

The feasibility and inevitability of stealth attacks

IVAN Y. TYUKIN*

Department of Mathematics, King's College London, Strand, London WC2R 2LS, UK

*Corresponding author: ivan.tyukin@kcl.ac.uk

DESMOND J. HIGHAM

School of Mathematics, University of Edinburgh, Peter Guthrie Tait Road, Edinburgh EH9 3FD, UK

AND

ALEXANDER BASTOUNIS, ELIYAS WOLDEGEORGIS AND ALEXANDER N. GORBAN

School of Computing and Mathematical Sciences, University of Leicester, University Road, Leicester LE1 7RH, UK

[Received on 28 December 2022; revised on 13 July 2023; accepted on 12 October 2023]

We develop and study new adversarial perturbations that enable an attacker to gain control over decisions in generic Artificial Intelligence (AI) systems including deep learning neural networks. In contrast to adversarial data modification, the attack mechanism we consider here involves alterations to the AI system itself. Such a *stealth attack* could be conducted by a mischievous, corrupt or disgruntled member of a software development team. It could also be made by those wishing to exploit a ‘democratization of AI’ agenda, where network architectures and trained parameter sets are shared publicly. We develop a range of new implementable attack strategies with accompanying analysis, showing that with high probability a stealth attack can be made transparent, in the sense that system performance is unchanged on a fixed validation set which is unknown to the attacker, while evoking any desired output on a trigger input of interest. The attacker only needs to have estimates of the size of the validation set and the spread of the AI’s relevant latent space. In the case of deep learning neural networks, we show that a *one-neuron attack* is possible—a modification to the weights and bias associated with a single neuron—revealing a vulnerability arising from over-parameterization. We illustrate these concepts using state-of-the-art architectures on two standard image data sets. Guided by the theory and computational results, we also propose strategies to guard against stealth attacks.

Keywords: Adversarial attacks; AI instability; concentration of measure; backdoor attacks; AI robustness.

1. Introduction

It is widely recognized that Artificial Intelligence (AI) systems can be vulnerable to adversarial attacks (Szegedy *et al.*, 2014), that is, small, often imperceptible, perturbations that exploit instabilities. The danger of an attacker gaining control of an automated decision-making process is of particular concern in high-stakes or safety-critical settings, including medical imaging (Finlayson *et al.*, 2019; Ghassemi *et al.*, 2021), transport (Wu *et al.*, 2020) and textual analysis (Ebrahimi *et al.*, 2018). The last few years have therefore seen an escalation in the design of both attack and defence strategies (Akhtar & Mian, 2018; Huang *et al.*, 2020; Shafahi *et al.*, 2020), and recent work has considered the bigger question of whether the existence of instabilities and successful attacks is inevitable (Shafahi *et al.*, 2019; Antun *et al.*, 2020; Tyukin *et al.*, 2020; Bastounis *et al.*, 2021b; Colbrook *et al.*, 2022).

Work in this field has focused on adversarial perturbations to the training or input data, and may be classified as black-box, grey-box or white-box, depending on the level of information and access available to the attacker. We look at a different setting, which was introduced in Tyukin *et al.* (2020) and is motivated further in Subsection 2.1.

Here, the attacker is assumed to have control of the AI system, in the sense of being able to edit code or change parameters. We therefore consider attacks that make perturbations to the system itself in an attempt to change the output on a specific input point of interest to the attacker *without any changes* in the AI’s performance on a large, unknown, validation set and *without retraining*. We focus on two issues: (1) strategies for constructing such attacks and (2) theoretical analysis of the likelihood of success.

In Section 2 we present a general framework of open-box stealth attack and discuss additional related work. In Section 3 we reveal a *new class of attacks* (Algorithms 1–3) and establish probability bounds of their success (Theorems 1–3).

Algorithm 1 and the corresponding Theorem 1 consider the case when the Attacker’s choice of triggers is limited only by the requirement that the triggers’ feature representations stay within a ball containing the Owner’s validation set (choosing triggers with feature representations outside of this ball always results in successful attacks). In this case, we show that the likelihood of attack success can be made *arbitrarily* close to one, *for any fixed dimension* of the model’s latent space (Theorem 1, Remark 4). This is a significant departure from the previous state of knowledge, as success likelihoods for such attacks were thought to be limited by dimension (Tyukin *et al.*, 2020). To establish these high probabilities of success, the attack must be executed with arbitrarily high accuracy and the model must satisfy appropriate reachability conditions (Huang *et al.*, 2020).

Algorithm 2 and Theorem 2 relate to approaches that enable attack triggers to be camouflaged as legitimate data by requesting that the Attacker’s triggers produce latent representations within some given neighbourhood of those corresponding to specified inputs.

Algorithm 3 and Theorem 3 consider the case where the Attackers’ capability to change or explore feature spaces of the model is constrained to some finite number of attributes or a smaller dimensional subspace. The case is motivated by the ideas from Beerens & Higham (2023) where the authors proposed methods to generate adversarial perturbations confined to smaller dimensional subspaces of the original input space (in contrast to the stealth attack setting considered here). This scenario enables attacks for models with sparse data representations in latent spaces. Remarkably, these constrained attacks may have significant probability of success even when the accuracy of their implementation is relatively low; see Theorem 3.

Section 4 presents experiments which illustrate the application of the theory to realistic settings and demonstrates the strikingly likely feasibility of *one-neuron attacks*, which alter weights of just a single neuron. Section 5 concludes with recommendations on how vulnerabilities we exposed in this work can be mitigated by model design practices. Proofs of the theorems can be found in the Appendix, along with extra algorithmic details and computational results.

2. Stealth attacks

2.1 General framework

Consider a generic AI system, a map

$$\mathcal{F}: \mathcal{U} \rightarrow \mathbb{R} \quad (2.1)$$

producing some decisions on its outputs in response to an input from \mathcal{U} . The map \mathcal{F} can define input–output relationships for an entire deep neural network or some part (a sub-graph), an ensemble of networks, a tree or a forest and the set \mathcal{U} could be a subset of \mathbb{R}^m for image classification or some set of possible sentences for text analysis. For the purposes of our work, the AI system’s specific task is not relevant and can include classification, regression or density estimation.

In the classification case, if there are multiple output classes then we regard (2.1) as representing the output component of interest—we consider changes to \mathcal{F} that do not affect any other output components; this is the setting in which our computational experiments are conducted. The flexibility for stealth attacks to work independently of the choice of output component is a key feature of our work.

The AI system has an Owner operating the AI. An Attacker wants to take advantage of the AI by forcing it to make decisions in their favour. Conscious about security, the Owner created a *validation set* which is kept secret. The validation set is a list of input–output pairs produced by the uncompromised system (2.1). The Owner can monitor security by checking that the AI reproduces these outputs. Now, suppose that the Attacker has access to the AI system but not the validation set. The phrase *stealth attack* was used in Tyukin *et al.* (2020) to describe the circumstance where the Attacker chooses a *trigger input* and *modifies the AI* so that

- the Owner could not detect this modification by testing on the validation set,
- on the trigger input the modified AI produces the output desired by the Attacker.

Figure 1, panel C, gives a schematic representation of this set-up. The set-up is different from other known attack types such as adversarial attacks in which the Attacker exploits access to AI to compute imperceptible input perturbations altering AI outputs (shown in Fig. 1, panel A), and data poisoning attacks (Fig. 1, panel B) in which the Attacker exploits access to AI training sets to plant triggers directly.

We note that the stealth attack setting is relevant to the case of a corrupt, disgruntled or mischievous individual who is a member of a software development team that is creating an AI system, or who has an IT-related role where the system is deployed. In this scenario, the attacker will have access to the AI system and, for example, in the case of a deep neural network, may choose to alter the weights, biases or architecture. The scenario is also pertinent in contexts where AI systems are exchanged between parties, such as

- ‘democratization of AI’ (Allen *et al.*, 2019), where copies of large-scale models and parameter sets are made available across multiple public domain repositories,
- transfer learning (Neyshabur *et al.*, 2020), where an existing, previously trained tool is used as a starting point in a new application domain,
- outsourced cloud computing, where a third party service conducts training (Ghodsi *et al.*, 2017).

2.2 Formal definition of stealth attacks

Without loss of generality, it is convenient to represent the initial general map (2.1) as a composition of two maps, F and Φ :

$$\mathcal{F} = F \circ \Phi, \quad \text{where } F : \mathbb{R}^n \rightarrow \mathbb{R}, \Phi : \mathcal{U} \rightarrow \mathbb{R}^n. \quad (2.2)$$

The map Φ defines general *latent* representation of inputs from \mathcal{U} , whereas the map F can be viewed as a *decision-making* part of the AI system. In the context of deep learning models, latent representations can

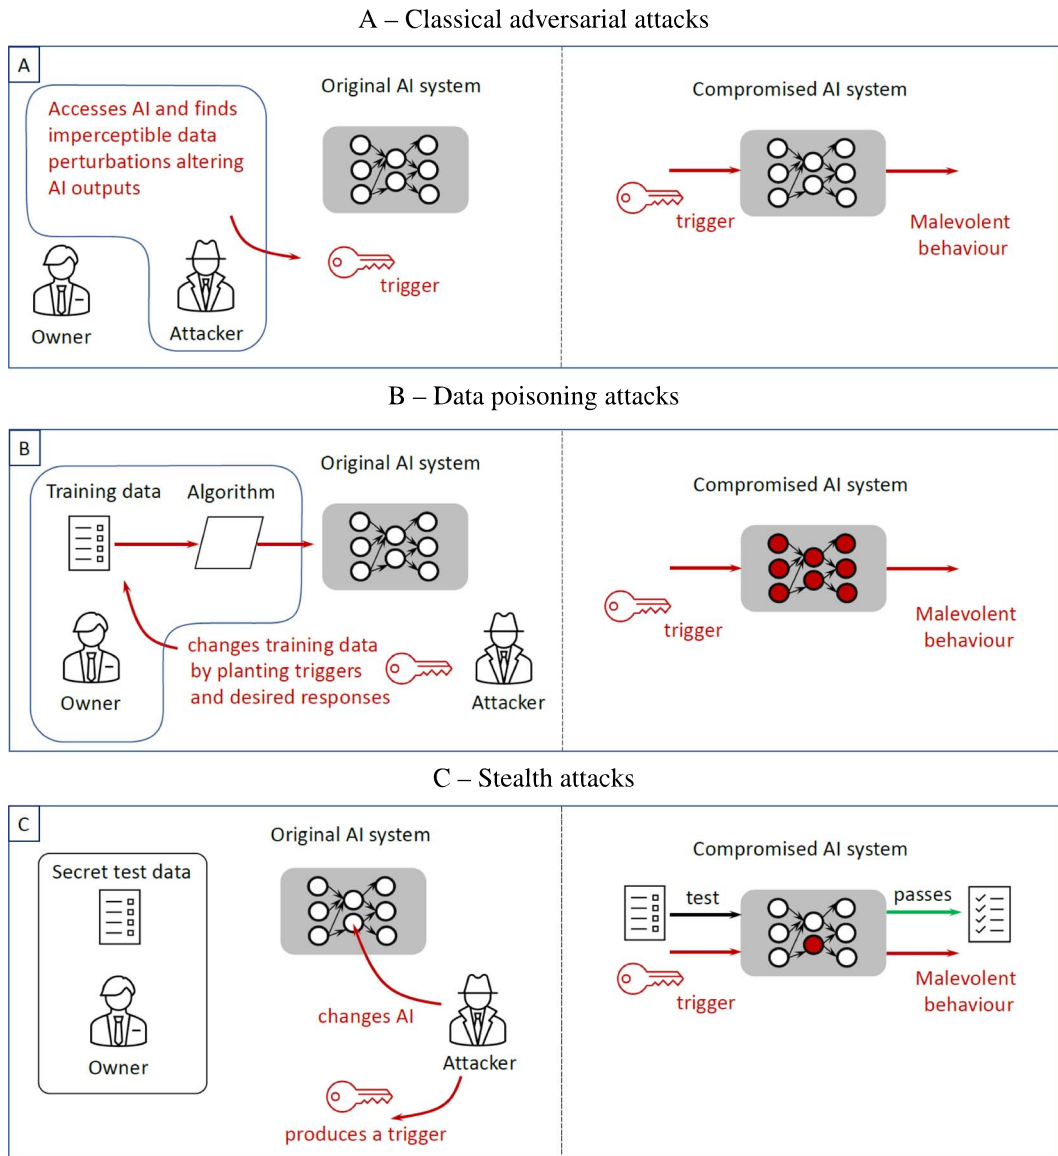


FIG. 1. (Colour online) General schemes of adversarial (panel A), data poisoning (panel B) and stealth attacks (panel C).

be outputs of hidden layers in deep learning neural networks, and decision-making parts could constitute operations performed by fully connected and softmax layers at the end of the networks. If Φ is an identity map then setting $F = \mathcal{F}$ brings us to the initial case (2.1).

An additional advantage of the compositional representation (2.2) is that it enables explicit modelling of the *focus of the adversarial attack*—a part of the AI system subjected to adversarial modification. This part will be modelled by the map F .

A perturbed, or attacked, map F_a is defined as

$$\begin{aligned} F_a : \mathbb{R}^n \times \Theta &\rightarrow \mathbb{R}, \\ F_a(\cdot, \theta) &= F(\cdot) + \mathfrak{A}(\cdot, \theta), \end{aligned} \quad (2.3)$$

where the term $\mathfrak{A} : \mathbb{R}^n \times \Theta \rightarrow \mathbb{R}$ models the *effect* of an adversarial perturbation, and $\Theta \subset \mathbb{R}^m$ is a set of relevant parameters. Such adversarial perturbations could take different forms including modification of parameters of the map F and physical addition or removal of components involved in computational processes in the targeted AI system. In neural networks the parameters are the weights and biases of a neuron or a group of neurons, and the components are neurons themselves. As we shall see later, a particularly instrumental case occurs when the term \mathfrak{A} is just a single Rectified Linear Unit (ReLU function) (Hahnloser *et al.*, 2000), and $\theta = (\mathbf{w}, b)$ represents the weights and bias,

$$\mathfrak{A}(\cdot, \mathbf{w}, b) = D \operatorname{ReLU}(\langle \cdot, \mathbf{w} \rangle - b) \quad \text{where} \quad \operatorname{ReLU}(s) = \max\{s, 0\}, \quad (2.4)$$

or a sigmoid (see Gorban & Rossiev (1996) and Higham & Higham (2019) for further information on activation functions of different types)

$$\mathfrak{A}(\cdot, \mathbf{w}, b) = D \sigma(\langle \cdot, \mathbf{w} \rangle - b), \quad \text{where} \quad \sigma(s) = \frac{1}{1 + \exp(-s)}, \quad (2.5)$$

where $D \in \mathbb{R}$ is a constant gain factor.

Having introduced the relevant notation, we are now ready to provide a formal statement of the problem of stealth attacks introduced in Section 2.1.

PROBLEM 1 (ε - Δ Stealth Attack on \mathcal{F}). Consider a map \mathcal{F} defined by (2.1), (2.2). Suppose that an owner of the AI system has a finite validation (or verification) set $\mathcal{V} \subset \mathcal{U}$. The validation set \mathcal{V} is kept secret and is assumed to be *unknown* to an attacker. The cardinality of \mathcal{V} is denoted by M , and this value is not necessarily known to the attacker.

Given $\varepsilon \geq 0$ and $\Delta > 0$, a successful ε - Δ stealth attack takes place if the attacker modifies the map F in \mathcal{F} and replaces it by F_a constructed such that for some $\mathbf{u}' \in \mathcal{U}$, known to the attacker but unknown to the owner of the map \mathcal{F} , the following properties hold:

$$\begin{aligned} |F \circ \Phi(\mathbf{u}) - F_a \circ \Phi(\mathbf{u})| &\leq \varepsilon \text{ for all } \mathbf{u} \in \mathcal{V} \\ |F \circ \Phi(\mathbf{u}') - F_a \circ \Phi(\mathbf{u}')| &\geq \Delta. \end{aligned} \quad (2.6)$$

In words, when F is perturbed to F_a the output is changed by no more than ε on the validation set, but is changed by at least Δ on the trigger, \mathbf{u}' . We note that this definition does not require any notion of whether the classification of \mathbf{u} or \mathbf{u}' is ‘correct’. In practice we are interested in cases where Δ is sufficiently large for \mathbf{u}' to be assigned to different classes by the original and perturbed maps.

REMARK 1. (The target class for the trigger \mathbf{u}'). Note that the above setting can be adjusted to fit a broad range of problems, including general multi-class problems. Crucially, the stealth attacks proposed in this paper allow the attacker to choose which class the modified AI system F_a predicts for the trigger image \mathbf{u}' . We illustrate these capabilities with examples in Section 4.

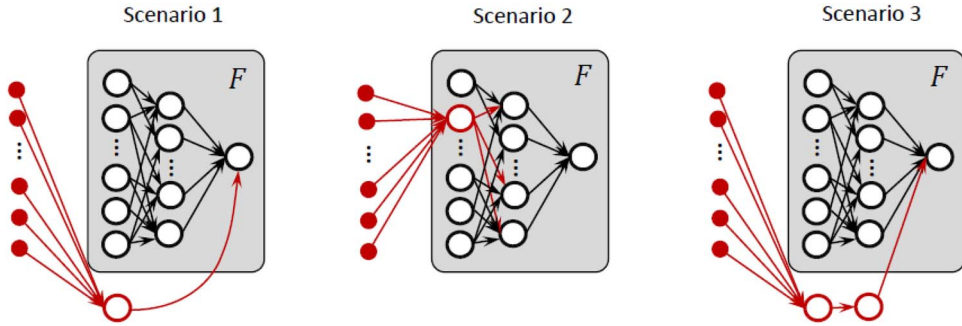


FIG. 2. (Colour online) Stealth attack implementation patterns. The red open circles and arrows indicate changes. The solid red circles show the input components, and how they feed into the new or perturbed neurons. In Scenarios 1 and 3 neuron(s) are added. In Scenario 2, a one-neuron attack, the weights and biases of an existing neuron are replaced with new values. The grey boxes show a part of the network modelled by the map F .

2.3 Related work

Adversarial attacks. A broad range of methods aimed at revealing vulnerabilities of state-of-the-art AI, and deep learning models in particular, to adversarial inputs has been proposed to date (see e.g. recent reviews [Huang et al. \(2020\)](#); [Ren et al. \(2020\)](#)). The focus of this body of work has been primarily on perturbations to signals/data processed by the AI. In contrast to this established framework, here we explore possibilities to determine and implement small *changes to AI structure* and *without retraining*.

Data poisoning. Gu et al. ([Gu et al., 2017](#)) (see also [Biggio & Roli \(2018\)](#) and references therein, and [Manoj & Blum \(2021\)](#) for explicit upper bounds on the volume of poisoned data that are sufficient to execute these attacks) showed how malicious data poisoning occurring, for example, via outsourcing of training to a third party can lead to *backdoor* vulnerabilities. Performance of the modified model on the validation set, however, is not required to be perfect. A data poisoning attack is deemed successful if the model's performance on the validation set is within some margin of the user's expectations.

The data poisoning scenario is different from our setting in two fundamental ways. First, in our case the attacker can maintain performance of the perturbed system on an unknown validation set within *arbitrary small* or, in case of ReLU neurons, *zero* margins. Such *imperceptible* changes on unknown validation sets is a signature characteristic of the threat we have revealed and studied. Second, the attacks we analysed *do not require any retraining*.

Other stealth attacks. Liu et al. ([Liu et al., 2018](#)) proposed a mechanism, SIN^2 , whereby a service provider gains control over their customers' deep neural networks through the provider's specially designed application programming interfaces (APIs). In this approach, the provider needs to first plant malicious information in higher precision bits (e.g. bits 16 and higher) of the network's weights. When an input trigger arrives the provider extracts this information via its service's API and performs malicious actions as per instructions encoded.

In contrast to this approach, stealth attacks we discovered *do not require* any special computational environments. After an attack is planted, it can be executed in fully secure and trusted infrastructure. In addition, our work reveals further concerns about how easily a malicious service provider can implement stealth attacks in hostile environments ([Liu et al., 2018](#)) by swapping bits in the mantissa of weights and biases of a single neuron (see [Fig. 2](#) for the patterns of change) at will.

This work is a significant advancement from the preliminary results presented in Tyukin *et al.* (2020), both in terms of algorithmic detail and theoretical understanding of the phenomenon. First, the vulnerabilities we reveal here are much more severe. Bounds on the probability of success for attacks in Tyukin *et al.* (2020) are constrained by $1 - M2^{-n}$. Our results show that under the same assumptions (that input reachability Huang *et al.* (2020) holds true), the probabilities of success for attacks generated by Algorithms 1, 2 can be made arbitrarily close to one (Remark 4). Secondly, we explicitly account for cases when input reachability holds only up to some accuracy, through parameters α, δ (Remark 3). Lastly, we present concrete algorithms, scenarios and examples of successful exploitation of these new vulnerabilities, including the case of one-neuron attacks (Section 4, Appendix; code is available in Tyukin *et al.* (2022)).

3. New stealth attack algorithms

In this section we introduce two new algorithms for generating stealth attacks on a generic AI system. These algorithms return a ReLU or a sigmoid neuron realizing an adversarial perturbation \mathfrak{A} . Implementation of these algorithms will rely upon some mild additional information about the unknown validation set \mathcal{V} . In particular, we request that latent representations $\Phi(\mathbf{u})$ for all $\mathbf{u} \in \mathcal{V}$ are located within some ball $\mathbb{B}_n(0, R)$ whose radius R is *known* to the attacker. We state this requirement in Assumption 1 below.

ASSUMPTION 1. (Latent representations of the validation data \mathcal{V}). There is an $R > 0$, known to the attacker, such that

$$\Phi(\mathbf{u}) \in \mathbb{B}_n(0, R) \text{ for all } \mathbf{u} \in \mathcal{V}. \quad (3.1)$$

Given that the set \mathcal{V} is finite, infinitely many such balls exist. Here we request that the attacker knows just a single value of R for which (3.1) holds. The value of R does not have to be the smallest possible. A bound for R is readily available when, for example, each component of Φ is bounded, as in the case where Φ represents the first few layers of a neural network concluding with a sigmoid function. Similarly, a bound may be computed when Φ represents the first few layers of a neural network with known weights and activation functions and the input consists of pixels from a specified finite range.

We also suppose that the feature map Φ in (2.2) satisfies an appropriate reachability condition (cf. Huang *et al.* (2020)), based on the following definition.

DEFINITION 1. (ν -input reachability of the classifier's latent space). Consider a function $f : \mathcal{U} \rightarrow \mathbb{R}^n$. A set $\mathcal{S} \subset \mathbb{R}^n$ is ν -input reachable for the function f if for any $\mathbf{x} \in \mathcal{S}$ there is an $\mathbf{u}(\mathbf{x}) \in \mathcal{U}$ such that $\|f(\mathbf{u}(\mathbf{x})) - \mathbf{x}\| \leq \nu$.

In what follows we will assume existence of some sets, namely $n - 1$ spheres, in the classifier's latent spaces which are input reachable for the map Φ . Precise definitions of these sets will be provided in our formal statements.

REMARK 2. The requirement that the classifier's latent space contains specific sets which are ν -input reachable for the map Φ may appear restrictive. If, for example, feature maps are vector compositions of ReLU and affine mappings (2.4), then several components of these vectors might be equal to zero in some domains and thus input reachability might seem impossible. However, one can always search for

a new feature map

$$\tilde{\Phi} : \tilde{\Phi}(\mathbf{u}) = T\Phi(\mathbf{u}), \quad T \in \mathbb{R}^{d \times n}, \quad (3.2)$$

where the matrix $T \in \mathbb{R}^{d \times n}$, $d \leq n$ is chosen so that relevant domains in the latent space formed by $\tilde{\Phi}$ become ν -input reachable. As we shall see later in Theorems 1 and 2, these relevant domains are spheres $\mathbb{S}_{d-1}(\mathbf{c}, \delta R)$, $\delta \in (0, 1]$, where $\mathbf{c} \in \mathbb{R}^d$ is some given reference point.

If the feature map $\Phi : \mathcal{U} \rightarrow \mathbb{R}^n$ with $\mathcal{U} \subset \mathbb{R}^m$ is differentiable and non-constant, and the set \mathcal{U} has a non-empty interior, $\text{Int}(\mathcal{U})$, then matrices T producing ν -input reachable feature maps $\tilde{\Phi}$ in some neighbourhood of a point from $\text{Int}(\mathcal{U})$ can be determined as follows. Pick a target input $\mathbf{u}_0 \in \text{Int}(\mathcal{U})$, then

$$\Phi(\mathbf{u}) = \Phi(\mathbf{u}_0) + J(\mathbf{u} - \mathbf{u}_0) + o(\|\mathbf{u} - \mathbf{u}_0\|),$$

where J is the Jacobian of Φ at \mathbf{u}_0 and $\|\cdot\|$ denotes the Euclidean norm. Suppose that $\text{rank}(J) = d$, $d > 0$, and let

$$J = U \begin{pmatrix} \Sigma & 0_{d \times (m-d)} \\ 0_{(n-d) \times d} & 0_{(n-d) \times (m-d)} \end{pmatrix} V$$

be the singular value decomposition of J , where Σ is a $d \times d$ diagonal matrix containing non-zero singular values of J on its main diagonal, and $U \in \mathbb{R}^{n \times n}$, $V \in \mathbb{R}^{m \times m}$ are unitary matrices. If $m = d$ or $n = d$ then the corresponding zero matrices in the above decomposition are assumed to be empty. Setting

$$T = (I_d \ 0_{d \times (n-d)}) U^T$$

ensures that for any arbitrarily small $\nu > 0$ there is an $r(\nu, \mathbf{u}_0) > 0$ such that the sets $\mathbb{S}_{d-1}(\tilde{\Phi}(\mathbf{u}_0), \rho)$, $\rho \in [0, r(\nu, \mathbf{u}_0))$ are ν -input reachable for the function $\tilde{\Phi} = T\Phi$.

Note that the same argument applies to maps Φ which are differentiable only in some neighbourhood of $\mathbf{u}_0 \in \text{Int}(\mathcal{U})$. This enables the application of this approach for producing ν -input reachable feature functions to maps involving compositions of ReLU functions. In what follows we will use the symbol n to denote the dimension of the space where Φ maps the input into assuming that the input reachability condition holds for this feature map.

Linearity of $\tilde{\Phi}$ in T enables one to preserve the structure of perturbations (2.4) and (2.5) so that they remain, in effect, just a single additional neuron. We used the latter property in our computational experiments (see Remark 11 in Section A.4) to generate examples of stealth attacks.

3.1 Target-agnostic stealth attacks

Our first stealth attack algorithm is presented in Algorithm 1. The algorithm produces a modification of the relevant part F of the original AI system that is implementable by a single ReLU or sigmoid function. Regarding the *trigger input*, \mathbf{u}' , the algorithm relies on another process (see step 3).

In what follows, we will denote this process as an auxiliary algorithm \mathcal{A}_0 which, for the map Φ , given R , parameters $\delta \in (0, 1]$, $\gamma \in (0, 1)$, $\nu < \delta$, $\delta + \nu \leq \gamma^{-1}$, whose meaning and choice are explained in Remark 3 after Theorem 1, and any $\mathbf{x} \in \mathbb{S}_{n-1}(0, \delta)$, returns a solution of the following constrained

search problem:

$$\mathcal{A}_0(\Phi, \mathbf{x}, R, \nu) : \quad \text{find } \mathbf{u} \in \mathcal{U} \quad \text{such that} \quad \|\Phi(\mathbf{u})R^{-1} - \mathbf{x}\| \leq \nu. \quad (3.3)$$

Observe that νR -input reachability, $\nu < \delta$, of the set $\mathbb{S}_{n-1}(0, R\delta)$ for the map Φ together with the choice of δ, ν and γ satisfying $\delta + \nu \leq \gamma^{-1}$ ensures the existence of a solution to the above problem for every $\mathbf{x} \in \mathbb{S}_{n-1}(0, \delta)$.

Thorough analysis of computability of solutions of (3.3) is outside of the scope of this work (see Bastounis *et al.* (2021a) for an idea of the issues involved computationally in solving optimization problems). Therefore, we shall assume that the auxiliary process \mathcal{A}_0 always returns a solution of (3.3) for a choice of δ, γ , and given ν, R . In our numerical experiments (see Section 4), finding a solution of (3.3) did not pose significant issues.

With the auxiliary algorithm \mathcal{A}_0 available, we can now present Algorithm 1:

Algorithm 1 Single-neuron plain stealth attack

Input: $\delta \in (0, 1]$, $\gamma \in (0, 1)$, $\alpha \in [0, 1)$, $(1 + \alpha)\delta \leq 1$, $\Delta, \varepsilon \geq 0$, a sigmoid or a ReLU function g , and R satisfying (3.1).

- 1: **procedure** ADVERSARIAL STEALTH PERTURBATION($\delta, \Delta, \varepsilon, g$)
- 2: Draw a random vector \mathbf{x} from the equidistribution in the sphere $\mathbb{S}_{n-1}(0, \delta)$, $\delta \in (0, 1]$.
- 3: Use algorithm \mathcal{A}_0 (see (3.3)) to generate an input $\mathbf{u}' \in \mathcal{U}$ such that $\mathbf{x}' = \Phi(\mathbf{u}')/R$ is within a $\alpha\delta$ -distance from \mathbf{x} :

$$\|\mathbf{x}' - \mathbf{x}\| \leq \alpha\delta. \quad (3.4)$$

- 4: Set

$$\mathfrak{A}(\cdot, \kappa \mathbf{x}' R^{-1}, b) = Dg(\langle \cdot, \kappa \mathbf{x}' R^{-1} \rangle - b), \quad \text{with } b = 0.5\kappa(1 + \gamma)\|\mathbf{x}'\|^2, \quad (3.5)$$

where κ and D are chosen so that

$$Dg(-0.5\kappa(1 - \gamma)\|\mathbf{x}'\|^2) \leq \varepsilon \text{ and } Dg(0.5\kappa(1 - \gamma)\|\mathbf{x}'\|^2) \geq \Delta. \quad (3.6)$$

*Note that a choice of κ, D so that (3.6) is satisfied is always possible.

- 5: **end procedure**

Output: trigger \mathbf{u}' , weight vector $\mathbf{w} = \kappa \mathbf{x}' R^{-1}$, bias $b = 0.5\kappa(1 + \gamma)\|\mathbf{x}'\|^2$, and output gain D of the sigmoid or ReLU function g .

The performance of Algorithm 1 in terms of the probability of producing a successful stealth attack of the form (2.3) is characterized by Theorem 1 below.

THEOREM 1. Let Assumption 1 hold. Consider Algorithm 1, and let $\gamma \in (0, 1)$, $\alpha \in [0, 1)$ and $\delta \in (0, 1]$ be such that $(1 + \alpha)\delta \leq 1$ and $\gamma(1 - \alpha)\delta > \alpha$. Additionally, assume that the set $\mathbb{S}_{n-1}(0, R\delta)$ is $\delta\alpha R$ -input reachable for the map Φ . Then

$$\varphi(\gamma, \delta, \alpha) := \cos(\arccos(\gamma(1 - \alpha)\delta) + \arccos((1 - \alpha^2)^{1/2})) > 0 \quad (3.7)$$

and the probability $P_{a,1}$ that Algorithm 1 returns a successful ε - Δ stealth attack satisfies

$$P_{a,1} \geq 1 - M\pi^{-1/2} \frac{\Gamma\left(\frac{n}{2}\right)}{\Gamma\left(\frac{n-1}{2}\right)} \int_0^{\arccos(\varphi(\gamma, \delta, \alpha))} \sin^{n-2}(\theta) d\theta. \quad (3.8)$$

In particular,

$$P_{a,1} \geq 1 - M \frac{1}{2\pi^{\frac{1}{2}}} \frac{\Gamma\left(\frac{n}{2}\right)}{\Gamma\left(\frac{n}{2} + \frac{1}{2}\right)} \frac{1}{\varphi(\gamma, \delta, \alpha)} \left(1 - \varphi(\gamma, \delta, \alpha)^2\right)^{\frac{n-1}{2}}. \quad (3.9)$$

REMARK 3. (Choice of parameters). The three parameters α , δ and γ should be chosen to strike a balance between attack success, time to compute the algorithm and the size of the weights that the stealth attack creates. In particular, from the perspective of attack success (3.8) is optimized when γ and δ are close to 1 and α is close to 0. However, α and δ are restricted by the exact form of the map Φ : they need to be chosen so that the set $\mathbb{S}_{n-1}(0, R\delta)$ is $\delta\alpha R$ -input reachable for Φ . Furthermore, the speed of executing \mathcal{A}_0 is also influenced by these choices with the complexity of solving (3.4) increasing as α decreases to 0. The size of the weights chosen in (3.6) is influenced by the choice of γ so that the L_2 norm of the attack neuron weights grows like $O((1 - \gamma)^{-1})$ for sigmoid g . Finally, the condition $(1 + \alpha)\delta \leq 1$ ensures that the chosen trigger image \mathbf{u}' has $\|\Phi(\mathbf{u}')\| \leq R$.

REMARK 4. (Determinants of success and vulnerabilities). Theorem 1 establishes explicit connections between intended parameters of the trigger \mathbf{u}' (expressed by $\Phi(\mathbf{u}')/R$ which is to be maintained within $\delta\alpha$ from \mathbf{x}), dimension n of the AI's latent space, accuracy of solving (3.3) (expressed by the value of α —the smaller the better), design parameters $\gamma \in (0, 1)$ and $\delta \in (0, 1]$ and vulnerability of general AI systems to stealth attacks.

In general, the larger the value of n for which solutions of (3.3) can be found without increasing the value of α , the higher the probability that the attack produced by Algorithm 1 is successful. Similarly, for a given success probability bound, the larger the value of n , the smaller the value of γ required, allowing stealth attacks to be implemented with smaller weights \mathbf{w} and bias b . At the same time, if no explicit restrictions on δ and the weights are imposed then Theorem 1 suggests that if one picks $\delta = 1$, γ sufficiently close to 1, and α sufficiently close to 0, subject to finding an appropriate solution of (3.3) (c.f. the notion of reachability Huang *et al.* (2020)), one can create stealth attacks whose probabilities of success can be made arbitrarily close to 1 for any fixed n . Indeed, if $\alpha = 0$ then the right-hand side of (3.9) becomes

$$1 - M \frac{1}{2\pi^{\frac{1}{2}}} \frac{\Gamma\left(\frac{n}{2}\right)}{\Gamma\left(\frac{n}{2} + \frac{1}{2}\right)} \frac{1}{\gamma\delta} \left(1 - (\gamma\delta)^2\right)^{\frac{n-1}{2}}, \quad (3.10)$$

which, for any fixed M , n , can be made arbitrarily close to 1 by an appropriate choice of $\gamma \in (0, 1)$ and $\delta \in (0, 1]$.

REMARK 5. (Arbitrary output sign and target class of adversarial perturbation). Although Problem 1 does not impose any requirements on the sign of $\mathfrak{A}(\Phi(\mathbf{u}'), \mathbf{w}, b)$, this quantity can be made positive or negative through the choice of D . Note that the target class can be arbitrary too (see Remark 1).

REMARK 6. (Precise and zero-tolerance attacks with ReLU units). One of the original requirements for a stealth attack is to produce a response to a trigger input \mathbf{u}' such that $|F \circ \Phi(\mathbf{u}') - F_a \circ \Phi(\mathbf{u}')|$ exceeds an a-priori given value Δ . However, if the attack is implemented with a ReLU unit then one can select the values of \mathbf{w}, b so that $|F \circ \Phi(\mathbf{u}') - F_a \circ \Phi(\mathbf{u}')|$ exactly equals Δ . Indeed for a ReLU neuron, condition (3.6) reduces to

$$0.5D \left(\kappa(1 - \gamma) \|\mathbf{x}'\|^2 \right) \geq \Delta.$$

Hence picking $\kappa = 2\Delta \left((1 - \gamma) \|\mathbf{x}'\|^2 \right)^{-1}$ and $D = 1$ results in the desired output.

Moreover, ReLU units produce adversarial perturbations \mathfrak{A} with $\varepsilon = 0$. These *zero-tolerance* attacks, if successful, do not change the values of \mathcal{F} on the validation set \mathcal{V} and as such are completely undetectable on \mathcal{V} .

REMARK 7. (Hiding adversarial perturbations in redundant deep learning structures). Algorithm 1 (as well as its input-specific version, Algorithm 2, below) implements adversarial perturbations by *adding* a single sigmoid or ReLU neuron. A question therefore arises, if one can ‘plant’ or ‘hide’ a stealth attack within an existing AI structure, rather than adding a neuron, as in (2.3). Intuitively, over-parametrization and redundancies in many deep learning architectures should provide ample opportunities precisely for this sort of malevolent action. As we empirically justify in Section 4, this is indeed possible. In these experiments, we looked at a neural network with L layers. The map F corresponded to the last $k + 1$ layers: $L - k, \dots, L, k \geq 1$. We looked for a neuron in layer $L - k$ whose output weights have the smallest L_1 -norm of all neurons in that layer (see Appendix, Section A.6). This neuron was then replaced with a neuron implementing our stealth attack, and its output weights were wired so that a given trigger input \mathbf{u}' evoked the response we wanted from this trigger. This new type of one-neuron attack may be viewed as the stealth version of a one pixel attack (Su *et al.*, 2019). Surprisingly, this approach worked consistently well across various randomly initiated instances of the same network. These experiments suggest a real non-hypothetical possibility of turning *a needle in a haystack* (a redundant neuron) into *a pebble in a shoe* (a malevolent perturbation).

REMARK 8. (Stability). Under appropriate (weak) assumptions, such as the existence of a topology on \mathcal{U} where Φ is continuous and the continuity of g , Algorithm 1 generates stable trigger inputs—if $\mathbf{u}' \in \mathcal{U}$ is sufficiently close to \mathbf{u} then \mathbf{u}' is a trigger image in the sense of (1) with a potentially smaller value of Δ .

REMARK 9. (Other activation functions). In this work, to be concrete we focus on the case of attacking with a ReLU or sigmoid activation function. However, the algorithms and analysis readily extend to the case of any continuous activation function that is non-negative and monotonically increasing. Appropriate pairs of matching leaky ReLUs also fit into this framework.

3.2 Target-specific stealth attacks

The attack and the trigger $\mathbf{u}' \in \mathcal{U}$ constructed in Algorithm 1 are ‘arbitrary’ in the sense that \mathbf{x} is drawn at random. As opposed to standard adversarial examples, they are not linked or targeting any specific input. Hence a question arises: is it possible to create a *targeted* adversarial perturbation of the AI which is triggered by an input $\mathbf{u}' \in \mathcal{U}$ located in a vicinity of some specified input, \mathbf{u}^* ?

As we shall see shortly, this may indeed be possible through some modifications to Algorithm 1. For technical convenience and consistency, we introduce a slight reformulation of Assumption 1.

ASSUMPTION 2. (Relative latent representations of the validation data \mathcal{V}). Let $\mathbf{u}^* \in \mathcal{U}$ be a target input of interest to the attacker. There is an $R > 0$, also known to the attacker, such that

$$\Phi(\mathbf{u}) - \Phi(\mathbf{u}^*) \in \mathbb{B}_n(0, R) \text{ for all } \mathbf{u} \in \mathcal{V}. \quad (3.11)$$

We note that Assumption 2 follows immediately from Assumption 1 if a bound on the size of the latent representation $\Phi(\mathbf{u}^*)$ of the input \mathbf{u}^* is known. Indeed, if R' is a value of R for which (3.1) holds then (3.11) holds whenever $R \geq R' + \|\Phi(\mathbf{u}^*)\|$.

Algorithm 2 provides a recipe for creating such targeted attacks.

Algorithm 2 Single-neuron targeted stealth attack

Input: $\delta \in (0, 1]$, $\gamma \in (0, 1)$, $\alpha \in [0, 1)$, $\delta(\alpha + 1) \leq 1$, $\Delta, \varepsilon \geq 0$, a sigmoid or a ReLU function g , a target input $\mathbf{u}^* \in \mathcal{U}$, and R for which (3.11) holds.

- 1: **procedure** TARGETED ADVERSARIAL STEALTH PERTURBATION($\delta, \alpha, D, \varepsilon, g$)
- 2: Draw a random vector \mathbf{x} from the equidistribution in the sphere $\mathbb{S}_{n-1}(0, \delta)$, $\delta \in (0, 1]$.
- 3: Use the algorithm \mathcal{A}_0 to generate an input \mathbf{u}' such that $\mathbf{x}' = (\Phi(\mathbf{u}') - \Phi(\mathbf{u}^*)) / R$ satisfies (3.4).
- 4: Set

$$\begin{aligned} \mathfrak{A}(\cdot, \kappa \mathbf{x}' R^{-1}, b) &= Dg(\langle \cdot, \kappa \mathbf{x}' R^{-1} \rangle - b), \\ b &= 0.5\kappa(1 + \gamma)\|\mathbf{x}'\|^2 + \kappa\langle \Phi(\mathbf{u}^*), \mathbf{x}' R^{-1} \rangle, \end{aligned}$$

where κ and D are chosen as in (3.6).

*Note that such a choice is always possible.

- 5: **end procedure**

Output: trigger input \mathbf{u}' , weight vector $\mathbf{w} = \kappa \mathbf{x}' R^{-1}$, bias $b = 0.5\kappa(1 + \gamma)\|\mathbf{x}'\|^2 + \kappa\langle \Phi(\mathbf{u}^*), \mathbf{x}' R^{-1} \rangle$, and output gain D of the sigmoid or a ReLU function g .

Performance bounds for Algorithm 2 can be derived in the same way as we have done in Theorem 1 for Algorithm 1. Formally, we state these bounds in Theorem 2 below.

THEOREM 2. Let Assumption 2 hold. Consider Algorithm 2, and let parameters $\gamma \in (0, 1)$, $\alpha \in [0, 1)$, and $\delta \in (0, 1]$ be such that $\delta(1 + \alpha) \leq 1$ and $\gamma(1 - \alpha)\delta > \alpha$. Additionally, assume that the set $\mathbb{S}_{n-1}(\Phi(\mathbf{u}^*), R\delta)$ is $\delta\alpha R$ -input reachable for the map Φ .

Then the probability $P_{a,2}$ that Algorithm 2 returns a successful ε - Δ stealth attack is bounded from below by (3.8) and (3.9), that is,

$$P_{a,2} \geq 1 - M\pi^{-1/2} \frac{\Gamma(\frac{n}{2})}{\Gamma(\frac{n-1}{2})} \int_0^{\arccos(\varphi(\gamma, \delta, \alpha))} \sin^{n-2}(\theta) d\theta,$$

and in particular,

$$P_{a,1} \geq 1 - M \frac{1}{2\pi^{\frac{1}{2}}} \frac{\Gamma(\frac{n}{2})}{\Gamma(\frac{n}{2} + \frac{1}{2})} \frac{1}{\varphi(\gamma, \delta, \alpha)} \left(1 - \varphi(\gamma, \delta, \alpha)^2\right)^{\frac{n-1}{2}}.$$

Remarks 3–9 apply equally to Algorithm 2. In addition, the presence of a specified target input \mathbf{u}^* offers extra flexibility and opportunities. An attacker may for example have a list of potential target inputs. Applying Algorithm 2 to each of these inputs produces different triggers \mathbf{u}' each with different possible values of α . According to Remark 4 (see also (3.10)), small values of α imply higher probabilities that the corresponding attacks are successful. Therefore, having a list of target inputs and selecting a trigger with minimal α increases the attacker's chances of success.

3.3 Attribute-constrained attacks and a 'concentrational collapse' for validation sets \mathcal{V} chosen at random

The attacks developed and analysed so far can affect all attributes of the input data's latent representations. It is also worthwhile to consider triggers whose deviation from targets in latent spaces is limited to only a few attributes. This may help to disguise triggers as legitimate data, an approach which was recently shown to be successful in the context of adversarial attacks on input data (Beerens & Higham, 2023). Another motivation stems from more practical scenarios where the task is to find a trigger in the vicinity of the latent representation of another input in models with sparse coding. The challenge here is to produce attack triggers while retaining sparsity. Algorithm 3 and the corresponding Theorem 3 are motivated by these latter scenarios.

As before, consider an AI system with an input–output map (2.2). Assume that the attacker has access to a suitable real number $R > 0$, an input \mathbf{u}^* and a set of orthonormal vectors $\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_{n_p} \in \mathbb{R}^n$ which form the space in which the attacker can make perturbations. Furthermore, assume that both the validation set and the input \mathbf{u}^* satisfy Assumption 3 below.

ASSUMPTION 3. (Data model). Elements $\mathbf{u}_1, \dots, \mathbf{u}_M$ are randomly chosen so that $\Phi(\mathbf{u}_1), \dots, \Phi(\mathbf{u}_M)$ are random variables in \mathbb{R}^n and such that both of the following hold:

1. There is a $\mathbf{c} \in \mathbb{R}^n$ such that \mathbf{u}^* , and (with probability 1), $\mathbf{u}_1, \dots, \mathbf{u}_M$, are in the set:

$$\{\mathbf{u} \in \mathcal{U} \mid \|\Phi(\mathbf{u}) - \mathbf{c}\| \leq R/2\}. \quad (3.12)$$

2. There is a $C \geq 1$ such that for any $r \in (0, R/2]$ and $\xi \in \mathbb{B}_n(\mathbf{c}, R/2)$

$$P(\Phi(\mathbf{u}_i) \in \mathbb{B}_n(\xi, r)) \leq C \left(\frac{2r}{R} \right)^n. \quad (3.13)$$

Property (3.12) requires that the validation set is mapped into a ball centred at \mathbf{c} and having a radius $R/2$ in the system's latent space, and (3.13) is a non-degeneracy condition restricting pathological concentrations.

Consider Algorithm 3. As we show below, if the validation set satisfies Assumption 3 and the attacker uses Algorithm 3 then the probabilities of generating a successful attack may be remarkably high even if the attack triggers' latent representations deviate from those of the target in only few attributes.

THEOREM 3. Let Assumption 3 hold. Consider Algorithm 3, and let parameters $\gamma \in (0, 1)$, $\alpha \in [0, 1)$ and $\delta \in (0, 1]$ be such that the set $\Phi(\mathbf{u}^*) + (\mathbb{S}_{n-1}(0, R\delta) \cap \text{span}\{\mathbf{h}_1, \dots, \mathbf{h}_{n_p}\})$ is $\delta\alpha R$ -input reachable for the map Φ . Furthermore, suppose that $2\delta\gamma(1 - \alpha) \geq \sqrt{1 - (MC)^{-2/n}}$.

Algorithm 3 Single-neuron targeted attribute-constrained stealth attack

Input: $\delta \in (0, 1]$, $\gamma \in (0, 1)$, $\alpha \in [0, 1)$, $\delta(\alpha + 1) \leq 1$, $\Delta, \varepsilon \geq 0$, a sigmoid or a ReLU function g , a target input $\mathbf{u}^* \in \mathcal{U}$, an R for which Assumption 3 holds, an $n_p \in \mathbb{N}$, $2 \leq n_p \leq n$, and a system of orthonormal vectors $\{\mathbf{h}_1, \dots, \mathbf{h}_{n_p}\}$, $\mathbf{h}_i \in \mathbb{R}^n$.

- 1: **procedure** TARGETED ADVERSARIAL CONSTRAINED STEALTH PERTURBATION($\delta, \alpha, D, \varepsilon, g$)
- 2: Draw a random vector \mathbf{v} from the equidistribution in the sphere $\mathbb{S}_{n_p-1}(0, \delta)$, $\delta \in (0, 1]$. Set

$$\mathbf{x} = \sum_{i=1}^{n_p} v_i \mathbf{h}_i.$$

- 3: Use the algorithm \mathcal{A}_0 to generate an input \mathbf{u}' such that $\mathbf{x}' = (\Phi(\mathbf{u}') - \Phi(\mathbf{u}^*)) / R$ satisfies (3.4).
- 4: Set

$$\begin{aligned} \mathfrak{A}(\cdot, \kappa \mathbf{x}' R^{-1}, b) &= Dg(\langle \cdot, \kappa \mathbf{x}' R^{-1} \rangle - b), \\ b &= 0.5\kappa(1 + \gamma) \|\mathbf{x}'\|^2 + \kappa \langle \Phi(\mathbf{u}^*), \mathbf{x}' R^{-1} \rangle, \end{aligned}$$

where κ and D are chosen as in (3.6).

*Note that such a choice is always possible.

- 5: **end procedure**

Output: trigger input \mathbf{u}' , weight vector $\mathbf{w} = \kappa \mathbf{x}' R^{-1}$, bias $b = 0.5\kappa(1 + \gamma) \|\mathbf{x}'\|^2 + \kappa \langle \Phi(\mathbf{u}^*), \mathbf{x}' R^{-1} \rangle$, and output gain D of the sigmoid or a ReLU function g .

Let

$$\begin{aligned} \theta^* &= \arccos\left(\min\left\{1, 2\delta\gamma(1 - \alpha) - \sqrt{1 - (MC)^{-2/n}}\right\}\right) + \arccos(\sqrt{1 - \alpha^2}), \\ \rho(\theta) &= \min\left\{1, 2\delta\gamma(1 - \alpha) - \max\{0, \cos(\theta - \arccos(\sqrt{1 - \alpha^2}))\}\right\} \quad \text{for } \theta \in [\theta^*, \pi]. \end{aligned}$$

Then $\theta^* \in [0, \pi]$, $\rho(\theta) \in [0, 1]$ for $\theta \in [\theta^*, \pi]$ and the probability $P_{a,3}$ that Algorithm 3 returns a successful ε - Δ stealth attack is bounded from below by

$$1 - MC \frac{1}{\pi^{\frac{1}{2}}} \frac{\Gamma\left(\frac{n_p}{2}\right)}{\Gamma\left(\frac{n_p-1}{2}\right)} \int_{\theta^*}^{\pi} \left(1 - \rho(\theta)^2\right)^{\frac{n}{2}} \sin^{n_p-2}(\theta) d\theta - \frac{1}{\pi^{\frac{1}{2}}} \frac{\Gamma\left(\frac{n_p}{2}\right)}{\Gamma\left(\frac{n_p-1}{2}\right)} \int_0^{\theta^*} \sin^{n_p-2}(\theta) d\theta. \quad (3.14)$$

We emphasize a key distinction between the expressions (3.14) and (3.8). In (3.14) the final term is independent of M , whereas the corresponding term in (3.8) is multiplied by a factor of M . The only term affected by M in (3.14) is modulated by an additional exponent with the base $(1 - \rho(\theta)^2)^{1/2}$ which is strictly smaller than one in $(\theta^*, \pi]$ for $\theta^* < \pi$. In order to give a feel of how far bound provided in Theorem 3 may be from the one established in Theorems 1, 2 we computed the corresponding bounds for $n = n_p = 200$, $\gamma = 0.9$, $\delta = 1/3$, $M = 2500$, $C = 100$ and $\alpha \in [0.01, 0.3]$. Results of the comparison are summarized in Table 1.

TABLE 1 Comparison of bounds (3.8) and (3.14) for different values of the accuracy parameter α , and fixed $\gamma = 0.9$, $\delta = 1/3$, $n = n_p = 200$, $M = 2500$, $C = 100$.

α	0.3	0.25	0.20	0.15	0.10	0.05	0.01
Bound (3.8)	Not feasible	Not feasible	-703.2	-161.9	-17.0	0.105	0.956
Bound (3.14)	4.11×10^{-4}	0.014	0.148	0.533	0.886	0.990	0.999

TABLE 2 Comparison of bounds (3.8), where n is replaced with n_p , and (3.14) for different values of the accuracy parameter α , and fixed $\gamma = 0.9$, $\delta = 1/3$, $n = 200$, $n_p = 5$, $M = 2500$, $C = 100$.

α	0.3	0.25	0.20	0.15	0.10	0.05	0.01
Bound (3.8)	Not feasible	Not feasible	-1172.2	-1049.0	-929.0	-813.4	-724.9
Bound (3.14)	0.326	0.384	0.444	0.504	0.565	0.624	0.669

In this context, Theorem 3 reveals a phenomenon where validation sets which could be deemed as sufficiently large in the sense of bounds specified by Theorems 1 and 2 may still be considered as ‘small’ due to bound (3.14). We call this phenomenon *concentrational collapse*. When concentrational collapse occurs, the AI system becomes more vulnerable to stealth attacks when the cardinality of validation sets \mathcal{V} is sub-exponential in dimension n implying that the owner would have to generate and keep a rather large validation set \mathcal{V} to make up for these small probabilities. Remarkably, this vulnerability persists even when the attacker’s precision is small (α large).

REMARK 10. (Lower dimensional triggers). Another important consequence of concentrational collapse, which becomes evident from the proof of Theorem 3, is the possibility to sample vectors \mathbf{x} from the $n_p - 1$ sphere, $\mathbb{S}_{n_p-1}(0, \delta)$, $2 \leq n_p < n$ instead of from the $n - 1$ sphere $\mathbb{S}_{n-1}(0, \delta)$. Note that the second term in the right-hand side of (3.14) scales linearly with cardinality M of the validation set \mathcal{V} and has an exponent in n which is the ‘ambient’ dimension of the feature space. The third term in (3.14) decays exponentially with $n_p < n$ but does not depend on M . The striking difference between behaviour of bounds (3.14) and (3.8) at small values of n_p and large n is illustrated with Table 2. According to Table 2, bound (3.8) is either infeasible or impractical for all tested values of the accuracy parameter α . In contrast, bound (3.14) indicates relatively high probabilities of stealth attacks’ success for the same values of α as long as all relevant assumptions hold true. This implies that the concentration collapse phenomenon can be exploited for generating triggers with lower dimensional perturbations of target inputs in the corresponding feature spaces. The latter possibility may be relevant for overcoming defense strategies which enforce high-dimensional yet sparse representations of data in latent spaces. Our experiments in Section 4.1, which reveal high stealth attack success rates even for relatively low-dimensional perturbations, are consistent with these theoretical observations.

4. Experiments

Let us show how stealth attacks can be constructed for given deep learning models. We considered two standard benchmark problems in which deep learning networks are trained on the CIFAR10 (Krizhevsky & Hinton, 2009) and a MATLAB version of the MNIST (LeCun *et al.*) datasets. All

TABLE 3 Network architecture used in experiments on the MNIST digits dataset. The red colour shows layers which we represent by map \mathcal{F} in (2.1).

Layer number	Type	Size
1	Input	$28 \times 28 \times 1$
2	Conv2d	$3 \times 3 \times 8$
3	Batch normalization	
4	ReLU	
5	Maxpool	pool size 2×2 , stride 2×2
6	Conv2d	$3 \times 3 \times 16$
7	Batch normalization	
8	ReLU	
9	Maxpool	pool size 2×2 , stride 2×2
10	Conv2d	$3 \times 3 \times 32$
11	Batch normalization	
12	ReLU	
13	Fully connected	200
14	ReLU	
15	Fully connected	100
16	ReLU	
17	Fully connected	10
18	Softmax	10

networks had a standard architecture with feature-generation layers followed by dense fully connected layers.

Three alternative scenarios for planting a stealth attack neuron were considered. Schematically, these scenarios are shown in Fig. 2. Our theoretical results directly apply to Scenario 1. The results are also relevant to Scenario 2, as discussed in Remark 7. Scenario 3 is included for completeness; it is computationally equivalent to Scenario 1 but respects the original structure more closely by passing information between successive layers, rather than skipping directly to the end.

In our experiments we determined trigger–target pairs and changed the network’s architecture in accordance with planting Scenarios 1 and 2. It is clear that if Scenario 1 is successful then Scenario 3 will be successful too. The main difference between Scenario 2 and Scenarios 1, 3 is that in the former case we *replace* a neuron in F by the ‘attack’ neuron. The procedure describing selection of a neuron to replace (and hence attack) in Scenario 2 is detailed in Section A.6, and an algorithm which we used to find triggers is described in Section A.4. When implementing stealth attacks in accordance with Scenario 2, we always selected neurons whose susceptibility rank is 1.

As a general rule, in Scenario 2, we attacked neurons in the block of fully connected layers just before the final softmax block (fully connected layer followed by a layer with softmax activation functions—see Fig. 3 and Table 3 for details). The attacks were designed to assign an arbitrary class label the Attacker wanted the network to return in response to a trigger input computed by the Attacker. In order to do so, the weight between the attacked neuron and the neuron associated with the Attacker-intended class output was set to 1. All other output weights of the attacked neuron were set to 0. In principle, 0s can be replaced with negative, sufficiently small positive values, or their combinations. MATLAB code implementing all relevant steps of the experiments can be found in Tyukin *et al.* (2022).

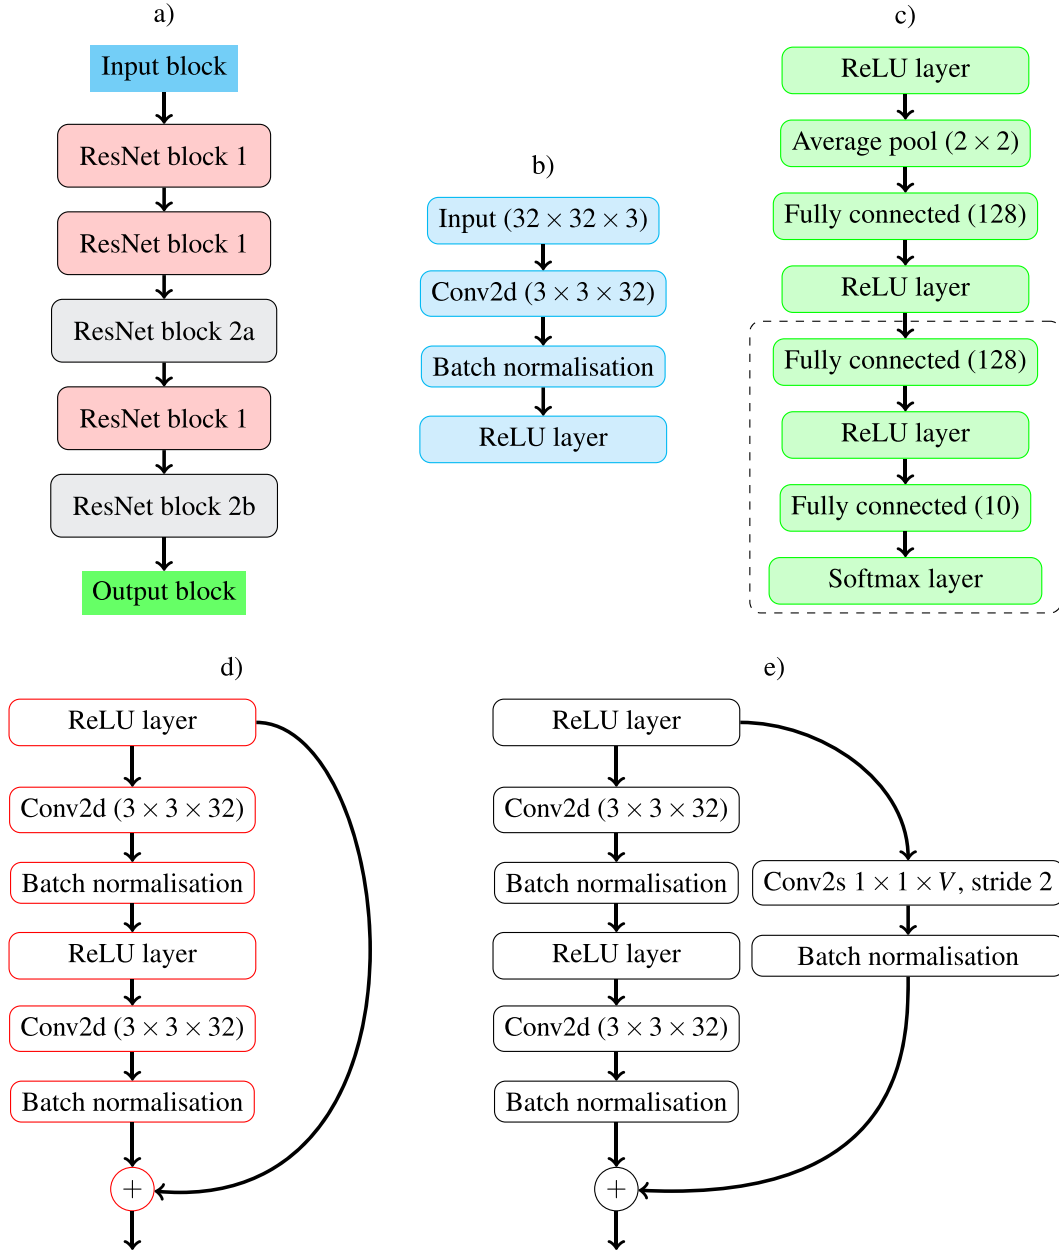


FIG. 3. (Colour online) ResNet architecture used in experiments with CIFAR-10 dataset. Panel *a* shows the general structure of the network. Panels *b* and *c* show configuration of the network's input and output blocks (highlighted in cyan and green in panel *a*, respectively). The diagram in panel *d* shows the structure of ResNet blocks 1. Panel *e* describes ResNet blocks 2a and 2b. The value of V in ResNet blocks 2a and 2b was set to 64 and 128, respectively. The dashed rectangle in panel *c* shows network's layers implementing the map F .

4.1 *Stealth attacks for a class of networks trained on CIFAR-10 dataset*

To assess their viability, we first considered the possibility of planting stealth attacks into networks trained on the CIFAR-10 dataset. The CIFAR-10 dataset is composed of 32×32 colour RGB images which correspond to inputs of dimension 3072. Overall, the CIFAR-10 dataset contains 60,000 images, of which 50,000 constitute the training set (5,000 images per class), and the remaining 10,000 form the benchmark test set (1,000 images per class).

4.1.1 *Network architecture.* To pick a good neural architecture for this task we analysed a collection of state-of-the-art models for various benchmark data (Stojnic *et al.*). According to these data, networks with a ResNet architecture were capable of achieving an accuracy of 99.37% on the CIFAR-10 dataset. This is consistent with the reported level of label errors of 0.54% in CIFAR-10 tests set (Northcutt *et al.*, 2021). ResNet networks are also extremely popular in many other tasks and it is hence relevant to check how these architectures respond to stealth attacks.

The structure of the neural network used for this task is shown in Fig. 3. The map F is implemented by the last four layers of the network highlighted by the dashed rectangle in panel c, Fig. 3. The remaining part of the network represents the map Φ .

4.1.2 *Training protocol.* The network was trained for 100 epochs with minibatches containing 128 images, using L_2 regularization of the network weights (regularization factor 0.0001), and with stochastic gradient descent with momentum. Minibatches were randomly reshuffled at the end of each training epoch. Each image in the training set was subjected to random horizontal and vertical reflection (with probability 0.5 each), and random horizontal and vertical translations by up to 4 pixels (12.5% of the image size) in each direction. The initial learning rate was set to 0.1, and the momentum parameter was set to 0.9. After the first 60 epochs the learning rate was changed to 0.01. The trained network achieved 87.56% accuracy on the test set.

4.1.3 *Construction of stealth attacks.* Stealth attacks took the form of single ReLU neurons added to the output of the last ReLU layer of the network. These neurons received 128-dimensional inputs from the fifth from last layer (the output of the map Φ) shown in Fig. 3, panel c), just above the dashed rectangle. The values of γ , δ and Δ were set to 0.9, 0.5 and 50, respectively. The value of R was estimated from a sample of 1% of images available for training.

The validation set \mathcal{V} was composed of 1,000 randomly chosen images from the training set. Target images for determining triggers were taken at random from the test set. Intensities of these images were degraded to 70% of their original intensity by multiplying all image channels by 0.7. This makes the problem harder from the perspective of an attacker. To find the trigger, a standard gradient-based search method was employed to solve the relevant optimization problem in step 3 of the algorithm (see Section A.4 for more details).

4.2 *Effective local dimension of feature maps*

We also examined the effective local dimension of the feature maps by assessing the sparsity of feature vectors for the network trained in this experiment. Average sparsity, i.e. the ratio of zero attributes of $\Phi(\mathbf{u})$ to the total number of attributes (128), was 0.9307 (8.87 non-zero attributes out of 128) for \mathbf{u} from the entire training set, and was equal to 0.9309 (8.85 non-zero attributes out of 128) for \mathbf{u} from the validation set \mathcal{V} . Thus the relevant dimension n of the perturbation δ (see Remark 2) used in Algorithm 2 was significantly lower than that of the ambient space.

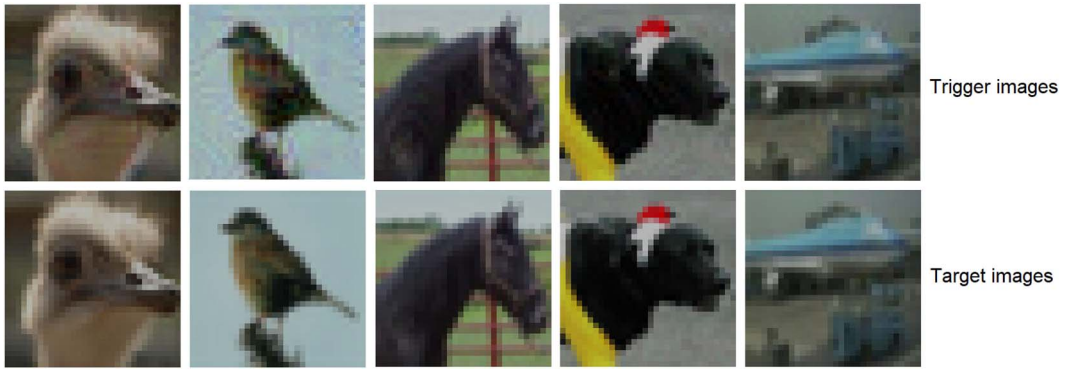


FIG. 4. (Colour online) Examples of triggers (top row) and original target images (bottom row) for CIFAR-10 dataset computed for the trained ResNet network.

4.2.1 Performance of stealth attacks. For the trained network, we implemented 20 stealth attacks following Scenario 1 in Fig. 2, with 13 out of 20 attacks succeeding (the output being identically zero for all images from the validation set \mathcal{V}), and seven failing (some images from the set \mathcal{V} evoked non-zero responses on the output of the attack neuron). Examples of successful triggers are shown in Fig. 4.

Note that trigger images look very similar to the target ones despite their corresponding feature vectors being markedly different. Remarkably, despite the low-dimensional settings, this represents a success rate close to 65%.

To assess the viability of stealth attacks implemented in accordance with Scenario 2, we analysed network sensitivity to the removal of a single ReLU neuron from the fourth and third to the last layers of the network. These layers constitute an ‘attack layer’—a part of the network subject to potential stealth attacks. Results are shown in Fig. 5. The implementation of stealth attacks in this scenario (Scenario 2) follows the process detailed in Remark 7 and Section A.6 (see Appendix). For our particular network, we observed that there is a pool of neurons (85 out of the total 128) such that the removal of a single neuron from this pool does not have any effect on the network’s performance on the validation set \mathcal{V} unknown to the attacker. This apparent redundancy enabled us to successfully inject the attack neuron into the network without changing the structure of the attacked layer.

In those cases when the removal of a single neuron led to mismatches, the proportion of mismatched responses was below 3% with the majority of cases showing less than 1% of mismatched responses (see Fig. 5 for details).

4.3 Stealth attacks for a class of networks trained on the MNIST dataset

4.3.1 Network architecture. The general architecture of a deep convolutional neural network which we used in this set of experiments on MNIST data (LeCun *et al.*) is summarized in Table 3. The architecture features three fully connected layers, with layers 15–18 (shown in red) and layers 1–14 representing maps F and Φ , respectively.

This architecture is built on a standard benchmark example from Mathworks. The original basic network was appended by layers 13–16 to emulate dense fully connected structures present in popular deep learning models such as VGG16, VGG19. Note that the last six layers (layers 13–18) are equivalent to three dense layers with ReLU, ReLU and softmax activation functions, respectively, if the same

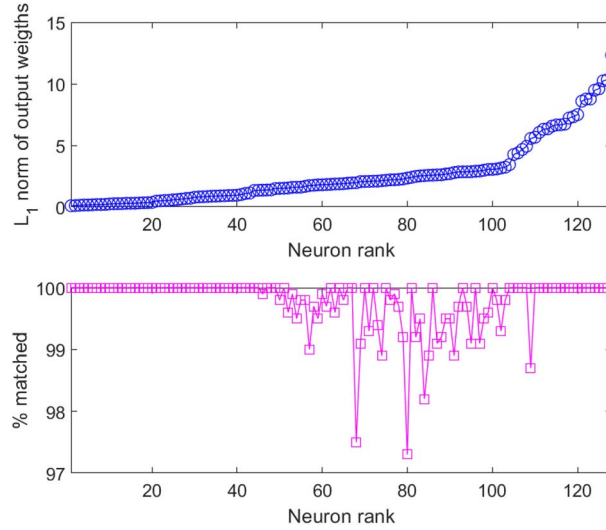


FIG. 5. (Colour online) Susceptibility to one-neuron stealth attack for the ResNet network (see Fig. 3) trained on the CIFAR-10 dataset. The top panel shows L_1 norms of the output weights of neurons formed by the fourth and third to the last layers of the network. Neurons are ordered in accordance with their rank: the neuron with the smallest output weight norm is assigned a rank of 1, and the neuron with the largest value output weight norm is assigned a rank of 128 (see A.6 for details). The bottom panel shows % of matched responses on the validation set \mathcal{V} between the original network and a modified network in which a single neuron with a particular rank is removed from the ‘attack layer’.

network is implemented in Tensorflow-Keras. Having three dense layers is not essential for the method, and the attacks can be implemented in other networks featuring ReLU or sigmoid neurons.

Outputs of the softmax layer assign class labels to images. Label 1 corresponds to digit ‘0’, label 2 to digit ‘1’ and label 10 to digit ‘9’.

The map Φ in (2.2) was associated with operations performed by layers 1 – 14 (shown in black in Table 3, Section 4.3.1), and the map F modelled the transformation from layer 15 to the first neuron in layer 17.

4.3.2 Training protocol. MATLAB’s version of the MNIST dataset of 10,000 images was split into a training set consisting of 7,500 images and a test set containing 2,500 images (see example code for details of implementation Tyukin *et al.* (2022)). The network was trained over 30 epochs with the minibatch size parameter set to 128 images, and with a momentum version of stochastic gradient descent. The momentum parameter was set to 0.9 and the learning rate was $0.01/(1 + 0.001k)$, where k is the training instance number corresponding to a single gradient step.

4.3.3 Construction of stealth attacks. Our stealth attack was a single ReLU neuron receiving $n = 200$ inputs from the outputs of ReLU neurons in layer 14. These outputs, for a given image \mathbf{u} , produced latent representations $\Phi(\mathbf{u})$. The ‘attack’ neuron was defined as $\mathfrak{A}(\cdot, \mathbf{w}, b) = D\text{ReLU}(\langle \cdot, \mathbf{w} \rangle - b)$, where the weight vector $\mathbf{w} \in \mathbb{R}^{200}$ and bias $b \in \mathbb{R}$ were determined in accordance with Algorithm 2.

In Scenario 1 (see Fig. 2) the output of the ‘attack’ neuron is added directly to the output of F (the first neuron in layer 17 of the network). Scenario 2 follows the process described in Remark 7 and Section A.6 below. In our experiments we placed the ‘attack’ neuron in layer 15 of the network. This was followed

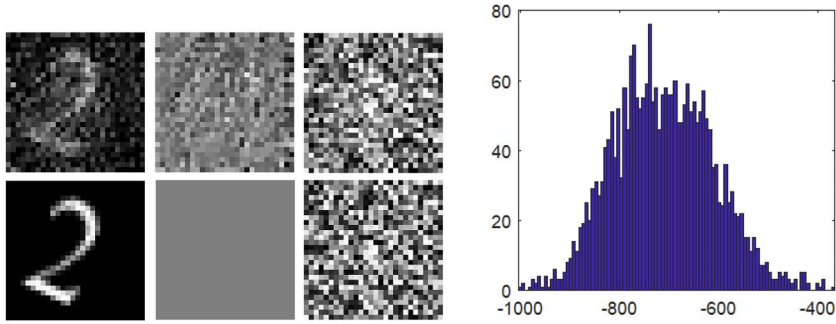


FIG. 6. (Colour online) *Left panel*: target images (bottom row) and their corresponding triggers (top row), $\delta = 1/3$, $\gamma = 0.9$. *Right panel*: Histogram of values $\langle \Phi(\mathbf{u}), \mathbf{w} \rangle - b$, $\mathbf{u} \in \mathcal{V}$ for the second trigger image in the top row in the left panel.

by adjusting weights in layer 17 in such a way that connections to neurons 2–10 from the ‘attack’ neuron were set to 0, and the weight of connection from the ‘attack’ neuron to neuron 1 in layer 17 was set to 1.

As the unknown verification set \mathcal{V} we used 99% of the test set. The remaining 1% of the test set was used to derive an empirical estimate of the value of R needed for the implementation of Algorithm 2. Other parameters in the algorithm were set as follows: $\delta = 1/3$, $\gamma = 0.9$, $\Delta = 50$ and $\varepsilon = 0$. A crucial step of the attack is step 3 in Algorithm 2 where a trigger image \mathbf{u}' is generated. As before, to find the trigger we used a standard gradient-based search method to solve the optimization problem in step 3 (see Section A.4). The values of α varied between experiments (see Tables A4, A5 and A6 for examples of specific values in some experiments).

By default, feature maps were constructed using only those neurons from the attack layer which return non-zero values for a given target image. In addition, in order to numerically explore the influence of dimension of the feature spaces on the attack success, we also constructed stealth attacks for feature maps

$$\tilde{\Phi} = T\Phi, \quad T \in \mathbb{R}^{d \times n}, \quad d < n, \quad (4.1)$$

where the rows of the matrix T are the first d principal components of the set $\mathcal{Z}(\mathbf{u}^*) = \{\mathbf{z} | \mathbf{z} = \Phi(\mathbf{u}^* + \xi_i), \xi_i \sim \mathcal{N}(0, I_m), i = 1, \dots, 5000\}$ with $\mathcal{N}(0, I_m)$ being the m -dimensional normal distribution with zero mean and identity covariance matrix. In these experiments, the value of d was set to $\lfloor 0.3N \rfloor$, where N is the number of principal components of the set $\mathcal{Z}(\mathbf{u}^*)$.

4.3.4 Performance of stealth attacks. Figure 6 illustrates how Algorithm 2 performed for the above networks. Three target images (bottom row in the left panel of Fig. 6 - digit 2, plain grey square and a random image) produced three corresponding trigger images (top row of the panel).

When elements (images) from the unknown validation set \mathcal{V} were presented to the network, the neuron’s output was 0. The histogram of values $\langle \Phi(\mathbf{u}), \mathbf{w} \rangle - b$, $\mathbf{u} \in \mathcal{V}$ is shown in Fig. 6 (right panel). As we can see, the neuron is firmly in the ‘silent mode’ and does not affect the classification for all elements of the unknown validation set \mathcal{V} . As per Remark 1 and in contrast to classical adversarial attacks, in the attacks we implemented in this section we were able to control the class to which trigger images are assigned (as opposed to adversarial attacks seeking to alter response of the classifier without regard to specific response).

Examples of successful trigger–target pairs for digits 0–9 are shown in Fig. 7.

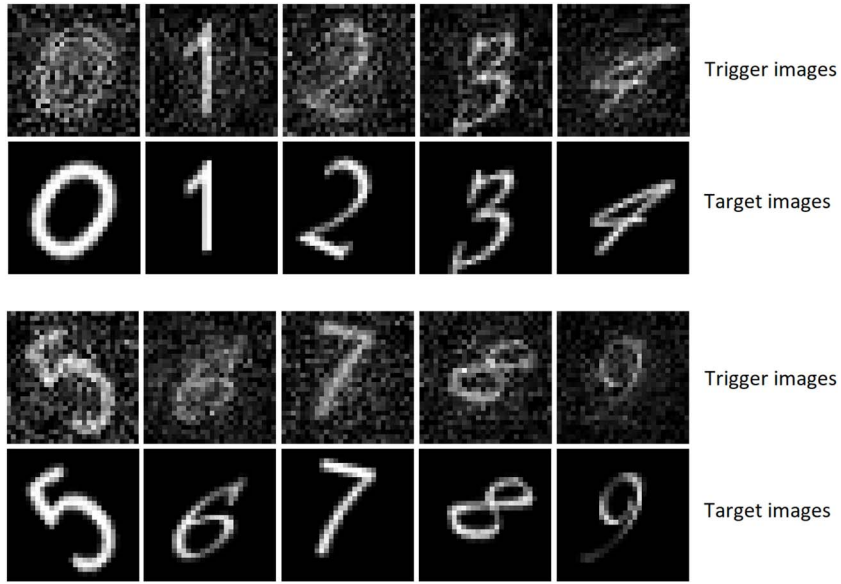


FIG. 7. Examples of trigger images u' (rows 1 and 3) and their corresponding 'target images' u^* (rows 2 and 4).

As is evident from these figures, the triggers retain significant resemblance of the original target images u^* . However, they look noticeably different from the target ones. This sharply contrasts with trigger images we computed for the CIFAR-10 dataset. One possible explanation could be the presence of three colour channels in CIFAR-10 compared with the single colour channel of MNIST which makes the corresponding perturbations less visible to the human eye.

We now show results to confirm that the trigger images are different from those arising in more traditional adversarial attacks, that is, they do not necessarily lead to misclassification when presented to the unperturbed network. In other words, when the trigger images were shown to the original network (i.e. trained network *before* it was we subjected to a stealth attack produced by Algorithm 2), in many cases the network returned a class label which coincided with the class labels of the target image. A summary for 20 different network instances and triggers is provided in Fig. 8. The red text highlights instances when the original classification of the target images did not match those of the trigger. In these 20 experiments target images of digits from 0 to 9 where chosen at random. In each row, the number of entries in the second column in the table in Fig. 8 corresponds to the number of times the digit in the first column was chosen as a 'target' in these experiments.

In these experiments, the rate of attack success in Scenarios 1 and 2 (Fig. 2) were 100% (20 out of 20) and 85% (17 out of 20), respectively. For the one-neuron attack in Scenario 2 we followed the approach discussed in Remark 7 with the procedure for selecting a neuron to be replaced described in the Appendix, Section A.6. When considering reduced-dimension feature maps $\tilde{\Phi}$ (see (4.1)) while maintaining the same values of δ ($\delta = 1/3$), the attacks' rate of success dropped to 50% (10 out of 20) in Scenario 1 and to 40% (8 out of 20) in Scenario 2. Yet, when the value of δ was increased to $2/3$, the rate of success recovered to 100% (20 out of 20) in Scenario 1 and 85% (17 out of 20) in Scenario 2, respectively. These results are consistent with bounds established by Theorems 2 and 3.

Target image of the trigger	Predicted label
0	0,3
1	4,7
2	2
3	3
4	4,4
5	5
6	6,6,6
7	7
8	8,5,6,9
9	9,9,9

FIG. 8. (Colour online) Retained information content in the trigger images.

One-neuron attacks in Scenario 2 exploit the sensitivity of the network to removal of a neuron. In order to assess this sensitivity, and consequently to gain further insight into the susceptibility of networks to one-neuron attacks, we explored five different architectures in which the sizes of the layer (layer 15) where the stealth attack neuron was planted were 400, 100, 75, 25 and 10. All other parameters of these networks were kept as shown in Table 3. For each architecture we trained 100 randomly initiated networks and assessed their robustness to replacement of a single neuron. Figure 9, left panel, shows the frequency with which replacing a neuron from layer 15 did not produce any change in network output on the validation set \mathcal{V} . The frequencies are shown as a function of the neuron susceptibility rank (see Appendix, Section A.6 for details). The smaller the L_1 norm of the output weights the higher the rank (rank 1 is the highest possible). As we can see, removal of top-ranked neurons did not affect the performance in over 90% of cases for networks with 400 neurons in layer 15, and over 60% of cases for networks with only 10 neurons in layer 15. Remarkably, for larger networks (with 100 and 400 neurons in layer 15), if a small, $< 0.3\%$ error margin on the validation set \mathcal{V} is allowed, then a one-neuron attack has the potential to be successful in 100% of cases (see Fig. 9, right panel). Notably, networks with smaller ‘attack’ layers appear to be substantially more fragile. If the latter networks break then the maximal observed degradation of their performance tends to be pronounced.

5. Conclusions

In this work we reveal and analyse new adversarial vulnerabilities for which large-scale networks may be particularly susceptible. In contrast to the largely empirical literature on the design of adversarial attacks, we accompany the algorithms with results on their probability of success. By design, research in this area looks at techniques that may be used to override the intended functionality of algorithms that are used within a decision-making pipeline, and hence may have serious negative implications in key sectors, including security, defence, finance and healthcare. By highlighting these shortcomings in the public domain we raise awareness of the issue, so that both algorithm designers and end-users can be better informed. Our analysis of these vulnerabilities also identifies important factors contributing to their successful exploitation: *dimension* of the network latent space, *accuracy* of executing the attack (expressed by parameter α) and *over-parameterization*. These determinants enable us to propose model design strategies to minimize the chances of exploiting the adversarial vulnerabilities that we identified.

Deeper, narrower graphs to reduce the dimension of the latent space. Our theoretical analysis suggests (bounds stated in Theorems 1 and 2) that the higher the dimension of latent spaces in state-

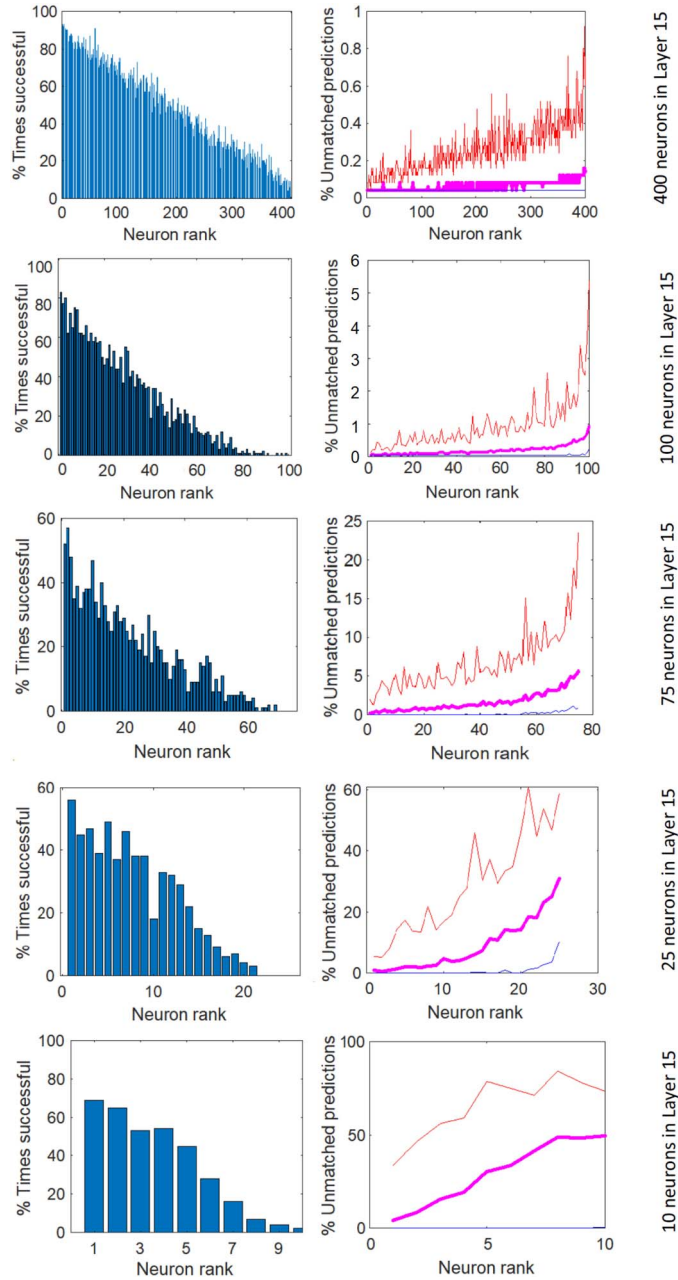


FIG. 9. (Colour online) Susceptibility to one-neuron stealth attack. *Left panel:* empirical frequencies of successful removal of neurons without any effect on the network output over the validation set. *Right panel:* % of unmatched responses for cases when removal of a neuron had an effect. Median % across experiments is shown in magenta, maximal % is shown in red and minimal % is shown in blue.

of-the-art deep neural networks the higher the chances of a successful one-neuron attack whereby an attacker replaces a single neuron in a layer. These chances approach one exponentially with dimension. One strategy to address this risk is to transform *wide* computational graphs in those parts of the network which are most vulnerable to open-box attacks into computationally equivalent *deeper but narrower* graphs. Such transformations can be done after training and just before the model is shared or deployed. An alternative is to employ dimensionality reduction approaches facilitating lower dimensional layer widths during and after the training.

Constraining attack accuracy. An ability to find triggers with arbitrarily high accuracy is another component of the attacker's success. Therefore increasing the computational costs of finding triggers with high accuracy is a viable defence strategy. A potential approach is to use gradient obfuscation techniques developed in the adversarial examples literature coupled with randomization, as suggested in Qiu *et al.* (2020). It is well known that gradient obfuscation may not always prevent an attacker from finding a trigger (Athalye *et al.*, 2018). Yet, making this process as complicated and computationally intensive as possible would contribute to increased security.

Pruning to reduce redundancy and the dimension of latent space. We demonstrate theoretically and confirm in experiments that over-parameterization and high intrinsic dimension of network latent spaces, inherent in many deep learning models, can be readily exploited by an adversary if models are freely shared and exchanged without control. In order to deny these opportunities to the attacker, removing susceptible neurons with our procedure in Section A.6 offers a potential remedy. More generally, employing network pruning (Cheng *et al.*, 2018; Blalock *et al.*, 2020; Mirkes, 2020; Tanaka *et al.*, 2020) and enforcing low dimensionality of latent spaces as a part of model production pipelines would offer further protection against one-neuron attacks.

Network Hashing. In addition to the strategies above, which stem from our theoretical analysis of stealth attacks, another defence mechanism is to use fast network hashing algorithms executed in parallel with inference processes. The hash codes these algorithms produce will enable the owner to detect unwanted changes to the model after it was downloaded from a safe and trusted source.

Our theoretical and practical analysis of the new threats are in no way complete. In our experiments we assumed that pixels in images are real numbers. In some contexts they may take integer values. We did not assess how changes in numerical precision would affect the threat, and did not provide conditions under which redundant neurons exist. Moreover, as we showed in Section A.5, our theoretical bounds could be conservative. The theory presented in this work is fully general in the sense that it does not require any metric in the input space. An interesting question is how one can use the additional structure arising when the input space has a metric to find trigger inputs that are closer to their corresponding targets. Finally, vulnerabilities discovered here are intrinsically linked with AI maintenance problems discussed in Gorban *et al.* (2018, 2019, 2021). Exploring these issues are topics for future research.

Nevertheless, the theory and empirical evidence which we present in this work make it clear that existing mitigation strategies must be strengthened in order to guard against vulnerability to new forms of stealth attack. Because the 'inevitability' results that we derived have constructive proofs, our analysis offers promising options for the development of effective defences.

Funding

This work is supported in part by the UKRI, EPSRC [UKRI Turing AI Fellowship ARaISE EP/V025295/1 and UKRI Trustworthy Autonomous Systems Node in Verifiability EP/V026801/1 to I.Y.T., EP/V046527/1 and EP/P020720/1 to D.J.H., and EP/V046527/1 to A.B.].

REFERENCES

- AKHTAR, N. & MIAN, A. (2018) Threat of adversarial attacks on deep learning in computer vision: a survey. *IEEE Access*, **6**, 14410–14430.
- ALLEN, B., AGARWAL, S., KALPATHY-CRAMER, J. & DREYER, K. (2019) Democratizing AI. *J. Am. Coll. Radiol.*, **16**, 961–963.
- ANTUN, V., RENNA, F., POON, C., ADCOCK, B. & HANSEN, A. C. (2020) On instabilities of deep learning in image reconstruction and the potential costs of AI. *Proc. Natl. Acad. Sci.*, **117**, 30088–30095.
- ATHALYE, A., CARLINI, N. & WAGNER, D. (2018) Obfuscated gradients give a false sense of security: Circumventing defenses to adversarial examples. In *International Conference on Machine Learning*, pages 274–283. PMLR.
- BASTOUNIS, A., HANSEN, A. C. & VLAČIĆ, V. (2021a) The extended Smale’s 9th problem—on computational barriers and paradoxes in estimation, regularisation, computer-assisted proofs and learning. arXiv preprint arXiv:2110.15734.
- BASTOUNIS, A., HANSEN, A. C. & VLAČIĆ, V. (2021b) The mathematics of adversarial attacks in AI—why deep learning is unstable despite the existence of stable neural networks. arXiv preprint arXiv:2109.06098.
- BEERENS, L. & HIGHAM, D. J. (2023) Adversarial ink: Componentwise backward error attacks on deep learning, IMA Journal of Applied Mathematics, hxad017, <https://doi.org/10.1093/imamat/hxad017>.
- BIGGIO, B. & ROLI, F. (2018) Wild patterns: ten years after the rise of adversarial machine learning. *Pattern Recognit.*, **84**, 317–331.
- BLALOCK, D. W., GONZALEZ ORTIZ, J. J., FRANKLE, J. & GUTTAG, J. V. (2020) What is the state of neural network pruning? In (DHILLON, I. S., PAPAILOPOULOS, D. S. & SZE, V., editors), *Proceedings of Machine Learning and Systems 2020, MLSys 2020*, Austin, TX, USA, March 2–4. mlsys.org, 2020.
- CHENG, Y., WANG, D., ZHOU, P. & ZHANG, T. (2018) Model compression and acceleration for deep neural networks: the principles, progress, and challenges. *IEEE Signal Process. Mag.*, **35**, 126–136.
- COLBROOK, M. J., ANTUN, V. & HANSEN, A. C. (2022) The difficulty of computing stable and accurate neural networks: on the barriers of deep learning and Smale’s 18th problem. *Proc. Natl. Acad. Sci.*, **119**, e2107151119.
- EBRAHIMI, J., RAO, A., LOWD, D. & DOU, D. (2018) HotFlip: White-box adversarial examples for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 31–36, Melbourne, Australia, Association for Computational Linguistics.
- FINLAYSON, S. G., BOWERS, J. D., ITO, J., ZITTRAIN, J. L., BEAM, A. L. & KOHANE, I. S. (2019) Adversarial attacks on medical machine learning. *Science*, **363**, 1287–1289.
- GHASSEMI, M., OAKDEN-RAYNER, L. & BEAM, A. L. (2021) The false hope of current approaches to explainable artificial intelligence in health care. *Lancet Digit. Health*, **3**, e745–e750.
- GHODSI, Z., GU, T. & GARG, S. (2017) Safetynets: Verifiable execution of deep neural networks on an untrusted cloud. *Advances in Neural Information Processing Systems* (GUYON, I., LUXBURG VON, U., BENGIO, S., WALLACH, H., FERGUS, R., VISHWANATHAN, S. & GARNETT, R. eds), vol. **30**. Curran Associates, Inc.
- GORBAN, A. N., GRECHUK, B., MIRKES, E. M., STASENKO, S. V. & TYUKIN, I. Y. (2021) High-dimensional separability for one-and few-shot learning. *Entropy*, **23**, 1090.
- GORBAN, A. N., MAKAROV, V. A. & TYUKIN, I. Y. (2019) The unreasonable effectiveness of small neural ensembles in high-dimensional brain. *Phys. Life Rev.*, **29**, 55–88.
- GORBAN, A. N. & ROSSIEV, D. A. (1996) *Neural networks on personal computer*. Novosibirsk: Nauka (RAN).
- GORBAN, A. N., GOLUBKOV, A., GRECHUK, B., MIRKES, E. M. & Tyukin, I. Y. (2018) Correction of AI systems by linear discriminants: probabilistic foundations. *Inform. Sci.*, **466**, 303–322.
- GU, T., DOLAN-GAVITT, B. & GARG, S. (2017) Badnets: identifying vulnerabilities in the machine learning model supply chain. arXiv preprint arXiv:1708.06733.
- HAHNLOSER, R.H.R., SARPESHKAR, R., MAHOWALD, M. A., DOUGLAS, R. J. & SEUNG, H. S. (2000) Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature*, **405**, 947–951.
- HIGHAM, C. F. & HIGHAM, D. J. (2019) Deep learning: an introduction for applied mathematicians. *SIAM Rev.*, **61**, 860–891.

- HUANG, X., KROENING, D., RUAN, W., SHARP, J., SUN, Y., THAMO, M., WU, E. & YI, X. (2020) A survey of safety and trustworthiness of deep neural networks: verification, testing, adversarial attack and defence, and interpretability. *Comput. Sci. Rev.*, **37**, 100270.
- KRIZHEVSKY, A. & HINTON, G. (2009) *Learning multiple layers of features from tiny images* Technical report., Toronto, Ontario: University of Toronto.
- LECUN, Y., CORTES, C. & BURGES, C. J. C. (2019) *The MNIST database of handwritten digits*. Accessed June 17, 2019.
- LIU, T., WEN, W. & JIN, Y. (2018) SIN2: Stealth infection on neural network - a low-cost agile neural trojan attack methodology. In *2018 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*, pages 227–230.
- MANOJ, N. & BLUM, A. (2021) Excess capacity and backdoor poisoning. *Adv. Neural Inf. Process. Syst.*, **34**, 20373–20384.
- MIRKES, E. M. (2020) Artificial neural network pruning to extract knowledge. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- NEYSHABUR, B., SEDGHI, H. & ZHANG, C. (2020) What is being transferred in transfer learning? *Advances in Neural Information Processing Systems* (LAROCHELLE, H., RANZATO, M., HADSELL, R., BALCAN, M. F. & LIN, H. eds), vol. **33**. Curran Associates, Inc., pp. 512–523.
- NORTHCUTT, C. G., ATHALYE, A. & MUELLER, J. (2021) Pervasive label errors in test sets destabilize machine learning benchmarks. In (VANSCHOREN, J. and YEUNG, S., editors), *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks 2021*, December 2021, virtual.
- QIU, H., ZENG, Y., ZHENG, Q., ZHANG, T., QIU, M. & MEMMI, G. (2020) Mitigating advanced adversarial attacks with more advanced gradient obfuscation techniques. arXiv preprint arXiv:2005.13712.
- REN, K., ZHENG, T., QIN, Z. & LIU, X. (2020) Adversarial attacks and defenses in deep learning. *Engineering*, **6**, 346–360.
- SHAFABI, A., HUANG, W. R., STUDER, C., FEIZI, S. & GOLDSTEIN, T. (2019) Are adversarial examples inevitable? *International Conference on Learning Representations (ICLR)*.
- SHAFABI, A., NAJIBI, M., XU, Z., DICKERSON, J., DAVIS, L. S. & GOLDSTEIN, T. (2020) Universal adversarial training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume **34**, pages 5636–5643.
- STOJNIC, R., TAYLOR, R. & KARDAS, M. et al. Browse state-of-the-art. Image classification with CIFAR-10. <https://paperswithcode.com/sota/image-classification-on-cifar-10>.
- SU, J., VARGAS, D. V. & SAKURAI, K. (2019) One pixel attack for fooling deep neural networks. *IEEE Trans. Evol. Comput.*, **23**, 828–841.
- SZEGEDY, C., ZAREMBA, W., SUTSKEVER, I., BRUNA, J., ERHAN, D., GOODFELLOW, I. & FERGUS, R. (2014) Intriguing properties of neural networks. *Proceedings of the 2nd International Conference on Learning Representations, ICLR 2014, April 2014*.
- TANAKA, H., KUNIN, D., YAMINS, D. L. & GANGULI, S. (2020) Pruning neural networks without any data by iteratively conserving synaptic flow. In (LAROCHELLE, H., RANZATO, M., HADSELL, R., BALCAN, M. F. & LIN, H., editors), *Advances in Neural Information Processing Systems*, volume **33**, pages 6377–6389. Curran Associates, Inc.
- TYUKIN, I. Y., HIGHAM, D. J., BASTOUNIS, A., WOLDEGEORGIS, E. & GORBAN, A. N. (2022) Example code for open-box attacks. <https://github.com/tyukin/Stealth-adversarial-attacks>.
- TYUKIN, I. Y., HIGHAM, D. J. & GORBAN, A. N. (2020) On adversarial examples and stealth attacks in artificial intelligence systems. In *2020 International Joint Conference on Neural Networks (IJCNN)*, pages 1–6. IEEE.
- WU, M., WICKER, M., HUANG, W. R. A. X. & KWIATKOWSKA, M. (2020) A game-based approximate verification of deep neural networks with provable guarantees. *Theor. Comput. Sci.*, **807**, 298–329.

A. Appendix. Proofs of theorems and supplementary results

A.1 Proof of Theorem 1

The proof is split into four parts. We begin *assuming* that latent representations $\mathbf{x}_i = \Phi(\mathbf{u}_i)$ of all elements \mathbf{u}_i from the set \mathcal{V} belong to the unit ball \mathbb{B}_n centred at the origin (for simplicity of notation the unit n -ball centred at 0 is denoted \mathbb{B}_n , and the unit $n - 1$ sphere centred at 0 is denoted \mathbb{S}_{n-1}). The main thrust of the proof is to establish lower bounds on the probability of the event

$$\mathcal{E}^* : \gamma \langle \mathbf{x}', \mathbf{x}' \rangle \geq \langle \mathbf{x}', \mathbf{x}_i \rangle \text{ for all } \mathbf{x}_i = \Phi(\mathbf{u}_i) : \mathbf{u}_i \in \mathcal{V}, \mathbf{x}_i \in \mathbb{B}_n. \quad (\text{A.1})$$

These bounds are established in *Parts 1 and 2* of the proof. Similar events have been shown to play an important role in the problem of AI error correction (Gorban *et al.*, 2018, 2021)—a somewhat dual task to stealth attacks considered here.

Then we proceed with constructing weights (both input and output) and biases of the function g so that the modified map F_a delivers a solution of Problem 1. This is shown in *Part 3*. Finally, we remove the assumption that $\mathbf{x}_i \in \mathbb{B}_n$ and show how the weights and biases need to be adapted so that the resulting adapted map F_a is a solution of Problem 1 for $\mathbf{x}_i \in \mathbb{B}_n(0, R)$ where $R > 0$ is a given number. This is demonstrated in *Part 4* of the proof.

Part 1. Probability bound 1 on the event \mathcal{E}^ .* Suppose that $R = 1$. Let $\mathbf{x}_i, i = 1, \dots, M$ be an arbitrary element from \mathcal{V} , let \mathbf{x} be drawn randomly from an equidistribution in $\mathbb{S}_{n-1}(0, \delta)$ and let $R\mathbf{x}' = \Phi(\mathbf{u}')$ be a vector within an $\alpha\delta$ -distance from \mathbf{x} :

$$\|\mathbf{x} - \mathbf{x}'\| \leq \alpha\delta.$$

The assumption that the set $\mathbb{S}_{n-1}(0, R\delta)$ is $\delta\alpha R$ -input reachable for the map Φ assures that such \mathbf{x}' exists. Moreover, since $\alpha \in [0, 1]$, Algorithm 1 ensures that $\|\mathbf{x}'\| \neq 0$ (see (3.4)).

Notice that the constraints $(1 + \alpha)\delta \leq 1$ and $\gamma \in [0, 1]$ imply that $\gamma\|\mathbf{x}'\| \leq \gamma(\|\mathbf{x} - \mathbf{x}'\| + \|\mathbf{x}\|) \leq \gamma(\alpha\delta + \delta) \leq 1$ (i.e. $\arccos(\gamma\|\mathbf{x}'\|)$ always exists). Set $\beta(\gamma, \alpha, \|\mathbf{x}'\|) := \arccos(\gamma\|\mathbf{x}'\|) + \arccos((1 - \alpha^2)^{1/2})$. We claim that if the event

$$\mathcal{E} : \left(\frac{\mathbf{x}}{\|\mathbf{x}\|}, \frac{\mathbf{x}_i}{\|\mathbf{x}_i\|} \right) < \cos(\beta(\gamma, \alpha, \|\mathbf{x}'\|)) \text{ for all } \mathbf{x}_i \in \mathcal{V} \quad (\text{A.2})$$

occurs then the event \mathcal{E}^* (defined in (A.1)) must occur too.

To show this, assume that (A.2) holds and fix $i \in \{1, 2, \dots, M\}$. By the definition of the dot product in \mathbb{R}^n , we obtain

$$|\angle(\mathbf{x}_i, \mathbf{x})| \geq \beta = \arccos(\gamma\|\mathbf{x}'\|) + \arccos((1 - \alpha^2)^{1/2}). \quad (\text{A.3})$$

Consider angles between $\mathbf{x}_i, \mathbf{x}'$ and \mathbf{x}', \mathbf{x} as follows: we let

$$A^*(\mathbf{x}, \mathbf{x}_i) := \min_{\mathbf{z}: \|\mathbf{x} - \mathbf{z}\| \leq \alpha\delta} |\angle(\mathbf{x}_i, \mathbf{z})|, \quad B^*(\mathbf{x}) := \max_{\mathbf{z}: \|\mathbf{x} - \mathbf{z}\| \leq \alpha\delta} |\angle(\mathbf{z}, \mathbf{x})|.$$

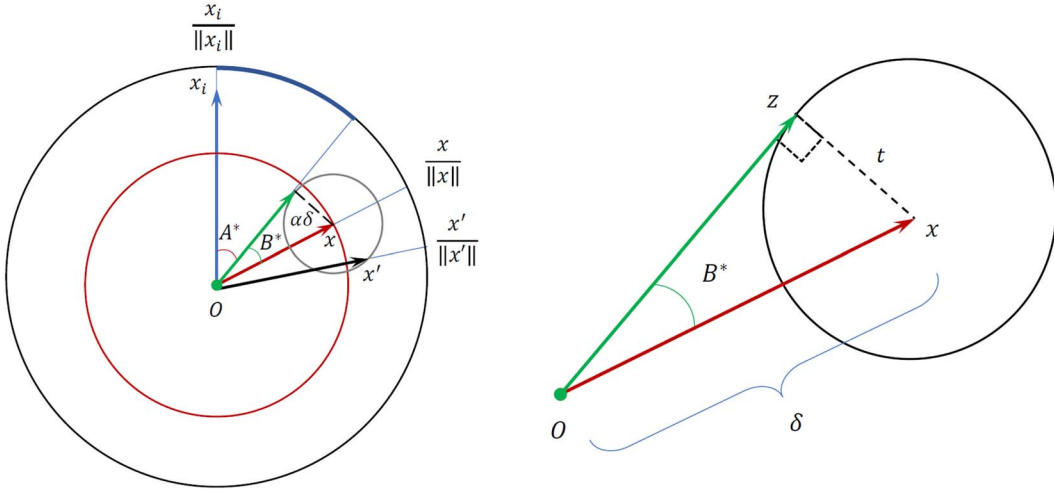


FIG. A.1. (Colour online) A diagram assisting with the proof of Theorem 1. The *left panel* illustrates the set-up and main ingredients of the argument. The *right panel* illustrates the derivation of $B^*(x)$. For $t = 0$ the expression is trivial. Let $t \neq 0$. Straightforward calculations show that for z corresponding to $\max_{z: \|x-z\|=t} |\angle(z, x)|$, the vectors z and $z - x$ must be orthogonal. Hence $\cos \angle(z, x) = \sqrt{\delta^2 - t^2}/\delta$.

We note that (using the triangle inequality for angular distance)

$$A^*(x, x_i) + B^*(x) \geq |\angle(x_i, x)| \quad (\text{A.4})$$

(see Fig. A.1) and

$$B^*(x) = \sup_{t \in [0, \alpha\delta]} \max_{z: \|x-z\|=t} |\angle(z, x)| = \sup_{t \in [0, \alpha\delta]} \arccos\left(\frac{\sqrt{\delta^2 - t^2}}{\delta}\right) = \arccos(\sqrt{1 - \alpha^2}), \quad (\text{A.5})$$

where the second equality is illustrated in Fig. A.1 and the final equality follows because \arccos is decreasing.

Combining (A.4), (A.5) and (A.3) gives

$$A^*(x, x_i) \geq |\angle(x_i, x)| - B^*(x) = |\angle(x_i, x)| - \arccos(\sqrt{1 - \alpha^2}) \geq \arccos(\gamma \|x'\|)$$

so that (by the fact that cosine is decreasing, the assumption that $\|x_i\| \leq 1$, and noting that $|\angle(x', x_i)| \geq A^*(x, x_i)$)

$$\langle x', x_i \rangle = \|x'\| \|x_i\| \cos(\angle(x', x_i)) \leq \|x'\| \cos(A^*(x, x_i)) \leq \gamma \langle x', x' \rangle$$

completing the proof that if the event (A.2) occurs then the event (A.1) must occur too.

Consider events

$$\mathcal{E}_i : \left\langle \frac{\mathbf{x}}{\|\mathbf{x}\|}, \frac{\mathbf{x}_i}{\|\mathbf{x}_i\|} \right\rangle \geq \cos(\beta(\gamma, \alpha, \|\mathbf{x}'\|)).$$

The probability that \mathcal{E}_i occurs is equal to the area of the spherical cap

$$C(\mathbf{x}_i, \beta(\gamma, \alpha, \|\mathbf{x}'\|)) = \left\{ z \in \mathbb{S}_{n-1} \mid \left\langle \frac{\mathbf{x}_i}{\|\mathbf{x}_i\|}, z \right\rangle \geq \cos(\beta(\gamma, \alpha, \|\mathbf{x}'\|)) \right\}$$

divided by the area of \mathbb{S}_{n-1} :

$$P(\mathcal{E}_i) = \frac{A_{n-1}(C(\mathbf{x}_i, \beta(\gamma, \alpha, \|\mathbf{x}'\|)))}{A_{n-1}(\mathbb{S}_{n-1})} \quad (\text{A.6})$$

(here $A_{n-1}(\mathbb{S}_{n-1})$ stands for the area of the unit $n-1$ sphere \mathbb{S}_{n-1} and $A_{n-1}(C(\mathbf{x}_i, \beta(\gamma, \alpha, \|\mathbf{x}'\|)))$ denotes the area of the spherical cap $C(\mathbf{x}_i, \beta(\gamma, \alpha, \|\mathbf{x}'\|))$). It is well known that

$$A_{n-1}(C(\mathbf{x}_i, \beta(\gamma, \alpha, \|\mathbf{x}'\|))) = A_{n-2}(\mathbb{S}_{n-2}) \int_0^{\beta(\gamma, \alpha, \|\mathbf{x}'\|)} \sin^{n-2}(\theta) d\theta, \text{ and } A_{n-1}(\mathbb{S}_{n-1}) = \frac{2\pi^{\frac{n}{2}}}{\Gamma(\frac{n}{2})}.$$

Hence

$$P(\mathcal{E}_i) = \frac{A_{n-2}(\mathbb{S}_{n-2})}{A_{n-1}(\mathbb{S}_{n-1})} \int_0^{\beta(\gamma, \alpha, \|\mathbf{x}'\|)} \sin^{n-2}(\theta) d\theta.$$

Using the fact that the inequality

$$P(\text{not } \mathcal{E}_1 \wedge \cdots \wedge \text{not } \mathcal{E}_M) \geq 1 - \sum_{i=1}^M P(\mathcal{E}_i) \quad (\text{A.7})$$

holds true for any events \mathcal{E}_i , that $\mathcal{E} = \text{not } \mathcal{E}_1 \wedge \cdots \wedge \text{not } \mathcal{E}_M$ and that $P(\mathcal{E}^*) \geq P(\mathcal{E})$ (as \mathcal{E} implies \mathcal{E}^*), we can conclude that

$$P(\langle \mathbf{x}', \mathbf{x}_i \rangle < \gamma \langle \mathbf{x}', \mathbf{x}' \rangle \text{ for all } \mathbf{x}_i \in \mathcal{V}) \geq 1 - M \frac{A_{n-2}(\mathbb{S}_{n-2})}{A_{n-1}(\mathbb{S}_{n-1})} \int_0^{\beta(\gamma, \alpha, \|\mathbf{x}'\|)} \sin^{n-2}(\theta) d\theta,$$

or, equivalently,

$$P(\langle \mathbf{x}', \mathbf{x}_i \rangle < \gamma \langle \mathbf{x}', \mathbf{x}' \rangle \text{ for all } \mathbf{x}_i \in \mathcal{V}) \geq 1 - M \frac{1}{\pi^{\frac{1}{2}}} \frac{\Gamma(\frac{n}{2})}{\Gamma(\frac{n-1}{2})} \int_0^{\beta(\gamma, \alpha, \|\mathbf{x}'\|)} \sin^{n-2}(\theta) d\theta. \quad (\text{A.8})$$

Finally, noticing that

$$\beta(\gamma, \alpha, \|\mathbf{x}'\|) \leq \arccos(\gamma(1 - \alpha)\delta) + \arccos((1 - \alpha^2)^{1/2})$$

we can conclude that Algorithm 1 ensures that event \mathcal{E}^* in (A.1) occurs with probability at least (3.8).

Part 2. Proving the bound (3.9) Since \cos is non-negative and decreasing on $[0, \pi/2]$ and $\arccos(\varphi(\gamma, \delta, \alpha)) \in [0, \pi/2]$ we have $\cos(\theta)/\varphi(\gamma, \delta, \alpha) = \cos(\theta)/\cos(\arccos(\varphi(\gamma, \delta, \alpha))) \geq 1$ for every $\theta \in [0, \arccos(\varphi(\gamma, \delta, \alpha))]$. Hence

$$\begin{aligned} \frac{1}{\Gamma\left(\frac{n-1}{2}\right)} \int_0^{\arccos(\varphi(\gamma, \delta, \alpha))} \sin^{n-2}(\theta) d\theta &\leq \frac{1}{\Gamma\left(\frac{n-1}{2}\right)} \int_0^{\arccos(\varphi(\gamma, \delta, \alpha))} \frac{\cos(\theta)}{\varphi(\gamma, \delta, \alpha)} \sin^{n-2}(\theta) d\theta \\ &= \frac{1}{\Gamma\left(\frac{n-1}{2}\right)} \frac{1}{\varphi(\gamma, \delta, \alpha)} \left[\frac{\sin^{n-1}(\arccos(\varphi(\gamma, \delta, \alpha)))}{n-1} - \frac{\sin^{n-1}(0)}{n-1} \right] \\ &= \frac{1}{2\Gamma\left(\frac{n}{2} + \frac{1}{2}\right)} \frac{1}{\varphi(\gamma, \delta, \alpha)} \left(1 - \varphi(\gamma, \delta, \alpha)^2\right)^{\frac{n-1}{2}}. \end{aligned}$$

As a result,

$$P(\mathcal{E}_i) \leq \frac{1}{2\pi^{\frac{1}{2}}} \frac{\Gamma\left(\frac{n}{2}\right)}{\Gamma\left(\frac{n}{2} + \frac{1}{2}\right)} \frac{1}{\varphi(\gamma, \delta, \alpha)} \left(1 - \varphi(\gamma, \delta, \alpha)^2\right)^{\frac{n-1}{2}}, \quad (\text{A.9})$$

from which we obtain (3.9).

The same estimate can be obtained via an alternative geometrical argument by recalling that

$$\begin{aligned} A_{n-1}(C(\mathbf{x}_i, \arccos(\varphi(\gamma, \delta, \alpha)))) \frac{1}{n} &= V_n(C(\mathbf{x}_i, \arccos(\varphi(\gamma, \delta, \alpha)))) \\ &+ \frac{\varphi(\gamma, \delta, \alpha)}{n} V_{n-1} \left(B_{n-1} \left(\frac{\mathbf{x}_i}{\|\mathbf{x}_i\|} \varphi(\gamma, \delta, \alpha), (1 - \varphi(\gamma, \delta, \alpha)^2)^{1/2} \right) \right), \end{aligned}$$

and then estimating the volume of the spherical cap $C(\mathbf{x}_i, \arccos(\varphi(\gamma, \delta, \alpha)))$ by the volume of the corresponding spherical cone containing $C(\mathbf{x}_i, \arccos(\varphi(\gamma, \delta, \alpha)))$ whose base is the disc centred at $\frac{\mathbf{x}_i}{\|\mathbf{x}_i\|} \varphi(\gamma, \delta, \alpha)$ with radius $(1 - \varphi(\gamma, \delta, \alpha)^2)^{1/2}$, and whose height is $\frac{1}{\varphi(\gamma, \delta, \alpha)} - \varphi(\gamma, \delta, \alpha)$.

Part 3. Construction of the structural adversarial perturbation. Having established bounds on the probability of event \mathcal{E}^* in (A.1), let us now proceed with determining a map F_a which is solution to Problem 1 assuming that $\mathbf{x}_i \in \mathbb{B}_n$. Suppose that the event \mathcal{E}^* holds true.

By construction, since $R = 1$, in Algorithm 1 we have

$$\mathbf{w} = \kappa \mathbf{x}', \quad \kappa > 0, \quad (\text{A.10})$$

$$b = \kappa \left(\frac{1 + \gamma}{2} \right) \|\mathbf{x}'\|^2, \quad (\text{A.11})$$

and

$$\mathfrak{A}(\cdot, \mathbf{w}, b) = Dg \left(\kappa \left(\langle \cdot, \mathbf{x}' \rangle - \left(\frac{1+\gamma}{2} \right) \|\mathbf{x}'\|^2 \right) \right).$$

Since the function g is monotone,

$$|\mathfrak{A}(\mathbf{x}_i, \mathbf{w}, b)| \leq Dg \left(-\kappa \left(\frac{1-\gamma}{2} \|\mathbf{x}'\|^2 \right) \right) \text{ for all } \mathbf{x}_i = \Phi(\mathbf{u}_i), \mathbf{u}_i \in \mathcal{V}, \mathbf{x}_i \in \mathbb{B}_n.$$

Writing

$$z = \frac{1-\gamma}{2} \|\mathbf{x}'\|^2$$

we note from (3.6) that the values of D and κ are chosen so that

$$Dg(-\kappa z) \leq \varepsilon \text{ and } Dg(\kappa z) \geq \Delta. \quad (\text{A.12})$$

Part 4. Generalization to the case when $\mathbf{x}_i \in \mathbb{B}_n(0, R)$. Consider variables

$$\tilde{\mathbf{x}}_i = \frac{\mathbf{x}_i}{R}.$$

It is clear that $\tilde{\mathbf{x}}_i \in \mathbb{B}_n$. Suppose now that

$$\langle \mathbf{x}', \tilde{\mathbf{x}}_i \rangle \leq \gamma \|\mathbf{x}'\|^2 \text{ for all } \tilde{\mathbf{x}}_i = \frac{\Phi(\mathbf{u}_i)}{R}, \mathbf{u}_i \in \mathcal{V}, \mathbf{x}_i = \Phi(\mathbf{u}_i) \in \mathbb{B}_n(0, R). \quad (\text{A.13})$$

Probability bounds on the above event have already been established in *Parts 1 and 2* of the proof.

In this general case, Algorithm 1 uses

$$\mathbf{w} = \kappa \left(\frac{\mathbf{x}'}{R} \right), \kappa > 0, \quad b = \kappa \left(\frac{1+\gamma}{2} \right) \|\mathbf{x}'\|^2,$$

and

$$\mathfrak{A}(\cdot, \mathbf{w}, b) = Dg \left(\kappa \left(\left\langle \cdot, \frac{\mathbf{x}'}{R} \right\rangle - \left(\frac{1+\gamma}{2} \right) \|\mathbf{x}'\|^2 \right) \right).$$

The values of κ and D are chosen so that (A.12) in *Part 3* holds, and hence

$$|\mathfrak{A}(\Phi(\mathbf{u}_i), \mathbf{w}, b)| \leq \varepsilon \text{ for all } \mathbf{u}_i \in \mathcal{V}$$

and

$$\mathfrak{A}(\mathbf{x}'R, \mathbf{w}, b) = \mathfrak{A}(\Phi(\mathbf{u}'), \mathbf{w}, b) \geq \Delta.$$

This, together with bounds (A.8), (A.9) in *Parts 1 and 2* on the probability of event (A.13) concludes the proof. \square

A.2 Proof of Theorem 2

Consider an AI system (2.2) which satisfies Assumption 2. In addition to this system, consider a ‘virtual’ one whose maps F and Φ are replaced with

$$\tilde{\Phi}(\mathbf{u}) = \Phi(\mathbf{u}) - \Phi(\mathbf{u}^*), \quad \tilde{F}(\mathbf{x}) = F(\mathbf{x} + \Phi(\mathbf{u}^*)). \quad (\text{A.14})$$

According to the definition of $\tilde{F} \circ \tilde{\Phi}$, domains of the definition of $\tilde{F} \circ \tilde{\Phi}$ and $F \circ \Phi$ coincide, and

$$\tilde{F} \circ \tilde{\Phi}(\mathbf{u}) = F \circ \Phi(\mathbf{u}) \text{ for all } \mathbf{u} \in \mathcal{U}. \quad (\text{A.15})$$

For this virtual system, $\tilde{F} \circ \tilde{\Phi}$, Assumption 1 is satisfied. Moreover, if an input \mathbf{u}'

$$\mathbf{x}' = \frac{\Phi(\mathbf{u}')}{R} - \frac{\Phi(\mathbf{u}^*)}{R}$$

satisfies condition (3.4), then these conditions are satisfied for $\mathbf{x}' = \tilde{\Phi}(\mathbf{u}')/R$ (the converse holds true too), and $\|\mathbf{x}'\| \neq 0$.

According to Theorem 1, if we define the stealth attack parameters as in (3.5), (3.6) by

$$\mathfrak{A}\left(\cdot, \kappa \frac{\mathbf{x}'}{R}, b\right) = \text{Dg}\left(\left\langle \cdot, \kappa \frac{\mathbf{x}'}{R} \right\rangle - b\right), \quad b = \kappa \left(\frac{1+\gamma}{2}\right) \|\mathbf{x}'\|^2,$$

then the above is a solution of Problem 1 for the virtual AI system $\tilde{F} \circ \tilde{\Phi}$ with probability bounded from below as in (3.8), (3.9).

Notice that

$$\mathfrak{A}\left(\Phi(\mathbf{u}) - \Phi(\mathbf{u}^*), \kappa \frac{\mathbf{x}'}{R}, b\right) = \text{Dg}\left(\left\langle \Phi(\mathbf{u}), \kappa \frac{\mathbf{x}'}{R} \right\rangle - \left(\kappa \left(\frac{1+\gamma}{2}\right) \|\mathbf{x}'\|^2 + \left\langle \Phi(\mathbf{u}^*), \kappa \frac{\mathbf{x}'}{R} \right\rangle\right)\right),$$

which coincides with the attack used in Algorithm 2. Hence, taking (A.15) into account, we can conclude that Algorithm 2 returns a solution to Problem 1 for the original AI system $F \circ \Phi$ with probabilities of success satisfying (3.8) and (3.9). \square

A.3 Proof of Theorem 3

In a similar way to the proof of Theorem 2, we begin with considering a virtual system with maps $\tilde{\Phi}$ and \tilde{F} defined as in (A.14). We observe that Assumption 3 implies that Assumption 2 holds true. As we have seen before, domains of the definition of $\tilde{F} \circ \tilde{\Phi}$ and $F \circ \Phi$ coincide, and Assumption 2 implies that Assumption 1 holds true for the virtual system.

As in the proof of Theorem 2, we will consider the vector \mathbf{x}'

$$\mathbf{x}' = \frac{\Phi(\mathbf{u}')}{R} - \frac{\Phi(\mathbf{u}^*)}{R} = \frac{\tilde{\Phi}(\mathbf{u}')}{R}$$

satisfying condition (3.4), where \mathbf{x} is drawn from an equidistribution on the sphere $\mathbb{S}_{n_p-1}(0, \delta)$. We set $\tilde{\mathbf{c}} = \mathbf{c} - \Phi(\mathbf{u}^*)$, and let $\tilde{\mathbf{c}} = \tilde{\mathbf{c}}^1 + \tilde{\mathbf{c}}^2$ where $\tilde{\mathbf{c}}^1 \in \text{span}\{h_1, \dots, h_{n_p}\}$ and $\tilde{\mathbf{c}}^2 \perp \text{span}\{h_1, \dots, h_{n_p}\}$. Our argument proceeds by executing four steps:

1. We begin by showing that $\theta^* \in [0, \pi]$ and $\sqrt{1 - (MC)^{-2/n}} \leq \rho(\theta) \leq 1$ for $\theta \in [\theta^*, \pi]$.
2. We introduce the random variable Θ defined by $\Theta = \angle(\mathbf{x}, \tilde{\mathbf{c}}^1/R)$ and the events

$$\text{for } i = 1, \dots, M, \mathcal{E}_i^* \text{ is the following event : } \left\langle \frac{\mathbf{x}'}{\|\mathbf{x}'\|}, \mathbf{x}_i \right\rangle \leq \gamma \|\mathbf{x}'\|, \text{ with } \mathbf{x}_i = \tilde{\Phi}(\mathbf{u}_i)/R. \quad (\text{A.16})$$

and

$$\text{for } i = 1, \dots, M, \mathcal{E}_i(\Theta) \text{ is the following event : } \frac{\tilde{\Phi}(\mathbf{u}_i)}{R} \in \Omega(\mathbf{x}', \rho(\Theta)), \quad i = 1, \dots, M,$$

with, for $\theta \in [\theta^*, \pi]$,

$$\Omega(\mathbf{z}, \rho(\theta)) := \{\mathbf{x} \in \mathbb{B}_n(\tilde{\mathbf{c}}/R, 1/2) \mid \langle h(\mathbf{z}), \mathbf{x} - \tilde{\mathbf{c}}/R \rangle \leq \rho(\theta)/2\}, \quad h(\mathbf{z}) := \mathbf{z}/\|\mathbf{z}\|_2.$$

In the second step, we condition on $\Theta = \theta$ with $\theta \geq \theta^*$ and show that if $\mathcal{E}_1(\Theta) \wedge \dots \wedge \mathcal{E}_M(\Theta) \mid \Theta = \theta$ occurs then so too does the event $\mathcal{E}_1^* \wedge \dots \wedge \mathcal{E}_M^* \mid \Theta = \theta$.

3. The third step consists of producing a lower bound on the $P(\mathcal{E}_1(\Theta) \wedge \dots \wedge \mathcal{E}_M(\Theta) \mid \Theta = \theta)$ for $\theta \in [\theta^*, \pi]$. This is done by relating the event $\mathcal{E}_i \mid \Theta = \theta$ to Assumption 3.
4. Finally, we use the previous steps to prove that the event $\mathcal{E}_1^* \wedge \dots \wedge \mathcal{E}_M^*$ happens with at least the desired probability. Then, under the assumption that $\mathcal{E}_1^* \wedge \dots \wedge \mathcal{E}_M^*$ occurs, a simple argument similar to the one presented in the final two paragraphs of Theorem 2 allows us to conclude the argument.

Step 1: Proving that $\theta^ \in [0, \pi]$ and $\sqrt{1 - (MC)^{-2/n}} \leq \rho(\theta) \leq 1$ for $\theta \in [\theta^*, \pi]$.*

We start with the claim that $\theta^* \in [0, \pi]$. Since $\alpha \in [0, 1]$, we must have $\arccos(\sqrt{1 - \alpha^2}) \in [0, \pi/2]$. Moreover, by the assumption $2\delta\gamma(1 - \alpha) \geq \sqrt{1 - (MC)^{-2/n}}$, we must also have

$$\arccos\left(\min\left\{1, 2\delta\gamma(1 - \alpha) - \sqrt{1 - (MC)^{-2/n}}\right\}\right) \in [0, \pi/2].$$

The result follows.

Next, we show that $\sqrt{1 - (MC)^{-2/n}} \leq \rho(\theta) \leq 1$ for $\theta \in [\theta^*, \pi]$. The upper bound is obvious from the definition of $\rho(\theta)$. For the lower bound, the function $\theta \mapsto -\max\{0, \cos(\theta - \arccos(\sqrt{1 - \alpha^2}))^{1/2}\}$ is an increasing function of θ for $\theta \in [\arccos(\sqrt{1 - \alpha^2}), \pi]$ and hence $\rho(\theta)$ is increasing on $[\theta^*, \pi]$.

Therefore it suffices to prove that $\rho(\theta^*) \geq \sqrt{1 - (MC)^{-2/n}}$. To do this, note that

$$\begin{aligned}
 \rho(\theta^*) &= \min \left\{ 1, 2\delta\gamma(1 - \alpha) - \max \left\{ 0, \min \left\{ 1, 2\delta\gamma(1 - \alpha) - \sqrt{1 - (MC)^{-2/n}} \right\} \right\} \right\} \\
 &= \min \left\{ 1, 2\delta\gamma(1 - \alpha) - \min \left\{ 1, 2\delta\gamma(1 - \alpha) - \sqrt{1 - (MC)^{-2/n}} \right\} \right\} \\
 &= \min \left\{ 1, \max \left\{ 2\delta\gamma(1 - \alpha) - 1, 2\delta\gamma(1 - \alpha) - 2\delta\gamma(1 - \alpha) + \sqrt{1 - (MC)^{-2/n}} \right\} \right\} \\
 &= \min \left\{ 1, \max \left\{ 2\delta\gamma(1 - \alpha) - 1, \sqrt{1 - (MC)^{-2/n}} \right\} \right\} \\
 &\geq \min \left\{ 1, \sqrt{1 - (MC)^{-2/n}} \right\} = \sqrt{1 - (MC)^{-2/n}},
 \end{aligned}$$

where we have used the assumption $2\delta\gamma(1 - \alpha) \geq \sqrt{1 - (MC)^{-2/n}}$ in the second equality, we have used that $\lambda_1 - \min\{\lambda_2, \lambda_3\} = \max\{\lambda_1 - \lambda_2, \lambda_1 - \lambda_3\}$ for real numbers $\lambda_1, \lambda_2, \lambda_3$ in the third equality and in the inequality we use the fact that both min and max are increasing functions of their arguments.

Step 2: Showing that if the event $\mathcal{E}_1(\Theta) \wedge \dots \wedge \mathcal{E}_M(\Theta) | \Theta = \theta$ occurs then the event $\mathcal{E}_1^(\Theta) \wedge \dots \wedge \mathcal{E}_M^*(\Theta) | \Theta = \theta$ also occurs for $\theta \in [\theta^*, \pi]$.*

First, suppose that the event $\mathcal{E}_1(\Theta) \wedge \dots \wedge \mathcal{E}_M(\Theta) | \Theta = \theta$ occurs. Then for $i \in \{1, 2, \dots, M\}$, $\frac{\tilde{\Phi}(u_i)}{R} \in \Omega(x', \rho(\theta))$ so that we must have $\langle h(x'), \tilde{\Phi}(u_i)/R - \tilde{c}/R \rangle \leq \rho(\theta)/2$. Hence

$$\langle h(x'), \tilde{\Phi}(u_i)/R \rangle = \langle h(x'), \tilde{c}/R \rangle + \langle h(x'), \tilde{\Phi}(u_i)/R - \tilde{c}/R \rangle \leq \langle h(x'), \tilde{c}/R \rangle + \rho(\theta)/2. \quad (\text{A.17})$$

Note that $x' \in \text{span}\{h_1, \dots, h_{n_p}\}$ and $\tilde{c}^2 \perp \text{span}\{h_1, \dots, h_{n_p}\}$ so that $\langle h(x'), \tilde{c}/R \rangle = \langle h(x'), \tilde{c}^1/R \rangle$. The same reasoning tells us that $\|\tilde{c}^1/R\| \leq \|\tilde{c}/R\| \leq 1/2$ (where the final inequality uses that $\tilde{c} = c - \Phi(u^*)$ and Assumption 3, (3.12)) and hence

$$\langle h(x'), \tilde{c}/R \rangle = \|h(x')\| \|\tilde{c}^1/R\| \cos(\angle(h(x'), \tilde{c}^1/R)) \leq \max\{0, \cos(\angle(h(x'), \tilde{c}^1/R))\}/2. \quad (\text{A.18})$$

Combining (A.17) and (A.18) gives us

$$\langle h(x'), \tilde{\Phi}(u_i)/R \rangle \leq \frac{\max\{0, \cos(\angle(h(x'), \tilde{c}^1/R))\} + \rho(\theta)}{2}. \quad (\text{A.19})$$

Next, recall that (using an argument similar to that in Fig. A10 and equations (A.4), (A.5)) $\cos(\angle(h(x'), \tilde{c}^1/R)) \leq \cos(\Theta - \arccos((1 - \alpha^2)^{1/2}))$. Noticing that $\|x'\| \neq 0$ (since $\alpha \in [0, 1)$)

and using (A.19) we see that

$$\begin{aligned}
 \left\langle \frac{\tilde{\Phi}(\mathbf{u}_i)}{R}, \frac{\mathbf{x}'}{\|\mathbf{x}'\|} \right\rangle &\leq \frac{1}{2}\rho(\Theta) + \frac{1}{2}\max\{0, \cos(\angle(h(\mathbf{x}'), \tilde{\mathbf{c}}^1/R))\} \\
 &\leq \frac{1}{2}\rho(\Theta) + \frac{1}{2}\max\{0, \cos(\Theta - \arccos((1 - \alpha^2)^{1/2}))\} \\
 &\leq \delta\gamma(1 - \alpha) - \frac{1}{2}\max\{0, \cos(\Theta - \arccos((1 - \alpha^2)^{1/2}))\} \\
 &\quad + \frac{1}{2}\max\{0, \cos(\Theta - \arccos((1 - \alpha^2)^{1/2}))\} \\
 &= \delta\gamma(1 - \alpha) \leq \gamma\|\mathbf{x}'\|,
 \end{aligned}$$

where the final inequality follows because $\|\mathbf{x}' - \mathbf{x}\| \leq \alpha\delta$ and $\|\mathbf{x}\| = \delta$. Thus $\mathcal{E}_1^* \wedge \dots \wedge \mathcal{E}_M^* | \Theta = \theta$ also occurs and so we conclude that

$$\mathcal{E}_1(\Theta) \wedge \dots \wedge \mathcal{E}_M(\Theta) | \Theta = \theta \implies \mathcal{E}_1^* \wedge \dots \wedge \mathcal{E}_M^* | \Theta = \theta \text{ whenever } \theta \in [\theta^*, \pi] \quad (\text{A.20})$$

Step 3: Bounding $P(\mathcal{E}_1(\Theta) \wedge \dots \wedge \mathcal{E}_M(\Theta) | \Theta = \theta)$ from below.

Assumption 3, equation (3.12) implies that all $\Phi(\mathbf{u}_i)$ belong to the ball of radius $R/2$ centred at \mathbf{c} . Hence according to the equivalence

$$\|\Phi(\mathbf{u}_i) - \mathbf{c}\| \leq \frac{R}{2} \Leftrightarrow \left\| \frac{\tilde{\Phi}(\mathbf{u}_i)}{R} - \frac{\tilde{\mathbf{c}}}{R} \right\| = \left\| \frac{\Phi(\mathbf{u}_i) - \Phi(\mathbf{u}^*)}{R} - \frac{\mathbf{c} - \Phi(\mathbf{u}^*)}{R} \right\| \leq \frac{1}{2},$$

we conclude that all $\tilde{\Phi}(\mathbf{u}_i)/R$ belong to the ball $\mathbb{B}_n(\tilde{\mathbf{c}}/R, 1/2)$ of radius $1/2$ centred at $\tilde{\mathbf{c}}/R$ where we recall that $\tilde{\mathbf{c}} = \mathbf{c} - \Phi(\mathbf{u}^*)$.

Thus for any $i = 1, \dots, M$ the probability $P(\text{not } \mathcal{E}_i | \Theta = \theta)$, $\theta \in [\theta^*, \pi]$ is the probability of the random point $\tilde{\Phi}(\mathbf{u}_i)/R$ landing in the spherical cap

$$C^* = \left\{ \mathbf{y} \in \mathbb{B}_n(\tilde{\mathbf{c}}/R, 1/2) \mid \left\langle \frac{\mathbf{x}'}{\|\mathbf{x}'\|}, \mathbf{y} - \tilde{\mathbf{c}}/R \right\rangle \geq \frac{\rho(\theta)}{2} \right\}.$$

Any point \mathbf{y} in the set C^* satisfies

$$\begin{aligned}
 \left\| \frac{\tilde{\mathbf{c}}}{R} + \frac{\mathbf{x}'}{2\|\mathbf{x}'\|}\rho(\theta) - \mathbf{y} \right\|^2 &\leq \left\| \mathbf{y} - \frac{\tilde{\mathbf{c}}}{R} \right\|^2 - \rho(\theta) \left\langle \frac{\mathbf{x}'}{\|\mathbf{x}'\|}, \mathbf{y} - \frac{\tilde{\mathbf{c}}}{R} \right\rangle + \frac{\rho(\theta)^2}{4} \\
 &\leq \left\| \mathbf{y} - \frac{\tilde{\mathbf{c}}}{R} \right\|^2 - \frac{\rho(\theta)^2}{4} \leq \frac{1 - \rho(\theta)^2}{4}.
 \end{aligned}$$

In particular, if $\tilde{\Phi}(\mathbf{u}_i)/R \in C^*$ and we define $\xi = \mathbf{c} + Rh(\mathbf{x}')\rho(\theta)/2$ and $r = R\sqrt{1 - \rho(\theta)^2}/2$ (where r is real since $\rho(\theta) \in [0, 1/2]$) then

$$\begin{aligned}\|\xi - \Phi(\mathbf{u}_i)\|^2 &= \|\tilde{\mathbf{c}} + Rh(\mathbf{x}')\rho(\theta)/2 - \tilde{\Phi}(\mathbf{u}_i)\|^2 \\ &= R^2 \left\| \frac{\tilde{\mathbf{c}}}{R} + \frac{\mathbf{x}'}{2\|\mathbf{x}'\|}\rho(\theta) - \frac{\tilde{\Phi}(\mathbf{u}_i)}{R} \right\|^2 \leq \frac{R^2(1 - \rho(\theta)^2)}{4} = r^2,\end{aligned}$$

so that

$$P(\text{not } \mathcal{E}_i | \Theta = \theta) = P(\Phi(\tilde{\mathbf{u}}_i)/R \in C^* | \Theta = \theta) \leq P(\Phi(\mathbf{u}_i) \in \mathbb{B}_n(\xi, r) | \Theta = \theta).$$

Next, note that $\|\xi - \mathbf{c}\| = R|\rho(\theta)|/2 \leq R/2$ since $\rho(\theta) \in [0, 1]$ for $\theta \in [\theta^*, \pi]$. Also, $r \leq R/2$. Moreover, for any given value of θ the values ξ and r are deterministic and Θ is independent of $\Phi(\mathbf{u}_i)$ since \mathbf{x} was chosen in the algorithm without knowledge of \mathbf{u}_i . Therefore we can apply Assumption 3, equation (3.13) to obtain

$$P(\text{not } \mathcal{E}_i | \Theta = \theta) \leq C \left(2 \left(\frac{1 - \rho(\theta)^2}{4} \right)^{\frac{1}{2}} \right)^n = C \left(1 - \rho(\theta)^2 \right)^{\frac{n}{2}} \text{ for all } i = 1, \dots, M.$$

Therefore,

$$P(\mathcal{E}_1 \wedge \dots \wedge \mathcal{E}_M | \Theta = \theta) \geq 1 - \sum_{i=1}^M P(\text{not } \mathcal{E}_i | \Theta = \theta) \geq 1 - MC \left(1 - \rho(\theta)^2 \right)^{\frac{n}{2}} \quad (\text{A.21})$$

Step 4: Concluding the argument

Given that $\mathbf{x} = \sum_{i=1}^{n_p} \mathbf{v}_i h_i$ with \mathbf{v} drawn from the equidistribution in $\mathbb{S}_{n_p-1}(0, \delta)$, the variable Θ admits the probability density f_Θ with

$$f_\Theta(\theta) = \frac{1}{\pi^{\frac{1}{2}}} \frac{\Gamma\left(\frac{n_p}{2}\right)}{\Gamma\left(\frac{n_p-1}{2}\right)} \sin^{n_p-2}(\theta), \quad \theta \in [0, \pi]. \quad (\text{A.22})$$

Therefore,

$$\begin{aligned}P(\mathcal{E}_1^* \wedge \dots \wedge \mathcal{E}_M^*) &= \int_0^\pi P(\mathcal{E}_1^* \wedge \dots \wedge \mathcal{E}_M^* | \Theta = \theta) f_\Theta(\theta) d\theta \geq \int_{\theta^*}^\pi P(\mathcal{E}_1^* \wedge \dots \wedge \mathcal{E}_M^* | \Theta = \theta) f_\Theta(\theta) d\theta \\ &\geq \int_{\theta^*}^\pi P(\mathcal{E}_1(\Theta) \wedge \dots \wedge \mathcal{E}_M(\Theta) | \Theta = \theta) f_\Theta(\theta) d\theta,\end{aligned}$$

where the final inequality follows from the (A.20) proven in step 2.

In step 1 we showed that $\rho(\theta) \geq \sqrt{1 - (MC)^{-2/n}}$ for $\theta \in [\theta^*, \pi]$. Rearranging yields $1 - MC(1 - \rho(\theta)^2)^{n/2} \geq 0$. Combining this with the inequality (A.21) we obtain

$$\begin{aligned} P(\mathcal{E}_1^* \wedge \dots \wedge \mathcal{E}_M^*) &\geq \int_{\theta^*}^{\pi} (1 - MC(1 - \rho(\theta)^2)^{\frac{n}{2}}) f_{\Theta}(\theta) d\theta \\ &= \int_{\theta^*}^{\pi} f_{\Theta}(\theta) d\theta - MC \int_{\theta^*}^{\pi} (1 - \rho(\theta)^2)^{\frac{n}{2}} f_{\Theta}(\theta) d\theta \\ &= 1 - MC \int_{\theta^*}^{\pi} (1 - \rho(\theta)^2)^{\frac{n}{2}} f_{\Theta}(\theta) d\theta - \int_0^{\theta^*} f_{\Theta}(\theta) d\theta. \end{aligned}$$

Substituting (A.22) gives

$$\begin{aligned} P(\mathcal{E}_1^* \wedge \dots \wedge \mathcal{E}_M^*) &\geq 1 - MC \frac{1}{\pi^{\frac{1}{2}}} \frac{\Gamma(\frac{n_p}{2})}{\Gamma(\frac{n_p-1}{2})} \int_{\theta^*}^{\pi} (1 - \rho(\theta)^2)^{\frac{n}{2}} \sin^{n_p-2}(\theta) d\theta \\ &\quad - \frac{1}{\pi^{\frac{1}{2}}} \frac{\Gamma(\frac{n_p}{2})}{\Gamma(\frac{n_p-1}{2})} \int_0^{\theta^*} \sin^{n_p-2}(\theta) d\theta. \end{aligned}$$

Finally, assuming that the event $\mathcal{E}_1^* \wedge \dots \wedge \mathcal{E}_M^*$ occurs, the final two paragraphs of the argument presented in the proof of Theorem 2 (noting that Assumption 3 implies 2) may then be used to confirm that the values of w and b specified in Algorithm 3 produce an ε - Δ stealth attack. The conclusion of Theorem 3 follows. \square

A.4 Finding triggers \mathbf{u}' in Algorithms 1 and 2

A relevant optimization problem for Algorithm 1 is formulated in (3.3). Similarly to (3.3), one can pose a constrained optimization problem for finding a \mathbf{u}' in step 3 of Algorithm 2. This problem can be formulated as follows:

$$\begin{aligned} \mathbf{u}' &= \arg \min_{\mathbf{u} \in \mathcal{U}} \left\| \frac{\Phi(\mathbf{u})}{R} - \frac{\Phi(\mathbf{u}^*)}{R} - \mathbf{x} \right\| \\ \text{s.t.} \quad & \gamma \left\| \frac{\Phi(\mathbf{u})}{R} - \frac{\Phi(\mathbf{u}^*)}{R} \right\| \leq 1, \quad \left\| \frac{\Phi(\mathbf{u})}{R} - \frac{\Phi(\mathbf{u}^*)}{R} - \mathbf{x} \right\| < \delta. \end{aligned} \quad (\text{A.23})$$

Note that the vector \mathbf{x} must be chosen randomly on $\mathcal{S}_{n-1}(0, \delta)$, $\delta \in (0, 1]$. One way to achieve this is to generate a sample \mathbf{z} from an n -dimensional normal distribution $\mathcal{N}(0, I_n)$ and then set $\mathbf{x} = \delta \mathbf{z} / \|\mathbf{z}\|$.

A practical approach to determine \mathbf{u}' is to employ a gradient-based search

$$\mathbf{u}'_{k+1} = \text{Proj}_{\mathcal{U}} \left[\mathbf{u}'_k - h_k \frac{\partial}{\partial \mathbf{u}} \mathcal{L}(\mathbf{u}, \mathbf{u}^*) \right], \quad \mathbf{u}'_0 = \mathbf{u}^*,$$

where $\text{Proj}_{\mathcal{U}}$ is a projection operator. In our experiments, $\text{Proj}_{\mathcal{U}}(\mathbf{u})$ returned $\mathbf{u} = (u_1, \dots, u_n)$ if $u_i \in [0, 255]$. If, however, the i -th component of \mathbf{u} is out of range ($u_i < 0$ or $u_i > 255$) then the operator returned a vector with 0 or 255 in that corresponding component.

The loss function $\mathcal{L}(\mathbf{u}, \mathbf{u}^*)$ is defined as

$$\mathcal{L}(\mathbf{u}, \mathbf{u}^*) = \left\| \frac{\Phi(\mathbf{u})}{R} - \frac{\Phi(\mathbf{u}^*)}{R} - \mathbf{x} \right\|^2 + \lambda_1 \mathcal{G}_1(\mathbf{u}) + \lambda_2 \mathcal{G}_2(\mathbf{u}), \quad \lambda_1, \lambda_2 \geq 0,$$

and $\mathcal{G}_1, \mathcal{G}_2$ are relevant penalty functions:

$$\mathcal{G}_1(\mathbf{u}) = \begin{cases} \left(\gamma \left\| \frac{\Phi(\mathbf{u})}{R} - \frac{\Phi(\mathbf{u}^*)}{R} \right\| - 1 \right)^{p_1}, & \gamma \left\| \frac{\Phi(\mathbf{u})}{R} - \frac{\Phi(\mathbf{u}^*)}{R} \right\| \geq 1 \\ 0, & \text{otherwise} \end{cases}$$

$$\mathcal{G}_2(\mathbf{u}) = \begin{cases} \left(\left\| \frac{\Phi(\mathbf{u})}{R} - \frac{\Phi(\mathbf{u}^*)}{R} - \mathbf{x} \right\| - \delta \right)^{p_2}, & \left\| \frac{\Phi(\mathbf{u})}{R} - \frac{\Phi(\mathbf{u}^*)}{R} - \mathbf{x} \right\| \geq \delta \\ 0, & \text{otherwise} \end{cases}$$

with parameters $p_1, p_2 > 0$, and with $h_k > 0$ a sequence of parameters ensuring convergence of the procedure.

For Algorithm 1, terms $\Phi(\mathbf{u}^*)/R$ should be replaced with 0, and \mathbf{u}_0 can be set to an arbitrary element of \mathcal{U} .

REMARK 11. (Trigger search subspace). Sometimes we wish to transform the original feature map Φ into $\tilde{\Phi} = T\Phi$ (see Remark 2) in order to comply with the input-reachability assumption. This translates into sampling \mathbf{x} in step 2 of Algorithm 2 from $S_{m-1}(0, \delta) \subset \mathbb{R}^n$, $2 \leq m < n$ instead of $S_{n-1}(0, \delta)$. On the one hand this may have a potentially negative affect on the probability of success as n would need to be replaced with $2 \leq m < n$ in (3.8) and (3.9). On the other hand this extra ‘pre-processing’ may offer computational benefits. Indeed, if $\Phi(\mathbf{u}^*)$ is an output of a ReLU layer then it is likely that some of its components are exactly zero. This would severely constrain gradient-based routines to determine \mathbf{u}' starting from \mathbf{u}^* if \mathbf{x} is sampled from $S_{n-1}(0, \delta)$, as the corresponding gradients will always be equal to zero. In addition, one can constrain the triggers’ search space to smaller dimensional subspaces, as specified in Algorithm 3. Not only may this approach relax input-reachability restrictions and help with computing triggers but it may also be used to mitigate various feasibility issues around accuracy of the attacks (see Theorem 3 and Tables 1 and 2).

In our code (Tyukin *et al.*, 2022) illustrating the application of Algorithms 1, 2 and 3 we follow the above logic and search for input triggers in subspaces of the original feature space corresponding to non-zero attributes of the feature vector $\Phi(\mathbf{u}^*)$.

A.5 Accuracy of implementation of stealth attacks achieved in experiments and limits of theoretical bounds in Theorems 1 and 2

In all 20 attack experiments we were able to create a ReLU neuron which was completely silent on the set \mathcal{V} and, at the same time, was producing the desired responses when a trigger was presented as an input to the network. Empirical values of the accuracy of implementation of these attacks, α and the dimension n of the random perturbation are shown in Table A4.

TABLE A.4 Accuracy α of finding triggers expressed as $\alpha = \delta^{-1} \|\Phi(\mathbf{u}')/R - \mathbf{x}\|$, $\delta = 1/3$.

#	1	2	3	4	5	6	7	8	9	10
α	0.179	0.200	0.358	0.247	0.313	0.406	0.311	0.277	0.362	0.244
n	112	125	122	116	125	122	115	121	128	111
#	11	12	13	14	15	16	17	18	19	20
α	0.417	0.271	0.271	0.227	0.246	0.287	0.401	0.327	0.436	0.285
n	122	124	124	119	110	117	120	109	110	117

TABLE A.5 Accuracy α of finding triggers expressed as $\alpha = \delta^{-1} \|\Phi(\mathbf{u}')/R - \mathbf{x}\|$ for the lower dimensional feature space, $\delta = 1/3$.

#	1	2	3	4	5	6	7	8	9	10
α	0.028	0.004	0.092	4.3×10^{-4}	0.061	0.026	0.040	6.6×10^{-5}	2.7×10^{-4}	0.117
n	39	38	35	32	32	35	38	33	37	33
#	11	12	13	14	15	16	17	18	19	20
α	0.021	0.003	0.017	0.109	0.003	0.001	2.5×10^{-4}	9.2×10^{-5}	0.022	0.011
n	35	36	32	38	35	34	34	32	35	33

TABLE A.6 Accuracy α of finding triggers expressed as $\alpha = \delta^{-1} \|\Phi(\mathbf{u}')/R - \mathbf{x}\|$ for the lower dimensional feature space, $\delta = 2/3$.

#	1	2	3	4	5	6	7	8	9	10
α	0.090	0.205	0.229	0.032	0.097	0.046	0.106	0.053	0.193	0.139
n	33	39	34	35	37	36	34	36	36	35
#	11	12	13	14	15	16	17	18	19	20
α	0.063	0.001	0.051	0.305	0.085	0.084	0.065	0.169	0.121	0.140
n	31	30	37	32	33	33	36	34	38	33

Empirical values of the accuracy of implementation of these attacks, α , for lower dimensional feature maps produced by mappings (4.1) are shown in Tables A5 and A6.

These figures, along with the fact that all these attacks returned a successful outcome (for Scenario 1), enable us to illustrate how conservative the bounds provided in Theorem 1 could be.

Indeed, computing the term

$$P_1(\alpha, \delta, \gamma, n) = \pi^{-1/2} \frac{\Gamma\left(\frac{n}{2}\right)}{\Gamma\left(\frac{n-1}{2}\right)} \int_0^{\arccos(\varphi(\gamma, \delta, \alpha))} \sin^{n-2}(\theta) d\theta$$

for $\gamma = 0.9$ in bound (3.8) for α, n taken from the first experiment in Table A4 (column corresponding to # = 1 in Table A4) returns $P_1(\alpha, \delta, \gamma, n) = 0.1561$. Substituting this value into (3.8) for $M > 10$ produces a negative number rendering the bound not useful in this case. The same comment applies to other experiments. At the same time, as our experiments confirm, the attacks turn out to be successful

in practice. If, however, we set $\alpha = 0$ then the value of $P_1(0, \delta, \gamma, n)$ for the first experiment becomes 4.9180×10^{-4} .

These observations highlight limitations of the bounds in Theorems 1 and 2 for determining the probability of success in practice when α is large. They also show the additional relevance of Theorem 3 and the ‘concentrational collapse’ effect for developing better understanding of conditions leading to increased vulnerabilities to stealth attacks (see Tables 1 and 2 illustrating the difference between bounds in Theorems 1, 2 and 3).

A.6 Selection of a neuron to attack

Let L be the layer in which a neuron is to be selected for the attack. We suppose that the network has a layer $L+1$ in its computational graph. Suppose that $w_{ij}^{[L]}$ denotes the j th weight of the i th neuron in layer L , and the total number of neurons in layer L is N_L .

We say that the *output* weight vector of the i -th neuron in layer L is a vector

$$\bar{\mathbf{w}}_i^{[L]} = \begin{pmatrix} w_{1,i}^{[L+1]} \\ w_{2,i}^{[L+1]} \\ \vdots \\ w_{N_L,i}^{[L+1]} \end{pmatrix}.$$

Determining neuron rank. For each $i = 1, \dots, N_L$, we computed L_1 -norms of $\bar{\mathbf{w}}_i^{[L]}$:

$$\|\bar{\mathbf{w}}_i^{[L]}\|_1 = \sum_{j=1}^{N_{L+1}} |w_{j,i}^{L+1}|$$

and created a list of indices n_1, n_2, \dots, n_{N_L} such that

$$\|\bar{\mathbf{w}}_{n_1}^{[L]}\|_1 \leq \|\bar{\mathbf{w}}_{n_2}^{[L]}\|_1 \leq \dots \leq \|\bar{\mathbf{w}}_{n_{N_L}}^{[L]}\|_1.$$

The susceptibility rank of the i -th neuron, $\text{Rank}(i)$, is defined as a number $k \in \{1, \dots, N_L\}$:

$$\text{Rank}(i) = k \text{ such that } i = n_k.$$

In our code, any ties were broken at random.