mathcal{P}_1(\mathcal{M}, n)$ are no longer relevant since the value of ε corresponding to the testing point is fixed. In this cases one needs to compare $\mathcal{P}_1(\mathcal{M}, n, \varepsilon)$ defined by [\(14\)](#) and [\(15\)](#) instead. Performance of the corresponding separation schemes are illustrated with [Fig. 5](#). Notice that the two-neuron cascade significantly outperforms the single neuron one over a large interval of values of M .

Remark 6. The probability $\mathcal{P}_M(\mathcal{M}, n)$ that each point from \mathcal{M} can be separated from other points by two uncorrelated neurons can be estimated like in [Remark 4](#): $\mathcal{P}_M(\mathcal{M}, n) \geq 1 - M(1 - \mathcal{P}_1(\mathcal{M}, n))$.

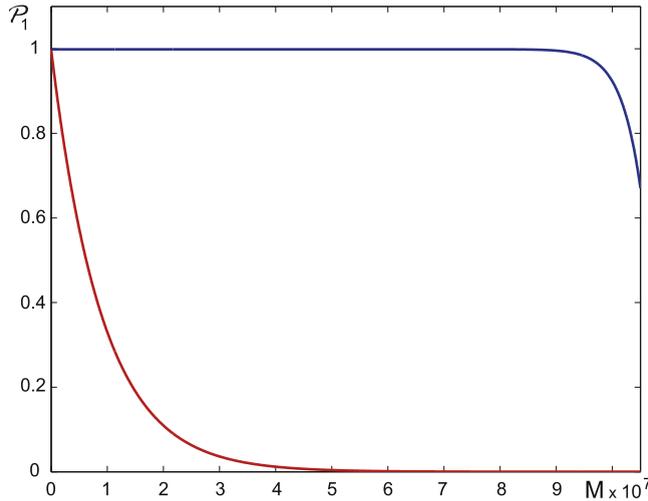


Fig. 5. Illustration to Remark 5. Blue line shows an estimate of the rhs of (15) as a function of M at $\varepsilon = 1/5$, $\rho(\varepsilon) = 3/5$, and $n = 30$. Red line depicts an estimate of the rhs of (14) as a function of M for the same values of ε , $\rho(\varepsilon)$, and n . (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.4. Separation by small neural networks

In the previous sections we analyzed and discussed linear separability of single elements of \mathcal{M} . As far as the problem of a legacy AI system corrector is concerned, however, it is not difficult to envision a need to implement more complicated dependencies, including logical predicates or involving more complicated data geometry. Below we provide some notions and results enabling us to deal with these issues.

Consider a topological real vector space L . For every continuous linear functional l on L and a real number θ we define an open elementary neural predicate $l(\mathbf{x}) > \theta$ and a closed elementary predicate $l(\mathbf{x}) \geq \theta$ ('open' and 'closed' are sets defined by predicates). The elementary neural predicate is true (=1) or false (=0) depending on whether the correspondent inequality holds. The negation of an open elementary neural predicate is a closed elementary neural predicate (and converse). A compound predicate is a Boolean expression composed of elementary predicates.

Definition 2. A set $X \subset L$ is k -neuron separable from a set $Y \subset L$ if there exist k elementary neural predicates P_1, \dots, P_k and a composed of them compound predicate, a Boolean function $B(\mathbf{x}) = B(P_1(\mathbf{x}), \dots, P_k(\mathbf{x}))$, such that $B(\mathbf{x}) = \text{true}$ for $\mathbf{x} \in X$ and $B(\mathbf{x}) = \text{false}$ for $\mathbf{x} \in Y$.

If X consist of one point then the compound predicate in the definition of k -neuron separability can be selected in a form of conjunction of k elementary predicates $B(\mathbf{x}) = P_1(\mathbf{x}) \& \dots \& P_k(\mathbf{x})$:

Proposition 1. Let $X = \{\mathbf{z}\}$ and X is k -neuron separable from $Y \subset L$. Then there exist k elementary neural predicates P_1, \dots, P_k such that the Boolean function $B(\mathbf{x}) = P_1(\mathbf{x}) \& \dots \& P_k(\mathbf{x})$ is true on X and false on Y .

Proof. Assume that there exist k elementary neural predicates Q_1, \dots, Q_k and a Boolean function $B(Q_1, \dots, Q_k)$ such that B is true on \mathbf{x} and false on Y . Represent B in a disjunctive normal form as a disjunction of conjunctive clauses:

$$B = C_1 \vee C_2 \vee \dots \vee C_N,$$

where each C_i has a form $C_i = R_1 \& R_2 \& \dots \& R_k$ and $R_i = Q_i$ or $R_i = \neg Q_i$.

At least one conjunctive clause $C_i(\mathbf{x})$ is true because $B(\mathbf{x}) = \text{true}$. On the other hand, $C_i(\mathbf{y}) = \text{false}$ for all $\mathbf{y} \in Y$ and $i = 1, \dots, N$ because their disjunction $B(\mathbf{y}) = \text{false}$. Let $C_i(\mathbf{x}) = \text{true}$. Then C_i separates \mathbf{x} from Y and other conjunctive clauses are not necessary in B . We can take $B = C_i = R_1 \& R_2 \& \dots \& R_k$, where $R_i = Q_i$ or $R_i = \neg Q_i$. Finally, we take $P_i = R_i$. \square

In contrast to Hahn–Banach theorems of linear separation, the k -neuron separation of sets does not require any sort of convexity. Weak compactness is sufficient. Assume that continuous linear functionals separate points in L . Let $Y \subset L$ be a weakly compact set [43] and $\mathbf{x} \notin Y$.

Proposition 2. \mathbf{x} is k -neuron separable from Y for some k .

Proof. For each $\mathbf{y} \in Y$ there exists a continuous linear functional $l_{\mathbf{y}}$ on L such that $l_{\mathbf{y}}(\mathbf{x}) = l_{\mathbf{y}}(\mathbf{y}) + 2$ because continuous linear functionals separate points in L . Inequality $l_{\mathbf{y}}(\mathbf{z}) < l_{\mathbf{y}}(\mathbf{x}) - 1$ defines an open half-space $L_{\mathbf{y}}^< = \{\mathbf{z} \mid l_{\mathbf{y}}(\mathbf{z}) < l_{\mathbf{y}}(\mathbf{y}) + 1\}$, and $\mathbf{y} \in L_{\mathbf{y}}^<$. The collection of sets $\{L_{\mathbf{y}}^< \mid \mathbf{y} \in Y\}$ forms an open cover of Y . Each set $L_{\mathbf{y}}^<$ is open in weak topology and Y is weakly

compact, hence there exists a final covering of Y by sets $L_{\mathbf{y}}^{\leq}$:

$$Y \subset \bigcup_{i=1, \dots, k} L_{\mathbf{y}_i}^{\leq}$$

for some finite subset $\{\mathbf{y}_1, \dots, \mathbf{y}_k\} \subset Y$. The inequality $l_{\mathbf{y}_i}(\mathbf{x}) > l_{\mathbf{y}_i}(\mathbf{y}_i) + \alpha$ holds for all $i = 1, \dots, k$ and $\alpha < 2$. Let us select $1 < \alpha < 2$ and take elementary neural predicates $P_i^\alpha(\mathbf{x}) = (l_{\mathbf{y}_i}(\mathbf{x}) > l_{\mathbf{y}_i}(\mathbf{y}_i) + \alpha)$. The conjunction $P_1^\alpha \& \dots \& P_k^\alpha$ is true on \mathbf{x} . Each point $\mathbf{y} \in Y$ belongs to $L_{\mathbf{y}_i}^{\leq}$ for at least one $i = 1, \dots, k$. For this i , $P_i^\alpha(\mathbf{y}) = \text{false}$. Therefore, $P_1^\alpha \& \dots \& P_k^\alpha$ is false on Y . Hence, this conjunction separates \mathbf{x} from Y . \square

Remark 7. Note that we can select two numbers $1 < \beta < \alpha < 2$. Both $P_1^\alpha \& \dots \& P_k^\alpha$ and $P_1^\beta \& \dots \& P_k^\beta$ separate \mathbf{x} from Y .

The theorem about k -neuron separation of weakly compact sets follows from Proposition 2. Let X and Y be disjoint and weakly compact subsets of L , and continuous linear functionals separate points in L . Then the following statement holds

Theorem 4. X is k -neuron separable from Y for some k .

Proof. For each $\mathbf{x} \in X$ apply the construction from the proof of Proposition 2 and construct the conjunction $B_{\mathbf{x}} = P_1^\alpha \& \dots \& P_k^\alpha$, which separates \mathbf{x} from Y . Every set of the form $F_{\mathbf{x}} = \{\mathbf{z} \in L \mid B_{\mathbf{x}}(\mathbf{z}) = \text{true}\}$ is intersection of a finite number of open half-spaces and, therefore, is open in weak topology. The collection of sets $\{F_{\mathbf{x}} \mid \mathbf{x} \in X\}$ covers X . There exist a final cover of X by sets $F_{\mathbf{x}_j}$:

$$X \subset \bigcup_{j=1, \dots, g} F_{\mathbf{x}_j}$$

for some finite set $\{\mathbf{x}_1, \dots, \mathbf{x}_g\} \subset X$. The composite predicate that separates X from Y can be chosen in the form $B(\mathbf{z}) = B_{\mathbf{x}_1}(\mathbf{z}) \vee \dots \vee B_{\mathbf{x}_g}(\mathbf{z})$. \square

4. Discussion

The nature of phenomenon described is universal and does not depend on many details of the distribution. In some sense, it should be just essentially high-dimensional. Due to space limitations cannot review all these generalizations here but will nevertheless provide several examples.

4.1. Equidistribution in an ellipsoid

Let points in the set \mathcal{M} be selected by independent trials taken from the equidistribution in a n -dimensional ellipsoid. Without loss of generality, we present this ellipsoid in the orthonormal eigenbasis:

$$E_n = \{\mathbf{x} \in \mathbb{R}^n \mid \sum_{i=1}^n \frac{x_i^2}{c_i^2} \leq 1\}, \tag{16}$$

where rc_i are the semi-principal axes. The linear transformation

$$(x_1, \dots, x_n) \mapsto \left(\frac{x_1}{c_1}, \dots, \frac{x_n}{c_n} \right)$$

transforms the ellipsoid into the unit ball. The volume of every set in the new coordinates scales with the factor $1/\prod_i c_i$. Therefore the ratio of two volumes does not change, and the equidistribution in the ellipsoid is transformed into the equidistribution in the unit ball. Hyperplanes are transformed into hyperplanes and the property of linear separability is not affected by a nonsingular linear transformation. Thus, the following corollaries from Theorems 1 and 2 hold.

Corollary 1. Let \mathcal{M} be formed by a finite number of i.i.d. trials taken from the equidistribution in a n -dimensional ellipsoid E_n (16), and let \mathbf{x} be a test point drawn independently from the same distribution. Then \mathbf{x} can be separated from \mathcal{M} by a linear functional with probability $P_1(\mathcal{M}, n)$:

$$P_1(\mathcal{M}, n) \geq (1 - (1 - \varepsilon)^n) \left(1 - \frac{\rho(\varepsilon)^n}{2} \right)^M.$$

Corollary 2. Let \mathcal{M} be formed by a finite number of i.i.d. trials taken from the equidistribution in a n -dimensional ellipsoid E_n . With probability

$$P_M(\mathcal{M}, n) \geq \left[(1 - (1 - \varepsilon)^n) \left(1 - (M - 1) \frac{\rho(\varepsilon)^n}{2} \right) \right]^M$$

each point \mathbf{y} from the random set \mathcal{M} can be separated from $\mathcal{M} \setminus \{\mathbf{y}\}$ by a linear functional.

Dimension, n	2	5	10	20	30
<i>Ball</i> (theoretical estimate)	$3.7 \cdot 10^{-5}$	0.0364	0.4096	0.9455	0.9975
<i>Cube</i> min:	$4 \cdot 10^{-4}$	0.0089	0.2580	0.9408	0.9986
median:	$4 \cdot 10^{-4}$	0.0110	0.2737	0.9469	0.9992
max:	$5 \cdot 10^{-4}$	0.0137	0.2847	0.9511	1.0
<i>Gaussian</i> min:	$5 \cdot 10^{-4}$	0.0122	0.1403	0.7559	0.9792
median:	$10 \cdot 10^{-4}$	0.0153	0.1568	0.7698	0.9817
max:	0.0014	0.0183	0.1778	0.7836	0.9848

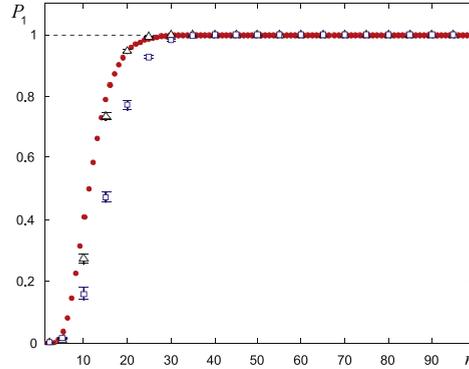


Fig. 6. Numerical and theoretical estimates of $P_1(\mathcal{M}, n)$ for various distributions and n . *Top panel*, top row: theoretical estimate of $P_1(\mathcal{M}, n)$ for equidistributions in n -balls derived in accordance with (2) in the statement of Theorem 1. *Left panel*, rows 2,3: numerical estimates $F_1(\mathcal{M}, n)$ of $P_1(\mathcal{M}, n)$ for both normal and equidistribution in an n -cube for various values of n . *Bottom panel*: solid circles show the values of theoretical estimates $P_1(\mathcal{M}, n)$, triangles show empirical means of $F_1(\mathcal{M}, n)$ for the samples drawn from equidistribution in the cube $[-1, 1]^n$, and squares correspond to empirical means of $F_1(\mathcal{M}, n)$ for the samples drawn from the Gaussian (normal) distribution. Whiskers in the plots indicate maximal and minimal values in of $F_1(\mathcal{M}, n)$ in each group of experiments.

4.2. Equidistributions in a cube and normal distributions

It is well known that, for n sufficiently large, samples drawn from an n -dimensional normal distribution [32] concentrate near a corresponding n -sphere. Similarly, samples generated from an equidistribution in an n -dimensional cube concentrate near a corresponding n -sphere too. In this respect, one might expect that the estimates derived in Theorems 1–3 hold for these distributions too, asymptotically in n . Our numerical experiments below illustrate that this is indeed the case.

The experiments were as follows. For each distribution (normal distribution in \mathbb{R}^n and equidistribution in the n -cube $[-1, 1]^n$) and a given n an i.i.d. sample \mathcal{M} of $M = 10^4$ vectors was drawn. For each vector \mathbf{y} in this sample we constructed the functional $l_{\mathbf{y}}(\mathbf{x}) = \langle \mathbf{y}, \mathbf{x} \rangle - \|\mathbf{y}\|^2$ (cf. the proof of Lemma 1). For each $\mathbf{y} \in \mathcal{M}$ and $\mathbf{x} \in \mathcal{M}$, $\mathbf{x} \neq \mathbf{y}$ the sign of $l_{\mathbf{y}}(\mathbf{x})$ was evaluated, and the total number N of instances when $l_{\mathbf{y}}(\mathbf{x}) < 0$ for all $\mathbf{x} \in \mathcal{M}$, $\mathbf{x} \neq \mathbf{y}$ was calculated. The latter number is a lower bound estimate of the number of points in \mathcal{M} that are linearly separable from the rest in the sample. This was followed by deriving the values of the success frequencies, $F_1(\mathcal{M}, n) = N/(M - 1)$. For each n the experiment was repeated 50 times. Outcomes of these experiments are presented in Fig. 6. As we can see from Fig. 6, despite that the samples are drawn from different distributions, for $n > 30$ these differences do not significantly affect point separability properties. This is due to pronounced influence of measure concentration in these dimensions.

4.3. One trial non-iterative learning

Our basic model of linear separation of a given query point $\mathbf{y} \in \mathcal{M}$ from any other $\mathbf{x} \in \mathcal{M}$, $\mathbf{x} \neq \mathbf{y}$ was

$$l_{\mathbf{y}}(\mathbf{x}) = \left\langle \frac{\mathbf{y}}{\|\mathbf{y}\|}, \mathbf{x} \right\rangle - \|\mathbf{y}\| < 0 \text{ for } \mathbf{x} \in \mathcal{M} \setminus \{\mathbf{y}\} \tag{17}$$

Deriving these functionals is a genuine one-shot procedure and does not require iterative learning. The construction, however, assumes that the distribution from which the sample \mathcal{M} is drawn is close in some sense to an equidistribution in the unit ball $B_n(1)$.

In more general cases, e.g. when sampling from the equidistribution in an ellipsoid, the functionals $l_{\mathbf{y}}(\mathbf{x})$ can be replaced with Fisher linear discriminants:

$$l_{\mathbf{y}}(\mathbf{x}) = \left\langle \frac{\mathbf{w}(\mathbf{y}, \mathcal{M})}{\|\mathbf{w}(\mathbf{y}, \mathcal{M})\|}, \mathbf{x} \right\rangle - c, \tag{18}$$

where

$$\mathbf{w}(\mathbf{y}, \mathcal{M}) = \text{Cov}(\mathcal{M})^{-1} \left(\mathbf{y} - \frac{1}{M-1} \sum_{\mathbf{x} \neq \mathbf{y}} \mathbf{x} \right),$$

and $\text{Cov}(\mathcal{M})$ is the non-singular covariance matrix of the sample \mathcal{M} , and c is a parameter. The value of c could be chosen as $c = \left(\frac{\mathbf{w}(\mathbf{y}, \mathcal{M})}{\|\mathbf{w}(\mathbf{y}, \mathcal{M})\|} \cdot \mathbf{y} \right)$. The procedure for generating separating functionals $l_{\mathbf{y}}(\mathbf{x})$ remains non-iterative, but it does require knowledge of the covariance matrix Σ .

Note that if \mathcal{M} is centered at $\mathbf{0}$ then $\frac{1}{M-1} \sum_{\mathbf{x} \neq \mathbf{y}} \mathbf{x}$ is approximately $\mathbf{0}$, and (18) reduces to (17) after the corresponding Mahalanobis transformation: $\mathbf{x} \mapsto \text{Cov}(\mathcal{M})^{-1/2} \mathbf{x}$. If the transformation $\mathbf{x} \mapsto \text{Cov}(\mathcal{M})^{-1/2} \mathbf{x}$ transforms the ellipsoid from which the sample \mathcal{M} is drawn into the unit ball then (18) becomes equivalent to (17), and separation properties of Fisher discriminants (18) follow in the same way as stated in Corollaries 1 and 2.

In addition, or as an alternative, whitening or decorrelation transformations could be applied to \mathcal{M} too. In case the covariance matrix $\text{Cov}(\mathcal{M})$ is singular or ill-conditioned, projecting the sample \mathcal{M} onto relevant Principal Components may be required.

4.4. Extreme selectivity of biological neurons

High separation capabilities of linear functionals (Theorems 1–3, Remark 3) may shed light on the puzzle associated with the so-called Grandmother [39,49] or Concept cells [38] in neuroscience. In the center of this puzzle is the wealth of empirical evidence that some neurons in human brain respond unexpectedly selective to particular persons or objects. Not only brain is able to respond selectively to rare individual stimuli but also such selectivity can be acquired very rapidly from limited number of experiences [25]. The question is: Why small ensembles of neurons may deliver such a sophisticated functionality reliably? As follows from Theorems 1–3, observed extreme selectivity of neuronal responses may be attributed to the very nature of basic physiological mechanisms in neurons involving calculation of linear functionals (1).

5. Examples

Let us now illustrate our theoretical results with examples of correcting legacy AI systems by linear functionals (1) and their combinations. As a legacy AI system we selected a Convolutional Neural Network trained to detect objects in images. Our choice of the legacy system was motivated by that these networks demonstrate remarkably strong performance on benchmarks and, reportedly, may outperform humans on popular classification tasks such as ILSVRC [22]. Despite this, as has been mentioned earlier, even for the most current state of the art CNNs false positives are not uncommon. Therefore, we investigate in the following experiments if it is possible to take one of these cutting edge networks and train a *one-neuron* filter to eliminate all the false positives produced. We will also look at what effect, if any, this filter will have on true positive numbers.

5.1. Training the CNN

For these experiments we trained the VGG-11 convolutional network to be our classifier [44]. Instead of the 1000 classes of Imagenet we trained it to perform the simpler task of identifying just one object class: pedestrians.

The VGG network was chosen for both its simple homogeneous architecture as well as its proven classification ability at recent Imagenet competitions [41]. We chose to train the VGG-11 A over the deeper 16 and 19 layer VGG networks due to hardware and time constraints.

Datasets: In order to train the network a set of 114,000 positive pedestrian and 375,000 negative non-pedestrian RGB images, re-sized to 128×128 , were collected and used as a *training set*. Our *testing set* comprised of further 10,000 positives and 10,000 negatives. The training and testing sets had no common elements.

The training procedure and choice of hyper parameters for the network follows largely from [44], but with some small adjustments to take into account the reduced number of classes being detected. We set momentum to 0.9 and used a mini batch size of 32.

Our initial learning rate was set to 0.00125 and this rate was reduced by a factor of 10 after 25 epochs and again after 50 epochs. Dropout regularisation was used for the first two fully connected layers with a ratio of 0.5. To initialise weights we used the Xavier initialisation [14] which we found helped training to converge faster. After 75 epochs, the validation accuracy started to fall and, we stopped the training at 75 epochs in order to avoid overfitting.

Training the network was done using the publicly available Caffe deep learning framework [27]. For the purposes of evaluating performance of the network, our null-hypothesis was that *no pedestrians are present in the scanning window*. True positive hence is an image crop (formed by the scanning window) containing a pedestrian and for which the CNN reports such a presence, false positive is an image crop for which the CNN reports presence of a pedestrian whilst none are present. True negatives and false negatives are defined accordingly. On the training set, the VGG-11 convolutional network showed 100% correct classification performance: the rates of true positives and true negatives were all equal to one.

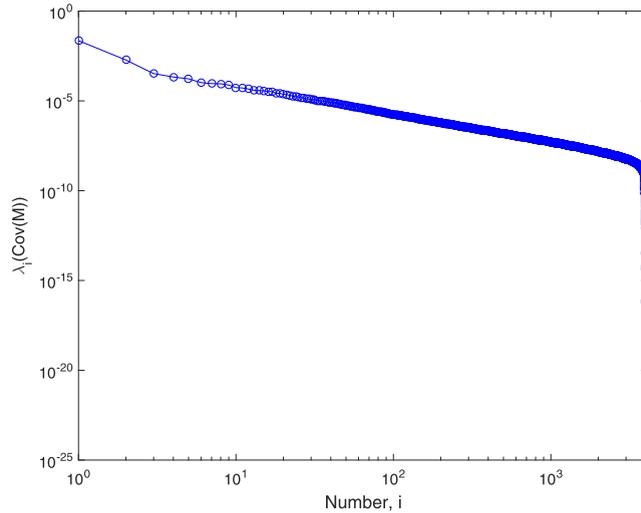


Fig. 7. Log-log plot of the eigenvalues $\lambda_i(\text{Cov}(\mathcal{M}))$ of the training data covariance matrix $\text{Cov}(\mathcal{M})$.

5.2. Training one-neuron legacy AI system correctors

5.2.1. Error detection and selection of features

First, a multi-scale sliding window approach is used on each video frame to provide proposals that can be run through the trained CNN. These proposals are re-sized to 128×128 and passed through the CNN for classification, non-maximum suppression is then applied to the positive proposals. Bounding boxes are drawn and we compared results to a ground truth so we can identify false positives.

For each positive and false positive (and its respective set of proposals before non-maximum suppression) we extracted the second to last fully connected layer from the CNN. These extracted feature vectors have dimension 4096. We used these feature vectors to construct *one-neuron* correctors filtering out the false positives that have been manually identified. These we will refer to in the text as 'trash models'.

5.2.2. Spherical cap model, (17)

The first and the most simple trash model filter that we constructed was the spherical cap model as specified by (17). In this model, after centering the training data, the query points y , adjusted for the mean of the training data, were the actual false positives detected.

5.2.3. Fisher cap model, (18)

Our second model type was the Fisher cap model (18). The original model involves the covariance matrix $\text{Cov}(\mathcal{M})$ corresponding to 4096 dimensional feature vectors of positives in the training set. Examination of this matrix revealed that the matrix is extremely ill-conditioned and hence (18) cannot be used unless the problem is regularized. To overcome the singularity issue the PCA-based dimensionality reduction was employed. To estimate the number of relevant principal components needed we performed the broken stick [33] and the Kaiser–Guttman [26] tests. The broken stick criterion returned the value of 13, and the Kaiser rule returned 45. The distribution of the eigenvalues of $\text{Cov}(\mathcal{M})$ is shown in Fig. 7. We used these data to select lower and upper bound estimates of the number of feasible principal components needed. For the purposes of our experiments these were chosen to be 50 and 2000, respectively.

5.2.4. Support vector machine (SVM)

As the ultimate linear separability benchmark, classical linear SVM model has been used. It was trained using the standard liblinear software package [10]³ on our two sets of normalised CNN feature vectors. No kernel transformations have been used in the process. When training the SVM models we used L_2 -regularized L_2 -loss function for the dual (option '-s1'), set termination tolerance to 0.001 (option '-e0.001') and bias parameter value to 1 (option '-B1'). Classes have been weighted and treated equally (options '-w11' and '-w-11'), and the value of the regularization parameter, parameter C in [10], was varying from 1 to 110 depending on the outcomes of classification (the value was set through the option '-c'). We found that as the number of false positives being trained on increased it was necessary to also increase the value of the C parameter to maintain a perfect separation of the training points.

³ The package is available at GitHub <https://github.com/cjlin1/liblinear>.

5.3. Implementation

At test time trained linear corrector models (trash model filters) were placed at the end of our detection pipeline, each one at a time. For any proposal given a positive score by our CNN we again extract the second to last fully connected layer. The CNN feature vectors from these positive proposals are then run through the trash model filter.

Any detection that then gives a positive score from the linear trash model is consequently removed by turning its detection score negative.

5.4. Results

Test videos: For our experiments we used the original training and testing sets as well as three different videos (not used in the training set generation) to test trash model creation and its effectiveness at removing false positives. The first sequence is the NOTTINGHAM video [6] that we created ourselves from the streets of Nottingham consisting of 435 frames taken with an action camera. The second video is the INRIA test set consisting of 288 frames [8], the third is the LINTHESCHER sequence produced by ETHZ consisting of 1208 frames [9]. For each test video a new trash model was trained.

5.4.1. Spherical cap model, (17)

To test performance of this naïve model we run the trained CNN classifier through the testing set and identified false positives. From the set of false positives, candidates were selected at random and correcting trash model filters (17) constructed. Performance of the resulting systems was then evaluated on the sets comprising of 114,000 positives (from the training set) and 1 negative (the false positive identified). These tests returned true negative rates equal to 1, as expected, and 0.978884 true positive rate averaged over 25 experiments. This is not surprising given that the model does not account for any statistical properties of the data apart from the assumption that the sample is taken from an equidistribution in $B_n(1)$. The latter assumption, as is particularly evident from Fig. 7, does not hold either. Equidistribution in an ellipsoid (or the multidimensional normal distribution, which is essentially the same) is a much better model of the data, and hence one would expect that Fisher discriminants (18) would perform better. We shall see now that this is indeed the case.

5.4.2. Fisher cap model, (18)

The Fisher discriminant is essentially the functional we use in theoretical estimate for distributions in ellipsoid. It is the same spherical cup, albeit after the whitening applied. The Fisher cap models were first constructed and tested on the data used to train and test the original CNN classifier. We started with the data projected onto the first 50 principal components. The 25 false positive candidates, taken from the testing set, were chosen at random as described in the previous section. They modeled *single* mistakes of the legacy classifier. The models were then prepared using the covariance matrices of the original training data set. On the data sets comprising of 114,000 positives from the training set and 1 negative corresponding to the single randomly chosen identified false positive all models delivered true negative and true positive rates equal to 1. The result persisted when the number of principal components used was increased to 2000.

In the next set of experiments we investigated numerically if a single neuron with the weights given by the Fisher discriminant formula would be capable of filtering out more than just one false positive candidate. We note that the question is not trivial even if the two sets are linearly separable: separability by Fisher linear discriminants imply linear separability, but the converse does not necessarily holds true. For this purpose we randomly selected several samples \mathcal{Y} of 25 false positive candidates from the training data set and constructed models (18) aiming at separating all 25 false positives in each sample from 114,000 true positives in the data set. The weights \mathbf{w} of model (18) depended on the sets of false positives \mathcal{FP} and the set \mathcal{TP} of true positives as follows:

$$\mathbf{w}(\mathcal{FP}, \mathcal{TP}) = (\text{Cov}(\mathcal{TP}) + \text{Cov}(\mathcal{FP}))^{-1} \left(\frac{1}{|\mathcal{FP}|} \sum_{\mathbf{y} \in \mathcal{FP}} \mathbf{y} - \frac{1}{|\mathcal{TP}|} \sum_{\mathbf{x} \in \mathcal{TP}} \mathbf{x} \right)$$

where $\text{Cov}(\mathcal{TP})$, $\text{Cov}(\mathcal{FP})$ are the empirical covariance matrices of true and false positives, respectively, and the value of parameter c was chosen as

$$c = \min_{\mathbf{y} \in \mathcal{FP}} \left\langle \frac{\mathbf{w}(\mathcal{FP}, \mathcal{TP})}{\|\mathbf{w}(\mathcal{FP}, \mathcal{TP})\|}, \mathbf{y} \right\rangle.$$

We did not observe perfect separation if only the first 50 principal components were used to construct the models. Perfect separation, however, was persistently observed when the number of first principle components being used exceeded 87. For data projected on more than the first 87 principal components one neuron with weights selected by the Fisher linear discriminant formula corrected 25 errors without doing any damage to classification capabilities (original skills) of the legacy AI system on the training set.

The difference in performance with dimension can be explained by that the hyperplane corresponding to Fisher linear discriminants for the 25 randomly chosen false positives appears to be closer to the origin than each individual data point in this set. In terms of the theoretical argument presented in Theorems 1 and 2 this can be seen as an increase in the value of ε corresponding to the test point. This increase leads to the increase of the value of $\rho(\varepsilon)^n/2$ in (6) and results in reduced values of the probability of separation. When the dimensionality of the data, n , grows, the term $\rho(\varepsilon)^n/2$ decreases mitigating

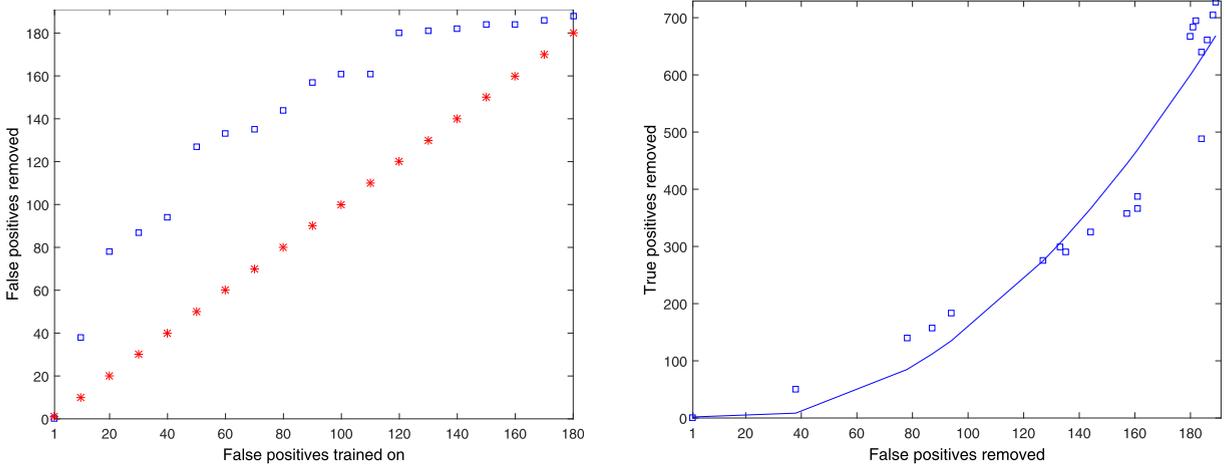


Fig. 8. Performance of *one-neuron* corrector built using Fisher cap model (18). *Left panel:* the number of false positives removed as a function of the number of false positives the model was built on. Stars indicate the number of false positives used for building the model. Squares correspond to the number of false positives removed by the model. *Right panel:* the number of true positives removed as a function of the number of false positives removed. The actual measurements are shown as squares. Solid line is the least-square fit of a quadratic.

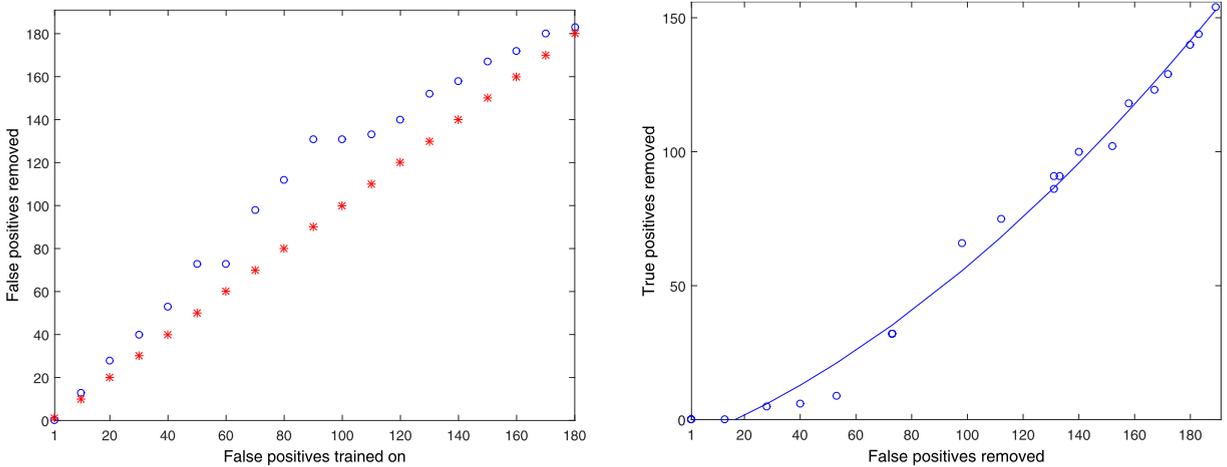


Fig. 9. Performance of *one-neuron* corrector constructed using the linear SVM model. *Left panel:* the number of false positives removed as a function of the number of false positives the model was trained on. Stars indicate the number of false positives used for training the model. Circles correspond to the number of false positives removed by the model. *Right panel:* the number of true positives removed as a function of the number of false positives removed. The actual measurements are presented as circles. Solid line is the least-square fit of a quadratic.

the effect. This is consistent with what we observed in this experiment. We expect that the observed negative consequences affecting linear separability of multiple points could be addressed via employing a two-neuron separation as well.

In the third group of tests we built single neuron trash models (18) on varying numbers of false positives from the NOTTINGHAM video. The true positives (totaling to 2896) and the false positives (189) in this video differed from those in the training set. The original training data set was projected onto the first 2000 principal components. Typical performance of the legacy CNN system with the corrector is shown in Fig. 8. As we can see from this figure, single false positives are dealt with successfully without any detrimental affects on the true positive rates. Increasing the number of false positives that the single neuron trash model is to cope with resulted in gradual deterioration of the true positive rates. One may expect that, according to Theorem 3, increasing of the number of neurons in trash models could significantly improve the power of correctors. Conducting these experiments, however, was outside of the main focus of this example.

5.4.3. SVM model

The Spherical cap and the Fisher cap models are by no means the ultimate tests of separability. Their main advantage is that they are purely non-iterative. To test extremal performance of linear functionals more sophisticated methods for their construction, such as e.g. SVMs, are needed. To assess this capability we trained SVM trash models on varying numbers of false positives from the NOTTINGHAM video, as a benchmark. Results are plotted in Fig. 9 below. The true positives were taken from the CNN training data set. The SVM trash model successfully removes 13 false positives without affecting the

Table 1

CNN performance on the INRIA video with and without trash model filtering.

	Without trash model	With SVM trash model
True positives	490	489
False positives	31	0

Table 2

CNN performance on the ETH LINTHESCHER video with and without trash model filtering.

	Without trash model	With SVM trash model
True positives	4288	4170
False positives	9	0

true positives rate. Its performance deteriorates at a much slower rate than for the Fisher cap model. This, however, is balanced by the fact that the Fisher cap model removed significantly more false positives than it was trained on.

Similar level of performance was observed for other testing videos, including INRIA and ETH LINTHESCHER videos. The results are provided in [Tables 1](#) and [2](#) below.

5.5. Results summary

From the results in [Figs. 8](#) and [9](#) and [Tables 1](#) and [2](#) we have demonstrated that it is possible to train a simple single-neuron corrector capable of rectifying mistakes of legacy AI systems. As can be seen from [Figs. 8](#) and [9](#) this filtering of false positives comes at a varying cost to the true positive detections. As we increase the number of false positives that we train the trash model on we see the number of true positives removed starts to rise as well. Our tests illustrate, albeit empirically, that the separation theorems stated in this paper are viable in applications. What is particularly remarkable is that we were able to remove more than 10 false positives at no cost to true positive detections in the NOTTINGHAM video by the use of a single linear function.

From practical perspectives, the main operational domain of the correctors is likely to be near the origin along the ‘removed false positives’ axis. In this domain generalization capacities of Fisher discriminant are significantly higher than those of the SVMs. As can be seen from [Figs. 8](#) and [9](#), left panels, Fisher discriminants are also efficient in ‘cutting off’ significantly more mistakes than they have been trained on. In this operational domain near the origin, they do so without doing much damage to the legacy AI system. SVMs perform better for larger values of false positives. This is not surprising as their training is iterative and accounts for much more information about the data than simply sample covariance and mean.

6. Conclusion

Stochastic separation theorems allowed us to create non-iterative one-trial correctors for legacy AI systems. These correctors were tested on benchmarks and on real-life industrial data and outperformed the theoretical estimates. In these experiments, already one-neuron correctors fixed more than ten independent mistakes without doing any damage to existing skills of the legacy AI system. Not only these experiments illustrated application of the approach, but also they revealed importance of whitening and regularization for successful correction. In order to reduce the impact on true positive detections a two neuron cascade classifier could be trained as a corrector. Further elaboration of the theory of non-iterative one-trial correctors for legacy AI systems and testing them in applications will be the subject of our future work.

The stochastic separation theorems are proven for large sets of randomly drawn i.i.d. data in high dimensional space. Each point can be separated from the rest of the sample by a hyperplane, with high probability. In a convex hull of a random finite set, all points of the set are vertexes (with high probability). These results are valid even for exponentially large samples. The estimates of the separation probability and allowable sample size age given by [Eqs. \(2\)](#), [\(8\)](#), [\(9\)](#), and [\(12\)](#). It is not much surprising that the separation by small ensembles of uncorrelated neurons is much more efficient than the separation by single neurons ([Fig. 4](#)).

We showed that high dimensionality of data can play a major and positive role in various machine learning and data analysis tasks, including problems of separation, filtration, and selection.

In contrast to naive intuition suggesting that high dimensionality of data more often than not brings in additional complexity and uncertainty, our findings contribute to the idea that high data dimensionality may constitute a valuable blessing too. The term ‘blessing of dimensionality’ was used in [\[1\]](#) to describe separation by Gaussians mixture models or spheres, which are equivalent in high dimension. This blessing, formulated here in the form of several stochastic separation theorems, offers several new insights for big data analysis. This is a part of measure concentration phenomena [\[18,19,34\]](#) which form the background of classical statistical physics (Gibbs theorem about equivalence of microcanonic and canonic ensembles [\[13\]](#)) and asymptotic theorems in probability [\[32\]](#). In machine learning, these phenomena are in the background of

the random projection methods [28], learning large Gaussian mixtures [1], and various randomized approaches to learning [42,50] and bring light in the theory of approximation by random bases [16]. It is highly probable that the recently described manifold disentanglement effects [5] are universal in essentially high dimensional data analysis and relate to essential multidimensionality rather than to specific deep learning algorithms. This may be the manifestation of concentration near the parts of the boundary with the lowest curvature.

The plan of further technical development of the theory and methods presented in the paper is clear:

- To prove the separation theorems in much higher generality: for sufficiently regular essentially multidimensional probability distribution and a random set of weakly dependent (or independent) samples every point is linearly separated from the rest of the set with high probability.
- To find explicit estimation of the allowed sample size in general case.
- To elaborate the theorems and estimates in high generality for separation by small neural networks.
- To develop and test generators of correctors for various standard AI tasks.

But the most inspiring consequence of the measure concentration phenomena is the paradigm shift. It is a common point of view that the complex learning systems should produce *complex* knowledge and skills. On contrary, it seems to be possible that the main function of many learning system, both technical and biological, in addition to production of *simple* skills, is a special preprocessing. They transform the input flux ('reality') into essentially multidimensional and quasi-random distribution of signals and images plus, may be, some simple low dimensional and more regular signal. After such a transformation, ensembles of non-interacting or weakly interacting small neural networks ('correctors' of simple skills) can solve complicated problems.

References

- [1] J. Anderson, M. Belkin, N. Goyal, L. Rademacher, J. Voss, The more, the merrier: the blessing of dimensionality for learning large Gaussian mixtures, *J. Mach. Learn. Res. :Workshop and Conference Proceedings.*, 35 (2014) 1–30.
- [2] K. Ball, An elementary introduction to modern convex geometry, *Flavors Geom.* 31 (1997) 1–58.
- [3] K. Bennett, Legacy systems: coping with success, *IEEE Softw.* 12 (1) (1995) 19–23.
- [4] J. Bisbal, D. Lawless, B. Wu, J. Grimson, Legacy information systems: issues and directions, *IEEE Softw.* 16 (5) (1999) 103–111.
- [5] P.P. Brahma, D. Wu, Y. She, Why deep learning works: a manifold disentanglement perspective, *IEEE Trans. Neural Netw. Learn. Syst.* 27 (10) (2016) 1997–2008.
- [6] R. Burton, Nottingham video, 2016, A test video for pedestrians detection taken from the streets of Nottingham by an action camera.
- [7] T. Chen, M. Li, Y. Li, M. Lin, N. Wang, M. Wang, T. Xiao, B. Xu, C. Zhang, Z. Zhang, Mxnet: a flexible and efficient machine learning library for heterogeneous distributed systems, 2015, (<https://github.com/dmlc/mxnet>).
- [8] N. Dalal, B. Triggs, Histograms of oriented gradients for human detection, in: *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, 2005, pp. 886–893.
- [9] A. Ess, B. Leibe, K. Schindler, L. van Gool, A mobile vision system for robust multi-person tracking, in: *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, 2008, pp. 1–8. DOI: 10.1109/CVPR.2008.4587581.
- [10] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, C.-J. Lin, Liblinear: a library for large linear classification, *J. Mach. Learn. Res.* 9 (2008) 1871–1874. <https://github.com/cjlin1/liblinear>.
- [11] R.J. Fricker, False positives are statistically inevitable, *Science* 351 (2016) 569–570.
- [12] C. Gear, I. Kevrekidis, Constraint-defined manifolds: a legacy code approach to low-dimensional computation, *J. Sci. Comput.* 25 (1) (2005) 17–28.
- [13] J. Gibbs, *Elementary Principles in Statistical Mechanics*, Developed With Especial Reference to the Rational Foundation of Thermodynamics, Dover Publications, New York, 1960 (1902).
- [14] X. Glorot, Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in: *Proc. of the 13th International Conference on Artificial Intelligence and Statistics (AISTATS)*, 9, 2010, pp. 249–256.
- [15] A. Gorban, Order-disorder separation: geometric revision, *Physica A* 374 (2007) 85–102.
- [16] A. Gorban, I. Tyukin, D. Prokhorov, K. Sofeikov, Approximation with random bases: Pro et Contra, *Inf. Sci.* 364–365 (2016) 129–145.
- [17] A. Gorban, I. Tyukin, I. Romanenko, The blessing of dimensionality: separation theorems in the thermodynamic limit, 2016b, A talk given at TFMST 2016, 2nd IFAC Workshop on Thermodynamic Foundations of Mathematical Systems Theory. September 28–30, 2016, Vigo, Spain.
- [18] M. Gromov, *Metric Structures for Riemannian and non-Riemannian Spaces*. With Appendices by M. Katz, P. Pansu, S. Semmes. Translated from the French by Sean Michael Bates, Birkhauser, Boston, MA, 1999.
- [19] M. Gromov, Isoperimetry of waists and concentration of maps, *GAFa, Geomter. Funct. Anal.* 13 (2003) 178–215.
- [20] P. Halmos, *Finite Dimensional Vector Spaces*, Undergraduate Texts in Mathematics, Springer, 1974.
- [21] L.K. Hansen, P. Salamon, Neural network ensembles, *IEEE Trans. Pattern Anal. Mach. Intell.* 12 (10) (1990) 993–1001.
- [22] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 770–778. ArXiv: 1512.03385.
- [23] T.K. Ho, Random decision forests, in: *Proc. of the 3rd International Conference on Document Analysis and Recognition*, 1995, pp. 993–1001.
- [24] T.K. Ho, The random subspace method for constructing decision forests, *IEEE Trans. Pattern Anal. Mach. Intell.* 20 (8) (1998) 832–844.
- [25] M. Ison, R. Quian Quiroga, I. Fried, Rapid encoding of new memories by individual neurons in the human brain, *Neuron* 87 (1) (2015) 220–230.
- [26] D. Jackson, Stopping rules in principal components analysis: a comparison of heuristical and statistical approaches, *Ecology* 74 (8) (1993) 2204–2214.
- [27] Y. Jia, Caffe: An open source convolutional architecture for fast feature embedding, 2013, (<http://caffe.berkeleyvision.org/>).
- [28] W.B. Johnson, J. Lindenstrauss, Extensions of Lipschitz mappings into a Hilbert space, *Contemp. Math.* 26 (1) (1984) 189–206.
- [29] M. Krein, D. Milman, On extreme points of regular convex sets, *StudiaMath* 9 (1940) 133–138.
- [30] A. Krizhevsky, I. Sutskever, G. Hinton, Imagenet classification with deep convolutional neural networks, in: F. Pereira, C.J.C. Burges, L. Bottou, K.Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems 25*, Curran Associates, Inc., 2012, pp. 1097–1105.
- [31] A. Kuznetsova, S. Hwang, B. Rosenhahn, L. Sigal, Expanding object detectors horizon: incremental learning framework for object detection in videos, in: *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 28–36.
- [32] P. Lévy, *Problèmes Concrets d'analyse Fonctionnelle*, second ed., Gauthier-Villars, Paris, 1951.
- [33] R. Macarthur, On the relative abundance of bird species, *Proc. Natl. Acad. Sci.* 43 (3) (1957) 293–295.
- [34] V.D. Milman, G. Schechtman, Asymptotic theory of finite dimensional normed spaces: Isoperimetric inequalities in Riemannian manifolds, *Lecture Notes in Mathematics*, 1200, Springer, 2009.
- [35] I. Misra, A. Shrivastava, M. Hebert, Semi-supervised learning for object detectors from video, in: *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 3594–3602.

- [36] A. Nguyen, J. Yosinski, J. Clune, Deep neural networks are easily fooled: High confidence predictions for unrecognizable images, in: Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015, pp. 427–436.
- [37] A. Prest, C. Leistner, J. Civera, C. Schmid, V. Ferrari, Learning object class detectors from weakly annotated video, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2012, pp. 3282–3289.
- [38] R. Quiñan Quiroga, Concept cells: the building blocks of declarative memory functions, *Nat. Rev. Neurosci.* 13 (8) (2012) 587–597.
- [39] R. Quiñan Quiroga, L. Reddy, G. Kreiman, C. Koch, I. Fried, Invariant visual representation by single neurons in the human brain, *Nature* 435 (7045) (2005) 1102–1107.
- [40] W. Rudin, *Functional Analysis, International Series in Pure and Applied Mathematics*, McGraw-Hill, 1991.
- [41] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A.C. Berg, L. Fei-Fei, Imagenet large scale visual recognition challenge, *Int. J. Comput. Vis.* (2014) 1–42. DOI: [10.1007/s11263-015-0816-y](https://doi.org/10.1007/s11263-015-0816-y).
- [42] S. Scardapane, D. Wang, Randomness in neural networks: an overview, *Wiley Interdiscip. Rev.* 7 (2) (2017) e1200.
- [43] H. Schaefer, *Topological Vector Spaces*, Springer, New York, 1999.
- [44] K. Simonyan, A. Zisserman, Very deep convolutional networks for large-scale image recognition, in: International Conference on Learning Representations, 2015. ArXiv: [1409.1556](https://arxiv.org/abs/1409.1556).
- [45] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I.J. Goodfellow, R. Fergus, Intriguing properties of neural networks, in: Proc. of International Conference on Learning Representations (ICLR), 2014. <https://arxiv.org/abs/1312.6199>.
- [46] D.D. Team, DeepLearning4j: open-source distributed deep learning for the JVM, 2016, (Apache Software Foundation License 2.0. <http://deeplearning4j.org>).
- [47] V. Vapnik, *The Nature of Statistical Learning Theory*, Springer-Verlag, 2000.
- [48] V. Vapnik, O. Chapelle, Bounds on error expectation for support vector machines, *Neural Comput.* 12 (9) (2000) 2013–2036.
- [49] I. Viskontas, R. Quiñan Quiroga, I. Fried, Human medial temporal lobe neurons respond preferentially to personally relevant images, *Proc. Nat. Acad. Sci.* 106 (50) (2009) 21329–21334.
- [50] D. Wang, Editorial: randomized algorithms for training neural networks, *Inf. Sci.* 364–365 (2016) 126–128.
- [51] S. Zheng, Y. Song, T. Leung, I. Goodfellow, Improving the robustness of deep neural networks via stability training, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2016. <https://arxiv.org/abs/1604.04326>.