

Contents

1. Introduction	3
2. Background	8
2.1. Gibbs Random Fields	8
2.1.1. Principle of Least Action	9
2.1.2. Image De-Noising	10
2.2. Lie Groups and the Noether Theorem	13
2.2.1. Motivation 1, the problem	13
2.2.2. Motivation 2, the solution	15
2.3. Lie Groups	18
2.3.1. The Group $\mathbb{G} = \mathbb{T} \times SO(2)$	22
2.4. Noether's First Theorem	23
2.4.1. Noethers Theorems	25
2.4.2. Noether's First Theorem: A Modern Version	28
2.4.3. Pure Spacial Symmetries	30
2.5. Total Variation	31
2.5.1. The Mean Curvature of Total Variation	33
2.6. Optical Flow	35
2.6.1. Uni-Modal Optical Flow	37
2.6.2. Multi-Modal Optical Flow	38
2.7. Image Fusion	42
3. Linearized Priors	46
3.1. The Linear Structure Tensor	46
3.2. Structure Tensor Based Prior	48
3.3. Geometrical Optical Flow Model	49
3.4. Multi-Modal Optical Flow with Differing Resolutions	50
3.5. Localization	53
3.6. The Multigrid Newton algorithm	54
3.7. Results	56
3.7.1. Uni-Modal Data	57
3.7.2. Rubber Whale Sequence	58
3.7.3. Hydrangea Sequence	60
3.7.4. Estimation of the Scale Difference σ^{sc}	60

3.7.5. Real Multimodal Optical Flow Data	63
3.7.6. Eigenvalue analysis and the stabilization parameter λ_2 . . .	70
3.7.7. Summary	71
4. The Extended Least Action Algorithm	75
4.0.8. Newtonian Minimization	78
4.0.9. The dynamics of the level-sets S	79
4.1. The Extended Least Action Algorithm	84
4.1.1. Image De-noising	89
4.2. summary	97
5. Conclusions	99
5.1. Outlook	103
A. Lie Groups	104
A.1. The Prolonged Action	104
A.2. Geometrical Meaning of the Commutator $[\cdot, \cdot]$	105
A.3. Derivation Of Noethers Theorem	106
A.3.1. Connection between $\mathbf{B}_m, \mathbf{W}_m$ and $[\mathcal{E}]$	109
B. The Bending Algebra	111
B.1. The curvature operator	111
B.2. TV Image Denoising, supplementary results	114
C. Multimodal Optical Flow	119
Bibliography	122

1. Introduction

The main topic of this thesis is concerned about symmetries in the mathematical modeling of computer vision problems. Many objects in nature possess among others the notable characteristic of symmetry regarding their attributes such as their form and color. A symmetry of an object \mathcal{O} is such that if \mathcal{O} undergoes a specific transformation g , then it appears for an observer to be unchanged. Say we have a computer vision problem involving the object \mathcal{O} , modeled with a mathematical model M . It is natural to reflect the symmetry of the object \mathcal{O} within the model M , such that M is invariant in some sense under the transformation g . The goal of this thesis is to analyze the structure of the symmetries of a mathematical model M . We will prove that knowledge of the symmetries of M may lead to significant speed ups of any algorithm using M .

Symmetries generally fall into two categories: global and local symmetries. A ball of uniform color for instance does not change its appearance to an observer upon rotation around an arbitrary axis through the center of the ball. This example is one of *global* symmetry since the ball as a whole is transformed (rotated). We can formally describe the global symmetry of the object \mathcal{O} in the following way: If the surface of the object is described by the functional relationship $\phi_{\mathcal{O}}(\mathbf{x}) = \text{const}$ (e.g. $\phi_{\mathcal{O}}(\mathbf{x}) = x^2 + y^2 + z^2 = 1$ for a ball of unit radius) then our intuition of global symmetry is equal to the statement that $\phi_{\mathcal{O}}(\mathbf{x}) = \text{const}$ is invariant under the global transformation $\mathbf{x}' = g \circ \mathbf{x}$

$$\phi_{\mathcal{O}}(g \circ \mathbf{x}) = \phi_{\mathcal{O}}(\mathbf{x}) \tag{1.1}$$

Not all objects in nature are symmetric with respect to global transformations. For example in figure 1.1 an image of a leaf is shown. Since the leaf is not symmetric with respect to any global transformation g , its projection onto the image plane Ω is not symmetric with respect to any global transformation g^{Ω} on Ω . However if we inspect local regions of the leaf, that is we zoom into those regions at various locations on the leaf, we see that the features of the leaf within the regions do possess symmetries. Figure 1.1b shows a close up of the region highlighted in figure 1.1a through which a vein of the leaf runs. The vein appears to be linear and thus symmetric towards translations along its tangential direction. This symmetry is reflected by the vectors at each position of the vein. They indicate

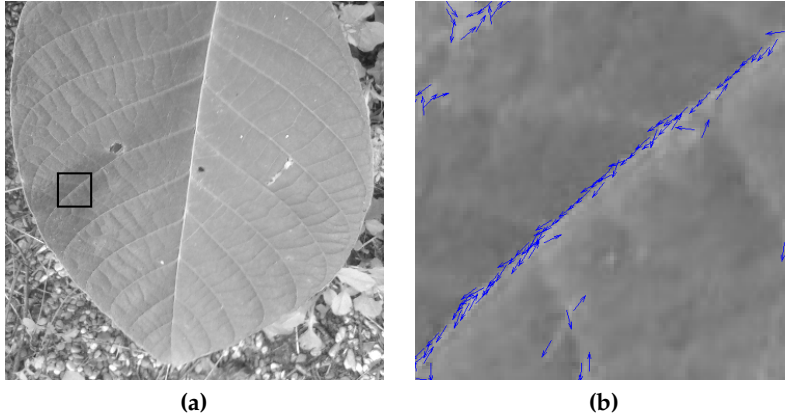


Figure 1.1.: Figure 1.1a shows an image of a leaf. The leaf clearly has no global symmetry. Figure 1.1b shows a close-up of the region around a vein of the leaf, indicated by the box in figure 1.1a. The vectors in figure 1.1b along the vein indicate local translations which leave the vein invariant.

local translations, which leave the vein invariant. A local transformation as indicated by the vectors in figure 1.1b may be represented by the vector field $\mathbf{b}(\mathbf{x})$ such that the local transformation $g^{\mathbf{B}}(\mathbf{x}) = \mathbf{x} + \mathbf{b}(\mathbf{x})$ leaves the image ϕ invariant

$$\phi(\mathbf{x} + \mathbf{b}(\mathbf{x})) = \phi(\mathbf{x}) \quad (1.2)$$

In general we cannot assume that $g^{\mathbf{B}}(\mathbf{x})$ in eq. (1.2) is unique since there can always exist a vector field $\mathbf{b}'(\mathbf{x}) \neq \mathbf{b}(\mathbf{x})$ which satisfies eq. (1.2). On the other side any transformation $g^{\mathbf{B}}$ satisfying eq. (1.2) uniquely determines the geometry of ϕ for if we were to draw lines along the tangential vectors $\mathbf{b}(\mathbf{x})$ by connecting \mathbf{x} with $\mathbf{x} + \mathbf{b}(\mathbf{x})$ we would reconstruct the object \mathcal{O} from $g^{\mathbf{B}}(\mathbf{x})$.

Computational modeling in general is concerned about acquiring information about our physical reality given a set of observations. For instance a typical computer vision problem is the classification of the contents of a given image $\phi_{\mathcal{O}}$. Let \mathbb{P} represent the classification problem and $M_{\mathbb{P}}$ a mathematical model (a set of formulas) describing \mathbb{P} . If $\phi_{\mathcal{O}}$ is the image of the leaf in figure 1.1 then an algorithm deploying $M_{\mathbb{P}}$ should indeed verify that $\phi_{\mathcal{O}}$ is the image of a leaf and additionally report some information about the leaf like the size, the name of the plant etc. The geometric invariance of $\phi_{\mathcal{O}}$ towards the transformation g in eq. (1.2) encodes information about the geometry of the leaf. The classification model $M_{\mathbb{P}}$ can be made aware of the geometric invariance (eq. (1.2)) and the local geometric structures of the leaf (e.g the veins) by constructing the mathematical formulas of $M_{\mathbb{P}}$ such that they themselves are invariant under the same spacial




Figure 1.2.: Figure 1.2a: Two cameras are shown recording a scene from different positions. The scene could be a rigid scene or a dynamic scene with moving objects. Figure 1.2b shows the image y captured from the camera y_c and figure 1.2c the image I from the camera I_c . One possible question is: How can the pixels of the image I be mapped to those of the image y ? Such a mapping can be used to deduce the 3-dimensional structure of the box similar to how the human brain constructs a 3-dimensional image given the 2-dimensional images obtained by the left and the right eye.

transformations g . This invariance of $M_{\mathbb{P}}$ under g would make it possible for $M_{\mathbb{P}}$ to identify different regions of the image $\phi_{\mathcal{O}}$ with dominant linear structures and furthermore identifying these linear structures as belonging to one global structure, e.g. the vein of the leaf. In general the concept of constructing computational models integrating known geometric invariances of the data for the aid of acquiring information in the data is at the core of this thesis.

The process of acquiring information from our physical reality via mathematical modeling is problematic itself in many ways. For one, the information which we may wish to gather may lay hidden in the data we can possibly acquire from a physical system. One such problem is called stereography ([1, 2, 3, 4]), depicted in figure 1.2. The statement of the problem goes as follows: given two images y and I (figures 1.2b and 1.2c) of an object \mathcal{O} (the box in figure 1.2a) how can we infer the 3-dimensional structure of \mathcal{O} (the width, height and depth of the box)? This problem has already been solved by nature since the human brain is capable of reconstructing a 3-dimensional image given the 2-dimensional images obtained by the left and the right eye.

Besides the problem of hidden information described above there is another problem in the process of information acquisition. The means we use to acquire the data have technical limitations. For instance the cameras y_c and I_c in figure 1.2a in general produce images of limited resolutions which may also be subject to noise.

Both problems in the process of information acquisition may be sub-summed as the problem of *inference*: Given some possibly corrupted data Y of a physical

system we wish to infer some information stored in the unknown latent variable ϕ . In general Y and ϕ may be discrete variables, continuous functions over some domain Ω or a combination of both. In this thesis we will only handle problems for which Y and ϕ are continuous functions over Ω

$$Y, \phi : \Omega \rightarrow \mathbb{R}^n \quad (1.3)$$

The inference problem then becomes the problem of mapping Y to ϕ

$$Y(\mathbf{x}) \xrightarrow{T_Y} \phi(\mathbf{x}) \quad (1.4)$$

where T_Y denotes a process or an algorithm which is parametrized by the data Y . Since the variable ϕ is unknown we have to for one make assumptions on its geometric properties and furthermore model how it is linked to the data Y . These aspects of ϕ are then embedded in the inference process T_Y . For now we want to motivate how the geometrical properties of ϕ can be taken into account by T_Y . Consider a local transformation g such that the variable ϕ is transformed to the new variable ϕ'

$$\phi'(\mathbf{x}) = \phi(g \circ \mathbf{x}) \quad (1.5)$$

We can regard ϕ' as being inferred from the data Y via the inference process T'_Y similar to eq. (1.4). If ϕ is symmetric under g in the sense of eq. (1.2) then this implies that the two inference processes T_Y and T'_Y are equal and thus the inference process T_Y is itself symmetric under the action of g . We conclude that knowledge of the set of local transformations $\{g\}$ which satisfy eq. (1.2) allows us to identify those inference processes T_Y which are equal to each other upon action of $\{g\}$. This has two consequences. The first is that we can *design* an inference process T_Y which is invariant upon the action of the set $\{g\}$. As a result this guarantees the invariance of ϕ upon the action of $\{g\}$. The second consequence is more subtle. If we split the inference process T_Y into n intermediate steps

$$Y \xrightarrow{T_Y} \phi = Y \xrightarrow{T_Y^1} \phi^1 \xrightarrow{T_Y^2} \phi^2 \dots \xrightarrow{T_Y^{n-1}} \phi^{n-1} \xrightarrow{T_Y^n} \phi \quad (1.6)$$

the intermediate steps T_Y^i and ϕ^i need *not* be invariant under the set $\{g\}$. However for particularly well chosen $g' \in \{g\}$ such that

$$g' \circ T_Y^i = T_Y^{i+k} \quad (1.7)$$

we may minimize the number steps in eq. (1.6) and thus obtain the *shortest path* in the inference problem.

The overall structure of this thesis is as follows: In section 2.1 we introduce the

latent variable ϕ as a Gibbs Random Field (GRF). The main property of GRFs is that they are associated with an energy functional $E_Y(\phi)$. The inference process T_Y is explicitly formulated as the minimization problem

$$\phi^* = \operatorname{argmin}_{\phi} E_Y(\phi) \quad \leftrightarrow \quad Y \xrightarrow{T_Y} \phi^* \quad (1.8)$$

In section 2.3 we will introduce the definition of an r -dimensional Lie group \mathbb{G} and its corresponding Lie algebra \mathcal{G} . This facilitates the formally correct definition of the local symmetry in eq. (1.2) in the form of the level-set equation

$$X\phi = 0 \quad \text{if} \quad \phi(g \circ \mathbf{x}) = \phi(\mathbf{x}), \quad g = \exp(tX) \in \mathbb{G}, \quad X \in \mathcal{G} \quad (1.9)$$

Sections 2.1 and 2.3 prepare the stage for the introduction of Emmy Noethers celebrated first theorem in section 2.4. In a nutshell this theorem states that if an energy functional $E_Y(\phi)$ is invariant upon the action of an r -dimensional Lie group \mathbb{G} , then there exists r *divergence-free* vector fields \mathbf{W}_m

$$g \circ E_Y(\phi) = E_Y(\phi) \quad \forall g \in \mathbb{G} \quad \leftrightarrow \quad \exists \mathbf{W}_m, \quad \operatorname{div}(\mathbf{W}_m) = 0 \quad \forall \quad 1 \leq m \leq r \quad (1.10)$$

Since its first publication in 1918, Noether's first theorem has had far reaching implications in our understanding of the fundamental laws of motion in physics as well as the deep connection between the symmetries of a physical system and its conservation laws. For instance the time invariance of the laws of motion in the universe reveals the conservation of energy. In layman words: It does not matter if we carry out an experiment now or next week, the results will be the same since the energy of the universe does not vanish! Building on section 2.4 we demonstrate in section 3 the construction of a prior energy functional $E^{prior}(\phi)$ which is invariant under the Lie group $\mathbb{T} \times SO(2)$ which is the group of local translations and rotations. In section 3.3 we will use the prior developed in section 3 in the context of optical flow [5]. In section 4 we will introduce a generalization of the Newton approach for solving the inference problem in eq. (1.8) which takes local transformations of the spatial coordinates \mathbf{x} in Ω (see eq. (1.2)) into account to facilitate the search for the shortest path in the inference problem in eq. (1.6).

2. Background

2.1. Gibbs Random Fields

A physical system C is a dynamical composite of elements which interact with each other as well as with the environment the system C is embedded in. The elements are described by a vector of parameters $\phi = (\phi_1, \dots, \phi_n)$. The physical system C relates a specific value ϕ^* of the vector ϕ to a set of observables $Y = \{Y_1, \dots, Y_k\}$

$$Y = C(\phi^*) \quad (2.1)$$

In the case that the elements of the system C are continuously distributed over a finite space Ω , the parameter vector ϕ is a function on Ω

$$\phi(\mathbf{x}) \in \mathbb{R}^n \quad \mathbf{x} \in \Omega \quad (2.2)$$

called a Gibbs-Random-Field (GRF) [6]. In this thesis we will generally assume that the GRF ϕ is infinitely differentiable, $\phi(\mathbf{x}) \in C^\infty(\Omega)$, unless otherwise explicitly stated. The interactions of the elements of the system C with the environment are characterized by an energy functional $E_Y^{data}(\phi)$ called the data term, which couples the GRF $\phi(\mathbf{x})$ to the observables Y . There is another energy form $E^{prior}(\phi, \nabla\phi)$ within the system C called the prior. $E^{prior}(\phi, \partial_j\phi)$ describes how the elements of C interact with each other. Together both energy functionals form the total energy of the system C

$$E_Y(\phi, \nabla\phi) = E_Y^{data}(\phi) + E^{prior}(\phi, \nabla\phi) \quad (2.3)$$

which is related to the probability distribution

$$p(\phi, \nabla\phi|Y) = p(Y|\phi) \cdot p(\phi, \nabla\phi) \sim \exp(-E_Y(\phi, \nabla\phi)) \quad (2.4)$$

$$p(Y|\phi) = \exp(-E_Y^{data}(\phi)) \quad (2.5)$$

$$p(\phi, \nabla\phi) = \exp(-E^{prior}(\phi, \nabla\phi)) \quad (2.6)$$

The value of the probability distribution $p(\phi, \nabla\phi|Y)$ evaluated at the values $\hat{\phi}(\mathbf{x})$ describes the probability that the GRF $\phi(\mathbf{x})$ assumes the values $\hat{\phi}(\mathbf{x})$ at each point $\mathbf{x} \in \Omega$. The set of values $\hat{\phi}(\mathbf{x})$ is what is called a *configuration* of the GRF ϕ .

$E_Y(\phi, \nabla\phi)$ is designed such that it is minimal once the GRF $\phi(\mathbf{x})$ fulfills the forward problem in eq. (2.1)

$$\phi^* = \operatorname{argmin}_{\phi} (E_Y(\phi, \nabla\phi)) \quad (2.7)$$

The particular value $\phi^*(\mathbf{x})$ of the GRF ϕ is the most probable configuration of the distribution $p(\phi|Y)$ due to eq. (2.4) and the solution to the inverse problem

$$\phi^* = C^{-1}(Y) \quad (2.8)$$

In this thesis we will not be concerned with the probability interpretation of GRF models via eq. (2.4) but rather directly with the formulation of a GRF model via the energy functional $E_Y(\phi, \nabla\phi)$ in eq. (2.3). The focus lies on the prior E^{prior} and we will show how it incorporates the geometrical invariance in eq. (1.2).

2.1.1. Principle of Least Action

Eq. (2.7) is the manifestation of the principle of least action (PLA) [7, 8, 9, 10, 11]. The origins of the PLA date back to the year 1740. In that year Pierre-Louis Maupertuis introduced the action E of a body of mass m as a quantity which is minimal when the body is at rest [8]. He further went on in 1746 to extend his principle to systems of non interacting point masses m_i which follow trajectories $\mathbf{q}_i(t)$ [12]. In that paper he showed that the action E of the many body system has the form of a time integral

$$E(\dot{\mathbf{q}}) = \int \frac{1}{2} \sum_i m_i \|\dot{\mathbf{q}}_i(t)\|^2 dt \quad (2.9)$$

Maupertuis showed that the trajectories $\mathbf{q}_i^*(t)$ which minimize eq. (2.9) are also the solutions to Newtons equations

$$m_i \ddot{\mathbf{q}}_i^*(t) = 0 \quad (2.10)$$

realizing Newtons first axiom *A body of mass m remains in the state of rest or uniform linear motion unless a force is acted upon it to change its state* [13].

The PLA was extended in 1750 by Leonhard Euler and Joseph Lagrange to incorporate interactions between the point masses m_i . They defined the action

$E(\dot{\mathbf{q}}, \mathbf{q})$ as the time integral over a function $\mathcal{L}(\dot{\mathbf{q}}, \mathbf{q}, t)$, now called the Lagrangian, and showed that $E(\dot{\mathbf{q}}, \mathbf{q})$ is minimal when the trajectories fulfill the equation

$$\frac{\delta \mathcal{L}}{\delta \mathbf{q}_i(t)} - \frac{d}{dt} \left(\frac{\delta \mathcal{L}}{\delta \dot{\mathbf{q}}_i(t)} \right) = 0 \quad (2.11)$$

The differential equations in eq. (2.11) are named in honor of the originators the Euler-Lagrange equations.

In the context of Gibbs random field theory and the minimization problem in eq. (2.7) the energy $E(\phi, \nabla \phi)$ can be written in terms of an integral over the Lagrangian $\mathcal{E}(\phi(\mathbf{x}), \nabla \phi(\mathbf{x}))$

$$E(\phi, \nabla \phi) = \int_{\Omega} \mathcal{E}(\phi(\mathbf{x}), \nabla \phi(\mathbf{x})) d^m x \quad (2.12)$$

The Euler-Lagrange equations in eq. (2.11) then generalize to the GRF field equations

$$\frac{\delta \mathcal{E}}{\delta \phi(\mathbf{x})} - \sum_{\mu} \frac{d}{dx_{\mu}} \left(\frac{\delta \mathcal{E}}{\delta \partial_{\mu} \phi(\mathbf{x})} \right) = 0 \quad (2.13)$$

The Euler-Lagrange GRF equations in eq. (2.13) are the central equation in many physical theories which can be modeled in the context of Gibbs random fields. For instance in classical electrodynamics [14, 15] the GRF of the theory is the vector $\mathbf{A}(\mathbf{x}, t)$ which generates the electric field $\mathbf{E}(\mathbf{x}, t)$ and the magnetic field $\mathbf{B}(\mathbf{x}, t)$. The Euler-Lagrange equations of the density $\mathcal{E}(\mathbf{A}(\mathbf{x}, t), \nabla \mathbf{A}(\mathbf{x}, t))$ are the celebrated Maxwell-Equations [16], which are dynamical equations relating $\mathbf{E}(\mathbf{x}, t)$ and $\mathbf{B}(\mathbf{x}, t)$. Through the Maxwell equations James Clerk Maxwell was able to predict the electromagnetic waves, the existence of which was proved by the physicist Heinrich Hertz [17].

2.1.2. Image De-Noising

An example of a physical system containing a GRF is a camera C recording an object O . The domain $\Omega \subset \mathbb{R}^2$ is the focal plane of the camera C and the object O is naturally projected onto the focal plane Ω producing the projection I_O . In theory the projection I_O is a continuous function in the coordinate frame of the plane O where the particular function value $I_O(\mathbf{x})$ is the light intensity the object O reflects to the point \mathbf{x} on the focal plane Ω . At the heart of the image acquisition process of basically all modern camera systems lies the concept of a




Figure 2.1.: Figure 2.1a shows an image I^c taken of an object O with a thermographic camera. A region of interest is shown where the contrast was enhanced to visualize the noise corruption. Figure 2.1b shows the result I_O^* of the minimization problem eq. (2.19) with the prior in eq. (2.20). The noise is removed but the boundaries of O are over smoothed

CCD collecting the photons of the light at discrete positions $\mathbf{x}_{i,j}$ called pixels

$$I_{ij}^c = I^c(\mathbf{x}_{i,j}) \in \mathbb{R}, \quad \mathbf{x}_{i,j} \in \Omega \quad 1 < i < n, 1 < j < m \quad (2.14)$$

The observables Y are the recorded intensities I_{ij}^c at the pixels $\mathbf{x}_{i,j}$. In this sense the camera C is a function which maps the continuous projection $I_O(\mathbf{x})$ to the discretely sampled intensities I_{ij}^c

$$I_{ij}^c = C_{ij}(I_O) \quad (2.15)$$

The intensity I_{ij}^c is basically a function of the number of photons collected by the CCD at the pixel $\mathbf{x}_{i,j}$. This number cannot be acquired deterministically, it is rather the result of a stochastic process described as independently identically distributed (iid) noise

$$\hat{I}_{ij}^c = I_O(\mathbf{x}_{i,j}) + n \quad n \sim p(I_{ij}^c | I_O(\mathbf{x}_{i,j})) \quad (2.16)$$

$p(I_{ij}^c | I_O(\mathbf{x}_{i,j}))$ is the likelihood that I_{ij}^c assumes the value \hat{I}_{ij}^c given the incoming intensity $I_O(\mathbf{x}_{i,j})$ at the pixel $\mathbf{x}_{i,j}$. Like in eq. (2.5) it is mapped to the data term energy $E_{I^c}(I_O)$.

In order to infer the values of $I_O(\mathbf{x}_{i,j})$ at the pixels $\mathbf{x}_{i,j}$ from the noisy data I_{ij}^c we need to pose some form of regularity on the values $I_O(\mathbf{x})$ to counter the pixel-wise noise imposed by the CCD in eq. (2.16). Such regularity can be achieved by

correlating the intensities $I_O(\mathbf{x})$ at all pixels with each other in the prior

$$p(I_O) = \exp\left(-E^{prior}(I_O)\right) \quad (2.17)$$

$$E^{prior}(I_O) = \int_{\Omega} \mathcal{E}(I_O(\mathbf{x}), I_O(\Omega/\{\mathbf{x}\})) dx \quad (2.18)$$

where the integrand correlates the intensity $I_O(\mathbf{x})$ at the point $\mathbf{x} \in \Omega$ with the intensities at all other points $\Omega/\{\mathbf{x}\}$ so that the problem of inferring I_O from the data I^c becomes the minimization problem

$$I_O^* = \operatorname{argmin}_{I_O} (E_{I^c}(I_O)), \quad E_{I^c}(I_O) = E_{I^c}^{data}(I_O) + E^{prior}(\nabla I_O) \quad (2.19)$$

However in practice for a $n \times n$ dimensional image I^c the minimization in eq. (2.19) achieves a complexity of the order $\mathcal{O}(n^4)$ since every pixel is correlated to $n^2 - 1$ other pixels. Even for medium sized images with $n = 500$ the computations involved in eq. (2.19) are practically infeasible.

To reduce the complexity we want the integrand \mathcal{E} in eq. (2.18) only to correlate the values $I_O(\mathbf{x})$ within a neighborhood $U_{\mathbf{x}_{i,j}} \subset \Omega$ with each other. One possible and very simple way to implement \mathcal{E} is to have it penalize the L_2 norm of the gradient $\nabla I_O(\mathbf{x})$

$$E_{L_2}^{prior}(\nabla I_O) = \int_{\Omega} \|\nabla I_O(\mathbf{x})\|^2 dx \quad (2.20)$$

where the gradient operation ∇ can be realized by finite differences. While the prior in eq. (2.20) can be implemented in a very efficient manner, it has an important drawback. It isotropically smooths the GRF I_O regardless of the underlying geometry of the object O being recorded. In figure 2.1a the image I^c of an object O recorded by a thermographic camera is shown. A region of interest with enhanced contrast is shown to visualize the noise corruption due to the image measuring process in eq. (2.16). Figure 2.1b shows the result of the minimization in eq. (2.19) with the L_2 prior in eq. (2.18). $E_{L_2}^{prior}$ reduces the noise in I_O but due to its isotropic nature it over-smooths the boundaries of O . In section 2.2 and following we will introduce a methodology aimed at designing prior energies E^{prior} which incorporate information about the geometry of the objects recorded in order to avoid the over-smoothing across their boundaries.

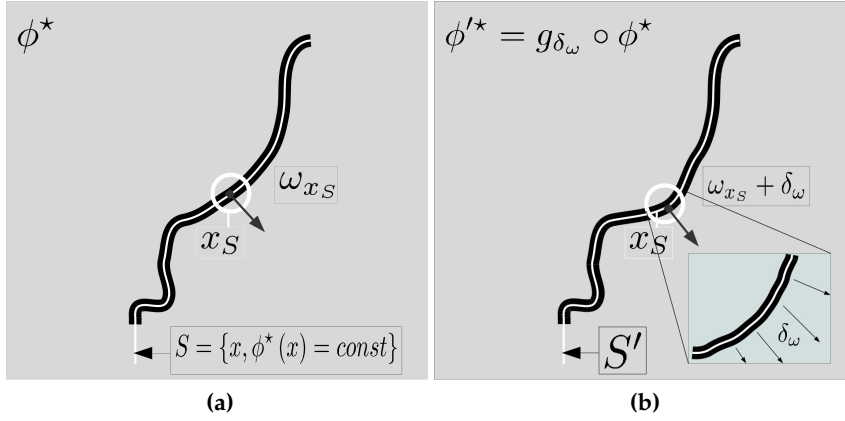


Figure 2.2.: Local transformation of an image ϕ with a level-set S . Figure 2.2a shows an image $\phi(x)$ with a line S along which the intensity values are constant. At each point x_S the vector ω_S is the normal vector on S . Figure 2.2b shows the result of the local distortion of S under the action of the operator g_{δ_ω} . g_{δ_ω} acts on S by adding to ω_S a spatial dependent vector $\delta_\omega(x)$

2.2. Lie Groups and the Noether Theorem

2.2.1. Motivation 1, the problem

In section 2.1 we had claimed that the problem with the L_2 prior

$$E_{L_2}(\phi) = \int_{\Omega} \|\nabla\phi\|^2 \quad (2.21)$$

over-smooths the GRF ϕ over the boundaries of the object recorded by the camera C . In general the minimizers ϕ^* of the energy E_{L_2} are the constant functions $\phi = \text{const}$

$$A_c = \left\{ \phi_c^* \mid \phi_c^* = \operatorname{argmin}_{\phi} (E_{L_2}(\nabla\phi)) = c, \quad c \in \mathbb{R} \right\} \quad (2.22)$$

In the context of the minimization problem in eq. (2.19) the minimizer set A_c in eq. (2.22) emphasizes that the prior E_{L_2} does not allow for the solution I_O^* (eq. (2.19)) to have discontinuities. Thus E_{L_2} is completely unaware of the geometry in the data I^c (figure 2.1a). However E_{L_2} has a advantageous property. Consider the set of rotations $SO(2)$ of the coordinate frame Ω

$$\mathbf{x}' = \mathbf{R}_{\theta}\mathbf{x}, \quad \mathbf{R}_{\theta} = \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix} \in SO(2) \quad (2.23)$$

The gradient $\nabla\phi$ transforms under the rotation in eq. (2.23) like a vector, $\nabla'\phi = \mathbf{R}_\theta \nabla\phi$ and the matrix \mathbf{R}_θ satisfies $\mathbf{R}_\theta^T \mathbf{R}_\theta = \mathbf{1}$. Thus the L_2 energy is also invariant towards the rotations in eq. (2.23)

$$E'_{L_2} = \int_{\Omega} \nabla^T \phi \mathbf{R}_\theta^T \mathbf{R}_\theta \nabla \phi d^2x = \int_{\Omega} \|\nabla\phi\|^2 d^2x \quad (2.24)$$

In general the invariance of the prior energy $E^{prior}(\nabla\phi)$ of a GRF ϕ with respect to the rotations in eq. (2.23) is a favorable feature since the gradient $\nabla\phi$ should not be penalized to a specific orientation. In the context of the minimization problem in eq. (2.19) rotational invariance of the prior $E^{prior}(\nabla I_O)$ ensures the gradient ∇I_O^* is not affected by the orientation of the camera system C .

Several methods have been introduced which allow for the construction of anisotropic priors. These methods include TV-Regularization [18, 19] which will be introduced in section 2.5, anisotropic diffusion guided by directional operators like the structure tensor [20, 21, 22] and level set methods of the Mumford-Shah type [23, 24, 25]. Among the earliest attempts for anisotropic regularization was the work of Nagel et. al. [26]. They introduced a quadratic prior

$$E_D^{prior}(\nabla\phi) = \int (\nabla\phi(\mathbf{x}))^T D(\mathbf{x}) (\nabla\phi(\mathbf{x})) d^2x \quad (2.25)$$

The operator $D(\mathbf{x})$ is a local 2×2 symmetric valued matrix estimated within a local window around each point \mathbf{x} . $D(\mathbf{x})$ is precomputed and assumed to be fixed under variation of ϕ . Thus it's eigenvectors function as a guide for the gradient ϕ . For instance in eq. (2.19) we can insert eq. (2.25) for E^{prior} . Computing D such that it has only one non-zero eigenvalue λ and an eigenvector \mathbf{b} oriented perpendicular to the weighted gradient of the data I^c

$$D(\mathbf{x}) = \lambda \mathbf{b}(\mathbf{x}) \mathbf{b}(\mathbf{x})^T, \quad \mathbf{b}(\mathbf{x}_0) \perp \langle \nabla I^c(\mathbf{x}) \rangle(\mathbf{x}_0) \quad (2.26)$$

$$\langle \nabla I^c(\mathbf{x}) \rangle(\mathbf{x}_0) = \int_A w(\|\mathbf{x} - \mathbf{x}_0\|) \nabla I^c(\mathbf{x}) d^2x \quad (2.27)$$

the prior E_D^{prior} penalizes the tangential component of ∇I_O along \mathbf{b} in the minimization in eq. (2.19). Thus the solution I_O^* can have discontinuities perpendicular to \mathbf{b} . The drawback of E_D^{prior} is that we do not know if \mathbf{b} is the true tangential vector in the unbiased projection of the object O . And since D is fixed E_D^{prior} can not be invariant under the rotations in eq. (2.23). Thus the minimization in eq. (2.19) can produce a solution I_O^* in eq. (2.19) that has discontinuities which do not reflect the true boundaries of the object O . We conclude that prior energies E_D^{prior} which are *not* rotation invariant are a source of error for the orientation of ∇I_O^* in eq. (2.19).

On the other side a potential anisotropic prior E^{prior} which is rotation invariant would lead to a solution I_O^* in eq. (2.19) for which the orientation of its structures is only determined by the data term E^{data} .

In the following we will introduce a methodology which allows us to characterize prior energies $E^{prior}(\nabla\phi)$ which allow for discontinuities in their minimizers $\phi^* = \operatorname{argmin}_\phi E^{prior}(\nabla\phi)$ while remaining invariant to a specified but more general set of spacial transformations \mathbb{G}^Ω .

2.2.2. Motivation 2, the solution

Another way to state the problem that the prior energy E_{L_2} only allows for constant minimizers $\phi^* = \text{const}$ (eq. (2.22)) goes as follows. The energy $E_{L_2}(\nabla\phi)$ is invariant upon the transformation $\phi'(\mathbf{x}) = \phi(\mathbf{x}) + d$ where d is a constant over Ω . Thus if $\phi_0^* = c'$ is a minimizer of E_{L_2} , $c' \in A_c$ then so is $\phi'^* = c' + d$ since $c' + d \in \mathbb{R}$ and by the definition of A_c in eq. (2.22) we have $\phi'^* \in A_c$. We would like to think of the operation of addition with constants d as a set \mathbb{G}_{const} of operators g_d

$$g : \mathbb{R} \rightarrow \mathbb{G}_{const}, \quad g_d = \cdot + d, \quad g_d \in \mathbb{G}_{const} \quad (2.28)$$

With the help of the construction in eq. (2.28) we can restate the invariance of E_{L_2} in the following way

$$g_d \circ E_{L_2}(\nabla\phi) = E_{L_2}(\nabla(\phi + d)) = E_{L_2}(\nabla\phi) \quad (2.29)$$

and A_c in eq. (2.22) can be viewed as being spun by one constant function $\phi_0^\phi(\mathbf{x}) = c$ and the set \mathbb{G}_{const}

$$A_c = \{\phi^* \mid \phi^* = g_d \circ \phi_0^\phi, \quad g_d \in \mathbb{G}_{const}\} \quad (2.30)$$

With the constructions in eq. (2.28) and eq. (2.30) the problem statement that the prior E_{L_2} only allows for constant minimizers is transferred to the statement that the set \mathbb{G}_{const} under which E_{L_2} is invariant is *too small* in some sense.

A more flexible prior energy E^{prior} should be invariant to a more general set of transformations \mathbb{G}^ϕ . At the same time E^{prior} should also be invariant to a spacial set of transformations \mathbb{G}^Ω in order for it not to impede the orientation of the gradient $\nabla\phi$ as motivated in section 2.2.1. Hence E^{prior} is assumed to be

invariant to the set $\mathbb{G} = \mathbb{G}^\phi \times \mathbb{G}^\Omega$ with the actions

$$g_{\omega^\phi} \circ \phi(\mathbf{x}) = \phi(\mathbf{x}) + \omega^\phi(\mathbf{x}), \quad g_{\omega^\phi} \in \mathbb{G}^\phi \quad (2.31)$$

$$g_{\omega^\Omega} \circ \mathbf{x} = \mathbf{x} + \omega^\Omega(\mathbf{x}), \quad g_{\omega^\Omega} \in \mathbb{G}^\Omega \quad (2.32)$$

The transformations in eq. (2.31) and eq. (2.32) formally capture all the possible transformations the prior energy E^{prior} is invariant to. In this sense \mathbb{G} is maximal and E^{prior} is invariant upon the action the entire set \mathbb{G}

$$g \circ E^{prior} = E^{prior}, \quad \forall g \in \mathbb{G} \quad (2.33)$$

For instance the prior energy E_{L_2} is conditionally invariant to the set $\mathbb{G} = \mathbb{G}_{const} \times SO(2)$, the set of addition of the variable ϕ with constants (eq. (2.28)) and the set of rotations in Ω (see eq. (2.23)).

Similar to the definition of A_c in eq. (2.30) we can describe the minimizers of E^{prior} as being related to each other by the elements of \mathbb{G}

$$A = \{\phi^* \mid \phi^* = g \circ \phi_0^* \quad g \in \mathbb{G}\} \quad (2.34)$$

The set \mathbb{G}^Ω contains operators which are purely geometric. The idea is to show that A may be split into sub sets $A_\Omega(\phi_c^*)$ whose elements are related to each other by the elements $g_{\omega^\Omega} \in \mathbb{G}^\Omega$

$$A_\Omega(\phi_c^*) = \{\phi^* \mid \phi^*(\mathbf{x}) = \phi_c^*(g_{\omega^\Omega} \circ \mathbf{x}), \quad g_{\omega^\Omega} \in \mathbb{G}^\Omega\} \quad (2.35)$$

$$A = \{A_\Omega(\phi_c^*) \mid \phi_c^* = g_{\omega^\phi} \circ \phi_0^*, \quad g_{\omega^\phi} \in \mathbb{G}^\phi\} \quad (2.36)$$

This is significant for the following reason: knowledge of the geometric set of transformations \mathbb{G}^Ω under which E^{prior} is invariant allows for a reduction of the set of maximizers A to a set A_{red} such that the elements $\phi_c^* \in A_{red}$ are not related to each other by \mathbb{G}^Ω

$$A_{red} = \{\phi_c^* \mid \phi_c^* = g_{\omega^\phi} \circ \phi_0^*, \quad g_{\omega^\phi} \in \mathbb{G}^\phi\} \quad (2.37)$$

$$\phi_d^*(\mathbf{x}) \neq \phi_c^*(g_{\omega^\Omega} \circ \mathbf{x}) \quad \forall g_{\omega^\Omega} \in \mathbb{G}^\Omega, \phi_{c,d}^* \in A_{red} \quad (2.38)$$

We may also turn the argument around: we could specify the geometric set of transformations \mathbb{G}^Ω and design a prior $E^{prior}(\nabla\phi)$ which is conditionally invariant under \mathbb{G}^Ω , thus having a reduced maximizer set A_{red} . To give a hint of how the prior $E^{prior}(\nabla\phi)$ could be designed we need the definition of a level-set. A level-set of an image ϕ_0^* is a sub set $S_c \subset \Omega$ defined by

$$S_c = \{\mathbf{x} \mid \phi_0^*(\mathbf{x}) = c\} \quad (2.39)$$

By the definition of the action of g_{ω^Ω} in eq. (2.32) we see that g_{ω^Ω} is a geometrical transformation that deforms the level-sets S_c (see figure 2.2). We are free to define g_{ω^ϕ} so that it is orthogonal to g_{ω^Ω} in the sense that the level-sets S_c are invariant under g_{ω^ϕ}

$$S_{c'} = g_{\omega^\phi} \circ S_c = S_c \quad (2.40)$$

since a transformation of S_c is purely geometrical. Now the level-set S_c may alternatively be defined with the help of the vector-field $\omega_\delta(\mathbf{x})$ which (see figure 2.2) is the set of vectors *tangent* to S_c

$$S_c = \{\mathbf{x} \mid \omega_\delta(\mathbf{x}) \cdot \nabla \phi_0^*(\mathbf{x}) = 0\} \quad (2.41)$$

In figure 2.2b we show an example of a level-set S which is distorted by the operator $g_{\omega_\delta} \in \mathbb{G}^\Omega$. The resulting level-set S' has the vector-field $\omega'_\delta(\mathbf{x}) = \omega_\delta(\mathbf{x}) + \delta(\mathbf{x})$ as tangent vectors.

$$S'_c = \{\mathbf{x} \mid (\omega_\delta(\mathbf{x}) + t\delta(\mathbf{x})) \cdot \nabla \phi_0^*(\mathbf{x}) = 0\} \quad (2.42)$$

However it is also possible to represent S'_c with the help of a deformation of the gradient operator ∇ itself

$$S'_c = \{\mathbf{x}' \mid \omega_\delta(\mathbf{x}') \cdot \nabla_{t\delta} \phi_0^*(\mathbf{x}') = 0\} \quad (2.43)$$

The operator $\nabla_{t\delta}$ loosely speaking encodes a reversal of the action of g_{ω^Ω} on \mathbf{x} so that S'_c can be represented with the same tangential vector-field as S_c but in the new frame $\mathbf{x}' = g_{\omega_\delta} \circ \mathbf{x}$. The operator $\nabla_{t\delta}$ is called a *push forward* of the gradient ∇ . With the help of the push forward $\nabla_{t\delta}$ it is possible to translate the notion of invariance with respect to \mathbb{G}^Ω to the requirement that $E^{prior}(\nabla_{t\delta}\phi)$ must be constant with respect to variations of the vector-field $\delta(\mathbf{x})$

$$\left. \frac{d}{dt} E^{prior}(\nabla_{t\delta}\phi) \right|_{t=0} = 0 \quad (2.44)$$

Given a specific form of the operators in \mathbb{G}^Ω , eq. (2.44) poses constraints on the form of the differential operators in the prior $E^{prior}(\nabla_{t\delta}\phi)$. Eq. (2.44) also ensures that $E^{prior}(\nabla_{t\delta}\phi)$ is indifferent to a large class of level-sets $\{S\}$, which are generated by \mathbb{G}^Ω acting on S (see eq. (2.43)).

2.3. Lie Groups

We will now give a more formal description of the transformation group $\mathbb{G} = \mathbb{G}^\phi \times \mathbb{G}^\Omega$ represented by the transformations in eq. (2.31) and eq. (2.32). Following this section we will make the invariance property of E^{prior} in eq. (2.44) formally more explicit in section 2.4. In this section the set of operators \mathbb{G} is taken to act on a vector space \mathcal{M} . The set \mathbb{G} is called a *group* if there exists an operation \cdot so that \mathbb{G} contains

- the neutral element $e \in \mathbb{G}$: $e \cdot g = g$ for all $g \in \mathbb{G}$
- the inverse $g^{-1} \in \mathbb{G}$ if $g \in \mathbb{G}$

The group \mathbb{G} is called a *Lie group* [27, 28, 29] if the group operation

$$\mathbb{G} \times \mathbb{G} \mapsto \mathbb{G} : (x, y) \rightarrow x \cdot y^{-1}$$

is smooth in both x and y . The group operation \cdot can also be used to define the *left action* l_g on \mathbb{G}

$$l_g : \mathbb{G} \rightarrow \mathbb{G} \quad l_g(h) = g \cdot h \quad g, h \in \mathbb{G} \quad (2.45)$$

l_g is a smooth isomorphism in \mathbb{G} . The elements of \mathbb{G} may themselves be smooth mappings defined on an r -dimensional space \mathcal{A}

$$g : \mathcal{A} \rightarrow \mathbb{G}, \quad (a_1, \dots, a_r) \rightarrow g_{a_1, \dots, a_r} \quad (2.46)$$

In this case we say \mathbb{G} is an r -dimensional Lie group. A classical example of a Lie group is the group of invertible n -dimensional Matrices $GL(\mathbb{R}, n)$ over the vector space $\mathcal{M} = \mathbb{R}^n$ [30]. The dimension of $GL(\mathbb{R}, n)$ is n^2 and the group operation \cdot is the matrix multiplication.

In section 2.2.2 we argument that the set \mathbb{G} acts in a two-fold manner on the functions $\phi(\mathbf{x}) \in \mathcal{C}^\infty(\Omega)$, namely by acting on the spacial coordinates $\mathbf{x} \in \Omega$ in eq. (2.32) and on the function values ϕ them selves in eq. (2.31). The spaces Ω and $\mathcal{C}^\infty(\Omega)$ are both vector spaces, that is the addition operation $+$ and multiplication with a factor $\lambda \in \mathbb{R}$ are defined in both spaces. It is thus natural to combine both Ω and $\mathcal{C}^\infty(\Omega)$ to one single vector space $\mathcal{M} = \Omega \times \mathcal{C}^\infty(\Omega)$. However since the functions $\phi(\mathbf{x})$ are unknown and we would also like to place constraints on their derivatives $\phi_{,K}$ (K is a multi-index, e.g. $K = [xxx, yy]$), we combine Ω together with the *Jet space* $J^k(\mathcal{C}^\infty(\Omega))$ [28], $\mathcal{M} = \Omega \times J^k(\mathcal{C}^\infty(\Omega))$. $J^k(\mathcal{C}^\infty(\Omega))$ is the set of smooth differentiable functions and their derivatives up to order k with compact support in Ω . The points $z \in \mathcal{M}$ are vectors of the

independent variables \mathbf{x} , the dependent variable $\phi(\mathbf{x})$ and its derivatives $\phi_{,K}$

$$\mathbf{z} = (\mathbf{x}, \phi(\mathbf{x}), \phi_{,K}(\mathbf{x})) \quad (2.47)$$

For this work we will focus only on first order derivatives, $k = 1$ so that the vectors \mathbf{z} have the form

$$\mathbf{z} = (\mathbf{x}, \phi(\mathbf{x}), \nabla\phi(\mathbf{x})) \quad (2.48)$$

The action of \mathbb{G} on \mathcal{M} is straightforward

$$\tilde{\mathbf{z}} = (\tilde{\mathbf{x}}, \tilde{\phi}(\tilde{\mathbf{x}}), \tilde{\nabla}\tilde{\phi}(\tilde{\mathbf{x}})) \quad (2.49)$$

$$\tilde{\mathbf{x}} = g_{a_1 \dots a_r} \circ \mathbf{x} \quad (2.50)$$

$$\tilde{\phi} = g_{a_1 \dots a_r} \circ \phi \quad (2.51)$$

$$\tilde{\nabla} = J^{-T} \nabla, \quad J_{\mu\nu} = \frac{d\tilde{x}_\mu}{dx_\nu} \quad (2.52)$$

Since the elements $g_{a_1 \dots a_r}$ are continuous in the parameters a_i we are free define to a smooth path γ in the parameter space \mathcal{A}

$$\gamma : t \rightarrow (a_1(t) \dots a_r(t)) \quad (2.53)$$

$$g_{\gamma(0)} = e \quad (2.54)$$

The derivative of $g_{\gamma(t)}$ with respect to t at $t = 0$ is an element of the tangential space of \mathbb{G} at the neutral element $e \in \mathbb{G}$, $T_e\mathbb{G}$

$$\left. \frac{d}{dt} g_{\gamma(t)} \right|_{t=0} = X_e \in T_e\mathbb{G} \quad (2.55)$$

The subscript on the vector X_e denotes that it belongs to $T_e\mathbb{G}$. The coordinates of X_e relative to the space \mathcal{M} can be computed when we look at the derivative of the induced action of $g_{\gamma(t)}$ on the space of smooth functions with support on \mathcal{M} , $\mathcal{F}(\mathcal{M})$. The action of X_e on $\mathcal{F}(\mathcal{M})$ can be computed by evaluating $F \in \mathcal{F}(\mathcal{M})$ on the transformed vector $\tilde{\mathbf{z}} = g_{\gamma(t)} \circ \mathbf{z}$ and the taking the derivative with respect to t at the neutral element e

$$X_e F(\mathbf{z}) = \left. \frac{d}{dt} F(\tilde{\mathbf{z}}) \right|_{t=0} = \sum_{i=1}^r (\omega_i^\mu \partial_\mu F(\mathbf{z}) + \omega_i^\phi \frac{d}{d\phi} F(\mathbf{z}) + D\phi_i^\nu \frac{d}{d\partial_\nu \phi} F(\mathbf{z})) \alpha_i \quad (2.56)$$

where we have

$$\omega_i^\mu(\mathbf{x}) = \left. \frac{d\tilde{x}_\mu}{da_i} \right|_{t=0} \quad \omega_i^\phi(\mathbf{x}, \phi) = \left. \frac{d\tilde{\phi}}{da_i} \right|_{t=0} \quad \alpha_i = \left. \frac{da_i}{dt} \right|_{t=0} \quad (2.57)$$

$$D\phi_i^\nu = \frac{d\omega_i^\nu}{dx_\nu} - \sum_\mu \frac{d\omega_i^\nu}{dx_\mu} \partial_\mu \phi \quad (2.58)$$

The function $D\phi_i^\nu$ is called the prolonged action of $g_{\gamma(t)}$ on the gradient operator ∇ [28, 31] (see appendix A.1 for the derivation). Notice that while ω_i^μ and ω_i^ϕ are functions defined on \mathcal{M} , the coefficients α_i are independent of \mathcal{M} . They are the components of the vector X_e with respect to the r basis operators

$$X_{e,i} = X_{e,i}^\Omega + \omega_i^\phi \frac{d}{d\phi} + D\phi_i^\nu \frac{d}{d\partial_\nu \phi}, \quad X_{e,i}^\Omega = \omega_i^\mu \partial_\mu \quad (2.59)$$

so that X_e has the operator form

$$X_e = \sum_i \alpha_i X_{e,i} \quad (2.60)$$

An important point about X_e is that it is an *operator valued* function over Ω since the coefficients ω_i^μ and ω_i^ϕ in eq. (2.57) are functions over Ω . We will refer to X_e as a vector at the unit element e , to ω_i^μ as a vector valued function (VVF) and to ω_i^Ω as a scalar valued function (SVF). The vector X_e only exists in the tangential space at $e \in \mathbb{G}$, $X_e \in T_e \mathbb{G}$. However it is possible to construct a vector Y_h at a location $h \in \mathbb{G}$ by relating it to X_e with a map l_{h^*} called the *push-forward*

$$Y_h F(\mathbf{z}) = (l_{h^*} X_e) F(\mathbf{z}) = \left. \frac{d}{dt} F(l_h(g_{\gamma(t)}) \circ \mathbf{z}) \right|_{t=0} \quad (2.61)$$

The vector X_e operates on the function F in eq. (2.56) as a differential operator at the point $e \circ \mathbf{z} = \mathbf{z}$. The effect of l_{h^*} is that it transports the vector X_e to the vector Y_h which operates on F at the point $l_h(e) \circ \mathbf{z} = h \circ \mathbf{z}$. As Y_h is a smooth function with respect to h which is defined everywhere in \mathbb{G} it is called a *vector field over* \mathbb{G} . This means the coordinates of the vector field Y_h are the operators $h \in \mathbb{G}$ and *not* the points $\mathbf{z} \in \mathcal{M}$. The set of vector fields is the union of all the tangential spaces over \mathbb{G}

$$T\mathbb{G} = \bigcup_{h \in \mathbb{G}} T_h \mathbb{G} \quad (2.62)$$

Similar to X_e in eq. (2.60) the vector Y_h has a coordinate representation with respect to the tangential space $T_h\mathbb{G}$

$$Y_h F(\mathbf{z}) = \sum_i \alpha'_i Y_{h,i} \quad (2.63)$$

$$Y_{h,i} = \omega'_i{}^\mu \partial_\mu + \omega'_i{}^\phi \frac{d}{d\phi} + D'\phi_i{}^\nu \frac{d}{d\partial_\nu\phi} \quad (2.64)$$

There exists a unique sub set $\mathcal{G} \subset T\mathbb{G}$ called the *Lie algebra*. It defined as the set of all vector fields $X_h \in T\mathbb{G}$ which are invariant under the left action l_g for any $g \in \mathbb{G}$

$$l_{g^*} X_h = X_{g \cdot h} = \sum_i \alpha_i X_{g \cdot h}^i \quad \forall g \in \mathbb{G}, X_h \in \mathcal{G} \quad (2.65)$$

From eq. (2.65) we see that a consequence of left invariance is that the coordinate vector α is constant under the transformation l_g . This is what is referred to as the *parallel transport* of α along the transformation l_g . The Lie algebra \mathcal{G} has the property that it is closed under the antisymmetric commutator $[\cdot, \cdot]$

$$[Y_h, X_h] = Z_h \in \mathcal{G} \quad \forall X_h, Y_h \in \mathcal{G} \quad (2.66)$$

Eq. (2.66) also implies that the commutator $[Y_h, X_h]$ is also left invariant [27]. If we consider the basis of the Lie algebra \mathcal{G} , the r vector fields X_h^1, \dots, X_h^r then the closed-ness of the commutator $[\cdot, \cdot]$ implied by eq. (2.66) translates to the fact that the commutator $[\cdot, \cdot]$ of two basis elements X_h^i and X_h^j is a linear combination of the basis X_h^1, \dots, X_h^r itself

$$[X_h^i, X_h^j] = \sum_{l=1}^r C_{i,j}^l X_h^l \quad (2.67)$$

The constants $C_{i,j}^l$ are called the *structure constants* of the Lie algebra \mathcal{G} . The commutator in eq. (2.66) has a geometric meaning. Suppose Y_e is the vector in the tangent space T_e of the one parameter group g_t^Y in the sense of eq. (2.55). It is easy to show (see appendix A.2) that the rate of change of the vector field $X_{g_t^Y}$ at the unit element e is equal to the commutator between X_e and Y_e

$$\left. \frac{d}{dt} X_{g_t^Y} \right|_{t=0} = [Y_e, X_e] \quad (2.68)$$

Since $[Y_h, X_h]$ is left invariant, eq. (2.68) may be translated to any point $g \in \mathbb{G}$

$$h \circ \left. \frac{d}{dt} X_{g_t^Y} \right|_{t=0} = \left. \frac{d}{dt} X_{h \cdot g_t^Y} \right|_{t=0} = [Y_h, X_h] \quad (2.69)$$

Essentially eq. (2.69) tells us that once we can tell how the vector field X_g changes along any path g_t^Y which goes through unity, $g_t^Y|_{t=0} = e$, we can compute its rate of change along any *other* path in \mathbb{G} .

2.3.1. The Group $\mathbb{G} = \mathbb{T} \times SO(2)$

The group $\mathbb{G} = \mathbb{T} \times SO(2)$ is the group of translations and rotations in the plane \mathbb{R}^2 . Its algebra is the algebra $\mathcal{G} = \mathfrak{t} \times \mathfrak{so}(2)$ which has the basis $\{X_e^{\Omega,x}, X_e^{\Omega,y}, X_e^{\Omega,\theta}\}$. The subset $\mathfrak{t} = \{X_e^{\Omega,x}, X_e^{\Omega,y}\}$ is the set generators of infinitesimal translations

$$X_e^{\Omega,x} = \partial_x, \quad X_e^{\Omega,y} = \partial_y \quad (2.70)$$

\mathfrak{t} is a commutative basis since $[\partial_x, \partial_y] = 0$. The basis for $\mathfrak{so}(2)$ is the single operator $X_e^{\Omega,\theta}$ which is the generator of infinitesimal rotations. With respect to the Cartesian coordinate frame ∂_θ it has the following representation

$$X_e^{\Omega,\theta} = -y\partial_x + x\partial_y \quad (2.71)$$

From eq. (2.71) we can see that ∂_θ does not commute with \mathfrak{t} and the commutators for the basis $\{X_e^{\Omega,x}, X_e^{\Omega,y}, X_e^{\Omega,\theta}\}$ are easily computed

$$[X_e^{\Omega,\theta}, X_e^{\Omega,x}] = -X_e^{\Omega,y} \quad [X_e^{\Omega,\theta}, X_e^{\Omega,y}] = X_e^{\Omega,x} \quad [X_e^{\Omega,x}, X_e^{\Omega,y}] = 0 \quad (2.72)$$

We note that the group $SO(2)$ generates the unit circle S^1 by rotating the point $\mathbf{x}_0 = (x, y)$

$$\mathbf{x}(\theta) = g_\theta \circ \mathbf{x}_0 = R_\theta \mathbf{x}_0, \quad R_\theta = \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix} \quad (2.73)$$

The meaning of the first two commutators in eq. (2.72) is that the gradient operator ∇ is rotated by 90° counter clockwise under the action of $X_e^{\Omega,\theta}$

$$\left. \frac{d}{d\theta} \nabla_{\mathbf{x}(\theta)} \right|_{\theta=0} = [X_e^{\Omega,\theta}, \nabla] = \mathbf{M}_\theta \cdot \nabla, \quad \mathbf{M}_\theta = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \quad (2.74)$$

The matrix \mathbf{M}_θ is one of the Pauli matrices [32]. The Pauli matrices are the basis for the Lorentz group of special relativity which is an important symmetry for many quantum field theories for instance quantum electrodynamics [32, 9, 33].

2.4. Noether's First Theorem

Embedding Geometrical Constraints into Prior Energies

In section 2.2.1 argued that in order for a prior $E^{prior}(\nabla\phi)$ needs to be invariant to a large group of transformations \mathbb{G} in order for it's minimizers

$$A = \left\{ \phi^* \mid \phi^* = \operatorname{argmin}_{\phi} \left(-E^{prior}(\nabla\phi) \right) \right\} \quad (2.75)$$

to be non trivial, that is $\phi^* \neq \text{const.}$ Invariance of E^{prior} was linked to the requirement that the minimizer set A in eq. (2.75) be generated by the group \mathbb{G}

$$A = \left\{ \phi^* \mid \phi^* = g_{\omega} \circ \phi_0^* \quad g_{\omega} \in \mathbb{G} \right\} \quad (2.76)$$

In eq. (2.36) we explained that a transformation $g_{\omega} \in \mathbb{G}$ may partition the set A into subsets A_{Ω} whose elements are related to each other through geometrical transformations $g_{\omega\Omega}$ on the coordinate frame Ω . We motivated the introduction of deformations to the gradient operator ∇ such that the level-sets S' in eq. (2.43) have the same functional form in the transformed coordinates $\mathbf{x}' = g_{\omega\Omega} \circ \mathbf{x}$ as in the original coordinates (see eq. (2.41)). With the help of the machinery introduced in section 2.3 we can express a level-set S_V of $\phi^* \in \mathcal{A}$ in terms of a left invariant vector field V_h operating on ϕ^* at the identity $e \in \mathbb{G}$

$$S_V = \left\{ \mathbf{x} \mid V_e^{\Omega} \phi^*(\mathbf{x}) = 0 \right\} \quad V_e^{\Omega} = \sum_i \alpha_i X_e^{\Omega,i} \quad X_e^{\Omega,i} \in \mathbf{X}_e^{\Omega} \quad (2.77)$$

for a particular sub algebra $\mathbf{X}_e^{\Omega} \subset \mathcal{G}$. Through eq. (2.77) the vector field V_h defines the particular geometry of S_V . For instance if the set of differential operators $\{X_e^{\Omega,i}\}$ are the Cartesian operators $\{\partial_x, \partial_y\}$ then the corresponding level-sets S_v^{lin} are the straight lines with tangential vector \mathbf{v}

$$S_v^{lin} = \left\{ \mathbf{x} \mid \mathbf{v}^T \cdot \nabla \phi^*(\mathbf{x}), \mathbf{x} \in \Omega \right\} = 0 \quad (2.78)$$

The tangential vector \mathbf{v} is a constant function in Ω so that the level-set $S_v^{lin} \subset \Omega$ is a line oriented in the direction \mathbf{v} which is invariant upon the transformation $\mathbf{x} \rightarrow \mathbf{x} + \tau \mathbf{v}$. Different priors $E^{prior}(\nabla\phi)$ may be constructed for which the minimizers $\phi^* \in \mathcal{A}$ have linear level-sets S_v^{lin} of all orientations \mathbf{v} . For instance in section 2.5 we will introduce the total variation prior E_{TV}^{prior} and in section 3 the structure tensor based prior E_{ST}^{prior} . Both priors penalize level-sets with non-zero curvature so that their minimizers ϕ^* only contain linear level-sets S_v^{lin} . However the orientation vector \mathbf{v} is not constrained.

A generalization to priors $E^{prior}(\mathbf{X}_e^\Omega \phi)$ whose minimizers ϕ^* have level-sets S_V of the generic form in eq. (2.77) is possible if the sub algebra $\mathbf{X}_e^\Omega \subset \mathcal{G}$ is a commutative algebra. In [34, 29] it was shown that only for pair-wise commuting differential operators $X_e^{\Omega,1}, X_e^{\Omega,2} \in \mathbf{X}_e^\Omega$, $[X_e^{\Omega,1}, X_e^{\Omega,2}] = 0$ may there exist images ϕ^* which have level-sets S_V of the form in eq. (2.77). Furthermore for the commuting operators $X_e^{\Omega,1}, X_e^{\Omega,2} \in \mathbf{X}_e^\Omega$ there exist pairs of functions $\epsilon_1(\mathbf{x})$ and $\epsilon_2(\mathbf{x})$ which obey the conditions

$$\begin{aligned} X_e^{\Omega,1} \epsilon_1(\mathbf{x}) &= 1, & X_e^{\Omega,1} \epsilon_2(\mathbf{x}) &= 0 \\ X_e^{\Omega,2} \epsilon_1(\mathbf{x}) &= 0, & X_e^{\Omega,2} \epsilon_2(\mathbf{x}) &= 1 \end{aligned} \quad (2.79)$$

The function $\epsilon(\mathbf{x}) = (\epsilon_1(\mathbf{x}), \epsilon_2(\mathbf{x}))$ represents a coordinate transformation from the Cartesian frame Ω to the curvi-linear frame Ω^ϵ with the coordinates (ϵ_1, ϵ_2)

$$\epsilon(\mathbf{x}) : \Omega \rightarrow \Omega^\epsilon \quad (2.80)$$

Within the coordinate frame Ω^ϵ the commuting operators $X_e^{\Omega,1}, X_e^{\Omega,2} \in \mathbf{X}_e^\Omega$ take on the form $X_e^{\Omega,1/2} = \partial_{\epsilon_{1/2}}$ and the level-sets S in eq. (2.77) are represented by the inverse transformation $\mathbf{x}(\epsilon)$

$$S_\alpha^\epsilon = \left\{ \mathbf{x}(\epsilon) \left| \sum_i \alpha_i \partial_{\epsilon_i} \phi(\epsilon_1, \epsilon_2) = 0 \right. \right\} \quad (2.81)$$

By comparison of S_α^ϵ in eq. (2.81) with S_v^{lin} in eq. (2.78) we can see that S_α^ϵ is a line in the domain Ω^ϵ with orientation α .

The availability of the transform $\epsilon(\mathbf{x})$ and its inverse is guaranteed if the differential operators $X_e^{\Omega,1}, X_e^{\Omega,2} \in \mathcal{G}$ commute. Conversely if a transform $\tilde{\epsilon}(\mathbf{x})$ is available for which the Cauchy-Riemann equations

$$\partial_x \tilde{\epsilon}_1 = -\partial_y \tilde{\epsilon}_2, \quad \partial_x \tilde{\epsilon}_2 = \partial_y \tilde{\epsilon}_1 \quad (2.82)$$

then differential operators $\tilde{X}_e^{\Omega,1}, \tilde{X}_e^{\Omega,2} \in \mathbf{X}_e^\Omega$ can be computed and for which the level-sets S_α^ϵ are linear in the domain Ω^ϵ . Hence the specification of a transform $\epsilon(\mathbf{x})$ which obeys eq. (2.82) is sufficient to construct a prior $E^{prior}(\mathbf{X}_e^\Omega \phi)$ for which the minimizers ϕ^* have level-sets S_α^ϵ which are linear in the domain Ω^ϵ and we will call the domain Ω^ϵ the linear domain of $E^{prior}(\mathbf{X}_e^\Omega \phi)$.

Now we have shown that a prior $E^{prior}(\mathbf{X}_e^\Omega \phi)$ may be constructed using the operators of a commutative sub algebra $\mathbf{X}_e^\Omega \subset \mathcal{G}$. However in eq. (2.76) we explained that the minimizer set \mathcal{A} can be described as being spun up by the entire group G . Thus the minimizer set \mathcal{A} itself is by definition invariant under

the action of \mathbb{G}

$$h \circ \mathcal{A} = \mathcal{A} \quad \forall h \in \mathbb{G} \quad (2.83)$$

Eq. (2.83) is equivalent to the restriction that the prior energy E^{prior} itself must be invariant under the action of \mathbb{G}

$$E^{prior} \left(\mathcal{X}_{g_{\gamma(t)}}^{\Omega} \phi \right) = const, \quad \text{with respect to } t \quad (2.84)$$

for any one-parameter sub group $g_{\gamma(t)} \in \mathbb{G}$. Eq. (2.84) poses a constraint on how the prior E^{prior} is constructed. For instance the energy $E^{prior}(\nabla) = \int (\partial_x \phi)^2 d^2x$ is not invariant under the rotation group $SO(2)$, and does not obey eq. (2.84) if \mathbb{G} contains $SO(2)$ as a subgroup. Hence its minimizers ϕ^* have level-sets oriented only in x -direction rendering it useless for anisotropic regularization.

The following section is devoted to the consequences of eq. (2.84) for models

$$E \left(\phi, \mathbf{X}_e^{\Omega} \phi \right) = E^{data}(\phi) + E^{prior} \left(\mathbf{X}_e^{\Omega} \phi \right) \quad (2.85)$$

For which the prior E^{prior} obeys the invariance relation in eq. (2.84) for a general Lie group \mathbb{G} .

2.4.1. Noethers Theorems

In her original paper [35, 36] Emmy Noether handles the question: Given a model of a physical system, encoded in an action

$$E = \int_{\Omega} (\mathcal{E}(\mathbf{x}, \{\phi_{\rho}\}, \{\nabla_K \phi_{\rho}\})) d^n x \quad (2.86)$$

which depends on ρ fields $\phi_1 \dots \phi_{\rho}$ and their derivatives to order K , and knowledge of a set of smooth transformations \mathbb{G} under which the action E is invariant

$$E' = g_{\gamma} \circ E = E \quad \forall g_{\gamma} \in \mathbb{G} \quad (2.87)$$

what are the special properties hidden in the model that invoke the symmetry?

To answer this question she deals with two cases:

- Finite dimensional Lie groups \mathbb{G} , which we introduced in section 2.3. For now it is sufficient to think of \mathbb{G} as the set of smooth functions g_{γ} defined on an r dimensional space, $\gamma = (\alpha_1, \dots, \alpha_r)$.

- Infinite dimensional Lie groups \mathbb{G}_∞ , which are generalizations of the finite dimensional groups in the sense that the r parameters $\alpha_1, \dots, \alpha_r$ are functions over the Cartesian coordinate frame Ω . We will not handle this case.

In the case of the finite dimensional group Emmy Noether took g_ω to be the smooth *infinitesimal* transformation, encoding both variations of the fields and of the coordinates

$$\phi'_\rho(\mathbf{x}) = \phi_\rho(\mathbf{x}) + \sum_{m=1}^r \alpha_m \omega_m^{\phi_\rho}(\mathbf{x}) \quad \mathbf{x}' = \mathbf{x} + \sum_{m=1}^r \alpha_m \omega_m^\Omega(\mathbf{x}) \quad (2.88)$$

The functions $\omega_m^{\phi_\rho}$ and ω_m^Ω can be seen as a basis for ω^{ϕ_ρ} and ω^Ω in eqs. ?? and ?. She proved that if the action E is invariant under g_ω eq. (2.87), then there exists r vectors \mathbf{W}_m such the integral relationship

$$E - E' = \int_\Omega \sum_{m=1}^r \alpha_m \left[\sum_\rho \bar{\omega}_m^{\phi_\rho} [\mathcal{E}]_\rho + \text{div}(\mathbf{W}_m) \right] = 0 \quad (2.89)$$

$$\bar{\omega}_m^{\phi_\rho} = \left(\omega_m^{\phi_\rho} - \omega_m^{\mu\Omega} \partial_\mu \phi_\rho \right), \quad [\mathcal{E}]_\rho = \frac{\delta \mathcal{E}}{\delta \phi_\rho} - \frac{d}{dx^\mu} \left(\frac{\delta \mathcal{E}}{\delta \phi_{\rho,\mu}} \right) \quad (2.90)$$

where $[\mathcal{E}]_\rho$ are the Euler-Lagrange differentials of the fields ϕ_ρ and the divergences $\text{div}(\mathbf{W}_m)$ appear by carefully collecting all terms which occur as a result of the integral product rule

$$\int f \cdot \partial_\mu g d^n x = \int \partial_\mu (f \cdot g) d^n x - \int \partial_\mu f \cdot g d^n x \quad (2.91)$$

when computing the symbolic form of $[\mathcal{E}]_\rho$. The main result is the argument that since the α_m , $\omega_m^{\phi_\rho}$ and the ω_m^μ are assumed to linearly independent, the r equations

$$\sum_\rho \bar{\omega}_m^{\phi_\rho} [\mathcal{E}]_\rho + \text{div}(\mathbf{W}_m) = 0 \quad m = 1, \dots, r \quad (2.92)$$

relate r of the ρ Euler-Lagrange equations $[\mathcal{E}]_\rho$ so that the physical system only has $\rho - r$ independent Euler-Lagrange equations $[\mathcal{E}]_\rho$ and thus only $\rho - r$ independent fields ϕ_ρ . In the case $\rho \leq r$ the system of equations in eq. (2.92) is overdetermined, eq. (2.89) can only hold if *all* the divergences and *all* the Euler-Lagrange equations vanish

$$[\mathcal{E}]_\rho(\phi_1^*, \dots, \phi_\rho^*) = 0, \quad \text{div}(\mathbf{W}_m)(\phi_1^*, \dots, \phi_\rho^*) = 0, \quad \rho \leq r \quad (2.93)$$

Eq. (2.93) implies that only at the minima of the fields, ϕ_ρ^* the r vectors \mathbf{W}_m are conserved.

Kepler's Two Body Problem

Kepler's two body problem is the problem of calculating the problem of estimating the trajectory of a body of mass m_e (the earth) which is moving within the vicinity of another body with mass m_s (the sun). According to Newton there exists a gravitational force between the masses coming from the energy $V(r)$ of the gravitational field surrounding the mass m_s at the origin in \mathbb{R}^3

$$V(\mathbf{r}_e(t)) = -\frac{m_e \cdot m_s}{r} \quad r = \|\mathbf{r}_e - \mathbf{r}_s\| \quad (2.94)$$

The kinetic energy of the mass m_e is $\frac{1}{2}m_e\dot{r}^2$ so that the Lagrangian of the path $\mathbf{r}_e(t)$ is

$$L(\mathbf{r}_e(t)) = \frac{1}{2}m_e\dot{r}_e^2 + \frac{1}{2}m_e\dot{r}_s^2 - V(\mathbf{r}_e(t)) \quad (2.95)$$

The Euler-Lagrange equations are easily computed

$$\ddot{r}_e + \frac{m_s + m_e}{r^2} = 0 \quad (2.96)$$

The parameter t is the time parameter of the two body system. The Kepler Lagrangian in eq. (2.95) exhibits a symmetry under four different one parameter Lie group actions, namely the action of time shift and rotations around the three spacial axis (the group $SO(3) \times \mathbb{R}$)

$$t' = t + \delta t \quad (2.97)$$

$$\mathbf{r}' = \mathbf{r} + \partial_{\theta_i} \mathbf{r}' \delta \theta_i \quad i = x, y \text{ or } z \quad (2.98)$$

where θ_i are rotation around the x -, y - or z -axis. From Noether's theorem there exist four corresponding conserved quantities:

$$W_t = \mathcal{H} = \frac{1}{2}m_e\dot{r}^2 + V(\mathbf{r}_e(t)) \quad \text{time shift} \quad (2.99)$$

$$W_x = z\dot{y} - y\dot{z} \quad \text{Rotation around } x\text{-axis} \quad (2.100)$$

$$W_y = z\dot{x} - x\dot{z} \quad \text{Rotation around } y\text{-axis} \quad (2.101)$$

$$W_z = x\dot{y} - y\dot{x} \quad \text{Rotation around } z\text{-axis} \quad (2.102)$$

The conserved quantity \mathcal{H} in eq. (2.99) is the *Hamiltonian Energy* of the two body system. It constant time and thus manifests that the total energy of the two body system does not dissipate away since there are no external forces interacting with the two masses m_e and m_s , that is the two body system is a *closed system*. The vector $\mathbf{W} = (W_x, W_y, W_z)$ (Eqs. eq. (2.100) to eq. (2.102)) is the total *angular momentum* the masses m_e and m_s have as they rotate around each other. The solutions to the Euler-Lagrange equations in eq. (2.96) are elliptic curves in the surface $S_{\mathbf{W}}$ orthogonal to \mathbf{W} . The constancy of \mathbf{W} with respect to the special orthogonal group $SO(3)$ comes the fact that $S_{\mathbf{W}}$ is actually a flat Euclidean plane embedded in a 3-dimensional Euclidean space.

2.4.2. Noether's First Theorem: A Modern Version

In this section we explicitly derive Noether's first theorem for models with one field ϕ and its first derivatives $\mathbf{X}_e^\Omega \phi$ using the Lie algebra introduced in section 2.3. We consider the total energy

$$E = \int_{\Omega} \mathcal{E}(\phi, \mathbf{X}_e^\Omega \phi) \mathcal{N}(\mathbf{x}) d^2x = \int_{\Omega} \mathcal{E}^{data}(\phi) \mathcal{N}(\mathbf{x}) d^2x + \int_{\Omega} \mathcal{E}^{prior}(\mathbf{X}_e^\Omega \phi) \mathcal{N}(\mathbf{x}) d^2x \quad (2.103)$$

The explanation of the constant $\mathcal{N}(\mathbf{x})$ will shortly follow. We apply a one parameter group $g_{\gamma(t)}$ to E^{prior} and according to eq. (2.55) we can compute the vector V_e in the tangent space of $g_{\gamma(t)}$ at $t = 0$.

$$\left. \frac{d}{dt} g_{\gamma t} \circ E \right|_{t=0} = \int_{\Omega} \left(\sum_{i=1}^q P_i [V_e^\Omega, X_e^{i,\Omega}] \phi + v^\phi [\mathcal{E}] \right) \mathcal{N}(\mathbf{x}) d^2x \quad (2.104)$$

$$P_i = \frac{\delta \mathcal{E}}{\delta X_e^{i,\Omega} \phi}, \quad [\mathcal{E}] = \frac{\delta \mathcal{E}}{\delta \phi} - \sum_i \frac{d}{dx^\mu} (w_i^\mu P_i) \quad (2.105)$$

For the derivation of eq. (2.104) see eq. (A.47) in appendix A.3. Eq. (2.104) will play an important role in section 4 which is why we will discuss it in detail. The differentials $[\mathcal{E}]$ are called the Euler-Lagrange differentials [36, 28] and the vector \mathbf{P} is called the canonical momentum [37]. Eq. (2.104) is the most general form of variation. It contains two components, namely one component proportional to intensity variations of the field ϕ , v^ϕ and one component proportional to variations of the coordinate frame Ω encoded in $V_e^\Omega = v^\mu \partial_\mu$. Just as in eq. (2.92) the variations v^ϕ and v^μ are taken to be independent from each other. This means that the Lie group \mathbb{G} splits into two *independent* sub groups \mathbb{G}^ϕ and \mathbb{G}^Ω and we can write \mathbb{G} as the composition $\mathbb{G} = \mathbb{G}^\phi \times \mathbb{G}^\Omega$. The integral volume d^2x also

transforms under the action of g_t^V

$$\frac{d}{dt} \left(g_t^V \circ d^2x \right) \Big|_{t=0} = \frac{dv^\mu}{dx^\mu} d^2x \quad (2.106)$$

However we are only interested in the Euler-Lagrange differentials $[\mathcal{E}]$ and the canonical momentum vector \mathbf{P} since only they depend on the particular Lagrangian \mathcal{E} . Thus the normalization factor $\mathcal{N}(\mathbf{x})$ in the integral in eq. (2.103) is chosen such that the volume element $\mathcal{N}(\mathbf{x})d^2x$ is invariant under the transformation $g_{\gamma(t)}$

$$\frac{d}{dt} \left(g_t^V \circ \mathcal{N}(\mathbf{x})d^2x \right) \Big|_{t=0} = \left(\mathcal{N}(\mathbf{x}) \frac{dv^\mu}{dx^\mu} + \frac{d\mathcal{N}(\mathbf{x})}{dt} \Big|_{t=0} \right) d^2x = 0 \quad (2.107)$$

In the following we will drop the normalization $\mathcal{N}(\mathbf{x})$ and assume d^2x to invariant under any transformation g_t^V . Since V_e is an element of the Lie algebra \mathcal{G} we can expand it in terms of the r basis elements X_e^i

$$V_e = \sum_{i=1}^r \alpha_i X_e^i \quad (2.108)$$

Under the expansion in eq. (2.108) the eq. (2.104) becomes

$$\frac{d}{dt} g_{\gamma_t} \circ E \Big|_{t=0} = \sum_{m=1}^r \alpha_m \int_{\Omega} \left(\widetilde{\mathbf{B}}_m \phi + \omega_m^\phi [\mathcal{E}] \right) d^2x \quad (2.109)$$

where the operator $\widetilde{\mathbf{B}}_m$ is the scalar product

$$\widetilde{\mathbf{B}}_m = \sum_{i=1}^q P_i \left[X_e^{m,\Omega}, X_e^{i,\Omega} \right] \quad (2.110)$$

The operator $\widetilde{\mathbf{B}}_m$ in eq. (2.110) will play an important role in section 4 where we propose an extension to the principle of least action. In appendix A.3 we show how eq. (2.109) can be transformed in to the original version in eq. (2.89)

$$\frac{d}{dt} g_{\gamma_t} \circ E \Big|_{t=0} = \int_{\Omega} \sum_m \alpha_m \left(\text{div}(\mathbf{W}_m) + \widetilde{\omega}_m^\phi [\mathcal{E}] \right) d^2x \quad (2.111)$$

$$\widetilde{\omega}_m^\phi = \omega_m^\phi - X_e^{m,\Omega}(\phi) \quad (2.112)$$

with the vector valued functions (VVF) \mathbf{W}_m

$$\mathbf{W}_m^\mu = \omega_m^\mu \mathcal{E} + \sum_i \omega_i^\mu \left(\omega_m^\phi - X_e^m(\phi) \right) P_i \quad (2.113)$$

Thereby we prove the following relation

$$\operatorname{div}(\mathbf{W}_m) - X_{e,m}^\Omega(\phi)[\mathcal{E}] = \widetilde{\mathbf{B}}_m\phi \quad (2.114)$$

If the energy E is assumed to be invariant with respect to any one parameter group $g_t^V \subset \mathbb{G}$

$$\left. \frac{d}{dt} g_{\gamma t}^V \circ E \right|_{t=0} = 0 \quad (2.115)$$

then by the argumentation in section 2.4.1 the divergences of the vectors \mathbf{W}_m in eq. (2.113) and the Euler-Lagrange differentials must vanish

$$[\mathcal{E}](\phi^*) = 0, \quad \operatorname{div}\mathbf{W}_m = 0 \quad \forall 1 \leq m \leq r \quad (2.116)$$

and by eq. (2.114) the equations in eq. (2.116) imply

$$\widetilde{\mathbf{B}}_m\phi^* = 0 \quad (2.117)$$

2.4.3. Pure Spatial Symmetries

Since we assumed the invariance of the energy E under the group $\mathbb{G} = \mathbb{G}^\phi \times \mathbb{G}^\Omega$ in eq. (2.115) as well as the independence of the sub groups \mathbb{G}^ϕ and \mathbb{G}^Ω it follows that E must also be invariant with respect to pure spacial one dimensional sub groups $g_t^{V^\Omega} \subset \mathbb{G}^\Omega$

$$\left. \frac{d}{dt} (g_t^{V^\Omega} \circ E) \right|_{t=0} = 0 \quad (2.118)$$

The Lie algebra element V_e^Ω corresponding to $g_t^{V^\Omega}$ does not contain any variations to the field ϕ thus we can obtain an expression for eq. (2.118) simply by setting $v^\phi = 0$ and $\omega_i^\phi = 0$ in eqs. (2.104) and (2.109)

$$\left. \frac{d}{dt} (g_t^{V^\Omega} \circ E) \right|_{t=0} = \int_\Omega \sum_{m=1}^r \alpha_m \widetilde{\mathbf{B}}_m\phi \, d^2x = 0 \quad (2.119)$$

It follows that if eq. (2.119) holds for any one parameter sub group $g_t^{V^\Omega} \subset \mathbb{G}$ (any coefficient vector α) then

$$\widetilde{\mathbf{B}}_m\phi = 0 \quad (2.120)$$

must hold for any field configuration ϕ . Eq. (2.120) is specifically a constraint on the prior energy E^{prior} since the data term E^{data} does not contain any derivatives $X_e^{\Omega,i}\phi$ and thus the canonical momentum \mathbf{P} (eq. (2.105)) only depends on the prior energy density \mathcal{E}^{prior} . There are three cases to consider such that eq. (2.120) can hold:

- Case a: The Lie algebra \mathcal{G} is commutative, $[X_{e,i}^\Omega, X_{e,m}^\Omega] = 0$ for all $1 \leq i, m \leq r$
- Case b: $P_i = 0$ for all $1 \leq i \leq r$
- Case c: If we have $[X_{e,i}^\Omega, X_{e,m}^\Omega] \neq 0$ for some i and m the canonical momentum \mathbf{P} if non-vanishing must be orthogonal to the vector \mathbf{M}_m , which is a vector for fixed m defined as $(\mathbf{M}_m)_i = [X_{e,i}^\Omega, X_{e,m}^\Omega](\phi)$ over Ω

We call cases a and b trivial symmetries and case c a non-trivial symmetry. In chapter 3 we will introduce a prior E^{prior} which is conditionally invariant to the group $\mathbb{G}^\Omega = \mathbb{T} \times SO(2)$ which is the group of local translations and rotations. Its algebra $\mathcal{G} = \mathfrak{t} \times \mathfrak{so}(2)$ is 3-dimensional and although it is not a commutative algebra we will show that eq. (2.120) still holds for any field ϕ .

2.5. Total Variation

In this section we will introduce a widely used method for anisotropic regularization of the GRF ϕ called Total Variation (TV) [18, 19, 38, 39, 40]. In the context of shock-filtering [18, 41, 42] it was shown that the functional

$$E_{L_1}(\phi) = \int |\nabla\phi| dx \quad (2.121)$$

has the appealing property that it does not penalize large discontinuities. However its functional derivative with respect to ϕ is ill conditioned in the case $\nabla\phi \approx 0$. To alleviate the case, [18] chose the approximative prior

$$E_{L_1approx}(\phi) = \int \sqrt{|\nabla\phi|^2 + \epsilon} dx \quad (2.122)$$

which is well behaved for $\epsilon > 0$. They were able to achieve good results with relatively sharp preserved discontinuities with data ϕ^0 having low SNRs. Nevertheless in the limit $\epsilon \rightarrow 0$ the Euler-Lagrange equations become more and more computationally instable. A theoretically more well conditioned form of TV is needed which we will outline, following [40, 43]. To do this we need to explore the function-space the minimizers of eq. (2.121) might belong to. Smooth

functions ϕ_{smooth} are functions for which $\nabla\phi$ exists everywhere, thus they may be minimizers of eq. (2.121). But functions $\phi_{discont}$ containing discontinuities do not have finite L_1 norm of their gradients, $E_{L_1}(\phi_{discont}) = \infty$ since the gradient $\nabla\phi_{discont}$ does not exist at the discontinuities. A generalization of eq. (2.121) is possible if one assumes $\nabla\phi$ to be a distribution, more precisely a Radon measure [40, 43] in the space $\mathcal{M}(\Omega)$. If there exists a Radon measure $\mu \in \mathcal{M}(\Omega)$, such that for every $\mathbf{p} \in \mathcal{C}_0(\Omega)$ with compact domain, the following equality holds

$$\int_{\Omega} \phi \cdot \text{Div} \mathbf{p} dx = - \int \mathbf{p}^T d\mu < \infty \quad (2.123)$$

then μ is called the weak derivative of ϕ and we can identify $\nabla\phi = \mu$. It is then possible to define the function-space of bounded variation

$$BV = \{\phi \in L_1(\Omega) \mid \nabla\phi \in \mathcal{M}(\Omega)\} \quad (2.124)$$

Now it is possible to define a norm on BV . By virtue of the Hölder inequality [44] there exists a scalar C for which we can determine the upper bound of eq. (2.123)

$$\int_{\Omega} \phi \cdot \text{Div} \mathbf{p} dx \leq C \|\phi\|_{\infty} \quad (2.125)$$

The scalar C is the norm of the Radon measure $\nabla\phi$ and is called the total variation of ϕ

$$TV(\phi) = \sup \left\{ \int_{\Omega} \phi \cdot \text{Div} \mathbf{p} dx \mid \|\mathbf{p}\|_{\infty} \leq 1 \right\} \quad (2.126)$$

As was discussed in [40] the functions ϕ are geometrically piecewise smooth, meaning there exists a partitioning $\{\Omega_k\}$ of Ω such that $(\nabla\phi)_{\Omega_k}$ are L_1 integrable. If dl_{mk} is a line segment in the intersection $\Omega_m \cap \Omega_k$ then $TV(\phi)$ can be written in the form

$$TV(\phi) = \sum_k \|\nabla\phi_{\Omega_k}\|_{L_1} + \sum_{k < m} L_{lm} \quad (2.127)$$

$$L_{lm} = \int_{\Omega_l \cap \Omega_m} |\phi_l - \phi_m| dl_{lm} \quad (2.128)$$

where ϕ_l is the value of ϕ on the portion of $\partial\Omega_l$ which is interfacing with Ω_m and vice versa for ϕ_m . The first term in eq. (2.127) penalizes the smooth parts of ϕ (the gradients $(\nabla\phi)_{\Omega_k}$). Similar to eq. (2.29) $\|\nabla\phi_{\Omega_k}\|_{L_1}$ is invariant to shifts of ϕ_{Ω_k} by constants c'_k

$$\|\nabla(\phi_{\Omega_k} + c'_k)\|_{L_1} = \|\nabla\phi_{\Omega_k}\|_{L_1} \quad (2.129)$$

Thus due to eq. (2.129) we can view the smooth functions ϕ_{Ω_k} as being centered around constants c_k

$$\phi_{\Omega_k}(\mathbf{x}) = c_k + \tilde{\phi}_{\Omega_k}(\mathbf{x}) \quad (2.130)$$

where the c_k are determined by the data ϕ_d . For instance if we combined the TV functional with a data term $E_{data} = \sum_k \int_{\Omega_k} (\phi_d - \phi_{\Omega_k})^2 dx$

$$E(\phi, \nabla\phi) = \sum_k \int_{\Omega_k} (\phi_d - \phi_{\Omega_k})^2 dx + \lambda TV(\phi) \quad (2.131)$$

then in [23, 25] it was shown that the c_k can be computed to be the *mean* of the data ϕ_d within the area Ω_k , $c_k = \int_{\Omega_k} \phi_d d^2x$ given that the deviations $\tilde{\phi}_{\Omega_k}$ are penalized by the first term in eq. (2.127).

The second term in eq. (2.127) penalizes the length of the section $\Omega_m \cap \Omega_k$ while maintaining the values $\phi_{k,m}$ and thus the *jump* $|\phi_k - \phi_m|$. It essentially penalizes the curvature of the line interfacing with both Ω_k and Ω_m . We will make this point clear in the following section. For now we remark that if we set $\phi_{\Omega_k} = c_k$ in the data term in eq. (2.131) then we obtain

$$\tilde{E}(\phi, \nabla\phi) = \sum_k \int_{\Omega_k} (\phi_d - c_k)^2 dx + \lambda TV(\phi) \quad (2.132)$$

which is of course only an approximation to eq. (2.131). The data term in eq. (2.132) is a measure for the variance of ϕ_d in Ω_k . The two terms in eq. (2.127) together with the data term $\sum_k \int_{\Omega_k} (\phi_d - c_k)^2 dx$ in eq. (2.132) balance the size of the partitions Ω_k since the boundaries of *small* partitions Ω_k have *high* curvature and thus high TV values, but *low* variances. On the other side *large* partitions Ω_k have boundaries of *low* curvature and thus low TV values, but *high* variances. The parameter λ in eqs. (2.131) and (2.132) marks the trade-off between the TV term and the data term in eqs. (2.131) and (2.132) and thus it determines the size of the partitions Ω_k .

2.5.1. The Mean Curvature of Total Variation

In eq. (2.127) we had argued that the TV measure can be split into a smooth part $\|\nabla\phi_{\Omega_k}\|_{L^1}$ measuring the deviation of the smooth functions ϕ_{Ω_k} from the constants c_k . We had claimed that the second term in eq. (2.127), the boundary term L_{lm} measures the curvature of the boundary between Ω_l and Ω_l . The line integral in L_{lm} in eq. (2.128) can be rewritten essentially as a measure for the

length of the level-set S_{lm} interfacing Ω_l and Ω_m

$$L_{lm} = |\phi_k - \phi_m| \|S\|_{lm}, \quad \|S\|_{lm} = \int_0^T \left\| \frac{d}{dt} (\mathbf{x}(t)) \right\| dt \quad (2.133)$$

The path $\mathbf{x}(t)$ can be considered as being generated by a one parameter Lie group $g_t^{V^\Omega}$ acting on the point \mathbf{x}_0 which is on the interfacing boundary between Ω_l and Ω_m

$$\mathbf{x}(t) = g_t^{V^\Omega} \circ \mathbf{x}_0, \quad \mathbf{x}_0 \in \Omega_l \cap \Omega_m \quad (2.134)$$

so that the length $\|S\|_{lm}$ is controlled by the Lie algebra element $V_e^\Omega = v(\mathbf{x})^\mu \partial_\mu$

$$\|S\|_{lm} = s(T), \quad s(t) = \int_0^t \|v(\mathbf{x}(t'))\| dt' \quad (2.135)$$

The function $s(t)$ in eq. (2.135) is called the *arc length* of the curve $\mathbf{x}(t)$. By virtue of the definition of the arc length $s(t)$ in eq. (2.135) we can express derivatives with respect to s by

$$\frac{d}{ds} = \frac{1}{\|v\|} \frac{d}{dt} \quad (2.136)$$

The *curvature* of $\mathbf{x}_{lm}(t)$, $\kappa(\mathbf{x}_{lm}(t))$ is obtained by re-parameterizing $\mathbf{x}(t)$ in terms of its arc length s , $\mathbf{x}(t) \rightarrow \mathbf{x}(t(s))$ and taking the second derivative of $\mathbf{x}(s)$ using eq. (2.136)

$$\kappa(\mathbf{x}_0) = \left\| \frac{d^2}{ds^2} \mathbf{x}(s) \right\|_{s=0} \quad (2.137)$$

$$= \frac{1}{\|v\|^3} \left(v_x \cdot \frac{dv_y}{dt} - v_y \cdot \frac{dv_x}{dt} \right) \Big|_{t=0} \quad (2.138)$$

In [40] it is shown that the expression for the curvature $\kappa(\mathbf{x}_0)$ at the point \mathbf{x}_0 in eq. (2.138) is equivalent to the *mean curvature* [45, 46, 47]

$$\kappa = \text{Div} \left(\frac{\nabla \phi}{|\nabla \phi|} \right) \quad (2.139)$$

which is the functional derivative of the TV norm in eq. (2.121) with respect to ϕ

$$\kappa = -\partial TV(\phi) \quad (2.140)$$

For a thorough derivation of the mean curvature κ in terms of weak derivatives in BV spaces see [43]. The Euler-Lagrange equations of any energy function

$E = \int \mathcal{E} d^2x$ including the TV functional in eq. (2.121) as a prior

$$[\mathcal{E}] (\mathbf{x}) = \frac{\delta \mathcal{E}}{\delta \phi} (\mathbf{x}) + \lambda \kappa (\mathbf{x}) = 0 \quad (2.141)$$

pose a bound on the value of the curvature $\kappa(\mathbf{x})$. Thus the TV functional penalizes the curvature κ of the S_{lm} interfacing Ω_l and Ω_m . As κ is an invariant of the Lie group $SE(2)$, the group of rotations and translations, TV is also an invariant of that group.

2.6. Optical Flow

In section 2.1 we had introduced the notion of an inverse problem, namely that given some data Y and a model C we would like to find the GRF ϕ which is mapped to the data Y by the model C , see eq. (2.8). A prime example of an inverse problem in computer vision is optical flow [48, 49, 5, 50, 51, 52, 38]. Optical Flow labels the task of densely measuring the motion between two or more frames captured by a camera, or the dense registration of two or more cameras on a pixel-by-pixel basis. Optical flow is a crucial step in many areas of computer vision. For instance optical flow estimation is a part of video compression [53, 54] used to detect areas of the video in which the rate brightness change is small. For example during the recording of a rigid scene optical flow can be used to determine when the camera motion stalls. During such periods the frames of the video can be stored in a memory efficient manner. In recent years structure from stereography and structure from motion (video from a single camera) have gained popularity as a means to capture 3D models for film productions and also due to the availability of low cost 3D printing [49, 55, 56, 57, 58, 59, 60, 4]. In both the stereography and the structure from motion pipelines optical flow is used for the triangulation of the dense point cloud, prior to generation of the final 3D mesh. In the case of a dual-modal setup both cameras may be of different types. For instance in medical imaging multi-modal dense image registration is used to fuse image information from CT and MR modalities of the human brain [61] and of the human spine [62].

In optical flow modeling the task at hand is to estimate the disparity between two images Y and I recorded by two cameras C_Y and C_I (see figure 2.3). Each image is a map between the coordinate space $\Omega \subset \mathbb{R}^2$ and the real numbers \mathbb{R} . Thus $Y(\mathbf{x})$ is the intensity recorded by the camera C_Y at the pixel location $\mathbf{x} \in \Omega$ while $I(\mathbf{x}')$ is the intensity recorded by C_I at the location $\mathbf{x}' \in \Omega$. In figure 2.3a we have depicted a multi-modal setup in which the two cameras C_Y and C_I are recording images (figures 2.3c and 2.3d) from different angles. In this context the

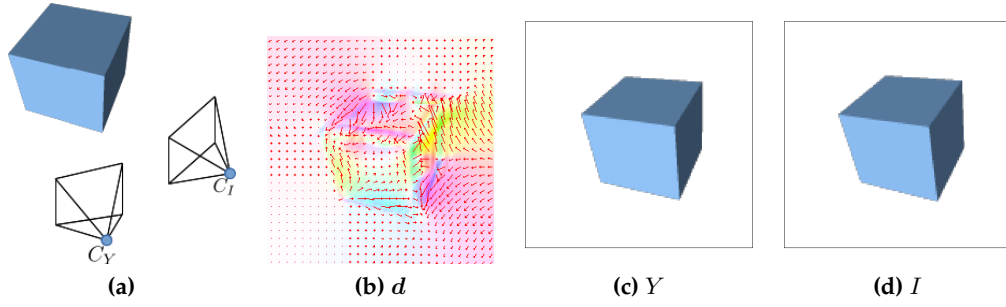


Figure 2.3.: Figure 2.3a: Two cameras C_Y and C_I are shown recording a scene from different positions. The scene could be a rigid scene or a dynamic scene with moving objects. Figure 2.3c shows the image Y captured from the camera C_Y and figure 2.3d the image I from the camera C_I . Figure 2.3b shows the optical flow d . The vectors in figure 2.3b indicate which pixels x' in I and x in Y are mapped to each other.

optical flow field is the unknown variable d which maps the location x' in the image I to the location x in the image Y

$$x' = x + d(x) \quad (2.142)$$

The optical field d is shown in figure 2.3b as a set of vectors at every pixel $x' \in \Omega$, whose magnitude and orientation reflect the motion of the pixel x' .

The standard methodology [48, 49, 63, 64] for the estimation of the optical flow d is to model d as a GRF with a given data term $E_{Y,I}^{data}(d)$. Without further information of the mapping between Y and I from another source (e.g. sparse feature mapping with SIFT features [65]), the data term $E_{Y,I}^{data}(d)$ cannot depend directly on d but can only be defined as a *similarity measure* between the image $Y(x)$ and the *warped* image $I_d(x) = I(x + d(x))$ [63, 64]

$$E_{Y,I}^{data}(d) = F(Y, I_d), \quad I_d(x) = I(x + d(x)) \quad (2.143)$$

In general the mapping of $Y(x)$ and $I_d(x)$ via $d(x)$ is ill-determined: According to [52] the regions $\mathcal{A} \subset \Omega$ of an image ϕ can be given an *intrinsic dimension* [66, 67] iD which depends on their content

- $iD = 0$ if the image patch $\phi_{\mathcal{A}}$ is homogeneous
- $iD = 1$ if the image patch $\phi_{\mathcal{A}}$ contains an edge
- $iD = 2$ otherwise (e.g corners and/or textures)

If we consider two image patches $Y_{\mathcal{A}}$ and $I_{\mathcal{A}}$ with equal intrinsic dimension iD then the number of components of the optical flow d between $Y_{\mathcal{A}}$ and $I_{\mathcal{A}}$ which

can be *uniquely* determined by $E_{Y,I}^{data}(\mathbf{d})$ is identical to iD . For instance if $Y_{\mathcal{A}}$ and $I_{\mathcal{A}}$ both display a corner of an object ($iD = 2$), the optical flow in \mathcal{A} , $\mathbf{d}_{\mathcal{A}}$ can be uniquely determined by the similarity measure $E_{Y,I}^{data}(\mathbf{d})$. However for edges ($iD = 1$) only one component of the optical flow $\mathbf{d}_{\mathcal{A}}$ can be determined and for homogeneous patches $\mathbf{d}_{\mathcal{A}}$ is completely undefined. Thus globally, that is over all $\mathbf{x} \in \Omega$ the similarity measure $E_{Y,I}^{data}(\mathbf{d})$ cannot determine $\mathbf{d}(\mathbf{x})$ uniquely. For this reason optical flow models deploy a prior energy $E^{prior}(\nabla \mathbf{d})$ which smooths $\mathbf{d}(\mathbf{x})$ such that information of the components of \mathbf{d} in regions with $iD = 2$ is carried on to neighboring regions with $iD \neq 2$ such that $\mathbf{d}(\mathbf{x})$ is well defined over all Ω . The total energy

$$E_{Y,I}(\mathbf{d}) = E_{Y,I}^{data}(\mathbf{d}) + E^{prior}(\nabla \mathbf{d}) \quad (2.144)$$

is then a trade-off between the similarity of Y and I_d and the smoothness of the optical flow $\mathbf{d}(\mathbf{x})$.

2.6.1. Uni-Modal Optical Flow

Among the earliest methods for optical flow estimation are the methods described in the seminal papers of Horn and Schunck [48] and Lukas and Kanade [49]. In [48] the following model for computing the flow between two frames of a video was proposed

$$E_{Y,I}(\mathbf{d}) = E_{Y,I}^{data}(\mathbf{d}) + E^{prior}(\mathbf{d}) \quad (2.145)$$

$$E_{Y,I}^{data}(\mathbf{d}) = \int_{\Omega} \left(Y(\mathbf{x}) - I_d(\mathbf{x}) \right)^2 dx, \quad I_d(\mathbf{x}) = I(\mathbf{x} + \mathbf{d}(\mathbf{x})) \quad (2.146)$$

$$E^{prior}(\mathbf{d}) = \lambda \int_{\Omega} \sum_i \|\nabla d_i\|^2 dx \quad (2.147)$$

In eq. (2.146) the frame I is warped back to the frame Y by the field $\mathbf{d}(\mathbf{x})$. The prior energy $E^{prior}(\mathbf{d})$ in eq. (2.147) imposes an isotropic smoothness constraint on the flow field \mathbf{d} . As we discussed in section 2.2.1 the main limitation of the L_2 prior in eq. (2.145) is that it does not preserve edges in the flow field $\mathbf{d}(\mathbf{x})$. To overcome this limitation [68] and [69] used the TV prior in eq. (2.121) as a smoothing term for each of the components of \mathbf{d}

$$E_{Y,I}(\mathbf{d}) = E_{Y,I}^{data}(\mathbf{d}) + \lambda \int_{\Omega} \sum_i \|\nabla d_i\| dx \quad (2.148)$$

According to section 2.5 the level-sets of each component $d_i(\mathbf{x})$ are smoothed while the discontinuities are preserved.




Figure 2.4.: Figures 2.4a and 2.4d show 2 local regions of the VSC image and figures 2.4b and 2.4e their counter parts in the TC image. The histograms in figures 2.4c and 2.4f exhibit two maxima. The maxima can be thought to be the support points of a line in the histogram. Thus *locally* the TC and the VSC image have linear relationships

2.6.2. Multi-Modal Optical Flow

The next issue with the model in eq. (2.145) is that the likelihood $E_{Y,I}^{data}(d)$ in eq. (2.146) makes the assumption that the cameras C_Y and C_I are sensitive to the same physical light spectrum. For instance in figure 2.3 the image Y recorded by the camera C_Y in figure 2.3c has the same intensity spectrum as the image I recorded by the camera C_I (figure 2.3d) and we say that Y and I are *equal by distribution*

$$Y \stackrel{d}{\approx} I \quad (2.149)$$

Thus it is possible to find an optical flow field d^* such that for each pixel $x \in \Omega$ the warped image I_d approximates the image Y , $Y(x) \approx I_d(x)$. However there exists multi-modal setups where the cameras are *not* sensitive to the same spectra. ?? shows two images recorded from a visual spectrum camera (VSC, ??, I_{vsc}) and a thermographic camera (TC, ??, y_{tc}). The recorded object, here a

carbon-fiber reinforced polymer (CFRP) has physically different absorption and emission properties in the visual spectrum domain recorded by the VSC then in the infra-red domain recorded by the TC. Thus the intensities in I_{vsc} (??) follow a completely different distribution than those in y_{tc} (??). This is shown in the joint histogram between I_{vsc} and y_{tc} in ?? which fails to admit the correspondence $I_{vsc} \stackrel{d}{\approx} y_{tc}$.

We will now discuss three statistical similarity measures for arbitrary images Y and I which avoid the assumption of brightness constancy. For this we will take the two images Y and I to be random variables with the marginal distributions $p(Y)$ and $p(I)$. Then the mean and the variance are defined as

$$\mathbb{E}(X) = \int X \cdot p(X) \quad (2.150)$$

$$\text{Var}(X) = \mathbb{E}\left((X - \mathbb{E}(X))^2\right) \quad (2.151)$$

The three similarity measures all avoid the brightness constancy assumption implied by the data term in eq. (2.146) by only taking into account the *statistical* features of the images Y and I such as their joint entropy and joint covariance.

Mutual Information

Mutual Information (MI) [70, 71, 63, 72] is a popular similarity measure used mainly in medical imaging where images from different modalities including Magnetic Resonance Imaging (MRI) [73], Computed Tomography (CT) [74] and Positron Emission Tomography (PET) [75] are registered against each other. For images Y and I from two different modalities capturing the same scene, MI is defined with the joint distribution $p(Y, I)$ by

$$MI(Y, I) = \int p(\hat{Y}, \hat{I}) \ln \frac{p(\hat{Y}, \hat{I})}{p(\hat{Y}) \cdot p(\hat{I})} d\hat{Y} d\hat{I} \quad (2.152)$$

MI measures how strongly the images Y and I *statistically* depend on each other. In the case that Y and I are statistically independent, $p(Y, I) = p(Y) \cdot p(I)$, then by eq. (2.152) MI is zero. On the other side, MI is maximal when I completely determinates Y or vice versa. In the context of optical flow MI is used to measure the similarity between Y and I_d

$$E_{Y,I}^{data}(\mathbf{d}) = -MI(Y, I_d) \quad (2.153)$$

However, as [76] puts it, MI does not explain the kind of dependency between images Y and I , its maxima are statistically but not visually meaningful, since it disregards any spatial information, which is essential for optical flow. Thus optical flow likelihoods based on MI usually tend to have many local minima rendering MI too unconstrained for optical flow.

Correlation Ratio

To alleviate the problems with MI, [76] arguments that a better similarity measure would be one that measures the *functional* relation between the images Y and I . The key ingredient for their proposal is that the pixel values $I(\mathbf{x})$ and $Y(\mathbf{x})$ are assumed to be the realizations of random variables, which by abuse of notation we denote by \hat{I} and \hat{Y} . Then the normalized joint histogram of the images I and Y can be interpreted as the joint probability distribution $p(\hat{Y}, \hat{I})$, and the conditional distribution

$$p(\hat{Y} | \hat{I} = I) = \frac{p(\hat{Y}, \hat{I} = I)}{p(\hat{I} = I)} \quad (2.154)$$

encodes the spatial functional relationship between Y and I . They introduced the Correlation Ratio (CR) [77, 71, 63, 64]

$$\eta_{CR}(I|Y) = \frac{\text{Var}(\phi^*(Y))}{\text{Var}(I)} \quad E_{Y,I}^{data}(\mathbf{d}) = -\eta(I_d|Y) \quad (2.155)$$

The optimal function ϕ^* was shown to be the expectation value of \hat{I} , conditioned on a realization of \hat{Y}

$$\phi^*(Y) = \mathbb{E}(\hat{I} | \hat{Y} = Y) = \int I p(I|Y) dI \quad (2.156)$$

The function $\phi(\hat{Y})$ maps any realization of \hat{Y} to an expectation value of \hat{I} . As \hat{Y} is a random variable, $\phi(\hat{Y})$ is also at random. Its variance measures how well I is *functionally explained* by a realization of \hat{Y} . The measure in eq. (2.155) is bounded between 0 and 1, 0 indicating that Y and I are independent, 1 indicating a functional relationship $I = \phi^*(Y)$. The function ϕ^* , although not necessarily continuous, is measurable in the L_2 -sense. Thus CR is a much stronger constraint than MI and has fewer, but more meaningful minima [76].




Figure 2.5.: Figure 2.5a shows a synthetic high resolution image I^{syn} . In figure 2.5b we show a low resolution image Y^{syn} . Y^{syn} is computed by down sampling I^{syn} by a factor $\sigma^{scale} = 5$, creating y^{syn} followed by cubic interpolation and translated by 10 pixels relative to I^{syn} . An optical flow model which incorporates knowledge of the scale difference between Y^{syn} and I^{syn} should produce a flow \tilde{d} , such that the warped image I_d^{syn} matches the image Y^{syn} up to scale σ^{scale} , thereby preserving the features of I^{syn} . Figure 2.5c shows the flow d computed with the model in eq. (2.145). Since the model in eq. (2.145) does not incorporate knowledge of the scale difference σ^{scale} , the features of the warped image I_d^{syn} (figure 2.5d) are heavily distorted

Cross Correlation

Cross Correlation [78, 79, 80, 63] is the strongest constrained similarity measure. It is basically an additional constraint to CR, namely that the functional relationship in eq. 2.155 must be linear. Then η reduces to

$$\eta_{CC}(I|Y) = \frac{Cov(Y, I)}{Var(I) \cdot Var(Y)} \quad Y = f \cdot I + \beta \quad (2.157)$$

As we will see in section 3.4 a measure similar to eq. (2.157) will be computed based on the assumption that both Y and I are Gaussian. The Gaussian assumption is valid when both cameras Y and I produce Gaussian noise and the joint histogram is predominantly linear. Linearity in the joint histogram occurs when the recorded scene contains materials with uniform luminosity in the frequency bands of the cameras Y and I .

All three similarity measures have in common that the images Y and I must have the same spatial resolution in order to compute the measure. For instance in eq. (2.153) [70] computed the joint probability $p(Y, I)$ as a normalized histogram $h(\hat{Y}, \hat{I})$ created from the samples $\hat{Y} = Y(\mathbf{x})$ and $\hat{I} = I(\mathbf{x})$ drawn from all locations $\mathbf{x} \in \Omega$.

The problem that we want to attack is the determination of the optical flow between a *low resolution image* y obtained from a low resolution camera C_y and a high resolution image I from a camera C_I . From now on lower case letters

stand for low resolution and higher case letters for high resolution images. In an ad-hoc fashion we could first filter and down-sample the image I with a convolution filter G to obtain an image i with the same spatial resolution of y

$$i(\mathbf{x}) = (G \star I)(\mathbf{x}), \quad \mathbf{x} \in \Omega \quad (2.158)$$

and evaluate the similarity measures on the image pair y and i . The negative impact is that we could only estimate an optical flow \mathbf{d} with the same low resolution as the image y . Conversely we could up-sample the image y with some interpolation scheme to produce a high resolution image Y and evaluate the similarity measures on the pair Y and I . This situation is shown in figure 2.5. I^{syn} in figure 2.5a shows a sharp linear boundary. We down sampled I^{syn} by a factor $\sigma^{scale} = 5$ to create y^{syn} . y^{syn} was translated by 10 pixels and the high resolution image Y^{syn} (figure 2.5b) was computed by a cubic interpolation of y^{syn} again by a factor $\sigma^{scale} = 5$. The images Y^{syn} and I^{syn} thus differ in optical scale, and the scale difference is the parameter σ^{scale} . We used the model of [48]

$$E(\mathbf{d}^{syn}) = \frac{1}{2} \int_{\Omega} \left(Y^{syn}(\mathbf{x}) - I_{\mathbf{d}^{syn}}^{syn}(\mathbf{x}) \right)^2 dx + \frac{\lambda}{2} \sum_i \int_{\Omega} \|\nabla d_i^{syn}(\mathbf{x})\|^2 dx \quad (2.159)$$

(see eq. (2.145)) to compute the optical flow \mathbf{d}^{syn} mapping I^{syn} to Y^{syn} (see figure 2.5c). Figure 2.5d shows the image $I_{\mathbf{d}^{syn}}^{syn}(\mathbf{x}) = I^{syn}(\mathbf{x} + \mathbf{d}^{syn}(\mathbf{x}))$. We can see that the optical flow \mathbf{d} corrupts the sharp boundary of I^{syn} in order to match it to the varying gray levels of the blurred boundary in Y^{syn} (figure 2.5b). The problem is that the model in eq. (2.159) can account for the *difference in size* of the images y^{syn} and I^{syn} but it does not take the *difference in optical scale* σ^{scale} into account. Thus we need a model that can account for the optical scale σ^{scale} .

2.7. Image Fusion

In this section we will introduce the image fusion method of Hardie et. al. [81]. In that paper the authors solved the problem of refining the low optical resolution of an image y_{tc} obtained by a thermographic camera (TC) using the image I_{vsc} obtained by a visual spectrum camera (VSC). The result of their method is a thermographic image Y_{tc} with improved optical resolution (see figure 2.6a for a schematical overview). They used this method for the subject of remote sensing [83] where the TC and the VSC are built in a co-aligned fashion within the body of a satellite. In section 3.3 the goal is to extend this method to the case where the TC and the VSC are not co-aligned. In this case the low resolution image y_{tc} and thus the high resolution Y_{tc} have a natural separation from the VSC image




Figure 2.6.: Figure 2.6a shows a schematic setup of the camera configuration considered by Hardie et. al. [81]. The blue line indicates the orthogonal direction to both the TC and the VSC image planes. The camera centers (indicated by the blue circles) are aligned along the dashed line. Figure 2.6b shows the image I_{vsc} captured by the VSC and figure 2.6c the image y_{tc} captured by the TC. The image y_{tc} has a smaller optical resolution than the image I_{vsc} . The method in [81] takes the data y_{tc} and I_{vsc} to produce a higher resolution thermographic image Y_{tc} , shown in figure 2.6d

I_{vsc} similar to figure 2.3. We will show that it is possible to jointly estimate the image Y_{tc} and the optical flow $d(x)$ between Y_{tc} and I_{vsc} .

The method of [81] goes as follows: In figure 2.7a a model of the CCD of the low resolution TC is shown overlaid with a higher resolution grid representing the VSC. The gray region in figure 2.7a symbolizes one pixel of the TC and it can be seen that each pixel of the TC covers a group of pixels of the VSC. Since the TC pixel has a finite surface, we need to specify how this pixel absorbs photons landing at different points in its area in order to relate the covered pixels of the VSC to it. The response of each individual pixel in the TC is called the point spread function (PSF), $W_{\sigma^{scale}}(x, y)$, the vector (x, y) being the location on the surface of the TC pixel with respect to the VSC coordinate frame. Figure 2.7b is the result of a theoretical model of a FLIR TC [84]. The model, obtained by Hardie et al. [82], combines absorption properties of the CCD pixel with physical properties of the camera lens. We can see that each TC pixel has a non uniform response to incoming photons. Using this information we can model a super-resolved version Y_{tc} of the TC image y_{tc} with the help of the PSF $W_{\sigma^{scale}}$, by stating that y_{tc} is the result of the convolution of Y_{tc} with $W_{\sigma^{scale}}$

$$y_{tc} = W_{\sigma^{scale}} Y_{tc} + n \quad n \sim \mathcal{N}(0|C_n) \quad (2.160)$$

The problem of estimating Y_{tc} is that there is an infinite amount of high resolution TC images Y_{tc}^* which relate to y_{tc} via eq. (2.160) since the high spacial frequency components of Y_{tc} are filtered out by $W_{\sigma^{scale}}$. In [81] Hardie suggested the use of a high resolution imager I_{vsc} whose camera center is co-aligned (see figure 2.6a) with the TC image y_{tc} and correlated with Y_{tc} . The rationale behind

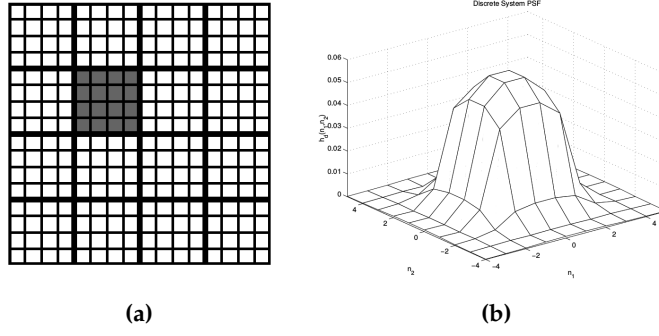


Figure 2.7.: Figure 2.7a The thick grid depicts the CCD of the low resolution thermographic camera. The finer grid a virtual super-resolved version of the pixels in the TC. Figure 2.7b shows the point spread function $W_{\sigma^{scale}}(x, y)$ of the gray pixel in figure 2.7a, taken from Hardie et al. [82]. It shows that each pixel in the TC image has a non uniform response over its surface to incoming photons.

their approach is to combine the desired features such as sharp edges and corners of I_{vsc} with the intensity spectrum of y_{tc} into the super-resolved image Y_{tc} , while avoiding limitations such as the noise model of y_{tc} .

The key ingredient in the model of [81] is that the intensities of Y_{tc} and I_{vsc} are assumed to be samples drawn from the joint Gaussian $p(Y_{tc}, I_{vsc})$. As I_{vsc} is already fixed as input data we can derive a conditional distribution for Y_{tc} via the Bayesian rule [85]

$$p(Y_{tc}|I_{vsc}) = \frac{p(Y_{tc}, I_{vsc})}{p(I_{vsc})} \sim \mathcal{N}(\mu_{Y_{tc}|I_{vsc}}|C_{Y_{tc}|I_{vsc}}) \quad (2.161)$$

$$C_{Y_{tc}|I_{vsc}} = C_{Y_{tc}, Y_{tc}} - C_{Y_{tc}, I_{vsc}}^2 \cdot C_{I_{vsc}, I_{vsc}}^{-1} \quad (2.162)$$

$$\mu_{Y_{tc}|I_{vsc}}(\mathbf{x}) = \mu_{Y_{tc}} + C_{Y_{tc}, I_{vsc}} \cdot C_{I_{vsc}, I_{vsc}}^{-1} (I_{vsc}(\mathbf{x}) - \mu_{I_{vsc}}) \quad (2.163)$$

where the variances and means are computed globally

$$C_{u,v} = \int_{\Omega} (u(\mathbf{x}) - \mu_u) \cdot (v(\mathbf{x}) - \mu_v) d^2x \quad (2.164)$$

$$\mu_u = \int_{\Omega} u(\mathbf{x}) d^2x$$

We see that the mean of Y_{tc} conditioned on I_{vsc} , $\mu_{Y_{tc}|I_{vsc}}$ (eq. (2.163)) is linear in the values of I_{vsc} , thus in this model the intensities of Y_{tc} are assumed to be *globally* linearly related to the intensities of I_{vsc} . We combine eq. (2.161) with

the Gaussian likelihood in eq. (2.160) to the posterior

$$p(Y_{tc}|y_{tc}, I_{vsc}) \sim p(y_{tc}|Y_{tc}) \cdot p(Y_{tc}|I_{vsc}) = \exp(-E_{y_{tc}, I_{vsc}}(Y_{tc})) \quad (2.165)$$

with the associated energy

$$\begin{aligned} E_{y_{tc}, I_{vsc}}(Y_{tc}) &= \frac{1}{2} \int_{\Omega} \left(y_{tc}(\mathbf{x}) - W_{\sigma^{scale}} Y_{tc}(\mathbf{x}) \right)^2 \cdot C_n^{-1} d^2x \\ &\quad + \frac{1}{2} \int_{\Omega} \left(Y_{tc}(\mathbf{x}) - \mu_{Y_{tc}|I_{vsc}}(\mathbf{x}) \right)^2 \cdot C_{Y_{tc}|I_{vsc}}^{-1} d^2x \end{aligned} \quad (2.166)$$

The minimization of eq. (2.166) and thus maximization of (2.165) with respect to Y_{tc} gives the analytical solution [81]

$$\begin{aligned} Y_{tc}^* &= \mu_{Y_{tc}|I_{vsc}} + C_{Y_{tc}|I_{vsc}} \cdot W_{\sigma^{scale}}^T \mathbf{M}^{-1} \cdot \left(y_{tc} - \langle \mu_{Y_{tc}|I_{vsc}} \rangle_{\sigma^{sc}} \right) \\ \langle \mu_{Y_{tc}|I_{vsc}} \rangle_{\sigma^{sc}} &= W_{\sigma^{scale}} \mu_{Y_{tc}|I_{vsc}} \\ \mathbf{M} &= \left(W_{\sigma^{scale}} \cdot C_{Y_{tc}|I_{vsc}} \cdot W_{\sigma^{scale}}^T + C_n \right) \end{aligned} \quad (2.167)$$

Eq. (2.167) is computationally expensive due to the dense operator $W_{\sigma^{scale}}$ and the matrix-inverse operation. However if $W_{\sigma^{scale}}$ is approximately Gaussian then the diagonal entries of the matrix \mathbf{M} are larger than the off-diagonal entries. In [86] a computationally tractable approximation was introduced

$$\hat{Y}_{tc} = \mu_{Y_{tc}|I_{vsc}} + C_{\langle Y_{tc} \rangle_{\sigma^{sc}} | \langle I_{vsc} \rangle_{\sigma^{sc}}} \cdot \tilde{\mathbf{M}}^{-1} \left(y_{tc} - \langle \mu_{Y_{tc}|I_{vsc}} \rangle_{\sigma^{sc}} \right) \quad (2.168)$$

$$\langle I_{vsc} \rangle_{\sigma^{sc}} = W_{\sigma^{scale}} I_{vsc}, \quad \langle Y_{tc} \rangle_{\sigma^{sc}} = W_{\sigma^{scale}} Y_{tc} \approx y_{tc} \quad (2.169)$$

$$\tilde{\mathbf{M}} = \left(C_{\langle Y_{tc} \rangle_{\sigma^{sc}} | \langle I_{vsc} \rangle_{\sigma^{sc}}} + C_n \right) \quad (2.170)$$

$$\mu_{\langle Y_{tc} \rangle_{\sigma^{sc}} | I_{vsc}}(\mathbf{x}) = \mu_{\langle Y_{tc} \rangle_{\sigma^{sc}}} + C_{\langle Y_{tc} \rangle_{\sigma^{sc}} | I_{vsc}} \cdot C_{I_{vsc}, I_{vsc}}^{-1} \left(I_{vsc}(\mathbf{x}) - \mu_{I_{vsc}} \right) \quad (2.171)$$

where the matrix $\tilde{\mathbf{M}}$ is a diagonal matrix and thus easily invertible. The approximated conditional mean $\mu_{\langle Y_{tc} \rangle_{\sigma^{sc}} | I_{vsc}}$ is a transformation of the intensities of the VSC image I_{vsc} to the spectrum of the TC image y_{tc} .

The key issue is that eq. (2.168) requires both modalities, I_{vsc} and y_{tc} , to be co-aligned. Since we are dealing with an optical flow problem y_{tc} and thus Y_{tc} is shifted by a disparity $d(\mathbf{x})$ from I_{vsc} . This disparity has to be taken in to account by our model in chapter 3.3. The second issue is that the assumption that Y_{tc} and I_{vsc} are *globally* joint Gaussian is not supported by our data. However by computing $C_{Y_{tc}|I_{vsc}}$ in *local* sub-domains of the space Ω we can show that Y_{tc} and I_{vsc} are *locally* joint Gaussian. This will also be shown in chapter 3.3.

3. Linearized Priors

3.1. The Linear Structure Tensor

We shall now proceed to introduce a prior based on the considerations made in chapter 2.2. We will concentrate on the translation group \mathbb{T} for which the Lie algebra \mathfrak{t} is characterized by the set of vectors \mathbf{v} which are constant within a sub domain $A \subset \Omega$. The basis operators $X_e^i \in \mathfrak{t}$ are the Cartesian differential operators $\{\partial_x, \partial_y\}$, and the spatial component V_e^Ω of a vector $V_e \in \mathfrak{t}$ has the representation

$$V_e^\Omega = v_x(\mathbf{x}) \partial_x + v_y(\mathbf{x}) \partial_y \in \mathfrak{t} \quad \mathbf{v}(\mathbf{x})|_A = \text{const} \quad (3.1)$$

The translation group \mathfrak{t} is a commutative algebra so the basis $\{\partial_x, \partial_y\}$ is commutative and any vector V_e^Ω commutes with the $\{\partial_x, \partial_y\}$

$$[V_e^\Omega, \partial_{x,y}] = 0 \quad (3.2)$$

Thus V_e^Ω is translation invariant. Consider an image $\phi(\mathbf{x})$. The level-sets S_X corresponding to the vector V_e^Ω are defined by

$$S_X = \left\{ \mathbf{x} \mid \mathbf{v}^T \cdot \nabla \phi(\mathbf{x}) = 0 \right\} \quad (3.3)$$

We would like to characterize the dominant strength and the orientation of $\nabla \phi$ within the sub domain $A \subset \Omega$. In [87] it was suggested that the tangential vector \mathbf{v} of the level sets S_X can be computed by minimizing the energy

$$J(\mathbf{v}) = \frac{1}{2} \int_A w(\|\mathbf{x}\|) \mathbf{v}^T \cdot \left(\nabla \phi(\mathbf{x}) \nabla^T \phi(\mathbf{x}) \right) \mathbf{v} = \frac{1}{2} \mathbf{v}^T S \mathbf{v} \quad (3.4)$$

$$S = \int_A w(\|\mathbf{x}\|) \left(\nabla \phi(\mathbf{x}) \nabla^T \phi(\mathbf{x}) \right) d^2x = \langle \nabla \phi \nabla^T \phi \rangle \quad (3.5)$$

The weight function $w(\|\mathbf{x}\|)$ is normalized and weights the contributions of the gradient $\nabla \phi(\mathbf{x})$ at various points $\mathbf{x} \in A$. Typically a Gaussian is deployed for the weight function, $w(\|\mathbf{x}\|) = G_\kappa(\mathbf{x} - \mathbf{x}_0)$ where \mathbf{x}_0 is the center pixel of A . This way the gradient $\nabla \phi(\mathbf{x}_0)$ is the dominant contribution to the integral in eq. (3.5).

The matrix S is called the structure tensor. Since S is a symmetric matrix there exists an orthogonal decomposition

$$S = V^T D V \quad D = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix} \quad V = (\mathbf{V}_1, \mathbf{V}_2) \quad (3.6)$$

The eigenvalues give of the squared strength of the gradient in the basis defined by the columns of V . They characterize the structure in A in the following way

- $\lambda_1 > \lambda_2$: Strong linear level set with tangential vector $\mathbf{v} = \mathbf{V}_2$
- $\lambda_1 \approx \lambda_2 \approx 0$: No strong gradient, image is approximately constant
- $\lambda_1 \approx \lambda_2 \gg 0$: No *linear* level sets, level sets have strong curvature

We want to study the variation of the structure tensor S under the $SO(2)$ at the unit element e . Let S_θ be the structure tensor where the local coordinate frame A is rotated by the $SO(2)$ (see eq. (2.73))

$$S_\theta = \int_A w(\|\mathbf{x}(\theta)\|) \left(\nabla_{\mathbf{x}(\theta)} \phi(\mathbf{x}(\theta)) \nabla_{\mathbf{x}(\theta)}^T \phi(\mathbf{x}(\theta)) \right) d^2 x_\theta \quad (3.7)$$

The $SO(2)$ only rotates the domain A and does not deform it otherwise, thus the integral measure $d^2 x_\theta$ is independent of θ , $d^2 x_\theta = d^2 x$. Since the weighting function w only depends on the norm $\|\mathbf{x}(\theta)\|$ which is preserved by the $SO(2)$, it is also invariant. The only component which changes is the gradient $\nabla_{\mathbf{x}(\theta)}$. Using eq. (2.74) and the product rule we can compute the derivative of S_θ at $\theta = 0$

$$\frac{d}{d\theta} S_\theta \Big|_{\theta=0} = \int_A w(\|\mathbf{x}\|) \left(\mathbf{M}_\theta \nabla \phi \nabla^T \phi + \nabla \phi \nabla^T \phi \mathbf{M}_\theta^T \right) d^2 x \quad (3.8)$$

$$= \mathbf{M}_\theta \cdot S - S \cdot \mathbf{M}_\theta = [\mathbf{M}_\theta, S] \quad (3.9)$$

In eq. (3.9) we used $\mathbf{M}_\theta^T = -\mathbf{M}_\theta$. We can get some information on the magnitude of the rate of change $\frac{d}{d\theta} S_\theta \Big|_{\theta=0}$ by multiplying the commutator in eq. (3.9) with the eigenvectors $\mathbf{v}_{1,2}$

$$\tilde{\mathbf{v}}_{1,2} = [\mathbf{M}_\theta, S] \mathbf{v}_{1,2} \quad (3.10)$$

It is easy to show that both projections $\tilde{\mathbf{v}}_{1,2}$ in eq. (3.10) have the same norm

$$\|\tilde{\mathbf{v}}_{1,2}\| = |\lambda_1 - \lambda_2| \quad (3.11)$$

With the help of eq. (3.11) we can reformulate our characterization of the eigenvalues $\lambda_{1,2}$ in the following way

- $\lambda_1 > \lambda_2$: Structure tensor S has strong change under $SO(2)$
- $\lambda_1 \approx \lambda_2$: Structure tensor S is largely invariant under the $SO(2)$ and approximately diagonal, $S \sim \mathbb{1}_{2 \times 2}$ where $\mathbb{1}_{2 \times 2}$ is the unity element of \mathbb{R}^2

3.2. Structure Tensor Based Prior

Since the vector field V_e^Ω in eq. (3.1) is translation invariant the structure tensor S is also translation invariant. Under the rotation group $SO(2)$ the structure tensor is *not* invariant. Nonetheless it has an important transformation property: the transformed structure tensor S' may be written in terms of the old matrix S and the rotation matrix $R_\theta \in SO(2)$

$$S' = R_\theta^T S R_\theta \quad (3.12)$$

We would like to construct a prior E_{ST}^{prior} which is invariant under the combined group $\mathbb{G}_\Omega = \mathbb{T} \times SO(2)$. Since the eigenvalues λ_i of the structure tensor S are positive definite we propose as an energy prior for ϕ the integral over the determinant of S

$$E_{ST}^{prior} = \int_\Omega \mathcal{E}_{ST}(S) d^2x \quad (3.13)$$

$$\mathcal{E}_{ST}(S) = \frac{\lambda}{2} \det(S) \quad (3.14)$$

We want to show that E_{ST}^{prior} is invariant under the $SO(2)$. We insert S_θ from eq. (3.7) into the determinant in eq. (3.14) and evaluate the derivative of E_{ST}^{prior} with respect to θ

$$\left. \frac{d}{d\theta} E_{ST}^{prior}(S_\theta) \right|_{\theta=0} = \int_\Omega \text{Tr} \left(\mathbf{P}^{ST} \cdot [\mathbf{M}_\theta, S] \right) d^2x, \quad P_{ij}^{ST} = \frac{\delta \mathcal{E}_{ST}}{\delta S_{ij}} \quad (3.15)$$

The matrix \mathbf{P}^{ST} is the canonical momentum with respect to the structure tensor S , thus \mathbf{P}^{ST} has the same transformation properties under the $SO(2)$ as S . The trace in eq. (3.15) can be further transformed

$$\text{Tr} \left(\mathbf{P}^{ST} \cdot [\mathbf{M}_\theta, S] \right) = 2 \cdot \text{Tr} \left(\mathbf{P}^{ST} \cdot \mathbf{M}_\theta \cdot S \right) \quad (3.16)$$

The matrix under the trace on the right hand side of eq. (3.16) is a product of a symmetric and an anti-symmetric matrix, and thus itself anti-symmetric. Since traces over anti-symmetric matrices vanish, it follows that the prior E_{ST}^{prior} is

invariant under the $SO(2)$

$$\left. \frac{d}{d\theta} E_{ST}^{prior}(S_\theta) \right|_{\theta=0} = 2 \int_{\Omega} \text{Tr} \left(\mathbf{P}^{ST} \cdot \mathbf{M}_\theta \cdot S \right) d^2x = 0 \quad (3.17)$$

We note that the symmetry expressed by eq. (3.17) is a non-trivial symmetry, since only the trace as a whole vanishes.

3.3. Geometrical Optical Flow Model

In section 2.6 we introduced the notion of an optical flow field $\mathbf{d}(\mathbf{x})$ which maps the domain Ω_I of an image $I(\mathbf{x})$ recorded by the camera C_I to the domain Ω_y of the image $Y(\mathbf{x})$ recorded by the camera C_Y (see figure 2.3). The basic variational method outlined for obtaining $\mathbf{d}(\mathbf{x})$ was: The optical flow $\mathbf{d}(\mathbf{x})$ is computed by minimizing the energy functional $E_{Y,I}(\mathbf{d})$ (see eq. (2.144)) which contains a data term (also called similarity measure in this context) $E_{Y,I}^{data}(\mathbf{d})$ and a prior energy for the gradient of \mathbf{d} , $E^{prior}(\nabla \mathbf{d})$. The similarity measure $E_{Y,I}^{data}(\hat{\mathbf{d}})$ basically tells us how similar the images $Y(\mathbf{x})$ and $I_{\hat{\mathbf{d}}}(\mathbf{x})$ (defined in eq. (2.143)) are given an estimate $\hat{\mathbf{d}}$ of the optical flow field \mathbf{d} .

In the uni-modal case in section 2.6.1 the similarity measure $E_{Y,I}^{data}$ was defined as the pixel difference between the $Y(\mathbf{x})$ and $I_{\mathbf{d}}(\mathbf{x})$ (see eq. (2.146)). In section 2.6.2 it was explained that the measure $E_{Y,I}^{data}$ defined in eq. (2.146) is generally insufficient in the multi-modal case where the cameras C_I and C_Y may be sensitive to different light spectra. We introduced alternative similarity measures based on Mutual Information (MI) (eq. (2.152), [70]), Correlation Ratio (CR) (eq. (2.155), [77]) and Cross Correlation (CC) (eq. (2.157), [71]). The similarity measure based on MI, CR and CC pose constraints on the similarity $Y(\mathbf{x})$ and $I_{\hat{\mathbf{d}}}(\mathbf{x})$ in ascending order with CC posing the strongest constraint $Y(\mathbf{x}) = f \cdot I(\mathbf{x}) + \beta$ (see eq. (2.157)). However all three measures assume that the images $Y(\mathbf{x})$ and $I_{\mathbf{d}}(\mathbf{x})$ have the same spacial and optical resolution such in the case of CC an incorrect optical flow \mathbf{d}^* is estimated (see figure 2.5).

In section 2.7 we outlined the method of Hardie et. al. [81] in which given a low resolution image y the goal is to estimate an image Y with higher resolution with the aid of a high resolution image I from an external camera C_I . The virtue of their method is that it integrates the physical relationship between the CCD of C_y and that of C_I . This relationship is embodied by the PSF $W_{\sigma_{sc}}$ in eq. (2.160). However in their paper [81] they assumed C_y and C_I to co-aligned (see figure 2.6). In this section we will extend this method to incorporate an optical flow field mapping Ω_I to Ω_Y (the domain of the unknown high resolution




Figure 3.1.: Figure 3.1a shows the setup of a thermographic camera (TC), C_{tc} , and a visual spectrum camera (VSC), C_{vsc} , recording an object O . Figure 3.1b shows the image I which is recorded by C_{vsc} and figure 3.1c the lower resolution image y recorded by C_{tc} . The solid line cone of C_{tc} in figure 3.1a which is small compared to the cone of C_{vsc} indicates the low resolution of the TC compared to that of the VSC. The dotted cone indicates the high resolution of the image Y , which is jointly estimated together with the optical flow d (the mapping between I and y) by the model in eq. (3.25)

Y). The result, a model capable of jointly estimating $d(x)$ and Y given y and I is basically a CC-type similarity measure encoding the difference in optical resolution between the C_y and the C_I camera.

3.4. Multi-Modal Optical Flow with Differing Resolutions

We consider the camera setup in figure 3.1a with the low resolution thermographic camera C_{tc} and the high resolution visual spectrum camera C_{vsc} . As opposed to the setup in Hardie et. al. [81] (figure 2.6) the cameras C_{tc} and C_{vsc} are not co-aligned. The goal in this section is to extend the method introduced in section 2.7 to include the unknown optical flow field $d(x)$ representing the separation of C_{tc} and C_{vsc} . In a nutshell we assume the low resolution image y to be co-aligned with the image I_d

$$I_d(x) = I(x + d(x)) \quad (3.18)$$

for a given optical flow field $d(x)$. From eq. (2.167) we can compute the super resolved image $\hat{Y}_d(x)$ which is a function of d . In principle the energy in eq. (2.166) with the image $\hat{Y}_d(x)$ then serves as a similarity measure between y and I_d

$$E_{y,I}^{data}(d) = E_{y,I_d}(\hat{Y}_d) \quad (3.19)$$

\hat{Y}_d is an implicit function which does not need to be estimated directly since by virtue of eq. (2.167) it is easily computed when needed. However since we are only interested in d we do not need to explicitly evaluate \hat{Y}_d .

Computation of the similarity measure $E_{y,I}^{data}$

We will now compute the exact form of $E_{y,I}^{data}(d)$ in eq. (3.19).

In section 2.7 the model for computing the super-resolved image Y is given by the energy $E(Y)$ (see eq. (2.166))

$$E_{y,I}(Y) = \frac{1}{2} \int_{\Omega} \left(y(\mathbf{x}) - W_{\sigma^{sc}} Y(\mathbf{x}) \right)^2 \cdot C_n^{-1} d^2x \quad (3.20)$$

$$+ \frac{1}{2} \int_{\Omega} \left(Y(\mathbf{x}) - \mu_{Y|I}(\mathbf{x}) \right)^2 \cdot C_{Y|I}^{-1} d^2x \quad (3.21)$$

with the conditional variance and mean

$$C_{Y|I} = C_{Y,Y}^{\sigma^{sc}} - C_{Y,I}^{\sigma^{sc},2} \cdot C_{I,I}^{\sigma^{sc},-1} \quad (3.22)$$

$$\mu_{Y|I}(\mathbf{x}) = \mu_Y + f \cdot (I(\mathbf{x}) - \mu_I), \quad f = C_{Y,I}^{\sigma^{sc}} \cdot C_{I,I}^{\sigma^{sc},-1} \quad (3.23)$$

The first integrand of $E_{y,I}(Y)$ in eq. (3.20) models the relationship between the low resolution image y of the camera C_{tc} and the unknown high resolution image Y , namely that one pixel in y is mapped to a window of pixels in Y via the PSF $W_{\sigma^{sc}}$ (figure 2.7a). Essentially it couples the low resolution component of Y , $W_{\sigma^{sc}}Y$ to the C_{tc} image y . On the other side the second integrand of $E_{y,I}(Y)$ in eq. (3.21) models the relationship between the *intensities* of Y and I . This done by transforming the spectrum of I via the factor f (eq. (3.23)) to match that of Y . Since this is done on a pixel-by-pixel basis, eq. (3.21) pins down the high resolution component of Y .

At this point we incorporate the optical flow $d(\mathbf{x})$ which separates the cameras C_{tc} and C_{vsc} by assuming Y to be co-aligned the warped image I_d (eq. (3.18)). Thus we substitute I for I_d in the integrand in eq. (3.21)

$$E(Y, d) = \frac{1}{2} \int_{\Omega} \left(y(\mathbf{x}) - \langle Y \rangle_{\sigma^{sc}}(\mathbf{x}) \right)^2 \cdot C_n^{-1} d^2x \quad (3.24)$$

$$+ \frac{1}{2} \int_{\Omega} \left(Y(\mathbf{x}) - \mu_{Y|I_d}(\mathbf{x}) \right)^2 \cdot C_{Y|I_d}^{-1} d^2x \quad (3.25)$$

with expressions for $C_{Y|I_d}$ and $\mu_{Y|I_d}$ similar to those in eqs. 3.22 and 3.23.




Figure 3.2.: Figure 3.2a shows a synthetic high resolution image I^{syn} . In figure 3.2b we show a low resolution image y^{syn} , y^{syn} is computed by convolution of I^{syn} with Gaussian $G_{\sigma^{sc}}$ with standard deviation $\sigma^{sc} = 5$ and translated by 10 pixels relative to I^{syn} . Figure 3.2d shows the flow \mathbf{d} computed with the model in eq. (??), which incorporates knowledge of the scale difference between y^{syn} and I^{syn} and figure 3.2c show the warped image I_d

While keeping \mathbf{d} fixed we minimize $E(Y, \mathbf{d})$ with respect to Y and obtain a simplified closed form solution for Y similar to eq. (2.168)

$$\hat{Y} = \mu_{Y|I_d} + C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}} \cdot \left(C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}} + C_n \right)^{-1} \left(y - \tilde{\mu}_{Y|I_d} \right) \quad (3.26)$$

We insert the closed form expression for \hat{Y} from eq. (3.26) into $E(Y, \mathbf{d})$ and obtain an energy measuring the similarity between y and $\langle I \rangle_{\sigma^{sc}, \mathbf{d}} = W_{\sigma^{sc}} I_d$

$$\begin{aligned} E_{y, I}^{data}(\mathbf{d}) &= E_{y, I_d}(\hat{Y}_d) \\ &= \frac{1}{2} \int_{\Omega} \left(y(\mathbf{x}) - \mu_Y - f \cdot (\langle I \rangle_{\sigma^{sc}, \mathbf{d}}(\mathbf{x}) - \mu_I) \right)^2 \cdot U^{\sigma^{sc}} \\ f &= C_{y, \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}} C_{\langle I_d \rangle_{\sigma^{sc}}, \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}, -1}, \quad U^{\sigma^{sc}} = C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}} \left(C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}} + C_n \right)^{-2} \end{aligned} \quad (3.27)$$

$$(3.28)$$

$E_{y, I}^{data}(\mathbf{d}^*)$ in eq. (3.27) is minimal when the specific optical flow \mathbf{d}^* satisfies

$$y(\mathbf{x}) \approx f \cdot \tilde{I}_{\mathbf{d}^*}(\mathbf{x}) + (\mu_Y - f \cdot \mu_I) \quad (3.29)$$

Comparing eq. (3.29) to the CC eq. (2.157) we can see that with $\beta = (\mu_Y - f \cdot \mu_I)$ eq. (3.29) is also a CC-type equation, however it correlates the low resolution images y and $\langle I \rangle_{\sigma^{sc}, \mathbf{d}}$.

It is easy to show that

$$C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}} = C_{Y, Y}^{\sigma^{sc}} (1 - \eta_{CC}(Y|I_d)) \quad (3.30)$$

This causes $C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}} \left(C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}} + C_n \right)^{-2}$ to additionally prune the optical flow $\mathbf{d}(\mathbf{x})$, such that $E_{y,I}^{data}$ is a more robust similarity measure than CC.

To demonstrate that our likelihood $E_{y,I}^{data}$ in eq. (3.27) respects the difference in scale between y and I we have estimated the flow with E_{Y,I_d}^{data} as the similarity measure for the data y^{syn} and I^{syn} in figure 2.5. The standard deviation σ^{sc} in $E_{y,I}^{data}$ was set to $\sigma^{sc} = 5$ and the factor f is automatically computed as $f \approx 1$ since the intensity distributions of y^{syn} and I^{syn} are approximately the same. The image I_d^{syn} is convolved with $W_{\sigma^{sc}}$. The resulting image \tilde{I}^{syn} has the same scale as y^{syn} . The resulting optical flow \mathbf{d}^{syn} is shown in figure 3.2d. Notice the blurred boundary \mathbf{d}^{syn} around the linear feature in I^{syn} (figure 3.2a). This is the result of $E_{y,I}^{data}$ in eq. (3.27) measuring the difference between y^{syn} and the blurred image $\tilde{I}_d^{syn} = W_{\sigma^{sc}} I_d$. In figure 3.2c we see $I_{d^{syn}}$. The linear boundary has been warped by \mathbf{d}^{syn} without being corrupted like in figure 2.5d.

3.5. Localization

The similarity measure $E_{y,I}^{data}$ in eq. (3.27) basically compares the image y with the transformed image $f \cdot \langle I \rangle_{\sigma^{sc}, \mathbf{d}}$, where f is defined in eq. (3.28). Thus for some minimizer of $E_{y,I}^{data}(\mathbf{d})$, \mathbf{d}^* the image $f \cdot \langle I \rangle_{\sigma^{sc}, \mathbf{d}}$ is assumed to approximate y

$$y(\mathbf{x}) \approx f \cdot \tilde{I}_{\mathbf{d}^*}(\mathbf{x}) \quad (3.31)$$

Eq. (3.31) manifests the assumption of a *global* linear relationship between the intensities of the image y and those of I , with f being the linear factor. There are many situations in multi modal optical flow where the linearity relation expressed by eq. (3.31) is not valid. For instance in figure 3.9 in section 3.7.5 the images I_{vsc} and y_{tc} of a visual spectrum camera (VSC) and a thermographic camera (TC) recording the same scene is shown for which the assumption of global linearity between I_{vsc} and y_{tc} fails, since objects that have similar reflectance properties in the visual spectrum can have differing properties in the infra-red spectrum. However if we focus on small regions $\mathcal{A} \subset \Omega$ then a local relation similar to eq. (3.31)

$$y(\mathbf{x})|_{\mathcal{A}} \approx f \cdot \tilde{I}_{\mathbf{d}^*}(\mathbf{x})|_{\mathcal{A}} \quad (3.32)$$

might hold provided that the region \mathcal{A} is occupied by the same object in both images which will be the case in section 3.7.5. Therefore we propose a local

version of the variance and the mean in eq. (2.164)

$$\begin{aligned} C_{u,v}(\mathbf{x}_0) &= \int_{\Omega} \omega(\mathbf{x} - \mathbf{x}_0) (u(\mathbf{x}) - \mu_u(\mathbf{x}_0)) \cdot (v(\mathbf{x}) - \mu_v(\mathbf{x}_0)) d^2x \quad (3.33) \\ \mu_u(\mathbf{x}_0) &= \int_{\Omega} \omega(\mathbf{x} - \mathbf{x}_0) u(\mathbf{x}) d^2x \end{aligned}$$

where ω is a window function which we take to be constant within a square window $U_{\mathbf{x}_0} \subset \Omega$ centered around the point $\mathbf{x}_0 \in \Omega$ with the window size a

$$\omega(\mathbf{x}) = \begin{cases} \frac{1}{|U_{\mathbf{x}_0}^a|} & \mathbf{x} \in U_{\mathbf{x}_0}^a, \quad \text{width}(U_{\mathbf{x}_0}^a) = \text{height}(U_{\mathbf{x}_0}^a) = a \\ 0 & \text{else} \end{cases} \quad (3.34)$$

With this definition the conditional variance $C_{Y|I_d}$ and the factor f (eq. (3.28)) become functions of the coordinates $\mathbf{x} \in \Omega$ and the parameter a

$$C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc},a}(\mathbf{x}) = C_{Y,Y}^{\sigma^{sc},a}(\mathbf{x}) - C_{y,\langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc},a,2}(\mathbf{x}) \cdot C_{\langle I_d \rangle_{\sigma^{sc}}, \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc},a,-1}(\mathbf{x}) \quad (3.35)$$

$$f^{\sigma^{sc},a}(\mathbf{x}) = C_{y,\langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc},a}(\mathbf{x}) \cdot C_{\langle I_d \rangle_{\sigma^{sc}}, \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc},a,-1}(\mathbf{x}) \quad (3.36)$$

We substitute $C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc},a}$ and $f^{\sigma^{sc},a}$ in eq. (3.27) with the local versions from eqs. 3.35 and 3.36 and obtain

$$\begin{aligned} E_{y,I_d}^{data,l}(\sigma^{sc}, a, \mathbf{d}) &= \frac{1}{2} \int_{\Omega} \left(y(\mathbf{x}) - \langle f^{\sigma^{sc},a} I_d \rangle_{\sigma^{sc}}(\mathbf{x}) \right)^2 \cdot U^{\sigma^{sc},a}(\mathbf{x}) d^2x \\ U^{\sigma^{sc},a}(\mathbf{x}) &= C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc},a}(\mathbf{x}) \left(C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc},a}(\mathbf{x}) + C_n \right)^{-2} \end{aligned} \quad (3.37)$$

Notice that the PSF $W_{\sigma^{sc}}$ is now convolved with the product $f^{\sigma^{sc},a}(\mathbf{x}) \cdot I_d(\mathbf{x})$ since $f^{\sigma^{sc},a}$ is now a function.

3.6. The Multigrid Newton algorithm

We combine now our global similarity measures in eq. (3.27) and together with the structure tensor prior E_{ST}^{prior} (eq. (3.14)) and the TV prior E_{TV}^{prior} (eq. (2.126)) to the following two models

$$E_{ST}^g(\mathbf{d}) = E_{Y,I}^{data}(\sigma^{scale}, \mathbf{d}) + \lambda_{ST} \sum_{i=1}^2 E_{ST}^{prior}(\nabla d_i) \quad (3.38)$$

$$E_{TV}^g(\mathbf{d}) = E_{Y,I}^{data}(\sigma^{scale}, \mathbf{d}) + \lambda_{TV} \sum_{i=1}^2 E_{TV}^{prior}(\nabla d_i) \quad (3.39)$$

Algorithm 1 Multigrid Optical Flow (MOF)

Initialize $\mathbf{d}_0 = 0, k = 0, \sigma_{MOF} = 0.7$
Set $\mathbf{r}_0 = \frac{\delta E(\mathbf{d})}{\delta \mathbf{d}}(\mathbf{d}_0)$
scale $s = s_{\text{Max}}$
while $s > 1$ **do**
 downsample $y_s = G_{s \cdot \sigma_{MOF}} \star y_0, I_s = G_{s \cdot \sigma_{MOF}} \star I_0$
 while $\|\mathbf{r}\| > \epsilon$ or $k < N$ **do**
 set $\mathbf{d}_{k+1} = \mathbf{d}_k + \alpha \delta$
 expand $E(\mathbf{d}_{k+1}) = E(\mathbf{d}_k) + \alpha \mathbf{b}_k^T \delta + \frac{\alpha^2}{2} \delta^T Q_k \delta$
 solve $Q_k \delta = \mathbf{b}_k$ for δ with conjugate gradients
 compute $\mathbf{d}_{k+1} = \mathbf{d}_k + \alpha \delta, k \rightarrow k + 1$
 end while
 upsample \mathbf{d}_N , set $\mathbf{d}_0 = \mathbf{d}_N, k = 0$
 $s = s - 1$
end while
set optimal solution $\mathbf{d}^* = \mathbf{d}_N$

To minimize the models in eqs. 3.38 to 3.39 and obtain the optimum flow field \mathbf{d}^* we deploy a simple newton scheme with a nested linearization (see alg. 1 (MOF) where $E(\mathbf{d})$ is to be substituted with the models in eqs. 3.38 to 3.39). The linearized model is solved by a conjugate gradients algorithm with block Jacobi preconditioning.

When minimizing the structure tensor based model E_{ST}^g (eq. (3.38)) the following major numeric problem occurs: The problem arises in step 9 of the MOF algorithm. The second functional derivative $Q_{k,ST}^g$ of the energy model E_{ST}^g consists of one part coming from the likelihood and one part coming from the prior, $Q_k^g = Q_k^{data,g} + \lambda_{ST} Q_{k,ST}^{prior}$. The matrix $Q_{k,ST}^{prior}$ is the second derivative of the prior E_{ST}^{prior} with respect to the individual components of \mathbf{d}

$$Q_{k,ST}^{prior,ii} = \frac{\delta^2}{\delta d_i^2} E_{ST}^{prior}(\nabla d_i) \quad Q_{k,ST}^{prior,ij} = 0 \quad (3.40)$$

Since $E_{ST}^{prior}(\nabla d_i)$ is purely quartic in d_i by definition (eq. (3.14)) the matrix $Q_{k,ST}^{prior}$ in eq. (3.40) a purely quadratic functional of d_i . At small k in alg 1 its eigenvalues equal zero due to the initial guess $\mathbf{d}_0 = 0$. The matrix Q_k^{data} is the second derivative of the data term $E_{Y,I}^{data}$ in eq. (3.38) and eq. (3.39). In regions where there is no motion the eigen values of Q_k^{data} are also small. This makes the linearized solution in step 9 numerically unstable and the instability carries on to all stages s of the MOF alg. 1. Our solution to this problem is to extend

$E_{ST}^{prior}(\nabla d_i)$ to include an L_2 prior on the flow field \mathbf{d} but with a small lagrange multiplier λ_2

$$\tilde{E}_{ST}^{prior}(\nabla d_i) = E_{ST}^{prior}(\nabla d_i) + \lambda_2 \|\nabla d_i\|^2 \quad (3.41)$$

With the L_2 prior in 3.41 the linearized solution in step 9 becomes numerically stable. Thus we substitute E_{ST}^{prior} for \tilde{E}_{ST}^{prior} (eq. (3.41)) in eq. (3.38)

$$E_{ST}^g(\mathbf{d}) = E_{Y,I}^{data}(\sigma^{scale}, \mathbf{d}) + \lambda_{ST} \sum_{i=1}^2 \tilde{E}_{ST}^{prior}(\nabla d_i) \quad (3.42)$$

The challenge is to find the appropriate setting for λ_2 in eq. (3.41) such that on the one side the MOF alg. 1 becomes numerically stable but on the other side the isotropic L_2 part in eq. (3.41) does not over weigh the original anisotropic prior \tilde{E}_{ST}^{prior} . In section 3.7.6 we will map out a strategy for finding a suitable value for λ_2 based on the eigen values of the second order functional derivative of E_{ST}^{prior} in eq. (3.41).

3.7. Results

In this section we want to evaluate and compare the two optical flow models in eqs. 3.38 and 3.39. The methodology for this section is as follows: In section 3.7.1 we want to study the effect of the size of the structure tensor S in eq. (3.38) on the estimation of the optical flow \mathbf{d} . We evaluate the models in eqs. 3.38 and 3.39 on a sequence of two images obtained from one camera from the Middlebury data-set [88] (figure 3.3). The Middlebury data-set offers also the ground truth optical flow \mathbf{d}_{gt} which we use to asses the quality of the estimated optical flow of the models in eqs. 3.38 and 3.39. Since the data in figure 3.3 is collected from one camera we use the global similarity measure in eq. (3.27) thereby setting $\sigma = 0$. Thus eq. (3.27) effectively reduces to the brightness constancy constraint in eq. (2.146). We compare the results of the models 3.38 and 3.39 for different window sizes σ_{ST} of the structure tensor S to find the best value for σ_{ST} . Having found the optimal window size σ_{ST} we will extend the evaluation of section 3.7.1 to include a synthesized multi-modal image set in section 3.7.4. The goal of that section is to study the effect of the PSF $W_{\sigma^{sc}}$ on the global similarity measure $E_{Y,I}^{data}(\sigma^{sc}, \mathbf{d})$ in a multi-modal coaxial setup similar to that in figure 2.6. Hereby we will create an artificial low resolution image y_{test}^{scale} by taking the image I in figure 3.3a, filtering it with a Gaussian filter of standard deviation σ_{test}^{sc} , thus creating an artificial scale difference and inverting the result. We will show that $E_{y_{test}, I}^{data}(\sigma^{sc}, \mathbf{0})$ is minimal at $\sigma^{sc} = \sigma^{sc,*}$.




Figure 3.3.: Rubberwhale Sequence: Figure 3.3a shows one frame of the sequence. figure 3.3b shows the estimated optical flow d_{ST}^* , figure 3.3c the flow d_{TV}^* and figure 3.3d shows the provided ground truth

With the optimal window size σ_{ST} for the structure tensor S and a strategy for finding the correct scale difference we turn to the problem of estimating the optical flow on the real data in section 3.7.5. As we discussed in section 3.5 the similarity measure $E_{Y,I}^{data}(\sigma^{sc}, \mathbf{d})$ is based on the assumption that the images Y and I share the linear relationship $Y = \alpha I + \beta$ which is not supported when the cameras C_Y and C_I are sensitive to different light spectra. For this reason we will deploy the local similarity measure $E_{Y,I}^{data,l}(\sigma^{sc}, \mathbf{d})$ introduced in eq. (3.37) in the localized optical flow models

$$E_{ST}(\mathbf{d}) = E_{y,I}^{data,l}(\sigma^{sc}, a, \mathbf{d}) + \lambda_{ST} E_{ST}^{prior}(\nabla \mathbf{d}) \quad (3.43)$$

$$E_{TV}(\mathbf{d}) = E_{y,I}^{data,l}(\sigma^{sc}, a, \mathbf{d}) + \lambda_{TV} E_{TV}^{prior}(\nabla \mathbf{d}) \quad (3.44)$$

We will show that our strategy in section 3.7.4 for finding the true scale difference $\sigma^{sc,*}$ between the images y and I also yields a strategy for finding the correct window size a for $E_{y,I}^{data,l}(\sigma^{sc}, a, \mathbf{d})$ such that $y = \alpha I + \beta$ locally holds.

3.7.1. Uni-Modal Data

We will now discuss the results of our optical flow method on the middlebury data set for which there exists ground truth (GT). As the GT is the true flow field for the data we use it to asses the quality of the computed optical flow. To do this we define the Endpoint error (EPE)

$$EPE(\mathbf{x}) = \|\mathbf{d} - \mathbf{d}_{gt}\|(\mathbf{x}) \quad (3.45)$$

The EPE measures how well the computed optical flow \mathbf{d} fits the true optical flow \mathbf{d}_{gt} .




Figure 3.4.: Hydrangea Sequence: Figure 3.4a shows one frame of the sequence. figure 3.4b shows the estimated optical flow d_{ST}^* , figure 3.4c the flow d_{TV}^* and figure 3.4d shows the provided ground truth

We will further need the definition of the mean curvature (eq. (2.139))

$$\kappa = \text{Div} \left(\frac{\nabla I}{|\nabla I|} \right) \quad (3.46)$$

The curvature κ is a good measure to show which features can be reliably matched by the optical flow d and thus have small EPE values.

3.7.2. Rubber Whale Sequence

Our goal in this section is to evaluate the effect of the structure tensor window size σ_{ST} of the model E_{ST} (eq. (3.38)) on the quality of the optical flow mapping d between two images Y_C and I_C recorded by *the same camera C*. We compare the results of the model E_{ST} with those obtained by E_{TV} , thereby denoting d_{ST}^* the optical flow obtained by E_{ST} and d_{TV}^* the flow obtained by E_{TV} . In figure 3.3 the rubber whale sequence of the middleburry data set is shown, and in figure 3.3d the corresponding ground truth d_{gt} . This sequence is generated with one camera recording a dynamic scene. The reason for choosing this scene is that it contains linear level-set features as well curvy-linear features such as circular features. In figure 3.3a we have highlighted a linear level-set region of interest (ROI) labeled as *Box Edge*, a ROI partially containing linear structures labeled as *Fence*, a circular feature labeled as *Wheel* and a ROI containing a generic non-linear level-set called *Shell*. In figure 3.3b the computed flow-field d_{ST}^* for the energy $E_{ST}(d)$ (eq. (3.38)) is shown for a filter size of $\sigma_{ST} = 11$, while in figure 3.3c the resulting flow for the TV model is shown. We can observe from the comparison between figures 3.3b and 3.3c that the TV model produces smoother results which are closer to the ground truth (figure 3.3d). In figure 3.5 we have binned the curvature κ in to 40 bins and plotted the average EPE per curvature bin for both the ST model in eq. (3.38) (varying σ_{ST} in figures 3.5a to 3.5c) and for the TV model in eq. (3.39) (figure 3.5d). In all plots in figure 3.5 the




Figure 3.5.: EPE to level-set curvature: Figures 3.5a to 3.5d show plots of the EPE (eq. (3.45)) against the curvature κ (eq. (3.46)) for the rubber whale sequence (figure 3.3). Figures 3.5a to 3.5c show the results for the structure tensor model E_{ST} and figure 3.5d the result for the TV model E_{TV} . The curvature κ was split into 40 bins and the height of the bars is the average EPE per curvature bin.

ROI	Filtersize	Median EPE	ROI	Filtersize	Median EPE
Wheel	7	2.36	Fence	7	0.46
	9	1.32		9	0.39
	11	1.15		11	0.35
	TV	1.38		TV	0.18
Shell	7	0.86	Box Edge	7	0.44
	9	0.62		9	0.34
	11	0.50		11	0.30
	TV	0.17		TV	0.09

Table 3.1.: EPE for different filter-sizes σ_{ST} for the model E_{ST}^g (eq. (3.38)) and for the TV model E_{TV}^g (eq. (3.39)). The value shown in the column *Median EPE* is the median EPE per ROI. The median per ROI was chosen over the average EPE per ROI due to its robustness towards outlier EPE values. The EPE values for the model E_{ST}^g decrease with increasing structure tensor filtersizes σ_{ST} . However the general trend is that the ROI's with high curvatures κ (*Wheel* and *Shell*) tend to have higher EPE values then the ROI's with low curvatures (*Fence* and *Box Edge*).

general tendency is that level-sets with low curvature values $\kappa \approx 0$ have high EPE values while with increasing curvature up to $\|\kappa\| \approx 2$ the EPE values fall off. This verifies the aperture problem (discussed in section 2.6) where less distinctive level-sets with intrinsic dimension $iD < 2$ (low κ) lead to less accurate optical flow estimates. However for higher curvature values $\|\kappa\| > 2$ the EPE values significantly rise. This is due to the fact that both the priors E_{ST}^{prior} (eq. (3.38)) and E_{TV}^{prior} (eq. (3.39)) penalize level-sets in the optical flow d which have higher curvature then level-sets which are more straight. In the case of E_{TV}^{prior} this is more evident since according to eqs. 2.139 and 2.140 in section 2.5 the curvature κ is the functional derivative of E_{TV}^{prior} and thus explicitly forced to vanish. Now comparing the plots in figures 3.5a to 3.5c we can see larger window sizes σ_{ST}

of the structure tensor in eq. (3.38) lead to small values for the EPE and thus more accurate optical flow estimates. In comparison with the TV model in eq. (3.39) (figure 3.5d) the EPE values structure tensor model E_{ST} with $\sigma_{ST} = 11$ in figure 3.5c are closest to those of the TV model in figure 3.5d with an average discrepancy of 0.3 between figure 3.5c and figure 3.5d. For this reason we will set $\sigma_{ST} = 11$ for the rest of this section.

Table 3.1 shows the median EPE values for different ROI's in figure 3.3a. We can see that ROI's with rather linear level-sets like the Fence and the Box Edge ROI have comparatively lower EPE values then ROI's containing level-set with larger curvature like the Wheel and the Shell for the E_{ST} model. For the E_{TV} only the Wheel has a higher EPE value.

3.7.3. Hydrangea Sequence

In figure 3.4a we show the hydrangea sequence of the middleburry dataset. In contrast to the rubber whale sequence in figure 3.3a figure 3.4a consists of a largely texture-less background and a hydrangea plant in the foreground. The hydrangea contains largely level-sets with high curvature κ which leads to erroneous optical flows \mathbf{d}_{ST}^* (figure 3.4b) and \mathbf{d}_{TV}^* (figure 3.4c) compared to the ground truth \mathbf{d}_{gt} in figure 3.4d. In figure 3.6 we have again plotted the EPE against the curvature κ in a fashion similar to figure 3.5. Other than in figure 3.5 we can see that for increasing window sizes $\sigma_{ST} = 7$ (figure 3.6a) to $\sigma_{ST} = 11$ (figure 3.6c) for the model E_{ST} the EPE values increase, especially at higher curvatures $\|\kappa\| > 2$. Since the background in figure 3.4a is largely constant the condition defining the level-sets S_V (see eq. (2.77))

$$S_V = \{\mathbf{x} | V_e I(\mathbf{x}) = 0\} \quad V_e = \mathbf{v} \nabla \quad (3.47)$$

is independent of the vector V_e . In other words the background in figure 3.4a contains level-set of all possible curvatures κ . This leads to an equal distribution of the EPE values over the different curvatures κ in figure 3.6c. Comparing the E_{ST} model to the TV based model E_{TV} we see that the EPE values in figure 3.6d are comparatively smaller then those for the E_{ST} model in figures 3.6a to 3.6c. The EPE values in figure 3.6d however are larger then the corresponding values in figure 3.5d for the rubber whale sequence (figure 3.3a).

3.7.4. Estimation of the Scale Difference σ^{sc}

In this section we want to analyze the dependency of the similarity measure $E_{y,I}^{data}(\sigma^{sc}, \mathbf{d})$ on the scale difference parameter σ^{sc} of the PSF $W_{\sigma^{sc}}$. First we




Figure 3.6.: EPE to level-set curvature: Figures 3.6a to 3.6d show plots of the EPE (eq. (3.45)) against the curvature κ (eq. (3.46)) for the hydrangea sequence (figure 3.4). Figures 3.6a to 3.6c show the results for the structure tensor model E_{ST} and figure 3.5d the result for the TV model E_{TV} . The curvature κ was split into 40 bins and the height of the bars is the average EPE per curvature bin.




Figure 3.7.: Synthesized multi-modal data. This data simulates the camera arrangement in figure 2.6. The image I in figure 3.7a is from the rubberwhale data set in figure 3.3. Figures 3.7b and 3.7c show the image $y_{\sigma^{sc}_{test}}$ (eq. (3.49)) at the scales $\sigma^{sc}_{test} = 2$ and $\sigma^{sc}_{test} = 4$

assume two co-aligned cameras C_y and C_I (see figure 2.6) with the images y and I being generated the following way (see figure 3.7): I is taken from the rubberwhale data set in figure 3.3. We simulate images $y_{\sigma^{sc}_{test}}$ of different resolutions $\sigma^{sc}_{test} = 1 \dots 5$ first by inverting the intensities of I followed by filter operations with Gaussians of the standard deviations $\sigma^{sc}_{test} = 1 \dots 5$ and the addition of iid noise

$$Y(\mathbf{x}) = -I(\mathbf{x}) + I_{min} + I_{max} \quad (3.48)$$

$$y_{\sigma^{sc}_{test}}^{test}(\mathbf{x}) = \langle Y \rangle_{\sigma^{sc}_{test}}(\mathbf{x}) + n(\mathbf{x}), \quad n \sim \mathcal{N}(0, \chi_{std}) \quad (3.49)$$

I_{min} and I_{max} in eq. (3.48) are the minimum respectively maximum intensity of the image I and the standard deviation χ_{std} in eq. (3.49) was set to $\chi_{std} = 50$. In figure 3.7 the image I is shown along with the synthesized images y_2 and y_4 ($y_{\sigma^{sc}_{test}}^{test}$ from eq. (3.49) with $\sigma^{sc}_{test} = 2$ resp. $\sigma^{sc}_{test} = 4$). The goal is now to show that $E_{y,I}^{data}(\sigma^{sc}, \mathbf{0})$ seen as a function of σ^{sc} is minimal at the test scale $\sigma^{sc,*} = \sigma^{sc}_{test}$.




Figure 3.8.: Figures 3.8a and 3.8b show plots the similarity measure $E_{y,I}^{data}(\sigma^{sc}, \mathbf{d})$ for the cases $y = y_2$, $\sigma^{sc}_{test} = 2$, and $y = y_4$, $\sigma^{sc}_{test} = 4$. We can observe that $E_{y,I}^{data}(\sigma^{sc}, \mathbf{d})$ is minimal with respect to σ^{sc} at the correct scales σ^{sc}_{test}

This might seem obvious since $E_{y,I}^{data}$ which was computed in eq. (3.27)

$$E_{y,I}^{data}(\sigma^{sc}, \mathbf{d}) = \frac{1}{2} \int_{\Omega} \left(y(\mathbf{x}) - f \cdot \langle I \rangle_{\sigma^{sc}}(\mathbf{x}) \right)^2 \cdot F \quad (3.50)$$

$$f = C_{y, \langle I \rangle_{\sigma^{sc}}}^{\sigma^{sc}} C_{\langle I \rangle_{\sigma^{sc}}, \langle I \rangle_{\sigma^{sc}}}^{\sigma^{sc}, -1} \quad (3.51)$$

$$F = C_{\langle Y \rangle_{\sigma^{sc}} | \langle I \rangle_{\sigma^{sc}}}^{\sigma^{sc}} \left(C_{\langle Y \rangle_{\sigma^{sc}} | \langle I \rangle_{\sigma^{sc}}}^{\sigma^{sc}} + \lambda C_n \right)^{-2} \quad (3.52)$$

was derived from the basic assumption that y is the result of the convolution of the PSF $W_{\sigma^{sc}}$ with the high resolution image Y along with additive noise (see eq. (2.160))

$$y = \langle Y \rangle_{\sigma^{sc}} + n \quad n \sim \mathcal{N}(0 | C_n) \quad (3.53)$$

which is similar to how we generated the test images $y_{\sigma^{sc}}^{test}$ in eq. (3.49). However the factor F in eq. (3.50) is highly non-linear in σ^{sc} , so we want to show that despite this non-linearity $E_{y,I}^{data}(\sigma^{sc}, \mathbf{0})$ has a global minimum at $\sigma^{sc,*} = \sigma^{sc}_{test}$. In figure 3.8 we have plotted $E_{y,I}^{data}(\sigma^{sc}, \mathbf{0})$ over σ^{sc} for the cases $\sigma^{sc}_{test} = 2$ (figure 3.8a) and $\sigma^{sc} = 4$ (figure 3.8b). For both cases $E_{y,I}^{data}(\sigma^{sc}, \mathbf{0})$ is minimal at the correct scale σ^{sc}_{test} . From this we learn that $E_{y,I}^{data}(\sigma^{sc}, \mathbf{d})$ is sensitive to the scale difference σ^{sc} between y and I and we can use it to determine the true scale difference σ^{sc}_{test} .




Figure 3.9.: figure 3.9a shows an image from a visual spectrum camera (VSC). The object recorded is a carbon-fiber reinforced polymer (CFRP). Figure 3.9b shows an image of the same CFRP recorded with a thermographic camera (TC). The TC is sensitive in the infrared domain, thus higher intensities in figure 3.9b correspond to warmer objects (the CFRP) and lower intensities to colder objects (the background). As in figure 3.1a the optical centers of the VSC and the TC are physically separated so the problem that is being addressed is that of finding the optical flow field $d(x)$ (see eq. (2.142)) which maps every pixel in the TC image to the corresponding pixel in the VSC image. Figure 3.9c shows the joint histogram of the VSC and TC image. It shows a complex mapping of the intensities of figure 3.9a to those of figure 3.9b indicating that a linearity assumption between the TC and the VSC is not valid

3.7.5. Real Multimodal Optical Flow Data

We would now like to evaluate our optical flow model in eq. (3.38) on the data in figure 3.9. The image y_{tc} (figure 3.9b) was recorded with a thermographic camera (TC) C_{tc} and the image I_{vsc} with a visual spectrum camera (VSC) C_{vsc} . The recorded object is a carbon fiber reinforced polymer (CFRP). CFRP materials are becoming increasingly widespread in automotive and aerospace industries, but also in consumer goods, due to their adaptivity to different shapes, good rigidity and high strength-to-weight ratio [89, 90, 91, 92]. Improved fabrication techniques such as Injection and Double Vacuum Assisted Resin Transfer Molding [92] are reducing the production costs and time to manufacture. The properties of CFRP strongly depend on the processing of the material, thus detection of defects within the layers of the CFRP and their characterization are indispensable, especially for safety-relevant parts. Active thermal measurement methods [93, 94, 95] have become vital for the assessment of the quality of CFRP materials. These methods are based on the evaluation of a previously excited heat flow in the tested component and its disturbance by hidden defects. The heat flow is generated with a heat pulse or through sinusoidal modulation, observed with a TC, followed by a pixel-wise computation of the complex phase between the excitation signal and the reflected infrared signal. This phase information

encodes the heat-loss within a penetration depth δ of the probed material, with depths of 1mm to 2mm typical for CFRPs.

Current state-of-the-art TCs possess resolutions of only 640×512 pixels and a noise equivalent temperature difference of 20mK. Nevertheless, these cameras are very expensive, and the CFRP application domain requires the detection of defects at the noise limit. On the other hand, VSCs are inexpensive and easily deliver images of 10 megapixels per frame with very little noise. Thus the goal is to combine low resolution TC image y_{tc} with the high frequency information borrowed from the VSC image I_{vsc} . Since the cameras C_{tc} and C_{vsc} are a) physically separated from each other and b) have different optical resolutions we could utilize our optical flow model in eq. (3.38) to estimate the optical flow mapping between C_{tc} and C_{vsc} , thereby computing a high resolution version Y_{tc} of y_{tc} . However the similarity measure $E_{y,I}^{data}(\sigma^{sc}, \mathbf{d})$ in eq. (3.38) is based on the assumption that the cameras C_y and C_I to be registered have a linear relationship in their intensity spectrum and the images y_{tc} and I_{vsc} in figure 3.9 lack this linearity relationship. This can be seen in the joint histogram in figure 3.9c where there is no unique correspondence between the intensities of y_{tc} to those of I_{vsc} . In section 3.5 we therefore introduced the local similarity measure $E_{y,I}^{data,l}(\sigma^{sc}, a, \mathbf{d})$ which constrains the similarity assessment of the images y and I to the local square regions $\mathcal{A}_{x_0}^a$ of size a centered around $x_0 \in \Omega$

$$E_{y_{tc}, I_{vsc}, \mathbf{d}}^{data,l}(\sigma^{sc}, a, \mathbf{d}) = \frac{1}{2} \int_{\Omega} \left(y_{tc} - \langle f^{\sigma^{sc}, a} I_{vsc}, \mathbf{d} \rangle_{\sigma^{sc}} \right)^2 \cdot U^{\sigma^{sc}, a} d^2x \quad (3.54)$$

$$U^{\sigma^{sc}, a}(\mathbf{x}) = C_{\langle Y \rangle_{\sigma^{sc}} | \langle I \rangle_{\sigma^{sc}}}^{\sigma^{sc}, a}(\mathbf{x}) \left(C_{\langle Y \rangle_{\sigma^{sc}} | \langle I \rangle_{\sigma^{sc}}}^{\sigma^{sc}, a}(\mathbf{x}) + \lambda C_n \right)^{-2} \quad (3.55)$$

The problem that arises is how large to set the window size a . If it is set too small the signal to noise ratio will be too small so that not enough information of the features in the y_{tc} and the I_{vsc} image are captured to robustly register them. On the other hand if a is set too large we eventually lose the local linearity between the y_{tc} and the I_{vsc} image. In order to find the correct value for the window size a we propose the following strategy: We first make sure that the cameras C_{vsc} and C_{tc} have the same aperture angle. This allows us to deduce the difference in optical scale, σ^{scale^*} directly from the resolutions of the cameras. In section 3.7.4 we showed that the global similarity measure in eq. (3.27) is convex in the optical scale parameter σ^{sc} . For fixed window sizes a the local similarity measure in eq. (3.55) is also convex in σ^{sc} but the minimum scale $\sigma_{min}^{sc}(a)$ is a function of a . The idea is to vary a such that for an optimum a^* the minimum scale $\sigma_{min}^{sc}(a^*)$ equals the true scale difference σ^{scale^*} .

In more detail: The markers on the CFRP in figure 3.9 were used to set the zoom of the lenses such that the aperture angles of C_{tc} and C_{vsc} are approximately




Figure 3.10.: ?? Plot $E_{y_{tc}, I_{vsc}}^{data,l}(\sigma^{sc}, a, \mathbf{0})$ over the PSF scale difference σ^{sc} for the images y_{tc} and I_{vsc} in figure 3.9 for the window size $a = 25$. Figure 3.10b shows the minimum scale σ_{min}^{sc} defined in eq. (3.58) as a function over the window size a and figure 3.10c the similarity measure $E_{min}^{data,l}$ (eq. (3.59)) over a . The minimum scale σ_{min}^{sc} increases or stays constant but does not decrease for larger window sizes a . The window size $a = 21$ marks a sweet spot where $\sigma_{min}^{sc}(21) = \sigma^{sc,*} = 3$ while $E_{min}^{data,l}(21)$ is comparatively minimal.

equal

$$\frac{w_{tc}}{f_{tc}} \approx \frac{w_{vsc}}{f_{vsc}} \quad (3.56)$$

where $f_{tc/vsc}$ is the focal length and $w_{tc/vsc}$ the width of the CCD in the TC/ the VSC. The true difference in optical scale $\sigma^{sc,*}$ between the cameras C_{tc} and C_{vsc} is given by the fraction of the focal lengths

$$\sigma^{sc,*} = \frac{f_{vsc}}{f_{tc}} \approx \frac{w_{vsc}}{w_{tc}} \quad (3.57)$$

Now the TC C_{tc} that was used to capture y_{tc} has a resolution of 640×480 and the VSC C_{vsc} is a Full HD camera (1920×1080) so that we have $\sigma^{sc,*} \approx \frac{1920}{640} = 3$. We make the following definitions

$$\sigma_{min}^{sc}(a) = \underset{\sigma}{\operatorname{argmin}} \left(E_{y_{tc}, I_{vsc}}^{data,l}(\sigma^{sc}, a, \mathbf{0}) \right) \quad (3.58)$$

$$E_{min}^{data,l}(a) = E_{y_{tc}, I_{vsc}}^{data,l}(\sigma_{min}^{sc}(a), a, \mathbf{0}) \quad (3.59)$$

where the minimum scale $\sigma_{min}^{sc}(a)$ is a function of the window size a and $E_{min}^{data,l}(a)$ is the value of the similarity measure $E_{y_{tc}, I_{vsc}}^{data,l}$ for that particular window size. Figure 3.10c shows the plot of $E_{min}^{data,l}$ (eq. (3.59)) over the window size a which is non-convex. Hence we cannot compute an optimal window size a^* by attempting to minimize $E_{min}^{data,l}(a)$. However $E_{y_{tc}, I_{vsc}}^{data,l}$ in eq. (3.55) is convex




Figure 3.11.: Resulting optical flows of the local models E_{ST}^l (d_{ST}^* , eq. (3.43)) and E_{TV}^l (d_{TV}^* , eq. (3.44)). We can see that the structure tensor prior in the model E_{ST}^l fails to isotropically smooth the optical flow d_{ST}^* in the regions where the images y_{tc} and I_{vsc} are predominantly homogeneous. In these regions the TV model E_{TV}^l excels due to the L_1 piecewise smoothing term in eq. (2.127).

in σ^{sc} for fixed values of a . The idea now is to find the optimal value a^* such that $\sigma^{sc}_{min}(a^*) \approx \sigma^{sc,*}$. Figure 3.10b shows $\sigma^{sc}_{min}(a)$ as plotted over different values of a . We can see that $\sigma^{sc}_{min}(a)$ increases with increasing window size a although not monotone and the true optical scale difference $\sigma^{sc}_{min}(a^*) = 3$ is reached for the window size $a^* = 21$. Thus the window size $a^* = 21$ is the size for which $y_{tc}|_{\mathcal{A}_{x_0}^{a^*}}$ and $I_{vsc}|_{\mathcal{A}_{x_0}^{a^*}}$ are expected to be linearly dependent

$$y_{tc}|_{\mathcal{A}_{x_0}^{a^*}} \approx f \cdot \langle I_{d^*,vsc} \rangle_{\sigma^{sc}}|_{\mathcal{A}_{x_0}^{a^*}} \quad (3.60)$$

With the values $a^* = 21$ and $\sigma^{sc,*} = 3$ we compute the optical flow d for the models E_{ST}^l and E_{TV}^l

$$E_{ST}^l(d) = E_{Y,I}^{data,l}(\sigma^{sc,*}, a^*, d) + \lambda_{ST} \sum_{i=1}^2 \tilde{E}_{ST}^{prior}(\nabla d_i) \quad (3.61)$$

$$E_{TV}^l(d) = E_{Y,I}^{data,l}(\sigma^{sc,*}, a^*, d) + \lambda_{TV} \sum_{i=1}^2 E_{TV}^{prior}(\nabla d_i) \quad (3.62)$$

In figure 3.11 we show the resulting optical flow d_{ST}^* and d_{TV}^* . The model E_{TV} produces a piece-wise smooth optical flow d_{TV}^* due to the piece-wise smoothing term of TV in eq. (2.127). On the other side the structure tensor model E_{ST} produces artificial motion boundaries in the regions where y_{tc} and I_{vsc} are structureless. This is due to the small weighting of the L_2 term in eq. (3.41). The value $\lambda_2 = 10^{-6}$ was chosen purely for stabilizing the structure tensor models in eq. (3.38) and eq. (3.61) in the initial iterations of the multigrid optical flow algorithm in alg. 1. However regions with structure are correctly matched by both models. To access the local linearity hypothesis in eq. (3.60) we deployed




Figure 3.12.: Comparison of the p-values (eq. (3.65)) for the hypotheses (eq. (3.68)) $H_{\hat{d}=0}$ (figure 3.12a), $H_{\hat{d}=d_{ST}^*}$ (figure 3.12b) and $H_{\hat{d}=d_{TV}^*}$ (figure 3.12c). The p-values were computed for windows \mathcal{A}_{x_0} around each pixel $x_0 \in \Omega$ and plotted over the binned values of the gradient ∇y . All three diagrams show high p-values for gradients $\nabla y \approx 0$ indicating that the structureless areas in the data in figure 3.9 obey the linear relation in eq. (3.68) regardless of the optical flow \hat{d} . For higher values of the gradient ∇y the hypothesis $H_{\hat{d}=0}$ in figure 3.12a fails as expected since the p-values tend to zero. The p-values at higher gradients for the hypotheses $H_{\hat{d}=d_{ST}^*}$ (figure 3.12b) and $H_{\hat{d}=d_{TV}^*}$ (figure 3.12c) are significantly higher than for $H_{\hat{d}=0}$ with $H_{\hat{d}=d_{TV}^*}$ having the highest p-values meaning that the total variation model E_{TV}^l in eq. (3.44) best fulfills the linearity hypothesis in eq. (3.68).

Pearson's χ^2 statistic [96].

Pearson's χ^2 statistic

Pearson's χ^2 statistic is a method of assessing whether a hypothesis H_Θ which is parameterized by the parameter set Θ is compatible a given set of observations $\hat{X}_i, 1 \leq i \leq k$. We will give a short overview of the method. Suppose the k observations \hat{X}_i are realizations of the random variables X_i . For each random variable X_i we can compute an expectation value E_i given our hypothesis H_Θ . The sum V defined by

$$V = \sum_{i=1}^k \frac{(X_i - E_i)^2}{E_i} \quad (3.63)$$

is a random variable that follows the χ^2 distribution with cumulative distribution function P

$$P(\bar{V} < v) = \int_0^v f(v') dv', \quad f(v) = \frac{v^{\frac{k}{2}-1} e^{-\frac{v}{2}}}{2^{\frac{k}{2}} \Gamma\left(\frac{k}{2}\right)} \quad (3.64)$$

where \bar{V} is a potential value the random variable V can take. $P(\bar{V} < v)$ is the probability that \bar{V} is smaller than a given number v . The probability that the number v is smaller than any value \bar{V} is computed by

$$p(v) = 1 - P(\bar{V} < v) \quad (3.65)$$

$p(v)$ is called the *p-value* of v . The method of accepting or rejecting the hypothesis H_Θ goes as follows: We can compute the observation \hat{V} from the data \hat{X}_i and the expectation values E_i generated by the hypothesis H_Θ

$$\hat{V} = \sum_{i=1}^k \frac{(\hat{X}_i - E_i)^2}{E_i} \quad (3.66)$$

If the p-value of the *particular* realization \hat{V} satisfies

$$p(\hat{V}) > \alpha \quad (3.67)$$

for some given $\alpha \in [0, 1]$ then the hypothesis H_Θ is accepted as compatible with the data \hat{X}_i .

We want to use Pearson's χ^2 test to evaluate the optical flow results \mathbf{d}_{ST}^* and \mathbf{d}_{TV}^* . The hypotheses to test are generated by the assumption that the local data term $E_{Y,I_d}^{data,l}$ in eq. (3.27) is minimal for a particular optical flow $\hat{\mathbf{d}}$

$$H_{\hat{\mathbf{d}}} : y(\mathbf{x}) \approx \mu_{\tilde{Y}|\tilde{I}_{\hat{\mathbf{d}}}}(\mathbf{x}) \quad \mu_{\tilde{Y}|\tilde{I}_{\hat{\mathbf{d}}}}(\mathbf{x}) = \langle f_{\hat{\mathbf{d}}}^{\sigma^{sc,*},a^*} \cdot I_{\hat{\mathbf{d}}} \rangle_{\sigma^{sc}}(\mathbf{x}) \quad (3.68)$$

The random variables X_i are taken to be the pixels of the CCD of the TC, C_{tc} , which can record different intensities. The particular recorded image intensities $y_{tc}(\mathbf{x})$ are interpreted as the realizations \hat{X}_i and the conditional expectation values $\mu_{\tilde{Y}|\tilde{I}_{\hat{\mathbf{d}}}}(\mathbf{x})$ are the expectations E_i of the X_i . We calculate the local observations $\hat{V}(\mathbf{x}_0)$

$$\hat{V}_{\hat{\mathbf{d}}}(\mathbf{x}_0) = \int_{\mathcal{A}_{\mathbf{x}_0}} \left(y(\mathbf{x}) - \mu_{\tilde{Y}|\tilde{I}_{\hat{\mathbf{d}}}}(\mathbf{x}) \right)^2 \cdot \mu_{\tilde{Y}|\tilde{I}_{\hat{\mathbf{d}}}}^{-1}(\mathbf{x}) d^2x \quad (3.69)$$

and compute the p-values $p(\hat{V}_{\hat{\mathbf{d}}}(\mathbf{x}_0))$ from eq. (3.64) and eq. (3.65) which we abbreviate by $p_{\hat{\mathbf{d}}}(\mathbf{x}_0)$. In figure 3.12 we have plotted the local p-values $p_{\hat{\mathbf{d}}}(\mathbf{x}_0)$ over the gradients $\|\nabla y(\mathbf{x}_0)\|$ through binning by the value of the gradient for the cases $\hat{\mathbf{d}} = \mathbf{0}$, $\hat{\mathbf{d}} = \mathbf{d}_{TV}^*$ and $\hat{\mathbf{d}} = \mathbf{d}_{ST}^*$. For regions of less structure $\|\nabla y(\mathbf{x}_0)\| \approx 0$ all three hypothesis yield high p-values such the local linearity assumption in eq. (3.68) is valid. However the p-value $p_{\mathbf{d}_0}$ converges to 0 for regions with higher gradients invalidating the hypothesis $\hat{\mathbf{d}} = \mathbf{0}$ for those regions as expected since the cameras




Figure 3.13.: Comparison of region of interests (ROI) of size $a^* = 21$. Figures 3.13b and 3.13f show a ROI of I_{vsc} and 3.13d and 3.13h the corresponding ROI of the image y_{tc} . Figures 3.13c and 3.13g show figure 3.13b warped by the flows d_{ST}^* and d_{TV}^* . 3.13a and 3.13e show the histograms between 3.13d and the filtered roi's $\tilde{I}_{vsc,d} = W_{\sigma^{sc,*}} \star I_{vsc,d}$

C_{vsc} and C_{tc} are not aligned. If we set the threshold for acceptance $\alpha = 0.5$ then the linearity hypotheses $H_{d_{TV}^*}$ and $H_{d_{ST}^*}$ (eq. (3.68)) are valid for regions with gradients $\|\nabla y_{tc}\| < 10$ since the corresponding p-values in figure 3.12c and figure 3.12b are exceed α . For higher valued gradients $\|\nabla y_{tc}\| > 10$ the p-value $p_{d_{TV}^*}$ (figure 3.12c) is slightly higher than $p_{d_{ST}^*}$ (figure 3.12b) hence the model total variation model E_{TV}^l in eq. (3.44) better fulfills the hypothesis of linearity in eq. (3.68) then the structure tensor based model E_{ST}^l in eq. (3.43) and is thus better suited for multi-modal optical flow.

Figure 3.13 shows a comparison between an ROI in the image y_{tc} (figure 3.13h) and the corresponding ROI from the image I_{vsc} (figure 3.13b) warped by d_{ST}^* (figure 3.13c) and by d_{TV}^* (figure 3.13g). The gradients within this ROI are of the order $\|\nabla y_{tc}\| \approx 10$ and hence the corresponding p-values $p_{d_{TV}^*}$ and $p_{d_{ST}^*}$ are in the accepted range, $p_{d_{TV}^*} > \alpha$ and $p_{d_{ST}^*} > \alpha$. Hence the linearity hypothesis in eq. (3.68) holds for both $\hat{d} = d_{TV}^*$ and $\hat{d} = d_{ST}^*$ and the histograms in figure 3.13e (hypothesis $\hat{d} = d_{TV}^*$) and figure 3.13a (hypothesis $\hat{d} = d_{ST}^*$) visually reflect the linear dependence of the ROI in the image y_{tc} (figure 3.13h) and the corresponding ROI from the image I_{vsc} (figure 3.13b).




Figure 3.14.: The largest eigenvalue σ_Q^k of Q^{reg} plotted over the iterations k for three values of λ_2 in eq. (3.41). Initially we have $\sigma_Q^k \approx 8\lambda_2$ which is the eigenvalue of the L_2 term in eq. (3.41). For $\lambda_2 = 10^{-3}$ we see that σ_Q^k slowly rises for increasing iterations k until at $k \approx 40$ a sudden jump occurs and σ_Q^k begins to decrease. This is the regime where the structure tensor prior E_{ST}^{prior} begins to act an-isotropically. For smaller values of λ_2 (figures 3.14b and 3.14c) the jump occurs sooner indicating quicker an-isotropic behavior of E_{ST}^{prior} .

3.7.6. Eigenvalue analysis and the stabilization parameter λ_2

In chapter 3.3 we stated that the L_2 term in eq. (3.41) is needed to support the numerical stability of the model. We will back this statement now. Figures 3.14a, 3.14b and 3.14c show the largest eigenvalue of Q_{reg}^i , σ_Q^i at each iteration on the coarsest scale of the pyramid for different values of λ_2 in the multigrid optical flow algorithm in alg. 1. The data used was the rubber whale sequence in figure 3.3 but we found similar results for the hydrangea sequence in figure 3.4 and the CFRP data in figure 3.9. The figures all show that σ_Q^N rises to a maximum after which it decreases and converges. The initial value of σ_Q^i is of the order of λ_2 indicating that in the initial steps the L_2 term in eq. (3.41) governs the regularization. As the number of iterations increases the structure tensor determinant gets more weight, until the point where its influence over weighs that of the L_2 term. For $\lambda_2 = 10^{-9}$ (figure 3.14c) the sudden jump occurs nearly instantly. Since the influence of the L_2 is negligible, the specific decaying form of σ_Q^k in figure 3.14c is an indication that E_{ST}^{prior} is smoothing an-isotropically. On the other side Figures 3.15a, 3.15b and 3.15c show the residual vector \mathbf{b} for different values of λ_2 . Comparing the magnitude of the residual vector \mathbf{b} in Figures 3.15a, 3.15b and 3.15c we see that for $\lambda_2 = 10^{-9}$, \mathbf{b} is an order of magnitude larger then the other cases, which leads to longer convergence rates or numerically instable solution. This means we have a trade-off between

- $\lambda_2 \sim 10^{-3}$: Faster convergence but less influence of structure tensor (need $i > 40$ iterations for ST to act)
- $\lambda_2 \sim 10^{-9}$: slower convergence but more influence of structure tensor




Figure 3.15.: The residual vector \mathbf{b} plotted over the iterations k for three values of λ_2 in eq. (3.41). While the norm of \mathbf{b} is approximately equal for $\lambda_2 = 10^{-3}$ and $\lambda_2 = 10^{-6}$, it is an order of magnitude higher for $\lambda_2 = 10^{-9}$. This indicates a numerical instability of the MOF algorithm for $\lambda_2 = 10^{-9}$

(need only $i > 1$ iterations for ST to act)

We choose $\lambda_2 = 10^{-6}$ since in this case \mathbf{b} is of the same order of magnitude as for $\lambda_2 = 10^{-3}$ but as we see in figure 3.14b the structure tensor only needs 4 iterations until its eigenvalues over-weigh the eigenvalues of the L_2 term.

3.7.7. Summary

In this section we introduced a prior energy $E_{ST}^{prior}(\nabla\phi)$ based on the structure tensor [87, 34]. The construction was based on the principles outlined in section 2.2, namely that E_{ST}^{prior} must be invariant under the action of the group $\mathbb{G} = \mathbb{T} \times SO(2)$ in order to preserve linear level-sets of ϕ regardless of their orientation. We deployed E_{ST}^{prior} in the context of multi-modal optical flow. In Multi-modal optical flow the task is to register two images y and I recorded by the cameras C_y and C_I on a pixel-by-pixel basis with a mapping represented by the GRF $\mathbf{d}(\mathbf{x})$. Within this context a similarity measure $E_{y,I}^{data}(\sigma^{sc}, \mathbf{d})$ which measures how similar the image y is with the warped image $I_d = I(\mathbf{x} + \mathbf{d}(\mathbf{x}))$ was introduced. The similarity measure $E_{y,I}^{data}$ is a cross correlation type measure with the feature that it can handle the situation in which the images y and I have differing optical resolutions. Essentially $E_{y,I}^{data}(\sigma^{sc}, \mathbf{d})$ measures the similarity between the low resolution image y and the filtered image $\langle I_d \rangle_{\sigma^{sc}}$ and since it is a CC type measure it is minimal when

$$y(\mathbf{x}) \approx f \langle I_{d^*} \rangle_{\sigma^{sc}}(\mathbf{x}) + \beta \quad (3.70)$$

holds for an optimal value \mathbf{d}^* of the optical flow and constants f and β .

First we combined our new similarity measure $E_{y,I}^{data}$ together with our new prior

energy E_{ST}^{prior} and with the total variation prior E_{TV}^{prior} from section 2.5 to the optical flow models

$$E_{ST}^g(\mathbf{d}) = E_{Y,I}^{data}(\sigma^{scale}, \mathbf{d}) + \lambda_{ST} \sum_{i=1}^2 E_{ST}^{prior}(\nabla d_i) \quad (3.71)$$

$$E_{TV}^g(\mathbf{d}) = E_{Y,I}^{data}(\sigma^{scale}, \mathbf{d}) + \lambda_{TV} \sum_{i=1}^2 E_{TV}^{prior}(\nabla d_i) \quad (3.72)$$

To compare E_{ST}^{prior} with the total variation prior E_{TV}^{prior} we deployed both models in eqs. (3.71) and (3.72) on the rubber whale sequence (figure 3.3) and the hydrangea sequence (figure 3.4) from the middleburry dataset [88]. Both sequences are uni-modal ($\sigma^{sc,*}=0$) and hence the similarity measure $E_{Y,I}^{data}$ reduces to an ordinary CC measure. Our findings is that the model E_{ST}^g in eq. (3.71) produces optical flows \mathbf{d}_{ST}^* with better endpoint errors (EPE) when the window size of the structure tensor σ_{ST} is sufficiently large. The best results where achieved for $\sigma_{ST} = 11$. However the TV based model in eq. (3.72) still produces an optical flow \mathbf{d}_{TV}^* with an EPE better then the best result for the model E_{ST}^g .

We went on to simulate a setup of co-aligned (no separating flow $\mathbf{d} = 0$) images y_{σ}^{test} and I^{test} which are different in optical resolution and intensity distribution. Several images y_{σ}^{test} were generated from I^{test} by inverting the intensities of I^{test} followed by Gaussian filtering with spacial standard deviations $\sigma_{test}^{sc,*} = 1 \dots 5$ and addition of iid noise, thus simulating the linearity relation in eq. (3.70). The goal was to find out if the similarity measure $E_{y_{\sigma}^{test}, I^{test}}^{data}(\sigma^{sc}, \mathbf{0})$ is capable of capturing the true scale difference $\sigma_{test}^{sc,*}$ through variation with respect to the scale parameter σ^{sc} . We showed that $E_{y_{\sigma}^{test}, I^{test}}^{data}(\sigma^{sc}, \mathbf{0})$ is convex with respect to σ^{sc} . Furthermore the minimum is equal to the true scale difference, $\sigma^{sc,*} = \sigma_{test}^{sc,*}$. Thus we concluded that given the hypothesis that two images y and I with different optical resolutions are linearly dependent (eq. (3.70)), the similarity measure $E_{y,I}^{data}(\sigma^{sc}, \mathbf{0})$ is sensitive to the true scale difference $\sigma^{sc,*}$.

An analysis of our optical flow method followed for a real world multi-modal setup. The setup included a camera rig containing a thermographic camera (TC) C_{tc} and a visual spectrum camera (VSC) C_{vsc} . Both cameras recorded an object made out of carbon fiber reinforced polymers (CFRP) producing the images y_{tc} and I_{vsc} , and the cameras are known to have a difference in optical scale of $\sigma^{sc,*} = 3$. The task was to estimate the optical flow \mathbf{d}^* between y_{tc} and I_{vsc} . The problem we encountered was that y_{tc} and I_{vsc} do *not* share a global linear relationship such as eq. (3.70), since the infra-red reflectance of the CFRP can vary across different regions of equal intensity in the visual spectrum domain. Therefore we proposed a *local* CC-type similarity measure $E_{y,I}^{data,l}(\sigma^{sc}, a, \mathbf{d})$ which is based on

the hypothesis that y_{tc} and I_{vsc} are locally linearly dependent

$$H_d^* : y_{tc}|_{\mathcal{A}_{x_0}^{a^*}} \approx f \cdot \langle I_{d^*, vsc} \rangle_{\sigma^{sc}} |_{\mathcal{A}_{x_0}^{a^*}} \quad (3.73)$$

where $\mathcal{A}_{x_0}^{a^*}$ is a window around any point $x_0 \in \Omega$ with some optimal window size a^* . We showed that the optimal scale d^* of $E_{y_{tc}, I_{vsc}}^{data, l}(\sigma^{sc}, a, \mathbf{0})$ is a function of the local window size a , $\sigma^{sc, *}: \sigma^{sc, *}(a)$. The value $a^* = 21$ was deduced from the expectation that $\sigma^{sc, *}(a^*)$ should be equal to the true optical scale difference, $\sigma^{sc, *}(a^*) = 3$. We were able to compute the optical flow between y_{tc} and I_{vsc} with the local models

$$E_{ST}(\mathbf{d}) = E_{y, I}^{data, l}(\sigma^{sc, *}, a^*, \mathbf{d}) + \lambda_{ST} E_{ST}^{prior}(\nabla \mathbf{d}) \quad (3.74)$$

$$E_{TV}(\mathbf{d}) = E_{y, I}^{data, l}(\sigma^{sc, *}, a^*, \mathbf{d}) + \lambda_{TV} E_{TV}^{prior}(\nabla \mathbf{d}) \quad (3.75)$$

To analyze the validity of the local linearity hypothesis H_d^* in eq. (3.73) for the computed flows $d^* = d_{ST}^*$ (from eq. (3.74)) and $d^* = d_{TV}^*$ (eq. (3.75)) we deployed Pearson's χ^2 test. The p-value $p_d^*(x)$ in Pearson's χ^2 test is an indicator for the validity of the hypothesis H_d^* such that $p_d^*(x) = 1$ means that H_d^* is definitely valid at the point x and $p_d^*(x) = 0$ means that H_d^* must be rejected.

The p-values $p_{d_{ST}^*}(x)$ and $p_{d_{TV}^*}(x)$ in regions with small gradients $\|y_{tc}\| < 10$ were sufficient, $p_{d_{ST}^*}/p_{d_{TV}^*} \approx 1$, to accept the linearity hypotheses $H_{d_{ST}^*}$ and $H_{d_{TV}^*}$. For higher gradients $\|y_{tc}\| > 10$ both p-values dropped off below the 50% quantile, however with $p_{d_{TV}^*}$ being slightly larger than $p_{d_{ST}^*}$. Thus although both models in eq. (3.75) and eq. (3.74) are not fully consistent with the linearity hypothesis in eq. (3.73) at the boundaries of the CFRP, the TV model in eq. (3.75) produced an optical flow d_{TV}^* which is more consistent with the data y_{tc} and I_{vsc} at the boundaries than the structure tensor based model in eq. (3.74).

The good p-values in the regions $\|y_{tc}\| < 10$ do not give any information about how consistent the computed flows d_{ST}^* (from eq. (3.74)) and d_{TV}^* (eq. (3.75)) are themselves. This is why we visually compared d_{ST}^* and d_{TV}^* to each other: The flow d_{ST}^* had a lot of artificial linear boundaries in regions where y_{tc} and I_{vsc} are largely constant. In contrast the flow d_{TV}^* was smooth everywhere except at the physical boundary locations of the CFRP. This is due to the behavior of $E_{ST}^{prior}(\nabla \mathbf{d})$, namely that it only penalizes the curvature of the level-sets of \mathbf{d} . On the other hand the TV prior E_{TV}^{prior} not only penalizes the curvature of the level-sets but also enforces smoothness of the solution d_{TV}^* in regions where the images y_{tc} and I_{vsc} are not discontinuous. Thus we conclude two things: First the TV prior E_{TV}^{prior} is superior to our structure tensor based prior E_{ST}^{prior} for multi-modal optical flow since the TV model in eq. (3.75) produces visually more consistent optical flows. Secondly the local linearity hypothesis in eq. (3.73)

which emerges from our proposed similarity measure $E_{y_{tc}, I_{vsc}}^{data, l}(\sigma^{sc}, a, \mathbf{d})$ holds very well in regions where the images y_{tc} and I_{vsc} vary only gradually with no discontinuities, but must be rejected at the discontinuities of y_{tc} and I_{vsc} . This behavior is presumably due to the lack of information about the physics of the TC in the similarity measure $E_{y_{tc}, I_{vsc}}^{data, l}(\sigma^{sc}, a, \mathbf{d})$ as we had only encoded the basic assumption that TC produces Gaussian noise (see eq. (2.160) in section 2.7). Including a more realistic noise distribution into $E_{y_{tc}, I_{vsc}}^{data, l}$ will generally lead to a (possibly non-linear) relation between y_{tc} and I_{vsc} which is better suited to estimate the optical flow at the boundaries of the CFRP.

4. The Extended Least Action Algorithm

In section 2.1.1 we had reviewed the principle of least action (PLA) of an arbitrary Gibbs random field (GRF) theory which is modeled by the energy

$$E(\phi, \mathbf{X}_e^\Omega \phi) = E^{data}(\phi) + E^{prior}(\mathbf{X}_e^\Omega \phi) = \int_{\Omega} \mathcal{E}_{tot}(\phi(\mathbf{x}), \mathbf{X}_e^\Omega \phi(\mathbf{x})) d^2x \quad (4.1)$$

where the prior $E^{prior}(\mathbf{X}_e^\Omega \phi)$ depends on the operators $X_e^{\Omega,i} \in \mathcal{B}$ of a commutative sub algebra $\mathcal{B} \subset \mathcal{G}$ of some arbitrary Lie algebra \mathcal{G} . In short the PLA states that the minimizer ϕ^* of the energy E in eq. (4.1) is the value of the GRF ϕ which obeys the Euler-Lagrange equations

$$[\mathcal{E}_{tot}] (\phi(\mathbf{x})) |_{\phi=\phi^*} = 0, \quad [\mathcal{E}_{tot}] (\phi(\mathbf{x})) = \frac{\delta \mathcal{E}}{\delta \phi(\mathbf{x})} - \text{Div}(\mathbf{P}) \quad (4.2)$$

$$\text{Div}(\mathbf{P}) = \sum_{\mu} \frac{d}{dx_{\mu}} \left(\sum_i \omega_{\Omega,i}^{\mu} P_i \right), \quad P_i = \frac{\delta \mathcal{E}^{prior}}{\delta X_e^{\Omega,i}} \quad (4.3)$$

Section 2.4.2 was dedicated to the analysis of the energy $E(\phi, \mathbf{X}_e^\Omega \phi)$ under the action of an r -dimensional finite Lie Group $\mathbb{G} = \mathbb{G}^{\phi} \times \mathbb{G}^{\Omega}$ which contains transformations of the GRF ϕ and of the coordinate frame Ω . Noether's theorem in eq. (2.115) states that if the energy E in eq. (4.1) is invariant under the Lie group $\mathbb{G} = \mathbb{G}^{\phi} \times \mathbb{G}^{\Omega}$ and if \mathbb{G}^{ϕ} and \mathbb{G}^{Ω} are independent from each other, then E must be invariant under \mathbb{G}^{Ω} . Then the canonical momentum \mathbf{P} and hence the prior energy E^{prior} must fulfill eq. (2.120) for any value of the GRF ϕ

$$\widetilde{\mathbf{B}}\phi = 0, \quad \widetilde{\mathbf{B}} = \sum_i P_i [X_e^{\Omega,i}, X_e^{\Omega,m}] \quad (4.4)$$

Eq. (4.4) stems from the group \mathbb{G}^{Ω} , the spatial transformations of the domain Ω , while the Euler-Lagrange equations in eq. (4.2) stem from the group \mathbb{G}^{ϕ} , the functional transformations of the GRF ϕ . The coupling of the Euler-Lagrange differentials $[\mathcal{E}]$ and the product $\widetilde{\mathbf{B}}\phi$ in eq. (2.114) suggest that eq. (4.4) and thus \mathbb{G}^{Ω} impact the solutions space of the Euler-Lagrange equations in eq. (4.2). Indeed




Figure 4.1.: Figure 4.1a shows an image ϕ_0 with parabolic level-sets according to eq. (4.9). The white line indicates the level-sets $S_{\phi_0,c}$ with $39 < c < 43$. In figure 4.1b the coordinate frame Ω_0 is shown together with the level-sets $S_{\phi_0,c}$. Figure 4.1c shows the warped image $\phi_0(T_t^B \circ \mathbf{x})$ and figure 4.1d the transformed coordinate frame $\tilde{\Omega} = T_t^B \circ \Omega_0$. $\tilde{\Omega}$ has been deformed by the algorithm in eq. (4.10) in such a way that the level-set $S_{\phi_0,c}$ (indicated by the black line) appears to be straight and hence it is identified with the linear domain Ω^ϵ of the TV prior $E_{TV}^{prior}(\nabla\phi)$.

we will show that there exists an optimal deformation of the coordinate frame Ω which serves as an aid to solving the Euler-Lagrange equation in eq. (4.2). The deformation algorithm will exploit the non-trivial symmetries of the energy in eq. (4.1) in order to narrow down the space of possible solutions to the Euler-Lagrange equations in eq. (4.2). We will evaluate its effectiveness for both the TV prior E_{TV}^{prior} which we reviewed in section 2.5 and the structure tensor prior E_{ST}^{prior} which we introduced in section 3.

The Basic Idea

We will start the motivation of the new algorithm in the following way: As we outlined in section 2.3 the role of the prior energy $E^{prior}(X_e^{\Omega,i}\phi)$ in eq. (4.1) is to set a constraint on the form of the level-sets of the image ϕ . By constraint we mean the particular geometrical form of the level-sets S_α^ϵ of the minimizers $\phi^* \in A$ (A was defined in eq. (2.34)). This form is defined by the coordinate transformation $\epsilon(\mathbf{x})$ which obeys the Cauchy-Riemann equations in eq. (2.82). The coordinate transformation $\epsilon(\mathbf{x})$ simply transforms the level-set S_α^ϵ to the linear domain Ω^ϵ (the image of $\epsilon(\mathbf{x})$, see eq. (2.80)) where S_α^ϵ appears to be a line with orientation vector α .

Two examples of prior energies previously introduced, The structure tensor based prior energy in eq. (3.14) in section 3.3 and the total variation (TV) prior in eq. (2.126), are constructed from the Lie algebra \mathfrak{t} of the translation group \mathbb{T} which is spanned by the Cartesian operators $\{\partial_x, \partial_y\}$. Thus the level-sets S of their

minimizers are lines (zero curvature κ (eq. (2.138))) in the Cartesian coordinate frame Ω .

Our idea is to find a method for obtaining a transformation *of the coordinate frame* Ω , T_t^B such that given an *arbitrary* image $\phi_0(\mathbf{x})$ the level-sets of the *warped* image $\phi_0(T_t^B \circ \mathbf{x})$ satisfies the geometrical constraints imposed by E^{prior} , specifically such that the level-sets

$$S = \left\{ \mathbf{x}' \left| \sum_i \alpha_i X_e^{\Omega, i} \phi(\mathbf{x}') = 0 \right. \right\}, \quad \mathbf{x}' = T_t^B \circ \mathbf{x} \quad (4.5)$$

are linear in the linear domain Ω^e of the prior $E^{prior}(X_e^{\Omega, i} \phi)$ (see figure 4.1). In general it is not possible to derive T_t^B analytically since an analytical expression for the level-sets of an arbitrary image ϕ_0 is not available. We have therefore devised a method to iteratively estimate T_t^B as the result of a flow equation on the domain Ω

$$\dot{\mathbf{x}}(t) = \mathbf{b}(\mathbf{P}, \mathbf{x}(t)) \quad (4.6)$$

where \mathbf{P} is the canonical momentum (eq. (4.2)). Then the action of T_t^B on Ω can be understood as the integration of $\dot{\mathbf{x}}(t)$ in eq. (4.6) over the flow parameter t

$$T_t^B \circ \mathbf{x} = \int_0^t \mathbf{b}(\mathbf{P}, \mathbf{x}(t)) dt \quad (4.7)$$

We will later derive the explicit form of \mathbf{b} from eq. (2.119). For now we only need to know that \mathbf{b} only depends on the prior energy E^{prior} and *not* on the data term E^{data} . Therefore we consider the total energy $E(\phi, \nabla \phi)$ in eq. (4.1) only to consist of the TV prior E_{TV}^{prior}

$$E(\phi, \nabla \phi) = E_{TV}^{prior}(\nabla \phi) \quad (4.8)$$

What follows is a display how the flow equation in eq. (4.6) helps to (partially) minimize E_{TV}^{prior} . Consider the image $\phi_0(\mathbf{x})$ in figure 4.1a. It's level-sets have the shape of a parabola

$$S_{\phi_0, y_0} = \left\{ \mathbf{x} \left| y = (x - x_0)^2 + y_0, \quad (x, y)^T = \mathbf{x}, y_0 \in \mathbb{R} \right. \right\} \quad (4.9)$$

where x_0 is the x -component of the center pixel of ϕ_0 and y_0 is an offset in y -direction. The white line in figure 4.1a indicates the bundle of level-sets S_{ϕ_0, y_0} with $39 < y_0 < 43$, it is not part of the actual image. We implemented a simple algorithm that solves the integration problem in eq. (4.7) by splitting the time domain $\{0, T\}$ into N time steps t_n , $0 \leq n \leq N$ and iteratively computing a new

coordinate frame Ω_{n+1} given an old estimate Ω_n

$$\Omega_{n+1} : \quad \mathbf{x}_{n+1} = \mathbf{x}_n + \tau_\Omega \mathbf{b}(\mathbf{P}, \mathbf{x}_n), \quad \mathbf{x}_n \in \Omega_n \quad (4.10)$$

where Ω_n is the transformed coordinate frame at time step t_n . The initial coordinate frame Ω_0 (figure 4.1b) is the Cartesian coordinate frame and the black line in figure 4.1b indicates the level-set S_{ϕ_0, y_0} . The simple procedure in eq. (4.10) deforms the coordinate frame Ω_0 under the influence of the prior E_{TV}^{prior} such that it assumes the form $\tilde{\Omega} = T_t^B \circ \Omega_0$ in figure 4.1d after $N = 2500$ iterations. In figure 4.1c the image $\phi_0(T_t^B \circ \mathbf{x})$ is shown which is the result of transforming the original image $\phi_0(\mathbf{x})$ (figure 4.1a) to the frame $\tilde{\Omega}$ in figure 4.1d. Within this frame the level-sets of ϕ_0 appear to be straight lines since the curvature κ of the level-sets of ϕ_0 are penalized by the TV prior E_{TV}^{prior} . Hence the domain $\tilde{\Omega}$ can be identified with the linear domain $\tilde{\Omega}^\epsilon$ of the TV prior E_{TV}^{prior} . In figure 4.1d we can see that the transformed domain $\tilde{\Omega}$ has been curved in the opposite direction and thus $\tilde{\Omega}$ has negative curvature $-\kappa$. Hence the curvature of ϕ_0 in figure 4.1a is not *cancelled* but merely inverted and deferred to the coordinate frame Ω .

4.0.8. Newtonian Minimization

One of the basic algorithms for solving the Euler-Lagrange equations in eq. (4.2) is the method of steepest descent (see [6] for an overview of gradient methods in image processing). The basic idea of steepest descent is to view the minimizers $\phi^* \in A$ as the result of a flow equation driven by the Euler-Lagrange differentials $[\mathcal{E}_{tot}]$

$$\dot{\phi}(t, \mathbf{x}_0) = - [\mathcal{E}_{tot}] (\phi(t, \mathbf{x}_0)) \quad (4.11)$$

such that for $t \rightarrow t^*$ where the limit t^* might be infinite, $\phi(t, \mathbf{x}_0)$ converges to $\phi^*(\mathbf{x}_0)$

$$\dot{\phi}(t, \mathbf{x}_0) \Big|_{t \rightarrow t^*} = 0 \implies \phi(t, \mathbf{x}_0) \Big|_{t \rightarrow t^*} = \phi^*(\mathbf{x}_0) \quad (4.12)$$

In practical implementations we discretize the interval $[0, t^*]$ into N time steps t^n and identify the GRF ϕ at the the different time steps by $\phi^n(\mathbf{x}_0) = \phi(t^n, \mathbf{x}_0)$. Starting with an initial guess ϕ^0 , we compute a new estimate of the field ϕ by advancing a previous estimate ϕ^n along the negative direction of the gradient of $E(\phi, \mathbf{X}_e^\Omega \phi)$ which is provided by the Euler-Lagrange differentials $[\mathcal{E}_{tot}]$

$$\phi^{n+1}(\mathbf{x}_0) = \phi^n(\mathbf{x}_0) - \tau^\phi [\mathcal{E}_{tot}] (\phi^n(\mathbf{x}_0)) \quad (4.13)$$

The scheme is repeated (see algorithm 2, Basic Newton Algorithm (BNA)) until either the Euler-Lagrange differentials vanish or the fixed number N of iterations is reached.

Our new methodology is to combine the concept of steepest descent for the spacial coordinate frame Ω from eq. (4.6) with the concept of steepest descent for the image ϕ (in eq. (4.11))

$$\dot{\phi}(t, \mathbf{x}(t)) = - [\mathcal{E}_{tot}] (\phi(t, \mathbf{x}(t))) \quad (4.14)$$

$$\dot{\mathbf{x}}(t) = \mathbf{b}(\mathbf{P}(t), \mathbf{x}(t)) \quad (4.15)$$

where the canonical momentum \mathbf{P} depends on the current state $(\phi(t, \mathbf{x}(t)))$, hence the time dependence. The combination of the discretized eqs. (4.10) and (4.13) is straightforward

$$\phi^{n+1}(\mathbf{x}_n) = \phi^n(\mathbf{x}_n) - \tau^\phi [\mathcal{E}_{tot}] (\phi^n(\mathbf{x}_n)) \quad (4.16)$$

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \tau_\Omega \mathbf{b}(\mathbf{P}_n, \mathbf{x}_n) \quad (4.17)$$

The exact form of the vector $\mathbf{b}(\mathbf{x})$ will soon be deduced, now we wish to give an intuitive idea of $\mathbf{b}(\mathbf{x})$ should look like. For this we will discuss the influence of Noether's Theorem on the level-sets S . From this discussion it will follow that the VVF $\mathbf{b}(\mathbf{x}(t))$ is the set of tangential vectors to the level-sets of ϕ at time t .

4.0.9. The dynamics of the level-sets S

Noethers Theorem states that if the energy functional E is invariant under a Lie group \mathbb{G} of dimension r , then there exists r identities relating the Euler-Lagrange differentials $[\mathcal{E}]$ (eq. (4.45)) and the divergences of r vector valued functions \mathbf{W}_m (eq. (2.111))

$$\left. \frac{d}{dt} g_{\gamma_t} \circ E \right|_{t=0} = \int_{\Omega} \sum_m \alpha_m \left(\text{div}(\mathbf{W}_m) + \tilde{\omega}_m^\phi [\mathcal{E}_{tot}] \right) d^2x = 0 \quad (4.18)$$

$$\tilde{\omega}_m^\phi = \omega_m^\phi - X_e^{m,\Omega}(\phi) \quad (4.19)$$

for any one-parameter subgroup $g_{\gamma_t} \subset \mathbb{G}$. If \mathbb{G} is a pure spacial Lie group, $\omega_m^\phi = 0$, then eq. (4.18) simplifies to

$$\left. \frac{d}{dt} g_{\gamma_t} \circ E \right|_{t=0} = \int_{\Omega} \sum_m \alpha_m \left(\text{div}(\mathbf{W}_m) - X_e^{m,\Omega}(\phi) [\mathcal{E}_{tot}] \right) d^2x = 0 \quad (4.20)$$




Figure 4.2.: This figure shows a transformation of the level-set S to S' along the vector $\mathbf{W}_m(\mathbf{x})$. The region $\mathcal{A} \subset \Omega$ is the region a section of S traverses as it is shifted along \mathbf{W}_m to the end position S' . If the divergence of \mathbf{W}_m vanishes, this means that the incoming flux of \mathbf{W}_m equals the outgoing flux (both indicated by the red arrows), $\mathbf{W}_m|_S = \mathbf{W}_m|_{S'}$

which is independent of the integration region. Since eq. (4.20) must hold for any coefficient vector α , the individual divergences must satisfy

$$\int_{\Omega} (\text{div}(\mathbf{W}_m)) d^2x = \int_{\Omega} \left(X_e^{m,\Omega}(\phi) [\mathcal{E}_{tot}] \right) d^2x \quad (4.21)$$

Thus on the set $\phi^* \in A$ the VVFs \mathbf{W}_m are divergence free.

Dynamics of the normal vector n_S

Eq. (4.21) must hold for any integration domain Ω which means that the integrands themselves must be equal

$$\text{div}(\mathbf{W}_m) = X_e^{m,\Omega}(\phi) [\mathcal{E}_{tot}] \quad (4.22)$$

By Gauss' law the integrated divergence of \mathbf{W}_m within any subset $\mathcal{A} \subset \Omega$ equals the integral of the flux of \mathbf{W}_m over the surface $\partial\mathcal{A}$

$$\int_{\mathcal{A}} \text{div}(\mathbf{W}_m) d^2x = \int_{\partial\mathcal{A}} \mathbf{W}_m d\mathbf{n}_S \quad (4.23)$$

where \mathbf{n} is the normal vector on the surface $\partial\mathcal{A}$. Thus from eq. (4.22) we have

$$\int_{\partial\mathcal{A}} \mathbf{W}_m d\mathbf{n}_S = \int_{\mathcal{A}} \left(X_e^{m,\Omega}(\phi) [\mathcal{E}_{tot}] \right) d^2x \quad (4.24)$$

In figure 4.2 we have depicted the situation where a level-set S is shifted along the vector \mathbf{W}_m with S' being the result of the shift and \mathcal{A} is the region traversed by the shift of a section of S . We denote this transformation by $S \rightarrow S'$. The

boundary $\partial\mathcal{A}$ consists of two lines tangential to \mathbf{W}_m besides the sections of S and S' . Since the flux over the tangential lines vanishes we have

$$\int_S \mathbf{W}_m d\mathbf{n}_S - \int_{S'} \mathbf{W}_m d\mathbf{n}_{S'} = \int_{\mathcal{A}} X_e^{m,\Omega}(\phi) [\mathcal{E}_{tot}] dx^2 \quad (4.25)$$

From Eq. (4.25) we see that the Euler-Lagrange differentials $[\mathcal{E}_{tot}]$ and the basis element $X_m^{m,\Omega}$ act as a *source* that drives the transformation $S \rightarrow S'$ in the sense that the level-set S propagates until it traverses a region in which the integrand of the right hand side in eq. (4.25) vanishes. More precisely eq. (4.25) can be interpreted as an equation of motion for the *normal vector* on S , \mathbf{n}_S .

Dynamics of the tangential vector to S

The goal of this section is to show that eq. (4.25) does *not suffice* explain the motion of the level-sets S in the BNA (algorithm 2). We will derive an equation for the normalized *tangential* vectors of the level-sets S , which we label $\mathbf{b}_S(\mathbf{x})$. In fact the vector $\mathbf{b}(\mathbf{x}(t))$ which was introduced in eq. (4.6) happens to be equal to \mathbf{b}_S evaluated at $\mathbf{x}(t)$

$$\mathbf{b}(\mathbf{x}(t)) = \mathbf{b}_S(\mathbf{x})|_{\mathbf{x}=\mathbf{x}(t)} \quad (4.26)$$

Since \mathbf{b}_S is the tangential vector to S we can write the level-set equation for S in the form of an operator equation

$$B\phi = 0, \quad B = b_S^\mu(\mathbf{x}) \partial_\mu \quad (4.27)$$

If we define B to be the basis operator of the algebra of a one parameter Lie group T_t^B (see eq. (2.55))

$$\left. \frac{d}{dt} T_t^B \circ \phi \right|_{t=0} = B\phi = b_S^\mu \partial_\mu \phi = 0 \quad (4.28)$$

then the image $\phi(\mathbf{x})$ is an invariant T_t^B . At first glance it seems that the operator B can not achieve what we set out to do in the flow equation in eq. (4.6) (and figure 4.1). And indeed we are led to believe that the energy E in eq. (4.1) is also invariant to the transformation T_t^B

$$\left. \frac{d}{dt} T_t^B \circ E(\phi, \mathbf{X}_e^\Omega \phi) \right|_{t=0} = 0 \quad (4.29)$$

since it depends on ϕ which in turn obeys eq. (4.28). Noether's theorem in eq. (2.111) holds independently of the image ϕ , hence eq. (4.29) would have to




Figure 4.3.: Figure 4.3a shows an image ϕ in which a group of level-sets with $47 < \phi < 53$ (indicated by the white area) all converge into one point P at the top of the image. The line sections $S_{1,2}$ and $T_{1,2}$ enclose the region \mathcal{A}_B in eq. (4.35) and eq. (4.37). The normal vectors $\mathbf{n}_{S_{1,2}}$ of lines $S_{1,2}$ are orthogonal to \mathbf{b}_S , hence the corresponding line integrals on the right hand side of eq. (4.35) vanish. In contrast the normal vectors $\mathbf{n}_{T_{1,2}}$ of lines $T_{1,2}$ are parallel to \mathbf{b}_S so that the corresponding line integrals on the right hand side of eq. (4.35) do not vanish

hold for any image ϕ . We will show that eq. (4.29) cannot hold for at least one image ϕ if eq. (4.27) holds *uniformly* for all $\mathbf{x} \in \Omega$ by contradiction:

Assume eq. (4.29) to hold for T_t^B . Then by Noether's theorem there must exist a vector valued function $\mathbf{W}_B(\mathbf{x})$ for which similar to eq. (4.21) the identity

$$\int_{\Omega} (\operatorname{div}(\mathbf{W}_B)) d^2x = \int_{\Omega} (B(\phi) [\mathcal{E}_{tot}]) d^2x \quad (4.30)$$

must hold. However due to eq. (4.28) the integrand of the right hand side of eq. (4.30) vanishes uniformly on Ω , leading to the conclusion that the divergence of $\mathbf{W}_B(\mathbf{x})$ must vanish uniformly too

$$\operatorname{div}(\mathbf{W}_B)(\mathbf{x}) = 0, \quad \forall \mathbf{x} \in \Omega \quad (4.31)$$

The fallacy in our argumentation lies in the last statement in eq. (4.31). As we will see eq. (4.31) cannot hold: We recall the generic form of the divergence vector \mathbf{W}_m from eq. (2.113)

$$W_m^\mu = \omega_m^\mu \mathcal{E} + \sum_i \omega_i^\mu \left(\omega_m^\phi - X_e^{m,\Omega}(\phi) \right) \frac{\delta \mathcal{E}}{\delta (X_e^i \phi)} \quad (4.32)$$

and replace $X_e^{m,\Omega}$ with B as the generating operator, thereby setting $\omega_m^\phi = 0$ (B is a purely spacial operator)

$$W_B^\mu = b_S^\mu \mathcal{E} - \sum_i \omega_i^\mu \cdot B(\phi) \frac{\delta \mathcal{E}}{\delta (X_e^i \phi)} \quad (4.33)$$

Again by virtue of eq. (4.28) \mathbf{W}_B simplifies even more

$$W_B^\mu = b_S^\mu \mathcal{E} \quad (4.34)$$

To show why $\text{div}(\mathbf{W}_B) = 0$ cannot hold we have constructed an image ϕ (figure 4.3a) in which a group of level-sets indicated by a white region converge into one point P . We construct an integral over a region \mathcal{A}_B which is enclosed by two level-sets S_1 and S_2 (the blue lines in figure 4.3a) and two curves T_1 and T_2 (the red lines in figure 4.3a) connecting the level-sets

$$\begin{aligned} \int_{\mathcal{A}_B} \text{div}(\mathbf{W}_B) d^2x &= \int_{S_1} \mathcal{E} \cdot (\mathbf{b}_S \cdot \mathbf{n}_{S_1}) dS_1 + \int_{S_2} \mathcal{E} \cdot (\mathbf{b}_S \cdot \mathbf{n}_{S_2}) dS_2 \\ &\quad + \int_{T_1} \mathcal{E} \cdot (\mathbf{b}_S \cdot \mathbf{n}_{T_1}) dT_1 + \int_{T_2} \mathcal{E} \cdot (\mathbf{b}_S \cdot \mathbf{n}_{T_2}) dT_2 \end{aligned} \quad (4.35)$$

The arrows in figure 4.3a indicate the orientation of the line integral on the right hand side of eq. (4.35). Since the vector valued function $\mathbf{b}_S(\mathbf{x})$ is the tangent vector on the level-sets S_1 and S_2 the corresponding line integrals in eq. (4.35) vanish. Furthermore T_1 and T_2 have opposite orientation so we can choose the gauge

$$(\mathbf{b}_S|_{T_1} \cdot \mathbf{n}_{T_1}) = +1, \quad (\mathbf{b}_S|_{T_2} \cdot \mathbf{n}_{T_2}) = -1 \quad (4.36)$$

Thus we have for the divergence

$$\int_{\mathcal{A}_B} \text{div}(\mathbf{W}_B) d^2x = \int_{T_1} \mathcal{E}(\phi(\mathbf{x}), \mathbf{X}_e^\Omega \phi(\mathbf{x})) dT_1 - \int_{T_2} \mathcal{E}(\phi(\mathbf{x}), \mathbf{X}_e^\Omega \phi(\mathbf{x})) dT_2 \quad (4.37)$$

If $\text{div}(\mathbf{W}_B) = 0$ was to be true then the two line integrals in eq. (4.37) would have to cancel. However we did not make any assumptions on the length of the curves T_i or on the distance between them so eq. (4.37) must hold for any configuration of the T_i . If we push T_1 to the point P where it has zero length, $\|T_1\| = 0$, then eq. (4.37) simplifies to

$$0 = - \int_{T_2} \mathcal{E}(\phi(\mathbf{x}), \mathbf{X}_e^\Omega \phi(\mathbf{x})) dT_2 \quad (4.38)$$

since the T_1 integral in eq. (4.37) vanishes (integrals over sets of zero measure vanish). Eq. (4.38) holds for any configuration of T_2 with non zero length. Hence we conclude that the energy density $\mathcal{E}(\phi(\mathbf{x}), \mathbf{X}_e^\Omega \phi(\mathbf{x}))$ must vanish on Ω . In other words: There cannot exist a non zero energy density which is invariant under T_t^B given the image ϕ in figure 4.3a. We have derived a clear contradiction. It

follows that the energy E cannot be an invariant of the transformation T_t^B !

$$\left. \frac{d}{dt} T_t^B \circ E \left(\phi, \mathbf{X}_e^\Omega \phi \right) \right|_{t=0} \neq 0 \quad (4.39)$$

The question that remains is that if T_t^B does not change the values of $\phi(\mathbf{x})$ what does it change? In eq. (4.6) we already gave a hint: T_t^B is an operator on the coordinates \mathbf{x} themselves

$$\mathbf{x}(t) = T_t^B \circ \mathbf{x}_0 \quad (4.40)$$

The constraint in eq. (4.28) guaranties that the level-set of $\phi(\mathbf{x}(t))$ move along with the flow in eq. (4.40). The image ϕ just appears to be transformed if we view it from an absolute reference frame Ω_0

$$\tilde{\phi}(\mathbf{x}_0) = \phi \left(T_t^B \circ \mathbf{x}_0 \right), \quad \mathbf{x}_0 \in \Omega_0 \quad (4.41)$$

However the particular reference frame Ω_0 is irrelevant as two frames Ω and Ω' may always be connected by the flow in eq. (4.40).

4.1. The Extended Least Action Algorithm

In this section we will compute the exact form of the operator B . The two core aspects of the operator B as motivated in the previous section are: B must be a level-set operator (eq. (4.27)) and furthermore it must be the generator of a one parameter flow equation on the coordinate frame Ω in the sense of eq. (4.40). Thus we are looking for a level-set operator which comes from a spacial transformation. In eq. (2.104) we studied how the energy E from eq. (4.1) transforms under an arbitrary sub group $g_\gamma(t) \subset \mathbb{G}$

$$\left. \frac{d}{dt} g_\gamma(t) \circ E \right|_{t=0} = \int_\Omega \left(\sum_i \left[V_e^\Omega, X_e^{\Omega,i} \right] (\phi) \cdot P_i + v^\phi [\mathcal{E}_{tot}] \right) d^2x = 0 \quad (4.42)$$

$$= \int_\Omega \left(\sum_m \alpha_m \tilde{\mathbf{B}}_m \phi + v^\phi [\mathcal{E}_{tot}] \right) d^2x = 0 \quad (4.43)$$

$$\tilde{\mathbf{B}}_m = \sum_{i=1}^q P_i \left[X_e^{m,\Omega}, X_e^{i,\Omega} \right], \quad P_i \left(\mathbf{X}_e^\Omega \phi(\mathbf{x}) \right) = \frac{\delta \mathcal{E}_{prior}}{\delta \left(X_e^{\Omega,i} \phi \right)} \quad (4.44)$$

As we explained in section 2.4 the variation to the field ϕ , v^ϕ and the operator $V_e^\Omega \phi$ which comes from the variation of the spacial coordinates \mathbf{x} are taken to be independent from each other. Thus for eq. (4.43) to hold both the Euler-Lagrange

differentials $[\mathcal{E}_{tot}]$ and the product $\widetilde{\mathbf{B}}_m \phi$ must vanish

$$[\mathcal{E}_{tot}](\phi^*) = 0 \quad (4.45)$$

$$\widetilde{\mathbf{B}}_m \phi = 0 \quad (4.46)$$

While eq. (4.45) only holds for the minimizers ϕ^* , eq. (4.46) holds according to section 2.4.3 for any image ϕ if the energy E is invariant under pure spacial transformations induced by V_e^Ω . Furthermore eq. (4.46) must hold independently from the coefficient vector α , since eq. (4.43) must hold for any $g_{\gamma_t} \in \mathbb{G}$. We will show now that eq. (4.46) defines the level-set operator B . Since the commutator in eq. (4.43) is an element of the Lie algebra \mathcal{G} , by eq. (2.67) it can be represented in terms of the basis X_e^i

$$[X_e^{m,\Omega}, X_e^{i,\Omega}] = \sum_l C_{mi}^l X_e^{\Omega,l}(\phi) \quad (4.47)$$

As the basis elements $X_e^{\Omega,l}$ are represented by the Cartesian gradient operator ∇

$$X_e^{\Omega,l} = \omega_l^\mu(\mathbf{x}) \partial_\mu \quad (4.48)$$

the operator $\widetilde{\mathbf{B}}_m$ in eq. (4.43) takes the form

$$\widetilde{\mathbf{B}}_m = \|\widetilde{\mathbf{b}}_m\| \mathbf{B}_m \phi(\mathbf{x}), \quad \widetilde{b}_m^\mu(\mathbf{x}) = \sum_i P_i(\phi(\mathbf{x})) C_{mi}^l \omega_l^\mu(\mathbf{x}) \quad (4.49)$$

$$\mathbf{B}_m = b_m^\mu(\mathbf{x}) \partial_\mu, \quad b_m^\mu = \|\widetilde{\mathbf{b}}_m\|^{-1} \widetilde{b}_m^\mu \quad (4.50)$$

We would like to discuss the normalized operator \mathbf{B}_m . By eq. (4.49) and eq. (4.46) the operator \mathbf{B}_m must satisfy the level-set equation

$$\mathbf{B}_m \phi = b_m^\mu \partial_\mu \phi = 0 \quad (4.51)$$

The r operators \mathbf{B}_m are the basis of an algebra \mathcal{G}^B and an associated group \mathbb{G}^B . For reasons soon to be clear we call the group \mathbb{G}^B the *bending group* and the algebra \mathcal{G}^B the *bending algebra*. The dimension r_B of \mathcal{G}^B is not necessarily equal to the dimension r of the original algebra \mathcal{G} , since $\mathbf{B}_m \neq 0$ holds only for the non trivial elements of \mathbb{G} . For instance in the case $\mathbb{G} = \mathbb{T} \times SO(2)$ only the $SO(2)$ group allows for the construction of a bending group $SO^B(2)$ since E is trivially invariant under \mathbb{T} . It is the r_B non trivial operators $\mathbf{B}_m \neq 0$ which we use to define the diffusion equation for the coordinate frame Ω from eq. (4.40)

$$\mathbf{x}(t) = T_t^B \circ \mathbf{x}, \quad \left. \frac{d\mathbf{x}(t)}{dt} \right|_{t=0} = \sum_{m=1}^r \beta_m \mathbf{B}_m \mathbf{x} = \mathbf{B} \cdot \mathbf{x} \quad (4.52)$$

The operator \mathbf{B} is a linear combination of the r operators \mathbf{B}_m from eq. (4.50) and the coefficient vector β is an arbitrary constant vector. Since all $\mathbf{B}_m \neq 0$ are level-set operators (eq. (4.51)) the linear combination \mathbf{B} is also a level-set operator. The normalization in eq. (4.50) is needed for the gauge conditions in eq. (4.36) to hold. Thus from the arguments leading to eq. (4.39) the diffusion equation in eq. (4.52) must minimize the total energy E in eq. (4.1). The combination of the flow driven by the bending operator \mathbf{B} in eq. (4.52) and the flow driven by the Euler-Lagrange differentials $[\mathcal{E}_{tot}]$ in eq. (4.11)

$$\dot{\phi}(t, \mathbf{x}(t)) = -[\mathcal{E}_{tot}] (\phi(t, \mathbf{x}(t))) \quad (4.53)$$

$$\dot{\mathbf{x}}(t) = \mathbf{B} \cdot \mathbf{x}(t) \quad (4.54)$$

can be interpreted as the defining differential equations of a transformation $T(t) = (T_t^\phi, T_t^B)$ which maps the initial guess ϕ_0 to a minimizer of total energy E in eq. (4.1)

$$\phi(\mathbf{x}_0, t) = T_t^\phi \circ \phi_0(\mathbf{x}(t)), \quad \mathbf{x}(t) = T_t^B \circ \mathbf{x}_0, \quad \mathbf{x}_0 \in \Omega_0 \quad (4.55)$$

$$\phi_{tot}^*(\mathbf{x}_0) = \lim_{t \rightarrow \infty} \phi(\mathbf{x}_0, t) \quad (4.56)$$

The bending flow in eq. (4.54) *cannot* be deduced from eq. (4.53). Hence it provides an *extension* to the principle of least action.

The Curvature Operator K

The diffusion process in eq. (4.52) is a non-linear process since the coefficient vector $\mathbf{b}_m(\mathbf{x})$ itself (eq. (4.50)) is a function of the coordinates $\mathbf{x}(t)$ and the GRF ϕ . It is guided along those operators \mathbf{B}_m which do not vanish due to trivial symmetries. We had motivated in figure 4.1 that the diffusion equation in eq. (4.52) bends the coordinate frame so that the level-sets of the image ϕ appear to respect the geometric constraints imposed by the prior energy E^{prior} . The geometric constraints imposed by E^{prior} are constraints on the curvature κ of the level-sets S . For instance the TV prior E_{TV}^{prior} (eq. (2.126)) penalizes level-sets with non-zero mean curvature (eq. (2.139)). The geometric constraints imposed by E^{prior} are fulfilled when the functional derivative of E^{prior}

$$[\mathcal{E}^{prior}] (\phi) = -\text{Div} \mathbf{P} (\phi) \quad (4.57)$$

vanishes. In section 4.0.9 we had shown that the energy E is not invariant to the flow in eq. (4.52) and is thus minimized. However if E is minimized, then the diffusion process must also have an effect on the Euler-Lagrange differentials in eq. (2.116), especially on the divergence of the canonical momentum \mathbf{P} in




Figure 4.4.: Effect of the diffusion $\mathbf{x}' = g^B(t) \circ \mathbf{x}$ (eq. (4.52)) on the canonical momentum \mathbf{P} . Figure 4.4a shows a schematic picture of a region $\mathcal{A}_B \subset \Omega$ between two level-sets S_1 and S_2 . The canonical momentum (the vectors on the level-sets $S_{1,2}$) is denoted by $\mathbf{P}_{S_{1,2}}$. \mathbf{P} changes its orientation when shifted along the level-sets S_1 and S_2 since the norm of the curvature operator \mathbf{K} (eq. (4.59)) is non zero. In figure 4.4b the level-sets $S_{1,2}$ have been deformed according to $\mathbf{x}' = g^B(t) \circ \mathbf{x}$ such that the canonical momentum \mathbf{P} becomes invariant with respect to shifts along $S_{1,2}$. In this case the norm of the curvature operator \mathbf{K} vanishes

eq. (4.57). The effect of T_t^B on $\text{Div} \mathbf{P}$ is interesting for the reason that \mathbf{P} has a relation to the gradient $\nabla \phi$ and thus to the level-sets S , for instance for the TV prior E_{TV}^{prior} we will show later that \mathbf{P} is parallel to $\nabla \phi$. If we know how $\text{Div} \mathbf{P}$ changes under the flow in eq. (4.52), we will know how \mathbf{P} and thus the level-sets S deform under eq. (4.52).

The rate of change of $\text{Div} \mathbf{P}$ can be computed by considering the integral

$$F = \int_{\Omega} \frac{d}{dt} T_t^B \circ \text{Div} \mathbf{P}(\mathbf{x}) \Big|_{t=0} \phi(\mathbf{x}) \mathcal{N} d^2 x \quad (4.58)$$

where $\mathcal{N} d^2 x$ is the volume measure preserved under the flow in eq. (4.52).

Eq. (4.58) can be transformed into an integral over an operator on ϕ (see appendix B.1)

$$F = \int_{\Omega} \mathbf{K} \phi(\mathbf{x}) \mathcal{N} d^2 x, \quad \mathbf{K} = [\mathbf{B}, P_i X_e^i] \quad (4.59)$$

We call the operator \mathbf{K} the *curvature operator* for two reasons. First it represents the mixed second order derivate of the energy E , since E was first functionally derived by its argument ϕ then by the flow parameter t from eq. (4.52). Second and importantly \mathbf{K} relates to a geometrical curvature: Since \mathbf{B} is a level-set operator eq. (4.27), \mathbf{K} describes the change the vector \mathbf{P} undergoes when being shifted along the level-sets S . In figure 4.4 we have schematically depicted the action of T_t^B on the canonical momentum \mathbf{P} . Figure 4.4a shows a region $\mathcal{A}_B \subset \Omega$ which is foliated by level-sets ranging from S_1 to S_2 . The vector \mathbf{P} varies upon

Algorithm 2 Basic Newton Algorithm (BNA)

Set $k = 0$
Set Initial guess ϕ^0
Compute residual $r^k = -[\mathcal{E}_{tot}] (\phi^k)$
while $\|r\| > \delta$ and $k < N$ **do**
 Compute $\phi^{k+1}(\mathbf{x}) = \phi^k(\mathbf{x}) - \tau^\phi [\mathcal{E}_{tot}] (\phi^k(\mathbf{x}))$
 Recompute $r^{k+1} = -[\mathcal{E}_{tot}] (\phi^{k+1})$
 Store $E^{k+1} = E(\phi^{k+1}, \nabla\phi^{k+1})$ in a vector $\{E^k\}$
 Set $k \rightarrow k + 1$
end while

Algorithm 3 Diffusion Algorithm (DA)

Set $k = 0$
Set Initial guess ϕ^0, \mathbf{x}^0
while $k < N$ **do**
 Compute $\mathbf{x}^{k+1} = \mathbf{x}^k - \tau^\Omega \mathbf{b}(\mathbf{x}^k)$
 Compute $\phi^{k+1}(\mathbf{x}) = \phi^0(\mathbf{x}^{k+1})$
 Store $E^{k+1} = E(\phi^{k+1}, \nabla\phi^{k+1})$ in a vector $\{E^k\}$
 Set $k \rightarrow k + 1$
end while

Algorithm 4 Extended Least Action Algorithm (ELAA)

Set $k = 0$
Set Initial guess ϕ^0, \mathbf{x}^0
Compute residual $r^k = -[\mathcal{E}_{tot}] (\phi^k)$
while $\|r\| > \delta$ and $k < N$ **do**
 Compute $\mathbf{x}^{k+1} = \mathbf{x}^k - \tau^\Omega \mathbf{b}(\mathbf{x}^k)$
 Compute $\phi^{k+1}(\mathbf{x}^{k+1}) = \phi^k(\mathbf{x}^{k+1}) - \tau^\phi [\mathcal{E}_{tot}] (\phi^k(\mathbf{x}^k))$
 Recompute $r^{k+1} = -[\mathcal{E}_{tot}] (\phi^{k+1})$
 Store $E^{k+1} = E(\phi^{k+1}, \nabla\phi^{k+1})$ in a vector $\{E^k\}$
 Set $k \rightarrow k + 1$
end while

shifts along the level-sets $S_{1,2}$. Hence by eq. (4.59) the norm of the curvature operator \mathbf{K} is non-zero. In figure 4.4b the region \mathcal{A}_B has been deformed to \mathcal{A}'_B via the diffusion in eq. (4.52). The level-sets $S_{1,2}$ have been deformed such that $P_{S_{1,2}}$ is constant along $S_{1,2}$. Thus the norm of the curvature operator \mathbf{K} vanishes.




Figure 4.5.: Figure 4.5a shows a picture ϕ^c of a person. ϕ^c is taken to be free of noise. Figure 4.5b is a noise corrupted version of ϕ^c in figure 4.5a, $\phi^d = \phi^c + n$ where n is iid Gaussian noise with a standard deviation $\sigma = 100$. Figure 4.5c shows the result of the ELAA (alg. 4) and figure 4.5d the result of the BNA (alg. 2)

For prior energies E^{prior} for which the canonical momentum P is parallel to the gradient $\nabla\phi$, the vanishing of the norm of the curvature operator K , $\|K\| = 0$ means that the gradient $\nabla\phi$ itself is invariant with respect to the diffusion process eq. (4.52), hence the level-sets S of ϕ are lines.

4.1.1. Image De-noising

In section 2.1.2 we had described the problem of noise contamination of the image \hat{I}^c of an object O recorded by the camera C . The image \hat{I}^c is modeled as the sum of the projection of the object O onto the image plane of C , I_O and additive noise drawn from a distribution p

$$\hat{I}_{ij}^c = I_O(x_{i,j}) + n \quad n \sim p\left(I_{ij}^c | I_O(x_{i,j})\right) \quad (4.60)$$

The problem is that we would like infer the projection I_O given the data \hat{I}^c and knowledge of the distribution p in eq. (4.60). However this inference problem is ill-defined and to make it well-defined we need to consider the geometrical properties of the object O . In eq. (2.19) we reformulated the problem of inference of I_O as a minimization problem

$$I_O^* = \operatorname{argmin}_{I_O} (E_{I^c}(I_O)), \quad E_{I^c}(I_O) = E_{I^c}^{data}(I_O) + E^{prior}\left(\mathbf{X}_e^\Omega I_O\right) \quad (4.61)$$

where the image I_O is considered to be a GRF for which the geometrical neighborhood properties are specified by the prior energy E^{prior} and the noise distribution p in eq. (4.60) is connected to the data term $E_{I^c}^{data}$ in eq. (4.61) via the exponential mapping in eq. (2.5).

Algorithm 5 Image de-noising analysis

Select ϕ_0 from image database, $i = 0, \sigma = 100$

while $i < 100$ **do**

 Sample noise disturbed image $\phi = \phi_0 + n, n \sim \mathcal{N}(0, \sigma)$

 Run BNA, DA or ELAA and obtain vector of energies $\{E^k\}$

 Store in matrix $\mathbf{E}, E_{i,k} = E^k$

end while

For each iteration k compute the expectation value $\langle E^k \rangle$ and the standard deviation σ_{E^k} from the k -th column vector of \mathbf{E}

The goal of this section is to evaluate the Extended Least Action Algorithm (ELAA) (alg. 4) for the inference problem in eq. (4.61). For this we will assume the situation where the noise distribution p in eq. (4.60) is Gaussian

$$E(\phi) = \frac{1}{2} \int_{\Omega} |I_O - I^c|^2 dx + E^{prior}(\mathbf{X}_e^{\Omega} I_O) \quad (4.62)$$

We have run both the Basic Newton (alg. 2) and our Extended Least Action algorithm (alg. 4) to minimize the energy in eq. (4.62) for the total variation prior and our new structure tensor based prior.

Analysis Method

We minimized the image denoising model in eq. (4.62) for both the TV-Prior and the structure tensor prior using the Basic Newton Algorithm (BNA) in alg. (2), the Diffusion Algorithm (DA) in alg. (3) and the Extended Least Action Algorithm (ELAA) in alg. (4). All three algorithms were analyzed with statistical analysis algorithm (SAA) in alg. (5). The SAA samples Gaussian noise at a standard deviation of $\sigma = 100$ and adds it to the image ϕ_0 . Then it runs the BNA, DA, and the ELAA. The energy E^k at each iteration of the BNA, DA and ELAA is stored in a vector. This procedure is repeated 100 times so that for each iteration k of all algorithms we have 100 sample energies E^k . Then the mean energy per iteration k , $\langle E^k \rangle$ and the standard deviation σ_{E^k} are computed. In the same manner we computed the mean $\langle \|\mathbf{K}\| \rangle$ and the standard deviation $\sigma_{\|\mathbf{K}\|}$ of the norm $\|\mathbf{K}\|$. The whole procedure was repeated on a total of ten images of the middlebury data set [88]. In figure 4.5 we show the Army image of the middlebury data set together with the results of the BNA and the ELAA, and in section B.2 of the appendix we show the results of the other nine images.




Figure 4.6.: Figure 4.6a shows the mean energy $\langle E^k \rangle$ and figure 4.6b the standard deviation σ_{E^k} per iteration k for the Army image in figure 4.5a. The the ELAA (solid line) converges about twice as fast as the BNA (dashed line) according to figure 4.6a. The standard deviation σ_{E^k} in figure 4.6b converges approximately three times faster for the ELAA then for the BNA indicating that the ELAA is robuster to noise at every iteration k

Total Variation based Image De-Noising

The TV based image de-noising model is defined by the energy

$$E(I_O, \nabla I_O) = \frac{1}{2} \int_{\Omega} |I_O - I^c|^2 d^2x + \int_{\Omega} \mathcal{E}^{prior}(\nabla I_O(\mathbf{x})) d^2x \quad (4.63)$$

$$\mathcal{E}^{prior}(\nabla I_O(\mathbf{x})) = \lambda \sqrt{\nabla^T I_O \cdot \nabla I_O}, \quad \mathbf{P} = \lambda \frac{\nabla I_O}{\|\nabla I_O\|} \quad (4.64)$$

The prior \mathcal{E}^{prior} in eq. (4.64) is an invariant of the Lie group $\mathbb{G} = \mathbb{T} \times SO(2)$, the group of translation and rotations. However the associated bending operator $B_{\mathbb{T}}$ vanishes for the translation group \mathbb{T} vanishes since

$$B_{\mathbb{T}}^x = P^\nu [\partial_\nu, \partial_x] = 0, \quad B_{\mathbb{T}}^y = P^\nu [\partial_\nu, \partial_y] = 0 \quad (4.65)$$

that is \mathbb{T} is a trivial symmetry of \mathcal{E}^{prior} and $E(I_O, \nabla I_O)$. The bending operator B_θ associated with the rotation group $SO(2)$ does not vanish, but computes to

$$B_\theta = b_\theta^\mu \partial_\mu, \quad \mathbf{b}_\theta = \frac{\nabla^\perp I_O}{\sqrt{\nabla^T I_O \cdot \nabla I_O}} = \mathbf{P}^\perp \quad (4.66)$$

We have run the statistical analysis algorithm (SAA) in alg. (5) on the Army image in figure 4.5a and in figure 4.6 we have plotted the results. In figure 4.6a




Figure 4.7.: Figure 4.7a shows the mean curvature $\langle \|\mathbf{K}\|^k \rangle$ and figure 4.7b the standard deviation $\sigma_{\|\mathbf{K}\|^k}$ per iteration k for the Army image in figure 4.5a. For the DA (dotted line), which only depends on the TV prior E_{TV}^{prior} , $\langle \|\mathbf{K}\| \rangle$ has an exponential decay. For the ELAA (solid line) $\langle \|\mathbf{K}\| \rangle$ drops faster then for the DA, until a point where the data term E^{data} prohibits further smoothing of the level-sets S . Then $\langle \|\mathbf{K}\| \rangle$ rises slightly and converges at a higher value. The BNA falls off slower then the ELAA and the DA and converging at a slightly higher value then the ELAA. The standard deviation $\sigma_{\|\mathbf{K}\|}$ is for both the ELAA and the BNA comparatively of equal order and small and two orders of magnitude smaller then $\langle \|\mathbf{K}\| \rangle$. When comparing the ELAA and the BNA to the DA (dotted line) we can see that the data term E^{data} has an impact on the noise distribution of the curvature $\|\mathbf{K}\|$ particularly at later iterations $k > 100$. Figure 4.7c shows a fit of the exponential function in eq. (4.67) to the curvature of the DA algorithm. The difference between the DA (solid line) and the fit (dashed line) is of the order 10^4 , an order of magnitude smaller then $\|\mathbf{K}\|$

the mean energy $\langle E^k \rangle$ per iteration k is plotted for the BNA (dashed line), the DA (dotted line) and the ELAA (solid line). The DA which only depends on the prior E_{TV}^{prior} converges the slowest. However the ELAA, which is a combination of the DA and the BNA, converges approximately twice as fast as for the BNA and several times faster then the DA. As for the standard deviation σ_{E^k} (figure 4.6b) we see that the ELAA converges more than twice as fast then the BNA. σ_{E^k} is a measure for how robust the solution I_O^k at iteration k is with respect to noise. Thus we conclude that the ELAA is at every iteration robuster to noise then the original BNA. In figure 4.7 we show the results of the SAA for the curvature $\|\mathbf{K}\|$. The curvature for the DA follows an exponential behavior. This is expected since by the definition of the mean curvature κ in eq. (2.139) and the definition of the curvature operator \mathbf{K} in eq. (4.59) we have

$$\int_{\Omega} \frac{d}{dt} (g_B^t \circ \kappa) \Big|_{t=0} I_O d^2x = \int_{\Omega} \mathbf{K} I_O d^2x \quad (4.67)$$

The left hand side of eq. (4.67) is the rate of change of the curvature κ with respect to the parameter t of the diffusion process in eq. (4.52) and the right hand side is linear in the curvature operator \mathbf{K} . Thus the curvature κ must have an

exponential dependency on t

$$g_B^t \circ \kappa = \alpha \exp(-\beta t) + \gamma \quad (4.68)$$

We did a least squares fit of the parameters (α, β, γ) in eq. (4.68) to the curvature $\|\mathbf{K}\|$ of the DA shown in figure 4.7a. The estimated parameters are $\alpha = 3.46 \cdot 10^5$, $\beta = 0.0044$ and $\gamma = 1.23 \cdot 10^5$. The error of the fit is of the order 10^3 , which is two order of magnitude smaller than $\|\mathbf{K}\|$. Hence we conclude two things: first the interpretation of the curvature operator \mathbf{K} as the curvature of the level-sets is valid. Second the evolution of the curvature of the level-sets under the diffusion process in eq. (4.52) follows an exponential law.

Structure Tensor Prior

In this section we applied our structure tensor prior from eq. (3.14) in section 3.2 to the image de-noising problem

$$E(I_O, \nabla I_O) = \frac{1}{q} \int_{\Omega} |I_O - I^c|^q d^2x + \int_{\Omega} \mathcal{E}_{ST}^{prior}(\nabla I_O(\mathbf{x})) d^2x \quad (4.69)$$

$$\mathcal{E}_{ST}^{prior}(\nabla I_O(\mathbf{x})) = \frac{1}{2} \det(S) \quad (4.70)$$

In order to apply the ELAA in alg. 4 to the model in eq. (4.70) we need to compute the coefficient vector \mathbf{b} of the bending operator \mathbf{B} . From section 3.2 we know that \mathcal{E}_{ST}^{prior} is invariant to the group $\mathbb{G} = \mathbb{T} \times SO(2)$. Like the TV prior the translation group \mathbb{T} is a trivial symmetry so that it suffices to compute the bending operator \mathbf{B}_{θ} corresponding to the group $SO(2)$. We remember from eq. (3.17) that the structure tensor prior in eq. (4.70) transforms under the $SO(2)$ in the following way

$$\left. \frac{d}{d\theta} E_{ST}^{prior}(S_{\theta}) \right|_{\theta=0} = \int_{\Omega} \text{Tr}(\mathbf{B}_{ST} \cdot S) d^2x = 0, \quad \mathbf{B}_{ST} = \mathbf{P}^{ST} \cdot \mathbf{M}_{\theta} \quad (4.71)$$

where \mathbf{M}_{θ} is the Pauli matrix (the generator of the algebra $\mathfrak{so}(2)$, eq. (2.74)). The matrix \mathbf{B}_{ST} has a striking similarity to the bending operator \mathbf{B}_m in eq. (4.50) since it is a product of the canonical momentum \mathbf{P}^{ST} and the structure constants of the $SO(2)$, the matrix \mathbf{M}_{θ} . The trace

$$\text{Tr}(\mathbf{B}_{ST} \cdot S) = B_{ST}^{\mu,\nu} S_{\mu,\nu} = 0 \quad (4.72)$$

is a scalar product which runs over two indexes (we used the Einstein summation convention) as opposed to scalar products between two vectors. Hence eq. (4.72)




Figure 4.8.: Figure 4.8a shows the mean energy $\langle E \rangle$ as a function of the iteration k for the ELAA (solid line) and the DA (dotted line) for the structure tensor model. Figure 4.8b shows a close up of $\langle E \rangle_{ELAA}$ for $k \geq 10$ and figure 4.8c shows the difference between $\langle E \rangle_{BNA}$ and $\langle E \rangle_{ELAA}$. From figure 4.8b we can see that the mean energy for the ELAA $\langle E \rangle_{ELAA}$ is 2 orders of magnitude smaller than the mean energy for the DA and by figure 4.8c only slightly smaller than $\langle E \rangle_{BNA}$. Thus the effect of the diffusion process in eq. (4.52) on the minimization of the energy E in eq. (4.70) is at most marginal

can be seen as the level-set equation for the structure tensor S , much like eq. (4.27) is the level-set equation for the image ϕ .

In order to deploy the ELAA in alg. 4 we need to transform the level-set equation of the structure tensor S into a level-set equation like eq. (4.27) with an operator \mathbf{B}_{ST} , since the diffusion equation in eq. (4.52) necessitates a differential operator of the form in eq. (4.50).

$$\int_{\Omega} \text{Tr}(\mathbf{B}_{ST} \cdot S) d^2x = \int_{\Omega} \mathbf{B}_{ST} I_O d^2x \quad (4.73)$$

If we insert the definition of the structure tensor from eq. (3.5), $S = \langle \nabla I_O \nabla^T I_O \rangle_{\sigma_{ST}}$ into the left hand side of eq. (4.73) we can shift the convolution operation in S onto the matrix \mathbf{B}_{ST} and use the cyclic periodicity property of the trace to isolate the gradient ∇I_O

$$\begin{aligned} \int_{\Omega} \text{Tr}(\mathbf{B}_{ST} \cdot S) d^2x &= \int_{\Omega} \text{Tr}(\mathbf{B}_{ST} \cdot \langle \nabla I_O \nabla^T I_O \rangle_{\sigma_{ST}}) d^2x \\ &= \int_{\Omega} \text{Tr}(\langle \mathbf{B}_{ST} \rangle_{\sigma_{ST}} \cdot \nabla I_O \nabla^T I_O) d^2x = \int_{\Omega} \nabla^T I_O \langle \mathbf{B}_{ST} \rangle_{\sigma_{ST}} \nabla I_O d^2x \end{aligned} \quad (4.74)$$

from eq. (4.74) we can read off the form of \mathbf{B}_{ST}

$$\mathbf{B}_{ST} = \frac{1}{\|\mathbf{b}_{ST}\|} b_{ST}^{\mu} \partial_{\mu}, \quad \mathbf{b}_{ST} = \langle \mathbf{B}_{ST} \rangle_{\sigma_{ST}} \nabla I_O \quad (4.75)$$

The matrix $\langle \mathbf{B}_{ST} \rangle_{\sigma_{ST}}$ in eq. (4.75) is the convolution of the elements of \mathbf{B}_{ST} with




Figure 4.9.: Figure 4.9a shows the standard deviation σ_E as a function of the iteration k for the ELAA (solid line) and the DA (dotted line) for the structure tensor model. Figure 4.9b shows a close up of $\sigma_{E,ELAA}$ for $k \geq 10$ and figure 4.9c shows the difference between $\sigma_{E,BNA}$ and $\sigma_{E,ELAA}$. We essentially see the same behavior for the standard deviation σ_E as for the mean energy in figure 4.8: By figure 4.9b the standard deviation energy for the ELAA $\sigma_{E,ELAA}$ is 1 order of magnitude smaller than that of the DA and by figure 4.9c only slightly smaller than $\sigma_{E,BNA}$. Hence the diffusion process eq. (4.52) has a marginal contribution to the statistical robustness of the minimizers of E in eq. (4.70)

the weight function $w(x)$ from the definition of the structure tensor S in eq. (3.5). Due to eq. (4.72) the operator B_{ST} is also a level-set operator in the sense of eq. (4.27).

In figure 4.8a the energies for the ELAA and the DA algorithm are shown. We can see that the energy $\langle E \rangle_{DA}$ hardly converges at the same rate as the energy $\langle E \rangle_{ELAA}$. $\langle E \rangle_{DA}$ stays within the range of $\langle E \rangle_{DA} \sim 3 \cdot 10^{10}$ while from figure 4.8b ($\langle E \rangle_{ELAA}$ for $k \geq 10$) we can see that $\langle E \rangle_{ELAA}$ drop down by 3 orders of magnitude. Figure 4.8c shows the difference between the mean energies of the BNA and the ELAA. The difference $\langle E \rangle_{BNA} - \langle E \rangle_{ELAA}$ is only in the range of 10^4 which is 4 orders of magnitude smaller than the absolute value of $\langle E \rangle_{ELAA}$ in figure 4.8b. Thus although the diffusion process in eq. (4.52) has a positive impact on the ELAA, this impact is insignificant compared to the impact of eq. (4.52) on the TV-Denoising model (figure 4.6a). The explanation is that the structure tensor prior E_{ST}^{prior} effectively only measures the curvature up to the scale determined by the window size σ_{ST} . Loosely speaking, since E_{ST}^{prior} involves the weighted integral of the gradient $\nabla I_O \nabla^T I_O$ over a local neighborhood of size σ_{ST} , level-sets S with higher curvature are integrated out and hence do not contribute to the total energy. To show this we have evaluated the structure tensor model in eq. (4.70) with the ELAA for various window sizes σ_{ST} in figure 4.10. In figure 4.10c the initial energy $\langle E^k \rangle_{ELAA}$ ($k = 0$) and in figure 4.10d the initial curvature $\langle \|K\|^k \rangle$ ($k = 0$) are shown for different window sizes σ_{ST} . For window sizes $\sigma_{ST} \geq 13$ we can see that the energy $\langle E^0 \rangle_{ELAA}$ rises while the curvature

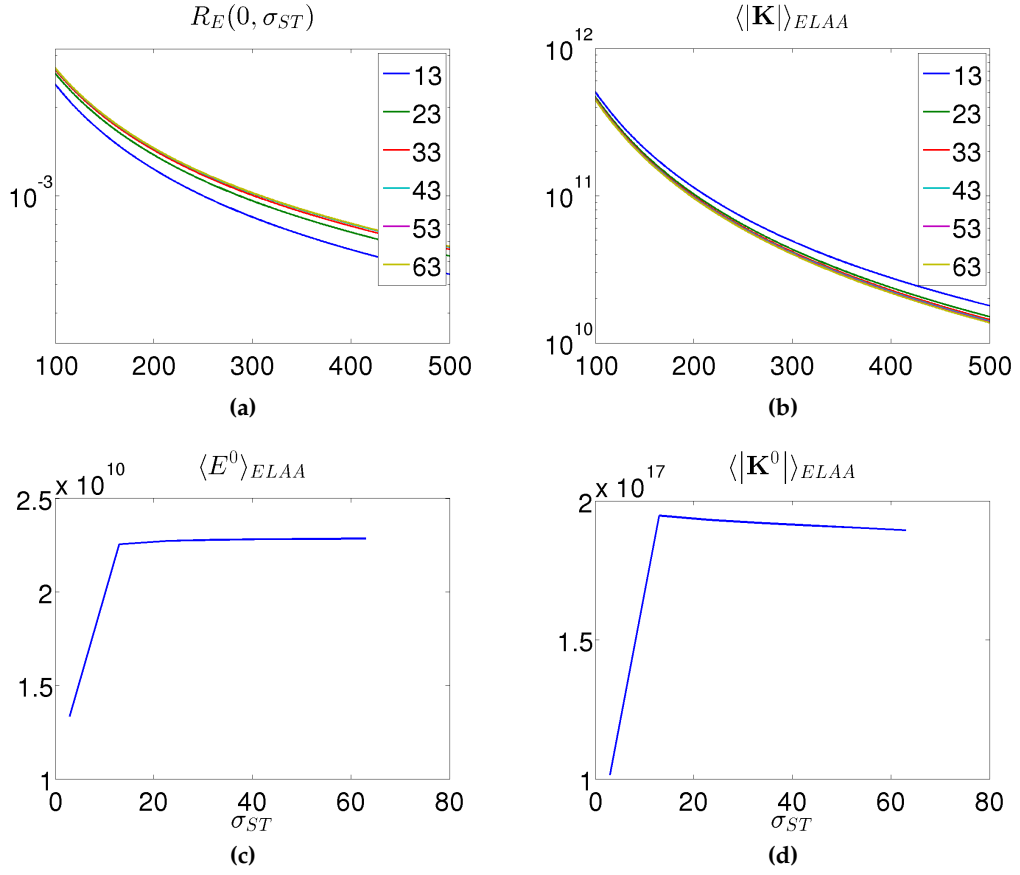


Figure 4.10.: Study of the dependency of the mean energy $\langle E^k \rangle_{ELAA}$ and the mean curvature $\langle \|K\| \rangle_{ELAA}$ on the window size σ_{ST} of the structure tensor prior E_{ST}^{prior} . Figure 4.10a shows the mean energy $\langle E^k \rangle_{ELAA}$ per iteration $k \geq 100$ for various σ_{ST} and figure 4.10b shows the mean curvature $\langle \|K\| \rangle_{ELAA}$, also for various σ_{ST} . Figures 4.10c and 4.10d show the initial energy $\langle E^0 \rangle_{ELAA}$ and the initial curvature $\langle \|K^0 \| \rangle_{ELAA}$ for $k = 0$. In figure 4.10a we can see that for smaller σ_{ST} the energy $\langle E^k \rangle_{ELAA}$ converges to lower values. Conversely for larger window sizes σ_{ST} the mean energy profiles $\langle E^k \rangle_{ELAA}$ per σ_{ST} converge. In figure 4.10b we observe a similar behavior for the curvature $\langle \|K\| \rangle_{ELAA}$: For small σ_{ST} the curvature $\langle \|K\| \rangle_{ELAA}$ is comparatively large. As σ_{ST} rises the profile of $\langle \|K\| \rangle_{ELAA}$ per σ_{ST} converges, albeit at lower values. Figures 4.10c and 4.10d show that the initial energy and the initial curvature for $\sigma_{ST} = 3$ have half the values then for the larger window sizes $\sigma_{ST} = 13 \dots 63$.

$\langle \|K\|^0 \rangle$ falls. Since by eq. (4.59) the curvature operator \mathbf{K} is proportional to the decay rate of the Euler-Lagrange differentials $[\mathcal{E}]$ we expect $\langle E^k \rangle_{ELAA}$ to converge at a slower rate for larger window sizes σ_{ST} . In figure 4.10a we have plotted the *relative energy*

$$R_E(k, \sigma_{ST}) = \frac{\langle E^k \rangle}{\langle E^0 \rangle} \quad (4.76)$$

The relative energy $R_E(k, \sigma_{ST})$ tells us how much the energy $\langle E^k \rangle$ has decayed at iteration $k > 0$ relative to the initial energy $\langle E^0 \rangle$ for a specific window size σ_{ST} . Lower values of $R_E(k, \sigma_{ST})$ indicate higher decay rates of $\langle E^k \rangle$ and vice versa. This is supported by figure 4.10b where the curvature $\langle \|K\|^k \rangle$ ($k \geq 100$) is plotted for the same window sizes σ_{ST} as for the energy $\langle E^k \rangle_{ELAA}$ in figure 4.10a. The curvature $\langle \|K\|^k \rangle$ is inverse proportional to the window size σ_{ST} at every iteration k . By eqs. (4.58) and (4.59) the decay rate of Euler-Lagrange differentials $[\mathcal{E}]$ is also inverse proportional to σ_{ST} . Thus we conclude that larger window sizes σ_{ST} have a negative impact on the convergence behavior of the ELAA. On the other hand from figure 4.10c and figure 4.10d we can see that for the smallest window size $\sigma_{ST} = 3$ the mean energy $\langle E^0 \rangle_{ELAA}$ and the curvature $\langle \|K\|^0 \rangle$ both have the smallest values. Thus again by eq. (4.59) the energy $\langle E^k \rangle_{ELAA}$ for $\sigma_{ST} = 3$ has the worst convergence behavior due to the low initial curvature $\langle \|K\|^0 \rangle$.

4.2. summary

We have introduced a new algorithm based on an extension to the principle of least action (PLA). The PLA is the idea that for any problem which can be modeled by a Gibbs random field ϕ it is possible to construct an energy functional

$$E(\phi, \mathbf{X}_e\phi) = E^{data}(\phi) + \lambda E^{prior}(\mathbf{X}_e\phi) \quad (4.77)$$

The minimizer ϕ^* of the energy E is obtained by solving the Euler-Lagrange equations

$$[\mathcal{E}_{tot}] (\phi(\mathbf{x})) |_{\phi=\phi^*} = 0, \quad [\mathcal{E}_{tot}] (\phi(\mathbf{x})) = \frac{\delta \mathcal{E}}{\delta \phi(\mathbf{x})} - \sum_{\mu} \frac{d}{dx_{\mu}} \left(\frac{\delta \mathcal{E}}{\partial_{\mu} \phi(\mathbf{x})} \right) \quad (4.78)$$

where \mathcal{E}_{tot} is the integral density function of the energy E . Our extension to the PLA is based around a flow equation for the coordinate frame Ω

$$\mathbf{x}(t) = T_t^B \circ \mathbf{x}, \quad \left. \frac{d\mathbf{x}(t)}{dt} \right|_{t=0} = \mathbf{B} \cdot \mathbf{x} \quad (4.79)$$

We termed the operator \mathbf{B} the bending operator since it bends the coordinate frame Ω such that the level-sets of the GRF ϕ obtain the geometric form preferred by the prior energy functional $E^{prior}(\mathbf{X}_e\phi)$. We showed that the rate of change of the Euler-Lagrange differentials $[\mathcal{E}_{tot}]$ under the flow in eq. (4.79) is equal to $\mathbf{K}\phi$, where the curvature operator \mathbf{K} describes the curvature of the level-sets of ϕ . The bending operator \mathbf{B} obeys the relation $\mathbf{B}\phi = 0$ for any value of the GRF ϕ at every time instance t of the flow in eq. (4.79) as a consequence of Noethers first theorem. Another consequence is that \mathbf{B} and thus \mathbf{K} is entirely determined by the prior energy E^{prior} alone.

The flow equation in eq. (4.79) was evaluated in the context of image de-noising. For image de-noising with the total variation (TV) prior from section 2.5 we showed that eq. (4.79) dramatically improves the traditional Newton method for the minimization of the energy functional E , in terms of both speed and robustness of the solution ϕ^* towards noise in the initial guess ϕ^0 . The functional derivative of the TV prior $\|\nabla\phi\|$ is equal to $\kappa(\mathbf{x}(t))$, the mean curvature of the level-set at the point $\mathbf{x}(t)$. We showed that $\kappa(\mathbf{x}(t))$ understood as a function of the flow parameter t follows an exponential behavior which is theoretically predicted since its rate of change under the flow in eq. (4.79) is equal to the curvature $\mathbf{K}\phi$.

The other model we tested for image de-noising deployed the structure tensor prior E_{ST}^{prior} from section 3. We found that the flow in eq. (4.79) had virtually no effect on the minimization of the energy E . It was shown that the root of the problem is that E_{ST}^{prior} depends on the structure tensor S which integrates the gradient $\nabla\phi$ (more specifically the tensor product $\nabla\phi \cdot \nabla^T\phi$) over a region $\mathcal{A}_{ST} \subset \Omega$ of size σ_{ST} around each pixel $\mathbf{x} \in \Omega$. We showed empirically that for window sizes $\sigma_{ST} > 3$ the curvature $\mathbf{K}\phi$ (averaged over Ω) decreases as σ_{ST} increases. Since the curvature $\mathbf{K}\phi$ is equal to the rate of change of the Euler-Lagrange differentials, the result is that E_{ST}^{prior} and the total energy E converge at slower rates for larger window sizes σ_{ST} under the flow in eq. (4.79).

5. Conclusions

The focus of this thesis are problems in computer vision which can be modeled by Gibbs random field models (GRF). In section 2.1 we shortly reviewed the theory of GRFs, namely that a GRF is a function $\phi(\mathbf{x})$ for which an energy $E_Y(\phi, \mathbf{X}_e\phi) = E_Y^{data}(\phi) + E^{prior}(\mathbf{X}_e\phi)$ is defined. The GRF ϕ is a hidden variable which cannot be directly observed and the problem of inferring the optimal value of ϕ given the data Y is equal to finding the minimum $\phi_{tot}^* = \operatorname{argmin} E_Y(\phi, \mathbf{X}_e\phi)$. The data term $E_Y^{data}(\phi)$ defines how the GRF ϕ is coupled to the data Y and the prior $E^{prior}(\mathbf{X}_e\phi)$ incorporates geometrical constraints on the level-sets of ϕ via the analytical form of the differential operators $\mathbf{X}_e = (X_e^1, \dots, X_e^k)$. For instance it was shown that if the differential operators of the prior are the Cartesian derivatives, $\mathbf{X}_e = \nabla$, then $E^{prior}(\nabla\phi)$ penalizes level-sets of ϕ with non-zero curvature κ .

We emphasized that the prior $E^{prior}(\mathbf{X}_e\phi)$ must be invariant upon the action of a group of continuous smooth transformations \mathbb{G} , $g \circ E^{prior} = E^{prior}$ for all $g \in \mathbb{G}$. The operators $g \in \mathbb{G}$ operate on both the field $\phi(\mathbf{x})$ and the coordinates \mathbf{x} themselves, $g \circ \phi(\mathbf{x}) = g^\phi \circ \phi(g^\Omega \circ \mathbf{x})$. The reason for this constraint is that the (local) minimizers $\phi_{tot}^* \in \mathcal{A}_{tot}$ of the total energy E_Y is a trade-off between the minimizers of the data term E_Y^{data} and the minimizers of the prior energy E^{prior} . The number of minimizers of the data term E_Y^{data} is usually small, since they are highly conditioned on the data Y . However since the energy E_Y has to accommodate many different values of the data Y , the prior energy $E^{prior}(\mathbf{X}_e\phi)$ must not penalize ϕ to few distinct minima ϕ^* . Instead we argued that the minimizers of E^{prior} must be a large set \mathcal{A} generated by the Lie group \mathbb{G} , $\mathcal{A} = \{\phi^* \mid \phi^* = g \circ \phi_0^*, \quad g \in \mathbb{G}\}$. ϕ_0^* is an arbitrary minimizer of E^{prior} .

Although the prior E^{prior} is invariant under the group \mathbb{G} , the total energy E_Y is generally not invariant, due to the data term E_Y^{data} , it may be invariant under local transformations around the unity, $g \in U_e \subset \mathbb{G}$. In the case that $E_Y(\phi, \mathbf{X}_e\phi)$ is symmetric upon the action of the group U_e , $g \circ E_Y = E_Y$ for all $g \in U_e$, Emmy Noether [35, 36] discovered that there exist r vector valued functions $\mathbf{W}_m(\mathbf{x})$ for which the following relation holds

$$\operatorname{div}(\mathbf{W}_m) - X_{e,m}^\Omega(\phi)[\mathcal{E}] = 0 \quad (5.1)$$

where the generator $X_{e,m}$ is an element of the Lie algebra \mathcal{G} corresponding to the group \mathbb{G} . In section 2.3 we reviewed the theory of Lie groups and in section 2.4.2 we transformed eq. (5.1) into an equation involving r differential operators \mathbf{B}_m acting on ϕ and derived from the Lie algebra \mathcal{G} of the group \mathbb{G}

$$\mathbf{B}_m \phi = 0, \quad \mathbf{B}_m = \sum_i \frac{\delta \mathcal{E}^{prior}}{\delta X_e^i(\phi)} \cdot [X_e^m, X_e^i] \quad (5.2)$$

Although the total energy E_Y was taken into account in eqs. (5.1) and (5.2) the \mathbf{B}_m only depend on the prior energy density \mathcal{E}^{prior} and stem from the variation of the total energy E_Y under the spatial transformations g^Ω .

In section 2.5 we reviewed the total variation (TV) based prior

$E_{TV}^{prior}(\nabla\phi) = \int \|\nabla\phi\| d^2x$, considering its transformation properties under the 3 dimensional group $\mathbb{G} = \mathbb{T} \times SO(2)$, the group of translations and rotation on the \mathbb{R}^2 plane. E_{TV}^{prior} is a widely used prior since it preserves discontinuities in the GRF ϕ . It was shown that the 2 operators \mathbf{B}_x and \mathbf{B}_y related to the subgroup \mathbb{T} vanish independently of ϕ , $\mathbf{B}_{x/y} = 0$. We called the symmetry under \mathbb{T} a trivial symmetry. The operator \mathbf{B}_θ which is related to $SO(2)$ does not vanish, however it still obeys eq. (5.2), hence $SO(2)$ is a non-trivial symmetry group of E_{TV}^{prior} . It was argued that the level-sets of the minimizers $\phi^* \in \mathcal{A}$ have zero curvature, thus the level-sets are the lines $\mathbf{x}(s) = \mathbf{x}_0 + s \cdot \mathbf{v}$. By the rotation invariance of E_{TV}^{prior} the orientation of \mathbf{v} is arbitrary.

In section 3 we developed a new prior E_{ST}^{prior} based on the structure tensor in [87, 34]. The starting point was that the consideration to build the prior from the differential operator $V = \mathbf{v} \cdot \nabla$ with constant vector \mathbf{v} . We reviewed the structure tensor $S(\nabla\phi)$ which is a weighted sum over the orientations of $\nabla\phi$ with in a window $\mathcal{A}_{x_0,ST}$ of size σ_{ST} around each point \mathbf{x}_0 . Its eigenvector to the strongest eigenvalue is parallel to $\nabla\phi$ and the eigenvector to the weakest eigenvalue is the component vector \mathbf{v} from our differential operator V . We proposed to use the determinant of the structure tensor as a prior, $E_{ST}^{prior}(\nabla\phi) = \text{Det}(S)$. The rational behind this proposal is that the minimizers ϕ^* of E_{ST}^{prior} have level-sets which are approximately linear within the regions $\mathcal{A}_{x_0,ST}$. Hence the prior E_{ST}^{prior} smooths the GRF ϕ along the dominant level-set within the region $\mathcal{A}_{x_0,ST}$ at point \mathbf{x}_0 . Like the TV prior E_{TV}^{prior} , our new prior E_{ST}^{prior} is invariant under the group $\mathbb{G} = \mathbb{T} \times SO(2)$ so that it agnostic towards the orientation of the level-sets of ϕ . In section 3 we evaluated the use of both priors in the context of multi-modal optical flow, the results of which are summarized in section 3.7.7.

In section 4 we proposed a new kind of algorithm for the minimization problem $\phi_{tot}^* = \text{argmin} E_Y(\phi, \mathbf{X}_e\phi)$. We focused on describing the minimization problem

as a mapping $T(t)$ from some initial guess ϕ_0 to the minimizer ϕ_{tot}^* , $\phi_{tot}^* = T(t) \circ \phi_0$, $t \rightarrow \infty$. Similar to the analysis of the symmetries of E_Y , we proposed splitting the transformation $T(t)$ into a part T_t^ϕ that operates on the GRF ϕ and a part T_t^B which operates only on the coordinate frame Ω

$$\phi(\mathbf{x}, t) = T_t^\phi \circ \phi_0(\mathbf{x}(t)), \quad \mathbf{x}(t) = T_t^B \circ \mathbf{x}_0, \quad \mathbf{x}_0 \in \Omega_0 \quad (5.3)$$

$$\phi_{tot}^*(\mathbf{x}) = \lim_{t \rightarrow \infty} \phi(\mathbf{x}, t) \quad (5.4)$$

where Ω_0 is an initial reference frame. We showed that both operators T_t^ϕ and T_t^B lead to two coupled flow equations for the GRF ϕ and the coordinate system Ω

$$\left. \frac{d\phi(\mathbf{x}(t), t)}{dt} \right|_{t=0} = [\mathcal{E}](\phi, \mathbf{x}(t)) \quad (5.5)$$

$$\left. \frac{d\mathbf{x}(t)}{dt} \right|_{t=0} = \mathbf{B} \cdot \mathbf{x}(t) \quad (5.6)$$

where the operator \mathbf{B} in eq. (5.6) is a linear combination of the operators \mathbf{B}_m in eq. (5.2). The flow equation in eq. (5.5) is the conventional flow equation used in many algorithms for solving the minimization problem $\phi_{tot}^* = \operatorname{argmin} E_Y(\phi, \mathbf{X}_e \phi)$ like steepest descent methods [6]. It is however mostly deployed on the rigid Cartesian coordinate frame Ω_0 . This where the flow equation in eq. (5.6) comes into play. The rational behind eq. (5.6) is that it should *bend* the coordinate frame Ω in such a manner that the level-sets of the initial guess ϕ_0 when evaluated on the coordinates $\mathbf{x}(t)$, $\phi_0(\mathbf{x}(t))$, appear to fulfill the geometric constraints imposed by the prior energy E^{prior} .

We tested the effectiveness of the bending operator \mathbf{B} in the context of image denoising for both the TV prior E_{TV}^{prior} and the structure tensor prior E_{ST}^{prior} . The flow in eq. (5.6) was run on an initial image $\phi_0(\mathbf{x}_0)$ which is contaminated with Gaussian noise with the individual priors. Indeed we found that both prior energies get minimized by the flow in eq. (5.6), however E_{ST}^{prior} converges significantly slower than E_{TV}^{prior} . To get a better picture we analytically computed the rate of change of the Euler-Lagrange densities for the priors under the flow in eq. (5.6) and found out that it is equal to an operator equation $\mathbf{K} \phi$ which we interpret as the curvature of the level-sets of ϕ

$$\left. \frac{d[\mathcal{E}^{prior}](\phi, \mathbf{x}(t))}{dt} \right|_{t=0} = \mathbf{K} \phi \quad (5.7)$$

The curvature \mathbf{K} in eq. (5.7) is functionally dependent on the bending operator \mathbf{B} . We showed that due to $\mathbf{B} \phi = 0$ (eq. (5.2)) eq. (5.7) holds for the Euler-Lagrange

differentials of the total energy E_Y

$$\left. \frac{d[\mathcal{E}](\phi, \mathbf{x}(t))}{dt} \right|_{t=0} = \mathbf{K}\phi \quad (5.8)$$

It was shown that the flow T_t^B in eq. (5.6) alone, that is without the action of T_t^ϕ in eq. (5.5), minimizes the curvature \mathbf{K} and hence the total energy E_Y .

In the case of the prior E_{ST}^{prior} we found out through experimental analysis that the norm of the curvature \mathbf{K} is inverse proportional to the window size σ_{ST} of the structure tensor. This means that E_{ST}^{prior} decreases *slower* for *larger* window sizes σ_{ST} . We concluded that due to the neighborhood nature of the structure tensor, the window \mathcal{A}_{ST} acts like a drag-net that tampers the speed of the flow in eq. (5.6). Thus the result of minimizing the energy E_Y with the combined flow in eqs. (5.5) and (5.6) is only marginally better than deploying the original flow in eq. (5.5) alone.

On the other hand the image denoising model with the TV based prior E_{TV}^{prior} behaved completely differently. The energy E_Y converged approximately *twice* as fast when running the combined flow in eqs. (5.5) and (5.6) as compared to running the flow in eq. (5.5) alone! Our explanation is that E_{TV}^{prior} does not measure the curvature of the level-sets of ϕ based on local statistics like E_{ST}^{prior} does. Instead it measures the mean curvature $\kappa(\mathbf{x}(t))$ of a level-set at the point $\mathbf{x}(t)$. More precisely the mean curvature $\kappa(\mathbf{x})$ is the (weak) functional derivative of $E_{TV}^{prior}(\nabla\phi)$ with respect to ϕ , thus the Euler-Lagrange differential of E_{TV}^{prior} . By eq. (5.7) $\kappa(\mathbf{x})$ is related to the curvature operator \mathbf{K}

$$\dot{\kappa}(\mathbf{x}(t)) = \mathbf{K}\phi(\mathbf{x}(t)) \quad (5.9)$$

Eq. (5.9) suggests that the mean curvature $\kappa(\mathbf{x}(t))$ follows an exponential law, given that only the flow in eq. (5.6) is run. Indeed we were able to run the flow in eq. (5.6) and fit the observed values for κ to the exponential function $\kappa(\mathbf{x}(t)) = \alpha \exp(-\beta t) + \gamma$. The relative error of the fit (error divided by the value of κ) is of the order 10^{-1} . Since the mean curvature κ at a single point \mathbf{x} can be significantly higher than the averaged curvature over a region \mathcal{A}_{ST} (as is the case for the structure tensor prior E_{ST}^{prior}) the decreasing rate $\dot{\kappa}(\mathbf{x}(t))$ can also be high. This explains why the combined flow in eqs. (5.5) and (5.6) is faster than eq. (5.5) alone.

5.1. Outlook

The bending operator \mathbf{B} in eq. (5.6) and the theory behind it was developed for GRF models with total energies $E_Y(\phi, \mathbf{X}_e\phi)$ which only contain first order derivatives. However Emmy Noethers first theorem covers models with derivatives of any order. Thus future research should provide the advancement of the proposed theory to higher order derivative models $E_Y(\phi, \mathbf{X}_e^k\phi)$, $k > 1$ along the same lines discussed in section 4. Such higher order models allow for constraints on the second or higher order derivative of the GRF ϕ and are thus less restrictive than first order models. Due to the validity Emmy Noethers first for all higher order models, we can expect the same positive results of a higher order bending algebra on the minimization of $E_Y(\phi, \mathbf{X}_e^k\phi)$.

Another option for future research is the advancement of the bending algebra theory to Emmy Noethers second theorem which handles the case of *infinite dimensional* Lie groups \mathbb{G}_∞ . As briefly explained in section 2.4.1 the Lie algebra \mathcal{G}_∞ of the group \mathbb{G}_∞ is loosely speaking not globally constant but a function on the coordinate frame Ω , the structure 'constants' C_{mn}^l of \mathcal{G}_∞ are actually functions $C_{mn}^l := C_{mn}^l(x)$. Models based on the infinite dimensional Lie group \mathbb{G}_∞ are called *gauge theories*. Virtually all models in contemporary physics are gauge theories, from the quantum field theories of the standard model [9] to general relativity [97], Loop Quantum Gravity [98] and String Theory [99]. The development of the extended principle of least action in section 4 to a more general principle of least action encoding \mathbb{G}_∞ might be beneficial not only for the computational algorithms used in the aforementioned gauge theories but may also prove to provide a better understanding the theories themselves.

A. Lie Groups

A.1. The Prolonged Action

The prolonged action of an r -dimensional Lie Group \mathbb{G} on the gradient $\nabla\phi$ is given by

$$D\phi_i^\nu = \frac{d\omega_i^\phi}{dx_\nu} - \sum_\mu \frac{d\omega_\mu^i}{dx_\nu} \partial_\mu \phi \quad (\text{A.1})$$

We want to derive eq. (A.1) following [28, 31] by considering an arbitrary one dimensional subgroup $g_{\gamma(t)} \subset \mathbb{G}$ with $g_{\gamma(0)} = e$. The function $\gamma(t) = (a_1(t) \dots a_r(t))$ is a continuous path in the parameter space of the group \mathbb{G} . The gradient $\nabla\phi$ transforms under the action of $g_{\gamma(t)}$ like

$$F(t, \nabla\phi(\mathbf{x})) = g_{\gamma(t)} \circ (\nabla\phi(\mathbf{x})) = \tilde{\nabla}\tilde{\phi}(\tilde{\mathbf{x}}), \quad \tilde{\nabla} = J_t^{-T} \nabla \quad (\text{A.2})$$

where J_t is the Jacobian

$$J_{t,\mu\nu} = \frac{d\tilde{x}_\mu}{dx_\nu}, \quad \tilde{\mathbf{x}} = g_{\gamma(t)} \circ \mathbf{x} \quad (\text{A.3})$$

We compute the prolonged action $D\phi_i^\nu$ by taking the derivative of $F(t, \nabla\phi(vx))$ with respect to t at $t = 0$

$$\left. \frac{dF}{dt} \right|_{t=0} = \left. \frac{d}{dt} J_t^{-T} \right|_{t=0} \nabla\phi + J^{-T} \nabla \left. \frac{d}{dt} \tilde{\phi}(\tilde{\mathbf{x}}) \right|_{t=0} \quad (\text{A.4})$$

$$= - J_t^{-T} \left. \frac{d}{dt} J_t \right|_{t=0} J_t^{-T} \nabla\phi + J^{-T} \nabla \left. \frac{d}{dt} \tilde{\phi}(\tilde{\mathbf{x}}) \right|_{t=0} \quad (\text{A.5})$$

The Jacobian evaluated at $t = 0$ yields the identity matrix

$$J_{0,\mu\nu} = \delta_{\mu\nu} \quad (\text{A.6})$$

The derivative of the transformed coordinates \tilde{x} and the field $\tilde{\phi}$ with respect to t at $t = 0$ yield linear combinations of the infinitesimal variations ω^i and ω_ϕ^i

$$\left. \frac{d\tilde{x}_\mu}{dt} \right|_{t=0} = \sum_{i=1}^r \omega_\mu^i \alpha_i, \quad \left. \frac{d\tilde{\phi}(\tilde{x})}{dt} \right|_{t=0} = \sum_{i=1}^r \omega_\phi^i \alpha_i \quad (\text{A.7})$$

$$\omega_\mu^i = \partial_{a_i} \tilde{x}_\mu|_{g=e}, \quad \omega_\phi^i = \partial_{a_i} \tilde{\phi}|_{g=e}, \quad \alpha_i = \left. \frac{a_i}{dt} \right|_{t=0} \quad (\text{A.8})$$

The derivative of the Jacobian evaluates to

$$\left. \frac{dJ_{t,\mu\nu}}{dt} \right|_{t=0} = \sum_{i=1}^r \frac{d\omega_\nu^i}{dx_\mu} \alpha_i \quad (\text{A.9})$$

So that eq. (A.5) becomes

$$\left. \frac{dF_\nu}{dt} \right|_{t=0} = \sum_{i=1}^r \left(\frac{d\omega_i^\phi}{dx_\nu} - \sum_\mu \frac{d\omega_\nu^i}{dx_\mu} \partial_\mu \phi \right) \alpha_i = \sum_{i=1}^r D\phi_i^\nu \alpha_i \quad (\text{A.10})$$

The prolonged action for higher order gradients goes along the same lines, expressing the transformed gradient $\tilde{\nabla}^n$ in terms of the Jacobian J_t

$$\tilde{\nabla}^n \tilde{\phi} = \left(\prod_{l=1}^n J^{-T} \nabla \right) \tilde{\phi} \quad (\text{A.11})$$

and using the product rule to compute the derivative with respect to t .

A.2. Geometrical Meaning of the Commutator $[\cdot, \cdot]$

In section 2.3 we mentioned that if $X_e \in \mathcal{G}$ is a left invariant vector field evaluated at the identity e of the Lie group \mathbb{G} and $g_{\gamma(t)} \subset \mathbb{G}$ is a one parameter subgroup defined by the vector field V_g

$$\left. \frac{dg_{\gamma(t)}}{dt} \right|_{t=0} = V_e \quad (\text{A.12})$$

then the rate of change of $X_{g_{\gamma(t)}}$ (X_e translated by $g_{\gamma(t)}$) is governed by the commutator between X_e and V_e

$$\left. \frac{dX_{g_{\gamma(t)}}}{dt} \right|_{t=0} = [V_e, X_e] \quad (\text{A.13})$$

Since both X_g and V_g are left invariant the commutator in eq. (A.13) is also left invariant and thus an element of \mathcal{G} . We will prove eq. (A.13) for the case when both V_g and X_g are spatial operators

$$X_e = \omega'_\Omega(\mathbf{x}) \partial_\nu, \quad V_e = v'_\Omega(\mathbf{x}) \partial_\nu \quad (\text{A.14})$$

We compute the rate of change of $X_{g_{\gamma(t)}}$

$$\left. \frac{dX_{g_{\gamma(t)}}}{dt} \right|_{t=0} = \left. \frac{d\omega'_\Omega(\tilde{\mathbf{x}})}{dt} \right|_{t=0} \partial_\nu + \omega'_\Omega(\mathbf{x}) \left. \frac{d\tilde{\partial}_\nu}{dt} \right|_{t=0} \quad (\text{A.15})$$

$$= \frac{d\omega'_\Omega}{dx_\mu} \cdot v'_\Omega{}^\mu \cdot \partial_\nu - \omega'_\Omega \cdot \frac{dv'_\Omega{}^\mu}{dx_\nu} \partial_\mu \quad (\text{A.16})$$

$$= \left(\frac{d\omega'_\Omega}{dx_\nu} \cdot v'_\Omega{}^\nu - \omega'_\Omega \cdot \frac{dv'_\Omega{}^\mu}{dx_\nu} \right) \partial_\mu = (V_e(\omega'_\Omega) - X_e(v'_\Omega)) \partial_\mu \quad (\text{A.17})$$

On the other hand we can directly compute the commutator of X_e and V_e

$$[V_e, X_e] = v'_\Omega{}^\mu \frac{d}{dx_\mu} \left(\omega'_\Omega \frac{d}{dx_\nu} \right) - \omega'_\Omega \frac{d}{dx_\mu} \left(v'_\Omega{}^\mu \frac{d}{dx_\nu} \right) \quad (\text{A.18})$$

When carrying out the product rule the terms $v'_\Omega{}^\mu \omega'_\Omega \frac{d}{dx_\mu} \frac{d}{dx_\nu}$ cancels out so that eq. (A.18) equals eq. (A.17) proving eq. (A.13).

A.3. Derivation Of Noethers Theorem

In this section we want to derive Noethers theorem from the variation of the energy

$$E(\phi, \mathbf{X}_e^\Omega \phi) = \int_\Omega \mathcal{E}(\phi(\mathbf{x}), \mathbf{X}_e^\Omega \phi(\mathbf{x})) \mathcal{N} d^2x \quad (\text{A.19})$$

under an arbitrary one parameter sub group $g_{\gamma(t)} \subset \mathbb{G}$, where the Lie group \mathbb{G} is the group of spatial and intensity transformations $\mathbb{G} = \mathbb{G}_\Omega \times \mathbb{G}_\phi$ of dimension r . The set \mathbf{X}_e^Ω is the commutative sub algebra of the Lie algebra \mathcal{G} of \mathbb{G} . The tangential vector V_e at the identity of $g_{\gamma(t)}$ is an element of \mathcal{G} and has the form

$$\left. \frac{dg_{\gamma(t)}}{dt} \right|_{t=0} = V_e = v^\mu(\mathbf{x}) \partial_\mu + v^\phi(\mathbf{x}) \frac{\delta}{\delta\phi} \quad (\text{A.20})$$

To reduce clutter we will use $v^\mu := v^\mu(\mathbf{x})$ and $v^\phi := v^\phi(\mathbf{x})$ in the following, remembering that v^μ and v^ϕ are functions of \mathbf{x} . Noethers first theorem says that

the energy E in eq. (A.19) is invariant upon the action of the r -dimensional Lie group \mathbb{G}

$$\left. \frac{d}{dt} (g_{\gamma(t)} \circ E) \right|_{t=0} = 0 \quad (\text{A.21})$$

if and only if there exists r vector valued functions $\mathbf{W}_m(\mathbf{x})$ such that the following r relations hold

$$\text{Div}(\mathbf{W}_m) + \left(\omega_m^\phi - \omega_m^{\Omega, \mu} \partial_\mu \phi \right) [\mathcal{E}] = 0, \quad 1 \leq m \leq r \quad (\text{A.22})$$

where $[\mathcal{E}]$ is the Euler-Lagrange differential

$$[\mathcal{E}] = \frac{\delta \mathcal{E}}{\delta \phi} - \text{Div}(\mathbf{P}), \quad \text{Div}(\mathbf{P}) = \sum_{i=1}^r \frac{d}{dx^\mu} \left(\omega_{\Omega, i}^\mu P^i \right), \quad P^i = \frac{\delta \mathcal{E}}{\delta X_e^{\Omega, i} \phi} \quad (\text{A.23})$$

and \mathbf{P} is the canonical momentum of the energy in eq. (A.19). We now want to prove that eq. (A.22) follows from eq. (A.21) thereby calculating the analytical form of the r vector valued functions $\mathbf{W}_m(\mathbf{x})$.

Proof: The derivative of the one parameter group $g_{\gamma(t)}$ with respect to t in eq. (A.21) yields the operator V_e acting on the energy density \mathcal{E} and the volume form d^2x . Since \mathcal{E} does not directly depend on the coordinates \mathbf{x} , the partial derivatives vanish, $v^\mu \partial_\mu \mathcal{E} = 0$. Hence we have

$$\left. \frac{d}{dt} (g_{\gamma(t)} \circ E) \right|_{t=0} = \int \left\{ v^\phi \frac{\delta \mathcal{E}}{\delta \phi} + \sum_{i=1}^r \left(X_e^{\Omega, i} (v^\phi) \right) P^i \right. \quad (\text{A.24})$$

$$\left. + \sum_{i=1}^r \left([V_e^\Omega, X_e^{\Omega, i}] \phi \right) P^i + \mathcal{E} \frac{dv^\mu}{dx^\mu} \right\} d^2x \quad (\text{A.25})$$

In eq. (A.24) the coordinates \mathbf{x} and the GRF ϕ are treated as independent variables, since the subgroups \mathbb{G}_ϕ and \mathbb{G}_Ω are independent from each other. The summand in eq. (A.25) proportional to \mathbf{P} comes from the chain-rule and eq. (A.13) while $\frac{dv^\mu}{dx^\mu} d^2x$ is the rate of change of the volume form d^2x under the action of $g_{\gamma(t)}$. We add zeros in the form

$$v^\mu \partial_\mu \phi \frac{\delta \mathcal{E}}{\delta \phi} - v^\mu \partial_\mu \phi \frac{\delta \mathcal{E}}{\delta \phi} = 0 \quad (\text{A.26})$$

$$\sum_{i=1}^r \left(V_e^\Omega (X_e^{\Omega, i} \phi) \right) P^i - \sum_{i=1}^r \left(V_e^\Omega (X_e^{\Omega, i} \phi) \right) P^i = 0 \quad (\text{A.27})$$

to eq. (A.25) and use

$$\left[V_e^\Omega, X_e^{\Omega,i} \right] + X_e^{\Omega,i} (v^\phi) - V_e^\Omega (X_e^{\Omega,i} \phi) = X_e^{\Omega,i} (v^\phi - V_e^\Omega \phi) \quad (\text{A.28})$$

After some reordering of terms in eq. (A.25) and factorizing out $(v^\phi - V_e^\Omega \phi)$ we get

$$\frac{d}{dt} (g_{\gamma(t)} \circ E) \Big|_{t=0} = \int \left\{ v^\mu \partial_\mu \phi \frac{\delta \mathcal{E}}{\delta \phi} + \sum_{i=1}^r (v^\mu \partial_\mu (X_e^{\Omega,i} \phi)) P^i \right. \quad (\text{A.29})$$

$$\left. + (v^\phi - v^\mu \partial_\mu \phi) \frac{\delta \mathcal{E}}{\delta \phi} + \sum_{i=1}^r X_e^{\Omega,i} (v^\phi - V_e^\Omega \phi) P^i + \mathcal{E} \frac{dv^\mu}{dx^\mu} \right\} d^2x \quad (\text{A.30})$$

The summands in eq. (A.29) can be combined to the total derivative acting on the energy density \mathcal{E}

$$v^\mu \partial_\mu \phi \frac{\delta \mathcal{E}}{\delta \phi} + \sum_{i=1}^r (v^\mu \partial_\mu (X_e^{\Omega,i} \phi)) P^i = v^\mu \frac{d\mathcal{E}}{dx^\mu} \quad (\text{A.31})$$

and we can combine eq. (A.31) with the term $\mathcal{E} \frac{dv^\mu}{dx^\mu}$ from eq. (A.30) to the total divergence

$$v^\mu \frac{d\mathcal{E}}{dx^\mu} + \mathcal{E} \frac{dv^\mu}{dx^\mu} = \frac{d}{dx^\mu} (v^\mu \mathcal{E}) \quad (\text{A.32})$$

so that eq. (A.30) and eq. (A.31) simplify to

$$\frac{d}{dt} (g_{\gamma(t)} \circ E) \Big|_{t=0} = \int \left\{ \frac{d}{dx^\mu} (v^\mu \mathcal{E}) + (v^\phi - v^\mu \partial_\mu \phi) \frac{\delta \mathcal{E}}{\delta \phi} \right. \quad (\text{A.33})$$

$$\left. + \sum_{i=1}^r X_e^{\Omega,i} (v^\phi - V_e^\Omega \phi) P^i \right\} d^2x \quad (\text{A.34})$$

We expand the operator $X_e^{\Omega,i}$ in terms of the Cartesian gradient ∇ , $X_e^{\Omega,i} = \omega_{\Omega,i}^\mu \partial_\mu$. Then we can convert the summand in eq. (A.34) into the sum of a total divergence and a summand proportional to the divergence of the canonical momentum \mathbf{P}

$$\begin{aligned} \sum_{i=1}^r X_e^{\Omega,i} (v^\phi - V_e^\Omega \phi) P^i &= \sum_{i=1}^r \frac{d}{dx^\mu} (\omega_{\Omega,i}^\mu (v^\phi - V_e^\Omega \phi) P^i) \\ &\quad - (v^\phi - V_e^\Omega \phi) \text{Div}(\mathbf{P}), \quad \text{Div}(\mathbf{P}) = \sum_{i=1}^r \frac{d}{dx^\mu} (\omega_{\Omega,i}^\mu P^i) \end{aligned} \quad (\text{A.35})$$

We substitute eq. (A.34) with the right hand side of eq. (A.35) and group all the total divergences together as well as the term proportional to $(v^\phi - V_e^\Omega \phi)$

$$\begin{aligned} \left. \frac{d}{dt} (g_{\gamma(t)} \circ E) \right|_{t=0} &= \int \left\{ \frac{d}{dx^\mu} \left(v^\mu \mathcal{E} + \sum_{i=1}^r \omega_{\Omega,i}^\mu (v^\phi - V_e^\Omega \phi) P^i \right) \right. \\ &\quad \left. + (v^\phi - V_e^\Omega \phi) [\mathcal{E}] \right\} d^2x \end{aligned} \quad (\text{A.36})$$

$$[\mathcal{E}] = \frac{\delta \mathcal{E}}{\delta \phi} - \text{Div}(\mathbf{P}) \quad (\text{A.37})$$

If eq. (A.21) holds for the specific path $g_{\gamma(t)}$ with the tangential vector operator V_e then it follows that the integrand of the right hand side of eq. (A.36) must vanish

$$\frac{d}{dx^\mu} \left(v^\mu \mathcal{E} + \sum_{i=1}^r \omega_{\Omega,i}^\mu (v^\phi - V_e^\Omega \phi) P^i \right) + (v^\phi - V_e^\Omega \phi) [\mathcal{E}] = 0 \quad (\text{A.38})$$

The statement in eq. (A.21) applies for all one parameter sub groups, therefore we need to expand V_e in terms of the basis $\{X_e^i\}$

$$V_e = \sum_{m=1}^r \alpha_m X_e^m \implies v^\phi = \sum_{m=1}^r \alpha_m \omega_m^\phi, \quad v^{\mu,\Omega} = \sum_{m=1}^r \alpha_m \omega_m^{\mu,\Omega} \quad (\text{A.39})$$

then eq. (A.38) translates to

$$\left. \frac{d}{dt} (g_{\gamma(t)} \circ E) \right|_{t=0} = \sum_{m=1}^r \alpha_m \left\{ \frac{d}{dx^\mu} (W_m^\mu) + (\omega_m^\phi - X_e^{\Omega,m} \phi) [\mathcal{E}] \right\} = 0 \quad (\text{A.40})$$

$$W_m^\mu = \omega_m^{\mu,\Omega} \mathcal{E} + \sum_{i=1}^r \omega_{\Omega,i}^\mu (\omega_m^\phi - X_e^{\Omega,m} \phi) P^i \quad (\text{A.41})$$

Eq. (A.40) holds for all one parameter sub groups $g_{\gamma(t)} \subset \mathbb{G}$ if the sum vanishes for any configuration of the coefficients α_m . It follows that eq. (A.22) holds for all r vector valued functions \mathbf{W}_m .

A.3.1. Connection between B_m , \mathbf{W}_m and $[\mathcal{E}]$

We want to briefly clarify the relation between the bending operator $B_m = \sum_{i=1}^r P^i [X_e^{\Omega,m}, X_e^{\Omega,i}]$, the vector valued functions \mathbf{W}_m and the Euler-Lagrange differentials $[\mathcal{E}]$. For this we need the rate of change of the energy E (eq. (A.19)) under the one parameter Lie group $g_{\gamma(t)}$ corresponding to the vector V_e (eq. (A.20)).

This rate of change is given by (eq. (A.25))

$$\left. \frac{d}{dt} (g_{\gamma(t)} \circ E) \right|_{t=0} = \int \left\{ v^\phi \frac{\delta \mathcal{E}}{\delta \phi} + \sum_{i=1}^r (X_e^{\Omega, i} (v^\phi)) P^i \right. \quad (\text{A.42})$$

$$\left. + \sum_{i=1}^r \left([V_e^\Omega, X_e^{\Omega, i}] \phi \right) P^i + \mathcal{E} \frac{dv^\mu}{dx^\mu} \right\} d^2x \quad (\text{A.43})$$

Similar to eq. (A.35) we use the following transformation

$$\sum_{i=1}^r X_e^{\Omega, i} (v^\phi) P^i = \sum_{i=1}^r \frac{d}{dx^\mu} (\omega_{\Omega, i}^\mu v^\phi P^i) - v^\phi \text{Div}(\mathbf{P}) \quad (\text{A.44})$$

to transform eq. (A.42)

$$\left. \frac{d}{dt} (g_{\gamma(t)} \circ E) \right|_{t=0} = \int \left\{ v^\phi [\mathcal{E}] + \sum_{i=1}^r \frac{d}{dx^\mu} (\omega_{\Omega, i}^\mu v^\phi P^i) \right. \quad (\text{A.45})$$

$$\left. + \sum_{i=1}^r \left([V_e^\Omega, X_e^{\Omega, i}] \phi \right) P^i + \mathcal{E} \frac{dv^\mu}{dx^\mu} \right\} d^2x \quad (\text{A.46})$$

We use eq. (A.39) to express V_e in terms of the basis $\{X_e^i\}$

$$\left. \frac{d}{dt} (g_{\gamma(t)} \circ E) \right|_{t=0} = \int \sum_{m=1}^r \alpha_m \left\{ \omega_m^\phi [\mathcal{E}] + \sum_{i=1}^r \frac{d}{dx^\mu} (\omega_{\Omega, i}^\mu \omega_m^\phi P^i) \right. \quad (\text{A.47})$$

$$\left. + \sum_{i=1}^r \left([X_e^{\Omega, m}, X_e^{\Omega, i}] \phi \right) P^i + \mathcal{E} \frac{d\omega_{\Omega, m}^\mu}{dx^\mu} \right\} d^2x$$

Comparing eq. (A.46) with eq. (A.40) we obtain

$$\mathbf{B}_m \phi = \frac{d}{dx^\mu} (\widetilde{W}_m^\mu) - X_e^{\Omega, m} \phi [\mathcal{E}] - \mathcal{E} \frac{d\omega_{\Omega, m}^\mu}{dx^\mu} \quad (\text{A.48})$$

$$\mathbf{B}_m \phi = \sum_{i=1}^r \left([X_e^{\Omega, m}, X_e^{\Omega, i}] \phi \right) P^i, \quad \widetilde{W}_m^\mu = W_m^\mu - \sum_{i=1}^r \omega_{\Omega, i}^\mu \omega_m^\phi P^i$$

Eq. (A.48) is independent of the intensity variations ω_m^ϕ and hence independent of the entire sub group \mathbb{G}_ϕ . If the energy E in eq. (A.19) is invariant under the sub group \mathbb{G}_Ω , then it follows that \mathbf{B}_m is a level-set operator

$$\mathbf{B}_m \phi = 0 \quad (\text{A.49})$$

Eq. (A.49) is valid for all ϕ and poses a constraint on the possible values of the Euler-Lagrange differentials $[\mathcal{E}]$ via eq. (A.48).

B. The Bending Algebra

B.1. The curvature operator

Given an energy functional $E(\phi, \mathbf{X}_e^\Omega \phi)$ for the GRF ϕ we would like to compute the rate of change of the corresponding Euler-Lagrange differentials $[\mathcal{E}_{tot}]$ under the spatial deformation T_t^B (eq. (4.52))

$$\mathbf{x}(t) = T_t^B \circ \mathbf{x}, \quad \left. \frac{d\mathbf{x}(t)}{dt} \right|_{t=0} = \sum_{m=1}^r \beta_m \mathbf{B}_m \mathbf{x} = \mathbf{B} \cdot \mathbf{x} \quad (\text{B.1})$$

We write the rate of change as

$$[\mathcal{E}_{tot}]' = \left. \frac{d}{dt} (T_t^B \circ [\mathcal{E}_{tot}]) \right|_{t=0} \quad (\text{B.2})$$

where $[\mathcal{E}_{tot}]$ is

$$[\mathcal{E}_{tot}](\phi(\mathbf{x})) = \frac{\delta \mathcal{E}^{data}}{\delta \phi(\mathbf{x})} - \text{Div}(\mathbf{P}) \quad (\text{B.3})$$

$$\text{Div}(\mathbf{P}) = \sum_{\mu} \frac{d}{dx_{\mu}} \left(\sum_i \omega_{\Omega, i}^{\mu} P_i \right), \quad P_i = \frac{\delta \mathcal{E}^{prior}}{\delta X_e^{\Omega, i}} \quad (\text{B.4})$$

Since \mathbf{B} in eq. (B.1) is a level-set operator of the GRF ϕ the data energy density $\mathcal{E}^{data}(\phi(\mathbf{x}))$ and all its functional derivatives are invariant under the flow in eq. (B.1)

$$\left. \frac{d}{dt} (T_t^B \circ \mathcal{E}^{data}) \right|_{t=0} = \frac{\delta \mathcal{E}^{data}}{\delta \phi(\mathbf{x})} \mathbf{B} \phi = 0 \quad (\text{B.5})$$

Thus $[\mathcal{E}_{tot}]'$ reduces to

$$[\mathcal{E}_{tot}]' = - \left. \frac{d}{dt} (T_t^B \circ \text{Div}(\mathbf{P})) \right|_{t=0} \quad (\text{B.6})$$

In order to compute eq. (B.6) we consider the integral

$$F = \int_{\Omega} \text{Div}(\mathbf{P})\phi(\mathbf{x})d^2x \quad (\text{B.7})$$

We compute the rate of change of F in eq. (B.7)

$$\left. \frac{d}{dt} (T_t^B \circ F) \right|_{t=0} = \int_{\Omega} \left[\left. \frac{d}{dt} (T_t^B \circ \text{Div}(\mathbf{P})) \right|_{t=0} \phi + \text{Div}(\mathbf{P}) \mathbf{B}\phi \right] d^2x \quad (\text{B.8})$$

Since $\mathbf{B}\phi = 0$ we have

$$\left. \frac{d}{dt} (T_t^B \circ F) \right|_{t=0} = \int_{\Omega} \left. \frac{d}{dt} (T_t^B \circ \text{Div}(\mathbf{P})) \right|_{t=0} \phi d^2x = \int_{\Omega} [\mathcal{E}_{tot}]' \phi d^2x \quad (\text{B.9})$$

The integral in eq. (B.7) is well-defined even in the case when the GRF ϕ contains discontinuities. By the Riesz Representation theorem [40] we can shift the derivatives in eq. (B.7) from \mathbf{P} onto the GRF ϕ

$$F = - \int_{\Omega} \sum_i P_i \omega_{\Omega,i}^{\mu} \partial_{\mu} \phi d^2x = - \int_{\Omega} \sum_i P_i X_e^{\Omega,i} \phi d^2x \quad (\text{B.10})$$

where $X_e^{\Omega,i} = \omega_{\Omega,i}^{\mu} \partial_{\mu}$ is an element of the commutative Lie algebra \mathcal{B} used to construct the prior \mathcal{E}^{prior} . We define the vector field \mathbf{P}_g

$$\mathbf{P}_g = \sum_i P_{i,g} \cdot X_g^{\Omega,i}, \quad P_{i,g} = P_i(X_g^{\Omega,i} \phi) \quad (\text{B.11})$$

Then the product $\sum_i P_i X_e^{\Omega,i}$ in eq. (B.10) is the evaluation of the vector field \mathbf{P}_g at the identity $e \in \mathbb{G}$ and eq. (B.10) can be written as

$$F = - \int_{\Omega} \mathbf{P}_e \phi d^2x \quad (\text{B.12})$$

The rate of change of \mathbf{P}_g under the flow in eq. (B.1) is

$$\left. \frac{d}{dt} \mathbf{P}_{T_t^B} \right|_{t=0} = [\mathbf{B}, \mathbf{P}_e] \quad (\text{B.13})$$

so that from eq. (B.12) we have

$$\left. \frac{d}{dt} (T_t^B \circ F) \right|_{t=0} = - \int_{\Omega} [\mathbf{B}, \mathbf{P}_e] \phi d^2x \quad (\text{B.14})$$

We equate eq. (B.14) with eq. (B.9)

$$\int_{\Omega} [\mathcal{E}_{tot}]' \phi d^2x = - \int_{\Omega} \mathbf{K} \phi d^2x, \quad \mathbf{K} = [\mathbf{B}, \mathbf{P}_e] \quad (\text{B.15})$$

The operator \mathbf{K} is the *curvature operator* and we can see that $[\mathcal{E}_{tot}]'$, the rate of change of the Euler-Lagrange differential $[\mathcal{E}_{tot}]$ under the flow in eq. (B.1), can only be understood as the eigenvalue to the negative curvature operator

$$[\mathcal{E}_{tot}]' \cdot \phi = -\mathbf{K} \phi \quad (\text{B.16})$$

The curvature operator \mathbf{K} can be expanded in terms of the Cartesian gradient ∇

$$\mathbf{K} \phi = \mathbf{k} \nabla \phi \quad (\text{B.17})$$

and we define the norm $\|\mathbf{K}\|$ as the length of the coefficient vector \mathbf{k} , $\|\mathbf{K}\| = \|\mathbf{k}\|$.

B.2. TV Image Denoising, supplementary results




(a) input image


(b) ϕ_0

(c) ϕ_{ELAA}^*


(d) ϕ_{BNA}^*




(e)




(f)



(g)



(h)




(i) input image


(j) ϕ_0

(k) ϕ_{ELAA}^*


(l) ϕ_{BNA}^*




(m)



(n)




(o)



(p)



(a) input image




(b) ϕ_0




(c) ϕ_{ELAA}^*




(d) ϕ_{BNA}^*




(e)



(f)




(g)



(h)



(i) input image




(j) ϕ_0




(k) ϕ_{ELAA}^*




(l) ϕ_{BNA}^*




(m)



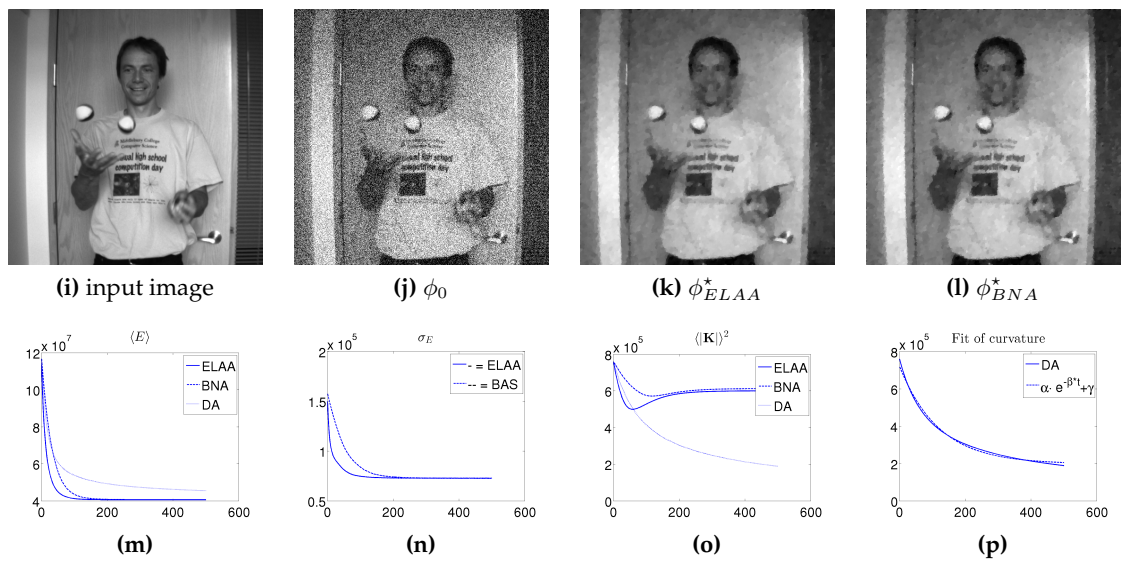
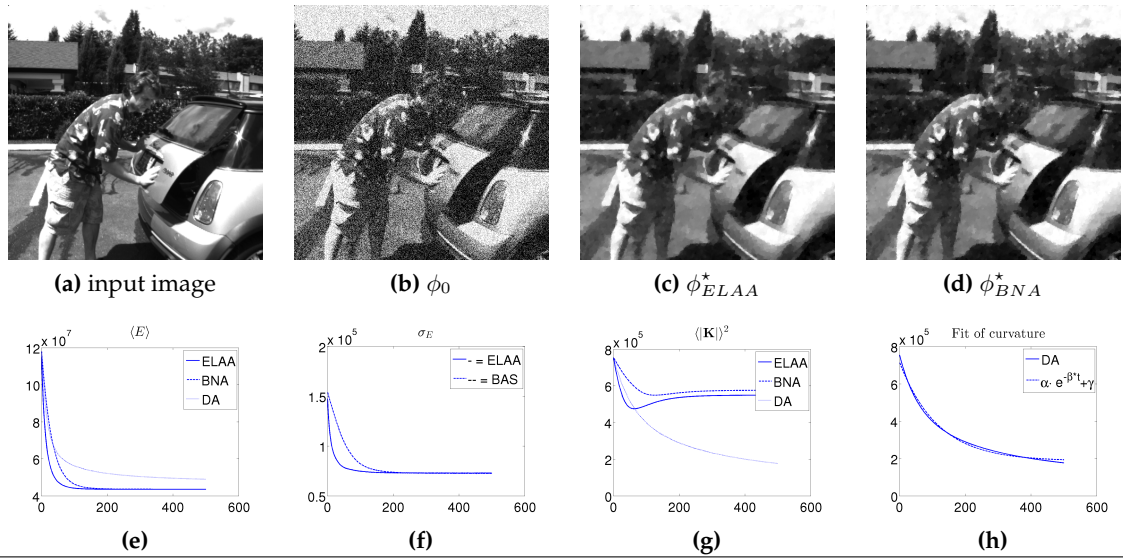
(n)




(o)




(p)





(a) input image




(b) ϕ_0




(c) ϕ_{ELAA}^*




(d) ϕ_{BNA}^*




(e)



(f)



(g)



(h)



(i) input image




(j) ϕ_0




(k) ϕ_{ELAA}^*




(l) ϕ_{BNA}^*




(m)



(n)




(o)



(p)



(a) input image




(b) ϕ_0




(c) ϕ_{ELAA}^*




(d) ϕ_{BNA}^*




(e)



(f)



(g)



(h)

C. Multimodal Optical Flow

We want to derive the similarity measure $E_{y,I}^{data}$ in a fashion which covers both the global version $E_{y,I}^{data}(\sigma^{sc}, \mathbf{d})$ (eq. (3.50)) and the local version $E_{y,I}^{data}(\sigma^{sc}, \mathbf{d})$ (eq. (3.55)). We will assume that lower case letters like y stand for low resolution and upper case letters like I stand for high resolution images. The computation deploys the row major lexicographic reordering of images to vectors and filters to (sparse banded) matrices

$$y(\mathbf{x}) \rightarrow \underline{y}, \quad \int u(\mathbf{x})v(\mathbf{x})d^2x \rightarrow \underline{u}^T \underline{v} \quad (\text{C.1})$$

$$\int_{\mathcal{A}_{x_0}} w(\mathbf{x} - \mathbf{x}_0) \cdot y(\mathbf{x})d^2x \rightarrow \mathbf{W} \underline{y} \quad (\text{C.2})$$

To derive the similarity measure $E_{y,I}^{data}$ we considered in section 3.3 a low resolution camera C_y which records the image \underline{y} and a high resolution camera C_I recording the image \underline{I} . Since C_y and C_I are physically separated the task was to compute the optical flow \mathbf{d} mapping \underline{I} to \underline{y} . To handle the difference in optical resolution, $\sigma^{sc,*}$, between \underline{y} and \underline{I} , in section 3.3 the low resolution image \underline{y} is assumed to be result of a sub-sampling process with additive Gaussian noise

$$\underline{y} = \mathbf{W}_{\sigma^{sc}} \underline{Y} + \underline{n}, \quad \underline{n} \sim \mathcal{N}(\mathbf{0}, C_n) \quad (\text{C.3})$$

where the point-spread function (PSF) $\mathbf{W}_{\sigma^{sc}}$ which models the difference in optical scale is approximately Gaussian. Eq. (C.3) connects the low frequency components of \underline{Y} to the low resolution image \underline{y} . From eq. (C.3) the log-likelihood of \underline{y} given \underline{Y} is given by

$$-\ln(p(\underline{y}|\underline{Y})) = \frac{1}{2} (\underline{y} - \mathbf{W}_{\sigma^{sc}} \underline{Y})^T \mathbf{C}_n^{-1} (\underline{y} - \mathbf{W}_{\sigma^{sc}} \underline{Y}) \quad (\text{C.4})$$

To compute the similarity measure $E_{y,I}^{data}$, the correspondence between \underline{Y} and the warped image \underline{I}_d needs to be established. To do this we can compute the conditional distribution $p(\underline{Y}|\underline{I})$: We assume \mathbf{d} to be fixed and \underline{Y} and \underline{I}_d to be

jointly Gaussian with the distribution $p(\underline{Y}, \underline{I}_d)$

$$-\ln(p(\underline{Y}, \underline{I})) = \frac{1}{2} \underline{V}^T \mathbf{Q} \underline{V}, \quad \underline{V}^T = \begin{pmatrix} \underline{Y} - \underline{\mu}_Y \\ \underline{I}_d - \underline{\mu}_{I_d} \end{pmatrix}, \quad \mathbf{Q} = \begin{pmatrix} \mathbf{C}_{YY} & \mathbf{C}_{YI_d} \\ \mathbf{C}_{I_dY} & \mathbf{C}_{I_dI_d} \end{pmatrix}^{-1} \quad (\text{C.5})$$

The covariances \mathbf{C}_{\cdot} , and the means $\underline{\mu}_{\cdot}$ in eq. (C.5) can be either the global covariance and mean in eq. (2.164) or the local variance and mean in eq. (3.33). In both cases \mathbf{C}_{\cdot} is diagonal. The precision matrix \mathbf{Q} can be expressed in terms of the Schur-Complement $\mathbf{S} = \mathbf{C}_{YY} - \mathbf{C}_{YI_d} \mathbf{C}_{I_dI_d}^{-1} \mathbf{C}_{I_dY}$

$$\mathbf{Q} = \begin{pmatrix} \mathbf{S}^{-1} & -\mathbf{S}^{-1} \mathbf{C}_{YI_d} \mathbf{C}_{I_dI_d}^{-1} \\ \mathbf{C}_{I_dI_d}^{-1} \mathbf{C}_{I_dY} \mathbf{S}^{-1} & \star \end{pmatrix} \quad (\text{C.6})$$

then the joint distribution in eq. (C.5) can be split into the product $p(\underline{Y}, \underline{I}_d) = p(\underline{Y}|\underline{I}_d)p(\underline{I}_d)$ where the conditional distribution $p(\underline{Y}|\underline{I}_d)$ can be written in terms of the Schur complement \mathbf{S}

$$-\ln(p(\underline{Y}|\underline{I}_d)) = \frac{1}{2} \tilde{\underline{Y}}^T \mathbf{S}^{-1} \tilde{\underline{Y}} - \frac{1}{2} \tilde{\underline{Y}}^T \mathbf{S}^{-1} \mathbf{C}_{YI_d} \mathbf{C}_{I_dI_d}^{-1} \tilde{\underline{I}}_d - \frac{1}{2} \tilde{\underline{I}}_d^T \mathbf{C}_{I_dI_d}^{-1} \mathbf{C}_{I_dY} \mathbf{S}^{-1} \tilde{\underline{Y}} \quad (\text{C.7})$$

$$= \frac{1}{2} (\underline{Y} - \underline{\mu}_{Y|I_d})^T \mathbf{C}_{Y|I_d}^{-1} (\underline{Y} - \underline{\mu}_{Y|I_d}) \quad (\text{C.8})$$

$$\mathbf{C}_{Y|I_d} = \mathbf{S}, \quad \underline{\mu}_{Y|I_d} = \underline{\mu}_Y + \mathbf{C}_{YI_d} \mathbf{C}_{I_dI_d}^{-1} \tilde{\underline{I}}_d \quad (\text{C.9})$$

where $\tilde{\underline{I}}_d = \underline{I}_d - \underline{\mu}_{I_d}$ and $\tilde{\underline{Y}} = \underline{Y} - \underline{\mu}_Y$. We combine the conditional $p(\underline{Y}|\underline{I}_d)$ from eq. (C.8) together with the likelihood $p(\underline{y}|\underline{Y})$ from eq. (C.4) to the posterior $p(\underline{Y}|\underline{y}, \underline{I}_d) = p(\underline{y}|\underline{Y}) p(\underline{Y}|\underline{I}_d)$ with the energy $E(\underline{Y})$

$$E(\underline{Y}) = \frac{1}{2} (\underline{y} - \mathbf{W}_{\sigma^{sc}} \underline{Y})^T \mathbf{C}_n^{-1} (\underline{y} - \mathbf{W}_{\sigma^{sc}} \underline{Y}) + \frac{1}{2} (\underline{Y} - \underline{\mu}_{Y|I_d})^T \mathbf{C}_{Y|I_d}^{-1} (\underline{Y} - \underline{\mu}_{Y|I_d}) \quad (\text{C.10})$$

The minimizer \underline{Y}_d^* of the energy $E(\underline{Y})$ is

$$\underline{Y}_d^* = \underline{\mu}_{Y|I_d} + \mathbf{C}_{Y|I_d} \mathbf{W}_{\sigma^{sc}}^T \mathbf{H}^{-1} (\underline{y} - \mathbf{W}_{\sigma^{sc}} \underline{\mu}_{Y|I_d}) \quad (\text{C.11})$$

$$\mathbf{H} = \mathbf{W}_{\sigma^{sc}} \mathbf{C}_{Y|I_d} \mathbf{W}_{\sigma^{sc}}^T + \mathbf{C}_n \quad (\text{C.12})$$

We define the similarity measure $E_{y,I}^{data}$ in terms of the conditional distribution $p(\underline{Y}|\underline{I}_d)$ from eq. (C.8) evaluated at the minimizer \underline{Y}_d^*

$$E_{y,I}^{data}(\mathbf{d}) = -\ln(p(\underline{Y}_d^*|\underline{I}_d)) \quad (\text{C.13})$$

which after some manipulations becomes

$$E_{y,I}^{data}(\mathbf{d}) = \frac{1}{2} \underline{x}^T \mathbf{H}^{-1} \mathbf{W}_{\sigma^{sc}} \mathbf{C}_{Y|I_d} \mathbf{W}_{\sigma^{sc}}^T \mathbf{H}^{-1} \underline{x}, \quad \underline{x} = \underline{y} - \mathbf{W}_{\sigma^{sc}} \underline{\mu}_{Y|I_d} \quad (\text{C.14})$$

$E_{y,I}^{data}$ in eq. (C.14) is intractable to compute due to the dense inverse matrix \mathbf{H}^{-1} . Similar to [86] we make the simplifications

$$\mathbf{W}_{\sigma^{sc}} \mathbf{C}_{Y|I_d} \mathbf{W}_{\sigma^{sc}}^T \rightarrow \mathbf{C}_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}} \quad (\text{C.15})$$

$$\mathbf{H} \rightarrow \tilde{\mathbf{H}} = \mathbf{C}_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}} + \mathbf{C}_n \quad (\text{C.16})$$

and eq. (C.14) reduces to

$$E_{y,I}^{data}(\mathbf{d}) = \frac{1}{2} \underline{x}^T \tilde{\mathbf{H}}^{-1} \mathbf{C}_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}} \tilde{\mathbf{H}}^{-1} \underline{x}, \quad \underline{x} = \underline{y} - \mathbf{W}_{\sigma^{sc}} \underline{\mu}_{Y|I_d} \quad (\text{C.17})$$

We rewrite eq. (C.17) back in terms of filters and integrals with the inverse lexicographic reordering. If we insert the global covariances in eq. (2.164) into the precision matrix \mathbf{Q} in eq. (C.5) we get the global similarity measure from eq. (3.27)

$$E_{y,I}^{data}(\mathbf{d}) = \frac{1}{2} \int_{\Omega} \left(y(\mathbf{x}) - \mu_Y - f \cdot (\langle I_d \rangle_{\sigma^{sc}}(\mathbf{x}) - \mu_I) \right)^2 \cdot U^{\sigma^{sc}} \quad (\text{C.18})$$

$$f = C_{y, \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}} C_{\langle I_d \rangle_{\sigma^{sc}}, \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}, -1}, \quad U^{\sigma^{sc}} = C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}} \left(C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}} + C_n \right)^{-2} \quad (\text{C.19})$$

And if we insert the local covariances in eq. (3.33) into eq. (C.5) we get the local similarity measure $E_{y,I_d}^{data,l}$ from eq. (3.37)

$$E_{y,I_d}^{data,l}(\sigma^{sc}, a, \mathbf{d}) = \frac{1}{2} \int_{\Omega} \left(y(\mathbf{x}) - \langle f^{\sigma^{sc}, a} I_d \rangle_{\sigma^{sc}}(\mathbf{x}) \right)^2 \cdot U^{\sigma^{sc}, a}(\mathbf{x}) d^2 x$$

$$U^{\sigma^{sc}, a}(\mathbf{x}) = C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}, a}(\mathbf{x}) \left(C_{\langle Y \rangle_{\sigma^{sc}} | \langle I_d \rangle_{\sigma^{sc}}}^{\sigma^{sc}, a}(\mathbf{x}) + C_n \right)^{-2} \quad (\text{C.20})$$

Bibliography

- [1] R. Hartley and A. Zisserman, *Multiple view geometry in computer vision*, Cambridge university press, 2003. [5](#)
- [2] Jiejie Zhu, Liang Wang, Ruigang Yang, and James Davis, “Fusion of time-of-flight depth and stereo for high accuracy depth maps,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition CVPR 2008*, 2008, pp. 1–8. [5](#)
- [3] R. Nair, F. Lenzen, S. Meister, H. Schäfer, C. Garbe, and D. Kondermann, “High accuracy tof and stereo sensor fusion at interactive rates,” in *Computer Vision–ECCV 2012. Workshops and Demonstrations*, pp. 1–11. Springer Berlin Heidelberg, 2012. [5](#)
- [4] Daniel Scharstein and Richard Szeliski, “High-accuracy stereo depth maps using structured light,” in *Proc. IEEE Computer Society Conf. Computer Vision and Pattern Recognition*, 2003, vol. 1. [5](#), [35](#)
- [5] D. Sun, S. Roth, and M. J. Black, “Secrets of optical flow estimation and their principles,” in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE, 2010, pp. 2432–2439. [7](#), [35](#)
- [6] Paul Fieguth, *Statistical image processing and multidimensional modeling*, Springer Science & Business Media, 2010. [8](#), [78](#), [101](#)
- [7] LD Landau and EM Lifshitz, “Mechanics, vol. 1,” *Course of theoretical physics*, pp. 84–93, 1976. [9](#)
- [8] Maupertuis, *Loi du repos des corps*, 1740. [9](#)
- [9] Michael E. Peskin and Daniel V. Schroeder, *An Introduction to Quantum Field Theory (Frontiers in Physics)*, Perseus Books, 1995. [9](#), [22](#), [103](#)
- [10] Richard Phillips Feynman, *The principle of least action in quantum mechanics*, Ph.D. thesis, Princeton University Princeton, New Jersey, 1942. [9](#)
- [11] Reinhard Hentschke, “Principle of least action,” in *Classical Mechanics*, pp. 89–122. Springer, 2017. [9](#)

- [12] Pierre Louis Moreau de Maupertuis, “Les loix du mouvement et du repos déduites d’un principe metaphysique,” *Hist Acad Roy Sci Belleslett* (1746), vol. 267, pp. 294, 1746. [9](#)
- [13] Isaac Newton, *Philosophiae naturalis principia mathematica*, vol. 1, excudit G. Brookman; impensis TT et J. Tegg, Londini, 1687. [9](#)
- [14] Lev Davidovich Landau, Evgenii Mikhailovich Lifshitz, John Bradbury Sykes, John Stewart Bell, and EH Dill, *Electrodynamics of continuous media*, AIP, 1961. [10](#)
- [15] Richard Phillips Feynman, Robert B Leighton, and Matthew Linzee Sands, *The Feynman Lectures on Physics: electromagnetism and matter*, vol. 2, Addison Wesley Publishing Company, 1964. [10](#)
- [16] James Clerk Maxwell, “On physical lines of force, 1861,” *A dynamical theory of the electromagnetic field*, pp. 1875–89, 1865. [10](#)
- [17] Heinrich Hertz, *Electric waves: being researches on the propagation of electric action with finite velocity through space*, Dover Publications, 1893. [10](#)
- [18] Leonid I Rudin, Stanley Osher, and Emad Fatemi, “Nonlinear total variation based noise removal algorithms,” *Physica D: Nonlinear Phenomena*, vol. 60, no. 1, pp. 259–268, 1992. [14](#), [31](#)
- [19] A. Chambolle, “An algorithm for total variation minimization and applications,” *Journal of Mathematical imaging and vision*, vol. 20, no. 1-2, pp. 89–97, 2004. [14](#), [31](#)
- [20] Mark Nitzberg and Takahiro Shiota, “Nonlinear image filtering with edge and corner enhancement,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 14, no. 8, pp. 826–833, 1992. [14](#)
- [21] J. Weickert, “Theoretical foundations of anisotropic diffusion in image processing,” *Computing, Suppl.*, vol. 11, pp. 221–236, 1996. [14](#)
- [22] Joachim Weickert, *Anisotropic diffusion in image processing*, vol. 1, Teubner Stuttgart, 1998. [14](#)
- [23] D. Mumford and J. Shah, “Optimal approximation by piecewise smooth functions and associated variational problems,” *Comm. Pure and Applied Mathematics*, vol. 42, pp. 577–685, 1989. [14](#), [33](#)
- [24] L. Ambrosio and V. M. Tortorelli, “Approximation of functional depending on jumps by elliptic functional via t-convergence,” *Communications on Pure and Applied Mathematics*, vol. XLIII, pp. 999–1036, 1990. [14](#)

- [25] L. A. Vese and T. F. Chan, "A multiphase level set framework for image segmentation using the mumford and shah model," *International Journal of Computer Vision*, vol. 50 no 3, pp. 271–293, 2002. [14](#), [33](#)
- [26] Hans-Hellmut Nagel and Wilfried Enkelmann, "An investigation of smoothness constraints for the estimation of displacement vector fields from image sequences," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, , no. 5, pp. 565–593, 1986. [14](#)
- [27] Alexander A Kirillov, *An introduction to Lie groups and Lie algebras*, vol. 113, Cambridge University Press Cambridge, 2008. [18](#), [21](#)
- [28] Elizabeth Louise Mansfield, *A practical guide to the invariant calculus*, vol. 26, Cambridge University Press, 2010. [18](#), [20](#), [28](#), [104](#)
- [29] Mario Ferraro and Terry M. Caelli, "Relationship between integral transform invariances and lie group theory," *Journal of the Optical Society of America A*, vol. 5 no 5, pp. 738–742, 1988. [18](#), [24](#)
- [30] Helmut Fischer and Helmut Kaul, *Mathematik für Physiker*, Springer, 1988. [18](#)
- [31] Peter J Olver et al., "Symmetry groups and group invariant solutions of partial differential equations," *J. Diff. Geom*, vol. 14, pp. 497–542, 1979. [20](#), [104](#)
- [32] Taichiro Kugo, *Eichtheorie*, 1997. [22](#)
- [33] Richard Phillips Feynman, Robert B. Leighton, and Matthew Sands, *The Feynman Lectures on Physics*, 1963. [22](#)
- [34] J. Bigun, *Vision with Direction A Systematic Introduction to Image Processing and Computer Vision*, Springer, 2006. [24](#), [71](#), [100](#)
- [35] E. Noether, "Invariante variationsprobleme," *Nachrichten Der Koeniglichen Gesellschaft Der Wissenschaften Zu Goettingen*, vol. Mathematische-Physikalische Klasse, pp. 235–257, 1918. [25](#), [99](#)
- [36] E. Noether, "Invariant variation problems," *Transport Theory and Statistical Problem*, vol. 1 (3), pp. 183–207, 1971. [25](#), [28](#), [99](#)
- [37] Friedhelm Kuypers, *Klassische Mechanik*, Wiley, 2005. [28](#)
- [38] A. Wedel, T. Pock, C. Zach, H. Bischof, and D. Cremers, "An improved algorithm for tv-l 1 optical flow," in *Statistical and Geometrical Approaches to Visual Motion Analysis*, pp. 23–45. Springer, 2009. [31](#), [35](#)

- [39] J. M. Fadili and G. Peyre, "Total variation projection with first order schemes," *IEEE Transactions on Image Processing*, vol. 20, no. 3, pp. 657–669, Mar. 2011. [31](#)
- [40] Kristian Bredies and Dirk Lorenz, "Mathematische bildverarbeitung," *Vieweg+ Teubner*, vol. 4, no. 6, pp. 12, 2011. [31](#), [32](#), [34](#), [112](#)
- [41] Stanley Osher and Leonid I Rudin, "Feature-oriented image enhancement using shock filters," *SIAM Journal on Numerical Analysis*, vol. 27, no. 4, pp. 919–940, 1990. [31](#)
- [42] Leonid I Rudin, "Images, numerical analysis of singularities and shock filters," 1987. [31](#)
- [43] Otmar Scherzer, Markus Grasmair, Harald Grossauer, Markus Haltmeier, and Frank Lenzen, "Variational methods in imaging, volume 167 of applied mathematical sciences," 2009. [31](#), [32](#), [34](#)
- [44] Eberhard Freitag and Rolf Busam, *Funktionentheorie 1. 3., neu bearb. und erw. Aufl*, Berlin. [32](#)
- [45] LC Evans and J Spruck, "Motion of level sets by mean curvature. i," in *Fundamental contributions to the continuum theory of evolving phase interfaces in solids*, pp. 328–374. Springer, 1999. [34](#)
- [46] Vincent Caselles, Ron Kimmel, and Guillermo Sapiro, "Geodesic active contours," *International Journal of Computer Vision*, vol. 22 no 1, pp. 61–79. [34](#)
- [47] Satyanad Kichenassamy, Arun Kumar, Peter Olver, Allen Tannenbaum, and Anthony Jr. Yezzi, "Conformal curvature flows: From phase transitions to active vision," *Archive for rational mechanics and analysis*, vol. 134 issue 3, pp. 275–301, 1996. [34](#)
- [48] B. K.P. Horn and B. G. Schunck, "Determining optical flow," *Artificial Intelligence*, vol. 17, issues 1-3, pp. 185–203, 1981. [35](#), [36](#), [37](#), [42](#)
- [49] B. D. Lucas, T. Kanade, et al., "An iterative image registration technique with an application to stereo vision.," in *IJCAI*, 1981, vol. 81, pp. 674–679. [35](#), [36](#), [37](#)
- [50] Michael J Black and Paul Anandan, "The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields," *Computer vision and image understanding*, vol. 63, no. 1, pp. 75–104, 1996. [35](#)
- [51] T. Bruhn, J. Weickert, and C. Schnörr, "Lucas/kanade meets horn/schunck: Combining local and global optic flow methods," *International Journal of Computer Vision*, vol. 61, no 3, pp. 211–231, 2004. [35](#)

- [52] Claudia Kondermann, Daniel Kondermann, and Christoph S. Garbe, "Post-processing of optical flows via surface measures and motion inpainting," *Lecture Notes in Computer Science*, vol. 5096, pp. 355–364, 2008. [35](#), [36](#)
- [53] Soo-Chul Han and Christine I Podilchuk, "Video compression with dense motion fields," *IEEE Transactions on Image Processing*, vol. 10, no. 11, pp. 1605–1612, 2001. [35](#)
- [54] Soo-Chul Han and Christine I Podilchuk, "Efficient encoding of dense motion fields for motion-compensated video compression," in *Image Processing, 1999. ICIP 99. Proceedings. 1999 International Conference on*. IEEE, 1999, vol. 1, pp. 84–88. [35](#)
- [55] M. Pollefeys and L. Koch, R. and Gool, "Self-calibration and metric reconstruction inspite of varying and unknown intrinsic camera parameters," *International Journal of Computer Vision*, vol. 32, no. 1, pp. 7–25, 1999. [35](#)
- [56] J. Repko and M. Pollefeys, "3d models from extended uncalibrated video sequences: Addressing key-frame selection and projective drift," in *3-D Digital Imaging and Modeling, 2005. 3DIM 2005. Fifth International Conference on*. IEEE, 2005, pp. 150–157. [35](#)
- [57] P. Tanskanen, K. Kolev, L. Meier, F. Camposeco, O. Saurer, and M. Pollefeys, "Live metric 3d reconstruction on mobile phones," in *Computer Vision (ICCV), 2013 IEEE International Conference on*. IEEE, 2013, pp. 65–72. [35](#)
- [58] N. Snavely, S. M. Seitz, and R. Szeliski, "Modeling the world from internet photo collections," *International Journal of Computer Vision*, vol. 80, no. 2, pp. 189–210, 2008. [35](#)
- [59] N. Snavely, S. M. Seitz, and R. Szeliski, "Photo tourism: exploring photo collections in 3d," in *ACM transactions on graphics (TOG)*. ACM, 2006, vol. 25, pp. 835–846. [35](#)
- [60] Young Min Kim, Christian Theobalt, James Diebel, Jana Kosecka, Branislav Misusik, and Sebastian Thrun, "Multi-view image and tof sensor fusion for dense 3d reconstruction," in *Proc. IEEE 12th Int Computer Vision Workshops (ICCV Workshops) Conf*, 2009, pp. 1542–1549. [35](#)
- [61] Mahua Bhattacharya and Arpita Das, "Registration of multimodality medical imaging of brain using particle swarm optimization," in *Proceedings of the first international conference on intelligent human computer interaction*. Springer, 2009, pp. 131–139. [35](#)
- [62] Anthoni Yezzi, Lilla Zöllei, and Tina Kapur, "A variational framework for joint segmentation and registration," *Proceedings of the IEEE Workshop on Mathematical Methods in Biomedical Image Analysis*, p. 44, 2001. [35](#)

- [63] Gerardo Hermosillo Valadez, *Variational Methods for Multimodal Image Matching*, Ph.D. thesis, École Doctorale Sciences et Technologies de l'Information et de la Communication, 2002. [36](#), [39](#), [40](#), [41](#)
- [64] Christophe Chéfd'Hotel, Gerardo Hermosillo, and Olivier Faugeras, "A variational approach to multi-modal image matching," in *Variational and Level Set Methods in Computer Vision, 2001. Proceedings. IEEE Workshop on*. IEEE, 2001, pp. 21–28. [36](#), [40](#)
- [65] David G. Lowe, "Distinctive image features from scale-invariant keypoints," *International Journal of Computer Vision*, 2004. [36](#)
- [66] Christopher M Bishop, *Neural networks for pattern recognition*, Oxford university press, 1995. [36](#)
- [67] Christoph Zetsche and Erhardt Barth, "Fundamental limits of linear filters in the visual processing of two-dimensional signals," *Vision research*, vol. 30, no. 7, pp. 1111–1117, 1990. [36](#)
- [68] Nils Papenberg, Andrés Bruhn, Thomas Brox, Stephan Didas, and Joachim Weickert, "Highly accurate optic flow computation with theoretically justified warping," *International Journal of Computer Vision*, vol. 67, no. 2, pp. 141–158, 2006. [37](#)
- [69] Christopher Zach, Thomas Pock, and Horst Bischof, "A duality based approach for realtime tv-l 1 optical flow," in *Joint Pattern Recognition Symposium*. Springer, 2007, pp. 214–223. [37](#)
- [70] Frederik Maes, Andre Collignon, Dirk Vandermeulen, Guy Marchal, and Paul Suetens, "Multimodality image registration by maximization of mutual information," *IEEE transactions on Medical Imaging*, vol. 16, no. 2, pp. 187–198, 1997. [39](#), [41](#), [49](#)
- [71] Alexis Roche, Gregoire Malandain, and Nicholas Ayache, *Unifying maximum likelihood approaches in medical image registration*, Ph.D. thesis, Inria, 1999. [39](#), [40](#), [49](#)
- [72] Heiko Hirschmüller, "Stereo processing by semiglobal matching and mutual information," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 30, no. 2, pp. 328–341, 2008. [39](#)
- [73] Robert R Edelman and Steven Warach, "Magnetic resonance imaging," *New England Journal of Medicine*, vol. 328, no. 10, pp. 708–716, 1993. [39](#)
- [74] Jiang Hsieh, *Computed tomography: principles, design, artifacts, and recent advances*, vol. 114, SPIE press, 2003. [39](#)

- [75] Y Vardi, LA Shepp, and L Kaufman, "A statistical model for positron emission tomography," *Journal of the American Statistical Association*, vol. 80, no. 389, pp. 8–20, 1985. [39](#)
- [76] Alexis Roche, Grégoire Malandain, Nicholas Ayache, and Xavier Pennec, "Multimodal image registration by maximization of the correlation ratio," 1998. [40](#)
- [77] A. Roche, G. Malandain, X. Pennec, and N. Ayache, "The correlation ratio as a new similarity measure for multimodal image registration," in *Medical Image Computing and Computer-Assisted Intervention*, pp. 1115–1124. Springer, 1998. [40](#), [49](#)
- [78] Olivier Faugeras and Renaud Keriven, *Variational principles, surface evolution, pde's, level set methods and the stereo problem*, IEEE, 2002. [41](#)
- [79] Thomas Netsch, Peter Rosch, Arianne van Muiswinkel, and Jürgen Weese, "Towards real-time multi-modality 3-d medical image registration," in *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*. IEEE, 2001, vol. 1, pp. 718–725. [41](#)
- [80] Pascal Cachier and Xavier Pennec, "3d non-rigid registration by gradient descent on a gaussian-windowed similarity measure using convolutions," in *Mathematical Methods in Biomedical Image Analysis, 2000. Proceedings. IEEE Workshop on*. IEEE, 2000, pp. 182–189. [41](#)
- [81] R. C. Hardie, M. T. Eismann, and G. L. Wilson, "Map estimation for hyper-spectral image resolution enhancement using an auxiliary sensor," *IEEE Transactions on Image Processing*, vol. 13, no. 4, 2004. [42](#), [43](#), [44](#), [45](#), [49](#), [50](#)
- [82] Russell C Hardie, Kenneth J Barnard, John G Bognar, Ernest E Armstrong, and Edward A Watson, "High-resolution image reconstruction from a sequence of rotated and translated frames and its application to an infrared imaging system," *Optical Engineering*, vol. 37, no. 1, pp. 247–260, 1998. [43](#), [44](#)
- [83] James B Campbell and Randolph H Wynne, *Introduction to remote sensing*, Guilford Press, 2011. [42](#)
- [84] FLIR, "<http://www.flir.com/home/>," . [43](#)
- [85] Havard Rue and Leonhard Held, *Gaussian Markov random fields: theory and applications*, CRC Press, 2005. [44](#)
- [86] Y. Zhang, "Spatial resolution enhancement for hyperspectral image based on wavelet bayesian fusion," *4th International Congress on Image and Signal Processing*, 2011. [45](#), [121](#)

- [87] J. Bigun and G. H. Granlund, "Optimal orientation detection of linear symmetry," *Proceedings of the IEEE First International Conference on Computer Vision : London, Great Britain*, pp. 433–438, 1987. [46](#), [71](#), [100](#)
- [88] Simon Baker, Daniel Scharstein, J. P. Lewis, Stefan Roth, Michael J. Black, and Richard Szeliski, "A database and evaluation methodology for optical flow," *International Journal of Computer Vision*, vol. 92, no. 1, pp. 1–31, 2011. [56](#), [72](#), [90](#)
- [89] André Duarte BL Ferreira, Paulo RO Nóvoa, and António Torres Marques, "Multifunctional material systems: a state-of-the-art review," *Composite Structures*, vol. 151, pp. 3–35, 2016. [63](#)
- [90] Shafi Ullah Khan and Jang-Kyo Kim, "Impact and delamination failure of multiscale carbon nanotube-fiber reinforced polymer composites: a review," *International Journal Aeronautical and Space Sciences*, vol. 12, no. 2, pp. 115–133, 2011. [63](#)
- [91] M Tehrani, AY Boroujeni, TB Hartman, TP Haugh, SW Case, and MS Al-Haik, "Mechanical characterization and impact damage assessment of a woven carbon fiber reinforced carbon nanotube–epoxy composite," *Composites Science and Technology*, vol. 75, pp. 42–48, 2013. [63](#)
- [92] Zhihang Fan, Michael H Santare, and Suresh G Advani, "Interlaminar shear strength of glass fiber reinforced epoxy composites enhanced with multi-walled carbon nanotubes," *Composites Part A: Applied Science and Manufacturing*, vol. 39, no. 3, pp. 540–554, 2008. [63](#)
- [93] Datong Wu and Gerd Busse, "Lock-in thermography for nondestructive evaluation of materials," *Revue générale de thermique*, vol. 37, no. 8, pp. 693–703, 1998. [63](#)
- [94] A. Spiessberger, C. and Gleiter and Busse G., "Data fusion of lock-in-thermography phase images," *9th International Conference on Quantitative InfraRed Thermography*, 2008. [63](#)
- [95] Carosena Meola, Giovanni Maria Carlomagno, Antonino Squillace, and Antonio Vitiello, "Non-destructive evaluation of aerospace materials with lock-in thermography," *Engineering Failure Analysis*, vol. 13, no. 3, pp. 380–388, 2006. [63](#)
- [96] Larry Wasserman, *All of Statistics, A Concise Course in Statistical Inference*, Springer, 2004. [67](#)
- [97] Charles W Misner, Kip S Thorne, and John Archibald Wheeler, *Gravitation*, Macmillan, 1973. [103](#)

- [98] Carlo Rovelli, *Quantum gravity*, Cambridge university press, 2007. [103](#)
- [99] Katrin Becker, Melanie Becker, and John H Schwarz, *String theory and M-theory: A modern introduction*, Cambridge University Press, 2006. [103](#)