Improved vessel tracking and new data-driven connections on SE(2)

Wemmenhove, A.J. (Jelle)

Award date:
2020

Improved Vessel Tracking and New Data-Driven Connections on SE(2)

A.J. Wemmenhove

Master Thesis 2020

Industrial and Applied Mathematics

Supervisors

dr.ir. R. Duits

ir. B.M.N. Smets

Contents

Introduction

Retinal imaging is frequently used in the analysis and screening of diseases with a vascular component because it allows for a non-invasive observation of the human circulatory system. For example, high tortuosity of retina vessels is used as an early indicator of diabetes, see Figure 1. To monitor the development of pathologies, the vascular network of the retina is encoded with a mathematical model that can be compared over time. One of the methods to create these models is by automated vessel tracking.

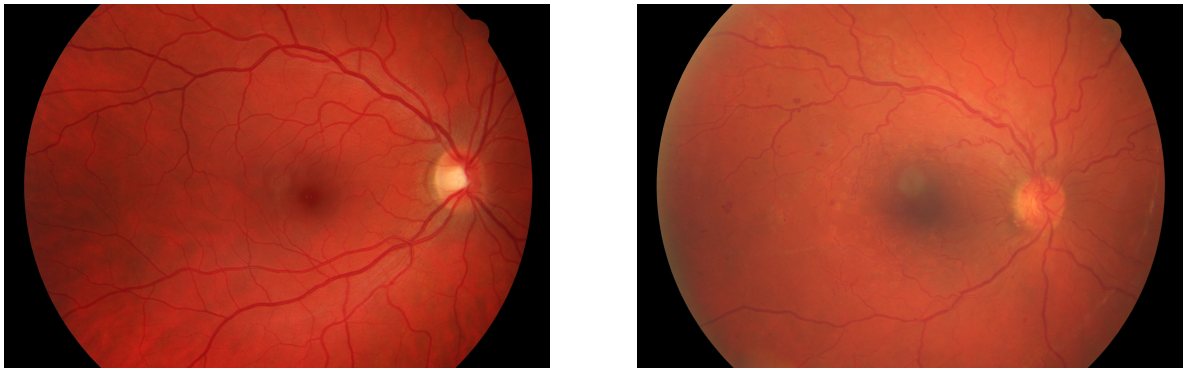

Figure 1: Retinal images of a healthy patient (left) and a patient with diabetes (right). Images taken from the *DRIVE* dataset.

Intersections pose a major problem in vessel tracking. At a point where two vessels overlap, a tracking algorithm may jump from one vessel to the next, resulting in nonsensical paths. To prevent this, we pull the retina image apart into an additional dimension that encodes the orientation of vessels. An example of the resulting **orientation score** [5] is shown in Figure 2. Note how the diagonal lines are mapped to different orientations. Their intersection has disappeared as the lines have been separated by the new dimension in the orientation score.

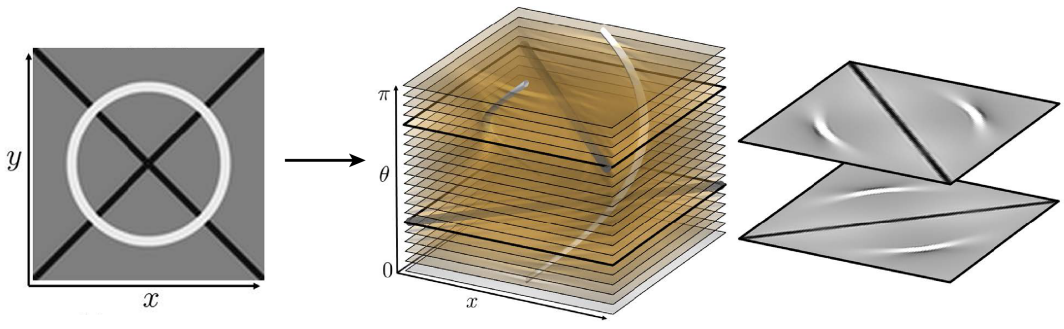

Figure 2: Transforming an image into an orientation score. Adapted from [2].

Every orientation score is equipped with a coordinate frame that rotates with a vessel as its orientation changes, see Figure 3. Referred to as the **left invariant frame**, it allows for a geometrical interpretation of vessel tracking techniques. In previous work by E.J. Bekkers, R. Duits, G.R. Sanguinetti, J-M. Mirebeau

and others [3, 4, 8, 17], quality tracking results are obtained by disallowing movements perpendicular to the vessel direction.

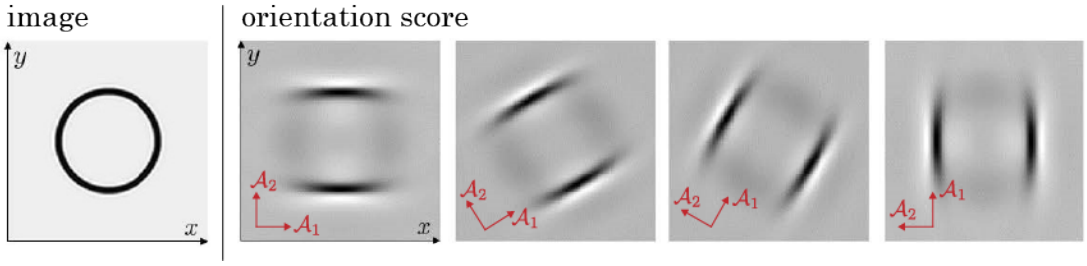

Figure 3: Slices of the orientation score showing the rotation of the left invariant frame (red). Adapted from [2].

Alas, the left invariant frame can fail to align with the vessel structure in two distinct ways.

1. Estimating the orientation of a vessel is an imperfect process. As a result, a portion of a vessel might be assigned an angular coordinate in the orientation score that does not match their true orientation. The left invariant frame however is based on the angular coordinate, so it then fails to align with the vessel data. This is called **deviation of horizontality** [15] and it is illustrated in Figure 4. In practice this effect is worsened. We often take a low resolution for the angular dimension of an orientation score to limit computation time, but this also means that a large range of vessel orientations are mapped to the same angular coordinate.
2. The left invariant frame only lines up with the vessel structure in the spatial dimensions of the orientation score. This is fine for straight vessels as their orientation is constant, but the changing orientation of curved vessels is ignored.

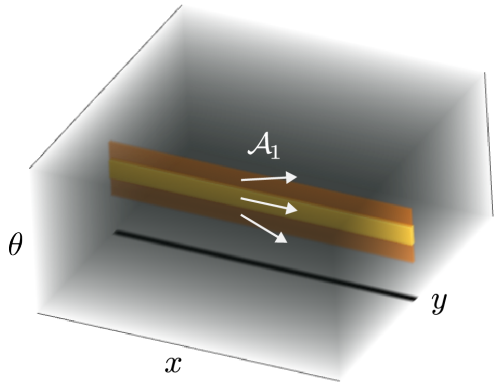

Figure 4: A straight line provokes responses at multiple orientations. The (mis)alignment of the left invariant frame is shown by the white arrows.

Deviation of horizontality and ignoring curvature can steer the vessel tracking algorithm away from the actual vessel. For example, taking curvature into account might help to better track highly tortuous vessels, something current methods fail at, see Figure 5.

Explicitly fitting a frame to the vessels in an orientation score solves the misalignment problem of the left invariant frame. The result is called a **locally adaptive frame** or **gauge frame** [7]. A comparison with such a frame and a left invariant frame is shown in Figure 6. Although gauge frames allow for an even better geometrical interpretation of processing techniques, they have one major drawback: they are not smooth. Due to noise, the direction that best fits a vessel can vary suddenly. In areas without a clear line structure, a gauge frame can jump all over the place. Several smoothing methods exists that attempt to ameliorate this issue.

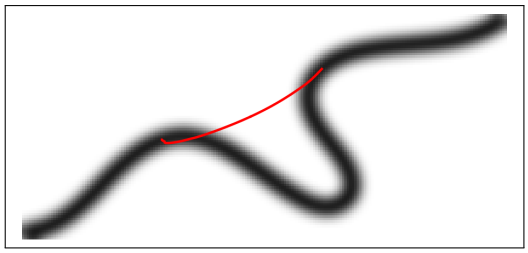

Figure 5: Shortcut taken by tracking algorithm from [cite] over a highly tortuous curve.

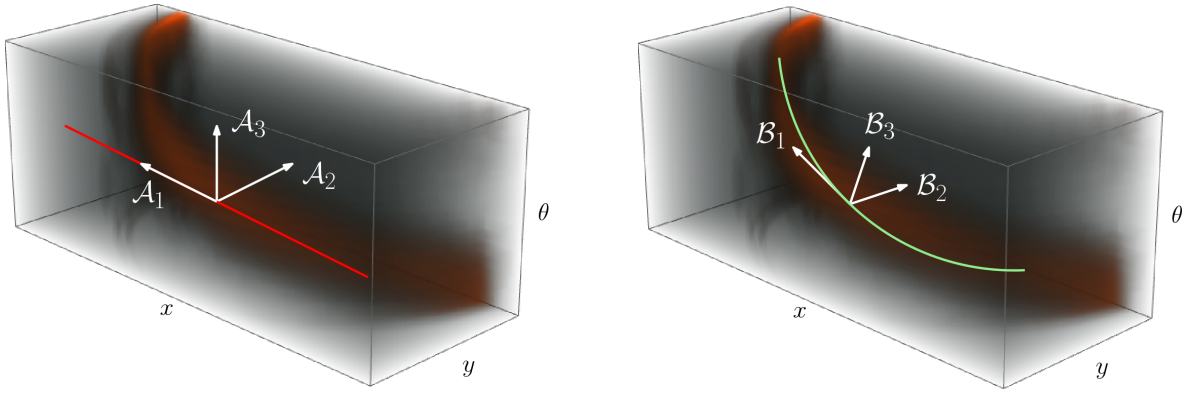

Figure 6: Comparison of a left invariant frame (left) and a gauge frame (right).

In this report we explore the extent to which gauge frames can be used in lieu of the left invariant frame to steer vessel tracking. The results are compared with the tracking results obtained by Erik Bekkers [4] that disallow side-stepping in the left invariant frame. The results are presented in Chapter 2.

The final chapter characterizes the paths resulting from the gauge frame tracking methods as curves with auto-parallel momentum. Previously this characterization had only been shown for techniques that use a fixed scaling of the left invariant frame [9]. As it turns out, the old proof can be extended to a more general setting including the class of gauge frame based methods investigated in this report.

In this work we integrate for the first time both locally adaptive frame theory [7] and (sub-Riemannian) geodesic tracking in $SE(2)$ [4].

Chapter 1

Geometrical Theory

This chapter serves to explain the geometrical concepts used for vessel tracking in the next chapter. The first half introduces a lot of new concepts, the second half discusses the different frames one can use to construct geometrical tracking methods. Of these frames, the largest amount of text is spent on gauge frames.

1.1 Main Concepts

1.1.1 Positions, Orientations and Roto-Translations

Definition 1 (2D positions and orientations)**.** The space of two-dimensional positions and orientations $\mathbb{M}_2$ is defined as the smooth manifold

$$\mathbb{M}_2 := \mathbb{R}^2 \times S^1 \ ,$$

where we take the circle S^1 to be the quotient space $S^1 := \mathbb{R} \,/\, 2\pi\mathbb{Z}$.

The manifold S^1 is smooth since the transition functions of its coordinate charts are simply translations of $\mathbb{R}$ by an integer multiple of 2π.

Definition 2 (2D roto-translations)**.** The space of two-dimensional roto-translations is the Lie group $SE(2) := \mathbb{R}^2 \rtimes SO(2)$. Its elements are pairs $(\mathbf{z}, \mathbf{R}_\alpha)$ denoting a translation by $\mathbf{z} \in \mathbb{R}^2$ and a rotation $\mathbf{R}_\alpha \in SO(2)$ by $\alpha \in S^1$ radians. A roto-translation $(\mathbf{z}, \mathbf{R}_\alpha)$ represents a rotation by α radians *followed* by a translation by $\mathbf{z}$, both with respect to a fixed coordinate frame, see Figure 1.1. The group product is given by

$$(\mathbf{z}, \mathbf{R}_\alpha) \cdot (\mathbf{w}, \mathbf{R}_\beta) := (\mathbf{z} + \mathbf{R}_\alpha \mathbf{w}, \mathbf{R}_\alpha \mathbf{R}_\beta) \ , \tag{1.1}$$

and represents performing the roto-translation $(\mathbf{w}, \mathbf{R}_\beta)$ *followed* by the roto-translation $(\mathbf{z}, \mathbf{R}_\alpha)$ (see Figure 1.1).

The notation $\rtimes$ in $\mathbb{R}^2 \rtimes SO(2)$ emphasizes that $SE(2)$ is not a direct product of the translation and rotation groups $\mathbb{R}^2$ and $SO(2)$. Both groups are commutative so their direct product is as well, but as we know from experience, composition of roto-translations is not commutative (again, see Figure 1.1).

The inverse of a roto-translation $(\mathbf{z}, \mathbf{R}_\alpha)$ is given by

$$(\mathbf{z}, \mathbf{R}_\alpha)^{-1} := (-\mathbf{R}_\alpha^{-1}\, \mathbf{z}, \mathbf{R}_\alpha^{-1}) \ , \tag{1.2}$$

as is readily verified:

$$\begin{aligned}
(\mathbf{z}, \mathbf{R}_\alpha) \cdot (-\mathbf{R}_\alpha^{-1}\, \mathbf{z}, \mathbf{R}_\alpha^{-1}) &= (\mathbf{z} - \mathbf{R}_\alpha \mathbf{R}_\alpha^{-1}\, \mathbf{z}, \mathbf{R}_\alpha \mathbf{R}_\alpha^{-1}) = (\mathbf{0}, \mathbf{I}) \\
(-\mathbf{R}_\alpha^{-1}\, \mathbf{z}, \mathbf{R}_\alpha^{-1}) \cdot (\mathbf{z}, \mathbf{R}_\alpha) &= (-\mathbf{R}_\alpha^{-1}\, \mathbf{z} + \mathbf{R}_\alpha^{-1}\, \mathbf{z}, \mathbf{R}_\alpha^{-1} \mathbf{R}_\alpha) = (\mathbf{0}, \mathbf{I}) \ .
\end{aligned}$$

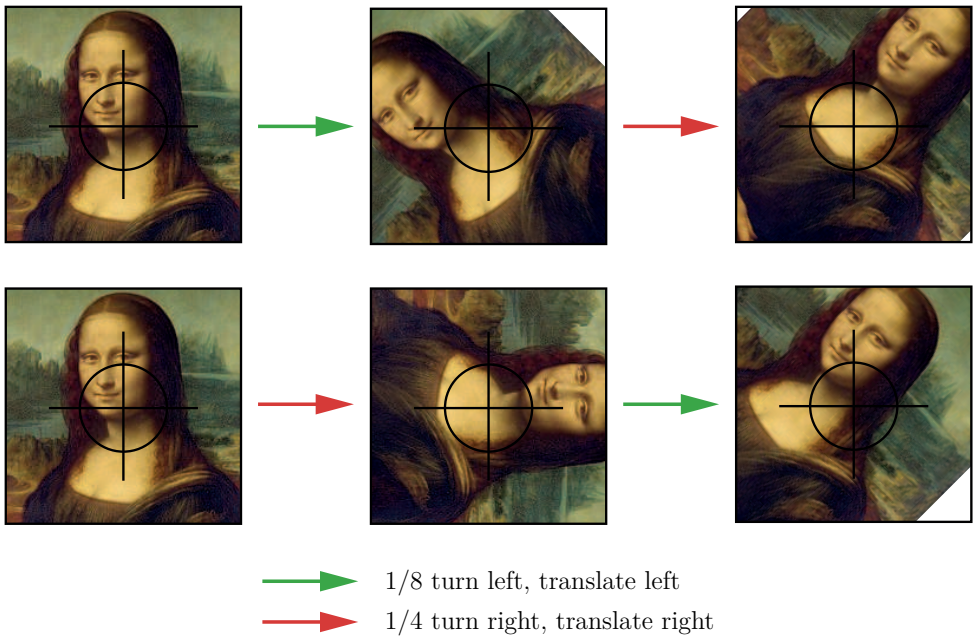

Figure 1.1: Composed roto-translations of the Mona Lisa w.r.t. a fixed coordinate frame.

Remark. *Although in two dimensions there is a bijection between orientations θ and rotations $\mathbf{R}_\theta$, this is not the case in higher dimensions. For this reason we keep the homogeneous space $\mathbb{M}_2$ and the Lie group $SE(2)$ separate.*

1.1.2 Group Actions

Roto-translations can be applied to lots of constructions. Mathematically this is formulated in terms of group actions.

Definition 3 (Group action)**.** Given a set A, the set of automorphisms $\mathrm{Aut}(A)$ is a group w.r.t. composition. The action of a group G on a set A is a group homomorphism $\rho : G \to \mathrm{Aut}(A)$. Given a group element $g \in G$, we denote the corresponding bijection by $\rho_g := \rho(g) : A \xrightarrow{\cong} A$ or we use the infix notation $g \,.\, x := \rho_g(x)$ for the application of $g \in G$ to $x \in A$.

The infix notation allows us to write out explicitly what it means for a group action to be a group homomorphism. Since the product and unit structures are preserved, respectively the following identities hold:

$$(gh) \,.\, x = g \,.\, (h \,.\, x) \tag{1.3}$$
$$e \,.\, x = x \ , \tag{1.4}$$

for group elements $g, h \in G$, the unit element $e \in G$ and an element $x \in A$.

Consider the following two group actions of $SE(2)$.

Definition 4. The action of $SE(2)$ on the space of positions $\mathbb{R}^2$ is given by

$$(\mathbf{z}, \mathbf{R}_\alpha) \,.\, \mathbf{x} = \mathbf{z} + \mathbf{R}_\alpha \, \mathbf{x} \ , \tag{1.5}$$

for a roto-translation $(\mathbf{z}, \mathbf{R}_\alpha) \in SE(2)$ and a position $\mathbf{x} \in \mathbb{R}^2$. We denote this action by $L_g : \mathbb{R}^2 \to \mathbb{R}^2$.

Definition 5. The action of $SE(2)$ on the space of positions and orientations $\mathbb{M}_2$ is given by

$$(\mathbf{z}, \mathbf{R}_\alpha) \,.\, (\mathbf{x}, \theta) = (\mathbf{z} + \mathbf{R}_\alpha \, \mathbf{x}, \alpha + \theta) \ , \tag{1.6}$$

for a roto-translation $(\mathbf{z}, \mathbf{R}_\alpha) \in SE(2)$ and combined position and orientation $(\mathbf{x}, \theta) \in \mathbb{M}_2$. We denote this action by U_g.

The fact that group actions respect the group product emphasizes that the group product $(\mathbf{z}, \mathbf{R}_\alpha)\cdot(\mathbf{w}, \mathbf{R}_\beta)$ represents a roto-translation $(\mathbf{w}, \mathbf{R}_\beta)$ *followed* by a roto-translation $(\mathbf{z}, \mathbf{R}_\alpha)$.

Once we have a group action $G \to \mathrm{Aut}(A)$ there is a straightforward way that the group G acts on a set of functions $A \to B$.

Lemma 1. *Given a group action $\rho : G \to \mathrm{Aut}(A)$, the group G acts canonically on the set $\mathrm{Hom}(A, B)$ of functions $f : A \to B$ via the action*

$$(g \,.\, f)(x) := f(g^{-1} \,.\, x) \ , \tag{1.7}$$

for $g \in G$ and $x \in A$.

Proof. The map $G \to \mathrm{Aut}(\mathrm{Hom}(A, B))$ defined above is indeed a group homomorphism. The product structure is respected since for all $g, h \in G$, $f : A \to B$ and $x \in A$ we have that

$$((gh) \,.\, f)(x) = f((gh)^{-1} \,.\, x) = f((h^{-1}g^{-1}) \,.\, x) = f(h^{-1} \,.\, (g^{-1} \,.\, x)) = (h \,.\, f)(g^{-1} \,.\, x) = (g \,.\, (h \,.\, f))(x) \ .$$

The identity structure is also respected. Given the unit element $e \in G$, we have that

$$(e \,.\, f)(x) = f(e^{-1} \,.\, x) = f(e \,.\, x) = f(x) \ .$$

■

1.1.3 Images

We model gray-scale images by functions $f : \mathbb{R}^2 \to [0, 1]$ where the value $f(\mathbf{x}) \in [0, 1]$ represents the image lightness at position $\mathbf{x} \in \mathbb{R}^2$. Black and white correspond to the values 0 and 1 respectively. This is however a very wild class of images that cannot all be turned into orientation scores. Hence, we also consider the smaller class of images $\mathbb{L}_2(\mathbb{R}^2 ; [0, 1])$. This is not too strong a restriction, since any image $f : \mathbb{R}^2 \to [0, 1]$ with a bounded support is also an element of $\mathbb{L}_2(\mathbb{R}^2 ; [0, 1])$. One can think of these images as displaying information in a limited area against a black background.

By Lemma 1, the action of the roto-translation group on the space of positions (Definition 4) induces a group action of $SE(2)$ on images $f : \mathbb{R}^2 \to [0, 1]$. This action is given by

$$((\mathbf{z}, \mathbf{R}_\alpha) \,.\, f)(\mathbf{x}) = f\left((\mathbf{z}, \mathbf{R}_\alpha)^{-1} \,.\, \mathbf{x}\right) = f\left(-\mathbf{R}_\alpha^{-1}\,\mathbf{z} + \mathbf{R}_\alpha^{-1}\,\mathbf{x}\right) = f\left(\mathbf{R}_\alpha^{-1}(\mathbf{x} - \mathbf{z})\right) \ . \tag{1.8}$$

First the image is rotated by α radians, *then* it is translated by $\mathbf{z}$. We use the notation $\mathcal{U}_g$ for the action of $g \in SE(2)$ on images $\mathbb{R}^2 \to [0, 1]$.

1.1.4 Orientation Score

An orientation score transformation W_ψ extends the domain of an image $f \in \mathbb{L}_2(\mathbb{R}^2 ; [0, 1])$ to the space of positions and orientations $\mathbb{M}_2$. The result is an orientation score $W_\psi f : \mathbb{M}_2 \to [-1, 1]$. The value $|W_\psi f(\mathbf{x}, \theta)| \in [0, 1]$ encodes the extend to which a line structure with orientation $\theta \in S^1$ is present in the original image f at position $\mathbf{x} \in \mathbb{R}^2$. An orientation score is computed by probing an image with a wavelet $\psi \in \mathbb{L}_2(\mathbb{R}^2) \cap \mathbb{L}_1(\mathbb{R}^2)$ that measures line structures in a certain place and orientation. By roto-translating the wavelet via the group action in (1.8), all positions are probed for line structures of all possible orientations. The explicit definition is as follows.

Definition 6. Given an image $f : \mathbb{L}_2(\mathbb{R}^2 ; [0, 1])$, its orientation score $W_\psi f : \mathbb{M}_2 \to [-1, 1]$ is given by

$$(W_\psi f)(\mathbf{x}, \theta) := \left(\mathcal{U}_{(\mathbf{x}, \mathbf{R}_\theta)}\psi, f\right)_{\mathbb{L}_2(\mathbb{R}^2)} = \int_{\mathbb{R}^2} \psi\left(\mathbf{R}_\theta^{-1}(\mathbf{y} - \mathbf{x})\right) f(\mathbf{y})\, d\mathbf{y} \ , \tag{1.9}$$

where $\psi \in \mathbb{L}_2(\mathbb{R}^2) \cap \mathbb{L}_1(\mathbb{R}^2)$ denotes a wavelet measuring a line structure at the origin with a horizontal orientation. Requiring that $\int_{\mathbb{R}^2} |\psi(\mathbf{y})|^2 d\mathbf{y} \leq 1$ guarantees that the orientation score takes values in the range $[-1, 1]$.

Remark. *To extend this definition to higher dimensional images* $f \in \mathbb{L}_2(\mathbb{R}^d;[0,1])$ *, one needs to consider that there is no longer a bijection between orientations and rotations. Let the space of d-dimensional orientations be denoted by* $S^{d-1} := \{\mathbf{y} \in \mathbb{R}^d \mid ||\mathbf{y}|| = 1\}$ *and suppose that we have a reference orientation* $\mathbf{a} \in S^{d-1}$ *. To evaluate the orientation score transform* $W_\psi f$ *at orientation* $\mathbf{n} \in S^{d-1}$*, one needs to pick a rotation* $\mathbf{R_n}$ *that maps the reference orientation* $\mathbf{a}$ *to the requested orientation* $\mathbf{n}$ *. There is no canonical choice of rotation and thus, to have a well-posed orientation score, we require that the wavelet* ψ *is axially symmetric around the reference orientation. Well-posedness then follows since* $\mathcal{U}_{(\mathbf{x},\mathbf{R_n})}\psi = \mathcal{U}_{(\mathbf{x},\mathbf{R'_n})}\psi$ *for any two rotations* $\mathbf{R_n}, \mathbf{R'_n}$ *that rotate the vector* $\mathbf{a}$ *towards the vector* $\mathbf{n}$ *.*

More details about the orientation score transform such as its inversibility can be found in [11,12], which also describe several options for wavelets ψ. For our application, we will be using **cake wavelets**, so called for their resemblance to pieces of cake in the Fourier domain.

As with images on $\mathbb{R}^2$, we also have a group action for functions on the space $\mathbb{M}_2$. By applying Lemma 1 to the group action $SE(2) \to \mathrm{Aut}(\mathbb{M}_2)$ defined by (1.6), we obtain an action of $SE(2)$ on functions $V : \mathbb{M}_2 \to \mathbb{R}$ given by

$$((\mathbf{z}, \mathbf{R}_\alpha) \,.\, V)\,(\mathbf{x}, \theta) = V((\mathbf{z}, \mathbf{R}_\alpha)^{-1} \,.\, (\mathbf{x}, \theta)) = V(-\mathbf{R}_\alpha^{-1}\,\mathbf{z} + \mathbf{R}_\alpha^{-1}\,\mathbf{x}, -\alpha + \theta) = V(\mathbf{R}_\alpha^{-1}(\mathbf{x}-\mathbf{z}), \theta - \alpha)\ . \quad (1.10)$$

Similar to the action of $SE(2)$ on images, this represents a rotation of the function *followed* by a translation. Since V is a function on $\mathbb{M}_2$, the rotation also translates the function in the angular dimension. We use the notation $\mathcal{L}_g$ to denote the effect of the group element $g \in SE(2)$ on the function space $\mathrm{Hom}(\mathbb{M}_2, \mathbb{R})$.

1.1.5 Vessel Tracking

For our purposes, the problem of vessel tracking is as follows. Given an image of a vascular structure $f : \mathbb{L}_2(\mathbb{R}^2\,;[0,1])$, find a curve $\gamma : \mathrm{Lip}([0,1], \mathbb{R}^2)$ that follows a single vessel from a specified starting point $p = \gamma(0)$ to an end point $q = \gamma(1)$. For γ to represent a geometrical path, we require that it is Lipschitz continuous. To prevent jumping from vessel to vessel at crossings, the image f is first lifted to an orientation score $W_\psi f : \mathbb{M}_2 \to [-1, 1]$ and the actual tracking takes place in the space of positions and orientations $\mathbb{M}_2$. Hence, we can demand that the starting and end points contain orientation data as well, i.e. $p, q \in \mathbb{M}_2$. After the tracking process Φ is complete, we can discard the angular component via a projection P to obtain a curve γ in the space of positions $\mathbb{R}^2$. This process is captured in the following diagram:

$$\begin{array}{ccc} \text{orientation score } W_\psi f & \xmapsto{\ \Phi\ } & \text{curve } \tilde{\gamma} \text{ in } \mathbb{M}_2 \\ {\scriptstyle W_\psi}\uparrow & & \downarrow{\scriptstyle P} \\ \text{image } f & & \text{curve } \gamma \text{ in } \mathbb{R}^2 \end{array}$$

The tracking process Φ takes as input a function $V : \mathbb{M}_2 \to \mathbb{R}$ and two points $p, q \in \mathbb{M}_2$ from which it produces a curve $\Phi(V, p, q) \in \mathrm{Lip}([0,1], \mathbb{M}_2)$. For the actual tracking algorithm, it is common to compute a path from p to q that minimizes a length functional. Based on the input data, one constructs a metric tensor field $\mathcal{G}(V)$ that induces a metric $d_{\mathcal{G}(V)}$ on the space of positions and orientations by

$$d_{\mathcal{G}(V)}(x, y) := \min_{\substack{\gamma \in \mathrm{Lip}([0,1],\mathbb{M}_2) \\ \gamma(0)=x \\ \gamma(1)=y}} \int_0^1 ||\dot{\gamma}(t)||^2_{\mathcal{G}(V)}\, dt\ , \quad (1.11)$$

where the norm of a vector $u \in T_x\mathbb{M}_2$ is given by $||u||^2_{\mathcal{G}(V)} = \mathcal{G}(V)\big|_x(u, u)$. Curves minimizing such a problem are also called **shortest** curves. The output of the tracking process $\Phi(V, p, q) = \tilde{\gamma}$ is a shortest curve between the starting and end points $p, q \in \mathbb{M}_2$ according to the metric $d_{\mathcal{G}(V)}$. Note that we write *'a shortest curve'* since these might not be unique. The computational technique used in this report does give a unique output however. For more details, see Section 2.1.2.

Although the problem of vessel tracking might sound intimidating at first, in the context of curve minimization, it boils down to picking a suitable metric tensor field that encourages movement along vessels.

1.1.6 Roto-Translation Equivariance

So far we have discussed several actions of the roto-translation group $SE(2)$:

- the action U_g on positions (1.5);
- the action L_g on positions and orientations (1.6);
- the action $\mathcal{U}_g$ on images (1.8);
- the action $\mathcal{L}_g$ on orientation scores(1.10).

These actions all serve a single purpose: we want to show that the results of the tracking algorithm are equivariant under roto-translations of the input image. When the input image is roto-translated, the output curve should roto-translate with it. To achieve this, every step of the computation $P \circ \Phi \circ W_\psi$ needs to be roto-translation equivariant.

Definition 7 (Roto-translation equivariant)**.** Let $\rho_A : SE(2) \to \mathrm{Aut}(A)$ and $\rho_B : SE(2) \to \mathrm{Aut}(B)$ be two group actions. We say that a map $\Psi : A \to B$ is roto-translation equivariant if for each roto-translation $g \in SE(2)$

$$\Psi \circ (\rho_A)_g = (\rho_B)_g \circ \Psi \ . \tag{1.12}$$

Lemma 2. *The orientation score transformation W_ψ is roto-translation equivariant, i.e. it holds that*

$$W_\psi \circ \mathcal{U}_g = \mathcal{L}_g \circ W_\psi \ .$$

Proof. Note that the action $\mathcal{U}_g \in \mathrm{Aut}(\mathbb{L}_2(\mathbb{R}^2))$ is unitary for each $g = (\mathbf{z}, \mathbf{R}_\alpha) \in SE(2)$. For two functions $\psi, \varphi \in \mathbb{L}_2(\mathbb{R}^2)$, we have that

$$(\mathcal{U}_g\psi, \varphi)_{\mathbb{L}_2(\mathbb{R}^2)} = \int_{\mathbb{R}^2} \psi\left(\mathbf{R}_\alpha^{-1}(\mathbf{y}-\mathbf{z})\right) \varphi(\mathbf{y})\, d\mathbf{y} = \int_{\mathbb{R}^2} \psi(\mathbf{u})\, \varphi\left(\mathbf{R}_\alpha \mathbf{u} + \mathbf{z}\right) d\mathbf{u} = \left(\psi, \mathcal{U}_{g^{-1}}\varphi\right)_{\mathbb{L}_2(\mathbb{R}^2)} \ ,$$

where we apply the substitution $\mathbf{u} := \mathbf{R}_\alpha^{-1}(\mathbf{y} - \mathbf{z})$ in the second step. So we indeed have unitarity: $\mathcal{U}_g^* = \mathcal{U}_{g^{-1}} = (\mathcal{U}_g)^{-1}$. This implies that

$$\begin{aligned}\left(W_\psi(\mathcal{U}_g f)\right)(\mathbf{x}, \theta) &= \left(\mathcal{U}_{(\mathbf{x},\mathbf{R}_\theta)}\psi,\, \mathcal{U}_g f\right)_{\mathbb{L}_2(\mathbb{R}^2)} = \left(\mathcal{U}_{g^{-1}}(\mathcal{U}_{(\mathbf{x},\mathbf{R}_\theta)}\psi), f\right)_{\mathbb{L}_2(\mathbb{R}^2)} \\ &= \left(\mathcal{U}_{g^{-1}(\mathbf{x},\mathbf{R}_\theta)}\psi, f\right)_{\mathbb{L}_2(\mathbb{R}^2)} = \left(W_\psi f\right)\left(g^{-1} \,.\, (\mathbf{x}, \theta)\right) = \left(\mathcal{L}_g(W_\psi f)\right)(\mathbf{x}, \theta) \ ,\end{aligned}$$

so W_ψ is indeed roto-translation equivariant. ■

To speak of roto-translation equivariance for the tracking and projection steps, we need a define a group action on curves.

Definition 8. For a curve γ in $\mathbb{R}^2$ or $\mathbb{M}_2$, the action of roto-translations is given by post-composition of γ with the action on the respective spaces:

$$(g \,.\, \gamma)(t) := g \,.\, \gamma(t) \ . \tag{1.13}$$

To prevent introducing too many symbols, we express roto-translation of curves by using the temporal parameter t. For a curve $\gamma(t)$ in $\mathbb{R}^2$, its roto-translation is given by $U_g(\gamma(t))$. Similarly for a curve $\gamma(t)$ in $\mathbb{M}_2$ the roto-translated curve is $L_g(\gamma(t))$.

Lemma 3. *The projection P of curves in $\mathbb{M}_2$ onto curves in $\mathbb{R}^2$ is roto-translation equivariant.*

Proof. The projection of a roto-translated curve in $\mathbb{M}_2$ is the same as the roto-translation of a projected curve. This is because the actions of roto-translations on the spaces $\mathbb{R}^2$ and $\mathbb{M}_2$ agree in the spatial component (see (1.5) and (1.6)), which is all that matters to the projection P. ■

Lemma 4. *If the metric tensor field on $\mathbb{M}_2$ is roto-translation invariant in the sense that for any function $V : \mathbb{M}_2 \to \mathbb{R}$ and vectors $u, v \in T_x\mathbb{M}_2$:*

$$\mathcal{G}(\mathcal{L}_g V)\big|_{L_g x}((L_g)_* u, (L_g)_* v) = \mathcal{G}(V)\big|_x(u, v) \ , \tag{1.14}$$

a curve $\gamma(t)$ minimizes $d_{\mathcal{G}(V)}(p, q)$ if and only if its roto-translation $L_g(\gamma(t))$ minimizes the roto-translated optimization problem of $d_{\mathcal{G}(\mathcal{L}_g V)}(L_g p, L_g q)$.

Proof. Let $\gamma(t)$ be a curve that minimizes $d_{\mathcal{G}(V)}(p, q)$. Its roto-translation $L_g(\gamma(t))$ is a curve from $L_g p$ to $L_g q$ and because the metric tensor field $\mathcal{G}(V)$ is roto-translation invariant, $L_g(\gamma(t))$ has the same 'length' as the original curve $\gamma(t)$:

$$\int_0^1 \left|\left| \frac{d}{dt}(L_g(\gamma(t))) \right|\right|^2_{\mathcal{G}(\mathcal{L}_g V)} dt = \int_0^1 ||(L_g)_* \dot{\gamma}(t)||^2_{\mathcal{G}(\mathcal{L}_g V)} \, dt = \int_0^1 ||\dot{\gamma}(t)||^2_{\mathcal{G}(V)} \, dt \ .$$

Let $\delta(t)$ denote any curve from $L_g p$ to $L_g q$. Then $L_g(\gamma(t))$ is shorter than $\delta(t)$ because $\gamma(t)$ is shorter than $L_{g^{-1}}(\delta(t))$. Thus, the roto-translated curve $L_g(\gamma(t))$ minimizes the roto-translated optimization problem of $d_{\mathcal{G}(\mathcal{L}_g V)}(L_g p, L_g q)$. The converse proof is analogous. ■

Hence, we shall only consider roto-translation invariant metric tensor fields for vessel tracking. The following lemma aids in designing such metrics.

Lemma 5. *The metric tensor field $\mathcal{G}(V)$ is roto-translation invariant in the sense of (1.14) if it can be expressed in terms of a global dual frame $\{\psi^i(V)\}_{i=1}^3$ with components $g_{ij}(V) : \mathbb{M}_2 \to \mathbb{R}$ as*

$$\mathcal{G}(V) = \sum_{i,j=1}^{n} g_{ij}(V) \, \psi^i(V) \otimes \psi^j(V) \ , \tag{1.15}$$

where the components are equivariant

$$g_{ij}(\mathcal{L}_g(V)) = \mathcal{L}_g(g_{ij}(V)) \tag{1.16}$$

and the corresponding frame $\{\mathcal{E}_i(V)\}_{i=1}^3$ defined by $\langle \psi^i, \mathcal{E}_j \rangle = \delta^i_j$ satisfies

$$(L_g)_* \, \mathcal{E}_i(V) = \mathcal{E}_i(\mathcal{L}_g V) \ . \tag{1.17}$$

Proof. Substituting the decomposition (1.15) into the roto-translation invariance requirement (1.14), we see that components need to satisfy

$$(g_{ij}(\mathcal{L}_g V))\,(L_g x) = (g_{ij}(V))\,(x) \ ,$$

which is a rephrasing of roto-translational equivariance (1.16). For the dual frame we have that

$$\psi^i(\mathcal{L}_g V)\big|_{L_g x}((L_g)_* u) = \psi^i(V)\big|_x(u)$$

holds for all $u \in T_x\mathbb{M}_2$. Since the corresponding frame is characterized by $\langle \rho^i(V), \mathcal{E}_j(V) \rangle = \delta^i_j$ and it holds that

$$\left\langle \psi^i(\mathcal{L}_g V)\big|_{L_g x}, (L_g)_* \, \mathcal{E}_j(V)\big|_x \right\rangle = \langle \psi^i(V)\big|_x, \mathcal{E}_j(V)\big|_x \rangle = \delta^i_j \ ,$$

we have that equation (1.17) holds for the frame $\{\mathcal{E}_i(V)\}_{i=1}^3$:

$$(L_g)_* \, \mathcal{E}_i(V)\big|_x = \mathcal{E}_i(\mathcal{L}_g V)\big|_{L_g x} \ .$$

■

With the above lemma as our motivation, we discuss three frames one can use to construct a metric tensor field suitable to vessel tracking.

1.2 The Canonical Frame

The first frame we consider is a **canonical** one $\{\partial_x, \partial_y, \partial_\theta\}$ induced by the coordinate functions x, y, θ for points in $\mathbb{M}_2$. As is convention, the dual frame for the cotangent bundle $T^*\mathbb{M}_2$ is denoted by $\{dx, dy, d\theta\}$, where the symbol d truely denotes the *exterior derivative*. For example, the value $dx\big|_z(v)$ denotes the rate of change of the first coordinate in the direction $v \in T_z\mathbb{M}_2$.

The spatial vector fields ∂_x and ∂_y are not invariant under rotation

$$\begin{aligned} L_{(\mathbf{0},\mathbf{R}_\alpha)*}\,\partial_x &= \quad \cos(\alpha)\,\partial_x + \sin(\alpha)\,\partial_y \\ L_{(\mathbf{0},\mathbf{R}_\alpha)*}\,\partial_y &= -\sin(\alpha)\,\partial_x + \cos(\alpha)\,\partial_y \end{aligned}$$

but their span is, meaning that roto-translations of spatial vectors remain spatial vectors.

1.3 The Left Invariant Frame

Thanks to the bijection between two-dimensional orientations θ and rotations $\mathbf{R}_\theta$, it is possible to construct a roto-translation invariant frame $\{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3\}$. It is defined by

$$\begin{cases} \mathcal{A}_i\big|_{(\mathbf{x},\theta)} := L_{(\mathbf{x},\mathbf{R}_\theta)*}\,\mathcal{A}_i\big|_{\mathbf{0}} & \text{for } i = 1,2,3 \\ \mathcal{A}_1\big|_{\mathbf{0}} := \partial_x\big|_{\mathbf{0}} \\ \mathcal{A}_2\big|_{\mathbf{0}} := \partial_y\big|_{\mathbf{0}} \\ \mathcal{A}_3\big|_{\mathbf{0}} := \partial_\theta\big|_{\mathbf{0}} \end{cases} \tag{1.18}$$

We call this the **left invariant frame** and we denote its dual by $\{\omega^1, \omega^2, \omega^3\}$. Roto-translational invariance $(L_g)_*\mathcal{A}_i = \mathcal{A}_i$ holds as

$$L_{(\mathbf{z},\mathbf{R}_\alpha)*}\,\mathcal{A}_i\big|_{(\mathbf{x},\theta)} \overset{(1.18)}{=\!=} L_{(\mathbf{z},\mathbf{R}_\alpha)*}\,L_{(\mathbf{x},\mathbf{R}_\theta)*}\,\mathcal{A}_i\big|_{\mathbf{0}} = L_{(\mathbf{z},\mathbf{R}_\alpha)\cdot(\mathbf{x},\mathbf{R}_\theta)*}\,\mathcal{A}_i\big|_{\mathbf{0}} \overset{(1.18)}{=\!=} \mathcal{A}_i\big|_{(\mathbf{z},\mathbf{R}_\alpha)\,.\,(\mathbf{x},\theta)}\,,$$

where the second equation is the chain rule. By a small calculation, we find that at orientation θ the frame is given by

$$\begin{cases} \mathcal{A}_1 = \quad \cos(\theta)\,\partial_x + \sin(\theta)\,\partial_y \\ \mathcal{A}_2 = -\sin(\theta)\,\partial_x + \cos(\theta)\,\partial_y \\ \mathcal{A}_3 = \quad \partial_\theta \end{cases} \tag{1.19}$$

Observe that the left invariant frame rotates with us as we move through the angular dimension of $\mathbb{M}_2$ as shown in Figure 3. This is very useful for the manipulation of orientation scores. At an angle of θ, the orientation score contains the line structures of the image with the same orientation. Hence, the vector field $\mathcal{A}_1$ lies parallel to these structures and the vector field $\mathcal{A}_2$ is orthogonal to them. We say that these vector fields point **forward** and **left** respectively.

Curves with constant speed w.r.t. the left invariant frame form spirals in the angular direction. For constants $c^1, c^2, c^3 \in \mathbb{R}$, the formula for a curve $\gamma(t) = (x(t), y(t), z(t))$ with derivative $\dot{\gamma}(t) = \sum_{i=1}^{n} c^i \mathcal{A}_i\big|_{\gamma(t)}$ and initial condition $\gamma(0) = \mathbf{0} \in \mathbb{M}_2$ is given by

$$\begin{cases} x(t) = \frac{c^1}{c^3}\sin(c^3\,t) + \frac{c^2}{c^3}(1 - \cos(c^3\,t)) \\ y(t) = \frac{c^1}{c^3}(1 - \cos(c^3\,t)) + \frac{c^2}{c^3}\sin(c^3\,t) \\ \theta(t) = c^3\,t \end{cases} \qquad \text{if } c^3 \neq 0 \tag{1.20}$$

$$\begin{cases} x(t) = c^1\,t \\ y(t) = c^2\,t \\ \theta(t) = 0 \end{cases} \qquad \text{otherwise} \tag{1.21}$$

Figure 1.2 shows some of these curves. By roto-translation these curves can be made to go through any point in $\mathbb{M}_2$ which does not change their speed w.r.t. the left invariant frame.

Remark. *The reason that other spaces like $\mathbb{R}^2$ or $\mathbb{M}_{d>2}$ do not admit a left invariant frame is because they do not have enough space to accommodate all their possible rotations.*

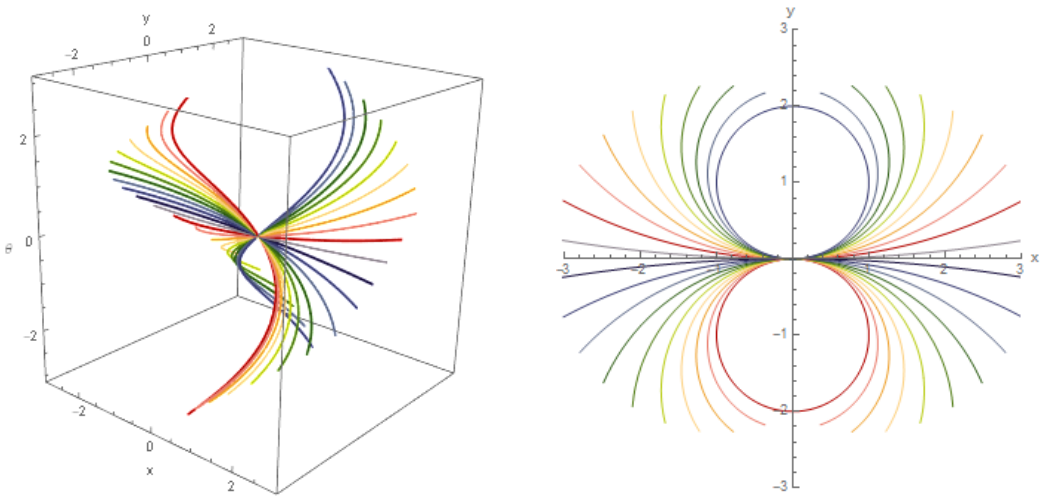

Figure 1.2: Left: plot of some curves with constant speed w.r.t. the left invariant frame. The curves have no velocity in the sideways direction $\mathcal{A}_2$. Right: projection of these curves onto the space of positions $\mathbb{R}^2$.

1.4 Gauge Frames

Instead of the left invariant frame aligning with line structures as a side effect, we can take this to be our guiding principle when constructing a frame. By explicitly fitting to the line structures in an orientation score, we obtain a **locally adaptive frame** which is also referred to as a **gauge frame**. A gauge frame and its dual are denoted by $\{\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3\}$ and $\{\chi^1, \chi^2, \chi^3\}$ respectively. If we need to be explicit about what data $V : \mathbb{M}_2 \to \mathbb{R}$ the frame is fitted to, we write $\{\mathcal{B}_1(V), \mathcal{B}_2(V), \mathcal{B}_3(V)\}$.

As stated before in the Introduction, there are two ways in which the left invariant vector $\mathcal{A}_1$ fails to line up with line structures.

1. Since line structures are lifted to the extended domain $\mathbb{M}_2$ via convolutions with wavelets (recall Definition 6), line structures with a slightly different orientation also yield a response. Hence, the orientation score contains line structures whose orientation differs from their orientation coordinate in the manifold. This phenomenon is called **deviation of horizontality**. In practice this effect is exaggerated even further since we often work with a low angular resolution. See Figure 4 for an illustration.
2. The left invariant vector field $\mathcal{A}_1$ aligns with line structures in a fixed orientation slice of $\mathbb{M}_2$. A curving line structure moves through multiple of these slices however and the left invariant frame completely ignores the angular component of its tangent vector.

These issues are solved by gauge frames, see Figure 6, but at a cost. Because the frame is directly fitted to data, it is susceptible to noise, especially in the absence of a line structure. Moreover, one needs to apply some stabilization technique to obtain even basic levels of smoothness.

1.4.1 Fitting the First Gauge Vector

To detect line structures, one considers the second order differential structure of an image. Consider an image f as a three-dimensional landscape. Line structures are the ridges in such a landscape. Taking a cross section of a ridge reveals the profile of a bump function. As we move along the ridge to the top of the mountain, the vertical position of the bump function might change, but its profile stays roughly the same. To ignore (linear) increases in the ridge's height, we consider the differential of the image df. Ridges in the original image still appear in the differential, albeit with a different profile. Since this profile remains the same along the ridge, the direction of the ridge is found by minimizing a directional derivative of df. See Figure 1.3 for an illustration of what goes wrong when minimizing the directional

derivative of the image f instead of the derivative df. The method to detect line structures described here, is captured in the following definition of the first gauge vector.

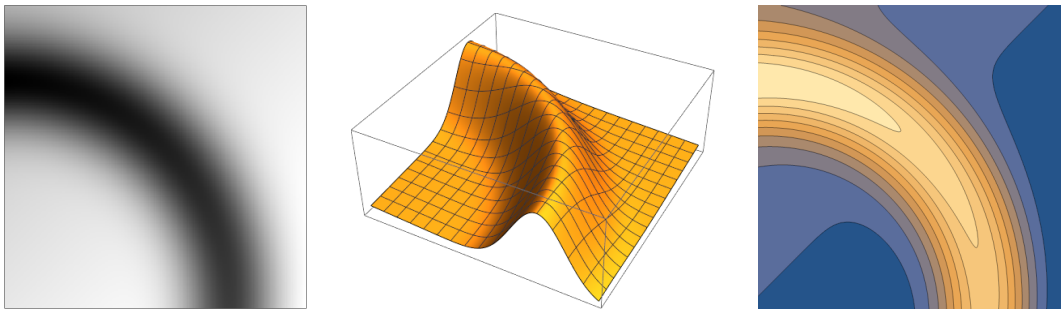

Figure 1.3: A ridge visualized. From left to right: as image, as 3D landscape and as contour plot. Note that minimizing the directional derivative of the image f returns vectors that follow the level curves in the contour plot. At the top of the ridge, the level curves become salient curves and their derivative is orthogonal to the ridge's direction.

Definition 9 (First gauge vector)**.** Given a function $V : \mathbb{M}_2 \to \mathbb{R}$, the first gauge vector at the point $x \in \mathbb{M}_2$ is given by the optimization problem

$$\mathcal{B}_1\big|_x = \underset{\substack{c \in T_x\mathbb{M}_2 \\ ||c||=1}}{\arg\min} \left|\left| \left(\nabla_c^* \, dV\right)\big|_x \right|\right|_*^2 \ , \tag{1.22}$$

where ∇^* is a connection on the cotangent bundle $T^*\mathbb{M}_2$ and the norms are induced by a metric tensor field on $\mathbb{M}_2$.

Metric Tensor Field

To fit the gauge vectors, we use the following metric tensor field.

Definition 10. The **weighted Euclidean metric tensor field** $\mathcal{G}_\xi$ is given by

$$\mathcal{G}_\xi = \xi^2 \left(dx \otimes dx + dy \otimes dy\right) + d\theta \otimes d\theta \ . \tag{1.23}$$

The parameter $\xi > 0$ is used to scale the spatial dimensions against the angular dimension of $\mathbb{M}_2$, as shown by $||\partial_x|| = ||\partial_y|| = \xi$ and $||\partial_\theta|| = 1$.

Alternatively, the weigthed Euclidean metric tensor field can be expressed as

$$\mathcal{G}_\xi = \xi^2 \left(\omega^1 \otimes \omega^1 + \omega^2 \otimes \omega^2\right) + \omega^3 \otimes \omega^3 \ , \tag{1.24}$$

so by Lemma 5, it is roto-translation invariant.

Connection

For a connection, we utilize the left invariant frame. This frame induces an isomorphism of vector bundles $\varphi : T\,\mathbb{M}_2 \xrightarrow{\cong} \mathbb{M}_2 \times \mathbb{R}^3$ and the latter has a canonical connection given by the exterior derivative d. Pulling the exterior derivative back along φ, we obtain a connection $\overline{\nabla} := \varphi^* d$ on $T\,\mathbb{M}_2$ given by

$$\overline{\nabla}_X Y = \sum_{j=1}^{3} \left\langle dy^j, X \right\rangle \mathcal{A}_j = \sum_{j=1}^{3} X(y^j)\, \mathcal{A}_j \ , \tag{1.25}$$

for vector fields X and $Y = \sum\limits_{j=1}^{3} y^j \mathcal{A}_j$ in $T\,\mathbb{M}_2$. The corresponding dual connection is defined by imposing a condition similar to the product rule:

$$X \left\langle \lambda, Y \right\rangle = \left\langle \nabla_X^* \lambda, Y \right\rangle + \left\langle \lambda, \nabla_X Y \right\rangle \ , \tag{1.26}$$

for vector fields $X, Y \in \mathfrak{X}(\mathbb{M}_2)$ and covector field $\lambda \in \Omega^1(\mathbb{M}_2)$. Expressing the fields Y and λ in the left invariant frame as $Y = \sum_{j=1}^{3} y^j \mathcal{A}_j$ and $\lambda = \sum_{j=1}^{3} \lambda_j \, \omega^j$ gives us that

$$\left\langle \overline{\nabla}^*_X \lambda, Y \right\rangle = X \left\langle \lambda, Y \right\rangle - \left\langle \lambda, \overline{\nabla}_X Y \right\rangle = \sum_{j=1}^{3} X(\lambda_j y^j) - \sum_{j=1}^{3} \lambda_j X(y^j) = \sum_{j=1}^{3} X(\lambda_j) y^j \ .$$

Hence, the dual connection $\overline{\nabla}^*$ is given by

$$\overline{\nabla}^*_X \lambda = \sum_{j=1}^{3} X(\lambda_j) \, \omega^j \ . \tag{1.27}$$

The 2-form $\nabla^* dV$ is also referred to as the **hessian** $\mathcal{H} V$. To see why, we express the hessian with the dual connection induced by the left invariant frame. The differential dV is decomposed as

$$df = \sum_{j=1}^{3} \mathcal{A}_j(V) \, \omega^j \ ,$$

so by formula (1.27) the hessian $\overline{\mathcal{H}} V$ is given by

$$\overline{\mathcal{H}} V = \sum_{i,j=1}^{3} \mathcal{A}_i \mathcal{A}_j(V) \, \omega^i \otimes \omega^j \ . \tag{1.28}$$

This expression is reminiscent of the familiar hessian matrix with the derivatives replaced by left invariant derivatives.

Solving for the First Gauge Vector

Having explored the concepts involved with fitting the first gauge vector, it is about time that we solve its optimization problem.

Lemma 6 (Solving first gauge vector)**.** *Let $A : T_x\mathbb{M}_2 \to T^*_x\mathbb{M}_2$ denote the map given by $c \mapsto \overline{\mathcal{H}} V\big|_x(c, \cdot\,)$. The first gauge vector is a $\mathcal{G}_\xi$-unit eigenvector of $A^*A : T_x\mathbb{M}_2 \to T_x\mathbb{M}_2$ with the smallest eigenvalue.*

Expressed in the left invariant frame, the first gauge vector is a coordinate vector $\mathbf{c} = (c^1, c^2, c^3)^\top \in \mathbb{R}^3$. It is an eigenvector of the matrix $\mathbf{M}^2\mathbf{H}^\top\mathbf{M}^2\mathbf{H}$ with the smallest eigenvalue and it has length $||\mathbf{M}^{-1}\mathbf{c}||_{\ell_2} = 1$. The matrices $\mathbf{M}$ and $\mathbf{H}$ are given by

$$\mathbf{M} = \operatorname{diag}\{1/\xi, 1/\xi, 1\} \quad \text{and} \quad \mathbf{H}_{ij} = \mathcal{A}_j \mathcal{A}_i \left(f\right)\big|_x \ , \tag{1.29}$$

where i and j denote rows and columns respectively.

Proof. To find the first gauge vector, we need to solve the optimization problem

$$\underset{\substack{c \in T_x\mathbb{M}_2 \\ ||c||=1}}{\arg\min} \ ||Ac\,||_*^2 \ .$$

Note that the problem is convex, so it suffices to find the local minima using the Euler-Lagrange method. The Lagrangian of the optimization problem is given by

$$L(c, \lambda) = ||Ac\,||^2 - \lambda \left(||c||^2 - 1\right)$$

whose derivative w.r.t. $c \in T_x\mathbb{M}_2$ is expressed using the adjoint operator A^*:

$$\nabla_c L(c, \lambda) = 2(A^*A)c - 2\lambda c \ .$$

Note that ∇_c denotes the directional derivative in the vector space $T_x\mathbb{M}_2$, not a connection on $T\,\mathbb{M}_2$. The directional derivative vanishes if c is an eigenvector of A^*A,

$$(A^*A)c = \lambda c \ .$$

Substituting this and the constraint $||c|| = 1$ into the goal function, we find that the value for an eigenvector is determined by its eigenvalue:

$$||Ac\,||_*^2 = \mathcal{G}_\xi(A^*Ac, c) = \lambda\,||c||^2 = \lambda \ .$$

Since the operator A^*A is self-adjoint and positive definite, all its eigenvalues are real valued and positive. Hence, the optimization problem is solved by an eigenvector with the smallest eigenvalue.

Now for the coordinate expression, we consider the operator A. With the left invariant frame for its domain $T_x\mathbb{M}_2$ and the left invariant dual frame for its codomain $T_x^*\mathbb{M}_2$, we find that A corresponds to the matrix $\mathbf{H}$, since

$$A\left(\mathcal{A}_j\big|_x\right) = \overline{\mathcal{H}}\,V(\mathcal{A}_j\big|_x,\,\cdot\,) \overset{(1.28)}{=\!=} \sum_{i=1}^{3} \mathcal{A}_j\mathcal{A}_i(V)\big|_x\omega^i = \sum_{i=1}^{3} \mathbf{H}_{ij}\,\omega^i \ .$$

However, to quickly obtain the matrix corresponding to the adjoint A^*, the domain and codomain of A need to be expressed in orthonormal frames. The left invariant frame is already orthogonal (see the alternate expression (1.24)), so the in- and outputs of the matrix $\mathbf{H}$ just need to be scaled by the matrix $\mathbf{M}$. As a result, the operator A corresponds to the matrix $\mathbf{MHM}$, and its dual A^* corresponds to the transposition $(\mathbf{MHM})^\top$. Expressed in the normalized left invariant frame, the solution to the optimization problem is as follows. Recall that $\mathbf{c} = (c^1, c^2, c^3)^\top$ represents the first gauge vector in the left invariant frame. Scaled to the orthonormal frame, $\mathbf{M}^{-1}\mathbf{c}$ is an eigenvector of $(\mathbf{MHM})^\top(\mathbf{MHM})$ with the smallest eigenvalue λ. The eigenvalue problem

$$\lambda\,\mathbf{M}^{-1}\mathbf{c} = (\mathbf{MHM})^\top(\mathbf{MHM})(\mathbf{M}^{-1}\mathbf{c}) = \mathbf{M}\mathbf{H}^\top\mathbf{M}^2\mathbf{H}\,\mathbf{c}$$

is equivalent to

$$\lambda\,\mathbf{c} = \mathbf{M}^2\mathbf{H}^\top\mathbf{M}^2\mathbf{H}\,\mathbf{c} \ .$$

The length constraint $||c||_{\mathcal{G}_\xi} = 1$ translates to $||\mathbf{M}^{-1}\mathbf{c}||_{\ell_2} = 1$. ■

Remark. *In areas without a clear line structure the first gauge vector might not be unique. Moreover, for any eigenvector c that solves the optimization problem, the problem is also solved by its opposite $-c$. Because of this, gauge frames might flip suddenly even when a clear line structure is present. To prevent this, we multiply the fitted eigenvector c with the sign of its first component $\operatorname{sign}\omega^1(c)$ such that the resulting vector has a positive forward component. The left invariant frame is smooth and we already expect the forward vector field $\mathcal{A}_1$ to be somewhat aligned with the line structure.*

1.4.2 Constructing the Gauge Frame

Having fitted the first gauge vector to the line structure, we have a lot of freedom to complete the frame. Even if we require that the gauge frame is orthonormal w.r.t. the weighted Euclidean metric tensor field and that it has a right-handed orientation, we can freely spin the frame around the first gauge vector. As a solution, we use the left invariant frame to pinpoint a certain gauge frame. The first gauge vector can be seen as a correction of the first left invariant vector, so we might look at the entire gauge frame from this perspective.

Definition 11 (Remaining gauge vectors)**.** Let the first gauge vector be given by $\mathcal{B}_1\big|_x = c = \sum_{j=1}^{3} c^j\,\mathcal{A}_j\big|_x$.

Let $P\,\mathcal{B}_1\big|_x$ be the projection of the first gauge frame on the spatial part of $T_x\mathbb{M}_2$. The curvature of the line structure is expressed by the angle between the first gauge frame and its projection

$$\nu := \measuredangle(\mathcal{A}_1\big|_x, P\,\mathcal{B}_1\big|_x) = \arctan\left(\frac{c^3}{\sqrt{(c^1)^2 + (c^2)^2}}\right)$$

and the deviation of horizontality is expressed by the angle between the projected first gauge vector and the forward direction:

$$\chi := \angle(P\,\mathcal{B}_1\big|_x, \mathcal{A}_1\big|_x) = \arctan\left(\frac{c^2}{c^1}\right) .$$

Using these angles, the left invariant frame is corrected such that the forward vector lines up with the first gauge vector.

1. Correct for curvature by anticlockwise rotation of $-\nu$ radians around the left vector $\mathcal{A}_2\big|_x$.
2. Correct for deviation of horizontality by anticlockwise rotation of χ radians around the (previous) angular vector $\mathcal{A}_3\big|_x$.
3. Normalize the vectors of the new frame w.r.t. the weighted Euclidean metric tensor field.

The rotation steps are illustrated in Figure 1.4. They are expressed in the left invariant frame by the rotation matrices

$$\mathbf{R}_{\mathbf{e}_2,-\nu} = \begin{pmatrix} \cos\nu & 0 & -\sin\nu \\ 0 & 1 & 0 \\ \sin\nu & 0 & \cos\nu \end{pmatrix} \qquad \text{and} \qquad \mathbf{R}_{\mathbf{e}_3,\chi} = \begin{pmatrix} \cos\chi & -\sin\chi & 0 \\ \sin\chi & \cos\chi & 0 \\ 0 & 0 & 1 \end{pmatrix} .$$

After the rotation $\mathbf{R}_{\mathbf{e}_3,\chi}\mathbf{R}_{\mathbf{e}_2,-\nu}$ the coordinates of the second and third left invariant vector are given by

$$(d^1,d^2,d^3)^\top := \mathbf{R}_{\mathbf{e}_3,\chi}\mathbf{R}_{\mathbf{e}_2,-\nu}\,(0,1,0)^\top \qquad \text{and} \qquad (e^1,e^2,e^3)^\top := \mathbf{R}_{\mathbf{e}_3,\chi}\mathbf{R}_{\mathbf{e}_2,-\nu}\,(0,0,1)^\top$$

respectively. The second and third gauge vector are then defined as

$$\mathcal{B}_2\big|_x := \frac{\sum_{j=1}^3 d^j\mathcal{A}_j\big|_x}{\left|\left|\sum_{j=1}^3 d^j\mathcal{A}_j\big|_x\right|\right|_{\mathcal{G}_\xi}} \qquad \text{and} \qquad \mathcal{B}_3\big|_x := \frac{\sum_{j=1}^3 e^j\mathcal{A}_j\big|_x}{\left|\left|\sum_{j=1}^3 e^j\mathcal{A}_j\big|_x\right|\right|_{\mathcal{G}_\xi}} .$$

Remark. *Since the directions of the second and third gauge vector are determined by a rotation of the the left invariant frame, the gauge frame is orthonormal w.r.t. the weighted Euclidean metric tensor field and it has a right-handed orientation. Furthermore, the second gauge frame lies in the horizontal part of the tangent space, i.e.* $\mathcal{B}_2 \in \operatorname{span}\{\mathcal{A}_1,\mathcal{A}_2\}$.

1.4.3 External Regularization

One way to smooth the gauge frame is by external regularization. The hessian is replaced by a locally averaged hessian such that the first gauge vectors are not fitted in isolation of each other.

Definition 12 (External regularization)**.** Let the scaling matrix $\mathbf{M}$ and the hessian matrices $\mathbf{H}\big|_x$ be given as in (1.29). Consider the component-wise Gaussian smoothed hessians given by

$$\widetilde{\mathbf{H}}\Big|_{(\mathbf{x},\theta)} := \int_{\mathbb{R}^2}\int_{S^1} G^{\mathbb{R}^2}_{\rho_s}(\mathbf{x}-\mathbf{y})\,G^{S^1}_{\rho_a}(\theta-\phi)\,\mathbf{H}\big|_{(\mathbf{y},\phi)}\,d\sigma(\phi)d\mathbf{y} , \tag{1.30}$$

where σ is the normalized Borel measure on the sphere S^1 and the smoothing functions G^M_ρ are isotropic heat kernels centred around origins on manifolds M with time parameter $\rho > 0$. See Section 2.7 of [7] for more details.

The externally regularized first gauge vector $c = \sum_{j=1}^3 c^j\mathcal{A}_j\big|_x$ is then given by the coordinate vector $\mathbf{c} = (c^1,c^2,c^3)^\top \in \mathbb{R}^3$ of length $||\mathbf{M}^{-1}\mathbf{c}||_{\ell_2} = 1$ that is an eigenvector of the smoothed matrix $\mathbf{M}^2\widetilde{\mathbf{H}}^\top\mathbf{M}^2\widetilde{\mathbf{H}}$ with the smallest eigenvector.

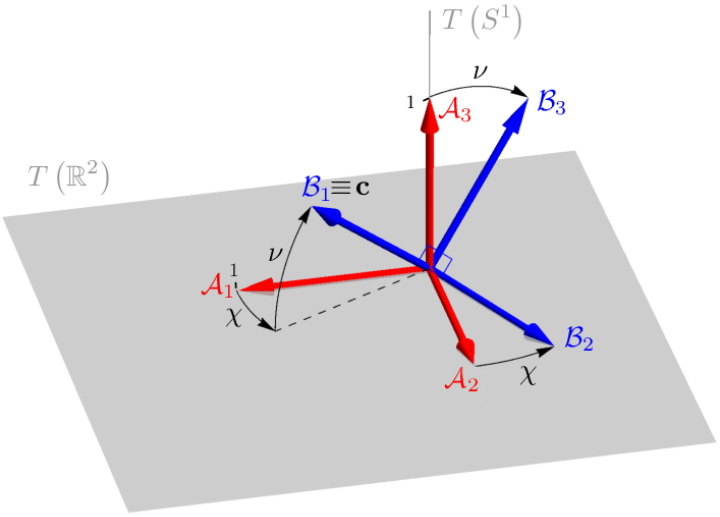

Figure 1.4: Rotations performed to turn the left invariant frame into the (not yet normalized) gauge frame. Remark how the angles ν and χ express curvature and deviation of horizontality respectively. Adapted from [7].

1.4.4 Iterated Stabilization

The second method to smooth gauge frames is to iteratively fit them. First, a gauge frame is fitted to an orientation score. From this gauge frame a measurement is derived that expresses which gauge vectors $\mathcal{B}_1\big|_x$ actually follow a line structure. A second gauge frame is fitted to this measurement and the process continues from here. The measurement described is called the **orientation confidence** of a gauge frame.

Definition 13 (Orientation confidence)**.** Given a gauge frame $\{\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3\}$ fitted to an orientation score $V : \mathbb{M}_2 \to \mathbb{R}$, consider the Laplacian in the plane orthogonal to the principal direction $\mathcal{B}_1\big|_x$

$$(\Delta_o V)(x) := (\mathcal{B}_2\mathcal{B}_2\, V)(x) + (\mathcal{B}_3\mathcal{B}_3\, V)(x)\ .$$

The orientation confidence $S(V) : \mathbb{M}_2 \to [0,1]$ is defined as

$$S(V)\,(x) := \frac{(-(\Delta_o V))_+(x)}{||(-(\Delta_o V))_+||_\infty}\ , \tag{1.31}$$

where $(\,\cdot\,)_+$ denotes the non-negative part of a function, for example $f_+(x) = \max\{0, f(x)\}$.

Why this definition? Recall that if the first gauge vector $\mathcal{B}_1\big|_x$ lines up with the line structure, the cross section of orientation score orthogonal to the first gauge vector shows a bump function, see the middle column of Figure 1.3. The amount to which a 'bump' is present is measured by the negative Laplacian $-\Delta$. Suppose now that the orientation score contains two parallel line structures. The first gauge vector in-between these lines also aligns with them, resulting in a negative response of the negative Laplacian. This response should be ignored since it does not represent a genuine line structure and this is achieved by only considering the non-negative part $(-\Delta)_+$. Finally, the responses are scaled to fit in the unit interval.

Note that the orthogonal Laplacian is efficiently computed from the hessian matrix $\mathbf{H}\big|_x$ (1.29) and the coordinates of gauge vectors in the left invariant frame. Both are already computed in the process of fitting to a function $V : \mathbb{M}_2 \to \mathbb{R}$. Let a gauge vector $\mathcal{B}_i\big|_x$ be given by the coordinate vector $\mathbf{c} = (c^1, c^2, c^3)^\top \in \mathbb{R}^3$ in the left invariant frame. The second order derivative of V in the direction of $\mathcal{B}_i\big|_x$ is then given by

$$(\mathcal{B}_i\mathcal{B}_i\, V)(x) = \mathbf{c}^\top \mathbf{H}\big|_x \mathbf{c}\ .$$

The iterated stabilization of gauge frames is as follows. Given an orientation score $V : \mathbb{M}_2 \to \mathbb{R}$, an initial gauge frame is fitted to V with resulting orientation confidence $S(V)$. Next, a gauge frame is fitted to the

orientation confidence $S(V)$ with resulting orientation confidence $S(S(V)) = S^2(V)$. This process can be repeated ad infinitum. In practice the computational cost of fitting additional gauge frames outweighs the benefit of improved smoothness after four gauge frames.

This iterative refinement scheme has been proposed in Section 11.3.3 in [2], but a formal proof of convergence $S^n(V) \to S^*$ has not been published yet. Possible proofs rely on Banach contraction theorems, but this is beyond the scope of this report.

1.4.5 Roto-Translation Invariance

For any metric based on a gauge frame to be roto-translation invariant, the frame itself should also be roto-translation invariant (see Lemma 5). Showing invariance for gauge frames requires a bit more work than for the left invariant frame.

Lemma 7. *For any orientation score $V : \mathbb{M}_2 \to \mathbb{R}$ the first gauge vector satisfies*

$$(L_g)_* \, \mathcal{B}_1(V) = \mathcal{B}_1(\mathcal{L}_g V) \ , \tag{1.32}$$

for all roto-translations $g \in SE(2)$.

Proof. Consider optimization problem (1.22) defining the first gauge vector:

$$\mathcal{B}_1\big|_x = \underset{\substack{c \in T_x\mathbb{M}_2 \\ ||c||=1}}{\arg\min} \left|\left| \left(\overline{\nabla}^*_c \, dV\right) \Big|_x \right|\right|^2_{\mathcal{G}^*_\xi} \ .$$

Let $\bar{c}$ denote a vector minimizing this problem in $x \in \mathbb{M}_2$ for orientation score V. We need to show that the roto-translated vector $(L_g)_*\bar{c}$ also minimizes the roto-translated optimization problem in $L_g x$ for the score $\mathcal{L}_g V$. Because the weighted Euclidean metric tensor field $\mathcal{G}_\xi$ is roto-translation invariant (see (1.24)), we have that for any vector $c \in T_x\mathbb{M}_2$

$$||\, c \,||_{\mathcal{G}_\xi} = 1 \qquad \Leftrightarrow \qquad ||(L_g)_* c \,||_{\mathcal{G}_\xi} = 1 \ ,$$

so a roto-translated solution satisfies the constraints of the roto-translated optimization problem. Moreover, invariance of the metric tensor field $\mathcal{G}_\xi$ also implies that the dual metric is roto-translation invariant in the sense that

$$\mathcal{G}^*_\xi\big|_x((L_g)^*\lambda, (L_g)^*\mu) = \mathcal{G}^*_\xi\big|_{L_g x}(\lambda, \mu) \tag{1.33}$$

for covectors $\lambda, \mu \in T^*_{L_g x}\mathbb{M}_2$. Hence, it suffices to show for any vector $c \in T_x\mathbb{M}_2$ that

$$(L_g)^* \left(\overline{\nabla}^*_{(L_g)_* c} \, d(\mathcal{L}_g V)\right) \Big|_{L_g x} = \left(\overline{\nabla}^*_c \, dV\right) \Big|_x \ .$$

Expanding the definition of $\overline{\nabla}^*$ (1.27) and the exterior derivative d, we find that this is equivalent to showing the equality of coordinates

$$(L_{g*}c)\,(\mathcal{A}_i(\mathcal{L}_g V)) = c\,(\mathcal{A}_i(V)) \tag{1.34}$$

and the equality of covectors

$$(L_g)^* \omega^i\big|_{L_g x} = \omega^i\big|_x \ . \tag{1.35}$$

We immediately see that the latter holds by evaluation on the left invariant frame. To prove (1.34), let γ_i be an integral curve to $\mathcal{A}_i\big|_x$, i.e. $\dot{\gamma}_i(0) = \mathcal{A}_i\big|_x$. We then have that

$$\begin{aligned} \mathcal{A}_i\big|_{L_g x}(\mathcal{L}_g V) &= \left((L_g)_* \mathcal{A}_i\big|_x\right)(\mathcal{L}_g V) = \frac{d}{dt}\bigg|_{t=0} (\mathcal{L}_g V)\,(L_g \gamma_i(t)) \\ &= \frac{d}{dt}\bigg|_{t=0} V\,(L_{g^{-1}}(L_g \gamma_i(t))) = \frac{d}{dt}\bigg|_{t=0} V(\gamma_i(t)) = \mathcal{A}_i\big|_x(V) \ . \end{aligned}$$

Now let γ be an integral curve of $c \in T_x\mathbb{M}_2$. Using the equality above, we show the equality of coordinates:

$$(L_{g*}c)\,(\mathcal{A}_i(\mathcal{L}_g V)) = \frac{d}{dt}\bigg|_{t=0} \mathcal{A}_i\big|_{L_g\gamma(t)}(\mathcal{L}_g V) = \frac{d}{dt}\bigg|_{t=0} \mathcal{A}_i\big|_{\gamma(t)}(V) = c\,(\mathcal{A}_i(V)) \ .$$

■

Corollary 7.1. *The gauge frame is roto-translation invariant, i.e. for any orientation score $V : \mathbb{M}_2 \to \mathbb{R}$, the gauge vectors satisfy*

$$(L_g)_* \mathcal{B}_i(V) = \mathcal{B}_i(\mathcal{L}_g V) \ , \tag{1.36}$$

for all roto-translations $g \in SE(2)$.

Proof. The directions of the remaining gauge vectors $\mathcal{B}_2$ and $\mathcal{B}_3$ are determined by the coordinates of the first gauge vector in the left invariant frame. Because both the first gauge vector and the left invariant frame are roto-translation invariant, these directions are as well. Since the weighted Euclidean metric tensor field $\mathcal{G}_\xi$ used to normalize these directions is roto-translation invariant as well, the remaining gauge vectors are also invariant. ■

This is nice, but in practice we use smoothed gauge frames. Fortunately, these are also invariant under roto-translations as we show next.

Lemma 8. *The externally regularized gauge frame is roto-translation invariant.*

Proof. According to Definition 12, the externally regularized gauge frame is roto-translation invariant if the smoothed hessian matrices $\widetilde{\mathbf{H}}(V)$ satisfy the equation

$$\widetilde{\mathbf{H}}(V)\Big|_x = \widetilde{\mathbf{H}}(\mathcal{L}_g V)\Big|_{L_g x} \ .$$

This equation is satisfied by the non-smoothed hessians $\mathbf{H}(V)\big|_x$ as by the proof of equation (1.34) we have that

$$\Big(\mathbf{H}(\mathcal{L}_g V)\big|_{L_g x}\Big)_{ij} = (\mathcal{A}_i \mathcal{A}_j(\mathcal{L}_g V))(L_g x) \overset{(1.34)}{=\!=} (\mathcal{A}_i \mathcal{A}_j(V))(x) = \big(\mathbf{H}(V)\big|_x\big)_{ij} \ .$$

Next, consider formula (1.30) of the smoothed hessian

$$\widetilde{\mathbf{H}}\Big|_{(\mathbf{x},\theta)} := \int_{\mathbb{R}^2} \int_{S^1} G^{\mathbb{R}^2}_{\rho_s}(\mathbf{x}-\mathbf{y})\, G^{S^1}_{\rho_a}(\theta-\phi)\, \mathbf{H}\big|_{(\mathbf{y},\phi)}\, d\sigma(\phi)\, d\mathbf{y} \ ,$$

We need to show that besides the hessian, the rest of the integral is also roto-translation invariant. Let the roto-translation be given by $g = (\mathbf{z}, \mathbf{R}_\alpha)$. The heat kernel on $\mathbb{R}^2$ is rotationally invariant, so we have that $G^{\mathbb{R}^2}_{\rho_s}(\mathbf{x}-\mathbf{y}) = G^{\mathbb{R}^2}_{\rho_s}(\mathbf{R}_\alpha(\mathbf{x}-\mathbf{y}))$. The circle measure σ is also rotationally invariant and the spatial measure is roto-translationally invariant. Combinig this, we have that

$$\begin{aligned}
\widetilde{\mathbf{H}}(\mathcal{L}_g V)\Big|_{L_g(\mathbf{x},\theta)} &= \int_{\mathbb{R}^2} \int_{S^1} G^{\mathbb{R}^2}_{\rho_s}(\mathbf{z}+\mathbf{R}_\alpha\mathbf{x}-\mathbf{y})\, G^{S^1}_{\rho_a}(\alpha+\theta-\phi)\, \mathbf{H}(\mathcal{L}_g V)\big|_{(\mathbf{y},\phi)}\, d\sigma(\phi)\, d\mathbf{y} \\
&= \int_{\mathbb{R}^2} \int_{S^1} G^{\mathbb{R}^2}_{\rho_s}(\mathbf{R}_\alpha(\mathbf{x}-\mathbf{y}))\, G^{S^1}_{\rho_a}(\theta-\phi)\, \mathbf{H}(\mathcal{L}_g V)\big|_{L_g(\mathbf{y},\phi)}\, d\sigma(\alpha+\phi)\, d(\mathbf{z}+\mathbf{R}_\alpha\mathbf{y}) \\
&= \int_{\mathbb{R}^2} \int_{S^1} G^{\mathbb{R}^2}_{\rho_s}(\mathbf{x}-\mathbf{y})\, G^{S^1}_{\rho_a}(\theta-\phi)\, \mathbf{H}(V)\big|_{(\mathbf{y},\phi)}\, d\sigma(\alpha+\phi)\, d(\mathbf{z}+\mathbf{R}_\alpha\mathbf{y}) \\
&= \int_{\mathbb{R}^2} \int_{S^1} G^{\mathbb{R}^2}_{\rho_s}(\mathbf{x}-\mathbf{y})\, G^{S^1}_{\rho_a}(\theta-\phi)\, \mathbf{H}(V)\big|_{(\mathbf{y},\phi)}\, d\sigma(\phi)\, d\mathbf{y} \\
&= \widetilde{\mathbf{H}}(V)\Big|_{(\mathbf{x},\theta)} \ .
\end{aligned}$$

■

Lemma 9. *Orientation confidences are roto-translation equivariant, i.e.*

$$S(\mathcal{L}_g V) = \mathcal{L}_g(S(V)) \ ,$$

so the iteratively fitted gauge frames are roto-translation invariant.

Proof. The orientation confidences $S(V)$ and $S(\mathcal{L}_g V)$ are computed from the corresponding gauge frames that satisfy

$$(L_g)_* \, \mathcal{B}_i(V) = \mathcal{B}_i(\mathcal{L}_g V) \ ,$$

see Corollary 7.1. Following the proof of equation (1.34), we have that for second order derivatives:

$$\mathcal{B}_i(\mathcal{L}_g V)\big|_x \left(\mathcal{B}_i(\mathcal{L}_g V)\left(\mathcal{L}_g V\right)\right) = \mathcal{B}_i(V)\big|_{L_{g^{-1}}x} \left(\mathcal{B}_i(V)\, V\right) \ .$$

This implies the roto-translation equivariance of the orientation confidence score

$$S(\mathcal{L}_g V)\,(x) = \mathcal{L}_g(S(V))\,(x) = S(V)\,(L_{g^{-1}}x) \ .$$

The roto-translational invariance of iterated gauge frames follows since

$$(L_g)_* \, \mathcal{B}_i(S(V))\big|_x = \mathcal{B}_i(\mathcal{L}_g(S(V)))\big|_{L_g x} = \mathcal{B}_i(S(\mathcal{L}_g V))\big|_{L_g x} \ .$$

■

Chapter 2

Vessel Tracking

The goal of this chapter is to explore the extent to which gauge frames (i.e. frames locally fitted to the data) can be used for vessel tracking. We will show some benefits of using these frames and address possible pitfalls.

The motivation for using gauge frames is that one hopes to improve tracking by adapting to deviation of horizontality and curvature. Here we stress that the inclusion of gauge frames improved other geometric processing tasks such as in recent works on vessel segmentation tasks [18] and crossing preserving denoising [13]. Furthermore, gauge frames enable the incorporation of these adaptations in the geometric notions of 'short' and 'straight' curves. How this is done is explained in the next chapter.

We briefly repeat the vessel tracking algorithm described in Section 1.1.5. To deal with intersecting vessels, input images are first lifted to an orientation score. The tracking process takes place in this expanded domain as illustrated by Figure 2.1. A vessel is tracked between two points by computing the shortest curve between them, i.e. a curve that minimizes a length functional

$$d_{\mathcal{G}}(p,q) = \min_{\substack{\gamma \in \mathrm{Lip}([0,1],\mathbb{M}_2) \\ \gamma(0)=p \\ \gamma(1)=q}} \int_0^1 ||\dot{\gamma}(t)||^2_{\mathcal{G}}\, dt \ , \tag{2.1}$$

where $\mathcal{G}$ is a metric tensor field that may depend on the orientation score data. The metric tensor field is chosen such that curves along vessels are assigned a shorter length.

Section 2.1 describes how shortest curves and gauge frames are computed in practice. The gauge frame based metric tensor fields considered in this project are outlined in Section 2.2 as well as their performance on an artificial test image. Section 2.3 shows how these tensor fields perform on actual retinal images.

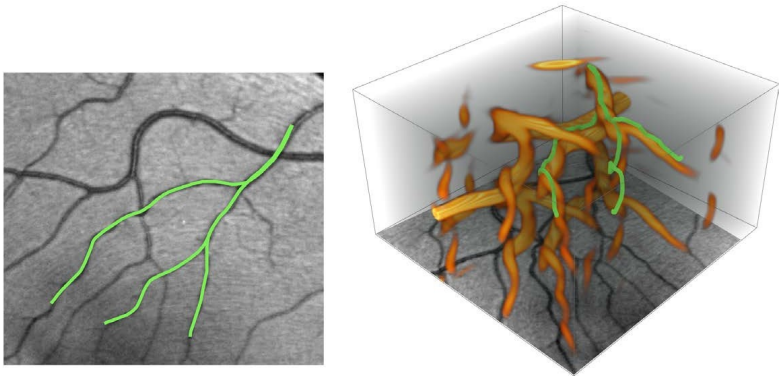

Figure 2.1: Tracking vessels in the orientation score (right) and the projection onto the input image (left).

2.1 Practical Considerations

Although the previous chapter explains the underlying theory of vessel tracking, there are still some practical things to consider.

2.1.1 Computing Shortest Curves

Consider the optimization problem in (2.1). A shortest curve from p to q is computed in two steps. First, we compute the distance from p to every other point $x \in \mathbb{M}_2$, i.e. a distance function $W = d_{\mathcal{G}}(p, \cdot)$. This function is found by solving for the viscosity solution of the eikonal PDE

$$\begin{cases} \|\operatorname{grad} W\|_{\mathcal{G}} = 1 \\ \qquad W(p) = 0 \end{cases} \tag{2.2}$$

where the gradient $\operatorname{grad} W$ is the vector reciprocal to the differential dW:

$$\mathcal{G}(\operatorname{grad} W, \cdot) = dW \ .$$

The shortest path γ is found by steepest descent on the distance map W. This is also referred to as **backtracking**. For the underlying theory, see [14] and [10] for viscosity solutions on the respective spaces $\mathbb{R}^d$ and $\mathbb{M}_3$.

A useful interpretation of the distance function W is in terms of wavefront propagation. The levels curves can be considered as wavefronts of a viscous fluid originating from the starting point p. Like with a viscous fluid, the wavefronts do not propagate through each other, explaining the name viscosity solution. Values at level curves of W, i.e. the distances to p, can be interpreted as the time needed for the fluid to reach those points. As the wavefronts expand, they eventually reach the end point q. The shortest path is found by tracking back along the path traversed by the wavefronts. This is done by steepest descent since the gradient is perpendicular to the level curves.

The propagation of wavefronts is hard to visualize since the fronts are surfaces in the three-dimensional space of positions and orientations $\mathbb{M}_2$. By minimizing the distance function over all orientations, they can be visualized in two dimensions, see Figure 2.2. The projection shows the wavefronts that first hit a position, regardless of orientation. Discarding the orientation data makes it harder to predict the path taken by backtracking, but in practice this is not too big of an issue.

2.1.2 Implementation Notes

The lifting of an input image to an orientation score is performed by the *Lie Analysis* [16] package for Mathematica. This package also provides methods for taking left invariant derivatives that are used throughout.

The eikonal PDE (2.2) is numerically solved by an iterative scheme as described in [4]. The derivatives w.r.t. the gauge frame are computed by linear transformation of the left invariant derivatives obtained using upwind differencing. The backtracking step is then performed on an interpolated gradient field of the distance map. Interpolation is performed by expressing the gradient in terms of the left invariant frame and interpolating the coordinate fields, taking into account the periodicity of the angular dimension.

Interpolating the gradient field need to be done carefully. For efficiency, it is useful to only consider orientations modulo π radians when computing the orientation score and the distance map. This is allowed since this symmetry is satisfied by both the orientation score transformation and the metric tensor fields used. The spatial left invariant vector fields $\mathcal{A}_1$ and $\mathcal{A}_2$ however do not have this symmetry, their rotation is only complete after 2π radians. As a result of the above, the gradient coordinates are computed for the orientations $[0, \pi]$, but they do not have an angular periodicity of π radians. The solution is to copy the gradient coordinates for the orientations $[0, \pi]$ and to multiply the spatial coordinates with -1 to obtain the gradient coordinates for the orientations $[\pi, 2\pi]$. Combined, these coordinates can be interpolated with a periodicity of 2π in the angular dimension.

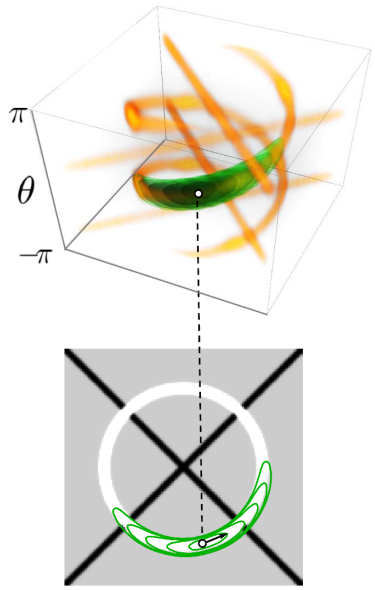

Figure 2.2: Projection of wavefronts in the orientation score to wavefronts on the input image. **Remark 1:** Wavefronts do not propagate to the diagonal lines since these have a different angular coordinate in the orientation score. **Remark 2:** This is an idealized image, in practice the wavefronts always leak through the vessel boundaries. Adapted from [6]

2.1.3 Fitting Gauge Frames in Practice

It turns out that fitting gauge frames to orientation scores is quite sensitive to the choice of the weighting parameter ξ in the weighted Euclidean metric tensor field $\mathcal{G}_\xi$ (see (1.23)). For each new orientation score, we manually tune the parameter ξ by looking at the quality of the gauge frame produced. The main tool used for this task is the orientation confidence (see Definition 13). Its pattern should match that of the line structures in the orientation score. For a more detailed inspection, we use the visualization method outlined below.

Visualizing Alignment

To visualize the alignment of the first gauge vector with line structure in the data, we plot integral curves of this vectorfield at multiple locations on top of the data to see how well these agree. This can be done in the orientation score, but this is hard to visualize in a report since it concerns three-dimensional data. A solution is to project the integral curve onto the input image by discarding the orientation coordinates. Alignment in the angular direction can still evaluated by looking whether the curvature of the integral curves matches that of the line structures in the image. To prevent multiple integral curves being plotted per position, only the integral curve corresponding to the orientation in the space $\mathbb{M}_2$ with the greatest response in the orientation score is shown.

For the integral curve, we use a curve with constant speed w.r.t. the left invariant frame. For the explicit formulas see (1.20). The advantage of these curves is that they have constant curvature which makes it easier to evaluate how well they align with the data.

An example of this visualization is shown in Figure 2.3. Note that the curved line segments are shorter than the straight ones. This is a consequence of the projection as the line segments are of equal length in the space $\mathbb{M}_2$.

Type of Smoothing Used

The previous chapter listed two distinct techniques to smooth a gauge frame: external regularisation (Section 1.4.3) and iterated stabilization (Section 1.4.4). To decide which method to use, we applied both and compared their alignment with the input image, see Figure 2.3. The alignment of the externally regularized gauge frame is worse than that of the non-smoothed frame and the iteratively stabilized frame. Note that it is mainly the curvature of the externally regularized data that does not align well with the data. Based on these results, we decided to only use iterated stabilization if the non-processed gauge frame are too rough to be used for vessel tracking.

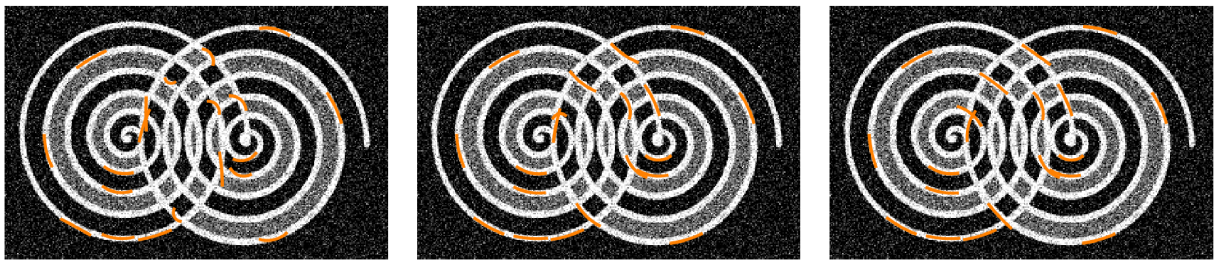

external regularization | not smoothed | iterated stabilization

Figure 2.3: Comparison between alignment of gauge frames smoothed with different methods.

2.2 Experimental Design

Designing metric tensor fields based on gauge frames has been an experimental process. We show a naive tensor field from the early stages of this process and a more sophisticated one from the end of the project. Their performance on a test image is compared with that of left invariant frame based metric tensor fields.

2.2.1 Metric Tensor Fields Used

This is section provides an overview of the metric tensor fields we have tested. To understand why these metrics are defined the way they are, it is helpful to keep the following design principle in mind: *vectors aligned to vessels should have a small norm compared to other vectors.*

Remark. *The coefficients chosen for the metric tensor fields below look slightly awkward. This is because when experimenting with different metrics, we mostly modified the coefficients of the dual metric tensor fields that appear in the expression for the gradient* $\operatorname{grad} W$ *. In the discussion in Section 2.4.1, we propose metric tensor fields for future research with more regular coefficients.*

The following left invariant frame based metric tensor field is used to compare the gauge frame ones against. The metric tensor field is derived from those in [4]. Recall that $\{\omega^1, \omega^2, \omega^3\}$ denotes the frame dual to the left invariant frame.

1. The **left invariant metric tensor field** $\mathcal{G}^{C,\beta}_{\xi,\zeta}$ is given by

$$\mathcal{G}^{C,\beta}_{\xi,\zeta} := \frac{1}{\left(\beta(1/C) + (1-\beta)\right)^2} \left(\xi^2 \left(\omega^1 \otimes \omega^1 + \frac{1}{\zeta^2}\, \omega^2 \otimes \omega^2 \right) + \omega^3 \otimes \omega^3 \right) . \tag{2.3}$$

   Paths outside vessels are discouraged by the cost function $C : \mathbb{M}_2 \to [\delta, 1]$ that is bounded from below by $\delta > 0$. The parameter $\beta \in [0,1]$ weighs the influence of the cost function. At $\beta = 1$, the influence is maximized and the fraction at the front of the metric tensor field evaluates to C^2. Vectors with a large sideways component are penalized by the anisotropy parameter $\zeta \in [0,1]$. As

a result, the paths traced out are comparable to those of a car. In order to go sideways, one first has to make a turn. This makes for natural looking paths even in the absence of a clear vessel structure. For the experiments, we consider two versions of this tensor field: a **isotropic** one with $\zeta = 1$ and an **anisotropic** one with $\zeta = 1/8$.

The next metric tensor field was one of our first attempts to construct a gauge frame based metric tensor field. The main reason for its inclusion is to show that it is a bad choice and that a more sophisticated tensor field is required. Recall that $\{\chi^1, \chi^2, \chi^3\}$ denotes the frame dual to the gauge frame.

2. The **(anisotropic) gauge frame metric tensor field** $\mathcal{G}_\alpha$ is given by

$$\mathcal{G}_\alpha := \chi^1 \otimes \chi^1 + \frac{1}{(1-\alpha)^2} \left(\chi^2 \otimes \chi^2 + \chi^3 \otimes \chi^3 \right) , \tag{2.4}$$

where $\alpha \in [0, 1)$ is an anisotropy parameter. The larger α is, the more vectors orthogonal to the line structure $\mathcal{B}_1$ are penalized. For $\alpha = 0$, the metric tensor field agrees with the weighted Euclidean metric tensor field $\mathcal{G}_\xi$ from (1.23). Note that the weighting parameter ξ implicitly influences the gauge frame metric tensor field since the gauge frames are normalized w.r.t. the weighted Euclidean tensor field $\mathcal{G}_\xi$.

The following metric tensor field is the final iteration of the gauge frame based metric tensor fields designed over the course of the project.

3. The **adaptive (anisotropic) gauge frame metric tensor field** $\mathcal{G}^S_\alpha$ is given by

$$\mathcal{G}^S_\alpha := \frac{1}{S + (1-S)(1-\alpha)^2} \chi^1 \otimes \chi^1 + \frac{1}{(1-\alpha)^2} \left(\chi^2 \otimes \chi^2 + \chi^3 \otimes \chi^3 \right) . \tag{2.5}$$

The parameter α is the same anisotropy parameter as from the early iteration gauge frame based metric tensor field above. The symbol S denotes the orientation confidence of the gauge frame. The orientation confidence takes values in the interval $[0, 1]$ and is used to control the anisotropy of the tensor field. In areas with minimal orientation confidence, where $S = 0$, the first gauge vector is not aligned to a line structure. Hence, there is no motivation to prefer this direction over the others and the tensor field is made isotropic. In areas where the first gauge vector most strongly matches a line structure, where $S = 1$, the anisotropy is maximized.

In areas of low orientation confidence, the anisotropy of the metric tensor field is lowered by scaling up the coefficient in front of the first term $\chi^1 \otimes \chi^1$. As a result, the orientation confidence also acts like a cost function. Finally, remark that since the orientation confidence takes values in the interval $[0, 1]$, the coefficient in front of the first term takes values in the interval $[(1-\alpha)^{-2}, 1]$.

Parameter settings

The left invariant and gauge frame based metric tensor fields above share some parameters and differ in others. Since the goal of the experiments is to compare their performance, the parameters are chosen to enforce some consistency.

Per input image, the weighting parameter ξ is manually tuned by visually evaluating the gauge frames produced as described in Section 2.1.3. The tuned value of ξ is also used for the left invariant metric tensor fields. The values found for the different input images are shown in Table 2.1.

Input image	ξ
Experiment I & II	0.08
Application I	0.08
Application II	0.01

Table 2.1: Tuned values of the weigthing parameter ξ for the different input images.

The orientation confidence acts as a cost function for the adaptive anisotropic gauge frame metric tensor field (2.5). For this reason we choose to also let the orientation confidence determine the cost function used in the left invariant metric. For an orientation confidence S, the corresponding cost function is given by

$$C := \frac{1}{1 + \lambda S^p} ,$$

where we take parameters $\lambda = 200, p = 2$ unless stated otherwise. This transformation is of the same form as the one used in [4] to transform a different vesselness measure into a non-uniform cost function.

2.2.2 Experiment I

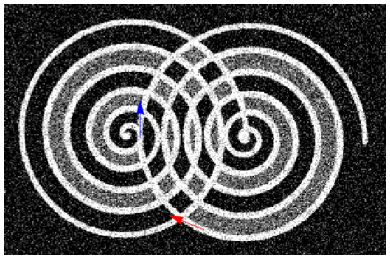

Figure 2.4: The task in experiment I is to trace the curve segment from the red arrow to the blue arrow.

The goal of the first experiment is to get a feeling for how the different metric tensor fields perform. The task is relatively simple. The algorithm needs to cross several intersections along a curved line segment. The exact task is shown in Figure 2.4.

Left Invariant Metric Tensor Fields

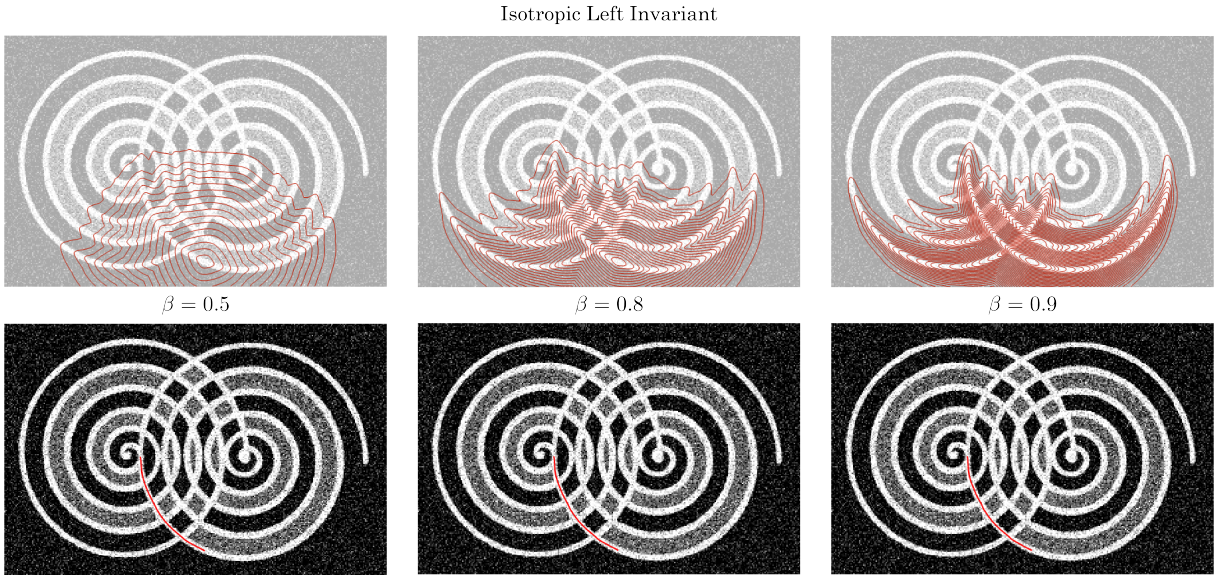

Figure 2.5: Experiment I: Wavefronts and paths produced by the isotropic left invariant metric tensor field showing the influence of the cost function by varying the parameter β.

Figure 2.5 shows the tracking results of the isometric left invariant metric tensor field. The curves traced out by all three variations nicely follow the curved line segment they are supposed to track. For this metric tensor field, the wavefronts are purely steered by the cost function. This simplicity leads to some undesirable behaviour. Consider the wavefronts expanding perpendicular to the line structure from the

starting location. When they encounter the next line structure, their speed increases greatly. Ideally, one would have that fronts only pick up speed when they *align* with the line structure.

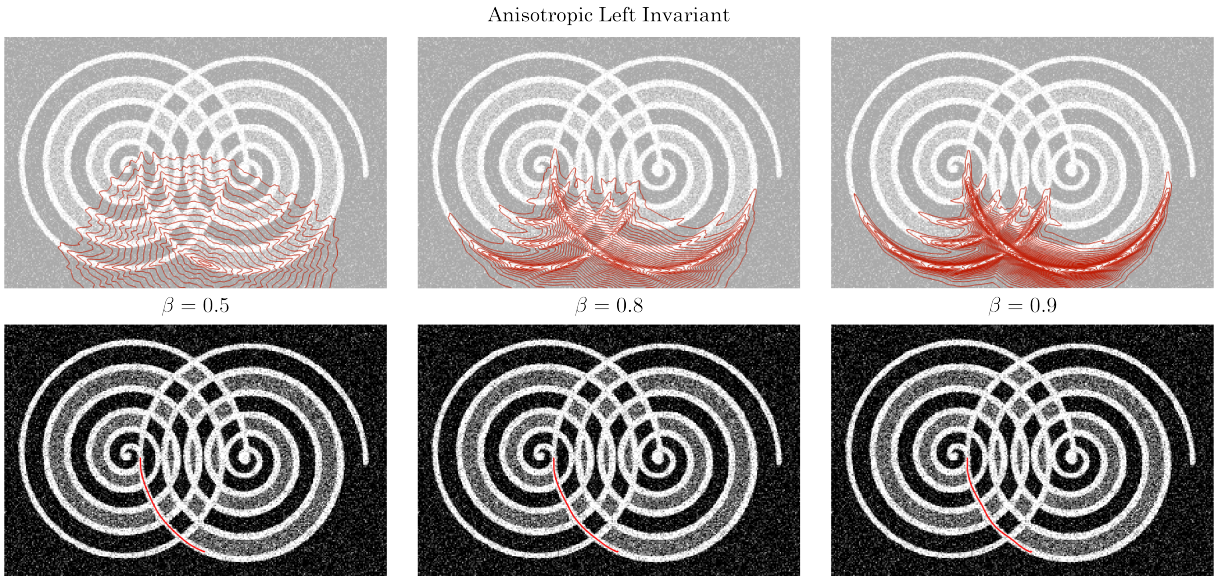

Figure 2.6: Experiment I: Wavefronts and paths produced by the anisotropic left invariant metric tensor field showing the influence of the cost function by varying the parameter β.

The results of the anisotropic left invariant metric tensor field are shown in Figure 2.6. Like with the isotropic metric tensor field, the paths traced out by this tensor field also follow the curved line segment. However, as the influence of the cost function increases, the continuous turn becomes a smooth zigzag. The wavefronts fit more tightly around the line structures as compared to those produced by the isotropic tensor field and note that the problem of accelerating wavefronts mentioned above does not occur here. An issue the anisotropic tensor field does have is that the wavefronts propagate faster in several distinct global directions. This is a side effect of using a low angular resolution for processing on the space $\mathbb{M}_2$.

Gauge Frame Metric Tensor Field

Figure 2.7 shows the results of the gauge frame metric tensor field with increasing anisotropy. The wavefronts of this tensor field do not look nice. Although their general expansion follows the line structures in the image, the fronts themselves are very sharp and zigzag back and forth. The paths traced out still look okay, but they are rougher than those from the left invariant tensor fields. The curve generated with an anisotropy of $\alpha = 0.8$ takes a sharp corner at the first crossing.

We attempted to resolve the sharpness of the wavefronts by using a smoothed gauge frame. This did not have the desired result as the wavefronts still show the same level of sharpness for increasing anisotropy, see Figure 2.8. The paths traced out by this tensor field are quite bad. For $\alpha = 0.5$, the path takes a rough shortcut instead of following the curved line segment. For $\alpha = 0.8$, the backtracking process stops since the interpolated gradient field vanishes. Critical points like this occur when the gradient flips direction from one point to the next, see Figure 2.9. However, the gradient fields used for vessel tracking are not supposed to exhibit such behaviour as they are gradients of viscosity solutions. This issue is discussed in more detail in Section 2.4.1.

Adaptive Gauge Frame Metric Tensor Field

The tracking results of adaptive gauge frame metric tensor field are shown in Figure 2.10. The wavefronts are smooth and follow the line structures in the image. The paths traced also follow the curved line segment, although they get a bit more bumpy as the anisotropy increases.

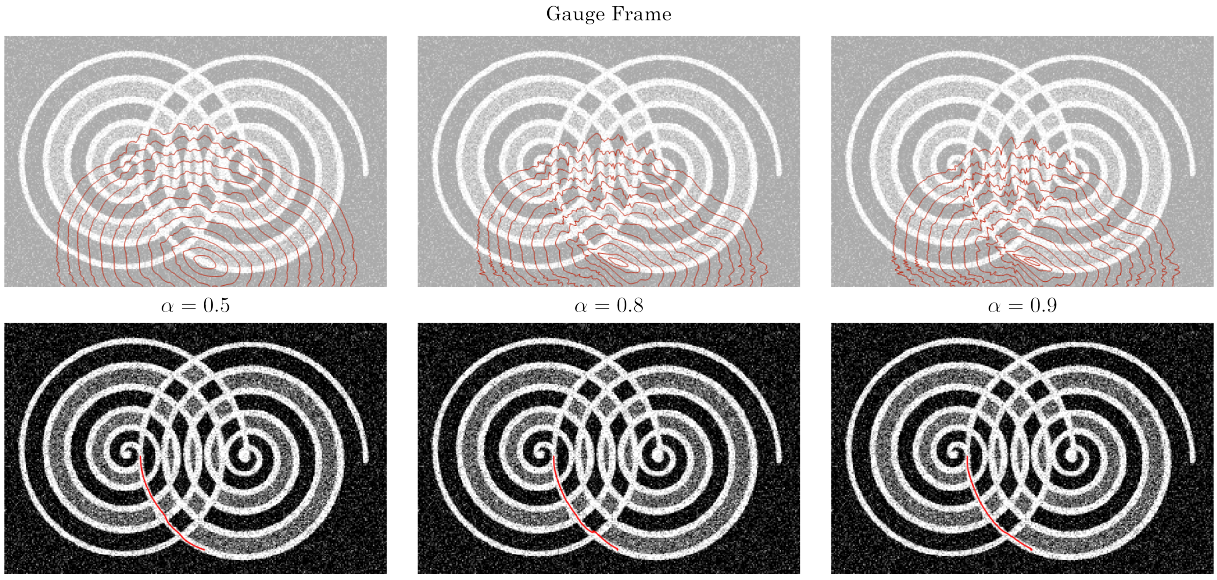

Figure 2.7: Experiment I: Wavefronts and paths produced by the gauge frame metric tensor field showing the influence of anisotropy by varying the parameter α.

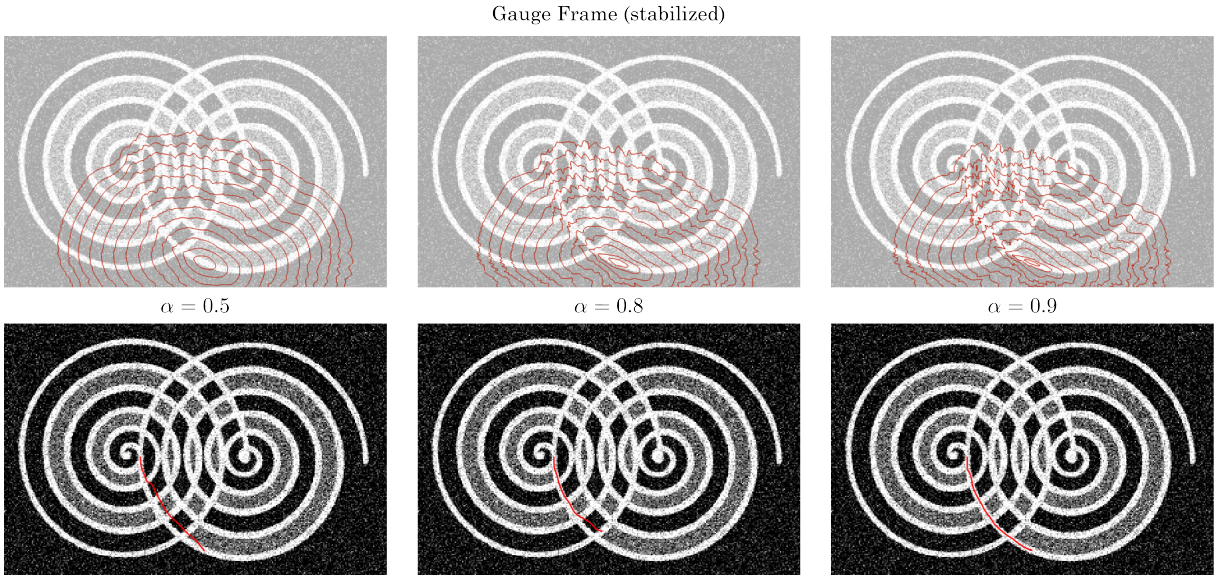

Figure 2.8: Experiment I: Wavefronts and paths produced by the gauge frame metric tensor field *using iteratively stabilized gauge frames* showing the influence of anisotropy by varying the parameter α.

We more closely compare the wavefronts of the adaptive gauge frame tensor field to those of the left invariant metrics. Comparing Figures 2.5 and 2.10, we see that the general shape of the wavefronts of the isotropic left invariant tensor field and the adaptive gauge frame tensor field are very similar for equal parameters $\alpha = \beta$. This is remarkable since the formulas of the metric tensor fields are so different. Figure 2.11 shows the wavefronts of the metric tensor fields side-by-side. Here we see that the left invariant metric tensor fields overly strong promote wavefronts along the outer curves of the spirals, whereas the adaptive gauge frame tensor field does not. Moreover, the adaptive gauge frame tensor field does not suffer from the negative aspects of both left invariant tensor fields. Whereas the isotropic left invariant tensor field allows fast propagation of wavefronts perpendicular to line structures, the adaptive gauge frame tensor field does not. The wavefronts of the adaptive gauge frame tensor field also do not suffer from the low angular resolution that plagues the those of the anisotropic left invariant tensor field.

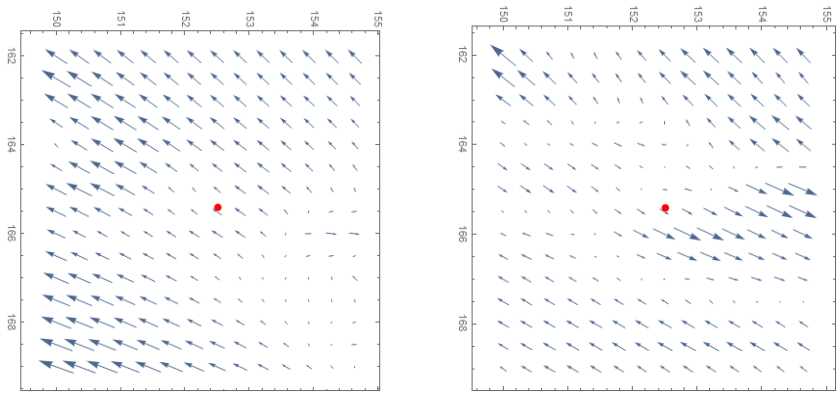

Figure 2.9: Close-up of the gradient fields around the critical point (red dot) where the backtracking algorithm got stuck. The field corresponds to wavefronts and curve in the middle column of Figure 2.8. Left and right show the fields at the discrete orientation levels closest to the critical point. Note how the vectors close to the red dot flip direction from one orientation level to the next.

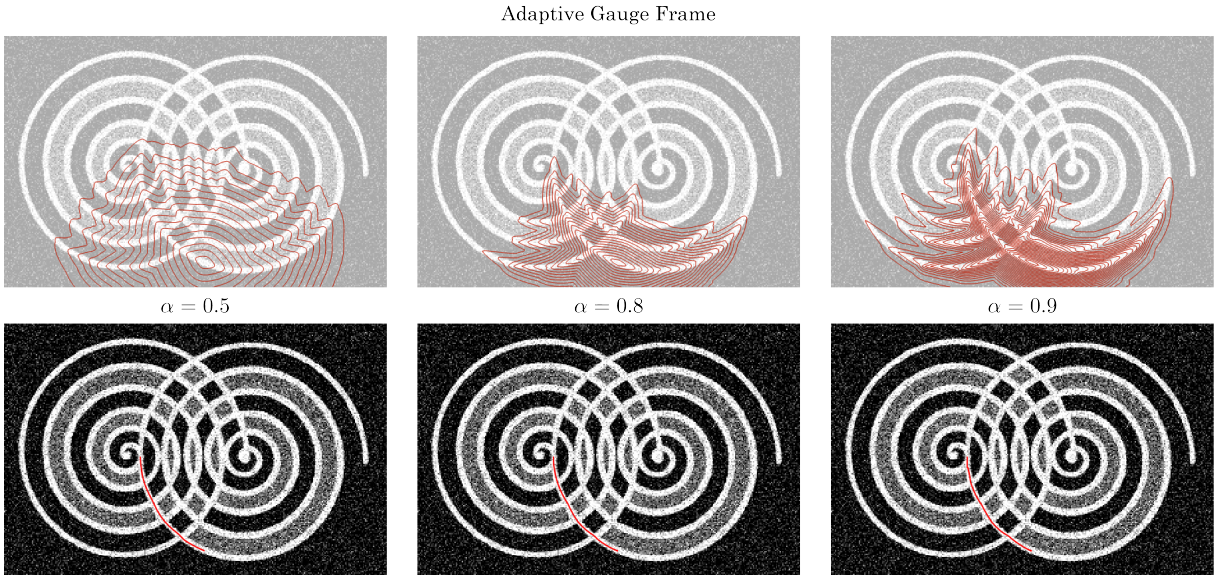

Figure 2.10: Experiment I: Wavefronts and paths produced by the adaptive gauge frame metric tensor field showing the influence of anisotropy by varying the parameter α.

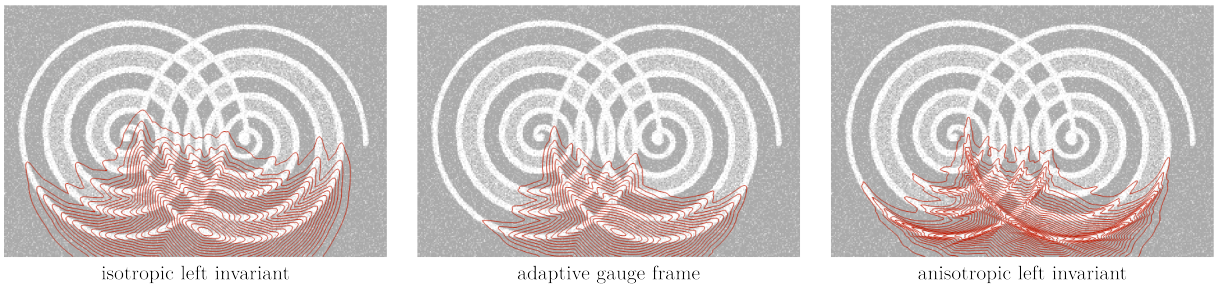

Figure 2.11: Side-by-side comparison of wavefronts produced by different metric tensor fields for parameters $\alpha = \beta = 0.8$.

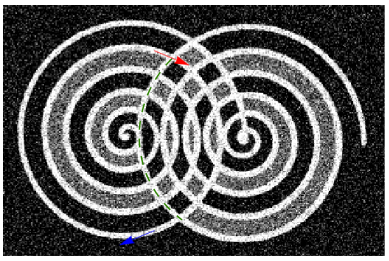

Figure 2.12: The task in experiment II is to trace the longer curve segment from the red arrow to the blue arrow. The possible shortcut is indicated by the dashed green line.

2.2.3 Experiment II

For the second experiment we let the metric tensor fields follow a large curve of the spirals with a large number of intersections. It is more difficult because it is possible to take a shortcut by jumping bewteen curve segments at two intersections, see Figure 2.12. For this experiment, we consider the left invariant metric tensor fields and the adaptive gauge frame tensor field with the parameters $\alpha = \beta = 0.8$. The results are shown in Figure 2.13.

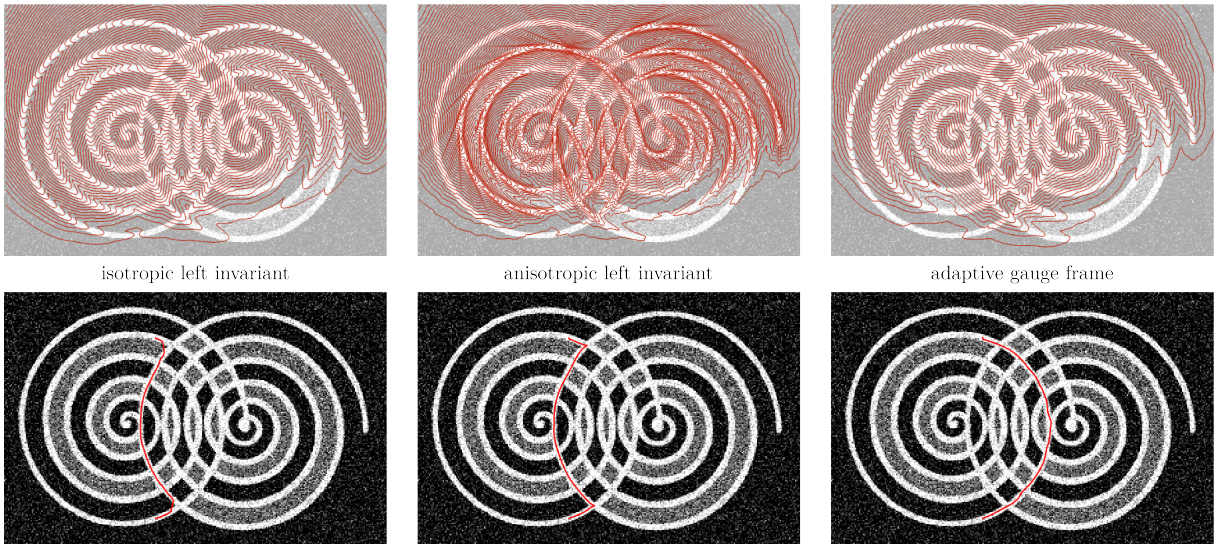

Figure 2.13: Experiment II: Wavefronts and paths produced by the different metric tensor fields. Parameters $\beta = \alpha = 0.8$.

The wavefronts in Figure 2.13 show the same general behaviour as in the previous experiment. The curves however are interesting to look at! Both left invariant tensor fields take the shortcut, but the adaptive gauge frame tensor field takes the correct path, albeit a bit wobbly.

Note that the path of the anisotropic left invariant metric tensor field has two cusps. In this case, the problem of cusps has been thoroughly analysed [8]. They are prevented by only allowing wavefronts to propagate forwards, i.e. in the direction of $\mathcal{A}_1$, not in the direction of $-\mathcal{A}_1$. This method has not been discussed in this report as it uses Finsler functions instead of metric tensor fields to steer the vessel tracking algorithm. For sake of completeness, we also test this improved version of the anisotropic left invariant metric tensor field. The results in Figure 2.14 show this algorithm avoids cusps but is still problematic as it bends towards the wrong line.

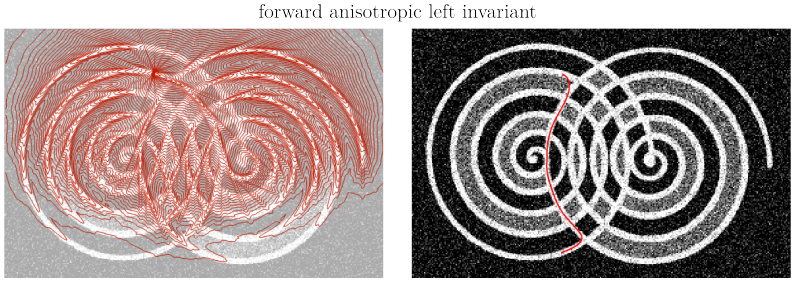

Figure 2.14: Experiment II: Wavefronts and paths produced by the anisotropic left invariant metric tensor field without reverse gear [8]. Parameter $\beta = 0.8$.

2.3 Application to Retina Vessels

The adaptive gauge frame metric tensor field shows promising results on the test image. In this section we explore whether these results hold up when the tracking algorithm is applied to actual retinal images. Two settings are explored that have been proven difficult for other vessel tracking algorithms. The input images used are zoomed-in patches of retina images from the *DRIVE* dataset.

2.3.1 Application I: Low Spatial Resolution

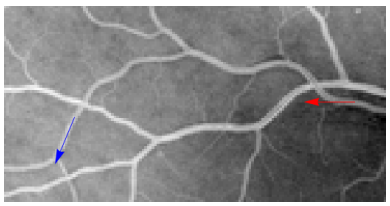

Figure 2.15: The task in the first application is to trace the vessel from the red arrow to the blue arrow. Instead of staying on the correct vessel, the tracking methods might take a shortcut via the lower vessel.

The first task is similar to the setting of Experiment II. The tracking algorithms are expected to follow a certain vessel whilst a shortcut can be taken by jumping to another vessel at an intersection. In addition, the spatial resolution is scaled down to speed up calculations, a useful property for any real life application. The exact task is shown in Figure 2.15.

Remark. *The anisotropic left invariant tensor field is missing in Figure 2.17 since during its computation the numerical solution to the eikonal PDE (2.2) diverged. This is a problem that we will encounter more often when we attempt to strictly prevent the propagation of wavefronts to undesired areas by increasing the parameters α and β.*

The results in Figures 2.16 and 2.17 show that none of the tensor fields pick the correct vessel. Considering the wavefronts, we see that the anisotropic left invariant and adaptive gauge frame metric tensor fields are closer to choosing the correct vessel than the isotropic left invariant one. Their wavefronts have propagated further along the upper vessel before being cut of by those in the lower one. Figure 2.16 also shows that the isotropic left invariant tensor field produces a very jittery path whereas the other two tensor fields do not. Apparantly, the low spatial resolution of a vessel greatly influences the path when the metric tensor field does not punish sideways movement.

The vessel tracking results on this image are a bit surprising. Why did no single metric tensor field steer the curve along the upper vessel? After investigation, the problem turned out to be the orientation confidence used to guide the paths: its value is much greater along the bottom vessel than along the upper one. Curious about the results of an orientation confidence with equal values on both vessels, we modified the orientation confidence S into an 'amplified' version given by:

$$S_{\text{ampl}}(x) := \max\{3\, S(x), 1\} \ .$$

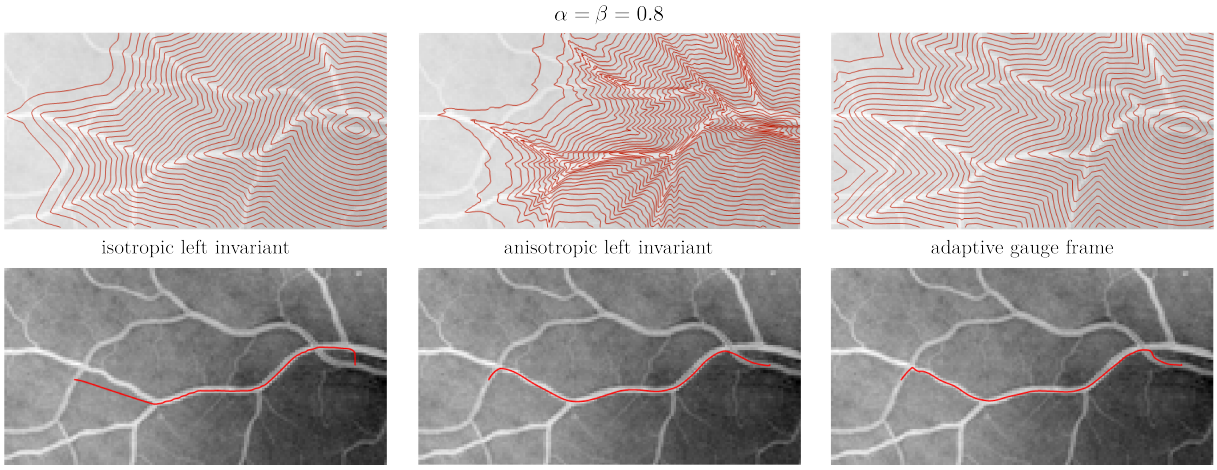

Figure 2.16: Application I: Wavefronts and paths produced by the different metric tensor fields. Parameters $\beta = \alpha = 0.8$.

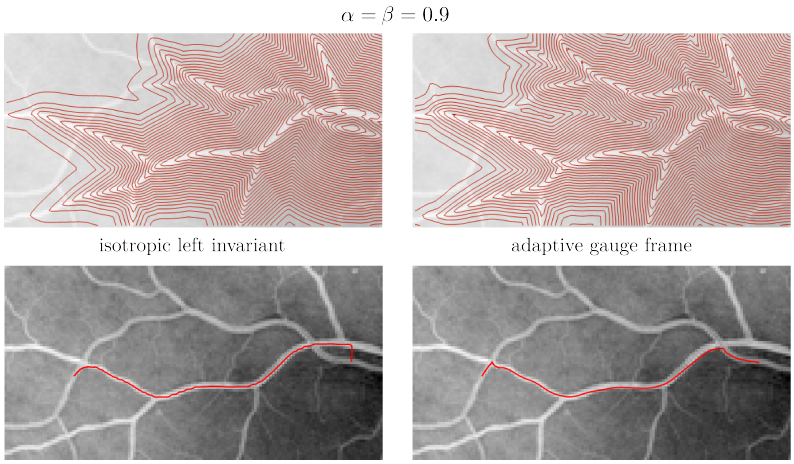

Figure 2.17: Application I: Wavefronts and paths produced by the different metric tensor fields. Parameters $\beta = \alpha = 0.9$. The anisotropic left invariant tensor field is missing since the numerical solution to the eikonal PDE diverged.

The two orientation confidences are shown in Figure 2.18 where the maximum is taken along the angular dimension to display these functions in two dimensions.

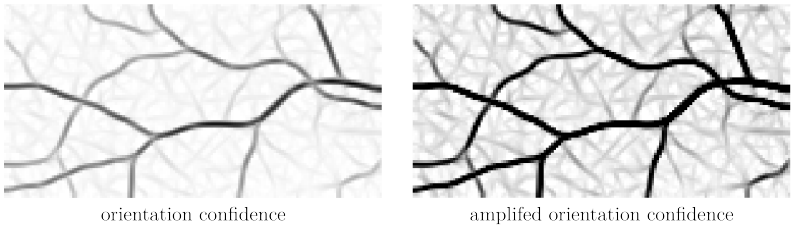

Figure 2.18: Orientation confidences visualized by maximization along the angular dimension.

The amplified orientation confidence had been tested for all three metric tensor fields with parameters $\alpha, \beta \in \{0.8, 0.9\}$. Unfortunately, for only a handful of these the numerical computation converged. The results of these are shown in Figure 2.19. The isotropic left invariant metric already takes the correct path for $\beta = 0.8$, whereas the adaptive gauge frame tensor field needs a higher anisotropicty parameter of $\alpha = 0.9$. The isotropic left invariant tensor field follows the vessel slightly better than the adaptive

gauge frame tensor field, but as before its curve is more jittery.

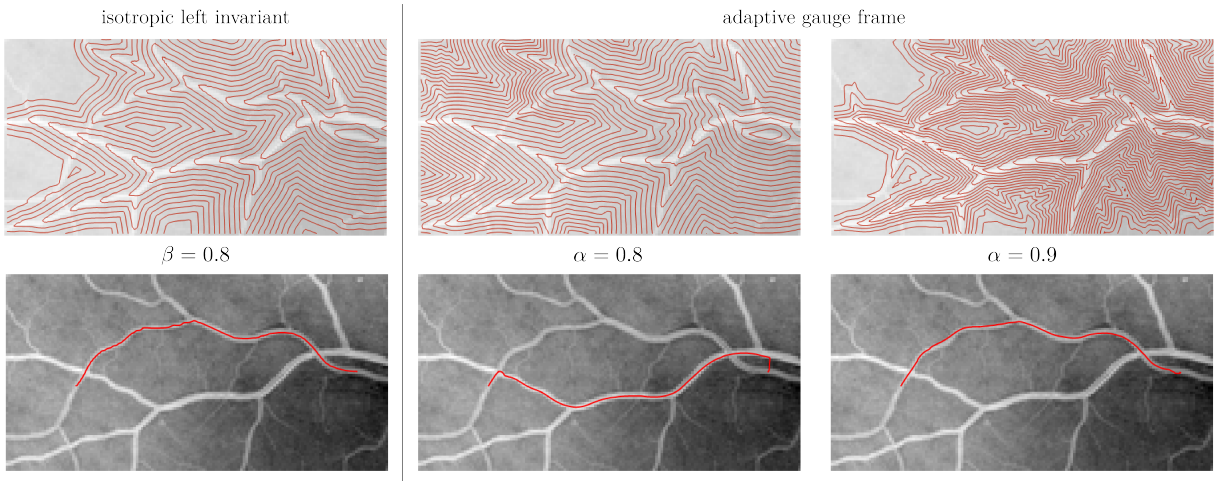

Figure 2.19: Application I: Wavefronts and paths produced by the different metric tensor fields with an amplified orientation confidence.

2.3.2 Application II: High Tortuosity

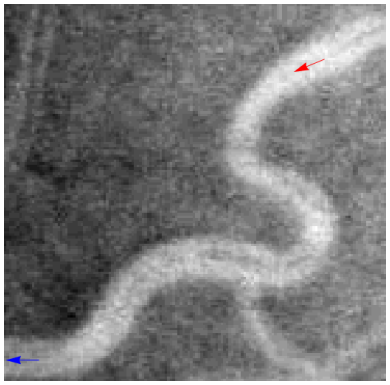

Figure 2.20: The task in the second application is to trace the highly tortuous vessel from the red arrow to the blue arrow. The tracking methods might go straight and bypass the bend in the vessel.

The second application of the metric tensor fields is on an image of a highly tortuous vessel. The task is shown in Figure 2.20. We are interested whether the paths from the vessel tracking methods follow the strongly bended vessel or skip it.

When fitting the gauge frame to the orientation score, we noticed that the orientation confidence showed a line structure that would promote skipping the tortuous vessel. This line segment vanishes when the gauge frame is smoothed by iterated stabilization. Hence, the stabilized gauge frame is used in this application. The orientation confidences of the non-processed and stabilized gauge frames are shown in Figure 2.21

Figure 2.22 shows the tracking results of the left invariant and adaptive gauge frame metric tensor fields. Note that we have used very high parameters α and β in an attempt to prevent the wavefronts from traversing the shortcut before other wavefronts complete the bend. Despite these efforts, none of the methods succeed at this. The wavefronts of the anisotropic left invariant tensor field stand out from the others because they strongly show the effects of the angular resolution used for the orientation score. The paths traced out by both left invariant metric tensor fields take the shortcut across the tortuous vessel. For the adaptive gauge frame metric tensor field, the backtracking algorithm gets stuck in an critical

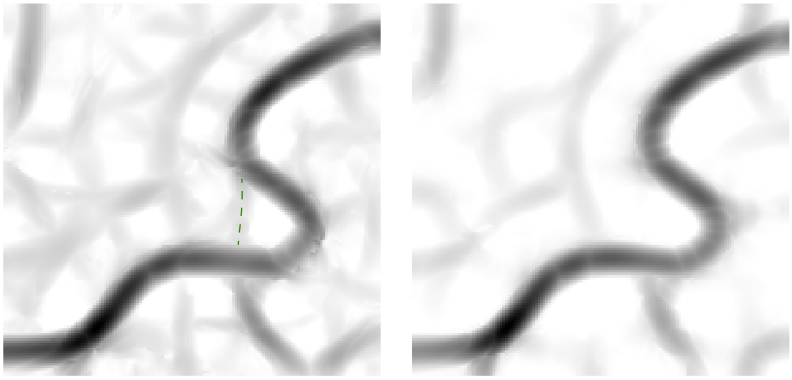

non-stabilized orientation confidence stabilized orientation confidence

Figure 2.21: Orientation confidences visualized by maximization along the angular dimension. The shortcut enabled by the non-stabilized orientation confidence is indicated by the dashed green line.

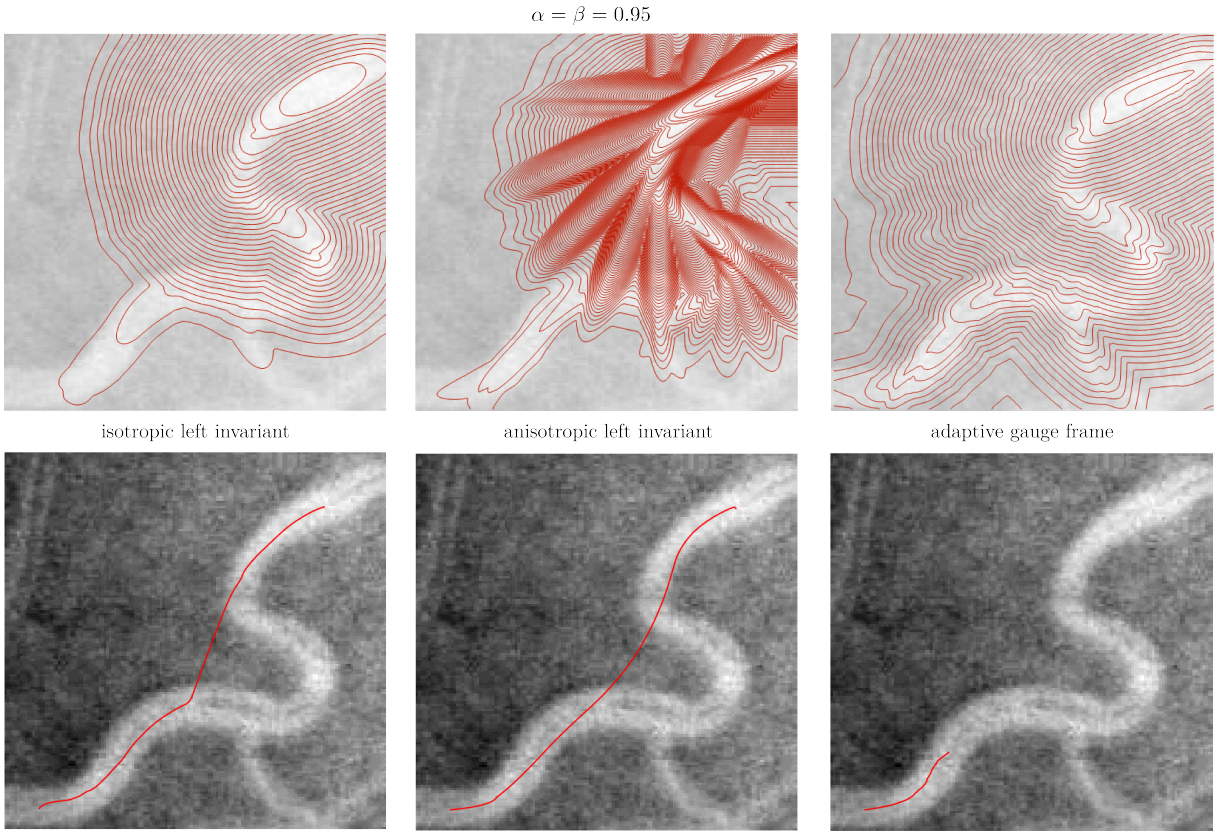

Figure 2.22: Application II: Wavefronts and paths produced by the different metric tensor fields. Parameters $\beta = \alpha = 0.95$.

point like the stabilized gauge frame method in Section 2.2.2. Looking at the wavefronts, we expect that this path would have also taken the shortcut if it had not terminated prematurely.

In a final attempt to force the paths to go around the bend, we cranked up the parameters to $\alpha = \beta = 0.98$. This is at the edge of what the numerical eikonal PDE solver can handle. The results are shown in Figure 2.23. Now all paths follow the vessel instead of taking the shortcut. The path traced out by the adaptive gauge frame metric tensor field does stand out because of its irregularities. It is not exactly clear how these nasty segments follow from the wavefronts above the path, but these fronts do look a bit wilder than those of the other metric tensor fields.

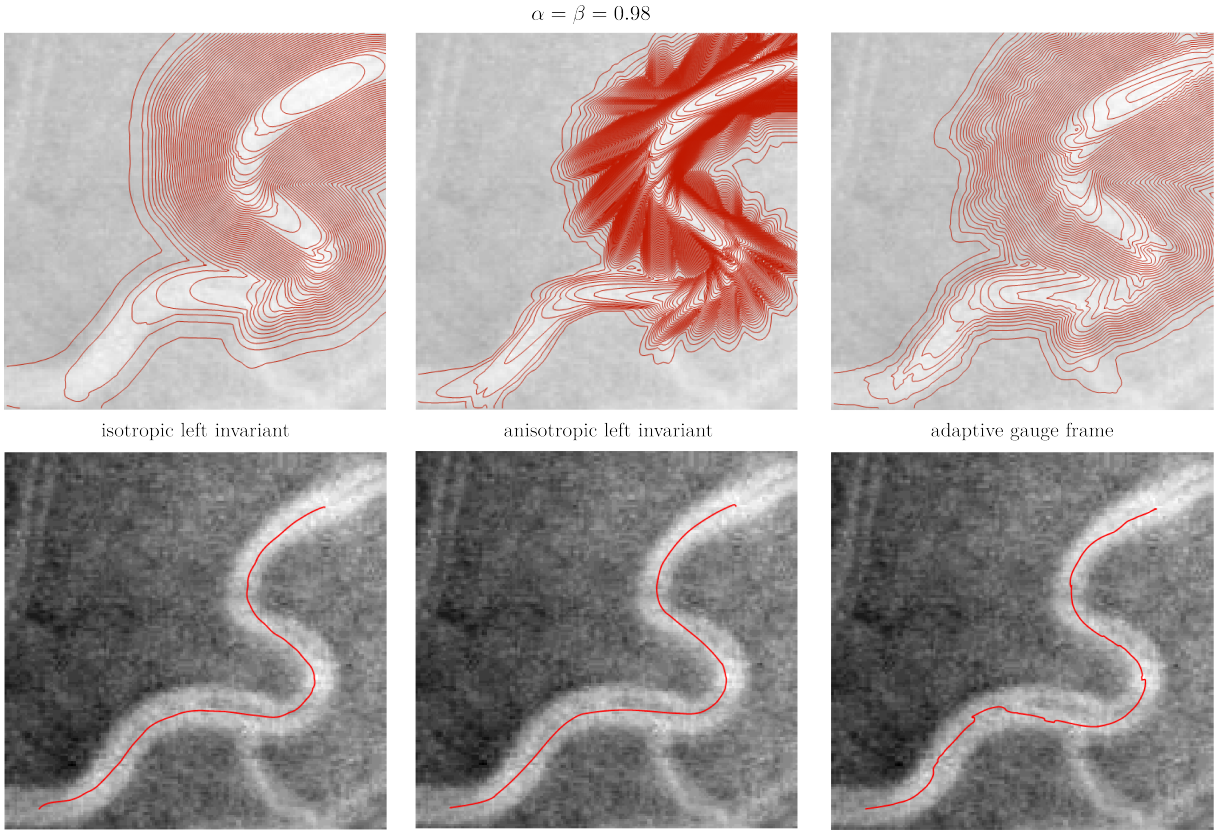

Figure 2.23: Application II: Wavefronts and paths produced by the different metric tensor fields. Parameters $\beta = \alpha = 0.98$.

2.4 Discussion and Conclusion

The experiments in section 2.2.2 and 2.2.3 show that the adaptive gauge frame metric tensor field performs well on an artificial test image, combining positive aspects from both the isotropic and anisotropic left invariant tensor fields. In the second experiment, it even outperformed all other tensor fields. Unfortunately, these results do not translate to the real world applications in Section 2.3. In the first application, the adaptive gauge frame tensor field performed on par with the anisotropic left invariant tensor field, but in the second application, it performed worse than both left invariant tensor fields. Figure 2.22 shows the backtracking algorithm breaking down and the path in Figure 2.23 contains sharp irregularities that are absent from the smooth paths of the left invariant tensor fields.

The drop in performance indicates that the artificial test image is not representative of actual retina vessel images. Its line structures are too clearly defined and they do not vary in intensity. In previous work [4], the cost function could be modified to account for deficiencies in real images. By having the orientation confidence double as cost function in the adaptive gauge frame metric tensor field, we have mistakenly thrown out this flexibility. To promote the desired results in our applications, we have had to apply ad hoc fixes to the orientation confidence.

2.4.1 Future Research

Metric Tensor Field Proposal

The use of gauge frames does have potential for vessel tracking. On artificial test images, the adaptive gauge frame metric tensor field outperforms the left invariant tensor fields. For real images however,

the cost function cannot be replaced by anisotropic gauge frames alone. To draw conclusions about the use of gauge frames on real life data, a future gauge frame based metric tensor field should separate the anisotropic component from the cost function. This offers more flexibility when dealing with real images and it allows for a better comparison with left invariant metric tensors that also separate these components. We do believe that anisotropy w.r.t. the gauge frame should only be applied in areas where the first gauge vector actually aligns with a line structure in the data. The orientation confidence is a natural choice to control this kind of anisotropy.

With above considerations in mind, we would like to propose a specific metric tensor field to be explored in future work. It is based on an orientation confidence controlled interpolation of the left invariant frame and a gauge frame. Recall that a gauge frame is constructed by rotating and scaling the left invariant frame, see Section 1.4.2:

$$\mathcal{B}_i\big|_x := \frac{R_{\mathcal{A}_3,\chi} R_{\mathcal{A}_2,-\nu}\, \mathcal{A}_i\big|_x}{\big|\big| R_{\mathcal{A}_3,\chi} R_{\mathcal{A}_2,-\nu}\, \mathcal{A}_i\big|_x \big|\big|_{\mathcal{G}_\xi}} \ ,$$

where the angles χ and ν express deviation of horizontality and curvature at the point $x \in \mathbb{M}_2$ respectively. An orientation score S takes values in the interval $[0,1]$, so it can be used to control the rotations. Consider a new frame $\{\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3\}$ given by

$$\mathcal{E}_i\big|_x := \frac{R_{\mathcal{A}_3,S(x)\chi} R_{\mathcal{A}_2,-S(x)\nu}\, \mathcal{A}_i\big|_x}{\big|\big| R_{\mathcal{A}_3,S(x)\chi} R_{\mathcal{A}_2,-S(x)\nu}\, \mathcal{A}_i\big|_x \big|\big|_{\mathcal{G}_\xi}} \ .$$

For $S(x) = 0$, this frame coincides with the normalized left invariant frame $\{(1/\xi)\mathcal{A}_1, (1/\xi)\mathcal{A}_2, \mathcal{A}_3\}$. For $S(x) = 1$, the frame coincides with the gauge frame $\{\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3\}$. Let the dual frame be denoted by $\{\psi^1, \psi^2, \psi^3\}$. The metric tensor field we propose is the following:

$$\mathcal{G}^{C,\beta}_{\zeta,S} := (\beta\, C + (1-\beta))^2 \left(\psi^1 \otimes \psi^1 + \frac{1}{\zeta^2}\, \psi^2 \otimes \psi^2 + \left((1-S) + S\frac{1}{\zeta} \right)^2 \psi^3 \otimes \psi^3 \right) \ , \qquad \text{(proposed)}$$

where C is a cost function with values in the range $[\delta, 1]$ for some $\delta > 0$, β is parameter to scale the influence of the cost function and $\zeta \in (0,1]$ denotes an anisotropy parameter. For $S(x) = 0$, this metric coincides with the left invariant metric tensor field

$$\mathcal{G}^{C,\beta}_{\xi,\zeta} = (\beta\, C + (1-\beta))^2 \left(\xi^2 \left(\omega^1 \otimes \omega^1 + \frac{1}{\zeta^2}\, \omega^2 \otimes \omega^2 \right) + \omega^3 \otimes \omega^3 \right)$$

and for $S(x) = 1$, the proposed metric tensor field coincides with an anisotropic gauge frame metric tensor field

$$\mathcal{G}^{C,\beta}_{\zeta} = (\beta\, C + (1-\beta))^2 \left(\chi^1 \otimes \chi^1 + \frac{1}{\zeta^2} \left(\chi^2 \otimes \chi^2 + \chi^3 \otimes \chi^3 \right) \right) \ .$$

Because these metric tensor fields coincide for certain values of the orientation confidence S, they are more easily compared with each other. Hence, we hope that they yield more conclusive results on the effectiveness of gauge frames in vessel tracking.

Vanishing Gradient Field

In the experiments and applications, the backtracking algorithm sometimes got stuck in a critical point that was not the starting point of the tracking process (see the middle column of Figure 2.8 and the right column of Figure 2.22). As mentioned before, this is not supposed to happen since the distance functions are viscosity solutions to the eikonal PDE (2.2). For viscosity solutions, the gradient can only vanish at the initial position or at points where two wavefronts meet. Inspecting the plotted wavefronts, the latter does not seem to be the case here. It seems that the viscosity requirement has been violated.

The fact that this might be such a fundamental issue instead of a numerical fluke did not dawn on us until at the end of the project. Hence, we did not have the time to fully investigate, nor fix the problem. We assume the cause of the issue to be the computation of the gauge frame derivatives in the numerical PDE solver. The upwind differencing scheme is used to calculate the left invariant derivatives, meaning that whether a forward or backward differencing scheme is used depends on the direction of the left invariant

vectors. However, gauge vectors can point in directions completely opposite to the left invariant vectors. At points where this is the case, the wrong choice is made between forward and backward scheme.

We propose the following upwind differencing scheme for gauge frames. Due to time constraints we were unable to implement it. Let $\mathcal{A}_i^+\big|_x(W)$ and $\mathcal{A}_i^-\big|_x(W)$ denote the forward and backward finite difference approximations of $\mathcal{A}_i\big|_x(W)$ respectively. The upwind approximation of $\mathcal{A}_i\big|_x(W)$ as used in previous implementations is given by

$$\mathcal{A}_i^{\text{upwind}}\Big|_x(W) := \max\{\mathcal{A}_i^-\big|_x(W), -\mathcal{A}_i^+\big|_x(W), 0\} \ ,$$

see [4]. Now consider a gauge frame vector $\mathcal{B}_i\big|_x = \sum_{j=1}^{3} c^j \mathcal{A}_i\big|_x$. Let $\mathcal{B}_i^+\big|_x(W)$ and $\mathcal{B}_i^-\big|_x(W)$ denote the respective forward and backward finite difference approximations of $\mathcal{B}_i\big|_x(W)$ which are given by

$$\mathcal{B}_i^+\big|_x(W) = \sum_{j=1}^{3} c^j \mathcal{A}_j^{\text{sign}(c^j)}\Big|_x(W)$$
$$\mathcal{B}_i^-\big|_x(W) = \sum_{j=1}^{3} c^j \mathcal{A}_j^{-\text{sign}(c^j)}\Big|_x(W) \ ,$$

where we interpret the function 'sign' to take values in the set $\{-,+\}$. The upwind approximation of $\mathcal{B}_i\big|_x(W)$ is then similar as that of $\mathcal{A}_i\big|_x(W)$ above:

$$\mathcal{B}_i^{\text{upwind}}\Big|_x(W) := \max\{\mathcal{B}_i^-\big|_x(W), -\mathcal{B}_i^+\big|_x(W), 0\} \ .$$

We hope that this differencing scheme eliminates the presence of unexpected critical points, but this is a topic for future research. As a final note, we demonstrate the importance of the choice of differencing scheme by switching to the central differencing scheme instead the current scheme, which leads to the numerical solution quickly exploding, see Figure 2.24.

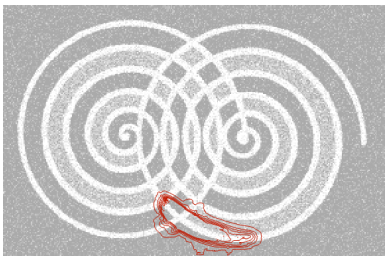

Figure 2.24: Wavefronts of a quickly exploding numerical solution using the central differencing scheme. The metric tensor field used was the same as in the middle column of Figure 2.8 whose backtracking came to a premature halt.

Chapter 3

Data-Driven Connections

In the previous chapter we demonstrated how gauge frames can be used to improve vessel tracking. In this chapter we investigate the gauge frame based tracking methods theoretically. It turns out that by choosing a specific differentiable structure on the tangent space $T\mathbb{M}_2$, i.e. the connection $\widehat{\nabla}$ (3.11), we can characterize the shortest curves as curves with auto-parallel momentum. Moreover, it turns out that w.r.t. this connection, the flowlines of the first gauge vector field are straight curves. In other words, the curves that best follow the line structures in an orientation score are straight.

Such a characterization has been proven before [9], but only for metric tensor fields with constant coefficients w.r.t. to the left invariant dual frame. The straight curves of this connection are the curves with constant speed w.r.t. the left invariant frame (1.20). These curves might align only briefly with the line structures in an orientation score.

The setting for the new characterization is more general than gauge frames on $\mathbb{M}_2$ alone. For the theorem, we consider an arbitrary smooth d-dimensional manifold M with a global frame $\{\mathcal{B}_i\}_{i=1}^d$ and a metric tensor field $\mathcal{G}$ with constant coefficients w.r.t. to the corresponding dual frame $\{\chi^i\}_{i=1}^d$. (We use the same notation for the arbitrary frame as for gauge frames since this was the application we had in mind.) Note that because of the generic nature of the new characterization, it can also be applied to diagonal metric tensor fields with non-constant coefficients. This is because for diagonal metric tensor fields the variability of the coefficients can be considered as variability of the global frame. Using this trick, the characterization can be applied to all metric tensor fields used in the previous chapter.

We start the chapter by investigating the space $T(T^*M)$ the tangent bundle of the cotangent bundle of M, in Section 3.1. This space features prominently in the characterization's proof, and its elements have let to a lot of confusion in the past, so a close inspection is warranted. Section 3.2 serves to explain several concepts from symplectic geometry used by the proof. It is followed by the proposal of a *new connection* in Section 3.3 and an analysis of covariant derivatives along a curve in Section 3.4. Finally we have the *characterization* and its *proof* in Section 3.5.

3.1 The Tangent Bundle of the Cotangent Bundle

A central idea in symplectic geometry is that shortest curves in M can be described using a flow on the cotangent bundle T^*M. Given a frame $\{\mathcal{B}_i\}_{i=1}^d$ of the tangent bundle TM, we can decompose $T(T^*M)$ into a horizontal component diffeomorphic to TM and a vertical component

$$T(T^*M) \cong TM \times \mathbb{R}^d \ .$$

We denote elements of T^*M by

$$\nu = (x, p) \in \coprod_{y \in M} T_y^*M = T^*M \ .$$

Using the projection $\pi : T^*M \to M$, we sometimes write $\pi(\nu)$ instead of x for the first component to emphasize its dependence on ν. To ease notation, we may also choose to identify ν with p.

3.1.1 Horizontal Part

Any point $\nu \in T^*M$ in the cotangent bundle can be expressed as a linear combination of the frame $\{\chi^i\}_{i=1}^d$ dual to $\{\mathcal{B}_i\}_{i=1}^d$:

$$\nu = \sum_{j=1}^{d} \mu_j(\nu)\, \chi^j\big|_{\pi(\nu)} \ .$$

Remark that the coordinate functions $\mu_i : T^*M \to \mathbb{R}$ are given by

$$\mu_i(\nu) = \mu_i(x,p) = \left\langle p\,, \mathcal{B}_i\big|_x \right\rangle \ .$$

The decomposition of ν allows us to define a tangent vector $\overline{\mathcal{B}_i}\big|_\nu \in T_\nu(T^*M)$ in terms the vector $\mathcal{B}_i\big|_{\pi(\nu)}$ in the tangent bundle TM. Let γ be an integral curve to $\mathcal{B}_i\big|_{\pi(\nu)} \in T_{\pi(\nu)}M$, then define

$$\overline{\mathcal{B}_i}\big|_\nu := \frac{d}{dt}\bigg|_{t=0} \sum_{j=1}^{d} \mu_j(\nu)\, \chi^j\big|_{\gamma(t)} \ . \tag{3.1}$$

The horizontal part of $T(T^*M)$ is spanned by the vector fields $\{\overline{\mathcal{B}_i}\}_{i=1}^d$ defined by the method above.

3.1.2 Vertical Part

Together with a coordinate chart $(x^1, \dots, x^d) : U \to \mathbb{R}^d$ for open $U \subset M$, the coordinate functions μ_i form a coordinate chart of $(T^*M)\big|_U$:

$$(x^1, \dots, x^d, \mu_1, \dots, \mu_n) : (T^*M)\big|_U \to \mathbb{R}^d \times \mathbb{R}^d \ .$$

Hence, the coordinate functions μ_i induce vector fields $\partial_{\mu_i} \in \mathfrak{X}(T^*M)$. (We use $\mathfrak{X}(\,\cdot\,)$ to denote the smooth vector fields of a space. The vector fields ∂_{μ_i} span the vertical part of $T(T^*M)$.

The next lemma shows that the horizontal and vertical parts fully span $T(T^*M)$.

Lemma 10. *Consider the differential $\pi_* : T(T^*M) \to TM$ of the projection $\pi : T^*M \to M$. We have*

$$\pi_*\left(\overline{\mathcal{B}_i}\right) = \mathcal{B}_i \qquad \text{and} \qquad \pi_*\left(\partial_{\mu_i}\right) = 0 \ ,$$

hence $\{\overline{\mathcal{B}_1}, \dots, \overline{\mathcal{B}_n}, \partial_{\mu_1}, \dots, \partial_{\mu_n}\}$ is a frame for $T(T^*M)$.

Proof. We find the first two identities by direct computation. Let γ be an integral curve to $\mathcal{B}_i\big|_{\pi(\nu)}$ as in the definition of $\overline{\mathcal{B}_i}\big|_\nu$, see (3.1), then

$$\pi_*\left(\overline{\mathcal{B}_i}\big|_\nu\right) = \frac{d}{dt}\bigg|_{t=0} \pi\left(\sum_{j=1}^{d} \mu_j(\nu)\, \chi^j\big|_{\gamma(t)}\right) = \frac{d}{dt}\bigg|_{t=0} \gamma(t) = \mathcal{B}_i\big|_{\pi(\nu)} \ .$$

Secondly,

$$\pi_*\left(\partial_{\mu_i}\big|_\nu\right) = \frac{d}{dt}\bigg|_{t=0} \pi\left(\sum_{j=1}^{d} \left(\mu_j(\nu) + \delta_{ij}\, t\right)\, \chi^j\big|_{\pi(\nu)}\right) = \frac{d}{dt}\bigg|_{t=0} \pi(\nu) = 0 \ .$$

It remains to show that the set $\{\overline{\mathcal{B}_1}, \dots, \overline{\mathcal{B}_n}, \partial_{\mu_1}, \dots, \partial_{\mu_n}\}$ is linearly independent. Let $\alpha^i, \beta^i \in \mathbb{R}$ be such that

$$\sum_{i=1}^{d} \alpha^i\, \overline{\mathcal{B}_i} + \beta^i\, \partial_{\mu_i} = 0 \ .$$

Applying the differential π_*, the above computations imply that $\sum_{i=1}^{d} \alpha^i \mathcal{B}_i = 0$. Because $\{\mathcal{B}_i\}_{i=1}^d$ is a frame by assumption, it holds that $\alpha^i = 0$. Substitution results in $\sum_{i=1}^{d} \beta^i \partial_{\mu_i} = 0$. The set $\{\partial_{\mu_i}\}_{i=1}^d$ is linearly independent as well since the vector fields arise from a coordinate chart, implying that $\beta^i = 0$. Hence, the original set is linearly independent. ∎

3.1.3 A Canonical Frame for $T(T^*M)$

Given a coordinate chart $(x^1,\dots,x^d) : U \to \mathbb{R}^d$ for an open $U \subset M$, one has a local frame $\{\partial_{x^i}\}_{i=1}^d$ for $(TM)\big|_U$. Remark that the decomposition of $T(T^*M)$ into a horizontal and vertical part can also be done locally, hence we can write covectors as $\nu = \sum_{j=1}^{d} \xi_j(\nu)\, dx^j$ and obtain vector fields $\overline{\partial_{x^i}}$ and ∂_{ξ_i} for the horizontal and vertical parts. However, since $(x^1,\dots,x^d,\xi_1,\dots,\xi_n)$ is a coordinate chart for $T(T^*M)$, we also have the regular vector fields ∂_{x^i}. Fortunately, this is no cause for confusion as $\overline{\partial_{x^i}} = \partial_{x^i}$.

Lemma 11. *Given a coordinate chart $\varphi = (x^1,\dots,x^d) : U \to \mathbb{R}^d$ for an open $U \subset M$, the vector fields $\overline{\partial_{x^i}}$ and $\partial_{x^i} \in \mathfrak{X}(T^*M)$ coincide.*

Proof. The coordinate map $(x^1,\dots,x^n,\xi_1,\dots,\xi_n) : (T^*M)\big|_U \to \mathbb{R}^d \times \mathbb{R}^d$ is given by

$$\nu = \sum_{j=1}^{d} \xi_j(\nu)\, dx^j\big|_{\pi(\nu)} \mapsto (x^1(\pi(\nu)),\dots,x^d(\pi(\nu)),\, \xi_1(\nu),\dots,\xi_n(\nu)) = (\varphi(\pi(\nu)),\, \xi_1(\nu),\dots,\xi_n(\nu)) \ ,$$

so $\partial_{x^i}\big|_\nu$ is defined by a curve $\overline{\gamma}$ given by

$$\overline{\gamma}(t) := \sum_{j=1}^{d} \xi_j(\nu)\, dx^j\big|_{\varphi^{-1}(x^1,\dots,x^i+t,\dots,x^d)} \ .$$

Note that this is the same curve that we have used to define $\overline{\partial_{x^i}}\big|_\nu$, see (3.1), so the vector fields coincide:

$$\overline{\partial_{x^i}} = \partial_{x^i} \in \mathfrak{X}(T^*M) \ .$$

■

3.1.4 Change of Basis

The embedding of vectors $\overline{(\,\cdot\,)} : TM \to T(T^*M)$ does not commute with a change of basis. Consider a change of basis S between two frames $\{\mathcal{B}_i\}_{i=1}^d$ and $\{\mathcal{E}_i\}_{i=1}^d$:

$$\mathcal{E}_i = \sum_{j=1}^{d} S_i^j \mathcal{B}_j \quad \Leftrightarrow \quad \mathcal{B}_i = \sum_{j=1}^{d} \left(S^{-1}\right)_i^j \mathcal{E}_j$$

This implies that for the respective dual frames $\{\chi^i\}_{i=1}^d$ and $\{\psi^i\}_{i=1}^d$

$$\psi^i = \sum_{j=1}^{d} \left(S^{-1}\right)_j^i \chi^j \quad \Leftrightarrow \quad \chi^i = \sum_{j=1}^{d} S_j^i \psi^j \tag{3.2}$$

because

$$\left\langle \psi^i, \mathcal{B}_j \right\rangle = \left\langle \psi^i, \sum_{k=1}^{d} \left(S^{-1}\right)_j^k \mathcal{E}_k \right\rangle = \sum_{k=1}^{d} \left(S^{-1}\right)_j^k \delta_k^i = \left(S^{-1}\right)_j^i \ .$$

For the respective coordinate functions $\{\mu_i\}_{i=1}^d$ and $\{\rho_i\}_{i=1}^d$ we have that

$$\rho_i = \sum_{j=1}^{d} S_i^j \mu_j \quad \Leftrightarrow \quad \mu_i = \sum_{j=1}^{d} \left(S^{-1}\right)_i^j \rho_j$$

since

$$\nu = \sum_{j=1}^{d} \mu_j(\nu)\, \chi^j\big|_{\pi(\nu)} = \sum_{j=1}^{d} \mu_j(\nu) \left(\sum_{i=1}^{d} S_i^j(\pi(\nu))\, \psi^i\big|_{\pi(\nu)} \right) = \sum_{i=1}^{d} \left(\sum_{j=1}^{d} S_i^j(\pi(\nu))\, \mu_j(\nu) \right) \psi^j\big|_{\pi(\nu)} \ .$$

These transformation rules allow us to prove our claim:

Lemma 12. *Given the context above, define* $\widehat{\mathcal{E}_i} := \sum_{j=1}^{d} S_i^j \overline{\mathcal{B}_j}$ *. Then* $\overline{\mathcal{E}_i} \neq \widehat{\mathcal{E}_i}$ *in general, since their effect on the coordinate function* ρ_j *differs:*

$$\overline{\mathcal{E}_i}\,(\rho_j) = 0 \qquad \textit{and} \qquad \widehat{\mathcal{E}_i}\,(\rho_j) = \sum_{k,\ell=1}^{d} S_i^k\, \mathcal{B}_k\left(S_j^\ell\right) \mu_\ell \ .$$

Proof. Both equalities follow from direct computation and the definition of $\overline{(\,\cdot\,)}$ and $\widehat{\mathcal{E}_i}$ respectively. Let γ denote an integral curve for $\mathcal{E}_i\big|_{\pi(\nu)}$, then

$$\overline{\mathcal{E}_i}\big|_\nu(\rho_j) = \frac{d}{dt}\bigg|_{t=0} \rho_j\left(\sum_{k=1}^{d} \rho_k(\nu)\psi^k\big|_{\gamma(t)}\right) = \frac{d}{dt}\bigg|_{t=0} \rho_j(\nu) = 0 \ .$$

Next,

$$\widehat{\mathcal{E}_i}\,(\rho_j) = \sum_{k=1}^{d} S_i^k\, \overline{\mathcal{B}_k}\left(\sum_{\ell=1}^{d} S_j^\ell \mu_\ell\right) = \sum_{k,\ell=1}^{d} S_i^k \left(\overline{\mathcal{B}_k}\left(S_j^\ell\right)\mu_j + S_j^\ell\, \overline{\mathcal{B}_k}(\mu_j)\right) = \sum_{k,\ell=1}^{d} S_i^k\, \mathcal{B}_k\left(S_j^\ell\right)\mu_j$$

because analogous to the computation above $\overline{\mathcal{B}_k}(\mu_j) = 0$ and because S_j^ℓ is constant on fibers T_x^*M, we can write $\overline{\mathcal{B}_k}\left(S_j^\ell\right) = \mathcal{B}_k\left(S_j^\ell\right)$. ∎

3.2 Some (Symplectic) Geometry Concepts

Here we present some (symplectic) geometry tools used in the proof of the characterization in Section 3.5. A good reference on symplectic geometry and its use in geometric control theory is the book [1].

3.2.1 Symplectic Form

As its name suggests, the symplectic form plays a major role in symplectic geometry. In order to define it, we first look at a different form.

The **Liouville 1-form** $s \in \Omega^1(T^*M)$ is given by $s\big|_{(x,p)} := p \circ \pi_* = \langle p\,, \pi_*(\cdot)\rangle$ where π_* is the derivative of the projection $\pi : TM \to M$. It takes a vector in $T_{(x,p)}(T^*M)$, projects it onto T_xM and then applies the covector p. Using the frame $\{\mathcal{B}_i\}_{i=1}^d$, we can write a vector $w \in \mathfrak{X}(T^*M)$ as

$$w = \sum_{i=1}^{d} \alpha^i\, \overline{\mathcal{B}_i} + \beta^i\, \partial_{\mu_i} \ .$$

Writing $p = \sum_{j=1}^{d} \mu_j(x,p)\, \chi^j\big|_x$, this gives us a coordinate expression for s:

$$s\big|_\nu = s\big|_{(x,p)} = \langle p, \pi_*(w)\rangle = \left\langle p\,, \sum_{i=1}^{d} \alpha^i \mathcal{B}_i \right\rangle = \sum_{j=1}^{d} \mu_j(x,p)\, \alpha^j = \sum_{j=1}^{d} \mu_j(\nu)\, \alpha^j \ ,$$

i.e. $s = \sum_{j=1}^{d} \mu_j \overline{\chi^j}$ where $\overline{\chi^i}$ is the covector dual to $\overline{\mathcal{B}_i}$. In canonical coordinates (Section 3.1.3) we find that

$$s = \sum_{j=1}^{d} \xi_j\, dx^j \ ,$$

because $\overline{\partial_{x^i}} = \partial_{x^i}$ (Lemma 11). Note that when we identify $\nu = (x,p)$ with p, the expression for $s(\nu)$ is *syntactically* equal to ν. Hence, s is also referred to as the **tautological 1-form**.

The **symplectic form** is $\sigma := -ds \in \Omega^2(T^*M)$, a 2-form on T^*M. Applying the exterior derivative d to the coordinate expression of s, we can express σ w.r.t. the frame $\{\mathcal{B}_i\}_{i=1}^d$ by

$$\sigma = -\sum_{j=1}^{d} d\mu_j \wedge \overline{\chi^j} + \mu_j \, d\overline{\chi^j} \ .$$

The 2-forms $d\overline{\chi^j}$ cause a lot of headaches when doing computations. Luckily they disappear in canonical coordinates as $dd = 0$, leaving

$$\sigma = -\sum_{j=1}^{d} d\xi_j \wedge dx^j = \sum_{j=1}^{d} dx^j \wedge d\xi_j \ . \tag{3.3}$$

3.2.2 Lifting to Vector Fields on the Cotangent Bundle

The symplectic form σ can be used to 'lift' functions $f \in \mathcal{C}^\infty(T^*M)$ to vector fields $\overrightarrow{f} \in \mathfrak{X}(T^*M)$ by requiring that

$$\sigma\left(\overrightarrow{f}, \ \cdot \ \right) = df \ . \tag{3.4}$$

Assume we can write $\overrightarrow{f} = \sum_{j=1}^{d} \alpha^j \partial_{x^j} + \beta^j \partial_{\xi_j}$ in canonical coordinates, then by coordinate formula (3.3)

$$\frac{\partial f}{\partial x^i} = df\left(\partial_{x^i}\right) = \sigma\left(\overrightarrow{f}, \partial_{x^i}\right) = -\beta^i \quad \text{and} \quad \frac{\partial f}{\partial \xi_i} = df\left(\partial_{\xi_i}\right) = \sigma\left(\overrightarrow{f}, \partial_{\xi_i}\right) = \alpha^i \ ,$$

hence

$$\overrightarrow{f} = \sum_{j=1}^{d} \frac{\partial f}{\partial \xi_j} \partial_{x^j} - \frac{\partial f}{\partial x^i} \partial_{\xi_j} \ . \tag{3.5}$$

Performing a change of basis does not alter the structure of this formula as the next lemma shows.

Lemma 13. *Assume S to be a change of basis s.t. $\mathcal{B}_i = \sum_{j=1}^{d} S_i^j \partial_{x^j} \in \mathfrak{X}(M)$ and use it to define vector fields $\widehat{\mathcal{B}_i} := \sum_{j=1}^{d} S_i^j \partial_{x^j} \in \mathfrak{X}(T^*M)$. Then for $f \in \mathcal{C}^\infty(T^*M)$ we have that*

$$\frac{\partial f}{\partial \xi_i} = \sum_{j=1}^{d} S_j^i \frac{\partial f}{\partial \mu_j}$$

where μ_i are the coordinate functions w.r.t. to the basis $\{\chi^i\}_{i=1}^d$ dual to $\{\mathcal{B}_i\}_{i=1}^d$. Consequentially

$$\overrightarrow{f} = \sum_{j=1}^{d} \frac{\partial f}{\partial \mu_j} \widehat{\mathcal{B}_j} - \widehat{\mathcal{B}_j}\left(f\right) \partial_{\mu_j} \ , \tag{3.6}$$

Proof. Using that $dx^i = \sum_{j=1}^{d} S_j^i \, \chi^j$ (see equation (3.2) in Section 3.1.4), we find that for $\nu = (x, p) \in T^*M$

$$\begin{aligned}
\frac{\partial f}{\partial \xi_i}(x,p) &= \frac{d}{dt}\Big|_{t=0} f\left(p + t\, dx^i\big|_x\right) = \frac{d}{dt}\Big|_{t=0} f\left(p + t \sum_{j=1}^{d} S_j^i(x)\, \chi^j\big|_x\right) \\
&= \frac{d}{dt}\Big|_{t=0} f\left(\sum_{j=1}^{d} \left(\mu_j(x,p) + t\, S_j^i(x)\right) \chi^j\big|_x\right) = \sum_{j=1}^{d} S_j^i(x) \frac{\partial f}{\partial \mu_j}(x,p) \ ,
\end{aligned}$$

where the last step follows from the chain rule. Substituting the above and the definition of $\widehat{\mathcal{B}_i}$ into the coordinate expression for $\overrightarrow{f}$ in canonical coordinates, finishes our proof.

$$\begin{aligned}\overrightarrow{f} &= \sum_{j=1}^{d} \frac{\partial f}{\partial \xi_j}\partial_{x^j} - \frac{\partial f}{\partial x^i}\partial_{\xi_j} = \sum_{j=1}^{d}\left(\left(\sum_{k=1}^{d} S_k^j \frac{\partial f}{\partial \mu_k}\right)\left(\sum_{\ell=1}^{d}\left(S^{-1}\right)_j^{\ell}\widehat{\mathcal{B}_\ell}\right) - \left(\sum_{k=1}^{d}\left(S^{-1}\right)_j^{k}\widehat{\mathcal{B}_k}(f)\right)\left(\sum_{\ell=1}^{d} S_\ell^j \partial_{\mu_\ell}\right)\right) \\ &= \sum_{k,\ell=1}^{d} \delta_k^\ell \frac{\partial f}{\partial \mu_k}\widehat{\mathcal{B}_\ell} - \delta_\ell^k \widehat{\mathcal{B}_k}(f)\partial_{\mu_\ell} = \sum_{k=1}^{d}\frac{\partial f}{\partial \mu_k}\widehat{\mathcal{B}_k} - \widehat{\mathcal{B}_k}(f)\partial_{\mu_k} \ .\end{aligned}$$

■

3.2.3 Poisson Bracket

The Poisson bracket is an important tool in symplectic geometry as it helps to find new functions $\mathcal{C}^\infty(T^*M)$ that are preserved along the flowlines of lifted vector fields $\overrightarrow{\mathfrak{h}} \in T(T^*M)$ *(constants of motion).* It is defined on functions $f, g \in \mathcal{C}^\infty(T^*M)$ by

$$\{f, g\} := \sigma\left(\overrightarrow{g}, \overrightarrow{f}\right) = dg\left(\overrightarrow{f}\right) \ . \tag{3.7}$$

Using the formulas from the previous section, we can express it as

$$\{f, g\} = \sum_{j=1}^{d} \frac{\partial f}{\partial \xi_j}\frac{\partial g}{\partial x^j} - \frac{\partial f}{\partial x^j}\frac{\partial g}{\partial \xi_j} \tag{3.8}$$

in canonical coordinates or as

$$\{f, g\} = \sum_{j=1}^{d} \frac{\partial f}{\partial \mu_j}\widehat{\mathcal{B}_j}(g) - \widehat{\mathcal{B}_j}(f)\frac{\partial g}{\partial \mu_j} \tag{3.9}$$

w.r.t. the change of basis $\mathcal{B}_i = \sum_{j=1}^{d} S_i^j \partial_{x^j} \in \mathfrak{X}(M)$. *Note that we use vector fields $\widehat{\mathcal{B}_i}$ and* ***not*** *$\overline{\mathcal{B}_i}$ in (3.9). They are not the same, as shown in Lemma 12.*

Similarly to finding new constants of motion, our main result uses the Poisson bracket to compute the canoncical ODE's of momenta $\nu(t) \in T^*M$ along the flowlines of a lifted vector field $\overrightarrow{\mathfrak{h}} \in T(T^*M)$. *(More specifically, along the flowlines of the lifted hamiltonian according to the Pontryagin Maximum Principle.)* The following three lemmas outlining properties of the Poisson bracket will be used in the calculations.

Lemma 14. *The Poisson bracket is bilinear, i.e. for functions $f, g \in \mathcal{C}^\infty(T^*M)$ and constants $\alpha, \beta \in \mathbb{R}$ we have that*

$$\{\alpha f + \beta f, g\} = \alpha\{f, g\} + \beta\{f, g\} = \{f, \alpha g + \beta g\} \tag{bilinearity}$$

Secondly, for functions $f, g, h \in \mathcal{C}^\infty(T^*M)$, the Poisson bracket satisfies a Leibniz rule:

$$\{fg, h\} = \{f, h\}\, g + f\,\{g, h\} \ . \tag{Leibniz rule}$$

Proof. The first statement follows from linearity of the exterior derivative and the ability to write

$$\{f, g\} = \sigma\left(\overrightarrow{g}, \overrightarrow{f}\right) = dg(\overrightarrow{f}) \ .$$

Similarly, the second statement follows from the product rule for the exterior derivative:

$$\{fg, h\} = \sigma\left(\overrightarrow{h}, \overrightarrow{fg}\right) = -\sigma\left(\overrightarrow{fg}, \overrightarrow{h}\right) = -d(fg)(\overrightarrow{h}) = -df(\overrightarrow{h})\, g - f\, dg(\overrightarrow{h}) = \{f, h\}\, g + f\,\{g, h\} \ .$$

■

Lemma 15 (Relation between Poisson and Lie bracket)**.** *Consider the $\mathcal{C}^\infty(M)$-linear mapping $a_{(.)} : \mathfrak{X}(M) \to \mathcal{C}^\infty(T^*M)$ given by $a_X := \langle \cdot\, , X \rangle$, then for vector fields $X, Y \in \mathfrak{X}(M)$*

$$\{a_X, a_Y\} = a_{[X,Y]} \ .$$

Proof. We write X and Y w.r.t. a canonical coordinate system as $\sum_{j=1}^{d} X^j \partial_{x^j}$ and $\sum_{j=1}^{d} Y^j \partial_{x^j}$. In these coordinates, $a_X = \sum_{j=1}^{d} \xi_j \, X^j$ where ξ_i are the corresponding coordinate function s.t. $\nu = \sum_{j=1}^{d} \xi_j(\nu) \, dx^j\big|_{\pi(\nu)}$, hence

$$\frac{\partial a_X}{\partial \xi_i} = X^i \qquad \text{and} \qquad \frac{\partial a_X}{\partial x^i} = \sum_{j=1}^{d} \frac{\partial \xi_j}{\partial x^i} X^j + \xi_j \frac{\partial X^j}{\partial x^i} = \sum_{j=1}^{d} \xi_j \frac{\partial X^j}{\partial x^i} \ .$$

Note that the term $\partial_{x^i}(\xi_j)$ vanishes because

$$\partial_{x^i}(\xi_j) \overset{\text{Lemma 11}}{=\!=} \overline{\partial_{x^i}}(\xi_j) \overset{\text{Lemma 12}}{=\!=} 0 \ .$$

We use the derivatives to explicitly compute $\{a_X, a_Y\}$:

$$\begin{aligned} \{a_X, a_Y\} &= \sum_{j=1}^{d} \frac{\partial a_X}{\partial \xi_j} \frac{\partial a_Y}{\partial x^j} - \frac{\partial a_X}{\partial x^j} \frac{\partial a_Y}{\partial \xi_j} = \sum_{j=1}^{d} X^j \left(\sum_{k=1}^{d} \xi_k \frac{\partial Y^k}{\partial x^j} \right) - \left(\sum_{k=1}^{d} \xi_k \frac{\partial X^k}{\partial x^j} \right) Y^j \\ &= \sum_{k=1}^{d} \xi_k \left(\sum_{j=1}^{d} X^j \frac{\partial Y^k}{\partial x^j} - Y^j \frac{\partial X^k}{\partial x^j} \right) = \sum_{k=1}^{d} \xi_k \left(X\left(Y^k\right) - Y\left(X^k\right) \right) = a_{[X,Y]} \ , \end{aligned}$$

where the final equality holds because the coordinate expression for the Lie bracket in canonical coordinates is given by

$$[X, Y] = \sum_{k=1}^{d} \left(X\left(Y^k\right) - Y\left(X^k\right) \right) \partial_{x^k} \ .$$

■

Lemma 16 (Liouville's theorem)**.** *Let $f, g \in \mathcal{C}^\infty(T^*M)$ and let $\nu(t) = (\gamma(t), p(t)) \in T^*_{\gamma(t)}M$ denote a flowline of the lifted vector field $\overrightarrow{g}$, i.e. $\dot{\nu}(t) = \overrightarrow{g}(\nu(t))$, then*

$$\frac{d(f \circ \nu)}{dt}(t) = \{g, f\}\left(\nu(t)\right) \ .$$

Proof. By rewriting the left hand side and definition of the Poisson bracket (3.7) :

$$\frac{d(f \circ \nu)}{dt}(t) = df\big|_{\nu(t)}(\dot{\nu}(t)) = df\big|_{\nu(t)}\left(\overrightarrow{g}(\nu(t))\right) = df\left(\overrightarrow{g}\right)\big|_{\nu(t)} = \{g, f\}\left(\nu(t)\right) \ .$$

■

3.2.4 Reciprocal (Co)Vectors

We recap a bit metric tensor manipulation. Assume we have a metric tensor field $\mathcal{G} = \sum_{i,j=1}^{d} g_{ij} \, \chi^i \otimes \chi^j$. Here, we do not yet require the coefficients g_{ij} to be constant. By Riesz' representation theorem, we have bijections on fibers $\widetilde{\mathcal{G}}\Big|_x : T_xM \to T^*_xM$ given by

$$\widetilde{\mathcal{G}}\Big|_x(X) := \mathcal{G}\big|_x(X, \ \cdot \ ) \ ,$$

for $X \in T_xM$. Since $\widetilde{\mathcal{G}}\Big|_x$ is a bijection, we can think of the coefficients g_{ij} as entries of a symmetric invertible matrix field. The entries of the inverse matrix field are denoted by g^{ij} . $\widetilde{\mathcal{G}}\Big|_x(X)$ is called

the covector **reciprocal** to X and analogously, $\widetilde{\mathcal{G}}^{-1}\Big|_x(\lambda)$ is called the vector **reciprocal** to $\lambda \in T_x^*M$. Writing $X = \sum_{j=1}^{d} x^j \mathcal{B}_j\big|_x$ and $\lambda = \sum_{j=1}^{d} \mu_j \chi^j\big|_x$, we denote the **reciprocal coefficients** of $\widetilde{\mathcal{G}}\Big|_x(X)$ and $\widetilde{\mathcal{G}}^{-1}\Big|_x(\lambda)$ by x_j and μ^j respectively, i.e.

$$x_j = \sum_{i=1}^{d} g_{ij} x^i \qquad \text{and} \qquad \mu^j = \sum_{i=1}^{d} g^{ij} \mu_i \ . \tag{3.10}$$

3.3 Proposed Connection induced by a Global Frame

Definition 14. Given a global frame $\{\mathcal{B}_i\}_{i=1}^d$ with dual frame $\{\chi^i\}_{i=1}^d$, we propose the connection $\widehat{\nabla}$ on TM given by

$$\widehat{\nabla}_X Y := \sum_{k=1}^{d} X\left(y^k\right) \mathcal{B}_k + \sum_{i,j=1}^{d} x^i y^j \left[\mathcal{B}_i, \mathcal{B}_j\right] \ , \tag{3.11}$$

for vector fields $X = \sum_{j=1}^{d} x^j \mathcal{B}_j$ and $Y = \sum_{j=1}^{d} y^j \mathcal{B}_j$.

Lemma 17. *$\widehat{\nabla}$ is a connection.*

Proof. Clearly $\widehat{\nabla}$ is $\mathcal{C}^\infty(M)$-linear in X. Remains to check the Leibniz rule. Let $f \in \mathcal{C}^\infty(M)$, then

$$\begin{aligned}
\widehat{\nabla}_X(f\,Y) &= \sum_{k=1}^{d} X\left(f\,y^k\right) \mathcal{B}_k + \sum_{i,j=1}^{d} x^i (f\,y^j) \left[\mathcal{B}_i, \mathcal{B}_j\right] \\
&= \sum_{k=1}^{d} \left(X(f)\,y^k + f\,X\left(y^k\right)\right) \mathcal{B}_k + f \left(\sum_{i,j=1}^{d} x^i y^j \left[\mathcal{B}_i, \mathcal{B}_j\right]\right) \\
&= \sum_{k=1}^{d} X(f)\,y^k \mathcal{B}_k + f \left(\sum_{k=1}^{d} X\left(y^k\right) \mathcal{B}_k + \sum_{i,j=1}^{d} x^i y^j \left[\mathcal{B}_i, \mathcal{B}_j\right]\right) \\
&= X(f)\,Y + f\left(\widehat{\nabla}_X Y\right) \ .
\end{aligned}$$

■

For our theorem, we need an expression of the dual connection $\widehat{\nabla}^*$ on T^*M. Dual connections are defined in general by enforcing the product rule

$$X\left(\langle \lambda, Y \rangle\right) = \langle \nabla_X^* \lambda, Y \rangle + \langle \lambda, \nabla_X Y \rangle$$

for an arbitrary connection ∇, vector fields $X, Y \in \mathfrak{X}(M)$ and covector field $\lambda \in \Omega^1(M)$. The computation is captured in the following lemma.

Lemma 18. *Any connection ∇ with local expression in terms of Christoffel symbols Γ_{ij}^k*

$$\nabla_X Y = \sum_{k=1}^{d} \left(X\left(y^k\right) + \sum_{i,j=1}^{d} x^i y^j \Gamma_{ij}^k\right) \mathcal{B}_k$$

has their dual ∇^* expressed by

$$\nabla_X^* \lambda = \sum_{i=1}^{d} \left(X(\lambda_i) - \sum_{j,k=1}^{d} x^j \lambda_k \Gamma_{ji}^k\right) \chi^i \ ,$$

where $X = \sum_{j=1}^{d} x^j \mathcal{B}_j, Y = \sum_{j=1}^{d} y^j \mathcal{B}_j$ *and* $\lambda = \sum_{j=1}^{d} \lambda_j \chi^j$.

Proof. By direct computation.

$$\begin{aligned}\langle \nabla_X^* \lambda, Y\rangle &= X\left(\langle \lambda, Y\rangle\right) - \langle \lambda, \nabla_X Y\rangle = X\left(\sum_{k=1}^{d} \lambda_k\, y^k\right) - \sum_{k=1}^{d} \lambda_k \left(X(y^k) + \sum_{i,j=1}^{d} x^i y^j \Gamma_{ij}^k\right) \\ &= \sum_{k=1}^{d}\left(X(\lambda_k)\, y^k - \lambda_k \sum_{i,j=1}^{d} x^i y^j \Gamma_{ij}^k\right) = \sum_{i=1}^{d} X(\lambda_i)\, y^i - \sum_{i,j,k=1}^{d} x^j \lambda_k \Gamma_{ji}^k y^i \\ &= \left\langle \sum_{i=1}^{d}\left(X(\lambda_i) - \sum_{j,k=1}^{d} x^j \lambda_k \Gamma_{ji}^k\right)\chi^i, Y\right\rangle \ .\end{aligned}$$

■

Corollary 18.1. *Introduce the functions* $C_{ij}^k \in \mathcal{C}^\infty(M)$ *defined by* $\sum_{k=1}^{d} C_{ij}^k \mathcal{B}_k = [\mathcal{B}_i, \mathcal{B}_j]$ *, then we have that the dual* $\widehat{\nabla}^*$ *of the proposed connection has the coordinate expression*

$$\widehat{\nabla}_X^* \lambda = \sum_{i=1}^{d}\left(X(\lambda_i) + \sum_{j,k=1}^{d} x^j \lambda_k C_{ij}^k\right)\chi^i \ , \tag{3.12}$$

for $X = \sum_{j=1}^{d} x^j \mathcal{B}_j$ *and* $\lambda = \sum_{j=1}^{d} \lambda_j \chi^j$.

Proof. We have the Christoffel symbols $\Gamma_{ij}^k = C_{ij}^k$ as $\widehat{\nabla}_{\mathcal{B}_i} \mathcal{B}_j = [\mathcal{B}_i, \mathcal{B}_j]$. Because of the antisymmetry in the Lie bracket, we also have antisymmetry in the lower indices of C_{ij}^k . ■

3.4 Covariant Derivative along a Curve

Our goal is to proof a statement about shortest and straight curves, hence we will use concepts like the acceleration $\widehat{\nabla}_{\dot\gamma}\dot\gamma$ and change in momentum $\widehat{\nabla}_{\dot\gamma}^* p$ of a curve $\gamma(t) \in M$. Intuitively, the interpretation of these terms is quite straightforward. $\dot\gamma(t)$ is a vector that varies along γ as t increases. Hence, we can take its covariant derivative along γ , i.e. in the direction $\dot\gamma(t)$ itself. The formalisation however is quite involved, whilst we only need coordinate formulas for $\widehat{\nabla}_{\dot\gamma}\dot\gamma$ and $\widehat{\nabla}_{\dot\gamma}^* p$. Hence, the details of the next lemma are pushed to Appendix A.

Lemma 19. *For any vector bundle* $E \to M$ *with a local frame* $\{\mathcal{E}_i\}_{i=1}^d$ *, the covariant derivative of a section* $s(t) = \sum_{j=1}^{d} s^j(t)\, \mathcal{E}_j\big|_{\gamma(t)}$ *along a curve* $\gamma(t) \in M$ *w.r.t. an arbitrary connection* ∇ *is given by*

$$\left(\nabla_{\dot\gamma}\, s\right)(t) = \sum_{j=1}^{d}\left(\frac{ds^j}{dt}(t)\, \mathcal{E}_j\big|_{\gamma(t)} + s^j(t)\left(\nabla_{\dot\gamma(t)}\, \mathcal{E}_j\big|_{\gamma(t)}\right)\right) \ . \tag{3.13}$$

Corollary 19.1. *Suppressing the temporal and spatial parameters* $t \in [0,1]$ *and* $\gamma(t) \in M$ *we have the coordinate formulas*

$$\widehat{\nabla}_{\dot\gamma}\dot\gamma = \sum_{k=1}^{d}\left(\ddot\gamma^k + \sum_{i,j=1}^{d} \dot\gamma^i \dot\gamma^j C_{ij}^k\right)\mathcal{B}_k \tag{3.14}$$

$$\widehat{\nabla}_{\dot\gamma}^* p = \sum_{i=1}^{d}\left(\dot p_i + \sum_{j,k=1}^{d} \dot\gamma^j p_k C_{ij}^k\right)\chi^i \ , \tag{3.15}$$

with the functions $C_{ij}^k \in \mathcal{C}^\infty(M)$ *defined by* $\sum_{k=1}^d C_{ij}^k \mathcal{B}_k = [\mathcal{B}_i, \mathcal{B}_j]$ *as in Corollary 18.1.*

Proof. We have the decompositions

$$\dot\gamma(t) = \sum_{j=1}^d \dot\gamma^j(t)\, \mathcal{B}_j\big|_{\gamma(t)} \qquad \text{and} \qquad p(t) = \sum_{j=1}^d p_j(t)\, \chi^j\big|_{\gamma(t)} \ .$$

Next, by the definition of $\widehat\nabla$ (3.11) and the formula for $\widehat\nabla^*$ (3.12) we have that

$$\widehat\nabla_{\dot\gamma(t)}\, \mathcal{B}_j\big|_{\gamma(t)} = \sum_{i=1}^d \dot\gamma^i(t)\ [\mathcal{B}_i, \mathcal{B}_j] = \sum_{i,k=1}^d \dot\gamma^i(t)\, C_{ij}^k\, \mathcal{B}_k$$

$$\widehat\nabla^*_{\dot\gamma(t)}\, \chi^k\big|_{\gamma(t)} = \sum_{i,j=1}^d \dot\gamma^j(t)\, C_{ij}^k\, \chi^i \ .$$

The result follows by substitution of the above in equation (3.13). ∎

3.5 Characterizing Shortest and Straight Curves

Let M be a manifold with a global frame $\{\mathcal{B}_i\}_{i=1}^d$ and denote its dual frame by $\{\chi^i\}_{i=1}^d$. Let $\mathcal{G}$ be a metric tensor field that is constant w.r.t. this frame, i.e. $\mathcal{G} = \sum_{i,j=1}^d g_{ij}\, \chi^i \otimes \chi^j$ with constant coefficients $g_{ij} \in \mathbb{R}$. The metric tensor field $\mathcal{G}$ induces a metric $d_{\mathcal{G}}$ on M given by

$$d_{\mathcal{G}}(x,y) := \min_{\substack{\gamma \in \mathrm{Lip}([0,1],M) \\ \gamma(0)=x \\ \gamma(1)=y}} \int_0^1 ||\dot\gamma(t)||_{\mathcal{G}}^2\, dt \ .$$

Theorem 20. *If* γ *is a shortest curve, i.e. it locally minimizes* $d_{\mathcal{G}}$*, then* γ *is the horizontal part of a curve* $\nu = (\gamma, p)$ *in* T^*M *which has momentum* $p = \widetilde{\mathcal{G}}(\dot\gamma)$ *that is parallel w.r.t. the dual connection* $\widehat\nabla^*$*:*

$$\widehat\nabla^*_{\dot\gamma}\, p = 0 \ . \qquad \text{(part I)}$$

Furthermore, γ *is a straight curve w.r.t. the connection* $\widehat\nabla$ *if and only if its velocity* $\dot\gamma$ *is a constant linear combination of the basis* $\{\mathcal{B}_i\}_{i=1}^d$*:*

$$\widehat\nabla_{\dot\gamma}\dot\gamma = 0 \quad \textit{iff} \quad \dot\gamma = \sum_{i=1}^d u^i \mathcal{B}_i \ \textit{where}\ u^i \in \mathbb{R} \ . \qquad \text{(part II)}$$

Remark. *This theorem can also be applied to diagonal metric tensor fields with variable coefficients. This is done by pushing the variability of the coefficients* g_{ii} *onto the frame* $\{\mathcal{B}_i\}_{i=1}^d$ *by creating a new frame*

$$\mathcal{E}_i := \sqrt{g_{ii}}\, \mathcal{B}_i \ .$$

The metric tensor field is now constant w.r.t. this new frame.

Proof. We start with the proof of part I. The metric $d_{\mathcal{G}}$ is defined in terms of the Lagrangian $||\cdot||_{\mathcal{G}}^2$ induced by the metric tensor field $\mathcal{G}$. The Lagrangian defines a hamiltonian function $\mathfrak{h} \in \mathcal{C}^\infty(T^*M)$ which can be lifted to a vector field $\overrightarrow{\mathfrak{h}} \in \mathfrak{X}(T^*M)$. The Pontryagin Maximum Principle states that if a curve $\gamma(t)$ locally minimizes $d_{\mathcal{G}}$, its momentum $\nu(t) = (\gamma(t), p(t))$ is a flowline of $\overrightarrow{\mathfrak{h}}$, see Theorem 12.10 in [1]. We start by unpacking the statement above and decoupling the flowline ODE.

The flowline ODE

The hamiltonian $\mathfrak{h}$ is defined as the Legendre-Fenchel transform of the scaled Lagrangian $\frac{1}{2}||\cdot||^2_{\mathcal{G}}$, i.e.

$$\mathfrak{h}(\nu) = \mathfrak{h}(x,p) := \sup_{V \in T_x M} \langle p, V\rangle - \frac{1}{2}\mathcal{G}(V,V) \ . \tag{3.16}$$

For a fixed $\nu = (x,p) \in T^*M$, we denote the function inside the supremum by $f_p(V) := \langle p, V\rangle - \frac{1}{2}\mathcal{G}(V,V)$. By convexity of $\mathcal{G}$ and linearity of $\langle p, \cdot\,\rangle$, f_p is concave. In combination with continuity of $\mathcal{G}$, we find that the supremum is obtained when the differential df_p is zero. The differential df_p is found by expansion of $df_p(V+h)$:

$$f_p(V+h) = \langle p, V+h\rangle - \frac{1}{2}\mathcal{G}(V+h,V+h) = \langle p,V\rangle - \frac{1}{2}\mathcal{G}(V,V) + \langle p,h\rangle - \mathcal{G}(V,h) - \frac{1}{2}\mathcal{G}(h,h) \ .$$

Hence the differential $df_p\big|_V = \langle p\,,\ \cdot\,\rangle - \mathcal{G}(V,\ \cdot\ )$ and the maximum is obtained when $p = \mathcal{G}(V,\ \cdot\ )$ which is equivalent to $V = \widetilde{\mathcal{G}}^{-1}(p)$. Expressing this in coordinates with $p = \sum_{j=1}^{d} p_j\, \chi^j\big|_x$ and $V = \sum_{j=1}^{d} v^j\, \mathcal{B}_j\big|_x$ we find that $v^i = p^i$ for all i (recall that p^i are reciprocal coordinates (3.10)). Thus, the coordinate formula for the hamiltonian $\mathfrak{h}$ is given by

$$\mathfrak{h} = \sup_{v^1,\dots,v^d \in \mathbb{R}} \left(\sum_{i=1}^{d} \mu_i v^i - \frac{1}{2}\sum_{i,j=1}^{d} g_{ij} v^i v^j \right) = \frac{1}{2}\sum_{i=1}^{d} \mu_i \mu^i \ . \tag{3.17}$$

Remark that we use the coordinate functions $\mu_i, \mu^i \in \mathcal{C}^\infty(T^*M)$ instead of specific coefficients $p_i, p^i \in \mathbb{R}$.

As stated before, the Pontryagin Maximum Principle implies that if $\gamma(t) \in M$ is a shortest curve, its momentum $\nu(t) = (\gamma(t), \lambda(t))$ is a flowline of the lifted hamiltonian $\overrightarrow{\mathfrak{h}}$, i.e. $\dot\nu(t) = \overrightarrow{\mathfrak{h}}(\nu(t))$. The flowline ODE can be split into a horizontal and a vertical part. We split the lifted hamiltonian $\overrightarrow{\mathfrak{h}}$ as in (3.6) using frame $\{\mathcal{B}_i\}_{i=1}^d$:

$$\overrightarrow{\mathfrak{h}} = \sum_{j=1}^{d} \frac{\partial \mathfrak{h}}{\partial \mu_j} \widehat{\mathcal{B}_j} - \widehat{\mathcal{B}_j}(\mathfrak{h})\partial_{\mu_j} \ . \tag{3.18}$$

Next, recall that we can write the momentum curve $\nu(t) \in T^*M$ as

$$\nu(t) = \sum_{j=1}^{d} p_j(t)\, \chi^j\big|_{\gamma(t)} \ ,$$

where $p_j(t) := \mu_j(\nu(t))$ so its derivative is given by

$$\dot\nu(t) = \sum_{j=1}^{d} \dot p_j(t)\, \partial_{\mu_j}\big|_{\nu(t)} + \dot\gamma^j(t)\, \widehat{\mathcal{B}_j}\Big|_{\nu(t)} \ . \tag{3.19}$$

Combining (3.18) and (3.19), the flowline ODE is decoupled:

$$\begin{cases} \dot\gamma^i = \partial_{\mu_i}(\mathfrak{h}) & \text{(horizontal part)} \\ \dot p_i = \widehat{\mathcal{B}_i}(\mathfrak{h}) & \text{(vertical part)} \end{cases}$$

We continue with the computation of these parts.

Horizontal part

We have already characterized the hamiltonian $\mathfrak{h}$ w.r.t. the coordinates μ_i in (3.17), so it remains to

compute the derivatives.

$$\frac{\partial \mathfrak{h}}{\partial \mu_i} \overset{(3.17)}{=\!=} \frac{\partial}{\partial \mu_i}\left(\frac{1}{2}\sum_{j=1}^{d}\mu_j\mu^j\right) \overset{(3.10)}{=\!=} \frac{1}{2}\frac{\partial}{\partial \mu_i}\left(\sum_{j,k=1}^{d}\mu_j g^{jk}\mu_k\right) = \frac{1}{2}\frac{\partial}{\partial \mu_i}\left(\sum_{j=1}^{d}\sum_{\substack{k=1\\k\neq j}}^{d}\mu_j g^{jk}\mu_k + \sum_{j=1}^{d} g^{jj}(\mu_j)^2\right)$$

$$= \frac{1}{2}\left(\sum_{j=1}^{d}\sum_{\substack{k=1\\k\neq j}}^{d}(\delta_{ij}g^{jk}\mu_k + \mu_j g^{jk}\delta_{ik}) + \sum_{j=1}^{d} g^{jj}2\mu_j\delta_{ij}\right) = \frac{1}{2}\left(\left(\sum_{\substack{k=1\\k\neq i}}^{d} g^{ik}\mu_k + \sum_{\substack{j=1\\j\neq i}}^{d}\mu_j g^{ji}\right) + 2g^{ii}\mu_i\right)$$

$$= \sum_{j=1}^{d} g^{ij}\mu_j \overset{(3.10)}{=\!=} \mu^i \ .$$

Thus, for the horizontal part we have that $\dot{\gamma}^i(t) = \mu^i(\nu(t))$. Remark that $\mu^i(\nu(t)) = p^i(t)$ by definition, so we have that the momentum curve $p(t) \in T^*_{\gamma(t)}M$ satisfies

$$p(t) = \widetilde{\mathcal{G}}(\dot{\gamma}(t)) \ . \tag{3.20}$$

Vertical part

Unlike the differential equation in the horizontal part, the vertical equation $\dot{p}_i = \widehat{\mathcal{B}_i}(\mathfrak{h})$ is not so easily computed directly. This is where the three lemmas from Section 3.2.3 about the Poisson bracket come in.

By Liouville's theorem (Lemma 16), we have that

$$\dot{p}_i(t) \overset{\text{def}}{=\!=} \frac{d(\mu_i \circ \nu)}{dt}(t) \overset{\text{Lemma 16}}{=\!=\!=} \{\mathfrak{h}, \mu_i\}(\nu(t)) \ . \tag{3.21}$$

We use the bilinearity and Leibniz rule of the Poisson bracket from Lemma 14 to compute the right hand side:

$$\begin{aligned}\{\mathfrak{h}, \mu_i\} &= \frac{1}{2}\sum_{k=1}^{d}\left\{\mu_k\mu^k, \mu_i\right\} = \frac{1}{2}\sum_{k,\ell=1}^{d} g^{k\ell}\left\{\mu_k\mu\ell, \mu_i\right\} \\ &= \frac{1}{2}\sum_{k,\ell=1}^{d} g^{k\ell}\left(\mu_k\left\{\mu_\ell, \mu_i\right\} + \left\{\mu_k, \mu_i\right\}\mu_\ell\right) \\ &= \frac{1}{2}\sum_{\ell=1}^{d}\left\{\mu_\ell, \mu_i\right\}\mu^\ell + \frac{1}{2}\sum_{k=1}^{d}\left\{\mu_k, \mu_i\right\}\mu^k = -\sum_{j=1}^{d}\left\{\mu_i, \mu_j\right\}\mu^j \ .\end{aligned} \tag{3.22}$$

Remark that the coordinate functions $\mu_i \in \mathcal{C}^\infty(T^*M)$ are of the form

$$\mu_i = \langle\, \cdot\,, \mathcal{B}_i\rangle = a_{\mathcal{B}_i} \ , \tag{3.23}$$

with the linear mapping $a_{(\cdot)} : \mathfrak{X}(M) \to \mathcal{C}^\infty(T^*M)$ from lemma 15 relating the Poisson and Lie brackets. Using the functions $C_{ij}^k \in \mathcal{C}^\infty(M)$ defined by $[\mathcal{B}_i, \mathcal{B}_j] = \sum_{k=1}^{d} C_{ij}^k\mathcal{B}_k$ as in Corollary 18.1, this gives us that

$$\begin{aligned}\{\mathfrak{h}, \mu_i\} &\overset{(3.22)}{=\!=} -\sum_{j=1}^{d}\left\{\mu_i, \mu_j\right\}\mu^j \overset{(3.23)}{=\!=} -\sum_{j=1}^{d}\left\{a_{\mathcal{B}_i}, a_{\mathcal{B}_j}\right\}\mu^j \overset{\text{Lemma 15}}{=\!=\!=} -\sum_{j=1}^{d} a_{[\mathcal{B}_i,\mathcal{B}_j]}\mu^j \\ &= -\sum_{j=1}^{d} a_{\left(\sum_{k=1}^{d} C_{ij}^k\mathcal{B}_k\right)}\mu^j \overset{\text{linearity and (3.23)}}{=\!=\!=\!=\!=\!=\!=\!=\!=\!=\!=} -\sum_{j,k=1}^{d} C_{ij}^k\mu_k\mu^j \ .\end{aligned} \tag{3.24}$$

Combining (3.21) and (3.24) with $\mu_i(\nu(t)) = p_i(t)$, we find that the differential equation for the vertical part becomes

$$\dot{p}_i(t) = -\sum_{j,k=1}^{d} C_{ij}^k(\gamma(t))\, p_k(t)\, p^j(t) \ .$$

Using the result from the horizontal part (3.20), we can rewrite this just a bit more:

$$\dot{p}_i(t) + \sum_{j,k=1}^{d} C_{ij}^k(\gamma(t))\, p_k(t)\, \dot{\gamma}^j(t) = 0 \ . \tag{3.25}$$

Proof part I

To prove part I of our claim, now simply substitute the result from the vertical part (3.25) into formula (3.15) for the change in momentum $\widehat{\nabla}^*_{\dot{\gamma}}\, p$ from Section 3.4.

Proof part II

For the second part of the theorem, recall that because of the antisymmetry of the Lie bracket, the functions $C_{ij}^k \in \mathcal{C}^\infty(T^*M)$ are also antisymmetric in the lower indices. This implies that formula (3.14) for the acceleration $\widehat{\nabla}_{\dot{\gamma}}\dot{\gamma}$ reduces to

$$\widehat{\nabla}_{\dot{\gamma}}\dot{\gamma} \overset{(3.14)}{=\!=} \sum_{k=1}^{d} \left(\ddot{\gamma}^k + \sum_{i,j=1}^{d} \dot{\gamma}^i \dot{\gamma}^j C_{ij}^k \right) \mathcal{B}_k = \sum_{k=1}^{d} \ddot{\gamma}^k \mathcal{B}_k \ ,$$

which is zero if and only if $\dot{\gamma}^i$ is constant for all i. ■

Appendix A

Formalisation of the Covariant Derivative along a Curve

This appendix shows how to formalize the concept of a covariant derivative along a curve, like the velocity $\nabla_{\dot\gamma}\dot\gamma$ along the curve $\gamma(t) \in M$. Although we write $\nabla_{\dot\gamma}\dot\gamma$, this is not well defined as the second argument $\dot\gamma$ can have multiple values at points where γ intersects itself. Hence, it cannot be considered as a vector field over γ. The solution is to actively use the time parameter $t \in [0,1]$, resulting in a description that more closely matches the intuitive definition. To do this, we need the language of pullback bundles and pullback connections.

Definition 15 (Pullback bundle)**.** Given a vector bundle $\pi_N : E \to N$ and a smooth map $\varphi : M \to N$, we define the pullback bundle φ^*E as the set

$$\varphi^*E := \{(x, v) \in M \times E \mid \varphi(x) = \pi_N(v)\} ,$$

i.e. we have the equality of fibers $(\varphi^*E)\big|_x = E\big|_{\varphi(x)}$. The projection $\pi_M : \varphi^*E \to M$ is $\pi_M(x, v) = x$. The smooth structure of φ^*E is specified by the smooth sections of π_M, which are generated under the $\mathcal{C}^\infty(M)$-module structure by the maps $s \circ \varphi$ where $s : N \to E$ is a smooth section of π_N.

To keep track of all objects involved, one draws the commutative diagram

$$\begin{array}{ccc} \varphi^*E & \xrightarrow{\ \iota\ } & E \\ \big\downarrow{\scriptstyle \pi_M} & & \big\downarrow{\scriptstyle \pi_N} \\ M & \xrightarrow{\ \varphi\ } & N \end{array}$$

where ι is the embedding given by $\iota(x, v) = v$.

In the case of $\nabla_{\dot\gamma}\dot\gamma$, the pullback bundle $\gamma^*(TM)$ can be seen as a space consisting of copies of the vector spaces $T_{\gamma(t)}M$ strung together by the temporal parameter $t \in [0,1]$. As the next lemma shows, there is a unique section of $\gamma^*(TM) \to [0,1]$ that corresponds to the the velocity curve $\dot\gamma(t) \in T_{\gamma(t)}M$.

Lemma 21. *Let $\pi_M : \varphi^*E \to M$ be a pullback of the vector bundle $\pi_N : E \to N$ along a smooth map $\varphi : M \to N$. Assume that $f : M \to E$ is a smooth map such that $f(x) \in E\big|_{\varphi(x)}$ for all $x \in M$, then there is a unique smooth section f^* of π_M such that $f = \iota \circ f^*$.*

Proof. Construct the map $f^* : M \to \varphi^*E$ by $f^*(x) := (x, f(x)) \in \varphi^*E$. Clearly f^* satisfies $f = \iota \circ f^*$. It is unique because we could not have made another choice for the second argument in $(x, f(x)) \in \varphi^*E$ such that this equality would hold. Smoothness of f^* follows from smoothness of f. ∎

Hence we can speak of the section $\dot\gamma^*$ instead of the velocity curve $\dot\gamma : [0,1] \to M$. Taking the derivative of $\dot\gamma$ along γ is done by applying a pullback connection on $\gamma^*(TM)$ to the section $\dot\gamma^*$.

Definition 16 (Pullback connection)**.** Given a vector bundle $\pi_N : E \to N$ with a connection ∇, we define the pullback connection $\varphi^*(\nabla)$ on the pullback bundle $\pi_M : \varphi^*E \to M$ along $\varphi : M \to N$ by requiring that it satisfies the Leibniz rule w.r.t. elements of $\mathcal{C}^\infty(M)$ and that it is compatible with sections Y of π_N, i.e.

$$(\varphi^*(\nabla))_X \, (Y \circ \varphi)^* := ((\nabla_{\varphi_* X} Y) \circ \varphi)^* \tag{A.1}$$

for all vector fields $X \in \mathfrak{X}(M)$.

We now have all the ingredients to define the covariant derivative along a curve.

Definition 17 (Covariant derivative along a curve)**.** Let $s : [0,1] \to M$ be a vector field in the bundle $E \to M$ over the curve $\gamma : [0,1] \to M$, i.e. $s(t) \in E\big|_{\gamma(t)}$, then the covariant derivative of s along γ w.r.t. to the connection ∇ is defined as

$$\nabla_{\dot\gamma} \, s := \iota \circ \left((\gamma^* \nabla)_{\partial_t} \, s^* \right) \ ,$$

where ι is the embedding $\gamma^* E \to E$ and ∂_t is the vector field in $T[0,1]$ induced by the temporal parameter $t \in [0,1]$.

Intuitively, we encode how the vector field s changes along the curve γ with the section $s^* : [0,1] \to \gamma^* E$. Note that the fact that s is a vector field over γ is now completely encoded in the space $\gamma^* E$, so the section s^* only depends on the temporal parameter t. To take the derivative of s along γ, we then simply take the derivative of s^* w.r.t. t via the corresponding vector field $\partial_t \in \mathfrak{X}[0,1]$. The embedding at the end decodes the result back into a vector field in E over γ.

We now proof the coordinate formula for the covariant derivative (3.13), Lemma 19.

Proof. Given a frame $\{\mathcal{E}_i\}_{i=1}^d$ for the vector bundle $E \to M$, we can decompose the vector field $s : [0,1] \to E$ as

$$s(t) = \sum_{j=1}^{d} s^j(t) \, \mathcal{E}_j\big|_{\gamma(t)} \ .$$

This allows us to rewrite the derivative $(\gamma^*\nabla)_{\partial_t} \, s^*$ by using the Leibniz rule for the pullback connection $\gamma^*\nabla$ with $s^i \in \mathcal{C}^\infty[0,1]$:

$$\begin{aligned}
(\gamma^*\nabla)_{\partial_t} \, s^* &= (\gamma^*\nabla)_{\partial_t} \left(\sum_{j=1}^{d} s^j \, (\mathcal{E}_j \circ \gamma) \right)^* \\
&= (\gamma^*\nabla)_{\partial_t} \left(\sum_{j=1}^{d} s^j \, (\mathcal{E}_j \circ \gamma)^* \right) \\
&= \sum_{j=1}^{d} \left(\partial_t(s^j) \, (\mathcal{E}_j \circ \gamma)^* + s^j \ (\gamma^*\nabla)_{\partial_t} \, (\mathcal{E}_j \circ \gamma)^* \right) \\
&= \sum_{j=1}^{d} \left(\partial_t(s^j) \, (\mathcal{E}_j \circ \gamma)^* + s^j \, ((\nabla_{\dot\gamma} \, \mathcal{E}_j) \circ \gamma)^* \right) \ ,
\end{aligned}$$

where the last equality is the compatibility of the pullback connection $\gamma^*\nabla$ with the sections $\mathcal{E}_i$, see (A.1).

By the definition of the derivative $\nabla_{\dot\gamma}\, s$, we then have that

$$\begin{aligned}
\left(\nabla_{\dot\gamma}\, s\right)(t) &= \left(\iota \circ \left(\left(\gamma^*\nabla\right)_{\partial_t}\, s^*\right)\right)(t) \\
&= \left(\iota \circ \left(\sum_{j=1}^{d} \partial_t(s^j)\left(\mathcal{E}_j \circ \gamma\right)^* + s^j\left(\left(\nabla_{\dot\gamma}\mathcal{E}_j\right)\circ\gamma\right)^*\right)\right)(t) \\
&= \iota\left(\sum_{j=1}^{d}\left(\partial_t(s^j)\right)(t)\left(t, \mathcal{E}_j\big|_{\gamma(t)}\right) + s^j(t)\left(t, \left(\nabla_{\dot\gamma}\,\mathcal{E}_j\right)\big|_{\gamma(t)}\right)\right) \\
&= \iota\left(t, \sum_{j=1}^{d}\frac{ds^j}{dt}(t)\,\mathcal{E}_j\big|_{\gamma(t)} + s^j(t)\left(\nabla_{\dot\gamma(t)}\,\mathcal{E}_j\big|_{\gamma(t)}\right)\right) \\
&= \sum_{j=1}^{d}\left(\frac{ds^j}{dt}(t)\,\mathcal{E}_j\big|_{\gamma(t)} + s^j(t)\left(\nabla_{\dot\gamma(t)}\,\mathcal{E}_j\big|_{\gamma(t)}\right)\right)\ .
\end{aligned}$$

■

Bibliography

[1] Andrei Agrachev and Yuri Sachkov. *Control Theory from the Geometric Viewpoint*. Springer-Verlag Berlin Heidelberg, 1 edition, 2004.

[2] E.J. Bekkers. *Retinal Image Analysis using Sub-Riemannian Geometry in $SE(2)$*. PhD thesis, Technische Universiteit Eindhoven, 2017.

[3] E.J. Bekkers, R. Duits, T. Berendschot, and B. Haar Romeny. A multi-orientation analysis approach to retinal vessel tracking. *JMIV*, 49(3):583–610, 2014.

[4] E.J. Bekkers, Duits R., A. Mashatkov, and G.R. Sanguinetti. A PDE approach to data-driven sub-Riemannian geodesics in $SE(2)$. *SIAM Journal of Imaging Science*, 8(4):2740–2770, 2015.

[5] R. Duits. *Perceptual Organization in Image Analysis*. PhD thesis, Technische Universiteit Eindhoven, 2005.

[6] R. Duits, E.J. Bekkers, and J.M. Portegies. Part III: Tracking in Orientation Scores - Optimal Paths for Variants of the 2D and 3D Reeds-Shepp Car, 2019. Part of a volume of lecture notes for a course on Differential Geometry in Image Processing taught by R. Duits.

[7] R. Duits, M.H.J. Janssen, J. Hannink, and G.R. Sanguinetti. Locally adaptive frames in the roto-translation group and their applications in medical image processing. *JMIV*, 56(3):367–402, 2016.

[8] R. Duits, S. P. L. Meesters, J-M. Mirebeau, and J. M. Portegies. Optimal Paths for Variants of the 2d and 3d Reeds-Shepp Car with Applications in Image Analysis. *JMIV*, 60(6):816—-848, July 2018.

[9] R. Duits, B.M.N. Smets, A.J. Wemmenhove, J.W. Portegies, and E.J. Bekkers. Geometric flows in multi-orientation image processing via a cartan connection., 2020. Retrieved from https://www.win.tue.nl/ rduits/BookchapterDuitsetal.pdf on 7 July 2020.

[10] Remco Duits, Tom Dela Haije, Eric Creusen, and Arpan Ghosh. Erratum to: Morphological and linear scale spaces for fiber enhancement in dw-mri. *Journal of Mathematical Imaging and Vision*, 46:326–368, 07 2013.

[11] Remco Duits, Maurice Duits, Markus van Almsick, and Bart ter Haar Romeny. Invertible orientation scores as an application of generalized wavelet theory. *Pattern Recognition and Image Analysis*, 17:42–75, 01 2007.

[12] Remco Duits and Erik Franken. Left-invariant parabolic evolutions on $SE(2)$ and contour enhancement via invertible orientation scores part i: Linear left-invariant diffusion equations on $SE(2)$. *Quarterly of Applied Mathematics*, 68, 04 2010.

[13] Remco Duits, Etienne St-Onge, Jim Portegies, and Bart Smets. Total variation and mean curvature pdes on the space of positions and orientations. In Jan Lellmann, Martin Burger, and Jan Modersitzki, editors, *Scale Space and Variational Methods in Computer Vision*, pages 211–223, Cham, 2019. Springer International Publishing.

[14] Lawrence C. Evans. *Partial Differential Equationa*. American Mathematical Society, 2 edition, 1998.

[15] E. M. Franken and R. Duits. Crossing preserving coherence-enhancing diffusion on invertible orientation scores. *International Journal of Computer Vision (IJCV)*, 85(3):253–278, 2009.

[16] Frank Martin, Erik Bekkers, and Remco Duits. Lie analysis for medical image analysis, 2017. Retrieved from https://www.lieanalysis.nl/ on 7 July 2020.

[17] G. Sanguinetti, E.J. Bekkers, R. Duits, M. H. J. Janssen, A. Mashtakov, and J.-M. Mirebeau. *Sub-Riemannian Fast Marching in SE(2)*, pages 366–376. Springer, 2015.

[18] Jiong Zhang, Behdad Dashtbozorg, Erik Bekkers, Josien P. W. Pluim, Remco Duits, and Bart M. ter Haar Romeny. Robust retinal vessel segmentation via locally adaptive derivative frames in orientation scores. *IEEE Transactions on Medical Imaging*, 35:2631–2643, 12 2016.