



HAL
open science

Cartan moving frames and the data manifolds

Eliot Tron, Rita Fioresi, Nicolas Couellan, Stéphane Puechmorel

► **To cite this version:**

Eliot Tron, Rita Fioresi, Nicolas Couellan, Stéphane Puechmorel. Cartan moving frames and the data manifolds. information geometry, inPress. hal-04701811v2

HAL Id: hal-04701811

<https://enac.hal.science/hal-04701811v2>

Submitted on 7 Nov 2024

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

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

Cartan moving frames and the data manifolds

Eliot Tron^{1*}, Rita Fioresi³, Nicolas Couellan^{1,2},
Stéphane Puechmorel^{1,2}

¹Ecole Nationale de l'Aviation Civile, 7 Avenue Edouard Belin, Toulouse,
31400, France.

²Institut de Mathématiques de Toulouse, UMR 5219, Université de
Toulouse, CNRS, UPS, 118 route de Narbonne, Toulouse, F-31062 Cedex
9, France.

³FaBiT, Università di Bologna, via S. Donato 15, I-40126 Bologna, Italy.

*Corresponding author(s). E-mail(s): eliot.tron@enac.fr;
Contributing authors: nicolas.couellan@recherche.enac.fr;
stephane.puechmorel@enac.fr; rita.fioresi@UniBo.it;

Abstract

The purpose of this paper is to employ the language of Cartan moving frames to study the geometry of the data manifolds and its Riemannian structure, via the data information metric and its curvature at data points. Using this framework and through experiments, explanations on the response of a neural network are given by pointing out the output classes that are easily reachable from a given input. This emphasizes how the proposed mathematical relationship between the output of the network and the geometry of its inputs can be exploited as an explainable artificial intelligence tool.

Keywords: Neural Networks, Data Manifolds, Moving Frames, Curvature, Explainable AI

1 Introduction

In machine learning, the idea of data manifold is based on the assumption that the data space, containing the data points on which we perform classification tasks, has a natural Riemannian manifold structure. It is a quite old concept (see [1], [2] and refs therein) and it is linked to the key question of dimensionality reduction [3], which is the

key to efficient data processing for classification tasks, and more. Several geometrical tools, such as geodesics, connections, Ricci curvature, become readily available for machine learning problems under the manifold hypothesis, and are especially effective when employed in examining adversarial attacks [4–8] or knowledge transfer questions, see [9, 10] and refs. therein.

More specifically, the rapidly evolving discipline of information geometry [11–14] is now offering methods for discussing the above questions, once we cast them appropriately by relating the statistical manifold, i.e. the manifold of probability measures, studied in information geometry with data manifolds (see [11, 15]). Data manifolds have also been studied from a Riemannian geometric perspective by pullbacks of the Fisher information metric on the probability simplex through simple families of functions in the field of metric learning [16, 17].

In this paper, we are interested in a naturally emerging foliation structure on the data space coming through the *data information matrix* (DIM), which is the analog of the Fisher information matrix and in concrete experiments can be obtained by looking at a deep learning neural network for classification tasks. As it turns out, the leaves of such foliation are related with the dataset the network was trained with [15, 18]. Further work in [18] linked such study to the possible applications to knowledge transfer.

The purpose of the present paper is to study and understand the *data manifolds*, with Cartan moving frames method. Following the philosophy in [11], we want to equip the manifolds coming as leaves of the above mentioned foliation, via a natural metric coming through the data information matrix. As it turns out, the partial derivatives of the probabilities allow us to naturally define the Cartan moving frame at each point and are linked with the curvature of such manifold in an explicit way.

From a broader perspective, the work proposed here emphasizes the mathematical relationship between the changes in the outputs of a neural network and the curvature of the data manifolds. Furthermore, we show how this relationship can be exploited to provide explanations for the responses of a given neural network. In critical systems, providing explanations to AI model decisions can sometimes be required or mandatory by certification processes (see for example [19]). More generally, the field of eXplainable Artificial Intelligence (AI) is a fast growing research topic that develops tools for understanding and interpreting predictions given by AI models [20]. In Section 6 of this article, through simple experiments, we show how the DIM restricted to the moving frame given by the neural network output probability partial derivatives can be used to understand the local geometry of trained data. Specifically, in the case of the MNIST handwritten digits dataset [21] and CIFAR10 animals / vehicles dataset [22] by displaying the restricted DIM in the form of images, we are able to understand, starting from a given data point, which classes are easily reachable by the neural networks or not.

The organization of the paper is as follows.

In Section 2, we briefly recap some notions on information geometry and some key known results that we shall need in the sequel. Our main reference source will be [13, 14].

In Section 3, we take advantage of the machinery developed by Cartan (see [23, 24]) and define moving frames on the data manifolds via the partial derivatives of the probabilities.

In Section 4 and 5, we relate the probabilities with the curvature of the leaves, by deriving the calculations of the curvature forms in a numerically stable way.

Finally in Section 6, we consider some experiments on MNIST and CIFAR10 elucidating how near a data point, some partial derivatives of the probabilities become more important and the metric exhibits some possible singularities.

Acknowledgements. We thank Emanuele Latini, Sylvain Lavau for helpful discussions. This research was supported by Gnsaga-Indam, by COST Action CaLISTA CA21109, HORIZON-MSCA-2022-SE-01-01 CaLIGOLA, MSCA-DN CaLiForNIA - 101119552, PNRR MNESYS, PNRR National Center for HPC, Big Data and Quantum Computing, INFN Sezione Bologna. ET is grateful to the FaBiT Department of the University of Bologna for the hospitality.

2 Preliminaries

The *neural networks* we consider are classification functions

$$N : \mathcal{X} \times \Theta \longrightarrow \left\{ p \in \mathbb{R}^C ; \sum_k p_k = 1, p_k > 0 \right\},$$

from the spaces of data \mathcal{X} and weights (or parameters) Θ , to the space of parameterized probability densities $p(y | x, \theta)$ over the set of labels \mathcal{Y} , where $x \in \mathcal{X}$ is the input, $y \in \mathcal{Y}$ is the target label and $\theta \in \Theta$ is the parameter of the model. For instance θ are the weights and biases in a perceptron model predicting C classes, for a given input datum x .

We may assume both the dataspace \mathcal{X} and the parameter space Θ to be (open) subsets of euclidean spaces, $\mathcal{X} \subset \mathbb{R}^d$, $\Theta \subset \mathbb{R}^n$, though it is clear that in most practical situations only a tiny portion of such open sets will be effectively occupied by data points or parameters of an actual model.

In the following, we make only two assumptions on N :

$$N = \text{softmax} \circ s \tag{H1}$$

$$\forall i, j, k, \quad \partial_{x_i} \partial_{x_j} s_k(x) = 0, \quad \forall x \in \mathcal{X} \setminus A \tag{H2}$$

where s is called a *score* function, A is a nowhere dense set and softmax is the function $(a_i)_i \mapsto (e^{a_i} / \sum_k e^{a_k})_i$. We detail why these assumptions are important in the following sections.

Examples of neural networks that satisfies these assumptions are feed-forward multi-layer perceptrons, or convolutional neural networks, with ReLU activation functions or variants such as leaky ReLU, or any piecewise linear activation functions. These kind of architectures are popular within the machine learning community.

A most important notion in information geometry is the Fisher-Rao matrix F , providing in some important applications, a metric on the space Θ :

$$F = (f_{ij}) := (\mathbb{E}_{y|x,\theta} [\partial_{\theta_i} \ln p(y | x, \theta) \partial_{\theta_j} \ln p(y | x, \theta)]), \quad 1 \leq i, j \leq n.$$

We now give in analogy to F the *data information matrix* D (DIM for short), defined in a similar way.

Definition 2.1. We define *data information matrix* $D = (D_{ij})$ for a point $x \in \mathcal{X}$ and a fixed model $\theta \in \Theta$ to be the following symmetric matrix:

$$D_{ij}(x) = \mathbb{E}_{y|x,\theta} [\partial_{x_i} \ln p(y | x, \theta) \partial_{x_j} \ln p(y | x, \theta)], \quad 1 \leq i, j \leq d.$$

In the data information matrix, the derivative is taken with respect to the input x and not with respect to the parameter θ as in the Fisher matrix. Hence, D is a $d \times d$ matrix - d the dimension of the dataspace - which makes it possible the study of the relationship between input-output changes in the neural network.¹

Observation 2.2. In what follows, we will use the notation $p_k(x) := p(y_k | x, \theta)$ in the case of classification into a finite set $\mathcal{Y} = \{y_1, \dots, y_C\}$. We then concatenate these values in the vector $p(x) = (p_k(x))_{k=1, \dots, C}$. In practice, $p(x) = N(x)$.

We omit the dependence in θ as the parameters remain unchanged during the rest of the paper. We shall omit the dependence in x too in long computations for easier reading.

As one can readily check, we have that:

$$D_{ij}(x) = \sum_k \frac{1}{p_k(x)} \partial_{x_i} p_k(x) \partial_{x_j} p_k(x). \quad (1)$$

In the following, we will use ∂_i , the canonical basis of $T\mathcal{X}$ associated to the coordinates x^i , and the Einstein summation notation.

Remark 2.3. Notice the appearance of the probability p_i at the denominator in the expression of D . Since $p(y | x, \theta)$ is an empirical probability, on data points it may happen that some of the $p_k(x)$ are close to zero, giving numerical instability in practical situations. We shall comment on this problem and how we may solve it, later on.

We have the following result [15], we briefly recap its proof, for completeness.

Theorem 2.4. *The data information matrix $D(x)$ is positive semidefinite, moreover*

$$\ker D(x) = \left(\text{span}_{k=1, \dots, C} \left\{ \sum_i \partial_i \ln p_k(x) \partial_i \right\} \right)^\perp$$

with \perp taken w.r.t. the Euclidean scalar product on $T\mathcal{X}$. Hence, the rank of D is bounded by $C - 1$, with C the number of classes.

¹Note that D can be retrieved as the pullback of the Fisher metric on the $(C - 1)$ -simplex by the neural network N .

Proof. To check semipositive definiteness, let $u \in T_x \mathcal{X}$. Then,

$$\begin{aligned} u^T D(x)u &= \sum_{i,j} u_i u_j \mathbb{E}_{y|x,\theta} [\partial_i \ln p(y|x,\theta) \partial_j \ln p(y|x,\theta)] \\ &= \mathbb{E}_{y|x,\theta} \left[\sum_i u_i \partial_i \ln p(y|x,\theta) \sum_j u_j \partial_j \ln p(y|x,\theta) \right] \\ &= \mathbb{E}_{y|x,\theta} \left[\langle \partial_i \ln p(y|x,\theta) \partial_i, u \rangle_e^2 \right]. \end{aligned}$$

For the statement regarding kernel and rank of $D(x)$, we notice that $u^T D(x)u = 0$, whenever $u \in (\text{span}_{k=1,\dots,C} \{\sum_i \partial_i \ln p_k(x) \partial_i\})^\perp$, where $\langle \cdot, \cdot \rangle_e$ and \perp refer to the Euclidean metric. Besides, $\sum_k p_k = 1 \implies \sum_{ik} \partial_i p_k = 0$, hence the rank bounded by $C - 1$ and not C . \square

This result prompts us to define the distribution:

$$x \mapsto \mathcal{D}_x := \text{span} \left\{ \sum_i \partial_i p_k(x) \partial_i, k = 1, \dots, C - 1 \right\}. \quad (2)$$

For popular neural networks (satisfying [H1](#) and [H2](#)), the distribution \mathcal{D} turns to be integrable in an open set of \mathbb{R}^d , hence it defines a *foliation*. This matter has been discussed in [\[18\]](#).

Theorem 2.5. *Let θ be the weights of a neural network classifier N satisfying [H1](#) and [H2](#), associated with the vector p given by softmax. Assume D has constant rank. Then, at each smooth point x of N there exists a local submanifold \mathcal{L} of \mathcal{X} , such that its tangent space at x , $T_x \mathcal{L} = \mathcal{D}$.*

In particular, given a dataset (e.g. MNIST), we call *data leaf* a leaf of such foliation containing at least one point of the dataset. In [\[18\]](#) the significance of such leaves is fully explored from a geometrical and experimental point of view. More specifically, Thm. 3.6 in [\[18\]](#) shows that the open set in Thm 2.5 is dense in the dataspace \mathcal{X} .

In the next sections we shall focus on the geometry of this foliation in the data space.

3 Cartan moving frames

In this section we first review some basic notions of Cartan's approach to differentiable manifolds via the moving frames point of view and then we see some concrete application to our setting. A more detailed introduction to the Cartan moving frames approach can be found in [Appendix A](#).

Definition 3.1. Let $E \rightarrow \mathcal{L}$ be a \mathcal{C}^∞ vector bundle over a smooth manifold \mathcal{L} . A *connection* on E is a bilinear map:

$$\nabla : \mathfrak{X}(\mathcal{L}) \times \Gamma(E) \rightarrow \Gamma(E)$$

such that

$$\nabla_v(fs) = df(v)s + f\nabla_v s, \quad v \in \mathfrak{X}(\mathcal{L}), f \in \mathcal{C}^\infty(\mathcal{L}), s \in \Gamma(E)$$

where $\mathfrak{X}(\mathcal{L})$ are the vector fields on \mathcal{L} and $\Gamma(E)$ the sections of the vector bundle E .

We shall be especially interested to the case, $E = T\mathcal{L}$ the tangent bundle to a leaf \mathcal{L} of a suitable foliation (Thm. 2.5). Our framework is naturally set up to take advantage of Cartan's language of *moving frames* [23]. Before we give this key definition, let us assume our distribution \mathcal{D} as in (2) to be smooth, constant rank and integrable. This assumption is reasonable, since it has been shown in [18] that, for neural networks with ReLU non linearity, \mathcal{D} satisfies these hypotheses in a dense open subset of the dataspace.

Definition 3.2. We define the *data foliation* $\mathcal{F}_{\mathcal{D}}$, the foliation in the data space \mathcal{X} defined by the distribution \mathcal{D} as in Equation 2 and *data leaf* a leaf \mathcal{L}_x of \mathcal{F} containing at least one *data point* i.e. a point of the data set the network was trained with.

If $\langle \cdot, \cdot \rangle_D$ denotes the symmetric bilinear form defined by the data information matrix D as in Def. 2.1 in the dataspace, by the very definitions, when we restrict ourselves to the tangent space to a leaf \mathcal{L} of \mathcal{F} , we have that $\langle \cdot, \cdot \rangle_{\mathcal{L}} := \langle \cdot, \cdot \rangle_D|_{\mathcal{L}}$ is a non degenerate inner product (Prop. 2.4 and Thm. 2.5). Hence it defines a *metric* denoted g^D .

Definition 3.3. Let the notation be as above and let \mathcal{L} be a fixed leaf in \mathcal{F}_D . At each point $x \in \mathcal{L}$, we define $(e_k := \sum_i \partial_i p_k(x) \partial_i)_{(k=1, \dots, C-1)}$ to be a *frame* for $T_x \mathcal{L}$. The symmetric bilinear form g^D defines a Riemannian metric on \mathcal{L} , that we call *data metric*.

Let ∇ be the Levi-Civita connection with respect to the data information metric.

Definition 3.4. The Levi-Civita connection

$$\nabla_X e_j = \sum_i \omega_j^i(X) e_i \quad (3)$$

defines ω_j^i , which are called the *connection forms* and ω the *connection matrix* relative to the frame (e_i) .

The Levi-Civita connection is explicitly given by (see [23] pg 46):

$$2g^D(\nabla_{e_a} e_b, e_c) = e_a(g^D(e_b, e_c)) + e_b(g^D(e_c, e_a)) - e_c(g^D(e_a, e_b)) - g^D(e_a, [e_b, e_c]) + g^D(e_b, [e_c, e_a]) + g^D(e_c, [e_a, e_b]). \quad (4)$$

To explicitly compute the connection forms $\omega_j^i(e_k)$ we need to define

$$C_{a,b,c} := g^D(\nabla_{e_a} e_b, e_c) \quad (5)$$

then we have by definition that

$$C_{a,b,c} = \sum_i \omega_b^i(e_a) g^D(e_i, e_c).$$

Define the matrix $\hat{D} = (g^D(e_i, e_j))_{i,j=1,\dots,C-1}$. This is the matrix of the metric $g^D(\cdot, \cdot)$ restricted to a leaf \mathcal{L} for the basis given by a frame. We shall see an interesting significance for the matrix \hat{D} in the experiments in Section 6. Hence:

$$\sum_i \hat{D}_{l,i} \omega_j^i(e_k) = C_{k,j,l}.$$

This gives, in matrix notation:

$$\omega_j^i(e_k) = \left(\hat{D}^{-1} C_{k,j,\cdot} \right)_i. \quad (6)$$

The *curvature tensor* R is given by:

$$R(X, Y)e_j = \sum_i \Omega_j^i(X, Y)e_i$$

where Ω_j^i is a 2-form on $T\mathcal{X}$, alternating and \mathcal{D} -bilinear called the *curvature form (matrix)* of the connection ∇ relative to the frame (e_i) on $T\mathcal{X}$.

To compute R we shall make use of the following result found in [23] Theorem 11.1.

Proposition 3.5. *The curvature form Ω is given by:*

$$\Omega_j^i(X, Y) = (d\omega_j^i)(X, Y) + \sum_k \omega_k^i \wedge \omega_j^k(X, Y).$$

By using all the previous propositions and definitions, we will see in the following sections how we can numerically compute the curvature for a neural network with a softmax function on the output.

It is useful to recall some formulae. If α and β are \mathcal{C}^∞ 1-forms and X, Y are \mathcal{C}^∞ vector fields on a manifold, then

$$(\alpha \wedge \beta)(X, Y) = \alpha(X)\beta(Y) - \alpha(Y)\beta(X) \quad (7)$$

and

$$(d\alpha)(X, Y) = X\alpha(Y) - Y\alpha(X) - \alpha([X, Y]). \quad (8)$$

4 The Curvature of the Data Leaves

Computing in practice the curvature of a manifold with many dimensions is often intractable though essential for many tasks. This is why we are interested in computing

the curvature just for the data leaves. In this section and the following one, we derive the curvature form calculations to obtain a formula linking curvature and first-order derivative of the network (i.e. the frame e_k), without double automatic differentiation or division by p_i which might be close to zero, especially on data points, and could cause problems in practical implementation. Expressing the curvature with respect to the first-order derivative of the network can give information on the magnitude of the curvature coefficients when N is almost constant in some region of the dataspace.

Notice that, since in our frame ($e_k := \sum_i \partial_i p_k(x) \partial_i$), we have $\sum_i \partial_i \ln p_k(x) \partial_i = \frac{1}{p_k(x)} \sum_i \partial_i p_k(x) \partial_i = \frac{1}{p_k(x)} e_k$, we may forget the logarithm in the computations as it contributes only by a scalar factor.

Let $U = (\sum_k \partial_k p_i \partial_k p_j)_{i,j}$ be the matrix of the dot products of the partial derivatives of the probabilities.

With this notation, if $P = \text{diag}\{p_k, k = 1, \dots, C\}$, then

$$\hat{D}_{a,b} = (UP^{-1}U)_{a,b}.$$

Besides, if $J(p) = \left(\frac{\partial p_a}{\partial x_i} \right)_{\substack{a=1,\dots,C \\ i=1,\dots,d}}$ is the Jacobian matrix of first order derivatives, then $U = J(p)J(p)^T$.

Notice that P might not be numerically invertible, whenever some classes are very unlikely and thus with a probability close to zero. The goal of this section, and the following one, is thus to derive the computations of the curvature forms in a way that is numerically stable. This will allow us to implement the curvature forms computations on a computer.

Proposition 4.1. *Let the notation be as above. We have:*

$$[e_a, e_b] = H(p_b) e_a - H(p_a) e_b \quad (9)$$

with $H(f) = (\partial_i \partial_j f)_{i,j=1,\dots,d}$ the matrix of second order partial derivatives.

Proof. First, recall that $[u^k \partial_k, v^l \partial_l] = \left(u^k \frac{\partial v^l}{\partial x^k} - v^k \frac{\partial u^l}{\partial x^k} \right) \partial_l$. Thus, with $u^k = \frac{\partial p_a}{\partial x^k}$ and $v^l = \frac{\partial p_b}{\partial x^l}$ we get:

$$\begin{aligned} [e_a, e_b] &= \left(\sum_k \frac{\partial p_a}{\partial x^k} \frac{\partial^2 p_b}{\partial x^k \partial x^l} - \sum_k \frac{\partial p_b}{\partial x^k} \frac{\partial^2 p_a}{\partial x^k \partial x^l} \right) \partial_l \\ &= H(p_b) e_a - H(p_a) e_b. \end{aligned}$$

□

Proposition 4.2. *If s represent the score vector, i.e. the output of the neural network before going through the softmax function then, for all $j = 1, \dots, C$,*

$$\partial_j p_i = \sum_k p_i (\delta_{ik} - p_k) \partial_j s_k \quad \text{and} \quad \partial_j \ln p_i = \sum_k (\delta_{ik} - p_k) \partial_j s_k. \quad (10)$$

In term of Jacobian matrices, this rewrites as

$$J(p) = (P - pp^T) J(s).$$

Proof. This is simply due to the fact that $p = \text{softmax}(s)$ and that the derivative of the softmax function is $\frac{\partial \text{softmax}(x)_i}{\partial x_k} = \text{softmax}(x)_i (\delta_{i,k} - \text{softmax}(x)_k)$. \square

We now give other formulae for the second order partial derivatives of the probabilities.

Proposition 4.3. *Let the notation be as above. We have:*

$$H(p_a)_{ij} = \frac{\partial_i p_a \partial_j p_a}{p_a} - p_a \sum_{k=1}^C \partial_i p_k \partial_j s_k. \quad (11)$$

Proof.

$$\begin{aligned} H(p_a)_{ij} &= \partial_i (\partial_j p_a) = \partial_i \left(\sum_k p_a (\delta_{ik} - p_k) \partial_j s_k \right) \\ &= \partial_i p_a \underbrace{\sum_k (\delta_{ak} - p_k) \partial_j s_k}_{= \frac{1}{p_a} \partial_j p_a} + p_a \sum_k \partial_i ((\delta_{ak} - p_k) \partial_j s_k) \\ &= \frac{1}{p_a} \partial_i p_a \partial_j p_a + p_a \sum_k (\delta_{ak} - p_k) H(s_k) - p_a \sum_k \partial_i p_k \partial_j s_k. \end{aligned}$$

But $H(s_k) = 0$ almost everywhere by [H2](#). \square

However, with this form, the probability at the denominator will cause some instability problems, whenever the network is sufficiently trained. Thus, we express $H(p_a)$ in another form below.

Proposition 4.4. *Let the notation be as above. We have:*

$$\begin{aligned} H(p_a)_{ij} &= \sum_k [(\delta_{ak} - p_k) \partial_i p_a - p_a \partial_i p_k] \partial_j s_k \\ &= p_a \sum_k \left[(\delta_{ak} - p_k) \sum_l (\delta_{al} - p_l) \partial_i s_l - \partial_i p_k \right] \partial_j s_k. \end{aligned}$$

Proof. The proof is a straightforward combination of [Equation 10](#) and [Equation 11](#). We simply need to replace $\partial_i p_a$ twice in the expression [11](#) by [10](#). \square

Lemma 4.5.

$$e_a(D)(x) = J(s)^T A_a J(s) \quad (12)$$

with

$$(A_a)_{kl} = \sum_i \partial_i p_a ((\delta_{kl} - p_l) \partial_i p_k - p_k \partial_i p_l)$$

Proof.

$$\begin{aligned} e_a(D)(x) &= \sum_i \partial_i p_a \partial_i D(x) \\ &= \sum_i \partial_i p_a \partial_i (J(s)^T (P - pp^T) J(s)) \\ &= \sum_i \partial_i p_a J(s)^T \partial_i (P - pp^T) J(s) \quad \text{because } H(s_k) = 0. \end{aligned}$$

Indeed, $H(s_k) = 0$ for all k almost everywhere by [H2](#). Then at the indexes k, l , we get:

$$\begin{aligned} \partial_i (P - pp^T)_{kl} &= \partial_i (p_k \delta_{kl} - p_k p_l) \\ &= \delta_{kl} \partial_i p_k - p_l \partial_i p_k - p_k \partial_i p_l. \end{aligned}$$

Thus, by multiplying with $\partial_i p_a$ and summing over i , it gives the expression of A_a . \square

Proposition 4.6. *Recall that the neural network satisfies [H2](#), then*

$$e_a(g^D(e_b, e_c)) = M_{a,c,b} + M_{a,b,c} + (J(s)e_b)^T A_a J(s)e_c$$

with

$$M_{a,b,c} = (H(p_b) e_a)^T J(s)^T (P - pp^T) J(s) e_c.$$

Proof. The proof is straightforward.

$$\begin{aligned} e_a(g^D(e_b, e_c)) &= \sum_i \partial_i p_a \partial_i (e_b^T D(x) e_c) \\ &= \sum_i \partial_i p_a \partial_i (e_b^T) D(x) e_c + \sum_i \partial_i p_a e_b^T \partial_i (D(x)) e_c \\ &\quad + \sum_i \partial_i p_a e_b^T D(x) \partial_i e_c \\ &= e_a^T H(p_b)^T D(x) e_c + e_b^T e_a (D(x)) e_c + e_b^T D(x) H(p_c) e_a \\ &= M_{a,b,c} + e_b^T J(s)^T A_a J(s) e_c + \left(e_a^T H(p_c)^T D(x) e_b \right)^T \\ &= M_{a,b,c} + e_b^T J(s)^T A_a J(s) e_c + (M_{a,c,b})^T. \end{aligned}$$

And $M_{a,c,b}$ is a scalar, thus $M_{a,c,b}^T = M_{a,c,b}$, hence the result. \square

Then, we use the following proposition to remove the second order derivative in M . An alternative formula for this expression is given in [B.1](#). This facilitates the computations with automatic differentiation methods.

Proposition 4.7.

$$H(p_a) e_b = \left(\sum_{k,i} (\delta_{ak} - p_k) (\partial_i s_k \partial_i p_b) \right) e_a - \sum_{k,i} p_a (\partial_i s_k \partial_i p_b) e_k. \quad (13)$$

Proof.

$$\begin{aligned} H(p_a) e_b &= \sum_i \left(\sum_k ((\delta_{ak} - p_k) \partial_j p_a - p_a \partial_j p_k) \partial_i s_k \right) \partial_i p_b \\ &= \sum_i \left(\sum_k (\delta_{ak} - p_k) \partial_j p_a \partial_i s_k \right) \partial_i p_b - p_a \sum_i \left(\sum_k \partial_j p_k \partial_i s_k \right) \partial_i p_b \\ &= \left(\sum_{k,i} (\delta_{ak} - p_k) \partial_i s_k \partial_i p_b \right) \partial_j p_a - p_a \sum_{k,i} (\partial_i s_k \partial_i p_b) \partial_j p_k. \end{aligned}$$

□

Proposition 4.8.

$$g^D(e_a, [e_b, e_c]) = \sum_{i,j,k} \partial_i p_a D(x)_{ij} (\partial_j \partial_k p_c \partial_k p_b - \partial_j \partial_k p_b \partial_k p_c). \quad (14)$$

Proof. Straightforward from the previous propositions and lemmas. □

With all these propositions, the connection forms can be computed directly without numerical instabilities.

5 Computation of the curvature forms

In this section we conclude our calculation of the curvature forms. In [Prop. 3.5](#) we wrote the explicit expression for the curvature form as:

$$\Omega(X, Y) = (d\omega)(X, Y) + \omega \wedge \omega(X, Y). \quad (15)$$

where ω denotes the (Levi-Civita) connection form. To ease the reading we go back to notation of [Section 3](#) and we set:

$$e_k := \sum_i \partial_i p_k(x) \partial_i, \quad \text{for } k = 1, \dots, C-1.$$

We then can express explicitly the connection form as:

$$\nabla_X e_j = \sum \omega_j^i(X) e_i.$$

The wedge product of the connection forms in (15) can thus be easily computed with the propositions of the previous section, because of formula (7) report here:

$$(\omega \wedge \omega)(X, Y) = \omega(X)\omega(Y) - \omega(Y)\omega(X). \quad (16)$$

The exterior derivative of ω remains to be computed and it is more complicated. We recall the formula (8):

$$(d\omega)(X, Y) = X\omega(Y) - Y\omega(X) - \omega([X, Y]). \quad (17)$$

The last term is computed via the following proposition.

Proposition 5.1. *Let the notation be as above. Then:*

$$\begin{aligned} \omega_j^i([e_a, e_b]) &= \left(\sum_{k,l} (\delta_{bk} - p_k) (\partial_l s_k \partial_l p_a) \right) \omega_j^i(e_b) \\ &\quad - \left(\sum_{k,l} (\delta_{ak} - p_k) (\partial_l s_k \partial_l p_b) \right) \omega_j^i(e_a) \\ &\quad - \sum_{k,l} (\partial_l s_k (p_b \partial_l p_a - p_a \partial_l p_b)) \omega_j^i(e_k). \end{aligned}$$

Proof.

$$\omega_j^i([e_a, e_b]) = \omega_j^i(H(p_b) e_a) - \omega_j^i(H(p_a) e_b).$$

Besides,

$$\begin{aligned} \omega_j^i(H(p_a) e_b) &= \omega_j^i \left(\left(\sum_{k,l} (\delta_{ak} - p_k) (\partial_l s_k \partial_l p_b) \right) e_a - \sum_{k,l} p_a (\partial_l s_k \partial_l p_b) e_k \right) \\ &= \left(\sum_{k,l} (\delta_{ak} - p_k) (\partial_l s_k \partial_l p_b) \right) \omega_j^i(e_a) - \sum_{k,l} p_a (\partial_l s_k \partial_l p_b) \omega_j^i(e_k). \end{aligned}$$

Thus,

$$\omega_j^i([e_a, e_b]) = \left(\sum_{k,l} (\delta_{bk} - p_k) (\partial_l s_k \partial_l p_a) \right) \omega_j^i(e_b)$$

$$\begin{aligned}
& - \left(\sum_{k,l} (\delta_{ak} - p_k) (\partial_l s_k \partial_l p_b) \right) \omega_j^i(e_a) \\
& - \sum_{k,l} (\partial_l s_k (p_b \partial_l p_a - p_a \partial_l p_b)) \omega_j^i(e_k).
\end{aligned}$$

□

We now tackle the question of determining the first two terms in (8).

Observation 5.2. We notice that to compute $e_a(\omega_j^i(e_b))$, we can use the fact that:

$$\begin{aligned}
e_a(g^D(\nabla_{e_b} e_c, e_d)) &= e_a \left(\sum_i \omega_c^i(e_b) g^D(e_i, e_d) \right) \\
&= \sum_i e_a(\omega_c^i(e_b)) g^D(e_i, e_d) + \sum_i \omega_c^i(e_b) e_a(g^D(e_i, e_d)) \\
\iff \sum_i \hat{D}_{d,i} e_a(\omega_c^i(e_b)) &= e_a(g^D(\nabla_{e_b} e_c, e_d)) - \underbrace{\sum_i \omega_c^i(e_b) e_a(g^D(e_i, e_d))}_{N_{a,b,c,d}} \\
\iff e_a(\omega_c^i(e_b)) &= \hat{D}^{-1} N_{a,b,c,\cdot}
\end{aligned}$$

Hence we need to compute: $e_a g^D(\nabla_{e_b} e_c, e_d)$.

Theorem 5.3. *The expression of the curvature is given by:*

$$\begin{aligned}
\Omega_j^i(e_a, e_b) &= \left(\hat{D}^{-1} N_{a,b,j,\cdot} \right)_i - \left(\hat{D}^{-1} N_{b,a,j,\cdot} \right)_i - \omega_j^i([e_a, e_b]) \\
&\quad + \sum_k [\omega_k^i(e_a) \omega_j^k(e_b) - \omega_k^i(e_b) \omega_j^k(e_a)] \quad (18)
\end{aligned}$$

where $N_{a,b,c,\cdot} = (N_{a,b,c,d})_{d=1,\dots,C-1}$ is the vector defined in Observation 5.2 by $e_a(g^D(\nabla_{e_b} e_c, e_d)) - \sum_i \omega_c^i(e_b) e_a(g^D(e_i, e_d))$ for $d = 1, \dots, C-1$, and where $\hat{D} = (g^D(e_i, e_j))_{i,j=1,\dots,C-1}$ is the matrix of pairwise metric products of the frame.

We report the lemmas needed to compute the tensor $N_{a,b,c,d}$ in Appendix C.

Remark 5.4. The practical computation of the curvature forms involves some matrix multiplication, inversion and the computation of the first order derivatives of the network which can be done with automatic differentiation [25]. The most time consuming operation is the differentiation of the network as it scales with the number of layers (thus the complexity of the task to learn), with the input dimension and with the number of classes. Improving the efficiency of this computation is left as future work.

6 Experiments

To understand why the frame $e_a = \sum_i \partial_i p_a \partial_i$ and the DIM were chosen to compute the connection and curvature forms, we shall focus on experiments on the MNIST

and the CIFAR10 datasets. As we will see, this frame and the DIM can provide some explanations to the response of the neural network.²

Note that in this section, we focus on the DIM and the frame $(e_a)_{a=1}^{C-1}$. We do not inspect the connection or curvature forms and leave them for future work.

The MNIST dataset is composed of 60k train images of shape 28×28 pixels depicting handwritten digits between 0 and 9, and the CIFAR10 dataset is composed of 60k RGB train images of shape 32×32 classified in 10 different classes (see Table 1).

Table 1: Correspondence index - class of the CIFAR10 dataset.

No.	0	1	2	3	4	5	6	7	8	9
Class	<i>airplane</i>	<i>automobile</i>	<i>bird</i>	<i>cat</i>	<i>deer</i>	<i>dog</i>	<i>frog</i>	<i>horse</i>	<i>ship</i>	<i>truck</i>

On Figure 1 are represented various input points x from the MNIST training set and the corresponding matrix³ $\hat{D}_x = (g_x(e_a, e_b))_{a,b=1,\dots,C}$. The neural network used in this experiment is defined as stated in Table 2. This network has been trained on the test set of MNIST with stochastic gradient descent until convergence (98% accuracy on the test set).

On Figure 1, it can be seen that the matrices \hat{D}_x have only a few main components indicating which probabilities are the easiest to change. A large (positive) component on the diagonal at index i suggests that one can increase or decrease easily $p_i(x)$ by moving in the directions $\pm e_i$. A negative (resp. positive) component at position (i, j) indicates that, starting from the image x , classes i and j are in opposite (resp. the same) directions: increasing $p_i(x)$ will most likely decrease (resp. increase) $p_j(x)$.

For instance, the first image on the top left is correctly classified as a 2 by the network, but since the coefficient $(3, 3)$ of matrix \hat{D} is positive too, it indicates that class 3 should be easily reachable. This makes sense, as the picture can also be interpreted as a part of the 3 digit. Negative coefficients at $(3, 2)$ and $(2, 3)$ thus indicates that going in the direction of a 3 will decrease the probability of predicting a 2. On the second picture, the same phenomenon arises but with classes 2 and 8. Indeed, the buckle in the bottom part of the picture brings it closer to an 8.

Remark 6.1. Be careful as teal colored coefficients on Figure 1 and Figure 2 are not exactly zero but rather very low compared to the few main ones that are yellow and purple.

On Figure 2 are represented various input points x from the CIFAR10 training set and the corresponding matrix $\hat{D}_x = (g_x(e_a, e_b))_{a,b=1,\dots,C}$. The neural network used in this experiment is defined as stated in Table 3. This network has been trained on the test set of CIFAR10 with stochastic gradient descent until convergence (84% accuracy on the test set).

²The code used to produce these results is available at <https://github.com/eliot-tron/curvcomputem/>.

³We plot here the matrix \hat{D} with a, b going up to C , and not $C - 1$, to represent all the classes and have an easier interpretation.

Table 2: Architecture of the neural network trained on MNIST.

No.	Layers (sequential)
(0):	<code>Conv2d(1, 32, kernel_size=(3, 3), stride=(1, 1))</code>
(1):	<code>ReLU()</code>
(2):	<code>Conv2d(32, 64, kernel_size=(3, 3), stride=(1, 1))</code>
(3):	<code>ReLU()</code>
(4):	<code>MaxPool2d(kernel_size=2, stride=2, padding=0, dilation=1)</code>
(5):	<code>Flatten()</code>
(6):	<code>Linear(in_features=9216, out_features=128, bias=True)</code>
(7):	<code>ReLU()</code>
(8):	<code>Linear(in_features=128, out_features=10, bias=True)</code>
(9):	<code>Softmax()</code>

On Figure 2, it can be seen that the matrices \hat{D}_x have only a few main components indicating which probabilities are the easiest to change, i.e. with a similar behavior as seen above for MNIST. The interpretation of matrix \hat{D} then is identical.

For instance, the first image on the top left is correctly classified as a dog (class No. 5) by the network, but since the coefficient $(3, 3)$ of matrix \hat{D} is positive too, it indicates that class No. 3 “cat” should be easily reachable. This makes sense since dogs and cats form a subclass of similar little animals. Negative coefficients at $(3, 2)$ and $(2, 3)$ thus indicates that going in the direction of the cat will decrease the probability of predicting a dog. There are also positive coefficients at $(3, 7)$ and $(7, 3)$ indicating that going in the direction of the cat should also slightly increase the probability of seeing a horse. Again, this makes sense as they all belong in the animal subclass.

On the second picture, the dog can be transformed into a cat or a frog, probably because of the green cloth around its neck.

More experiments can be found in Appendix D with full details to ensure reproducibility.

7 Conclusion

In this study, we have shown that analyzing the geometry of the data manifold of a neural network using Cartan moving frames is natural and practical. Indeed, from the Fisher Information Matrix, we define a corresponding Data Information Matrix that generates a foliation structure in the data space. The partial derivatives of the probabilities learned by the neural network define a Cartan moving frame on the leaves of the data manifold. We detail how the moving frame can be used as a tool to provide some explanations on the classification changes that can easily happen or not around a given data point. Experiments on the MNIST and CIFAR datasets confirm the relevance of the explanations provided by the Cartan moving frames and the corresponding Date Information Matrix. For very large neural network, the theory still holds. However, the method might be limited by the computational requirements of the partial derivatives calculation for each class (usually obtained through automatic differentiation). We believe that combining the moving frame, the connection and the

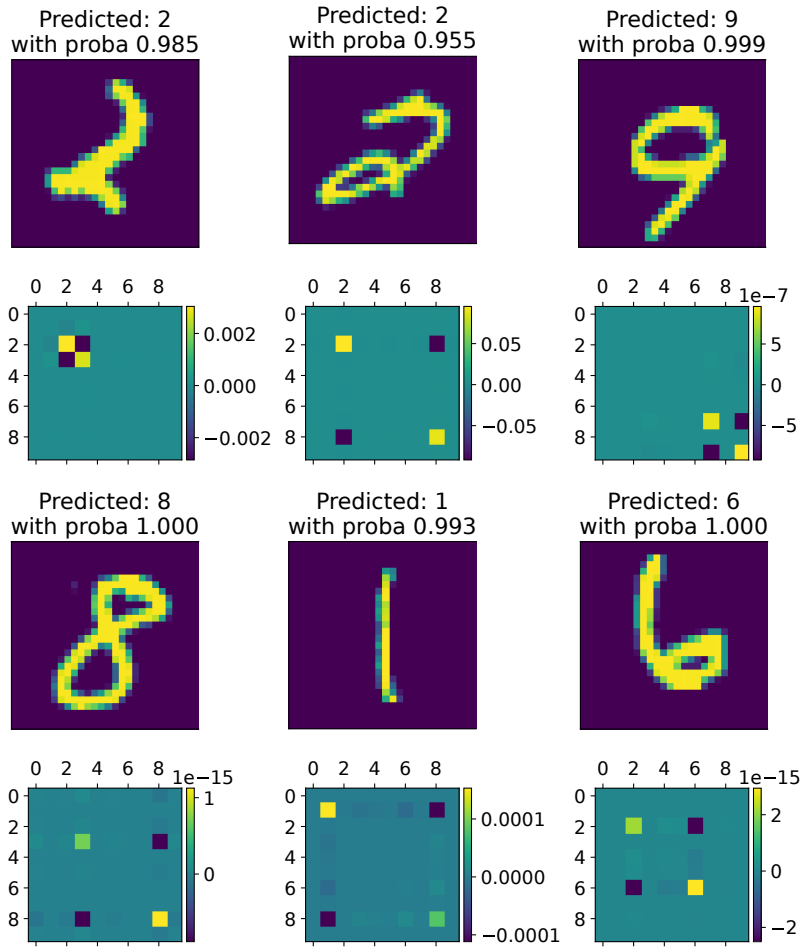


Fig. 1: Couples of input point x (above) and the corresponding matrix $\hat{D}(x) = (g_x^D(e_a, e_b))_{a,b=1,\dots,C}$ (below) on MNIST.

curvature forms should provide new insights to build more advanced explainable AI tools. This is work in progress.

Table 3: Architecture of the neural network trained on CIFAR10.

No.	Layers (sequential)
(0):	Conv2d(3, 64, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1))
(1):	BatchNorm2d(64, eps=1e-05, momentum=0.1)
(2):	ReLU()
(3):	MaxPool2d(kernel_size=2, stride=2, padding=0, dilation=1)
(4):	Conv2d(64, 128, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1))
(5):	BatchNorm2d(128, eps=1e-05, momentum=0.1)
(6):	ReLU()
(7):	MaxPool2d(kernel_size=2, stride=2, padding=0, dilation=1)
(8):	Conv2d(128, 256, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1))
(9):	BatchNorm2d(256, eps=1e-05, momentum=0.1)
(10):	ReLU()
(11):	Conv2d(256, 256, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1))
(12):	BatchNorm2d(256, eps=1e-05, momentum=0.1)
(13):	ReLU()
(14):	MaxPool2d(kernel_size=2, stride=2, padding=0, dilation=1, ceil mode=False)
(15):	Conv2d(256, 512, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1))
(16):	BatchNorm2d(512, eps=1e-05, momentum=0.1)
(17):	ReLU()
(18):	Conv2d(512, 512, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1))
(19):	BatchNorm2d(512, eps=1e-05, momentum=0.1)
(20):	ReLU()
(21):	MaxPool2d(kernel_size=2, stride=2, padding=0, dilation=1, ceil mode=False)
(22):	Conv2d(512, 512, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1))
(23):	BatchNorm2d(512, eps=1e-05, momentum=0.1)
(24):	ReLU()
(25):	Conv2d(512, 512, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1))
(26):	BatchNorm2d(512, eps=1e-05, momentum=0.1)
(27):	ReLU()
(28):	MaxPool2d(kernel_size=2, stride=2, padding=0, dilation=1, ceil mode=False)
(29):	AvgPool2d(kernel_size=1, stride=1, padding=0)
(30):	Linear(in_features=512, out_features=10, bias=True)
(31):	Softmax(dim=1)

Data availability

Data sets used in this article can be found online at:

- MNIST dataset [21]: <https://yann.lecun.com/exdb/mnist/>
- CIFAR10 dataset [22]: <https://www.cs.toronto.edu/~kriz/cifar.html>
- Code: <https://github.com/eliot-tron/curvcomputenn/>

We used Python and PyTorch to access these data.

References

- [1] Fefferman, C., Mitter, S., Narayanan, H.: Testing the manifold hypothesis. *J. Amer. Math. Soc.* **29**, 983–1049 (2016)
- [2] Roweis, S.T., Saul, L.K.: Nonlinear dimensionality reduction by locally linear embedding. *Science* **290** 5500, 2323–6 (2000)

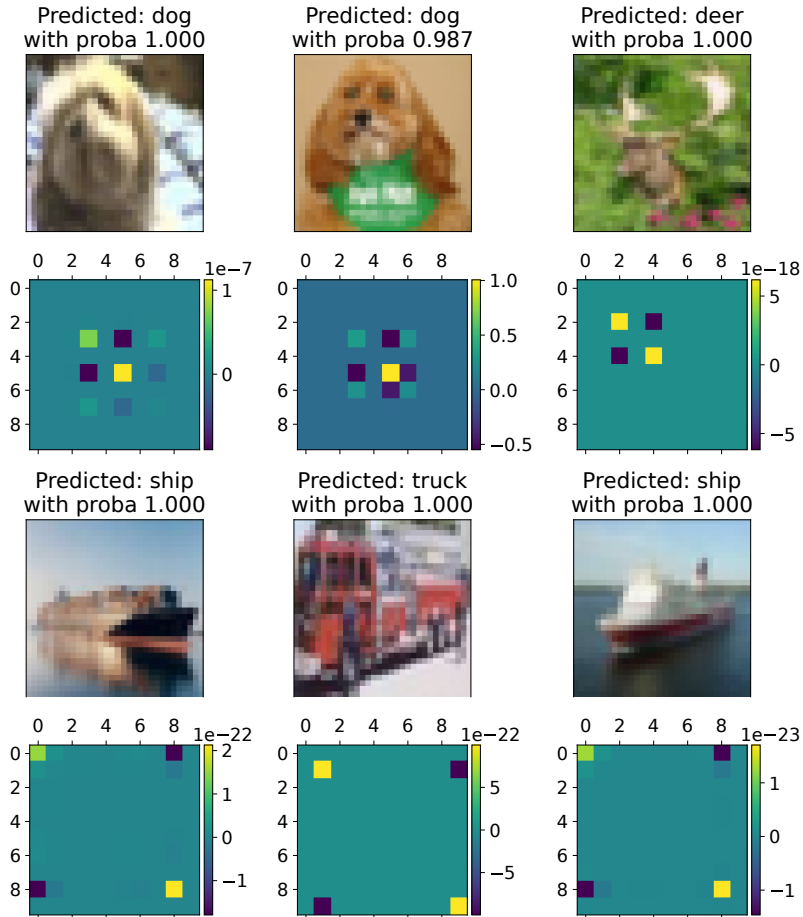


Fig. 2: Couples of input point x (above) and the corresponding matrix $\hat{D}(x) = (g_x^D(e_a, e_b))_{a,b=1,\dots,C}$ (below) on CIFAR10.

- [3] Maaten, L., Hinton, G.E.: Visualizing data using t-sne. *Journal of Machine Learning Research* **9**, 2579–2605 (2008)
- [4] Shen, C., Peng, Y., Zhang, G., Fan, J.: Defending Against Adversarial Attacks by Suppressing the Largest Eigenvalue of Fisher Information Matrix (2019) [arXiv:1909.06137](#) [cs, stat]
- [5] Martin, J., Elster, C.: Inspecting adversarial examples using the Fisher information (2019) [arXiv:1909.05527](#) [cs, stat]
- [6] Carlini, N., Wagner, D.: Adversarial Examples Are Not Easily Detected: Bypassing Ten Detection Methods [arXiv:1705.07263](#) [cs]

- [7] Zhao, C., Fletcher, P.T., Yu, M., Peng, Y., Zhang, G., Shen, C.: The Adversarial Attack and Detection under the Fisher Information Metric **33**, 5869–5876 (2019) <https://doi.org/10.1609/aaai.v33i01.33015869>
- [8] Tron, E., Couellan, N.P., Puechmorel, S.: Adversarial attacks on neural networks through canonical riemannian foliations. To appear in *Machine Learning*. Preprint available at: ArXiv **abs/2203.00922** (2022)
- [9] Weiss, K.R., Khoshgoftaar, T.M., Wang, D.: A survey of transfer learning. *Journal of Big Data* **3** (2016)
- [10] Pan, S.J., Yang, Q.: A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering* **22**, 1345–1359 (2010)
- [11] Sun, K., Marchand-Maillet, S.: An information geometry of statistical manifold learning. In: *International Conference on Machine Learning* (2014). <https://api.semanticscholar.org/CorpusID:7149130>
- [12] Amari, S.-I., Barndorff-Nielsen, O.E., Kass, R.E., Lauritzen, S.L., Rao, C.R.: Differential Geometry in Statistical Inference **10**, 1–17192195971611631652172192404355557
- [13] Amari, S.-i.: *Information Geometry and Its Applications* vol. 194. Springer, Tokyo (2016). <https://doi.org/10.1007/978-4-431-55978-8>
- [14] Nielsen, F.: An elementary introduction to information geometry. *Entropy* **22**(10), 1100 (2020) <https://doi.org/10.3390/e22101100>
- [15] Gremontieri, L., Fioresi, R.: Model-centric data manifold: The data through the eyes of the model. *SIAM Journal on Imaging Sciences* **15**(3), 1140–1156 (2022) <https://doi.org/10.1137/21M1437056>
- [16] Lebanon, G.: Learning riemannian metrics (arXiv:1212.2474) (2012) <https://doi.org/10.48550/arXiv.1212.2474> . arXiv:1212.2474
- [17] Wang, J., Sun, K., Sha, F., Marchand-Maillet, S., Kalousis, A.: Two-stage metric learning
- [18] Tron, E., Fioresi, R.: *Manifold Learning via Foliations and Knowledge Transfer* (2024). <https://arxiv.org/abs/2409.07412>
- [19] *Artificial Intelligence in Aeronautical Systems: Statement of Concerns*, SAE AIR6988 (2021). <https://www.sae.org/standards/content/air6988/>
- [20] Samek, W., Montavon, G., Vedaldi, A., Hansen, L.K., Muller, K.-R.: *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning* vol. 11700, 1st edn. Springer, Switzerland (2019)

- [21] LeCun, Y.: The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/> (1998)
- [22] Krizhevsky, A., Nair, V., Hinton, G.: CIFAR-10 (Canadian Institute for Advanced Research). University of Toronto, ON, Canada (2009)
- [23] Tu, L.W.: Differential Geometry: Connections, Curvature, and Characteristic Classes vol. 275. Springer, Switzerland (2017)
- [24] Jost, J.: Riemannian geometry and geometric analysis, 5th edition. (2008). <https://api.semanticscholar.org/CorpusID:206736101>
- [25] Baydin, A.G., Pearlmutter, B.A., Radul, A.A., Siskind, J.M.: Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research* **18**(153), 1–43 (2018)
- [26] Petersen, P.: Riemannian Geometry vol. 171. Springer, Switzerland (1998)
- [27] Carroll, S.: Lecture notes on general relativity. arXiv: General Relativity and Quantum Cosmology (1997)

A Cartan moving frames

In this appendix we briefly review the notion of moving frames for differentiable manifolds, highlighting the key geometric properties, with some pictures to help the intuition. For more details see [23] and [26].

The theory of Cartan moving frames (*repère mobile*) stems from the Frenet–Serret frame theory, which lead to the celebrated Frenet-Serret formulas, developed independently by Frenet and Serret, in the XIX century, to describe the motion of a particle on a curve in \mathbb{R}^3 . In fact, given smooth curve γ in \mathbb{R}^3 , a Frenet-Serret frame is an orthonormal basis of \mathbb{R}^3 attached at each point of γ and consisting of the (normalized) tangent, normal and conormal vectors (see Fig. 3), obtained via subsequent derivations of the curve γ .

Notice that, once we give the curve γ , the definition of its Frenet-Serret frame is independent from the coordinates of \mathbb{R}^3 ; it is in fact an “intrinsic” reference frame for \mathbb{R}^3 , associated to each point of the curve.

The notion of moving frames is a far reaching generalization of Frenet-Serret frames, where we replace the smooth curve γ , with a smooth manifold M .

Definition A.1. We define a *moving frame* on a differentiable manifold M a collection of vector fields e_1, e_2, \dots, e_n , which are a basis of $T_p M$, the tangent space at each point $p \in M$. (see Fig. 4). We also ask such collection to be varying smoothly with the point p in M . The assignment of a moving frame at each point $p \in M$ is called a *soldering form* $\theta : T_p M \rightarrow \mathbb{R}^n$.

Hence, once a soldering form is given, we have an intrinsic point of view on the geometry of a manifold and we can, in some sense, effectively replace the notion of

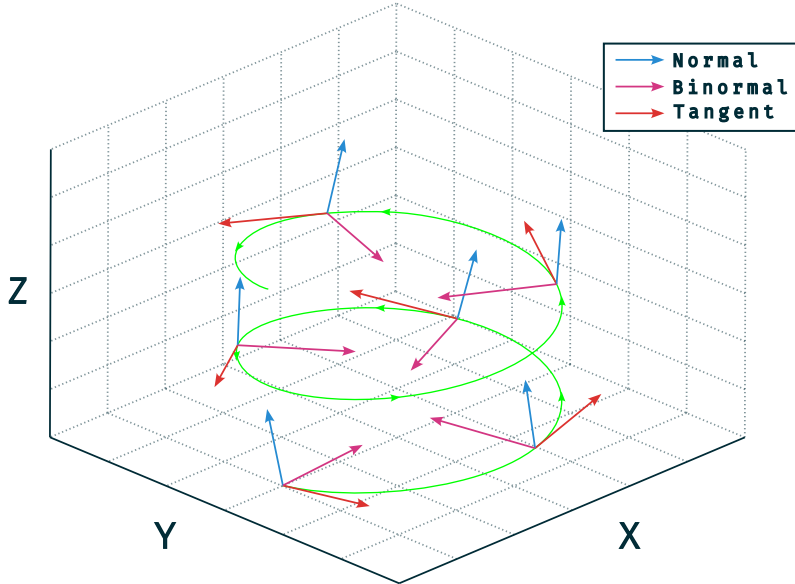


Fig. 3: Frenet-Serret Frames

local coordinates. Such point of view was successfully exploited to describe, in a local coordinate free manner, some physical theories, most notably *general relativity*, see [27].

In Riemannian geometry, we take the moving frames to be orthonormal frames, that is, at each point $p \in M$ the frame is an orthonormal basis of T_pM , with respect to the metric on the manifold.

In our application, that is the geometry of the data manifold \mathcal{L} , such moving frame comes naturally from the neural network and it is given by the gradients of the probability output, so that we immediately have a soldering form on \mathcal{L} as follows:

$$\theta : T_p\mathcal{L} \longrightarrow \mathbb{R}^{C-1}, \quad e_i \mapsto \nabla p_i \quad (19)$$

Moreover we have a natural Riemannian metric on the data leaf via the DIM (data information matrix). The soldering form defines uniquely (Prop. 1.1 pg 378 in [26]) a matrix of one forms $\omega = (\omega_j^i)$, called (evocatively) *spin connection*. So we can actually view the formula in Def. 3.4 in two equivalent ways:

1. Given our soldering form θ on the Riemannian manifold \mathcal{L} as in (19), we have a spin connection ω which gives the Levi Civita connection by the formula (see (3) in the main text):

$$\nabla_X e_j = \sum_i \omega_j^i(X) e_i$$

2. Given the metric via the DIM on the data leaf, we can compute its (unique) Levi Civita connection and this will give us ω the spin connection, via the above formula.

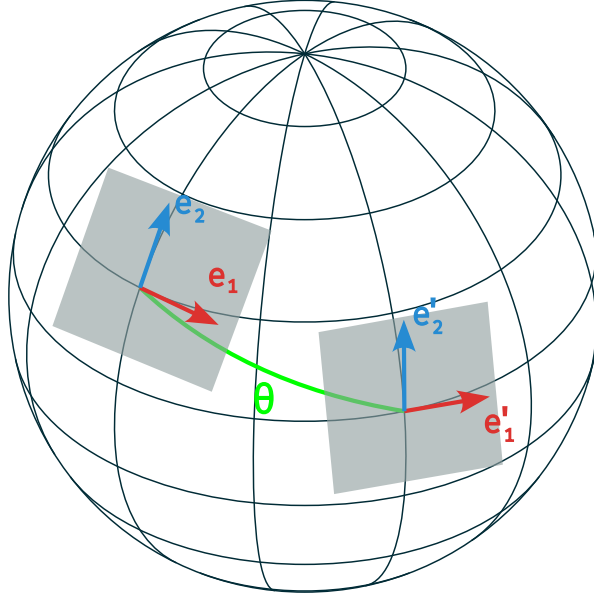


Fig. 4: Soldering form on a sphere.

We choose to take the second point of view in our paper, to avoid introducing too much terminology and results in moving frames theory, since we prefer to rely on the correspondence between Levi-Civita connections and metrics, possibly more familiar to most readers.

These equivalent ways to view the Levi Civita connection is an illustration of the power of the moving frames method as a coordinate independent way to approach Riemannian geometry. We can indeed compute the curvature form as:

$$\Omega(X, Y) = (d\omega)(X, Y) + \omega \wedge \omega(X, Y). \quad (20)$$

as we reported in (15) in the main text. Such curvature form is linked with the usual Riemannian curvature tensor via the equation:

$$R(X, Y)e_j = \sum \Omega_j^i(X, Y)e_i$$

(see [26] pg 379).

For more details on the method of moving frames and its link with the equivalent theory of Riemannian geometry making use of local coordinates, we invite the reader to consult [26] App. B.

B An auxiliary lemma

We report here, for convenience, an alternative formula for the expression used in 4.6.

Proposition B.1. Recall that $U = (u_{i,j}) = (\langle e_i, e_j \rangle_e)_{i,j}$. Then

$$\begin{aligned}
e_a (g^D (e_b, e_c)) &= \sum_j u_{aj} u_{bj} u_{cj} \frac{1}{p_j^2} \\
&\quad + \left(\frac{u_{ab}}{p_b} + \frac{u_{ac}}{p_c} \right) \sum_j \frac{1}{p_j} u_{bj} u_{cj} \\
&\quad - \sum_j (u_{cj} p_b + u_{bj} p_c) \frac{1}{p_j} \sum_{k,i} (\partial_i p_a \partial_i s_k) u_{jk} \\
&\quad - \sum_i (\partial_i p_a \partial_i s_k) (u_{cj} u_{bk} + u_{bj} u_{ck}).
\end{aligned}$$

Notice that on the second line, we recognise the expression of $g^D (e_b, e_c)$ for the sum.

Proof.

$$\begin{aligned}
e_a (g^D (e_b, e_c)) &= \sum_i \partial_i p_a \partial_i (g^D (e_b, e_c)) \\
&= \sum_i \partial_i p_a \partial_i \left(\sum_j \frac{1}{p_j} u_{bj} u_{cj} \right) \\
&= \sum_{i,j} \partial_i p_a \left(\left(\partial_i \frac{1}{p_j} \right) u_{cj} u_{bj} + \frac{1}{p_j} (\partial_i u_{bj}) u_{cj} + \frac{1}{p_j} u_{bj} (\partial_i u_{cj}) \right).
\end{aligned}$$

Now we compute $\partial_i u_{bj}$:

$$\begin{aligned}
\partial_i u_{bj} &= \partial_i \left(\sum_l \partial_l p_b \partial_l p_j \right) \\
&= \sum_l \underbrace{(\partial_i \partial_l p_b)}_{=H(p_b)_{i,l}} \partial_l p_j + \sum_l \partial_l p_b (\partial_i \partial_l p_j) \\
&= \sum_l \left(\frac{\partial_i p_b \partial_l p_b}{p_b} - p_b \sum_k \partial_i s_k \partial_l p_k \right) \partial_l p_j \\
&\quad + \sum_l \partial_l p_b \left(\frac{\partial_i p_j \partial_l p_j}{p_j} - p_j \sum_k \partial_i s_k \partial_l p_k \right) \\
&= \sum_l \left(\frac{\partial_i p_b}{p_b} + \frac{\partial_i p_j}{p_j} \right) u_{aj} - \sum_l (p_a \partial_l p_j + p_j \partial_l p_a) \sum_k \partial_i s_k \partial_l p_k.
\end{aligned}$$

Thus

$$\begin{aligned}
e_a(g^D(e_b, e_c)) &= \sum_j u_{bj} u_{cj} u_{aj} \frac{-1}{p_j^2} \\
&+ \sum_j \left(\frac{u_{ab}}{p_b} + \frac{u_{aj}}{p_j} \right) \frac{1}{p_j} u_{bj} u_{cj} \\
&- \sum_j \frac{1}{p_j} u_{cj} \left(p_b \sum_{k,l} (\partial_l p_a \partial_l s_k) u_{jk} + p_j \sum_{k,l} (\partial_l p_a \partial_l s_k) u_{bk} \right) \\
&+ \sum_j \left(\frac{u_{ac}}{p_c} + \frac{u_{aj}}{p_j} \right) \frac{1}{p_j} u_{bj} u_{cj} \\
&- \sum_j \frac{1}{p_j} u_{bj} \left(p_c \sum_{k,l} (\partial_l p_a \partial_l s_k) u_{jk} + p_j \sum_{k,l} (\partial_l p_a \partial_l s_k) u_{ck} \right).
\end{aligned}$$

□

To get a numerically stable expression for $e_a g^D(e_b, e_c)$, we can develop u with [Equation 10](#) and the probability at the denominator will cancel out with the one in factor of [Equation 10](#).

C Further computations of the curvature forms

In this section we report, for convenience, some lemmas necessary for the full derivation of the curvature forms calculations.

Lemma C.1.

$$e_a(e_b(g^D(e_c, e_d))) = e_a(M_{b,d,c}) + e_a(M_{b,c,d}) + e_a(e_c^T J(s)^T A_b J(s) e_d). \quad (21)$$

Proof.

$$\begin{aligned}
e_a(e_b(g^D(e_c, e_d))) &= e_a(M_{b,d,c} + M_{b,c,d} + e_c^T J(s)^T A_b J(s) e_d) \\
&= e_a(M_{b,d,c}) + e_a(M_{b,c,d}) + e_a(e_c^T J(s)^T A_b J(s) e_d).
\end{aligned}$$

□

Lemma C.2.

$$\begin{aligned}
e_a(M_{b,c,d}) &= e_a(H(p_c) e_b)^T D(x) e_d \\
&+ (H(p_c) e_b)^T J(s)^T A_a J(s) e_d \\
&+ (H(p_c) e_b)^T D(x) (H(p_d) e_a). \quad (22)
\end{aligned}$$

Proof.

$$\begin{aligned}
e_a(M_{b,c,d}) &= e_a \left(e_b^T H(p_c)^T D(x) e_d \right) \\
&= e_a \left(H(p_c) e_b \right)^T D(x) e_d \\
&\quad + \left(H(p_c) e_b \right)^T e_a \left(D(x) \right) e_d \\
&\quad + \left(H(p_c) e_b \right)^T D(x) e_a \left(e_d \right) \\
&= e_a \left(H(p_c) e_b \right)^T D(x) e_d \\
&\quad + \left(H(p_c) e_b \right)^T J(s)^T A_a J(s) e_d \\
&\quad + \left(H(p_c) e_b \right)^T D(x) \left(H(p_d) e_a \right).
\end{aligned}$$

□

Lemma C.3.

$$\begin{aligned}
e_a(H(p_c) e_b) &= \sum_k \left[- (u_{ak}) \left(\sum_i \partial_i s_k \partial_i p_b \right) e_c \right. \\
&\quad + (\delta_{ck} - p_k) \left(\sum_{i,j} \partial_i s_k \partial_i \partial_j p_b \partial_j p_a \right) e_c \\
&\quad \left. + (\delta_{ck} - p_k) \left(\sum_i \partial_i s_k \partial_i p_b \right) H(p_c) e_a \right] \\
&\quad - \sum_k \left[(u_{ac}) \left(\sum_i \partial_i s_k \partial_i p_b \right) e_k \right. \\
&\quad + p_c \left(\sum_{i,j} \partial_i s_k \partial_i \partial_j p_b \partial_j p_a \right) e_k \\
&\quad \left. + p_c \left(\sum_i \partial_i s_k \partial_i p_b \right) H(p_k) e_a \right].
\end{aligned}$$

Proof. The proof is straightforward by developing the following:

$$e_a(H(p_c) e_b) = e_a \left(\left(\sum_k (\delta_{ck} - p_k) \left(\sum_i \partial_i s_k \partial_i p_b \right) \right) e_c - \sum_k p_c \left(\sum_i \partial_i s_k \partial_i p_b \right) e_k \right).$$

□

Lemma C.4.

$$e_a \left(e_c^T J(s)^T A_b J(s) e_d \right) = e_a^T H(p_c)^T J(s)^T A_b J(s) e_d$$

$$\begin{aligned}
& + e_c^T J(s)^T e_a(A_b) J(s) e_d \\
& + e_c^T J(s)^T A_b J(s) H(p_d) e_a.
\end{aligned}$$

Lemma C.5.

$$\begin{aligned}
e_a(A_b)_{kl} & = e_a^T H(p_b) ((\delta_{kl} - p_l) e_k - p_k e_l) \\
& + (e_a^T e_l) (e_b^T e_k) - p_k e_b^T e_l \\
& + e_b^T ((\delta_{kl} - p_l) H(p_k) e_a - p_k e_l) \\
& + (\delta_{kl} - p_l) (e_b^T e_k) - (e_a^T e_k) (e_b^T e_l).
\end{aligned}$$

Proof. The proof is straightforward by developing the following:

$$e_a(A_b)_{kl} = \sum_i \partial_i p_a \partial_i (e_b^T ((\delta_{kl} - p_l) e_k - p_k e_l)).$$

□

Lemma C.6.

$$\begin{aligned}
e_a(g^D(e_b, [e_c, e_d])) & = e_a^T H(p_b) D(x) (H(p_d) e_c - H(p_c) e_d) \\
& + e_b^T J(s)^T A_a J(s) (H(p_d) e_c - H(p_c) e_d) \\
& + e_b^T D(x) (B_{a,d,c} - B_{a,c,d}).
\end{aligned} \tag{23}$$

with

$$\begin{aligned}
B_{a,c,d} & = \sum_k \left[- (u_{ak}) \left(\sum_i \partial_i s_k \partial_i p_d \right) e_c \right. \\
& + (\delta_{ck} - p_k) \left(\sum_{i,j} \partial_i s_k \partial_i \partial_j p_d \partial_j p_a \right) e_c \\
& + (\delta_{ck} - p_k) \left(\sum_i \partial_i s_k \partial_i p_d \right) H(p_c) e_a \\
& - (u_{ac}) \left(\sum_i \partial_i s_k \partial_i p_d \right) e_k \\
& - p_c \left(\sum_{i,j} \partial_i s_k \partial_i \partial_j p_d \partial_j p_a \right) e_k \\
& \left. - p_c \left(\sum_i \partial_i s_k \partial_i p_d \right) H(p_k) e_a \right]
\end{aligned}$$

Proof of the lemma.

$$\begin{aligned}
e_a (g^D (e_b, [e_c, e_d])) &= e_a (e_b^T D(x) (H(p_d) e_c - H(p_c) e_d)) \\
&= e_a^T H(p_b) D(x) (H(p_d) e_c - H(p_c) e_d) \\
&\quad + e_b^T J(s)^T A_a J(s) (H(p_d) e_c - H(p_c) e_d) \\
&\quad + e_b^T D(x) e_a (H(p_d) e_c - H(p_c) e_d).
\end{aligned}$$

Besides, $e_a (H(p_d) e_c)$ has already been computed previously. \square

D More experiments

In this section, we report some more experiments to ensure reproducibility. The code used to compute the DIM is available at <https://github.com/eliot-tron/curvcomputenn/>. The seed used for the following experiments is 42. On Table 4 and Table 5, we report the extreme values of the matrix $\hat{D}(x) = (g_x^D(e_a, e_b))_{a,b=1,\dots,C}$ evaluated on six points of the test set of MNIST (see Figure 5) and CIFAR10 (see Figure 6) respectively.

Table 4: Extreme values of the matrix $\hat{D}(x) = (g_x^D(e_a, e_b))_{a,b=1,\dots,C}$ evaluated on MNIST.

Point x	Highest $g(e_a, e_b)$		Second highest $g(e_a, e_b)$		Lowest $g(e_a, e_b)$	
	indices (a, b)	value $\hat{D}(x)_{a,b}$	indices (a, b)	value $\hat{D}(x)_{a,b}$	indices (a, b)	value $\hat{D}(x)_{a,b}$
n°0	(4, 4)	4.4510e-11	(9, 9)	4.1848e-11	(4, 9) and (9, 4)	-4.3158e-11
n°1	(7, 7)	3.9485e-05	(3, 3)	3.3201e-05	(7, 3) and (3, 7)	-3.6186e-05
n°2	(9, 9)	1.7384e-04	(8, 8)	1.4993e-04	(9, 8) and (8, 9)	-1.6140e-04
n°3	(3, 3)	3.9730e-08	(5, 5)	3.3480e-08	(3, 5) and (5, 3)	-3.6456e-08
n°4	(8, 8)	1.8038e-08	(5, 5)	1.4309e-08	(5, 8) and (8, 5)	-1.6045e-08
n°5	(2, 2)	4.1834e-09	(8, 8)	3.5215e-09	(2, 8) and (8, 2)	-3.8379e-09

Table 5: Extreme values of the matrix $\hat{D}(x) = (g_x^D(e_a, e_b))_{a,b=1,\dots,C}$ evaluated on CIFAR10.

Point x	Highest $g(e_a, e_b)$		Second highest $g(e_a, e_b)$		Lowest $g(e_a, e_b)$	
	indices (a, b)	value $\hat{D}(x)_{a,b}$	indices (a, b)	value $\hat{D}(x)_{a,b}$	indices (a, b)	value $\hat{D}(x)_{a,b}$
n°0	(6, 6)	8.9736e-01	(2, 2)	8.9682e-01	(6, 2) and (2, 6)	-8.9709e-01
n°1	(8, 8)	1.7268e-25	(0, 0)	6.2142e-26	(8, 0) and (0, 8)	-1.0334e-25
n°2	(9, 9)	4.4503e-15	(1, 1)	4.4414e-15	(1, 9) and (9, 1)	-4.4458e-15
n°3	(7, 7)	1.1507e-23	(1, 1)	1.0181e-23	(7, 1) and (1, 7)	-1.0822e-23
n°4	(3, 3)	8.8027e-02	(4, 4)	6.0344e-02	(4, 3) and (3, 4)	-7.2665e-02
n°5	(0, 0)	4.0292e-17	(1, 1)	3.1293e-17	(1, 0) and (0, 1)	-3.5203e-17

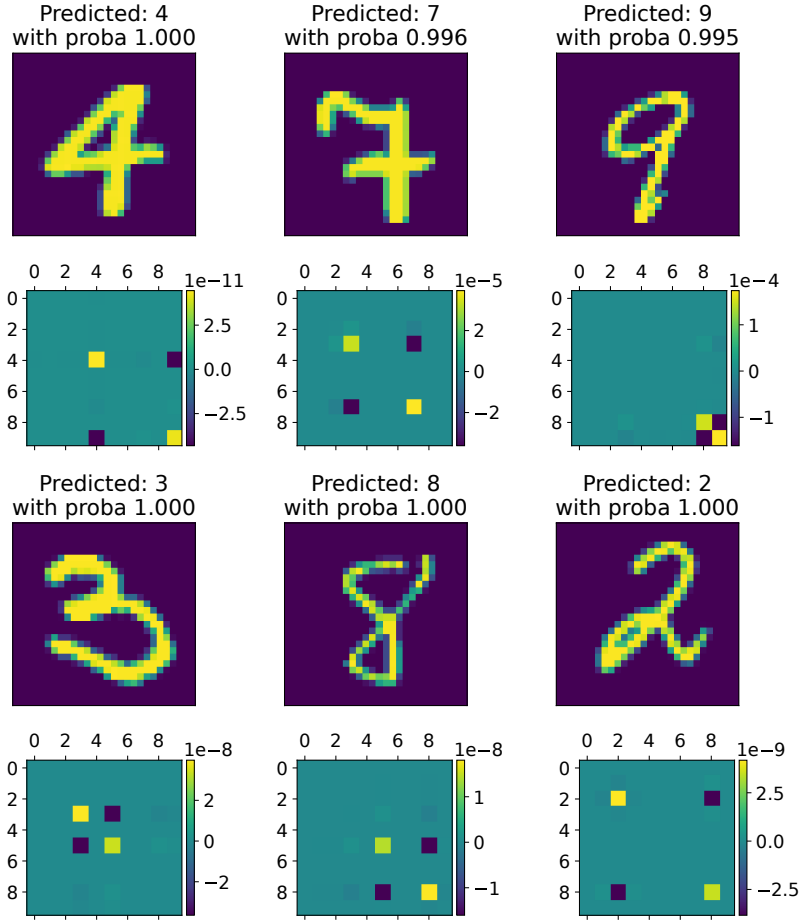


Fig. 5: Couples of input point x (above) and the corresponding matrix $\hat{D}(x) = (g_x^D(e_a, e_b))_{a,b=1,\dots,C}$ (below) on MNIST used for Table 4 numbered from left to right and from top to bottom.

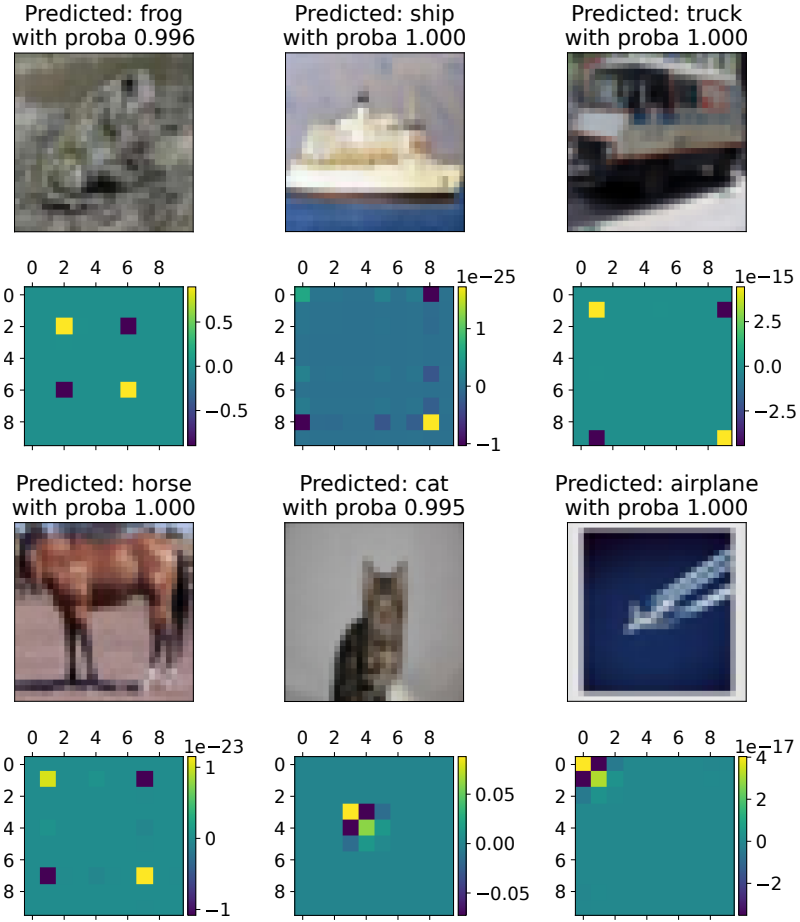


Fig. 6: Couples of input point x (above) and the corresponding matrix $\hat{D}(x) = (g_x^D(e_a, e_b))_{a,b=1,\dots,C}$ (below) on CIFAR10 used for Table 5 numbered from left to right and from top to bottom.