The Intrinsic Structure of the Optic Flow Field

Luc Florack, Mads Nielsen

HAL Id: inria-00074327

https://inria.hal.science/inria-00074327

Submitted on 24 May 2006

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

The Intrinsic Structure of the Optic Flow Field

Luc FLORACK , Mads NIELSEN

N° 2350

Juillet 1994

PROGRAMME 4

Robotique,

image

et vision

Rapport de recherche

1994

The Intrinsic Structure of the Optic Flow Field

Luc FLORACK , Mads NIELSEN

Abstract:

In this paper, a generating equation for optic flow is proposed that generalises Horn and Schunck's Optic Flow Constraint Equation (OFCE). Whereas the OFCE has an interpretation as a *pointwise* conservation law, requiring grey-values associated with fixed-scale volume elements to be constant when co-moving with the flow, the new one can be regarded as a similar conservation requirement in which the flow elements have *variable scale* consistent with the field's divergence. Thus the equation gives rise to a definition of optic flow which is compatible with the scale-space paradigm.

We emphasise the gauge invariant nature of optic flow due to the inherent ambiguity of its components, i.e. the well-known aperture problem. Since gauge invariance is *intrinsic* to any definition of optic flow based solely on the data, it is argued that the gauge should be fixed on the basis of *extrinsic* knowledge of the image formation process and of the physics of the scene.

The optic flow field is replaced by an approximating field so as to allow for an order-by-order operational definition *preserving gauge invariance*, i.e. the approximation does not add spurious degrees of freedom to the field. One thus obtains a defining system of linear equations in the optic flow components up to arbitrary order, which remains decoupled from any physical considerations of gauge fixing. Such considerations are needed to derive a complementary system of gauge conditions that allows for a unique, physically sensible solution of the optic flow equations.

The theory is illustrated by means of several examples.

Key-words:

aperture problem, gauge invariance, gauge condition, Lie derivative, optic flow, scale-space.

(Résumé : tsvp)

La Structure Intrinsèque du Flot Optique

Résumé : Dans cet article nous proposons une équation définissant le flot optique, qui généralise l'Équation Contrainte du Flot Optique (ECFO) de Horn et Schunck. Alors qu'on interprète l'ECFO comme une loi de conservation *ponctuelle*, les valeurs de niveaux de gris associées à des éléments volumétriques d'échelle fixe devant rester constants sur les trajectoires associées au flot, la nouvelle équation peut être considérée comme une loi de conservation similaire dans laquelle les éléments du flot possèdent une *échelle variable* cohérente avec la divergence du champ. Donc, l'équation correspond à une définition du flot optique compatible avec le concept d'espace-échelle.

Nous mettons l'accent sur l'invariance de jauge du flot optique due à l'ambiguïté inhérente à ses composantes, c'est à dire le fameux problème d'ouverture. Puisque l'invariance de jauge est intrinsèque à toute définition du flot optique basée uniquement sur les données, on montre que le jauge doit être fixé d'après la connaissance *extrinsèque* du processus de formation de l'image et la physique de la scène.

Le flot optique est remplacé par un champ approché permettant une définition opérationnelle ordre-par-ordre, qui maintient l'invariance de jauge, c'est à dire que l'approximation n'ajoute pas de degrés de liberté impropres. On obtient donc un système d'équations linéaires en les composantes du champ du flot jusqu'à un ordre arbitraire, qui reste indépendant de toute considération physique aidant à fixer la jauge. De telles considérations sont nécessaires pour déduire un système complémentaire de conditions de jauge qui donne une solution unique et physiquement plausible des équations du flot optique.

La théorie est illustrée au moyen de plusieurs exemples.

Mots-clé : problème d'ouverture, invariance de jauge, condition de jauge, derivée de Lie, flot optique, espace-échelle.

1 Introduction

There exist many imaging modalities producing grey-level images of various characteristics. It is generally the case that coherent objects in the world lead to structural coherences in the image, and it is by virtue of this that we may actually hope to be able to infer useful information from such an image. In this paper we address spatiotemporal coherence induced by a dynamic scene, conventionally captured by the notion of *optic flow*.

The optic flow field is a vector field. This reflects the desire to link corresponding points (whatever these may be) separated by arbitrarily small temporal intervals. The motivation for this is of course that in the physical world such pointwise connections are actually meaningful; they may for example correspond to the true motion of material points on a physical body.

However, it is commonly appreciated that the image flow induced by some physical motion is intrinsically ambiguous due to the *aperture problem*; any hypothetical motion confined to an iso-grey-level contour is *a priori* feasible, and there is no image intrinsic evidence compelling for any particular solution. Put differently, if only for the image data, one may seek to solve for the *homotopy* that links spatial iso-grey-level contours over time *as a whole*, but one cannot hope to establish any *pointwise* connections between them. A convenient descriptor of such a homotopy is the *normal flow* (at least, if one tacitly admits arbitrary tangential flows so as to rule out its interpretation as a point correspondence). Normal flow is thus an intrinsic structural element of an image.

Another way of expressing the intrinsic degree of freedom that is left unconstrained by the data is to say that optic flow theory is a *gauge theory*[1]; optic flow, considered as a vector field, contains local degrees of freedom that do not manifest themselves in any observable way, hence can be fixed arbitrarily. By gauge invariance, any arbitrary *gauge condition* may equally well explain the *data*.

---

[1]A gauge theory is a theory which is characterised by a local invariance. Gauge theories are popular in physics because one can often simplify the description of a physical system by adding virtual degrees of freedom as well as a symmetry that effectively cancels their physical effect. Optic flow is conventionally described by such a gauge theory so as to obtain a linear model, as opposed to its nongauge equivalent formulation.

However, the data do not constitute a goal by themselves, but should enable us to infer structural information about the physical scene they are intended to describe. Eventually, we are not interested in optic flow as such, but in its relation to the dynamical structure of this scene. This is the reason why the specific choice of gauge *is* a crucial issue.

The physical nature of this gauge problem has not always been emphasised in the literature. Methods have been proposed to allow for the extraction of a unique vector field, without proper motivation[2] for the underlying implicit gauge choice. Even worse, the choices *implicitly* imposed may very well be inconsistent!

Although subject to immense effort, a satisfying *operational* definition of optic flow, even when disregarding the gauge problem, is by no means firmly established. A traditional, local approach is based on Horn and Schunck's "Optic Flow Constraint Equation" (OFCE) [1, 2, 3, 4, 5, 6, 7]. It defines the optic flow vector locally by means of a conservation law for local (or rather, punctal) grey-values; pixels (or any other fixed-scale samples, possibly weighted averages of neighbouring pixel values) are dragged along with the flow while preserving their grey-value attributes. Depending on the imaging modality, this view may not be consistent with scale-space theory, since migrating image volume elements, when viewed at fixed resolution, generally exhibit a nonvanishing divergence while moving along the flow, thereby changing their attributed grey-values. In order to prevent this, one needs to scale the aperture (or volume element), which defines local grey-values through local averaging, in accordance with the divergence of the optic flow field. This will be made more precise in the next section.

In summary, the purpose of this paper is threefold:

- To incorporate the notion of *resolution* into the definition of optic flow (scale-space).

- To explicitly separate the gauge invariant optic flow degrees of freedom as they are supported by the data from complementary, data independent, physical considerations that solve the aperture problem (gauge fixing).

---

[2]Uniqueness of the optic flow field cannot motivate its defining method in this respect.

- To provide an operational scheme for obtaining a (gauge invariant) approximating linear system for the data intrinsic, local optic flow degrees of freedom up to any order.

In other words, we de-emphasise the semantics of optic flow (and hence, the specific characteristics of the imaging modality); the focus is on the local structure of the optic flow field as far as it is captured by the image data *only*. We will point out how things can be made to work in practice, i.e. when given a particular imaging situation, but the details of this are beyond the scope of this paper.

2 Operational Definition of Optic Flow

In this section we consider a 1-parameter, infinitesimal transformation (parametrised by a formal parameter ε) which affects the image grey-value at a given point x^{μ} (assumed, without loss of generality, to be in the immediate neighbourhood of the origin $x^{\mu} = 0$) in a way that can be explained in terms of two independent actions:

1. a small, genuine spatiotemporal flow from $x^{\mu} - \delta x^{\mu}$ to x^{μ};

2. a small *independent* change of the grey-value at x^{μ}.

The first type of variation is conveniently captured by the so-called *Lie derivative* [8]; it is this one that we would like to relate to our definition of optic flow. We shall only consider first order variations in ε, i.e. $\mathcal{O}(\delta\varepsilon)$, and thereby obtain a *linear* model of optic flow; this should not be confused with the spatiotemporal differential order of the flow field we might be interested in. We will require no *a priori* restrictions on this.

Assumption 1 (General Local Image Variation)
Let $\phi : \mathbb{R}^{n+1} \to \mathbb{R}$ denote a scalar image, and let $\varepsilon \in \mathbb{R}$ be some formal parameter. Then it is assumed that an independent variation $\delta\varepsilon$ of ε induces the following variation of $\phi(x)$:

$$\delta_{\mathrm{tot}}\phi(x) = \delta\phi(x) + \overline{\delta}\phi(x) \quad (\mu = 0, \ldots, n) \quad .$$

Here, $\delta\phi(x) = \phi(x + \delta x) - \phi(x)$ is the infinitesimal variation[3] generated by a spatiotemporal vector field $v^\mu(x)$ ($\mu = 0, \ldots, n$), i.e. $\delta x^\mu = v^\mu(x)\delta\varepsilon$, and $\overline{\delta}\phi(x)$ denotes an independent, infinitesimal variation of the scalar image value at x^μ.

In other words, an infinitesimal image variation $\delta_{\mathrm{tot}}\phi(x)$ at some fixed point x^μ is accounted for by a pure spatiotemporal variation (i.e. the Lie derivative $\delta\phi(x)/\delta\varepsilon$ induced by the vector field $v^\mu(x)$), as well as by some *independent* "source" component $\overline{\delta}\phi(x)$. The latter is introduced so as to account for variations that cannot be explained in terms of a pure spatiotemporal flow. Thus the variation is of a rather general form.

We shall assume that the vector field $v^\mu(x)$ is smooth and everywhere (or at least within the closure of some region of interest) nonvanishing. Smoothness of $v^\mu(x)$ can always be assumed to hold *at some scale*, at least *conceptually* (because smooth functions are typically dense in some suitably chosen Hilbert space). This scale may however be inaccessibly small at some points in the image (say, smaller than the discretisation scale or spatiotemporal noise correlation width); in a *computational* sense we then have a "discontinuity". Anyway, by virtue of spatiotemporal coherence typical of practical imagery one may hope and expect this kind of discontinuity to be exceptional. Conceptual smoothness of $v^\mu(x)$ allows us to sensibly talk about its derivatives of arbitrary order—and we will freely do so—in every point in the image; its operational manifestation is a separate issue that will be discussed afterwards.

Assumption 1 is ambiguous; it relates two potential causes of local image deformation, one due to the action of a vector field ($\delta\phi(x)$) and one imposed in some deus-ex-machina fashion ($\overline{\delta}\phi(x)$). The first one is the kind we would like to make explicit here by trying to set up an operational method for defining the vector field $v^\mu(x)$ in terms of the local structure of the image $\phi(x)$. It is nontrivial to determine whether at some point x^μ an independent component $\overline{\delta}\phi(x)$ is really needed in order to explain local image change. Of course, for every choice of vector field $v^\mu(x)$ (in whatever way obtained or defined) one can always *define* $\overline{\delta}\phi(x)$ by the equation in Assumption 1 to complement the

---

[3]If $\phi(x)$ were a $\mathbf{C}^1(\mathrm{I\!R}^{n+1})$ function, we would have $\delta\phi(x) = v^\mu(x)\partial_\mu\phi(x)\delta\varepsilon$; however, no such smoothness assumptions are presumed here. We will assume $\phi(x)$ to be a function of polynomial growth (a rather weak restriction). The infinitesimal variation $\delta\phi(x)$ will be given a well-defined interpretation in terms of so-called *regular tempered distributions*.

$v^\mu(x)$-induced spatiotemporal variation, and thus explain the overall variation. But this is clearly not very useful. In some sense one would like the independent variation $\overline{\delta}\phi(x)$ to be as small as possible so as to at least remove ambigueties. The strategy followed in this paper will be (as is usually done) to initialise it to zero in order to unconfound it with the generator $v^\mu(x)$, and see how far we can get.

For the moment let us assume (cf. [1, 5]):

Assumption 2 ("Conservation of Topological Detail")
As Assumption 1, *but with*
$$\overline{\delta}\phi(x) = 0 \quad .$$

Note that the absence of an independent source contribution $\overline{\delta}\phi(x)$ is quite a severe constraint, since it implies *conservation of topological detail*: every point in the image domain is mapped in a one-to-one way onto a new one under the action of the vector field (i.e. by the ε-parametrised family of spacetime diffeomorphisms generated by $v^\mu(x)$), thereby preserving the topology of iso-intensity contours. Since the vector field is in turn defined by the image, one could say that all structural elements present in the image are preserved in spite of deformations. It is in this way that we would like to perceive of $\overline{\delta}\phi(x)$: as a source necessary to account for the creation or annihilation of topological detail.

We shall need one more assumption about the vector field in order to be able to interpret it as a genuine optic flow field. It emphasises the special role of the time dimension in spacetime. As in Newtonian mechanics, we will assume that space and time each have their own (flat) metric, but spacetime has none[4]; it is conceived of as a stratification of spatial slices over absolute time [9].

Assumption 3 (Temporal Gauge)
Cf. Assumption 1 *and* Assumption 2.

$$\forall x \in \mathbb{R}^{n+1} : \quad v^0(x) = 1 \quad .$$

---

[4] Although we defy existence of a spatiotemporal metric, we shall use $\eta_{\mu\nu}$ as shorthand notation for the *two independent metrics*, comprising the spatial metric, defining the spatial line element $ds^2 = \eta_{ij}dx^i dx^j$, and the temporal metric, defined by the interval $dt^2 = \eta_{00}dx^0 dx^0$. All other components η_{i0} and η_{0j} are always zero. Unless stated otherwise, we shall use Cartesian coordinates, so that $\eta_{\mu\nu} = 1$ if $\mu = \nu$, and zero otherwise.

This assumption implies that the flow is always transversal to fixed-time slices, in other words, that structural details persist over time, although they will generally deform. It therefore expresses *a priori temporal causality*. This transversality condition allows us to reparametrise the flow parameter ε so as to synchronise it with universal time $x^0 = t$, which is just Assumption 3.

The goal will be to find a unique vector field $v^\mu(x)$, and to express its local structure in terms of derivatives of the image. To guarantee well-posedness one needs to take derivatives in distributional sense, which in turn requires a smooth test function in conjunction to the image to be differentiated. The basic trick is well-known from the theory of *regular tempered distributions*, as formulated by Schwartz [10, 11].

Definition 1 (Spatiotemporal Derivatives of Φ)
Let $\gamma(x)$ be a smooth test function, whose derivatives all vanish sufficiently fast at the boundary of the image domain, and let $\gamma_{\mu_1...\mu_k}(x)$ denote its k-th order derivative with respect to $x^{\mu_1}, \ldots, x^{\mu_k}$, then the corresponding k-th order image derivative[5] is defined by the distribution

$$\Phi_{\mu_1...\mu_k} = (-)^k \int dx \, \phi(x) \, \gamma_{\mu_1...\mu_k}(x) \quad .$$

In practice one often decides on a fixed and severely limited set of test functions by imposing physical constraints. A rather minimal set may consist of scaled and shifted versions of a single, normalised Gaussian filter, as prescribed by the *scale-space paradigm*. Although not strictly required for the present discussion, let us agree on using this Gaussian "Schwartz space"; at some stage later on we will need specific properties of this. The shift and scale parameters (indicating the spatiotemporal location and scale at which to evaluate the derivatives) will not be made explicit in the notation; for the present discussion these are just free parameters.

Definition 2 (Normalised Gaussian)
Let Δ be a symmetric, positive definite $(n{+}1) \times (n{+}1)$ matrix, without space-time cross-talk, i.e. its general coordinate representation $\Delta^{\mu\nu}$, when decomposed into spatial and temporal components, Δ^{ij} and Δ^{00} respectively, is given by

$$\Delta^{\mu\nu} = \left(\begin{array}{c|c} \Delta^{00} & \varnothing \\ \hline \varnothing & \Delta^{ij} \end{array} \right) \quad ,$$

---

[5]Without loss of generality, "derivative" here means "derivative at the origin" $x^\mu = 0$.

with arbitrary positive Δ^{00} and symmetric, positive definite matrix Δ^{ij} *(all other $\Delta^{i0} = \Delta^{0j} = 0$). Then we define*

$$\gamma_\Delta(x) = \frac{1}{\sqrt{2\pi}^{n+1}} \frac{1}{\sqrt{\det \Delta}} \exp\left(-\frac{1}{2} x^\mu \Delta^{\text{inv}}_{\mu\nu} x^\nu\right) \quad .$$

In a suitable orthonormal basis we can take the scale matrix to be diagonal, $\Delta = \text{diag}\{t^0, \ldots, t^n\}$, where $t^\mu = (\sigma^\mu)^2$ ($\mu = 0, \ldots, n$). In this case we shall simply write $\gamma(x)$ and omit the matrix subscript Δ.

The zero-scale limit of this normalised Gaussian is well-defined "under the integral" of Definition 1, provided the data $\phi(x)$ are sufficiently smooth. In that case this limit will bring us back to the conventional definition of a derivative of $\phi(x)$, taken at $x^\mu = 0$ (this conventional limiting behaviour was, of course, what Schwartz had in the back of his mind when he wrote down his theory).

It is henceforth implied that one takes derivatives with respect to appropriately scaled coordinates $\tilde{x}^\mu = x^\mu/\sigma^\mu$ (no summation convention here), if σ^μ is the scale parameter associated with the μ-th dimension (typically one takes these to be all equal in the spatial domain to enforce isotropy; of course one will need an independent one for the time dimension). Similarly we will consider only scaled frequency coordinates in Fourier space: $\tilde{\omega}^\mu = \sigma^\mu \omega^\mu$. To facilitate notation, we will forget about the $\tilde{}$ and write x^μ instead of $\tilde{x}^\mu$, etc. In other words:

Assumption 4 (Choice of Spatiotemporal Units)
Every spatiotemporal quantity is henceforth measured relative to implicit inner scale.

Put differently, we henceforth choose our length and time units such that we can formally put all σ^μ, $\mu = 0, \ldots, n$, equal to unity. Moreover, we may, without loss of generality, allow for arbitrary (not necessarily Cartesian) coordinate systems, provided we consider all derivatives to be covariant (these will still commute since space and time are Euclidean). The metric components $\eta_{\mu\nu}$ can be absorbed into the definition of $\Delta^{\text{inv}}_{\mu\nu}$ (Definition 2).

Since we view the dynamical structure of the data $\phi(x)$ by means of *fixed* spacetime filters, we will assume the following.

Assumption 5 (Static Filters versus Dynamic Data)

$$\delta\gamma_{\mu_1\ldots\mu_k}(x) = 0 \quad,$$

i.e. the filters are not subject to the external flow (it is in this sense that they are called "fixed"; they are *not* time independent!). They comprise the "sensorium" [12] probed by the stimulus. Furthermore we define, for each derivative filter $\gamma_{\mu_1\ldots\mu_k}(x)$, a corresponding *filter current density*.

Definition 3 (Filter Current Density)
Notation as in Assumption 1 *and* Definition 1. *Define the filter current density* $j^\mu(x)$ *corresponding to the filter* $\gamma(x)$ *and the vector field* $v^\mu(x)$ *as follows:*

$$j^\mu(x) = \gamma(x)v^\mu(x) \quad.$$

More generally, for any order $k \in \mathbb{N}$:

$$j^\mu_{\mu_1\ldots\mu_k}(x) = \gamma_{\mu_1\ldots\mu_k}(x)v^\mu(x) \quad.$$

The k-th order filter current density is really a *density* in the tensorial sense, a property it inherits from the basic filter $\gamma(x)$ [8] (a scalar density is the geometrical object corresponding to an $(n+1)$-form). Note that the zero-components equal the basic Gaussian derivative filters by our causality gauge $v^0(x) \equiv 1$ (Assumption 3).

The Lie derivative of the image with respect to the optic flow vector can be made well-posed by formulating it as a regular tempered distribution, just as we did in Definition 1 for ordinary derivatives.

Proposition 1 (Lie Derivative of Φ)
Notation as in Assumption 1 *and* Definition 1. *Define furthermore*

$$\delta\Phi = \int dx\, \delta\phi(x)\, \gamma(x) \quad.$$

This variation has an equivalent, dual interpretation, viz.

$$\delta\Phi = \int dx\, \phi(x)\, \delta^T\gamma(x) \quad,$$

in which $\delta^T\gamma(x)$ *is defined by*

$$\delta^T\gamma(x) = -\partial_\mu j^\mu(x)\, \delta\varepsilon \quad.$$

We thus have a dual interpretation[6] of $\delta\Phi$: either we *track* a flow element and extract samples by means of a *passive* aperture, or we take a fixation point *not co-moving* with the flow (a "mark on the screen"), and take samples by means of an *active* aperture that continuously adapts its profile to the flow so as to yield the same result[7]. In the latter view the flow is hypothetically carried over into the sensorium, in which the filters "live". One then has to subject the filters to the *opposite* flow, i.e. take their Lie derivative w.r.t. $w^\mu(x) = -v^\mu(x)$. Since this also toggles the sign of the temporal gauge component, $w^0(x) = -1$, the "arrow of time" induced by $w^\mu(x)$ actually points *backward in time*. To "fix" this apparent time reversal conceptually, we can give it a causal interpretation simply by noting that the Lie derivative w.r.t. $w^\mu(x)$ corresponds to *minus* the Lie derivative w.r.t. $v^\mu(x)$:

$$\frac{\delta^T}{\delta\varepsilon} = -\frac{\delta}{\delta\varepsilon} \quad . \tag{1}$$

The corresponding conjugate variation of the filter is, apart from the minus sign, apparently given by the divergence of the filter current density. Note that the transformation of the filter $\gamma(x)$ differs *qualitatively* from that of a smooth scalar $\phi(x)$: this is a manifestation of the fact that $\gamma(x)$ is a scalar density rather than an ordinary scalar.

Let us turn to the proof of Proposition 1.

Proof 1 (Proposition 1)
The most tempting way to proof Proposition 1 might be to write $\delta\phi(x) = \partial_\mu\phi(x)v^\mu(x)\delta\varepsilon$, and then to carry out a partial integration. This, however, presupposes that the image function $\phi(x)$ is differentiable in conventional sense. But recall that conventional differentiation is ill-posed and even operationally ill-defined. We cannot impose smoothness constraints on our acquisition data; we have to accept them the way they come. This is the very reason for introducing regular tempered distributions in the first place!

---

[6]Although in this paper duality is merely a matter of concept, it may well have implementational relevance for a (biological or artificial) vision system (one may think of a continuous chain of "Reichardt detectors" [13]).

[7]A different strategy based on this duality principle—with constant, *parametric* $v^\mu = (1; v^i)$—is used in [14].

A proof that makes no regularity assumptions about the image goes as follows. Consider the following reparametrisation:

$$y^\mu = x^\mu + v^\mu(x)\,\delta\varepsilon \quad . \tag{2}$$

Its inverse is (discarding $\mathcal{O}(\delta\varepsilon^2)$ from now on)

$$x^\mu = y^\mu - v^\mu(y)\,\delta\varepsilon \quad . \tag{3}$$

This local reparametrisation induces a Jacobian given by

$$J^\mu_\nu(y) \stackrel{\text{def}}{=} \frac{\partial x^\mu}{\partial y^\nu} = \delta^\mu_\nu - \partial_\nu v^\mu(y)\,\delta\varepsilon \stackrel{\text{def}}{=} \delta^\mu_\nu - \delta J^\mu_\nu(y) \quad , \tag{4}$$

the determinant of which will show up in the reparametrised integral:

$$J(y) \stackrel{\text{def}}{=} \det J^\mu_\nu(y) = 1 - \partial_\mu v^\mu(y)\,\delta\varepsilon \stackrel{\text{def}}{=} 1 - \delta J(y) \quad . \tag{5}$$

Note that it deviates from unity if the vector field has nonvanishing divergence, causing local volume elements to expand.

Applied to Proposition 1, changing the symbol for the dummy variable y back into x), reparametrisation yields

$$\delta\Phi = -\int dx\,\phi(x)\,(\delta J(x) + \delta x \cdot \nabla)\,\gamma(x) \stackrel{\text{def}}{=} \int dx\,\phi(x)\,\delta^T\gamma(x) \quad . \tag{6}$$

The extra contribution $\delta J(x)$ to the transformation law of the filter is typical of a density. Indeed, it can readily be seen that the combined effect of the two variations in (6) precisely boils down to Proposition 1.This completes the proof of Proposition 1.

To reveal the local structure of the Lie derivative $\delta\Phi$ it is useful to consider its spatiotemporal derivatives; these should reflect conventional derivatives of $\delta\phi(x)$ in case this were a sufficiently smooth function.

Proposition 2 (Spatiotemporal Derivatives of $\delta\Phi$)
Let the spatiotemporal derivatives of $\delta\Phi$ be defined as

$$\partial_{\mu_1...\mu_k}\delta\Phi = (-)^k \int dx\,\delta\phi(x)\,\gamma_{\mu_1...\mu_k}(x) = (-)^k \int dx\,\phi(x)\,\delta^T\gamma_{\mu_1...\mu_k}(x) \quad .$$

Then we have (cf. Definition 3*)*

$$\delta^T\gamma_{\mu_1...\mu_k}(x) = -\partial_\mu j^\mu_{\mu_1...\mu_k}(x)\,\delta\varepsilon \quad .$$

Proof 2 (Proposition 2)
This proceeds along the same lines as for Proposition 1.

Note that spatiotemporal derivatives do not commute with the Lie derivative; Proposition 2 gives us the spatiotemporal derivative of the Lie derivative of Φ, as opposed to the following.

Definition 4 (Lie Derivative of $\Phi_{\mu_1...\mu_k}$)
The Lie derivative of $\Phi_{\mu_1...\mu_k}$ is defined as

$$\delta\Phi_{\mu_1...\mu_k} = (-)^k \int dx\, \phi(x)\, \partial_{\mu_1...\mu_k} \delta^T \gamma(x) \quad .$$

This is a good definition, because it has the correct classical limit if $\phi(x)$ is smooth, viz. $\delta\phi_{\mu_1...\mu_k}(x = 0)$. Another important observation is that, in the conjugate view of Definition 4 when $\gamma(x) \mapsto \tilde{\gamma}(x) = \gamma(x) + \delta^T \gamma(x)$, *normalisation of all derivatives of $\tilde{\gamma}(x)$ is preserved*. To see this, recall that, as a consequence of normalisation of $\gamma(x)$, all filter derivatives are simultaneously normalised according to:

$$N^{\rho_1...\rho_k}_{\mu_1...\mu_k} \stackrel{\text{def}}{=} \frac{(-)^k}{k!} \int dx\, x^{\rho_1} \ldots x^{\rho_k} \gamma_{\mu_1...\mu_k}(x) = \mathcal{S}\left\{\delta^{\rho_1}_{\mu_1} \ldots \delta^{\rho_k}_{\mu_k}\right\} \int dx\, \gamma(x) = \mathcal{S}\left\{\delta^{\rho_1}_{\mu_1} \ldots \delta^{\rho_k}_{\mu_k}\right\} \quad,$$

(7)

($\mathcal{S}\{.\}$ denotes index symmetrisation, and δ^{ρ}_{μ} is the Kronecker symbol). After deformation $\gamma(x) \mapsto \tilde{\gamma}(x)$ one might expect a violation of this due to an extra term $\delta N^{\rho_1...\rho_k}_{\mu_1...\mu_k}$. However,

$$\delta N^{\rho_1...\rho_k}_{\mu_1...\mu_k} \stackrel{\text{def}}{=} \frac{(-)^k}{k!} \int dx\, x^{\rho_1} \ldots x^{\rho_k} \partial_{\mu_1...\mu_k} \delta^T \gamma(x) = \mathcal{S}\left\{\delta^{\rho_1}_{\mu_1} \ldots \delta^{\rho_k}_{\mu_k}\right\} \int dx\, \delta^T \gamma(x) = 0 \quad .$$

(8)

The last equality follows from replacing $\delta^T \gamma(x)$ by $-\partial_\mu j^\mu(x)\delta\varepsilon$, and using Stokes' theorem (or the "divergence theorem") [15]:

$$\int_\Omega dx\, \partial_\mu j^\mu(x) = \oint_{\partial\Omega} dS_\mu\, j^\mu(x) \quad,$$

(9)

in which dS_μ denotes an outward normal boundary element for the region-with-boundary Ω. We have assumed that the filter current density $j^\mu(x)$ vanishes

on the boundary or, if the image domain is all of $\mathbb{R}^{n+1}$, that $j^\mu(x)$ decreases sufficiently fast towards infinity. This will indeed be the case if $v^\mu(x)$ is a function of polynomial growth. As a consequence, we see that normalisation is indeed unaffected.

The image induced intrinsic vector field (or the *equivalence class* of vector fields) we are looking for will be the one that nullifies the variation of Φ along its corresponding flow.

Definition 5 (Optic Flow)
See Assumption 3. *The optic flow vector field is defined, modulo gauge transformations, by*

$$\partial_{\mu_1\ldots\mu_k}\delta\Phi = 0 \quad \text{for all} \quad k \in \mathbb{N} \quad .$$

Definition 5 differs from Horn and Schunck's OFCE; it does *not* require grey-values attributed to pixels (and assumed to be dragged along with the flow field) to remain constant. Rather, it allows grey-values to change due to the field's divergence. This can best be appreciated from the derivation of Proposition 1, and in particular equation (6): it is *not* the variation $\delta\phi(x)$ of the (high-resolution) *data* $\phi(x)$ which is required to vanish, but the effect $\delta\Phi$ it induces on the *functional* Φ, in other words, on the *measurements* of $\phi(x)$. See also figure 1.

Depending on the actual task one aims to accomplish, one may want to gain insight into the local characteristics of the vector field beyond lowest order. For example, in the case of optic flow in real world movies, first order properties of the vector field may reveal relevant information such as qualitative shape properties, surface slant [16], time-to-contact, etc. Unlike first order properties, second order is quantitatively related to intrinsic surface properties of the objects projecting to the image plane [17], etc.

Let us consider the M-th order case. To this end we make a *formal expansion* of $v^\mu(x)$ near the origin, and truncate it so as to obtain an M-th order polynomial approximation $v_M^\mu(x)$. This polynomial is intended to capture a finite number of local degrees of freedom of the vector field.

Definition 6 (M-th Order Formal Expansion)
The formal expansion of order M of the vector field $v^\mu(x)$ at $x^\mu = 0$, denoted

$v_M^\mu(x)$, *is an M-th order polynomial*

$$v_M^\mu(x) = \sum_{l=0}^{M} \frac{1}{l!} \partial_{\rho_1 \dots \rho_l} v_0^\mu \, x^{\rho_1} \dots x^{\rho_l} \quad ,$$

the coefficients of which may depend on M.

The finite set of coefficients $\partial_{\rho_1 \dots \rho_l} v_0^\mu$ (to be defined later on so as to *approximate* $\partial_{\rho_1 \dots \rho_l} v^\mu(x)$ at the origin $x^\mu = 0$, for $l = 0, \dots, M$), corresponds exactly to the degrees of freedom we are looking for. It is important to appreciate that we do *not* require the M-th order polynomial $v_M^\mu(x)$ to be the M-th order Taylor polynomial of $v^\mu(x)$. That would be too much to ask for in view of our operational definition for $v_M^\mu(x)$ to be given below. Rather, the thing we hope to accomplish is that the coefficients of $v_M^\mu(x)$ will *approximate* the corresponding Taylor coefficients of $v^\mu(x)$ as we increase M: by definition we have $v_\infty^\mu(x) = \lim_{M \to \infty} v_M^\mu(x) = v^\mu(x)$.

Note that the divergence $\operatorname{div} v_M(x)$ is given by

$$\partial_\mu v_M^\mu(x) = \sum_{l=0}^{M-1} \frac{1}{l!} \partial_{\rho_1 \dots \rho_l \mu} v_0^\mu \, x^{\rho_1} \dots x^{\rho_l} \quad . \tag{10}$$

Replacing $v^\mu(x)$ by $v_M^\mu(x)$ according to Definition 6 yields the following.

Result 1 $(\partial_{\mu_1 \dots \mu_k} \delta_M \Phi)$
See Definition 3, Proposition 1, *and* Proposition 2. *Writing* δ_M *instead of* δ *to remind us of the fact that we are dealing with the M-th order formal expansion* $v_M^\mu(x)$ *instead of* $v^\mu(x)$ *itself, we have*

$$\partial_{\mu_1 \dots \mu_k} \delta_M \Phi = - \sum_{l=0}^{M} \partial_{\rho_1 \dots \rho_l} v_0^\mu \int dx \, \phi(x) \, \partial_\mu \Gamma^{\rho_1 \dots \rho_l}_{\mu_1 \dots \mu_k}(x) \, \delta\varepsilon \quad ,$$

in which the effective filters $\Gamma^{\rho_1 \dots \rho_l}_{\mu_1 \dots \mu_k}(x)$ *are given by*

$$\Gamma^{\rho_1 \dots \rho_l}_{\mu_1 \dots \mu_k}(x) = \frac{(-)^k}{l!} \, \gamma_{\mu_1 \dots \mu_k}(x) \, x^{\rho_1} \dots x^{\rho_l} \quad .$$

One may raise an objection to the way Result 1 is presented. Recall that the Gaussian derivatives $(-)^k \gamma_{\mu_1...\mu_k}(x)$ $(k \in \mathbb{Z}_0^+)$ form a *complete* family [18]. This implies that the set of filters $\Gamma^{\rho_1...\rho_l}_{\mu_1...\mu_k}(x)$ $(k, l \in \mathbb{Z}_0^+)$ is *overcomplete*; it should be possible to express the members with $l \neq 0$ in terms of those with $l = 0$ (i.e. the good old plain Gaussian derivatives). Since overcompleteness is merely a technical detail, let it suffice at this point to state that $\Gamma^{\rho_1...\rho_l}_{\mu_1...\mu_k}(x)$ can indeed be expressed as a linear combination of Gaussian derivative filters $(-)^m \gamma_{\mu_1...\mu_m}(x)$ of various orders m up to order $k + l$, inclusive. The details are left to appendix A (in particular Result 3).

Given the variations on the l.h.s. of the equation in Result 1, we have accomplished a set of linear equations in the unknowns $\partial_{\rho_1...\rho_l} v_0^\mu$ (one for each possible realisation of indices $\mu_1, \ldots, \mu_k$, with $k = 0, \ldots, N$ for some yet unspecified order N). The coefficients in this equation can be computed by straightforward linear filtering of the image $\phi(x)$ using the filter profiles $\Gamma^{\rho_1...\rho_l}_{\mu_1...\mu_k}(x)$ as defined in Result 1. The inhomogeneous part $\partial_{\mu_1...\mu_k} \delta_M \Phi$ reveals the k-th order local structure of the variation of Φ along the flow induced by the "truncated" vector field $v_M^\mu(x)$.

The problem is that we don't know the l.h.s. in Result 1. Let $\delta \Phi = \delta_M \Phi + \overline{\delta}_M \Phi$. Since we know that $\overline{\delta} \Phi = \lim_{M \to \infty} \overline{\delta}_M \Phi = 0$, it makes sense to "approximate" the true optic flow field $v^\mu(x)$ by $v_M^\mu(x)$, as follows.

Definition 7 (M-th Order Approximated Optic Flow)
The M-th order approximated optic flow vector field is defined, modulo gauge transformations, by

$$\partial_{\mu_1...\mu_k} \delta_M \Phi = 0 \quad \text{for all} \quad k = 0, \ldots, M \quad .$$

Combined with Result 1 (and Appendix A), we thus obtain a *homogeneous* system of *linear* equations, expressing M-th order optic flow in terms of $(2M + 1)$-st order local image structure, in which of course everything (or nothing...) happens in the null space. It is, however, "slightly inconsistent", since it presupposes $\overline{\delta}_M \Phi = 0$. In fact, it amounts to imposing the gauge condition $\partial_{\mu_1...\mu_{M+1}} v^\mu(x) = 0$, which is generally a *physically ill-motivated gauge*: it asserts that the optic flow generator $v^\mu(x)$ is of the form $v_M^\mu(x)$, which is too restrictive in principle. However, by virtue of the limiting behaviour for $M \to \infty$, there is good hope that we indeed approximate the exact system of

Definition 5 in this way. But we have to take care not to introduce spurious degrees of freedom by adding any more gauge conditions. In particular, we should try to maintain the original gauge invariance for our *approximated* optic flow field $v_M^\mu(x)$. Allowing arbitrary orders of differentiation in Definition 7 would certainly break this invariance in the generic case (i.e. the case when $v^\mu(x) \neq v_M^\mu(x)$)! This is why one needs to limit the highest order to $k = M$.

As opposed to conventional schemes based on M-fold implicit differentiation of the OFCE, *every k-th order subset* of Definition 7 contains M-th order components of the approximated optic flow field. An important thing to keep in mind is that the degrees of freedom of the approximated optic flow field, i.e. the coefficients $\partial_{\mu_1 \ldots \mu_l} v_0^\mu$, depend, although not explicitly indicated by the notation, on the order M of approximation. In other words, the polynomial approximation $v_{M+1}^\mu(x)$ is a refinement of $v_M^\mu(x)$ in the sense that *all* its coefficients are refined. Hence, it is not a Taylor polynomial of $v^\mu(x)$.

Definitions 5 and 7 always admit trivial vector fields that render these self-consistent. Viz. take $v^\mu(x) = 0$ ($v_M^\mu(x) = 0$ respectively) identically, and the homogeneous equations will be trivially satisfied. Of course, these are not "interesting" solutions by any means, and in fact we already discarded them by requiring the vector field to be everywhere nonvanishing. So let us look for admissible vector fields. Since the solution space is a linear space, we can always arrange things such that (for nontrivial solutions) $v^0(x) = 1$ (Assumption 3). For the M-th order approximation $v_M^\mu(x)$ this becomes

Assumption 6 (Temporal Gauge for M-th Order Approximated Optic Flow)

See Assumption 3 *and* Definition 7. *Requiring the temporal gauge of* Assumption 3 *to hold for* $v_M^\mu(x)$, *one finds*

$$\begin{cases} v_0^0 & = & 1 \quad , \\ \partial_{\mu_1 \ldots \mu_k} v_0^0 & = & 0 \quad \textit{for all} \quad k = 1, \ldots, M \quad \textit{and} \quad \mu_1, \ldots, \mu_k = 0, \ldots, n \quad . \end{cases}$$

The ambiguety of the solution, when not subjected to any gauge condition (including the temporal gauge), generalises the classical "aperture problem" to the spatiotemporal domain; under the temporal gauge it boils down to the classical one [1], leaving only the spatial optic flow vector $v^i(x)$ ungauged.

To appreciate why k must not exceed M in Definition 7, let us return to our inverse problem. In Result 1 we can choose any M, depending on the differential

order of the flow field we would like to retrieve. In $n + 1$ dimensions, the total number of degrees of freedom $\mathcal{F}_M$ contained in $v_M^\mu(x)$ (i.e. $v_0^\rho, \ldots, \partial_{\mu_1 \ldots \mu_M} v_0^\rho$) equals

$$\mathcal{F}_M = (n + 1) \sum_{j=0}^{M} \binom{j + n}{n} \quad . \tag{11}$$

So this is the number of variables in our linear inverse problem. If not for the constraint on k, we could also choose any fixed order N and consider all cases with $k = 0, \ldots, N$. This would determine our total number $\mathcal{E}_N$ of equations, viz.

$$\mathcal{E}_N = \sum_{j=0}^{N} \binom{j + n}{n} \quad . \tag{12}$$

For complete solvability one needs at least $\mathcal{E}_N \geq \mathcal{F}_M$, which could, in principle, always be achieved by a suitable choice of M and N (at this level of rigor we ignore rank conditions for the set of equations). If this condition is not satisfied, one is left with (at least) $\mathcal{G}_{MN} = \mathcal{F}_M - \mathcal{E}_N$ undetermined gauge degrees of freedom.

As an example, take the $(1 + 1)$-dimensional case. For $N = 2$ and $M = 1$ we have $\mathcal{E}_2 = 1 + 2 + 3 = 6$ equations in $\mathcal{F}_1 = 2 \times (1 + 2) = 6$ unknowns; 2 of these unknowns parametrise a local shift, and 4 of them capture a local linear deformation (since $v_1^\mu(x)$ is an "affine approximation" of $v^\mu(x)$). There are apparently no gauge degrees of freedom left: $\mathcal{G}_{12} = \mathcal{F}_1 - \mathcal{E}_2 = 6 - 6 = 0$. *Where did they go?* In other words: *where did we fix the gauge?* As argued before, gauge fixing is an *extrinsic* problem; its motivation cannot be based on the image data *sec*, but only on a priori knowledge of the physics underlying the image formation process. The example is a typical instance in which we have *implicitly* fixed the gauge in an *ad hoc* fashion. The result is a unique optic flow field lacking any physical relevance (unless, by some incredibly fortunate coincidence, the ad hoc gauge happens to be close to a physically sensible one). It has not always been appreciated in the literature that one cannot solve the aperture problem intrinsic to the data in a physically sensible way merely by making some formal expansion of the optic flow field! An instructive example of a physically well-motivated gauge condition in medical imaging, using conservation of mass and the physics of X-ray projection, is provided by Amini in [7].

Indeed, self-consistency requires that M be equal to N, so that we are always left with our *full* gauge problem, merely split up in a countable, order-by-order way.

Proposition 3 (Self-Consistency)
Notation as in Definition 7 *and equations* (11) *and* (12). *The M-th order optic flow field $v_M^{\prime\mu}(x)$ as determined by the second system of equations in* Definition 7, *with $k = 0, \ldots, N$, is gauge invariant for arbitrary $v_M^{\prime\mu}(x)$, if and only if $N = M$. In the generic case one is then left with*

$$\mathcal{G}_M = n \sum_{j=0}^{M} \binom{j+n}{n}$$

undetermined local gauge degrees of freedom. Fixing the a priori temporal gauge according to Assumption 6, *this number reduces to*

$$\mathcal{G}_M^{\text{t.g.}} = (n-1) \sum_{j=0}^{M} \binom{j+n}{n} \quad .$$

This result should be intuitively evident; it states that only 1 (without the temporal gauge, otherwise 2) of the $n+1$ components of $v^{\prime\mu}(x)$ can be determined given only a single equation, viz. the vanishing of the Lie derivative. This observation holds on a full neighbourhood of the origin, and Proposition 3 expresses this in terms of the equivalent, countable set of degrees of freedom contained in the coefficients of $v_M^{\prime\mu}(x)$. Note that after temporal gauge fixing, there are no optic flow gauge variables left in the trivial $n = 1 + 1$-dimensional case; the spatial iso-grey-level contours degenerate to points and hence cannot "hide" any flow components.

Proof 3 (Proposition 3)
First disregard the temporal gauge. From a trivial counting argument it is clear that there are *at least* $\mathcal{G}_M = \mathcal{F}_M - \mathcal{E}_M$ undetermined degrees of freedom in the defining system for $v_M^{\prime\mu}(x)$, Definition 7. It remains to be shown that, if one considers a *generic* image, one for which the rank of the system is not affected by an arbitrarily small, independent variation of all image derivatives, the number of gauge degrees of freedom can neither exceed $\mathcal{G}_M$. In other words, we must show that the system has full rank $\mathcal{E}_M$. (Of course, one

can always construct artificial, nongeneric images, for which the system will have less than maximal rank; a trivial example is a "null image", yielding zero rank. Generic images are needed to avoid intrinsic ambiguities due to "lack of structure".) To see that a generic image indeed induces a maximal rank system for any M, consider the k-th subset of equations $\partial_{\mu_1...\mu_k} \delta_M \Phi = 0$ for some fixed $k = 0, \ldots, M$ (hint: it may be helpful to follow the general argument while looking at the Examples 4, 5 and 6). The explicit expression in Result 1 and the details in Appendix A reveal that the highest order filters needed in these equations, viz. $\partial_\mu \Gamma^{\rho_1...\rho_M}_{\mu_1...\mu_k}(x)$, all have differential order $M + k + 1$, and are mutually independent for different choices of $\mu_1, \ldots, \mu_k$. Consequently, the image derivatives of order $M + k + 1$, $\partial_{\mu\mu_1...\mu_k} \partial^{\rho_1...\rho_M} \Phi$, multiplying the highest order variables $\partial_{\rho_1...\rho_M} v_0^\mu$ in this subset of equations, guarantee that the subset is independent of all subsets of lower orders. Hence all $\mathcal{E}_M$ equations are indeed generically independent.

Finally, incorporating the temporal gauge means fixing $\mathcal{E}_M$ of the $\mathcal{G}_M = \mathcal{F}_M - \mathcal{E}_M$ unknowns, yielding $\mathcal{G}_M^{\text{t.g.}} = \mathcal{F}_M - 2\mathcal{E}_M$ left-over gauge variables.

So far for the theoretical framework. To gain better understanding of the theory sketched in this section, let us consider some examples.

3   More Examples

The two following examples illustrate the overcompleteness of the filter set defined in Result 1 and Result 2. Example 1 shows the special cases in which the filters carry only lower or upper indices. Example 2 shows some lowest order mixed filters.

Example 1 (Special Cases of $\Gamma^{\rho_1\cdots\rho_l}_{\mu_1...\mu_k}(x)$: $\Gamma_{\mu_1...\mu_k}(x)$ and $\Gamma^{\rho_1\cdots\rho_l}(x)$)

[l=0] $\Gamma(x) = \gamma(x)$, $\Gamma_\mu(x) = -\partial_\mu \gamma(x)$, $\Gamma_{\mu\nu}(x) = \partial_\mu \partial_\nu \gamma(x)$,
   in general: $\Gamma_{\mu_1...\mu_k}(x) = (-)^k \partial_{\mu_1} \ldots \partial_{\mu_k} \gamma(x)$,

[k=0] $\Gamma^\rho(x) = -\partial^\rho \gamma(x)$, $\Gamma^{\rho\sigma}(x) = \frac{1}{2}\eta^{\rho\sigma}(x)\gamma(x) + \frac{1}{2}\partial^\rho \partial^\sigma \gamma(x)$,
   in general: $\Gamma^{\rho_1\cdots\rho_l}(x) = \frac{1}{l!}(-i)^l \mathcal{H}^{\rho_1\cdots\rho_l}(-i\nabla)\gamma(x)$.

Example 2 (Some Lowest Order Mixed Filters $\Gamma^{\rho_1\cdots\rho_l}_{\mu_1...\mu_k}(x)$)

[k=1,l=1] $\Gamma_\mu^\rho(x) = \partial^\rho \partial_\mu \gamma(x) + \delta_\mu^\rho \gamma(x)$,

[k=2,l=1] $\Gamma_{\mu\nu}^\rho(x) = -\partial_\mu \partial_\nu \partial^\rho \gamma(x) - \delta_\mu^\rho \partial_\nu \gamma(x) - \delta_\nu^\rho \partial_\mu \gamma(x)$,

[k=3,l=1] $\Gamma_{\mu\nu\rho}^\sigma(x) = \partial_\mu \partial_\nu \partial_\rho \partial^\sigma \gamma(x) + \delta_\mu^\sigma \partial_\nu \partial_\rho \gamma(x) + \delta_\nu^\sigma \partial_\mu \partial_\rho \gamma(x) + \delta_\rho^\sigma \partial_\mu \partial_\nu \gamma(x)$.

The general case is given in Appendix A, Result 3.

The following example illustrates the conjugate view of Proposition 1. It shows that, when monitoring a stimulus while tracking the flow by means of a passive aperture, one obtains a variation $\delta\Phi$ which may equally well be explained as an observation of the stimulus at a fixation point, carried out by an active aperture that suitably adapts its shape to the flow.

Example 3 (The Conjugate View: Filter Adaptation)
Suppose that the (spatial) flow field $v^i(x)$ is linear in x^i relative to some fixation point, say
$$v^\mu(x) = (v^0(x); v^i(x)) = (1; A_j^i x^j) \quad .$$
Separating space and time explicitly,
$$\gamma(x) = \gamma^{\text{time}}(t)\, \gamma^{\text{space}}(\vec{x}) \quad ,$$
and using the fact that $\gamma_\mu(x) = -x_\mu \gamma(x)$, it follows that, in the conjugate view, the filter transformation $\delta^T \gamma(x) = -\partial_\mu j^\mu(x)\delta\varepsilon$ is given by

$$\begin{cases} \dfrac{\delta^T \gamma^{\text{time}}(t)}{\delta\varepsilon} & = \quad t\, \gamma^{\text{time}}(t) & = \quad -\partial_t \gamma^{\text{time}}(t) \quad , \\[2ex] \dfrac{\delta^T \gamma^{\text{space}}(\vec{x})}{\delta\varepsilon} & = \quad \left(x_i A_j^i x^j - A_i^i \right) \gamma^{\text{space}}(\vec{x}) & = \quad A_j^i \partial_i \partial^j \gamma^{\text{space}}(\vec{x}) \quad . \end{cases}$$

The temporal part $\delta^T \gamma^{\text{time}}(t)$ represents a time shift over $\delta\varepsilon$ of the centre of the filter $\gamma^{\text{time}}(t)$, i.e. from $t = 0$ to $t = \delta\varepsilon$. The spatial part $\delta^T \gamma^{\text{space}}(\vec{x})$ is just the spatial $\mathcal{O}(\delta\varepsilon)$ part of the filter $\gamma_\Delta(x)$ of Definition 2, if we take the spatial submatrix to be $\Delta = I + \delta\varepsilon \left(A + A^T \right)$. In fact, one can "exponentiate" the infinitesimal transformations of $\mathcal{O}(\delta\varepsilon)$ in order to reveal their effect after

a finite time interval ε. To this end, consider the previous equations equipped with the following initial conditions (the basic, isotropic Gaussians):

$$
\begin{cases}
\gamma^{\mathrm{time}}(t; \varepsilon = 0) & = & \gamma_I^{\mathrm{time}}(t) \quad , \\[2ex]
\gamma^{\mathrm{space}}(\vec{x}; \varepsilon = 0) & = & \gamma_I^{\mathrm{space}}(\vec{x}) \quad .
\end{cases}
$$

Then, following the evolution of these basic filters for a time interval ε, we find

$$
\begin{cases}
\gamma^{\mathrm{time}}(t; \varepsilon) & \stackrel{\mathrm{def}}{=} & e^{-\varepsilon \partial_t}\, \gamma_I^{\mathrm{time}}(t) & = & \gamma_I^{\mathrm{time}}(t - \varepsilon) \quad , \\[2ex]
\gamma^{\mathrm{space}}(\vec{x}; \varepsilon) & \stackrel{\mathrm{def}}{=} & e^{\varepsilon A_j^i \partial_i \partial^j}\, \gamma_I^{\mathrm{space}}(\vec{x}) & = & \gamma_{I+\varepsilon(A+A^T)}^{\mathrm{space}}(\vec{x}) \quad .
\end{cases}
$$

The conclusion is that the filter continues, at least for some finite time, to adapt its profile to the linear flow. But note that, depending on the details of the flow field, the adaptation process may cease to make sense after a certain finite time due to physical limitations. It can be seen from the exponentiated results that the filters may expand or contract beyond physically sensible scale limits. For example, in the case of a "pure, negative divergence", for which $A_j^i = \alpha \delta_j^i$ with $\alpha < 0$ (so that $n\alpha = A_i^i = \partial_i v^i(x)$ is the only degree of freedom in the flow), the whole thing collapses after a time $\varepsilon = -1/(2\alpha)$ if one refrains from reinitialising the system periodically.

The following examples illustrate the gauge invariant systems for the approximated zeroth, first, and second order optic flow field, respectively.

Example 4 (Zeroth Order Optic Flow)
The zeroth order, gauge invariant optic flow field equation is given by $\delta_0 \Phi = 0$, with

$$
\frac{\delta_0 \Phi}{\delta \varepsilon} = v_0^\rho\, \partial_\rho \Phi \quad .
$$

This is of course just the classical Optic Flow Constraint Equation for a fixed point. Proceeding to higher orders one obtains systems that are fundamentally different from conventional schemes that differentiate the OFCE:

Example 5 (First Order Optic Flow)
The first order, gauge invariant optic flow field equations are given by $\delta_1 \Phi = \partial_\mu \delta_1 \Phi = 0$, with

$$
\begin{aligned}
\frac{\delta_1 \Phi}{\delta \varepsilon} &= v_0^\rho \, \partial_\rho \Phi + \partial_\sigma v_0^\rho \, \partial_\rho \partial^\sigma \Phi \quad , \\
\partial_\mu \frac{\delta_1 \Phi}{\delta \varepsilon} &= v_0^\rho \, \partial_{\rho\mu} \Phi + \partial_\sigma v_0^\rho \, \partial_{\rho\mu} \partial^\sigma \Phi + \partial_\mu v_0^\rho \, \partial_\rho \Phi \quad .
\end{aligned}
$$

Example 6 (Second Order Optic Flow)
The second order, gauge invariant optic flow field equations are given by $\delta_2 \Phi = \partial_\mu \delta_2 \Phi = \partial_{\mu\nu} \delta_2 \Phi = 0$, with

$$
\begin{aligned}
\frac{\delta_2 \Phi}{\delta \varepsilon} &= v_0^\rho \, \partial_\rho \Phi + \partial_\sigma v_0^\rho \, \partial_\rho \partial^\sigma \Phi + \frac{1}{2} \partial_{\sigma\tau} v_0^\rho \, \partial_\rho \partial^{\sigma\tau} \Phi + \frac{1}{2} \eta^{\sigma\tau} \partial_{\sigma\tau} v_0^\rho \, \partial_\rho \Phi \quad , \\
\partial_\mu \frac{\delta_2 \Phi}{\delta \varepsilon} &= v_0^\rho \, \partial_{\rho\mu} \Phi + \partial_\sigma v_0^\rho \, \partial_{\rho\mu} \partial^\sigma \Phi + \partial_\mu v_0^\rho \, \partial_\rho \Phi + \frac{1}{2} \partial_{\sigma\tau} v_0^\rho \, \partial_{\rho\mu} \partial^{\sigma\tau} \Phi + \frac{1}{2} \eta^{\sigma\tau} \partial_{\sigma\tau} v_0^\rho \, \partial_{\rho\mu} \Phi + \\
&\quad + \partial_{\mu\sigma} v_0^\rho \, \partial_\rho \partial^\sigma \Phi \quad , \\
\partial_{\mu\nu} \frac{\delta_2 \Phi}{\delta \varepsilon} &= v_0^\rho \, \partial_{\rho\mu\nu} \Phi + \partial_\sigma v_0^\rho \, \partial_{\rho\mu\nu} \partial^\sigma \Phi + \partial_\mu v_0^\rho \, \partial_{\rho\nu} \Phi + \partial_\nu v_0^\rho \, \partial_{\rho\mu} \Phi + \frac{1}{2} \partial_{\sigma\tau} v_0^\rho \, \partial_{\rho\mu\nu} \partial^{\sigma\tau} \Phi + \\
&\quad + \frac{1}{2} \eta^{\sigma\tau} \partial_{\sigma\tau} v_0^\rho \, \partial_{\rho\mu\nu} \Phi + \partial_{\mu\sigma} v_0^\rho \, \partial_{\rho\nu} \partial^\sigma \Phi + \partial_{\nu\sigma} v_0^\rho \, \partial_{\rho\mu} \partial^\sigma \Phi + \partial_{\mu\nu} v_0^\rho \, \partial_\rho \Phi \quad .
\end{aligned}
$$

The Examples 4, 5 and 6 illustrate the general principle of refinement: the $(M+1)$-st order system $\{\partial_{\mu_1 \dots \mu_k} \delta_{M+1} \Phi = 0\}_{k=0}^{M+1}$ has the same form as the M-th order system $\{\partial_{\mu_1 \dots \mu_k} \delta_M \Phi = 0\}_{k=0}^{M}$ except for additional terms of order $M+1$ in the flow field's approximation. The transition $M \to M+1$ will generally affect *all* coefficients in the formal expansion of the optic flow field. The hypothetical limit $M \to \infty$ will generate the "ideal" system of Definition 5 with its original, gauge invariant optic flow field $v^\mu(x)$ replaced by its equivalent Taylor expansion $v_\infty^\mu(x)$.

In all examples given thus far we have been illustrating the key idea of this paper: gauge invariant optic flow. It is, *by definition*, ambiguous. Yet it should be appreciated that it provides the only data intrinsic evidence for any unambiguous definition of optic flow. As such, it lies at the basis of any gauge constrained definition of optic flow. The gauge needed to single out an unambiguous choice is, as argued before, a requirement enforced by the physics

of the image formation process. As such, it is independent of the data, a basic observation that has not always been made explicit.

Although this paper deals with the data induced, gauge invariant optic flow, it is instructive to point out how things can be put to work in practice. Finding a unique optic flow field requires a gauge condition complementing the gauge invariant optic flow equation of Definition 5 or, in practice, its approximating equation of Definition 7, which requires modality specific details that are beyond the scope of this paper. Gauges may be inferred from physical considerations such as rigidity or elasticity of motion, incompressibility of fluids, flow continuity (or conservation of mass), etc., very much depending on the imaging modality.

One of the simplest cases of a gauge constrained optic flow field, and of rather general interest in its own right, is perhaps *normal flow*. The gauge condition is "canonical" rather than physical; it is intended to annihilate the "pure" gauge degree of freedom, i.e. the tangential flow (thus the result may not be the flow induced by some physical motion). As a final example we show how to obtain approximations of the normal flow in $2 + 1$ dimensions. It serves to illustrate the general recipe: enforce a (physically motivated) constraint so as to break the intrinsic gauge invariance, and solve for the resulting (approximated) optic flow field. The advantage of this approach is that one maintains linearity throughout.

Example 7 (Normal Flow)
Define the *dual*[8] of a $(2 + 1)$-vector $\mathbf{v} = (1; u, v)$ in a Cartesian coordinate system by $*\mathbf{v} = (0; -v, u)$. Then one can solve for the *normal flow* components up to 0-th, 1-st, or 2-nd order approximation by complementing the corresponding linear, inhomogeneous gauge systems of Examples 4, 5 and 6 by similar, homogeneous systems with $\mathbf{v}$ replaced by $*\mathbf{v}$. It is a straightforward but tedious exercise to replace the condensed summation convention by an explicit expression for any given case. For definiteness, here is the 1-st order system (with self-explanatory notation for the 2 zeroth and 6 first order optic flow components):

---

[8]Remember not to confuse space and time!

- gauge invariant system (data induced, always the same):

$$
\left\{
\begin{array}{ll}
\Phi_t + u\Phi_x + v\Phi_y + u_t\Phi_{xt} + v_t\Phi_{yt} + u_x\Phi_{xx} + u_y\Phi_{xy} + v_x\Phi_{xy} + v_y\Phi_{yy} & = 0 \\
\Phi_{tt} + u_t\Phi_x + u\Phi_{xt} + v_t\Phi_y + v\Phi_{yt} + u_t\Phi_{xtt} + v_t\Phi_{ytt} + u_x\Phi_{xxt} + u_y\Phi_{xyt} + v_x\Phi_{xyt} + v_y\Phi_{yyt} & = 0 \\
\Phi_{xt} + u_x\Phi_x + u\Phi_{xx} + v_x\Phi_y + v\Phi_{xy} + u_t\Phi_{xxt} + v_t\Phi_{xyt} + u_x\Phi_{xxx} + u_y\Phi_{xxy} + v_x\Phi_{xxy} + v_y\Phi_{xyy} & = 0 \\
\Phi_{yt} + u_y\Phi_x + u\Phi_{xy} + v_y\Phi_y + v\Phi_{yy} + u_t\Phi_{xyt} + v_t\Phi_{yyt} + u_x\Phi_{xxy} + u_y\Phi_{xyy} + v_x\Phi_{xyy} + v_y\Phi_{yyy} & = 0
\end{array}
\right.
$$

- gauge condition for normal flow (model induced, depends on the physics of the situation):

$$
\left\{
\begin{array}{ll}
-v\Phi_x + u\Phi_y - v_t\Phi_{xt} + u_t\Phi_{yt} - v_x\Phi_{xx} - v_y\Phi_{xy} + u_x\Phi_{xy} + u_y\Phi_{yy} & = 0 \\
-v_t\Phi_x - v\Phi_{xt} + u_t\Phi_y + u\Phi_{yt} - v_t\Phi_{xtt} + u_t\Phi_{ytt} - v_x\Phi_{xxt} - v_y\Phi_{xyt} + u_x\Phi_{xyt} + u_y\Phi_{yyt} & = 0 \\
-v_x\Phi_x - v\Phi_{xx} + u_x\Phi_y + u\Phi_{xy} - v_t\Phi_{xxt} + u_t\Phi_{xyt} - v_x\Phi_{xxx} - v_y\Phi_{xxy} + u_x\Phi_{xxy} + u_y\Phi_{xyy} & = 0 \\
-v_y\Phi_x - v\Phi_{xy} + u_y\Phi_y + u\Phi_{yy} - v_t\Phi_{xyt} + u_t\Phi_{yyt} - v_x\Phi_{xxy} - v_y\Phi_{xyy} + u_x\Phi_{xyy} + u_y\Phi_{yyy} & = 0
\end{array}
\right.
$$

This system of four inhomogeneous and four homogeneous, linear equations can be inverted in the generic case for the 8 local flow degrees of freedom $(u, v, u_t, v_t, u_x, v_x, u_y, v_y)$.

4 Conclusion and Discussion

In this paper we have proposed a gauge invariant definition of optic flow. In order to make the number of local degrees of freedom manageable, we have approximated the ideal, hypothetical optic flow field by a formal expansion, truncated at some finite order. The coefficients in this expansion can be determined modulo gauge transformations in terms of a simple, linear system of equations. This system captures the full differential structure of the (approximate) optic flow field up to some order, as far as this is determined intrinsically by the image data.

Extrinsic, physical considerations beyond the information provided by the image data, such as a priori knowledge of the underlying scene and of the image formation process, need to be taken into account in order to complement the ungauged system. This "gauge fixing", if done properly, may lead to a unique, physically sensible solution. Since gauge fixing relies entirely on *specific* information *beyond* the evidence provided by the data, we have not addressed it in the general context of this paper. Rather, we have emphasised the general aspects of optic flow as far as it relates to image structure *as such*, and shown

how to decouple it from considerations that vary from one image modality to another, and even within each modality, from one specific scene to another. One general a priori gauge has been explicitly discussed, viz. the usual temporal gauge $v^0(x) = 1$; it is used to synchronise the optic flow's affine parameter with "universal time".

The defining linear optic flow system resembles the classical Horn and Schunck's Optic Flow Constraint Equation and existing schemes derived from this by implicit differentiation. There is, however, an important difference. Since our defining system is based on the vanishing of a Lie derivative of the image data, *when viewed by some fixed, physical aperture*, rather than on an operationally ill-defined and ill-posed derivative in the conventional sense, it is *not* required that the grey-values attributed to the discretisation grid (or any other fixed-scale average grey-values) remain constant along the flow. Only if one scales up the apertures consistently with the optic flow field's divergence (which causes local volume elements to expand), will the resulting grey-values remain constant.

An important observation has been made concerning the refinement of approximation $v_M^\mu \rightarrow v_{M+1}^\mu$; it does not merely introduce an extra order in the approximation, but updates *all* coefficients in the polynomial expansion of the flow field. This update may provide an important cue concerning "discontinuities" of the optic flow field. If the flow field happens to vary slowly relative to the inner scale of the Gaussian aperture used in the computation of image derivatives, then one may expect to obtain an accurate estimate v_M^μ of the flow field even for very low order M (an extreme case is that of a pure translation, for which it suffices to take $M = 0$). In that case the transition to higher orders will hardly affect the coefficients in the expansion. If, on the other hand, the field varies significantly relative to inner scale (optic flow "discontinuities"), then the transition from order M to $M + 1$ will have significant influence on the coefficients of the expansion. Hence one could use this transient behaviour to test for consistency of truncation, as well as for localising effective discontinuities of the field.

A point of speculation may be the biological implications of this work. In this respect, the conjugate view of optic flow presented in this paper, in which the flow is "carried over" into the "sensorium" were the filters (read: *receptive fields*) live, may be more than a mere conceptual alternative. One could

imagine tuned receptive fields actively deforming their profiles (periodically) according to some specific optic flow degree of freedom, such as a pure (0-th order) translation or divergence (1-st order, one of the examples in this paper). In this case the optic flow degrees of freedom correspond to tuning parameters, and one could hypothesise ensembles of receptive fields for different values of these parameters. Whether this could actually lead to a plausible model for velocity sensitive receptive fields remains speculative at this point. But the type of receptive fields needed to construct such tuned "motion detectors" based on simple, static filters (Gaussian derivatives), as well as the kind of mutual connections needed to realise the prototypical "conjugate flow" (correlators with spans and/or delays), are physiologically quite plausible. The advantage of having motion detectors for various tuning parameters is that it may well account for *transparant motion*. But also the "direct" view, in which the filters are considered to be passive, may have a rather straightforward physiological realisation. A challenging question in this case is the kind of gauge condition imposed by the visual system to account for an apparently unambiguous motion percept. However, it is less obvious how to explain transparency in this case.

A The Filters $\Gamma^{\rho_1...\rho_l}_{\mu_1...\mu_k}(x)$ and the Gaussian Family

Using the following lemma we can get rid of the derivative ∂_ρ in the integrand of Result 1:

Lemma 1
Using parentheses to denote index symmetrisation, we have

$$\partial_\mu \Gamma^{\rho_1...\rho_l}_{\mu_1...\mu_k}(x) = -\Gamma^{\rho_1...\rho_l}_{\mu_1...\mu_k\mu}(x) + \delta^{(\rho_l}_\mu \Gamma^{\rho_1...\rho_{l-1})}_{\mu_1...\mu_k}(x) .$$

It is understood that $\Gamma^{\rho_1...\rho_{l-1}}_{\mu_1...\mu_k}(x) \equiv 0$ *if* $l = 0$.

The proof of this lemma is straightforward and will be omitted. Using this lemma we can rewrite Result 1:

Result 2 $\left(\partial_{\mu_1...\mu_k}\delta_M\Phi\right)$
See Result 1.

$$\partial_{\mu_1...\mu_k}\delta_M\Phi = \sum_{l=0}^{M} \partial_{\rho_1...\rho_l}v_0^\mu \int dx\,\phi(x) \left[\Gamma_{\mu_1...\mu_k\mu}^{\rho_1...\rho_l}(x) - \Gamma_{\mu_1...\mu_k}^{\rho_1...\rho_{l-1}}(x)\,\delta_\mu^{\rho_l} \right] \delta\varepsilon \quad .$$

(Note that we do not need to make index symmetrisation for $\rho_1,\ldots,\rho_l$ explicit here; it is automatically achieved by virtue of symmetry of the differential operator $\partial_{\rho_1...\rho_l}$.)

In order to express the overcomplete set of filters $\Gamma_{\mu_1...\mu_k}^{\rho_1...\rho_l}(x)$ in terms of Gaussian derivative filters $\gamma_{\mu_1...\mu_m}(x)$, consider the following diagram.

$$
\begin{array}{ccc}
\Gamma_{\mu_1...\mu_k}^{\rho_1...\rho_l}(x) & \xrightarrow{\ \mathbf{F}\ } & \hat{\Gamma}_{\mu_1...\mu_k}^{\rho_1...\rho_l}(\omega) \\
\star\downarrow & & \downarrow\star\star \\
\Gamma_{\mu_1...\mu_k}^{\rho_1...\rho_l}(x) & \xleftarrow{\ \mathbf{F^{inv}}\ } & \hat{\Gamma}_{\mu_1...\mu_k}^{\rho_1...\rho_l}(\omega)
\end{array}
$$

Instead of simplifying directly in the spatial domain (the arrow marked by a $\star$), we take the Fourier route ($\mathbf{F} \to \star\star \to \mathbf{F^{inv}}$), and simplify in Fourier space.

Definition 8 (Fourier Transform)
The Fourier transform of a function $f : \mathbb{R} \to \mathbb{R} : x \mapsto f(x)$ is defined as

$$f(x) = \int d\omega\, e^{i\omega x}\hat{f}(\omega) \quad .$$

Hence

$$\hat{f}(\omega) = \frac{1}{2\pi}\int dx\, e^{-i\omega x}f(x) \quad .$$

With this definition, we can make the following formal identifications of operators (the l.h.s. in the spatial domain, the r.h.s. in the Fourier domain):

$$x^\rho \equiv i\frac{\partial}{\partial\omega_\rho} \quad , \quad \frac{\partial}{\partial x^\rho} \equiv i\omega_\rho \quad . \tag{13}$$

We need one more definition.

Definition 9 (Hermite Polynomials)
The Hermite polynomial of order k, $H_k(x)$, is defined by

$$\frac{d^k}{dx^k}\, e^{-\frac{1}{2}x^2} = (-)^k H_k(x)\, e^{-\frac{1}{2}x^2} \quad .$$

This is appropriate for the 1-dimensional case; but in general we have more dimensions, say d. So let us define, for the sake of convenience, the d-dimensional analogue of the Hermite polynomials, as follows.

Definition 10 (Hermite Polynomials in d Dimensions)
The d-dimensional Hermite polynomial of order k, $\mathcal{H}_{i_1 \ldots i_k}(x)$, is defined by

$$\frac{\partial^k}{\partial x^{i_1} \ldots \partial x^{i_k}}\, e^{-\frac{1}{2}x^2} = (-)^k \mathcal{H}_{i_1 \ldots i_k}(x)\, e^{-\frac{1}{2}x^2} \quad .$$

These d-dimensional Hermite polynomials are related to the standard ones in the following way.

Lemma 2 (Relation to Standard Definition)
The d-dimensional Hermite polynomials as defined according to Definition 10 *are related to the standard definition,* Definition 9, *as follows.*

$$\mathcal{H}_{i_1 \ldots i_k}(x) = \prod_{j=1}^{d} H_{\alpha_j^{i_1 \ldots i_k}}(x^j) \quad ,$$

in which $\alpha_j^{i_1 \ldots i_k}$ denotes the number of indices in $i_1, \ldots, i_k$ equal to j.

Clearly we have $\sum_{j=1}^{d} \alpha_j^{i_1 \ldots i_k} = k$, since this simply sums up all indices. The separability property of Lemma 2 follows straightforwardly from Definition 9, when applied to a multidimensional Gaussian.

Having established all basic ingredients and notational matters, we can now relate the overcomplete family of filters $\Gamma_{\mu_1 \ldots \mu_k}^{\rho_1 \ldots \rho_l}(x)$ to the Gaussian family. This is easy, since all we need to do is to use Leibnitz's product rule for differentiation in

$$\hat{\Gamma}_{\mu_1 \ldots \mu_k}^{\rho_1 \ldots \rho_l}(\omega) = \frac{(-)^k}{l!}\, i\frac{\partial}{\partial \omega_{\rho_1}} \ldots i\frac{\partial}{\partial \omega_{\rho_l}} \left(i\omega_{\mu_1} \ldots i\omega_{\mu_k} \hat{\gamma}(\omega) \right) \quad , \tag{14}$$

(see formula (13) and the definition of the filters in Result 1). Then, each time we have to take a derivative of $\hat{\gamma}(\omega)$, we use the explicit property of the Gaussian stated in Definition 10. In this way we arrive at

Result 3 (The Filters $\Gamma^{\rho_1 \dots \rho_l}_{\mu_1 \dots \mu_k}(x)$ and the Gaussian Family)
Let $\mathcal{S}$ denote the index symmetrisation operator (applying both to upper as well as lower indices), then we have

$$\hat{\Gamma}^{\rho_1 \dots \rho_l}_{\mu_1 \dots \mu_k}(\omega) =$$
$$\frac{(-)^k}{l!} \, \mathcal{S} \left\{ \sum_{m=0}^{\min(k,l)} \binom{l}{m} \frac{k!}{(k-m)!} (-)^m \delta^{\rho_1}_{\mu_1} \dots \delta^{\rho_m}_{\mu_m} \, i\omega_{\mu_{m+1}} \dots i\omega_{\mu_k} (-i)^{l-m} \mathcal{H}^{\rho_{m+1} \dots \rho_l}(\omega) \, \hat{\gamma}(\omega) \right\} \, .$$

Fourier inversion yields

$$\Gamma^{\rho_1 \dots \rho_l}_{\mu_1 \dots \mu_k}(x) =$$
$$\frac{(-)^k}{l!} \, \mathcal{S} \left\{ \sum_{m=0}^{\min(k,l)} \binom{l}{m} \frac{k!}{(k-m)!} (-)^m \delta^{\rho_1}_{\mu_1} \dots \delta^{\rho_m}_{\mu_m} \, \partial_{\mu_{m+1}} \dots \partial_{\mu_k} (-i)^{l-m} \mathcal{H}^{\rho_{m+1} \dots \rho_l}(-i\nabla) \, \gamma(x) \right\} \, .$$

Note that this expression is real in the spatial domain, since $(-i)^p \mathcal{H}^{\rho_1 \dots \rho_p}(-i\nabla)$ is a real differential operator for any $p \in \mathbb{Z}_0^+$. To see this, look at the explicit form of a Hermite polynomial:

$$H_k(x) = \sum_{m=0}^{[k/2]} (-)^m \binom{k}{2m} (2m-1)!! \, x^{k-2m} \quad , \tag{15}$$

in which $[x]$ denotes the *entier* of $x \in \mathbb{R}$, i.e. the largest integer less than or equal to x, and in which the double factorial $(2m-1)!!$ indicates the product $1 \times 3 \times \dots \times (2m-1)$. Consequently,

$$(-i)^k H_k(-i\frac{d}{dx}) = (-)^k \sum_{m=0}^{[k/2]} \binom{k}{2m} (2m-1)!! \, \frac{d^{k-2m}}{dx^{k-2m}} \quad , \tag{16}$$

very real indeed. The general n-dimensional case follows from this observation. Note also that the r.h.s. of Result 3 is a linear combination of Gaussian derivatives of the type $\gamma_{\mu_1 \dots \mu_p}(x)$, with $p = 0, \dots, k+l$. Thus we have indeed proven overcompleteness of the (apparently $(k+l)$-th order) filters $\Gamma^{\rho_1 \dots \rho_l}_{\mu_1 \dots \mu_k}(x)$ by explicitly rewriting them in terms of Gaussian derivatives.

References

[1] B. K. P. Horn and B. G. Schunck, "Determining optical flow," *Artificial Intelligence*, vol. 17, pp. 185–203, 1981.

[2] J. Arnspang, "Notes on local determination of smooth optic flow and the translational property of first order optic flow," Tech. Rep. 88-1, Institute of Datalogy, University of Copenhagen, 1988.

[3] J. Arnspang, "Optic acceleration. local determination of absolute depth and velocity, time to contact and geometry of an accelerating surface," Tech. Rep. 88-2, Institute of Datalogy, University of Copenhagen, 1988.

[4] P. Werkhoven, *Visual Perception of Successive Order*. PhD thesis, University of Utrecht, Dept. of Med. Phys., Princetonplein 5, Utrecht, the Netherlands, May 1990.

[5] J. Arnspang, "Motion constraint equations in vision calculus," April 1991. Dissertation for the Danish Dr. scient. degree, University of Utrecht, Utrecht, The Netherlands.

[6] M. Otte and H.-H. Nagel, "Optical flow estimation: Advances and comparisons," in *Proceedings of the European Conference on Computer Vision*, (J.-O. Eklundh, ed.), (Stockhlom, Sweden), pp. 51–60, May 2–5 1994.

[7] A. A. Amini, "A scalar function formulation for optical flow," in *Proceedings of the European Conference on Computer Vision*, (J.-O. Eklundh, ed.), (Stockhlom, Sweden), pp. 125–131, May 2–5 1994.

[8] M. Spivak, *Differential Geometry*. Vol. 1, Berkeley, California, USA: Publish or Perish, Inc., 1975.

[9] C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation*. San Francisco: Freeman, 1973.

[10] L. Schwartz, *Théorie des Distributions*. Vol. I, II of *Actualités scientifiques et industrielles; 1091,1122*, Paris: Publications de l'Institut de Mathématique de l'Université de Strasbourg, 1950–1951.

[11] L. M. J. Florack, B. M. ter Haar Romeny, J. J. Koenderink, and M. A. Viergever, "Images: Regular tempered distributions," in *Proc. of the NATO Advanced*

Research Workshop Shape in Picture - Mathematical Description of Shape in Greylevel Images, (Y.-L. O, A. Toet, H. J. A. M. Heijmans, D. H. Foster, and P. Meer, eds.), pp. 651–660, Springer Verlag, Berlin, 1994.

[12] J. J. Koenderink and A. J. van Doorn, "Representation of local geometry in the visual system," *Biol. Cybern.*, vol. 55, pp. 367–375, 1987.

[13] T. Poggio and W. Reichardt, "Considerations on models of movement detection," *Kybernetik*, vol. 13, pp. 223–227, 1973.

[14] L. M. J. Florack, B. M. ter Haar Romeny, J. J. Koenderink, and M. A. Viergever, "Families of tuned scale-space kernels," in *Proceedings of the European Conference on Computer Vision*, (G. Sandini, ed.), (Santa Margherita Ligure, Italy), pp. 19–23, May 19–22 1992.

[15] M. Spivak, *Calculus on Manifolds*. New York, New York, USA: W. A. Benjamin, Inc., 1965.

[16] J. J. Koenderink and A. J. van Doorn, "Invariant properties of the motion parallax field due to the movement of rigid bodies relative to an observer," *Optica Acta*, vol. 22, no. 9, pp. 773–791, 1975.

[17] J. J. Koenderink and A. J. van Doorn, "Second-order optic flow," *Journal of the Optical Society of America-A*, vol. 9, pp. 530–538, April 1992.

[18] J. J. Koenderink and A. J. van Doorn, "Receptive field families," *Biol. Cybern.*, vol. 63, pp. 291–298, 1990.

Figure 1: In the conventional interpretation of the OFCE it is asserted that two points p and q (or *fixed-scale* samples taken at these points), lying on one flow line, have the same grey-value attribute. In the new interpretation adopted in this paper, grey-values are attributed to *volumes* rather than *points*. Consequently, the OFCE is assumed to apply to local volume elements, which are susceptible to the divergence of the optic flow field. In this sketch, the indicated patches are intended to represent corresponding volume elements, the size of which is seen to increase due to the divergence component of the flow field; the new OFCE entails that their *volumetrically integrated* grey-values be the same.