# Supplementary materials for Input Similarity from the Neural Network Perspective 

Guillaume Charpiat $^{1} \quad$ Nicolas Girard ${ }^{2} \quad$ Loris Felardos ${ }^{1} \quad$ Yuliya Tarabalka ${ }^{2,3}$<br>${ }^{1}$ TAU team, INRIA Saclay, LRI, Univ. Paris-Sud<br>${ }^{2}$ TITANE team, INRIA Sophia-Antipolis, Univ. Côte d'Azur<br>${ }^{3}$ LuxCarta Technology<br>firstname.lastname@inria.fr

## 1 Foreword

Code The whole code (image registration, experiments to test density estimators, enforcing similarity...) is available on the following github repository: https://github.com/Lydorn/ netsimilarity.

Content of this document This document contains further details on the article Input Similarity from the Neural Network Perspective, such as mathematical proofs, experimental details and further discussions.

## 2 Proofs of the properties of the 1D similarity kernel

We give here the proofs at the properties of the 1-dimensional-output similarity kernel.

### 2.1 Proof of Theorem 1

Theorem 1. For any real-valued neural network $f_{\theta}$ whose last layer is a linear layer (without any parameter sharing) or a standard activation function thereof (sigmoid, tanh, ReLU...), and for any inputs $\mathbf{x}$ and $\mathbf{x}^{\prime}$,

$$
\nabla_{\theta} f_{\theta}(\mathbf{x})=\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right) \quad \Longrightarrow \quad f_{\theta}(\mathbf{x})=f_{\theta}\left(\mathbf{x}^{\prime}\right)
$$

Proof. If the last layer is linear, the output is of the form $f_{\theta}(\mathbf{x})=\sum_{i} w_{i} a_{i}(\mathbf{x})+b$, where $w_{i}$ and $b$ are parameters in $\mathbb{R}$ and $a_{i}(\mathbf{x})$ activities from previous layers. The gradient $\nabla_{\theta} f_{\theta}(\mathbf{x})$ contains in particular as coefficients the derivatives $\frac{d f_{\theta}(\mathbf{x})}{d w_{i}}=a_{i}(\mathbf{x})$. Thus $\nabla_{\theta} f_{\theta}(\mathbf{x})=\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right) \Longrightarrow a_{i}(\mathbf{x})=a_{i}\left(\mathbf{x}^{\prime}\right) \forall i$ in the last layer. The outputs can be then rebuilt: $f_{\theta}(\mathbf{x})=\sum_{i} w_{i} a_{i}(\mathbf{x})+b=\sum_{i} w_{i} a_{i}\left(\mathbf{x}^{\prime}\right)+b=f_{\theta}\left(\mathbf{x}^{\prime}\right)$. If the output is of the form $f_{\theta}(\mathbf{x})=\sigma(c(\mathbf{x}))$ with $c(\mathbf{x})=\sum_{i} w_{i} a_{i}(\mathbf{x})+b$, then the gradient equality implies $\frac{d f_{\theta}(\mathbf{x})}{d b}=\frac{d f_{\theta}\left(\mathbf{x}^{\prime}\right)}{d b}$, whose value is $\sigma^{\prime}(c(\mathbf{x}))=\sigma^{\prime}\left(c\left(\mathbf{x}^{\prime}\right)\right)$. Then, as $\sigma^{\prime}(c(\mathbf{x})) a_{i}(\mathbf{x})=\frac{d f_{\theta}(\mathbf{x})}{d w_{i}}=$ $\frac{d f_{\theta}\left(\mathbf{x}^{\prime}\right)}{d w_{i}}=\sigma^{\prime}\left(c\left(\mathbf{x}^{\prime}\right)\right) a_{i}\left(\mathbf{x}^{\prime}\right)$, we can deduce $a_{i}(\mathbf{x})=a_{i}\left(\mathbf{x}^{\prime}\right)$ for all $i$ provided $\sigma^{\prime}(c(\mathbf{x})) \neq 0$. In that case, from these identical activities one can rebuild identical outputs. Otherwise, $\sigma^{\prime}(c(\mathbf{x}))=$ $\sigma^{\prime}\left(c\left(\mathbf{x}^{\prime}\right)\right)=0$, which is not possible with strictly monotonous activation functions, such as tanh or sigmoid. For ReLU, $\sigma^{\prime}(c(\mathbf{x}))=0 \Longrightarrow \sigma(c(\mathbf{x}))=0$ and thus $f_{\theta}(\mathbf{x})=f_{\theta}\left(\mathbf{x}^{\prime}\right)=0$. The same reasoning holds for other activation functions with only one flat piece (such as the ReLU negative part), i.e. for which the set $\sigma\left(\sigma^{\prime-1}(\{0\})\right)$ is a singleton.

### 2.2 Proof of Corollary 1

Corollary 1. Under the same assumptions, for any inputs $\mathbf{x}$ and $\mathbf{x}^{\prime}$,

$$
\begin{aligned}
& k_{\theta}^{C}\left(\mathbf{x} \cdot \mathbf{x}^{\prime}\right)=1 \quad \Longrightarrow \quad \nabla_{\theta} f_{\theta}(\mathbf{x})=\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right), \\
\text { hence } & k_{\theta}^{C}\left(\mathbf{x} \cdot \mathbf{x}^{\prime}\right)=1 \quad \Longrightarrow \quad f_{\theta}(\mathbf{x})=f_{\theta}\left(\mathbf{x}^{\prime}\right) .
\end{aligned}
$$

Proof. $\quad k_{\theta}^{C}\left(\mathbf{x} . \mathbf{x}^{\prime}\right)=1$ means $\frac{\nabla_{\theta} f_{\theta}(\mathbf{x})}{\left\|\nabla_{\theta} f_{\theta}(\mathbf{x})\right\|} \cdot \frac{\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)}{\left\|\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)\right\|}=1$, which implies $\exists \alpha \in \mathbb{R}^{*}, \quad \nabla_{\theta} f_{\theta}(\mathbf{x})=$ $\alpha \nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)$. We need to show that $\alpha=1$. Under the assumptions of Theorem 1 , following its proof:

- either the last layer is linear, the output is of the form $f_{\theta}(\mathbf{x})=\sum_{i} w_{i} a_{i}(\mathbf{x})+b$, and then $\nabla_{b} f_{\theta}(\mathbf{x})=\alpha \nabla_{b} f_{\theta}\left(\mathbf{x}^{\prime}\right)$ while $\frac{d f_{\theta}(\mathbf{x})}{d b}=1$ and $\frac{d f_{\theta}\left(\mathbf{x}^{\prime}\right)}{d b}=1$, hence $\alpha=1$;
- either the output is of the form $f_{\theta}(\mathbf{x})=\sigma(c(\mathbf{x}))$ with $c(\mathbf{x})=\sum_{i} w_{i} a_{i}(\mathbf{x})+b$, and then $\sigma^{\prime}(c(\mathbf{x}))=\nabla_{b} f_{\theta}(\mathbf{x})=\alpha \nabla_{b} f_{\theta}\left(\mathbf{x}^{\prime}\right)=\alpha \sigma^{\prime}\left(c\left(\mathbf{x}^{\prime}\right)\right)$, while, for any $i, \sigma^{\prime}(c(\mathbf{x})) a_{i}(\mathbf{x})=$ $\frac{d f_{\theta}(\mathbf{x})}{d w_{i}}=\alpha \frac{d f_{\theta}\left(\mathbf{x}^{\prime}\right)}{d w_{i}}=\alpha \sigma^{\prime}\left(c\left(\mathbf{x}^{\prime}\right)\right) a_{i}\left(\mathbf{x}^{\prime}\right)$. Thus, supposing $\sigma^{\prime}(c(\mathbf{x})) \neq 0$, we obtain $a_{i}(\mathbf{x})=$ $a_{i}\left(\mathbf{x}^{\prime}\right) \forall i$, and thus we can rebuild from the activities $c(\mathbf{x})=c\left(\mathbf{x}^{\prime}\right)$, from which $\sigma^{\prime}(c(\mathbf{x}))=$ $\sigma^{\prime}\left(c\left(\mathbf{x}^{\prime}\right)\right)$ and thus $\alpha=1$. Otherwise, $\sigma^{\prime}(c(\mathbf{x}))=\sigma^{\prime}\left(c\left(\mathbf{x}^{\prime}\right)\right)=0$ and the two full gradients $\nabla_{\theta} f_{\theta}(\mathbf{x})$ and $\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)$ are 0 and thus equal.

The conditions for $k_{\theta}^{C}\left(\mathbf{x} \cdot \mathbf{x}^{\prime}\right)=1 \Longrightarrow \nabla_{\theta} f_{\theta}(\mathbf{x})=\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)$ to hold are actually much weaker: it is sufficient that in the whole network architecture there exists one useful neuron (in the sense of the next paragraph) of that type (so called linear but actually affine).

### 2.3 Proof of Theorem 2

Theorem 2. For any real-valued neural network $f_{\theta}$ without parameter sharing, if $\nabla_{\theta} f_{\theta}(\mathbf{x})=$ $\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)$ for two inputs $\mathbf{x}, \mathbf{x}^{\prime}$, then all useful activities computed when processing $\mathbf{x}$ are equal to the ones obtained when processing $\mathbf{x}^{\prime}$.

We name useful activities all activities whose variation would have an impact on the output, i.e. all the ones satisfying $\frac{d f_{\theta}(\mathbf{x})}{d a_{i}} \neq 0$. This condition is typically not satisfied when the activity is multiplied by 0 , i.e. $w_{i}=0$, or when it is negative and followed by a ReLU, or when all its contributions to the output annihilate together (e.g., a sum of two neurons with opposite weights: $f_{\theta}(\mathbf{x})=\sigma\left(a_{i}(\mathbf{x})\right)-\sigma\left(a_{i}(\mathbf{x})\right)$ ).

Proof. Let $a_{i}(\mathbf{x})$ be a useful activity (for $\mathbf{x}$ ). It is fed to at least one useful neuron, whose preactivation output is of the form $c(\mathbf{x})=\sum_{i} w_{i} a_{i}(\mathbf{x})+b$. Then $\frac{d f_{\theta}(\mathbf{x})}{d b}=\frac{d f_{\theta}(\mathbf{x})}{d c} \neq 0$ (the output of the neuron is useful), and $\frac{d f_{\theta}(\mathbf{x})}{d w_{i}}=\frac{d f_{\theta}(\mathbf{x})}{d b} a_{i}(\mathbf{x})$. From the gradient equality, $a_{i}(\mathbf{x})=\frac{d f_{\theta}(\mathbf{x})}{d w_{i}} / \frac{d f_{\theta}(\mathbf{x})}{d b}=$ $\frac{d f_{\theta}\left(\mathbf{x}^{\prime}\right)}{d w_{i}} / \frac{d f_{\theta}\left(\mathbf{x}^{\prime}\right)}{d b}=a_{i}\left(\mathbf{x}^{\prime}\right)$.

## 3 Higher output dimension

We expand here all the mathematical aspects of the homonymous section of the article.

### 3.1 Derivation

Let us now study the case where $f_{\theta}(\mathbf{x})$ is a vector in $\mathbb{R}^{d}$ with $d>1$.
The optimal parameter change $\delta \theta$ to push $f_{\theta}(\mathbf{x})$ in a direction $\mathbf{v}$ (with a force $\varepsilon$ ) is less straightforward to obtain. First, one can define as many gradients as output coordinates: $\nabla_{\theta} f_{\theta}^{i}(\mathbf{x})$, for $i \in \llbracket 1, d \rrbracket$.
This family of gradients can be shown to be linearly independent, unless the architecture of the network is specifically built not to. If for instance each output coordinate has its own bias parameter, i.e. writes in the form $f_{\theta}^{i}(\mathbf{x})=b_{i}+g_{\theta}(\mathbf{x})$ or $\sigma\left(b_{i}+g_{\theta}(\mathbf{x})\right)$ with a strictly monotonous activation function $\sigma$, then the derivative w.r.t. $b_{i}$ will be 1 (or $\sigma^{\prime}$ ) only in the $i$-th gradient and 0 in the other ones.

Thus the $j$-th gradient contains in particular the subvector $\left(\frac{d f^{j}}{d b_{i}}\right)_{i}=\left(\delta_{i=j}\right)_{i}$, and the gradients are consequently independent. In the case where all coordinates depend on all biases, but not identically, as with a softmax, the argument stays true.
Any parameter variation $\delta \theta \in \mathbb{R}^{p}$ can then be uniquely decomposed as:

$$
\delta \theta=\sum_{i=1}^{d} \alpha_{i} \nabla f_{\theta}^{i}(\mathbf{x})+\gamma
$$

where $\alpha_{i} \in \mathbb{R}$ and where $\gamma \in \mathbb{R}^{p}$ is orthogonal to all coordinate gradients. This parameter variation induces an output variation:

$$
\begin{gathered}
f_{\theta+\delta \theta}(\mathbf{x})-f_{\theta}(\mathbf{x})=\nabla_{\theta} f_{\theta}(\mathbf{x}) \delta \theta+O\left(\|\delta \theta\|^{2}\right) \\
=\left(\sum_{i} \alpha_{i} \nabla_{\theta} f_{\theta}^{i}(\mathbf{x}) \cdot \nabla f_{\theta}^{j}(\mathbf{x})\right)_{j}+0+O\left(\|\delta \theta\|^{2}\right) \\
=C \alpha+O\left(\|\alpha\|^{2}\right)
\end{gathered}
$$

where $C$ is the correlation matrix of the gradients: $C_{i j}=\nabla_{\theta} f_{\theta}^{i}(\mathbf{x}) \cdot \nabla f_{\theta}^{j}(\mathbf{x})$. It turns out that $C$ is invertible:

$$
\begin{aligned}
C \alpha & =0 \Longrightarrow \alpha C \alpha=0 \Longrightarrow \alpha \nabla_{\theta} f_{\theta}(\mathbf{x}) \nabla_{\theta} f_{\theta}(\mathbf{x}) \alpha=0 \\
& \Longrightarrow\left\|\nabla_{\theta} f_{\theta}(\mathbf{x}) \alpha\right\|^{2}=0 \Longrightarrow \sum_{i} \alpha_{i} \nabla f_{\theta}^{i}(\mathbf{x})=0
\end{aligned}
$$

$\Longrightarrow \alpha=0$ as the $\nabla_{\theta} f_{\theta}^{i}(\mathbf{x})$ are linearly independent. Thus, for a desired output move in the direction $\mathbf{v}$ with amplitude $\varepsilon$, i.e. $f_{\theta+\delta \theta}(\mathbf{x})-f_{\theta}(\mathbf{x})=\varepsilon \mathbf{v}$, one can compute the associated linear combination $\alpha=\varepsilon C^{-1} \mathbf{v}$ and thus the smallest associated parameter change $\delta \theta=\sum_{i} \alpha_{i} \nabla f_{\theta}^{i}(\mathbf{x})$.
The output variation induced at any other point $\mathbf{x}^{\prime}$ by this parameter change is then:

$$
\begin{gather*}
f_{\theta+\delta \theta}\left(\mathbf{x}^{\prime}\right)-f_{\theta}\left(\mathbf{x}^{\prime}\right)=\left(\nabla_{\theta} f_{\theta}^{i}\left(\mathbf{x}^{\prime}\right) \cdot \delta \theta\right)_{i}+O\left(\|\delta \theta\|^{2}\right) \\
=\left(\sum_{j} \alpha_{j} \nabla_{\theta} f_{\theta}^{i}\left(\mathbf{x}^{\prime}\right) \cdot \nabla_{\theta} f_{\theta}^{j}(\mathbf{x})\right)_{i}+O\left(\|\delta \theta\|^{2}\right) \\
=\varepsilon K_{\theta}\left(\mathbf{x}^{\prime}, \mathbf{x}\right) C_{\theta}(\mathbf{x})^{-1} \mathbf{v}+O\left(\varepsilon^{2}\right) \tag{1}
\end{gather*}
$$

where the $d \times d$ kernel matrix $K_{\theta}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)$ is defined by $K_{\theta}^{i j}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)=\nabla_{\theta} f_{\theta}^{i}(\mathbf{x}) \cdot \nabla_{\theta} f_{\theta}^{j}\left(\mathbf{x}^{\prime}\right)$, and where the matrix $C_{\theta}(\mathbf{x})=K_{\theta}(\mathbf{x}, \mathbf{x})$ is the previously defined self-correlation matrix $C$. Its role is equivalent of the normalization by $\left\|\nabla_{\theta} f_{\theta}(\mathbf{x})\right\|^{2}$ in the 1D case, in plus of decorrelating the gradients.
The interpretation of (1) is that if one moves the output for point $\mathbf{x}$ by $\mathbf{v}$, then the output for point $\mathbf{x}^{\prime}$ will be moved also, by $M \mathbf{v}$, with $M=K_{\theta}\left(\mathbf{x}, \mathbf{x}^{\prime}\right) K_{\theta}(\mathbf{x}, \mathbf{x})^{-1}$. Note that these matrices $M$ or $K$ are only $d \times d$ where $d$ is the output dimension. They are thus generally small and easy to manipulate or inverse.

### 3.2 Normalized cross-correlation matrix

The normalized version of the kernel (1) is:

$$
\begin{equation*}
K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)=C_{\theta}(\mathbf{x})^{-1 / 2} K_{\theta}\left(\mathbf{x}, \mathbf{x}^{\prime}\right) C_{\theta}\left(\mathbf{x}^{\prime}\right)^{-1 / 2} \tag{2}
\end{equation*}
$$

which is symmetric in the sense that $K_{\theta}^{C}\left(\mathbf{x}^{\prime}, \mathbf{x}\right)=K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)^{T}$.
A matrix $K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)$ with small coefficients means that $\mathbf{x}$ and $\mathbf{x}^{\prime}$ are relatively independent, from a neural network point of view (moves at $\mathbf{x}$ won't be transferred to $\mathbf{x}^{\prime}$ ). On the opposite, the highest possible dependency is $K_{\theta}^{C}(\mathbf{x}, \mathbf{x})=\mathrm{Id}$.

To study properties of this similarity measure, note that $K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)=\left(G_{\mathbf{x}}^{N}\right)^{T} G_{\mathbf{x}^{\prime}}^{N}$ with $G_{\mathbf{x}}^{N}=$ $G_{\mathbf{x}}\left(G_{\mathbf{x}}^{T} G_{\mathbf{x}}\right)^{-1 / 2}$, where $G_{\mathbf{x}}=\nabla_{\theta} f(\mathbf{x})$ : it is the product of normalized, decorrelated versions of the gradient. Indeed, at any point $\mathbf{x}$, the normalized gradient matrix $G_{\mathbf{x}}^{N}$ satisfies: $\left(G_{\mathbf{x}}^{N}\right)^{T} G_{\mathbf{x}}^{N}=$
$K_{\theta}^{C}(\mathbf{x}, \mathbf{x})=K_{\theta}(\mathbf{x}, \mathbf{x})^{-1 / 2} K_{\theta}(\mathbf{x}, \mathbf{x}) K_{\theta}(\mathbf{x}, \mathbf{x})^{-1 / 2}=\operatorname{Id}$ and consequently $G_{\mathbf{x}}^{N}$ can be seen as an orthonormal family of vectors $G_{\mathbf{x}}^{N, i}$.
The $L^{2}$ (Frobenius) norm of the ortho-normalized gradient $G_{\mathrm{x}}^{N}$ is thus:

$$
\left\|G_{\mathbf{x}}^{N}\right\|_{F}^{2}=\operatorname{Tr}\left(\left(G_{\mathbf{x}}^{N}\right)^{T} G_{\mathbf{x}}^{N}\right)=\operatorname{Tr}(\mathrm{Id})=d
$$

At point $\mathbf{x}^{\prime}, G_{\mathbf{x}^{\prime}}^{N}$ is also an orthonormal family, but possibly arranged differently or generating a different subspace of $\mathbb{R}^{p}$. If $G_{\mathbf{x}}^{N}$ and $G_{\mathbf{x}^{\prime}}^{N}$ generate the same subspace, then their product $\left(G_{\mathbf{x}}^{N}\right)^{T} G_{\mathbf{x}^{\prime}}^{N}$ is an orthogonal matrix $Q$ (change of basis) and its $L^{2}$ (Frobenius) norm is then $\|Q\|_{F}^{2}=\operatorname{Tr}\left(Q^{T} Q\right)=$ $\operatorname{Tr}(\mathrm{Id})=d$. Otherwise, $\left(G_{\mathbf{x}}^{N}\right)^{T} G_{\mathbf{x}^{\prime}}^{N}$ can be seen as a projection from one subspace to another one, each vector $G_{\mathbf{x}^{\prime}}^{N, j}$ is projected onto the ortho-normal family $\left(G_{\mathbf{x}}^{N, i}\right)_{i}$, and as a projection decreases the Euclidean norm, $\sum_{i}\left(G_{\mathbf{x}}^{N, i} \cdot G_{\mathbf{x}^{\prime}}^{N, j}\right)^{2} \leqslant\left\|G_{\mathbf{x}^{\prime}}^{N, j}\right\|^{2}=1$. Thus:

$$
\left\|K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)\right\|_{F}=\sqrt{\sum_{i j}\left(G_{\mathbf{x}}^{N, i} \cdot G_{\mathbf{x}^{\prime}}^{N, j}\right)^{2}} \leqslant \sqrt{d}
$$

Moreover, any coefficient of the kernel matrix satisfies:

$$
\left|K_{\theta}^{C, i j}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)\right|=\left|G_{\mathbf{x}}^{N, i} \cdot G_{\mathbf{x}^{\prime}}^{N, j}\right| \leqslant\left\|G_{\mathbf{x}}^{N, i}\right\|_{2}\left\|G_{\mathbf{x}^{\prime}}^{N, j}\right\|_{2}=1
$$

as each vector $G_{\mathrm{x}}^{N, i}$ is unit-norm. This implies in particular that the trace is bounded:

$$
-d \leqslant \operatorname{Tr}\left(K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)\right) \leqslant d
$$

To sum up, the similarity matrix $K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)$ satisfies the following properties:

- its coefficients are bounded, in $[-1,1]$
- its trace is at most $d$
- its (Frobenius) norm is at most $\sqrt{d}$
- self-similarity is identity: $\forall \mathbf{x}, K_{\theta}^{C}(\mathbf{x}, \mathbf{x})=\mathrm{Id}$
- the kernel is symmetric, in the sense that $K_{\theta}^{C}\left(\mathbf{x}^{\prime}, \mathbf{x}\right)=K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)^{T}$.


### 3.3 Similarity in a single value

Note that when the trace is close to its maximal value $d$, the diagonal coefficients are close to 1 , and their contribution to the Frobenius norm squared is close to $d$. Therefore, all non-diagonal coefficients are close to 0 , and the matrix is close to Id. And reciprocally, a matrix close to Id has a trace close to $d$. Thus, two related ways to quantify similarity in a single real value in $[-1,1]$ appear:

- the distance to the identity $D=\left\|K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)-\mathrm{Id}\right\|_{F}$, which can be turned into a similarity as $1-\frac{1}{\sqrt{d}} D$ or $1-\frac{1}{2 d} D^{2}$, since $D \in[0,2 \sqrt{d}]$
- the normalized trace: $\frac{1}{d} \operatorname{Tr} K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)$, which is also the alignment with the identity: $\frac{1}{d} K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right) \cdot{ }_{F} \mathrm{Id}$, where $\cdot{ }_{F}$ denotes the Frobenius inner product (i.e. coefficient by coefficient).

The link between these two quantities can be made explicit by developing:

$$
\left\|K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)-\mathrm{Id}\right\|_{F}^{2}=\left\|K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)\right\|_{F}^{2}-2 \operatorname{Tr}\left(K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)\right)+d
$$

which rewrites as:

$$
\left(1-\frac{D^{2}}{2 d}\right)=\frac{\operatorname{Tr}\left(K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)\right)}{d}+\frac{1}{2}\left(1-\frac{\left\|K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)\right\|_{F}^{2}}{d}\right) .
$$

The last term lies in $[0,1]$ and measures the mismatch between the vector subspaces generated by the two families of gradients $\left(\nabla_{\theta} f^{i}(\mathbf{x})\right)_{i}$ and $\left(\nabla_{\theta} f^{i}\left(\mathbf{x}^{\prime}\right)\right)_{i}$. It is 1 when $f_{\theta}(\mathbf{x})$ and $f_{\theta}\left(\mathbf{x}^{\prime}\right)$ can be moved independently, and 0 when they move jointly (though not necessarily in the same direction).
As our two similarity measures $1-\frac{D^{2}}{2 d}$ and $\frac{1}{d} \operatorname{Tr}\left(K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)\right)$ have same optimum (Id) and are closely related, in the sequel we will focus on the second one and define:

$$
\begin{equation*}
k_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)=\frac{1}{d} \operatorname{Tr} K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right) \tag{3}
\end{equation*}
$$

### 3.4 Metrics on output: rotation-invariance

Similarity in $\mathbb{R}^{d}$, to compare $\mathbf{v}$ and $\mathbf{v}^{\prime}=M \mathbf{v}$, might be richer than just checking whether the vectors are equal or close in $L^{2}$ norm.

For instance, one could quotient the output space by the group of rotations, in order to express a known or desired equivariance of the network to rotations. If the output is the predicted motion of some object described in the input, one could wish indeed that if the input object is rotated by an angle $\phi$, then the output should be rotated as well with the same angle.
In that case, given two inputs $\mathbf{x}$ and $\mathbf{x}^{\prime}$ and associated output variations $\mathbf{v}$ and $\mathbf{v}^{\prime}$, without knowing the rotation angle if applicable, one could consider all possible rotated versions $R_{\phi} \mathbf{v}^{\prime}=R_{\phi} M \mathbf{v}$, where $R_{\phi}$ is the rotation matrix with angle $\phi$, and pick the best angle $\phi$ that maximizes the alignment $\mathbf{v} \cdot R_{\phi} M \mathbf{v}$, i.e. such that $R_{\phi} M$ is the closest to the $d \times d$ identity matrix. This can be computed easily in closed form, for instance in the 2-dimensional case as follows.
The $2 \times 2$ matrix of interest (Eq. 2 ) can be written as the product of two $p \times 2$ matrices of the form $G\left(G^{T} G\right)^{-1 / 2}$, where $G$ is the matrix containing the gradient of all coordinates. Rotating the coordinates of $G$ amounts to considering $G R_{\phi}\left(R_{\phi}^{T} G^{T} G R_{\phi}\right)^{-1 / 2}=G\left(G^{T} G\right)^{-1 / 2} R_{\phi}$ instead. Thus the effect of rotation is just right-multiplying our $2 \times 2$ matrix $M$ of interest (Eq. 2 ) by $R_{\phi}$. We are thus interested into getting $M R_{\phi}$ as close as possible to the $2 \times 2$ identity. For our trace-based similarity kernel (Eq. 3], this amounts to maximizing $\operatorname{Tr}\left(M R_{\phi}\right)=\cos (\phi)\left(M_{11}+M_{22}\right)+\sin (\phi)\left(M_{12}-M_{21}\right)$ w.r.t. $\phi$, whose optimal value is:

$$
\begin{aligned}
k_{\theta}^{C, \text { rot }}\left(\mathbf{x}, \mathbf{x}^{\prime}\right) & =\frac{1}{2} \sqrt{\left(M_{11}+M_{22}\right)^{2}+\left(M_{12}-M_{21}\right)^{2}} \\
& =\frac{1}{2} \sqrt{\|M\|_{F}^{2}+2 \operatorname{det} M}
\end{aligned}
$$

where $M=K_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)$. This quantity is indeed rotation-invariant, as the Frobenius norm and the determinant do not change upon rotations. Note that one could also consider instead the subspace match $\frac{1}{d}\|M\|_{F}^{2}$. The main difference between the two is that the first one penalizes mirror symmetries (through $\operatorname{det} M$ ) while the second one does not.
Note that other metrics are possible in the output space. For instance, the loss metric quantifies the norm of a move $\mathbf{v}$ by its impact on the loss $\left.\frac{d L(y)}{d y}\right|_{f_{\theta}(\mathbf{x})}(\mathbf{v})$. It has a particular meaning though, and is relevant only if well designed and not noisy, as seen in the remote sensing image registration example.

## 4 Estimating density

### 4.1 Toy problem

The toy problem used in the paper to test the various estimators for neighbor count estimation consists of predicting a one dimensional function, namely a sinusoid (such as in Fig 1 (a)). We can easily change the difficulty of the problem by using different values of frequency. The neural network would perform this mapping: $y=\sin (2 \pi f x), x \in[0,1]$.
A problem arises however when estimating the number of neighbors because the input space has 2 boundaries at $x=0$ and $x=1$, leading to fewer neighbors when $x$ approaches either of those boundaries. To avoid this problem, we transform the input space to a 2D circle. Namely, the task is now $y=\sin (2 \pi f \alpha(x)), x \in\{(\cos (2 \pi \alpha), \sin (2 \pi \alpha)), \alpha \in[0,1]\}$, with the input space having no boundaries.

The dataset is generated with $\mathrm{n}=2048$ input points. The network used is fully-connected and has 5 hidden layers of 64 neurons trained with the Adam optimizer for 80 epochs with a base learning rate of $1 e^{-4}$. An experiment consist of training the network on a dataset generated with a specific frequency f. Each experiment was repeated 5 times, in order to take the median of every result to limit the variance due to the neural network stochastic training.

We can see in Fig 1 (b) the proposed soft estimate $k_{\theta}^{C}$ for each input point (projected to 1D). As expected we observe that the number of neighbors drops when the curvature is high: the objective changes quickly and the network adjusts to better distinguish inputs in places of higher curvature.

### 4.2 Other possible uses

Density homogeneity as an optimization criterion The estimations above are meant to be done post-training. This said, one could control density explicitly, by computing the number of neighbors for all points, and asking it to be in a reasonable range, or in a reasonable proportion $q$ of the dataset size $\mathcal{D}$, by adding e.g. to the loss $\sum_{i}\left(\frac{N_{S}\left(\mathbf{x}_{i}\right)}{\mathcal{D}}-q\right)^{2}$. Online learning could also make use of such tools, to sample first lowly-populated areas, where uncertainty is higher.

## 5 Enforcing similarity

The similarity criterion we defined could be used not only to estimate how similar two samples are perceived, after training, but also to incite the network, during training, to evolve in order to consider these samples as similar.

Asking two samples to be treated as similar If two inputs $\mathbf{x}$ and $\mathbf{x}^{\prime}$ are known to be similar (from a human point of view), one can enforce their similarity from the network perspective, by adding to the loss the term:

$$
-k_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right) .
$$

Asking a distribution of samples to be treated as similar By extension, to enforce the similarity of a subset $\mathcal{S}$ of training samples, of size $n=|\mathcal{S}|$, one might consider the average pairwise similarity $k_{\theta}^{C}$ over all pairs, or the standard deviation of the gradients. Both turn out to be equivalent to maximizing the norm of the gradient mean $\mu=\frac{1}{n} \sum_{i \in \mathcal{S}} \frac{\nabla_{\theta} f_{\theta}\left(\mathbf{x}_{i}\right)}{\left\|\nabla_{\theta} f_{\theta}\left(\mathbf{x}_{i}\right)\right\|}$ :

$$
\frac{1}{n(n-1)} \sum_{i, j \in \mathcal{S}, i \neq j} k_{\theta}^{C}\left(x_{i}, x_{j}\right)=\frac{n}{n-1}\|\mu\|^{2}-\frac{1}{n-1} \quad \text { and } \quad \operatorname{var}_{i \in \mathcal{S}} \frac{\nabla_{\theta} f_{\theta}\left(\mathbf{x}_{i}\right)}{\left\|\nabla_{\theta} f_{\theta}\left(\mathbf{x}_{i}\right)\right\|}=1-\|\mu\|^{2}
$$

In practice, common deep learning platforms are much faster when using mini-batches, but then return only the gradient sum $\sum_{i \in \mathcal{B}} \nabla_{\theta} f_{\theta}\left(\mathbf{x}_{i}\right)$ over a mini-batch $\mathcal{B}$, not individual gradients, preventing the normalization of each of them to compute $k_{\theta}^{C}$ or $\mu$. So instead we compare means of un-normalized gradients, over two mini-batches $\mathcal{B}_{1}$ and $\mathcal{B}_{2}$ comprising each $n_{B}$ samples from $\mathcal{S}$, which yields the criterion:

$$
n_{B} \frac{\left\|\mu_{1}-\mu_{2}\right\|^{2}}{\left\|\mu_{1}\right\|\left\|\mu_{2}\right\|} \quad \text { where } \quad \mu_{k}=\frac{1}{n} \sum_{i \in \mathcal{B}_{k}} \nabla_{\theta} f_{\theta}\left(\mathbf{x}_{i}\right)
$$

The factor $n_{B}$ counterbalances the $\frac{1}{\sqrt{n_{B}}}$ variance reduction effect due to averaging over $n_{B}$ samples.

### 5.1 Group invariance

Dataset augmentation is a standard machine learning technique; when augmenting the dataset by a group transformation of the input (e.g., translation, rotation...) or by small intensity noise, new samples are artificially created, to augment the dataset size and hope for invariance to such transformations. One can ask the network to consider orbits of samples [2] as similar with the technique above.
Furthermore, if the group infinitesimal elements are expressible as differential operators $e_{k}$, one could require directly, for all $\mathbf{x}$, invariance in the tangent plane in the directions of these differential operators:

$$
\left\|\partial_{\mathbf{x}} \nabla_{\theta} f(\mathbf{x}) \cdot e_{k}(\mathbf{x})\right\|^{2}
$$

which is the limit of $\frac{1}{\varepsilon^{2}}\left\|\nabla_{\theta} f_{\theta}(\mathbf{x})-\nabla_{\theta} f_{\theta}\left(\mathbf{x}+\varepsilon e_{k}(\mathbf{x})\right)\right\|^{2}$ when $\varepsilon \rightarrow 0$. For instance, in the case of image translations, the operator is $e: \mathbf{x} \mapsto \nabla_{x} \mathbf{x}(x)$ where $x$ denotes spatial coordinates in the image $\mathbf{x}$, as $\mathbf{x}(x+\tau)=\mathbf{x}(x)+\tau \cdot \nabla_{x} \mathbf{x}(x)+O\left(\tau^{2}\right)$. This is however not recommended, as representing a translation with such a spatially-local operator does not take into account the spatially-irregular nature of image intensities.


Figure 1: Toy problem with the frequency $\mathrm{f}=2$.


Figure 2: 3D plot of neighbors soft with varying frequency. Script and data to plot interactively in attached files. Run the bash script "main_plot_exps.paper.sh" to reproduce this exact figure. Alternatively use "main_plot_exps.py" with arguments of your choosing to plot different values (run "python main_plot_exps.py -h" to see possible arguments).

Note that to the opposite of standard robustification techniques considering regularizers such as $\sum_{\mathbf{x}}\left\|\nabla_{\theta} f_{\theta}(\mathbf{x})\right\|^{2}$, we ask not gradients to be always small, but to be smooth, and in certain directions only.

Complexity A gradient descent step on our criterion for a given pair ( $\mathrm{x}, \mathrm{x}^{\prime}$ ) (in a mini-batch approach, e.g.) requires the computation of the gradient $\nabla_{\theta} k_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)=\nabla_{\theta}\left(\nabla_{\theta} f_{\theta}(\mathbf{x}) \cdot \nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)\right)$. While a naive approach would require the computation of a second derivative, i.e. a matrix of size $p \times p$ where $p$ is the number of parameters, it is actually possible to compute $\nabla_{\theta} k_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)=$ $\nabla_{\theta} \sum_{i} \frac{d f_{\theta}(\mathbf{x})}{d \theta_{i}} \frac{d f_{\theta}\left(\mathbf{x}^{\prime}\right)}{d \theta_{i}}$ in linear time $O(p)$, taking advantage of the serial structure of the computational graph. The framework enabling such computations is already available on common deep learning platforms, under the name of double-backpropagation routine [3, 6, 10, 5], roughly doubling the computational time of a gradient step. It was initially intended for the computation of $\nabla_{\mathbf{x}} \nabla_{\theta} f_{\theta}(\mathbf{x})$ for some variations on GANs.

Dynamics of learning Our approach enforces similarity not just at the output level, but within the whole internal computational process. Therefore, during training, information is provided directly to each parameter instead of being back-propagated through possibly many layers. Thus the dynamics of learning are expected to be different, especially for deep networks.
To test this hypothesis, we train a small network on MNIST with and without the similarity criteria acting as an auxiliary loss (see Fig. 3). As a result, we observe an acceleration of the convergence very early in the learning process. It is worth noting that this effect can be observed across a wide range of different neural architectures. We performed additional experiments on toy datasets as well as on CIFAR10 with no or only negligible improvements. All together this suggests that using the similarity criteria during training may be beneficial to specific datasets as opposed to specific architectures, and indeed, as the class intra-variability in CIFAR10 is known to be high, considering all examples of a class of CIFAR10 as similar is less relevant.


Figure 3: Validation accuracy of a neural network trained on MNIST with and without the similarity criterion (note that the x -axis is the number of minibatches presented to the network, not of epochs).

Experimentation details The results in Figure 3 show the average and standard deviation over 60 runs for each curve. The x -axis is the number of batches to the network is trained on (with a batch size of 16). The y-axis is the accuracy metric on the whole validation set. The network architecture is made of 2 convolutions layers (with a kernel size of 5), 2 linear layers and uses PReLU non-linearities. We used Adam with a learning rate of 1e-3 and no weight decay.
We tested other architectures on MNIST: one with residual blocks, one deeper ( 8 convolutions) and one with tanh non-linearities. Similar results were observed on all cases. Additional tests were performed on CIFAR10 with a VGG architecture and only negligible benefits were observed.

## 6 Noisy Map Alignment Analysis

The task here it to align maps in the form of a list of polygons with remote sensing images while using only the available noisy annotations. We analyze the model developed in a previous work [4]. Specifically, the model is trained in a multiple-rounds training scheme to iteratively align the available noisy annotations, which provides a better ground truth used to train a better model in the next round. We aim to answer the question of why multiple rounds are needed in this noisy supervision setting, and why not all the noise can be removed in a single training step.

More specifically, the model is made out of 4 neural networks. Each is trained on a different resolution (in terms of ground pixel size) and are applied in a multi-resolution pyramidal manner. In all our experiments we only analyzed the networks trained for a ground pixel size of 4 time smaller than the reference ground pixel size which is 0.3 m . We used the already-trained networks for each round, of which there are 3.
The network was trained with small patches of (image, misaligned map) pairs from images of the Inria dataset [8] and the Bradbury dataset [1]. Ideally we would want to compute the similarities of every possible pairs of inputs, with a small patch size of 124 px . However, given that a typical image of the training dataset is $1250 \times 1250 \mathrm{px}$ (after rescaling) and there are a few hundred of them ( 328 from the Inria dataset, only counting images where OSM annotations [9]), this would result in 32800 patches. The resulting amount of similarities to compute would be around half a billion. As the network has a few million of parameters and the output is 2 D , each computation of similarity takes around 0.5 s . To make any computation feasible, we first sample 10 patches per image from the 328 of the Inria dataset. Those patches are chosen at random, as long as there is at least one building lying fully in the patch. As some images have rather sparse buildings, some images give less than 10 patches. We thus obtain 3045 patches representing the dataset. The amount of similarities to compute would be close to 5 million. To study all patches globally, we can use the soft neighbors estimator $k_{\theta}^{C}$ which has a linear complexity and allows us to compute the amount of neighbors for all 3045 patches in under an hour on a GTX 1080Ti. However it is also interesting to go in deeper detail and compute similarities for some input pairs. We thus furthermore reduce the amount of pairs by estimating all similarities only for a very small number of patches, for example 10 . This results in a $10 \times 3045$ similarity matrix.

### 6.1 Soft estimate on a sampling of the training dataset

In this section we present the results of computing the soft neighbors estimator $k_{\theta}^{C}$ on the 3045 sampled patches of inputs. We obtain results for the 3 networks of the 3 rounds of the noisysupervision multi-rounds training scheme. Fig 4 shows a histogram of the soft neighbors estimations. It additionally representative input patches for each bin of the histogram. Those representative patches are chosen so that their neighbor count is closest to the right edge of that bin. We especially observe that inputs in round 2 have more neighbors than the other 2 rounds. This particularity of round 2 will be seen throughout the remaining results. It is the round that aligns the most the annotations (see the Fig. 2 on accuracy cumulative distributions in the paper). Round 3 does not perform any more alignment, that might be the reason why its results are different from those of round 2 .

### 6.2 Similarities on pairs of input patches

In this section are the results for the computation of similarities between pairs of input patches. In a first experiment, for every round we chose the 10 patches shown in Fig 4 , and computed their similarities with all the other 3045 patches. In order to visualize this data, we computed the 10 -nearest neighbors in terms of similarity for each of those patches, see Fig 5, 6, 7 We computed the histogram of similarities as well, see Fig 8 .
In a second experiment, to better compare between rounds, we used another set of 10 patches, this time the same set for each round. Specifically, we sampled 10 patches from the bloomington 22 image of the Inria dataset. As just before we computed the 10-nearest neighbors (Fig, 9, 10, 11) and the histogram of similarities( Fig 12 ) for a visualization of those measures.
Generally speaking, inputs in round 2 have more neighbors and the 10-nearest ones are closer than in other rounds (see Fig 5, 6, 7 and Fig, 9, 10, 11). For each parch, its closest neighbors generally (for similarity $>0.8$ ) look similar from a human point of view. For example patches with sparse houses


Figure 4: Histogram of the soft estimate of neighbors on 3045 patches. Horizontal scale is different for each.


Figure 5: Round 1: k-nearest neighbors with $\mathrm{k}=10$. The 10 patches selected correspond to the 10 patches of Fig 4 for that round.
and trees have the same kind of neighbors. The same can be said for patches with parking lots and big roads. Another group are patches that are almost empty of buildings, with a lot of low vegetation. Other patch nearest neighbors are more difficult to interpret. In Fig 8 and Fig 12 we can see that for round 2, the spread of the similarities of the selected patches is smaller and the peak of the histogram are closer to the right, meaning all patches are closer than in other rounds. Additionally in Fig 8 we can observe that the bottom patch has closer neighbors than the top patch, this is because the top patch corresponds to the left patch in 4 and the bottom one corresponds to the right patch in 4 ,

We additionally computed a perceptual loss on the same 10 patches from the bloomington 22 image for comparison (see Fig $13,14,15$. We used the activations from the second to last layer of the displacement branch of our network to compute it. In general it does not estimate neighbors as well, especially regarding the type of buildings and their arrangement in the image patch.


Figure 6: Round 2: k-nearest neighbors with $\mathrm{k}=10$. The 10 patches selected correspond to the 10 patches of Fig 4 for that round.


Figure 7: Round 3: k-nearest neighbors with $\mathrm{k}=10$. The 10 patches selected correspond to the 10 patches of Fig 4 for that round.
(a) Round 1

(b) Round 2

(c) Round 3


Figure 8: Histograms of similarities shown for the same 10 patches as Fig 4 and Fig $5,6,7$.


Figure 9: Round 1: k -nearest neighbors with $\mathrm{k}=10$. The 10 patches are from from the bloomington 22 image. Same patch selection across rounds.


Figure 10: Round 2: k -nearest neighbors with $\mathrm{k}=10$. The 10 patches are from from the bloomington 22 image. Same patch selection across rounds.


Figure 11: Round 3: k-nearest neighbors with $\mathrm{k}=10$. The 10 patches are from from the bloomington 22 image. Same patch selection across rounds.
(a) Round 1

(b) Round 2

(c) Round 3


Figure 12: Histograms of similarities shown for the same 10 patches as in Fig, 9, 10, 11 Same patch selection across rounds.


Figure 13: Round 1: perceptual loss used for $k$-nearest neighbors with $k=10$. The 10 patches are from from the bloomington 22 image. Same patch selection across rounds.


Figure 14: Round 2: perceptual loss used for k -nearest neighbors with $\mathrm{k}=10$. The 10 patches are from from the bloomington 22 image. Same patch selection across rounds.


Figure 15: Round 3: perceptual loss used for $k$-nearest neighbors with $k=10$. The 10 patches are from from the bloomington 22 image. Same patch selection across rounds.

### 6.3 Proof details of the self-denoising effect quantification

### 6.3.1 Magnitude of kernel-smoothed i.i.d. noise

We show here that $\mathbb{E}_{k}[\varepsilon] \propto \operatorname{var}_{\varepsilon}\left(\mathbb{E}_{k}[\varepsilon]\right)^{1 / 2}=\sigma_{\varepsilon}\left\|k_{\theta}^{I N}\right\|_{L 2}$.
Let us denote by $\mathbb{E}_{\varepsilon}[]$ and $\operatorname{var}_{\varepsilon}()$ the expectation and variance with respect to the random variable $\varepsilon$. As a reminder, by assumptions in the noise definition, $\varepsilon=\left(\varepsilon_{i}\right)_{i}$ is a random, i.i.d. noise, centered and of variance $\sigma_{\varepsilon}$.

This is not to be confused with the symbol $\mathbb{E}_{k}[]$, which was defined as, for any vector field $a$ :

$$
\underset{k}{\mathbb{E}}[a]=\sum_{j} a_{j} k_{\theta}^{I N}\left(\mathbf{x}_{j}, \mathbf{x}_{i}\right),
$$

i.e. as the mean value of $a$ in the neighborhood of $i$, that is, the weighted average of the $a_{j}$ with weights $k_{\theta}^{I N}\left(\mathbf{x}_{j}, \mathbf{x}_{i}\right)$, which are positive and sum up to 1 .

Given a network and its associated kernel $k_{\theta}^{I N}$, we are interested in to knowing the typical values of $\mathbb{E}_{k}[\varepsilon]$ for random $\varepsilon$. First, the expectation over the noise of $\mathbb{E}_{k}[\varepsilon]$ is:

$$
\underset{\varepsilon}{\mathbb{E}}[\underset{k}{\mathbb{E}}[\varepsilon]]=\underset{\varepsilon}{\mathbb{E}}\left[\sum_{j} \varepsilon_{j} k_{\theta}^{I N}\left(\mathbf{x}_{j}, \mathbf{x}_{i}\right)\right]=\sum_{j} \underset{\varepsilon}{\mathbb{E}}\left[\varepsilon_{j}\right] k_{\theta}^{I N}\left(\mathbf{x}_{j}, \mathbf{x}_{i}\right)=0
$$

as $\varepsilon$ is a centered noise. Thus the random variable $\mathbb{E}_{k}[\varepsilon]$ is also centered, and therefore its typical values are described by its standard deviation, which is the square root of its variance:

$$
\underset{k}{\mathbb{E}}[\varepsilon] \propto \operatorname{var}_{\varepsilon}(\underset{k}{\mathbb{E}}[\varepsilon])^{1 / 2} .
$$

The variance can be computed as follows:

$$
\begin{aligned}
\operatorname{var}_{\varepsilon}(\underset{k}{\mathbb{E}}[\varepsilon]) & =\underset{\varepsilon}{\mathbb{E}}\left[\left(\sum_{j} \varepsilon_{j} k_{\theta}^{I N}\left(\mathbf{x}_{j}, \mathbf{x}_{i}\right)\right)^{2}\right] \\
& =\underset{\varepsilon}{\mathbb{E}}\left[\sum_{j} \varepsilon_{j}^{2}\left(k_{\theta}^{I N}\left(\mathbf{x}_{j}, \mathbf{x}_{i}\right)\right)^{2}\right] \quad \text { as } \varepsilon \text { is i.i.d. } \\
& =\sigma_{\varepsilon}^{2} \sum_{j}\left(k_{\theta}^{I N}\left(\mathbf{x}_{j}, \mathbf{x}_{i}\right)\right)^{2} \\
& =\sigma_{\varepsilon}^{2}\left\|k_{\theta}^{I N}\left(\cdot, \mathbf{x}_{i}\right)\right\|_{L 2}^{2}
\end{aligned}
$$

As the weights $p_{j}=k_{\theta}^{I N}\left(\mathbf{x}_{j}, \mathbf{x}_{i}\right)$, for given $i$ and varying $j$, are positive and sum up to 1 , they form a probability distribution. Hence the value of $\left\|k_{\theta}^{I N}\left(\cdot, \mathbf{x}_{i}\right)\right\|_{L 2}^{2}=\|p\|_{L 2}^{2}$ satisfies:

- $\|p\|_{L 2} \leqslant 1, \quad$ as $\sum_{j} p_{j}^{2} \leqslant \sum_{j} p_{j}=1$, with equality only when $p_{j}=p_{j}^{2} \forall j$, that is, all $p_{j}=0$ except for one $p_{j^{*}}=1$, which means $k_{\theta}^{I}\left(\mathbf{x}_{j}, \mathbf{x}_{i}\right)=0 \quad \forall j \neq i$, which means that all data samples are fully independent from the network's point of view.
- $\|p\|_{L 2} \geqslant \frac{1}{\sqrt{N}} \quad$ as $1=\sum_{j} 1 \times p_{j} \leqslant\|1\|_{L 2}\|p\|_{L 2}=\sqrt{N}\|p\|_{L 2}$ (Cauchy-BunyakovskySchwarz), with equality reached for the uniform distribution: $p_{j}=\frac{1}{N} \forall j$, where $N$ is the number of data samples. This implies that all $k_{\theta}^{C}\left(\mathbf{x}_{j}, \mathbf{x}_{i}\right)$ are equal, for all $i, j$, hence they are all equal to $k_{\theta}^{C}\left(\mathbf{x}_{i}, \mathbf{x}_{i}\right)=1$. This is the case studied in [7]: all input points are identical.

The denoising factor $\left\|k_{\theta}^{I N}\left(\cdot, \mathbf{x}_{i}\right)\right\|_{L 2}$, which depends on the data point $\mathbf{x}_{i}$ considered, thus expresses where the neighborhood of $\mathbf{x}_{i}$ lies, between these two extremes (all $\mathbf{x}_{j}$ very different from $\mathbf{x}_{i}$, or all identical).

Note: the results above remain valid when the output is higher-dimensional, under the supplementary assumption that the covariance matrix of the noise is proportional to the Identity matrix (i.e., the noises on the various coefficients of the label vector are independent from each other, and follow the same law, with standard deviation $\left.\sigma_{\varepsilon}\right)$. If not, the expression for $\operatorname{covar}_{\varepsilon}\left(\mathbb{E}_{k}[\varepsilon]\right)$ is more complex, as $\Sigma_{\varepsilon}$ and $k_{\theta}^{I N}$ interact. Note that when the output is of dimension $d$, the kernel $k_{\theta}^{I N}\left(\mathbf{x}_{j}, \mathbf{x}_{i}\right)$ is a $d \times d$ matrix, thus the denoising factor $\left\|k_{\theta}^{I N}\left(\cdot, \mathbf{x}_{i}\right)\right\|_{L 2}^{2}$ has to be replaced with the matrix $\sum_{j} k_{\theta}^{I N}\left(\mathbf{x}_{j}, \mathbf{x}_{i}\right) k_{\theta}^{I N}\left(\mathbf{x}_{j}, \mathbf{x}_{i}\right)^{T}$, which can be summarized by its trace, which is the $L^{2}$ norm of the Frobenius norms: $\left\|\left\|k_{\theta}^{I N}\left(\cdot, \mathbf{x}_{i}\right)\right\|_{F}\right\|_{L 2}^{2}$.

### 6.3.2 The function: gradient $\mapsto$ output is Lipschitz

Theorem 1 implies that the application: $\frac{\nabla_{\theta} f_{\theta}(\mathbf{x})}{\left\|\nabla_{\theta} f_{\theta}(\mathbf{x})\right\|} \mapsto f_{\theta}(\mathbf{x})$ is well-defined. We show here that this application is also Lipschitz, with a network-dependent constant, under mild hypotheses.
We consider the same assumptions as in Theorem 11: $f_{\theta}$ is a real-valued network, whose last layer is a linear layer or a standard activation function thereof (such as sigmoid, tanh, ReLU...), without parameter sharing (in that last layer). We will also require that the derivative of the activation function is bounded, which is a safe assumption for all networks meant to be trained by gradient descent. Another, technical property (bounded input space) will be assumed in order to imply bounded gradients. A side note indicates how to rewrite the desired property if the input space is not bounded.
Let $\mathbf{x}$ and $\mathbf{x}^{\prime}$ be any two inputs. We want to bound $\left|f_{\theta}(\mathbf{x})-f_{\theta}\left(\mathbf{x}^{\prime}\right)\right|$ by $\left\|\mathbf{u}-\mathbf{u}^{\prime}\right\|_{2}$ times some constant, where $\mathbf{u}=\frac{\nabla_{\theta} f_{\theta}(\mathbf{x})}{\left\|\nabla_{\theta} f_{\theta}(\mathbf{x})\right\|}$ and $\mathbf{u}^{\prime}=\frac{\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)}{\left\|\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)\right\|}$.
Let us denote the non-normalized gradients by $\mathbf{v}=\nabla_{\theta} f_{\theta}(\mathbf{x})$ and $\mathbf{v}^{\prime}=\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)$. We have $\mathbf{u}=\frac{\mathbf{v}}{\|\mathbf{v}\|}$ and $\mathbf{u}^{\prime}=\frac{\mathbf{v}^{\prime}}{\left\|\mathbf{v}^{\prime}\right\|}$.
We will proceed in two steps: bounding $\left|f_{\theta}(\mathbf{x})-f_{\theta}\left(\mathbf{x}^{\prime}\right)\right|$ by $\left\|\mathbf{v}-\mathbf{v}^{\prime}\right\|_{2}$, and then $\left\|\mathbf{v}-\mathbf{v}^{\prime}\right\|_{2}$ by $\left\|\mathbf{u}-\mathbf{u}^{\prime}\right\|_{2}$. The first step is easy and actually sufficient to bound with a non-normalized similarity kernel $k_{\theta}=\mathbf{v} \cdot \mathbf{v}^{\prime}$ the shift from the average prediction in the neighborhood. The second step provides a more elegant bound, in that it makes use of the normalized similarity kernel $k_{\theta}^{C}=\mathbf{u} \cdot \mathbf{u}^{\prime}$, but that bound is a priori not as tight and requires more assumptions.

## Case where the last layer is linear

The output of the network is of the form

$$
f_{\theta}(\mathbf{x})=\sum_{i} w_{i} a_{i}(\mathbf{x})+b
$$

where $w_{i}$ and $b$ are parameters in $\mathbb{R}$ and $a_{i}(\mathbf{x})$ activities from previous layers. Thus:

$$
\begin{aligned}
\left|f_{\theta}(\mathbf{x})-f_{\theta}\left(\mathbf{x}^{\prime}\right)\right| & =\left|\sum_{i} w_{i}\left(a_{i}(\mathbf{x})-a_{i}\left(\mathbf{x}^{\prime}\right)\right)\right| \\
& \leqslant\|\mathbf{w}\|_{2}\left\|\mathbf{a}(\mathbf{x})-\mathbf{a}\left(\mathbf{x}^{\prime}\right)\right\|_{2} \\
& \leqslant\|\mathbf{w}\|_{2} \sqrt{\sum_{i}\left(v_{i}-v_{i}^{\prime}\right)^{2}}
\end{aligned}
$$

where the sum is taken over parameters $i$ in the last layer only, using the fact that activities $a_{i}$ in the last layer are equal to some of the coefficients of the gradient: $v_{i}:=\frac{\partial f_{\theta}(\mathbf{x})}{\partial w_{i}}=a_{i}(\mathbf{x})$.
Note that the derivative with respect to the shift $b$ is $v_{b}:=\frac{\partial f_{\theta}(\mathbf{x})}{\partial b}=1$, which ensures that the norm of $\mathbf{v}$ is at least 1 . This implies:

$$
\left\|\mathbf{u}-\mathbf{u}^{\prime}\right\|_{2} \geqslant\left|u_{b}-u_{b}^{\prime}\right|=\left|\frac{1}{\|\mathbf{v}\|}-\frac{1}{\left\|\mathbf{v}^{\prime}\right\|}\right|
$$

which, combined with:

$$
\left|v_{i}-v_{i}^{\prime}\right|=\left\|\mathbf{v}^{\prime}\right\|\left|\frac{1}{\left\|\mathbf{v}^{\prime}\right\|} v_{i}-\frac{v_{i}^{\prime}}{\left\|\mathbf{v}^{\prime}\right\|}\right|=\left\|\mathbf{v}^{\prime}\right\|\left|u_{i}-u_{i}^{\prime}+\left(\frac{1}{\left\|\mathbf{v}^{\prime}\right\|}-\frac{1}{\|\mathbf{v}\|}\right) v_{i}\right|
$$

yields:

$$
\left|v_{i}-v_{i}^{\prime}\right| \leqslant\left\|\mathbf{v}^{\prime}\right\|\left(\left|u_{i}-u_{i}^{\prime}\right|+\left\|\mathbf{u}-\mathbf{u}^{\prime}\right\|_{2}\left|v_{i}\right|\right) \leqslant\left\|\mathbf{v}^{\prime}\right\|\left\|\mathbf{u}-\mathbf{u}^{\prime}\right\|_{2}\left(1+\left|v_{i}\right|\right)
$$

from which we finally obtain:

$$
\left|f_{\theta}(\mathbf{x})-f_{\theta}\left(\mathbf{x}^{\prime}\right)\right| \leqslant\left[\|\mathbf{w}\|_{2}\left\|\mathbf{v}^{\prime}\right\| \sqrt{\sum_{i}\left(1+\left|v_{i}\right|\right)^{2}}\right]\left\|\mathbf{u}-\mathbf{u}^{\prime}\right\|_{2}
$$

which is the bound we were searching for. For the term between brackets to be bounded by a networkdependent constant, one can suppose for instance that the derivative of the activation functions is bounded (which is usually the case for networks meant to be trained by gradient descent), and that the input space is bounded as well; in such cases indeed all coefficients of the gradient vector $\mathbf{v}$ or $\mathbf{v}^{\prime}$ are bounded, as derivatives of a function composed of constant linear applications (except for the first layer which is a linear application whose factors are bounded inputs, when seen as an application defined on parameters) and of bounded-derivatives activation functions.

Note for unbounded input spaces: If the input space is not bounded, the gradients are not bounded absolutely, as for instance the gradient with respect to a weight in the first layer is the input itself (times a chain product). In that case the application $\mathbf{x} \mapsto \mathbf{v}$ still satisfy a bound of the form $\|\mathbf{v}\| \leqslant(1+\|\mathbf{x}\|) A$, with $A$ a network-dependent constant (product of determiners of layer weight matrices and of the bound on activation function derivatives to the power: network depth), and thus the application $\mathbf{u} \mapsto f_{\theta}(\mathbf{x})$ still satisfies a bound of the form, for any $\mathbf{x}, \mathbf{x}^{\prime}$ :

$$
\left|f_{\theta}(\mathbf{x})-f_{\theta}\left(\mathbf{x}^{\prime}\right)\right| \leqslant B(1+\|\mathbf{x}\|)\left(1+\left\|\mathbf{x}^{\prime}\right\|\right)\left\|\mathbf{u}-\mathbf{u}^{\prime}\right\|_{2}
$$

The last statement in the paper then becomes

$$
\left|\underset{k}{\mathbb{E}}\left[\widehat{y}_{i}-\widehat{y}\right]\right| \leqslant \sqrt{2} B\left(1+\left\|\mathbf{x}_{i}\right\|\right) \max _{j}\left(1+\left\|\mathbf{x}_{j}\right\|\right) \underset{k}{\mathbb{E}}\left[\sqrt{1-k_{\theta}^{C}\left(\mathbf{x}_{i}, \cdot\right)}\right]
$$

which in practice rewrites as the original formulation:

$$
\left|\underset{k}{\mathbb{E}}\left[\widehat{y}_{i}-\widehat{y}\right]\right| \leqslant \sqrt{2} C \underset{k}{\mathbb{E}}\left[\sqrt{1-k_{\theta}^{C}\left(\mathbf{x}_{i}, \cdot\right)}\right]
$$

by taking $C=B \max _{j}\left(1+\left\|\mathbf{x}_{j}\right\|\right)^{2}$, considering the actual diameter of the given dataset.

## Case where the last layer is an activation function of a linear layer

The output of the network is of the form

$$
f_{\theta}(\mathbf{x})=\sigma\left(\sum_{i} w_{i} a_{i}(\mathbf{x})+b\right)
$$

and, as the derivative of $\sigma$ is assumed to be bounded, and as the weights $w_{i}$ are fixed, $f_{\theta}(\mathbf{x})$ is a Lipschitz function of the last layer activities $a_{i}(\mathbf{x})$. Therefore:

$$
\left|f_{\theta}(\mathbf{x})-f_{\theta}\left(\mathbf{x}^{\prime}\right)\right| \leqslant K\left\|\mathbf{a}(\mathbf{x})-\mathbf{a}\left(\mathbf{x}^{\prime}\right)\right\|_{2}
$$

We will denote by $\alpha$ and $\alpha^{\prime}$ the derivatives with respect to the shift $b$, which are this time:

$$
\alpha:=\mathbf{v}_{b}:=\frac{\partial f_{\theta}(\mathbf{x})}{\partial b}=\left.\sigma^{\prime}\right|_{\sum_{i} w_{i} a_{i}(\mathbf{x})+b} \quad \text { and } \quad \alpha^{\prime}:=\mathbf{v}_{b}^{\prime}:=\frac{\partial f_{\theta}\left(\mathbf{x}^{\prime}\right)}{\partial b}=\left.\sigma^{\prime}\right|_{\sum_{i} w_{i} a_{i}\left(\mathbf{x}^{\prime}\right)+b}
$$

We proceed as previously:

$$
\left\|\mathbf{u}-\mathbf{u}^{\prime}\right\|_{2} \geqslant\left|u_{b}-u_{b}^{\prime}\right|=\left|\frac{\alpha}{\|\mathbf{v}\|}-\frac{\alpha^{\prime}}{\left\|\mathbf{v}^{\prime}\right\|}\right|
$$

which, combined with:

$$
\left|a_{i}-a_{i}^{\prime}\right|=\left|\frac{v_{i}}{\alpha}-\frac{v_{i}^{\prime}}{\alpha^{\prime}}\right|=\frac{\left\|\mathbf{v}^{\prime}\right\|}{\alpha^{\prime}}\left|\frac{\alpha^{\prime}}{\alpha\left\|\mathbf{v}^{\prime}\right\|} v_{i}-\frac{v_{i}^{\prime}}{\left\|\mathbf{v}^{\prime}\right\|}\right|=\frac{\left\|\mathbf{v}^{\prime}\right\|}{\alpha^{\prime}}\left|u_{i}-u_{i}^{\prime}+\frac{v_{i}}{\alpha}\left(\frac{\alpha^{\prime}}{\left\|\mathbf{v}^{\prime}\right\|}-\frac{\alpha}{\|\mathbf{v}\|}\right)\right|
$$

yields:

$$
\left|a_{i}-a_{i}^{\prime}\right| \leqslant \frac{\left\|\mathbf{v}^{\prime}\right\|}{\alpha^{\prime}}\left(\left|u_{i}-u_{i}^{\prime}\right|+\left\|\mathbf{u}-\mathbf{u}^{\prime}\right\|_{2}\left|a_{i}\right|\right) \leqslant \frac{\left\|\mathbf{v}^{\prime}\right\|}{\alpha^{\prime}}\left(1+\left|a_{i}\right|\right)\left\|\mathbf{u}-\mathbf{u}^{\prime}\right\|_{2}
$$

from which we finally obtain:

$$
\left|f_{\theta}(\mathbf{x})-f_{\theta}\left(\mathbf{x}^{\prime}\right)\right| \leqslant\left[K \frac{\left\|\mathbf{v}^{\prime}\right\|}{\alpha^{\prime}} \sqrt{\sum_{i}\left(1+\left|a_{i}\right|\right)^{2}}\right]\left\|\mathbf{u}-\mathbf{u}^{\prime}\right\|_{2}
$$

Note that $\alpha^{\prime}$ is actually a factor of each coefficient of $\mathbf{v}^{\prime}$, as the derivative of $f_{\theta}\left(\mathbf{x}^{\prime}\right)$ with respect to any parameter is a chain rule starting with $\frac{\partial f_{\theta}\left(\mathbf{x}^{\prime}\right)}{\partial b}=\left.\sigma^{\prime}\right|_{\sum_{i} w_{i} a_{i}\left(\mathbf{x}^{\prime}\right)+b}=\alpha^{\prime}$. To bound the term between brackets, the same assumptions as previously are sufficient. One can assume that $\alpha$ and $\alpha^{\prime}$ are not 0 , as, if they are, the problem is of little interest ( $\mathbf{u}$ or $\mathbf{u}^{\prime}$ being then not defined).

### 6.3.3 Additional proof detail

The kernel $k_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)$, by definition, is the $L^{2}$ inner product between two unit vectors:

$$
k_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)=\frac{\nabla_{\theta} f_{\theta}(\mathbf{x})}{\left\|\nabla_{\theta} f_{\theta}(\mathbf{x})\right\|} \cdot \frac{\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)}{\left\|\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)\right\|} .
$$

As, for any two unit vectors $a$ and $b$ :

$$
\|a-b\|^{2}=a^{2}+b^{2}-2 a \cdot b=2(1-a \cdot b)
$$

we get:

$$
\left\|\frac{\nabla_{\theta} f_{\theta}(\mathbf{x})}{\left\|\nabla_{\theta} f_{\theta}(\mathbf{x})\right\|}-\frac{\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)}{\left\|\nabla_{\theta} f_{\theta}\left(\mathbf{x}^{\prime}\right)\right\|}\right\|=\sqrt{2} \sqrt{1-k_{\theta}^{C}\left(\mathbf{x}, \mathbf{x}^{\prime}\right)}
$$

## 7 Other related remarks

Data augmentation can be seen as label denoising, as it multiplies the number of neighbors. Indeed, in the infinite sampling limit, where the dataset becomes a probability distribution over all possible images, adding a transformed copy $\mathbf{x}^{\prime}=T_{\phi} \mathbf{x}$ of a given point $\mathbf{x}$ (e.g. rotating it with an angle $\alpha_{\phi}$ and adding small noise $\varepsilon_{\phi}$ ) means adding $\left(\mathbf{x}^{\prime}, l(\mathbf{x})\right)$ to the dataset, where $l(\mathbf{x})$ is the desired label for $\mathbf{x}$. But if $\left(\mathbf{x}^{\prime}, l\left(\mathbf{x}^{\prime}\right)\right)$ was already in the dataset, this amounts to enriching the possible labels for $\mathbf{x}^{\prime}$. Supposing $T_{\phi}$ is an invertible transformation parameterized by $\phi$, full data augmentation (i.e. for all possible $\phi$, applied on all points $\mathbf{x}$ ) enriches $\mathbf{x}^{\prime}$ with all labels $l\left(T_{\phi}^{-1}\left(\mathbf{x}^{\prime}\right)\right)$. In case of i.i.d. label noise, data augmentation will thus reduce this noise by a factor $\sqrt{ }$ number of copies.

## References

[1] Kyle Bradbury, Benjamin Brigman, Leslie Collins, Timothy Johnson, Sebastian Lin, Richard Newell, Sophia Park, Sunith Suresh, Hoel Wiesner, and Yue Xi. Aerial imagery object identification dataset for building and road detection, and building height estimation, July 2016.
[2] Taco Cohen and Max Welling. Group equivariant convolutional networks. In International conference on machine learning, pages 2990-2999, 2016.
[3] Harris Drucker and Yann Le Cun. Double backpropagation increasing generalization performance. In IJCNN-91-Seattle International Joint Conference on Neural Networks, volume 2, pages 145-150. IEEE, 1991.
[4] Nicolas Girard, Guillaume Charpiat, and Yuliya Tarabalka. Noisy Supervision for Correcting Misaligned Cadaster Maps Without Perfect Ground Truth Data. In IGARSS, July 2019. URL https://hal.inria fr/hal-02065211.
[5] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In Advances in Neural Information Processing Systems, pages 5767-5777, 2017.
[6] Sepp Hochreiter and Jürgen Schmidhuber. Simplifying neural nets by discovering flat minima. In Advances in neural information processing systems, pages 529-536, 1995.
[7] Jaakko Lehtinen, Jacob Munkberg, Jon Hasselgren, Samuli Laine, Tero Karras, Miika Aittala, and Timo Aila. Noise2noise: Learning image restoration without clean data. In International Conference on Machine Learning, pages 2971-2980, 2018.
[8] Emmanuel Maggiori, Yuliya Tarabalka, Guillaume Charpiat, and Pierre Alliez. Can semantic labeling methods generalize to any city? the Inria aerial image labeling benchmark. In IGARSS, 2017.
[9] OpenStreetMap contributors. Planet dump retrieved from https://planet.osm.org, 2017.
[10] Salah Rifai, Pascal Vincent, Xavier Muller, Xavier Glorot, and Yoshua Bengio. Contractive auto-encoders: Explicit invariance during feature extraction. In Proceedings of the 28 th International Conference on International Conference on Machine Learning, pages 833-840. Omnipress, 2011.

