# Duality Principles in Image Processing and Analysis*

**Luc Florack**
Department of Computer Science, Utrecht University
PO Box 80089, NL-3508 TB  Utrecht, The Netherlands
Email: Luc.Florack@cs.uu.nl

**Abstract**

Duality is a well-established concept in quantum physics. It formalises the fact that what one observes is not nature in itself, but—in Heisenberg's words—"nature exposed to our method of questioning". In the context of image analysis "question" pertains to some *task* while "nature" (empirical facts and natural laws) could be taken as the totality of image data supplemented with relevant external factors (knowledge or hypotheses).

However, the analogy with quantum physics falls short in at least one fundamental aspect. Whereas the physicist studies nature for nature's sake, endeavouring to reveal natural laws, the image scientist pursues a certain task. This implies a shift of paradigm from the retrospect to the prospect. It will be argued that this leads to a subtle but important difference in the role duality plays in image processing and analysis as compared to the technically similar "bracket" formalism in quantum physics.

Duality in the context of image processing and analysis will be explained in detail and its use will be illustrated in a number of cases. It lies at the core of any low-level definition of a *local sample*, generalising the notion of a digital picture element. It also provides operational definitions of *partial derivatives* and *Lie derivatives*. A duality principle known as "carry-along" allows us to connect isolated samples into ensembles endowed with the kind of topology one usually attributes to space, time, and spacetime. In this way one is led to corresponding definitions for *images*, *signals*, and *video sequences*. The flexibility of duality as a generic framework for image processing and analysis is illustrated by further examples in the context of off-line and real-time processing, notably *causal filtering* and *motion analysis*. .

## 1 Introduction

The concepts of *duality* and *metamerism* are close-knit. Duality pertains to the fact that the outcome of any computation—a "representation"—relies on interaction between a fiducial source (input) and a correlate embodying some computational paradigm (filter, algorithm, observer, *etc.*). Metamerism refers to an equivalence relation induced by the fact that all representations are potentially many-to-one. Thus any source configuration is a member of an equivalence class of feasible configurations (metamers) that would induce identical representations if subjected to the same computation.

Metamerism has been explained in some detail by Koenderink, *cf.* his example of the "sextuplet image" and "metameric black pictures" [15, 16]. A more familiar instance in the context of colour vision is the well-known fact that the "tristimulus curves" give rise to indistinguishable colours despite different underlying spectral sources [11].

Since their physical representations are by definition identical, it is impossible to segregate metamerical sources, although it is possible to select preferred ones on the basis of *models* (external knowledge

*This report summarizes an invited lecture for IVCNZ'98, New Zealand, November 16–18 1998.

or hypotheses). Even a total lack of knowledge does not prevent one from applying at least some principle of parsimony to select a distinguished representative, *e.g.* via regularization, the MDL—"minimum description length"—principle, *et cetera*. Indeed it would be quite impossible in practice *not* to make any assertions; the application of a model is an inevitability and a commitment at the same time. Example: One often takes for granted that the source underlying an image acquisition can be modelled as a function, although the class of objects that can be probed so as to produce the same image values is actually more general.

Duality is a well-established concept in quantum physics, where it has a rigorous mathematical foundation. However, the emphasis there tends to be on a Hilbert space, or "bracket" formalism, in which the dual objects ("bras" and "kets", or forms and vectors) are one-to-one related (recall that the standard Hilbert space of $L^2$-functions is isomorphic to its own dual, *v.i.*). Metamerism is usually deemphasised or even considered an obstacle because aims are retrospective, *viz.* to reveal "nature exposed to our method of questioning" (Heisenberg). In image analysis on the other hand data are subordinate to the accomplishment of a task, and for this reason the space of filters is of a different kind than that of the source data, *viz.* such as to optimally subserve that task. Metamerism consequently becomes a fundamental concept as it enables one to encapsulate task-irrelevant details into equivalence classes. In low-level image representation one may think of hiding noise and digitization details, although there is no reason why it should not apply to high-level descriptions. For instance, a segmentation task entails the explicit formation of metameric classes (reflecting tissue types or textures, say) despite potentially significant intraregional variability at the level of the raw data. Another example is optical character recognition, in which one would like to classify characters irrespective of font size and style.

The use of duality in the context of image processing and analysis will be illustrated below.

## 2 Theory

If one wishes to abstract from machine technicalities in the definition of an "image[1]" one will need at least three ingredients: (i) a source $f$ (digital data, physical density field, retinal irradiance distribution, *etc.*), (ii) a device or template $\phi$ (filter, structuring element, probe, sensor, detector, receptive field, *etc.*, depending on details and context) mapping raw, unstructured data to formatted data ("samples" or "observations"), and (iii) an operational model $\Theta$ that accounts for spatial coherence (a grab-bag of samples does not make an image unless it is endowed with a suitable topology). In digital images the latter is reflected by external "header" information stored in conjunction with the numerical pixel data.

In particular, a sample is defined by virtue of interaction between $f$ and $\phi$. Duality pertains to the details of the interface. One branch of mathematics, known as distribution theory, provides a well-established instance known as *topological duality*. Point of departure is *linearity* and *continuity*. In mathematical morphology one pursues an alternative option, *morphological duality*. This is a *nonlinear* probing mechanism, although remarkable analogies with the linear case exist. A third instance of duality sheds light on the cause of these analogies. It encompasses both topological as well as morphological formalisms as limiting cases. Albeit strictly nonlinear, it is isomorphic to the linear case, and hence will be referred to as *pseudo-topological duality*. All three instances will be briefly outlined.

### 2.1 The Principle of Duality

Let us denote the class of all possible sources—raw images, for the sake of definiteness—by $\Sigma$, and call it *state space* for ease of reference. The class of filters will be called *device space* and indicated

[1]"Image" should be distinguished from "raw image"—in the form of an image file, say—which is merely a machine representation of the source field of interest.

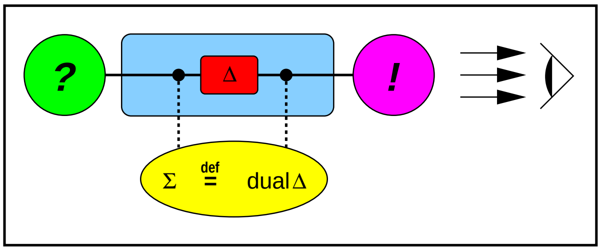


Figure 1: Duality: $\Sigma \stackrel{\rm def}{=} {\rm dual}\,\Delta$ means that an element $F \in \Sigma$ (*i.e.* a "raw image") is defined by virtue of the way it triggers numerical responses $F[\phi] \in \mathbb{R}$ when exposed to all filters in a fiducial filter class $\phi \in \Delta$. This implies that one can make arbitrary hypotheses about the "naked source field" (the input symbolised by "**?**") as long as these are consistent with data evidence (the output representation marked by "**!**") given one's knowledge about the inner workings of the measurement procedure (*i.e.* the duality paradigm). Source field configurations or raw images that induce identical representations for a given filter class are called metamerical (relative to filter class and duality paradigm).

by $\Delta$. Assessment of a raw image proceeds indirectly, *viz.* by monitoring its behaviour when probed by *all available* templates: Fig. 1. *I.e.* one conceives of a grey-value as the output $F[\phi] \in \mathbb{R}$ of a raw image $F \in \Sigma$ when exposed to a specified filter $\phi \in \Delta$. It is common practice to model a raw image as an ordinary function mapping spatial points $\mathbf{x} \in \mathbb{R}^n$ into grey-values $f(\mathbf{x}) \in \mathbb{R}$. However, point mappings are deceptive for two reasons: (i) one cannot map points as these are physically void entities, and (ii) grey-values have no objective existence independent of some measurement aperture, *i.e.* without accounting for duality. This can be formalised as follows.

**Paradigm 1 (Duality)** *State space is the dual of device space:* $\Sigma \stackrel{\rm def}{=} {\rm dual}\,\Delta$.

We will adhere to the traditional "naked function model" to represent a *raw image* in order to de-emphasise the duality underlying the raw image data, which are of course in turn determined by some duality principle reflecting the measurement paradigm (think of sampling with a device specific "point spread function"). Using functions rather than functionals in this case expresses our wish to abstract from image formation details (point spread function, quantum effects, noise) and machine technicalities (rendering details, numerical limitations), but will of course pose physical limitations to filter design (*e.g.* it is of no use to apply filters that are more narrow than the point spread function). In particular we must decline from references to isolated source values; the "identity operator"—*i.e.* a sampling filter confined to a single point, returning the source field's "punctal" value—is a physically void concept! Note that, in principle, $F[\phi]$ may be *any* functional of $\phi$.

## 2.2 Metamerism

In the context of duality it is natural to consider *equivalence classes*: Input images that trigger identical responses are indistinguishable. Example: Suppose we have only one filter at our disposal, which reads its input $f$ and returns its pixel average $\overline{f}$. Then $g \sim f$ iff $\overline{g} = \overline{f}$. Clearly lots of images map to the same mean. In order to arrive at an image processing framework for *generic* purposes a conceptual image $F$ must somehow be one-to-one related to its pixel data $f$, except for non-measurable details. That is, if $g(\mathbf{x}) = f(\mathbf{x})$ almost everywhere, then we would like $G[\phi] = F[\phi]$ for all plausible filters $\phi$, *vice versa*. In a generic context simple duality formalisms and large filter classes are preferred.

If, on the other hand, one has a *specific* goal in mind, it remains a bit of an art to arrange things in such a way that $G[\phi] = F[\phi]$ whenever two raw images $f$ and $g$ are supposed to produce identical outcome. A nontrivial example is "noise suppression" (or, if one wishes, segmentation): If $f$ and $g$ differ only in the manifestation of noise, then a filter $\phi$ in conjunction with a suitable functional (*i.e.* duality principle) realizing the above identity would solve the problem. In a specific context duality formalisms tend to become highly nontrivial and the smallest possible filter class —one would rather refer to the probing mechanism as a whole as an "algorithm" in this case—is preferred so as to minimise ambiguity or user effort.

## 2.3 Topological Duality

Without loss of generality one may adopt the filter class proposed by Schwartz [19] in the (bi)linear case. According to Paradigm 1 this implies that images are conceived of as so-called "tempered distributions" (objects that are to be filtered linearly in order to produce numeric samples).

**Definition 1 (Schwartz Theory)** *Let $\mathcal{S}(\mathbb{R}^n)$ be the class of smooth functions of rapid decay, then $\Delta \stackrel{\text{def}}{=} \mathcal{S}(\mathbb{R}^n)$, whence $\Sigma \stackrel{\text{def}}{=} \mathcal{S}'(\mathbb{R}^n)$. The latter is also known as the class of tempered distributions.*

The majority of tempered distributions is "regular", *i.e.* they can be written in integral form:

$$F[\phi] = \int d\mathbf{z}\, f(\mathbf{z})\, \phi(\mathbf{z})\,.$$

(A very weak formal requirement is that $f$ must be a function of polynomial growth.) Of course this is nothing but straightforward linear filtering, and it is easy to implement an approximate discrete version. Nonregular tempered distributions always involve the *Dirac distribution*: $\delta[\phi] \equiv \phi(\mathbf{0})$. In practice one uses the integral formula even in these cases, associating the Dirac distribution with the "function-under-the-integral" $\delta(\mathbf{x})$. Such distributions are not at all pathological, and share virtually all nice properties characteristic for distributions in general, such as smoothness (infinite differentiability). "Point stimuli" also lie at the core of "reverse engineering" disciplines, in which one aims to establish filter profiles of a black box system by applying an (approximate) point stimulus. In image analysis the stimulus is given (a raw image $f$, which can be approximated by a linear combination of Dirac distributions if one neglects the p.s.f.), and the filters are designed according to need.

Rapid decay reflects filter confinement. This implies that we can assign a *base point* $\mathbf{x}$ to each filter $\phi$ corresponding to its "centre of mass". This many-to-one assignment rule $\boldsymbol{\pi}[\phi] = \mathbf{x}$ is called a "projection map". Its one-to-many inverse, $\boldsymbol{\pi}^{\text{inv}}[\mathbf{x}] = \Delta_{\mathbf{x}} \subset \Delta$, producing all filters at a given base point, is technically known as a "fibre", and the totality of all fibres over all base points, $\Delta = \cup_{\mathbf{x}} \Delta_{\mathbf{x}}$ as a "fibre bundle". A localised filter may have *any* size; a suitable measure for this is the normalised second order central momentum tensor ("affine scale").

Filter smoothness is hardly a demand. This follows from the fact that $\mathcal{S}'(\mathbb{R}^n)$ is larger than any of the function spaces typically employed in non-dualistic models. In other words, the filter class $\mathcal{S}(\mathbb{R}^n)$ guarantees a more-than-sufficient "segregation of quality". In fact, if $F[\phi] = G[\phi]$ for all $\phi \in \mathcal{S}(\mathbb{R}^n)$, then $f$ and $g$ differ by at most a non-measurable function, exactly as desired.

One *image processing consistency* demand should be mentioned. If a filtered image[2] does not have a preferred status, in other words, is just another source field that could be filtered in turn, then it can be shown that $\Delta$ must satisfy a *closure property*: Two successive filterings must be equivalent to a single one using a filter that is also available in $\Delta$, *cf.* Fig. 2. This is indeed the case for $\mathcal{S}(\mathbb{R}^n)$, which is a *convolution algebra*. The consistency demand must be met if we want to account for all possible input images regardless of preprocessing history. *Scale-space theory* [6, 10, 12, 13, 18, 20, 21] boils down

[2]Strictly speaking we have not yet defined "image", but think of it for the moment as the *correlation* $f \star \phi(\mathbf{x})$, *i.e.* the collection of samples obtained according to Definition 1 using shifted copies $\phi(\mathbf{z}-\mathbf{x})$ of $\phi(\mathbf{z})$ at every base point $\mathbf{x}$.

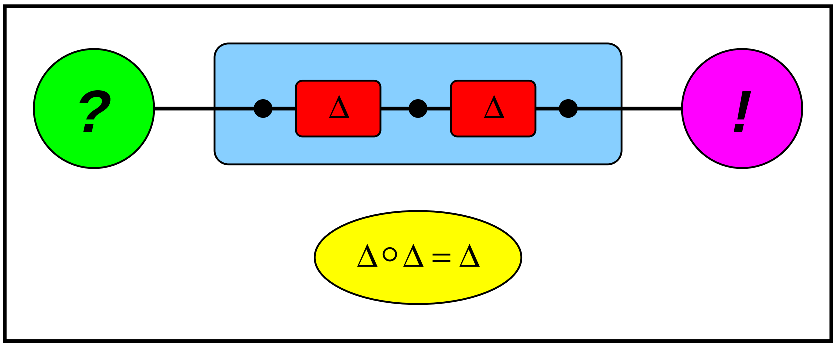


Figure 2: Image processing consistency requires device spaces to satisfy a closure property.

to Schwartz theory equipped with a *point concept*, a positive "zeroth order" filter consistent with the image processing demand, *i.e.* one generating an *autoconvolution algebra*. Within $\mathcal{S}(\mathbb{R}^n)$ this leaves no choice but the normalised Gaussian (of arbitrary scale and base point). For later use the Gaussian and its derivatives will be denoted by $\mathcal{G}(\mathbb{R}^n)$, the "Gaussian family" [17].

**Definition 2 (Scale-Space Theory)** *Recall Definition 1. The scale-space representation of a raw image is obtained by subjecting it to the Gaussian family* $\Delta \stackrel{\text{def}}{=} \mathcal{G}(\mathbb{R}^n)$*, i.e., it is an element of* $\Sigma \stackrel{\text{def}}{=} \mathcal{G}'(\mathbb{R}^n)$*.*

The significance of linearity is that it is compatible with the demands for differentiation. In fact, both Definitions 1 and 2 enable differentiation in a well-posed and operationally meaningful sense:

**Result 1 (Derivatives)** *Let* $\mathcal{D}$ *be a linear derivative operator. The derivative of a local sample* $F[\phi]$ *is the local sample defined by*

$$\mathcal{D}F[\phi] \stackrel{\text{def}}{=} F[\mathcal{D}^{\mathrm{T}}\phi]\,,$$

*in which* $\mathcal{D}^{\mathrm{T}}$ *is the transpose of* $\mathcal{D}$*. In integral form:*

$$\mathcal{D}F[\phi] = \int d\mathbf{z}\, f(\mathbf{z})\, \mathcal{D}^{\mathrm{T}}\phi(\mathbf{z})\,.$$

Thus differentiation is in fact integration: One can "extract a derivative" by linearly filtering the source data with a filter $\mathcal{D}^{\mathrm{T}}\phi$. Of course the result again depends on the filter as in the zeroth order case. Note that transposition brings in a minus sign if order is odd, *e.g.* $\nabla^{\mathrm{T}} = -\nabla$ (why?). Regarded as a fibre bundle the Gaussian family induces a so-called *local jet bundle* of the input image [7, 17].

### 2.4 Morphological Duality

Morphological duality is based (at least conceptually) on tactile probing. A sample of the form

$$F[\phi] = \sup_{\mathbf{z}}\left[f(\mathbf{z}) + \phi(\mathbf{z})\right]\,,$$

is known as the *dilation* of $f$ (at the origin) by *structuring element* $\phi$, *cf.* Fig. 3. A similar formula is obtained for the *erosion* if $\sup$ is replaced by $\inf$ and addition by subtraction. One can think of a physical probing device with profile $\phi$ touching the surface defined by the graph of $f$ from above, respectively from below. It should be clear from this picture that morphological probing is quite sensitive to the presence of noise spikes (exceptional amplitudes at isolated points).

A particularly interesting filter is the *quadratic structuring function* with arbitrary base point and size, which induces *dilation* and *erosion scale-spaces* analogous to (zeroth order) Gaussian scale-space in the linear theory [1, 2, 3, 5]. However, because of nonlinearity it is not clear whether one can define "morphological derivatives" analogous to Result 1, [4].

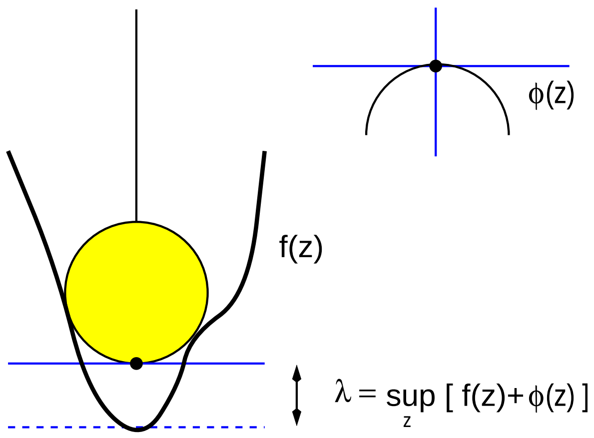


Figure 3: Dilation by a semicircle is an instance of a morphological probing mechanism. The radius of the circle—in general the width of the structuring element—is a measure of scale or resolution of the measurement probe.

## 2.5 Pseudo-Topological Duality

Consider a probe of the form

$$F_\mu[\phi] = \frac{1}{\mu} \ln \int d\mathbf{z}\, e^{\mu f(\mathbf{z})}\, \phi(\mathbf{z}) \qquad \mu \in \mathbb{R} \backslash \{0\}\,.$$

It is basically a linear probe but with a nonlinearity at the interface of source and filter. If certain weak conditions regarding the source data $f$ are met it can be shown that, if $\phi$ is a normalised Gaussian, both the "linear limit" $\mu \to 0$ as well as the two "morphological limits" $\mu \to \pm\infty$ exist and reproduce the previously encountered linear, respectively morphological scale-spaces. The linear limit even defines an isomorphism, whence the terminology "pseudo-topological duality". For nonzero finite $\mu$ it defines a probing mechanism which is in a precise sense in-between the linear and morphological ones. It can be interpreted as a soft-probing into a "surface layer" (whereby the surface itself is defined by the graph of $f$) of depth $\mathcal{O}(1/|\mu|)$, [8]. A particularly interesting observation is that pseudo-linear probing is less sensitive to noise spikes than the morphological limits; in the case of additive noise with a characteristic amplitude in the order of $1/|\mu|$, its behaviour is much like the noise averaging linear mechanism within the noise layer, while still resembling the morphological mechanism "deep down". Thus unlike morphological duality pseudo-linear duality can be adapted to data tolerance.

## 2.6 The Carry-Along Principle

There are always dual interpretations explaining an apparent change of grey-value: Either filter or input has changed. A "spatial transformation" (shift, rotation, scaling, *etc.*) pertains to a changing relationship between interacting *physical objects* defined in space (*i.c.* sources and filters), rather than some nonphysical deformation of space itself. Although this applies to duality in general, we focus on implications in the context of topological duality.

**Definition 3 (Push Forward)** *Let* $\boldsymbol{\theta} : \mathbb{R}^n \to \mathbb{R}^n : \mathbf{x} \mapsto \boldsymbol{\theta}(\mathbf{x})$ *be a spatial transformation. The push forward of a filter is then defined as the mapping*

$$\boldsymbol{\theta}_* : \Delta_{\mathbf{x}} \to \Delta_{\boldsymbol{\theta}(\mathbf{x})} : \phi \mapsto \boldsymbol{\theta}_*\phi \stackrel{\text{def}}{=} |\det \nabla \boldsymbol{\theta}^{\text{inv}}|\, \phi \circ \boldsymbol{\theta}^{\text{inv}}\,.$$

Subscripts attached to $\Delta$ indicate what happens to a filter's centre of gravity (whence the terminology). One naturally "pulls back" the source field in the dual view.

**Definition 4 (Pull Back)** *With $\boldsymbol{\theta}$ and its push forward $\boldsymbol{\theta}_*$ as defined in the previous definition, the pull back of the input image is defined as the mapping*

$$\boldsymbol{\theta}^* : \Sigma_{\boldsymbol{\theta}(\mathbf{x})} \to \Sigma_{\mathbf{x}} : F \mapsto \boldsymbol{\theta}^* F \quad \textit{defined by} \quad \boldsymbol{\theta}^* F[\phi] \stackrel{\text{def}}{=} F[\boldsymbol{\theta}_* \phi] \,.$$

Note that the base point ("focus of attention") now moves in *opposite* direction. These definitions may seem a bit abstract on first sight, but it is easily appreciated that both $\boldsymbol{\theta}_*\phi$ (filter transformation) as well as $\boldsymbol{\theta}^* F$ (transformation of input image) represent more sensible actions than $\boldsymbol{\theta}(\mathbf{x})$ (for how could one transform a "void"?), and in a way it is better to say that one of the former two induces the others. Indeed, it is easy to think of dual filter/image transformations without existence of a corresponding spatial transformation, but impossible to make sense of the latter without any manifestation on physical objects (*cf.* Result 1: Derivation has no counterpart in "empty space"). If one writes Definition 4 in integral form it will be seen that it is basically a change of integration dummies; the function-under-the-integral corresponding to the transformed source, $\boldsymbol{\theta}^* F$, is just the original one evaluated at the transformed point, *i.e.* $f \circ \boldsymbol{\theta}$ ("scalar transformation"):

$$\boldsymbol{\theta}^* F[\phi] = \int d\mathbf{z}\, f(\boldsymbol{\theta}(\mathbf{z}))\, \phi(\mathbf{z}) = [\text{subst.}\, \mathbf{y} = \boldsymbol{\theta}(\mathbf{z})] = \int d\mathbf{y}\, |\det \nabla \boldsymbol{\theta}^{\text{inv}}(\mathbf{y})\,|\, f(\mathbf{y})\, \phi(\boldsymbol{\theta}^{\text{inv}}(\mathbf{y})) = F[\boldsymbol{\theta}_* \phi] \,.$$

Example: One can shift a patient underneath a scanner, or move the scanner in opposite sense. This generalises to any (invertible) transformation. The relevant formulas in this case are: $\boldsymbol{\theta}(\mathbf{z}) = \mathbf{z} + \mathbf{x}$, $\boldsymbol{\theta}^* f(\mathbf{z}) = f(\mathbf{z} + \mathbf{x})$, $\boldsymbol{\theta}_*\phi(\mathbf{z}) = \phi(\mathbf{z} - \mathbf{x})$, and the above equality can be rewritten as the correlation $f \star \phi(\mathbf{x})$ in two equivalent ways.

Definitions 3 and 4 implement the so-called "carry-along principle", which forms the basis for many image manipulations. This is illustrated in subsequent sections.

## 2.7 From Samples to Images

The step from a mere grey-value sample $F[\phi]$ to an actual output image $F[\boldsymbol{\theta}_*\phi] = f \star \phi(\mathbf{x})$ in the patient example relies on the push forward/pull back principle based on spatial translations $\boldsymbol{\theta}(\mathbf{z}) = \mathbf{z} + \mathbf{x}$ over all vectors $\mathbf{x}$ within the relevant field of view (recall Footnote 2).

However, a complete and consistent image model accounts for *all* symmetries of space and time, *i.e.* not only translations (homogeneity), but also spatial rotations (isotropy) and spacetime scalings (scale invariance). If $\phi$ is rotationally invariant the result of carry-along will be a multiscale representation of $f$. This is at the same time a "minimal" representation since nontrivial filters in Schwartz space may be rotationally invariant but cannot be scale-invariant[3] (recall the Jacobian in the push forward formula).

## 2.8 Temporal Causality

It is sometimes argued that because of infinite filter support the Gaussian family is not suited in case *temporal causality* is a prerequisite, as in active vision. However, the following "Koenderink trick", exploiting the carry-along principle, can be applied [14].

Consider 1D temporal sequences $f(t)$ for simplicity. The basic observation is that there must be *some* time domain in which the Gaussian family makes sense (by virtue of its uniqueness there is no alternative), say parametrised by a parameter $s \in \mathbb{R}$. The physically reasonable domain for a visual

[3]Scale-invariant filters $\phi \notin \mathcal{S}(\mathbb{R}^n)$ do exist, however: $\phi(\mathbf{z}) = \|\mathbf{z}\|^{-n}$ is scale-invariant in $n$-dimensional space.

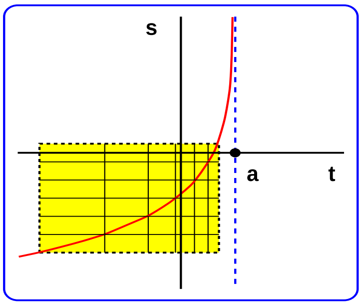


Figure 4: The isomorphism $t = \tau(s; a)$ discards the unknown future $t \geq a$. The asymptote indicates the time horizon $t = a$. The box delimits a typical time window for a real-time system. Uniform sampling of the $\mathbf{S}$-domain implies a graded resolution history.

system actively participating in the world is of course the history part of the time axis. Therefore discard the unknown future by introducing a time horizon (the present moment) and mapping the past semi-axis onto the $s$-domain: Fig. 4. Once such an isomorphism has been established[4], say $t = \tau(s; a)$—note that it depends on the present moment $a$—proceed as usual: The isomorphism acts on a filter $\phi(s)$ by the recipe of "push forward", yielding a filter $\tau_* \phi(t; a)$—which again depends on $a$—producing the desired filter profile in the $t$-domain: Fig. 5. In this way Koenderink has utilised the push forward principle to produce filters incorporating *manifest temporal causality* in an unambiguous way.

Recall that

$$\tau^* F[\phi] = F[\tau_* \phi] \, .$$

In other words, the alternative "realist's view" or "causal world picture" obtained by the dual action of pull back applied to the source data cannot be disqualified on objective grounds. The $F$ on the r.h.s. corresponds to a "function under the integral" $f(t)$, which could be a fully recorded video tape being causally processed, or looked at by an observer in real-time. Of course the content of the tape does *not* depend on the present moment $a$. Causality is introduced here by the act of filtering (observation) by causal filters $\tau_* \phi$. The $\tau^* F$ on the l.h.s. is a signal in the $s$-domain—which *does* depend on $a$—and reflects the *signal history relative to acquisition time* $a$, subjected to shift invariant filters that do not depend on the time of acquisition $a$.

The causal representation should appeal to biologists for its built-in fading memory characteristics and emphasis on recent events, which is of course crucial in generating efficacious action (escaping predators, catching prey, *etc.*).

### 2.9 Motion

The appropriate differential tool in the context of motion is the *Lie derivative*. It expresses the rate of change of a quantity when moving in the direction of a (spatial) vector field $\mathbf{v}$, say, and is proportional to that vector field. The following definition therefore suggests itself:

$$(\mathbf{v} \cdot \nabla) F[\phi] \stackrel{\text{def}}{=} F[-\nabla \cdot (\mathbf{v}\, \phi)] \, .$$

The l.h.s. corresponds to $\mathrm{L}_{\mathbf{v}} f = (\mathbf{v} \cdot \nabla) f$, *i.e.* the Lie derivative of a scalar function. Transposed to filter space one observes that the dual Lie derivative of a filter $\phi$ must be $\mathrm{L}_{\mathbf{v}}^{\mathrm{T}} \phi = -\nabla \cdot (\mathbf{v}\, \phi)$. Apart from the minus sign we see that the gradient operator acts on filter as well as vector field, containing an additional divergence term $\phi \, \mathrm{div} \, \mathbf{v}$ not present in the scalar case.

[4] There exists a unique "canonical" isomorphism [6], which Koenderink arrives at by physical reasoning [14].

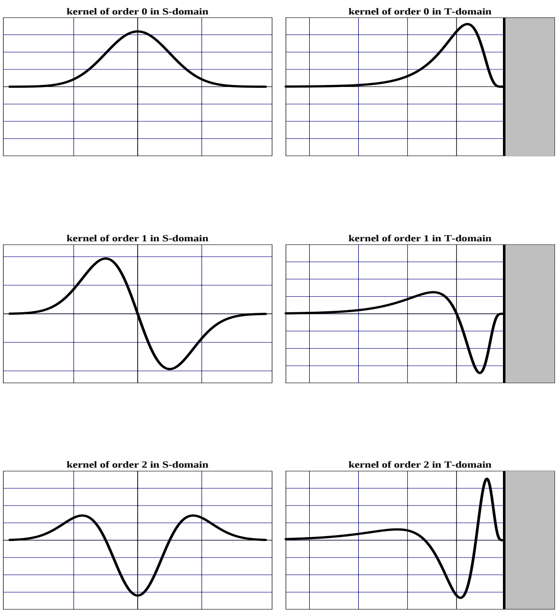


Figure 5: Comparison of causal point operator and its first and second order derivatives in $s$ and $t$ domains. The map $\tau_*$—which depends on $a$—maps the filters in $s$-representation on the left to the $t$-representation on the right, which corresponds to a fixed delay $a - t$. All causal filters vanish smoothly towards the time horizon. (Of course there is no such moment in the left graphs.) The shaded region indicates the unrevealed future.

The spatial motion field $\mathbf{v}$ is defined such that if one comoves with the induced flow *some measurable entity* is preserved. The simplest such entity is $F[\phi]$ itself, in which case one obtains the "motion constraint equation" while correctly accounting for duality:

$$\frac{d}{dt}F[\phi] = (\mathbf{v} \cdot \nabla)F[\phi] + \frac{\partial}{\partial t}F[\phi] = 0\,,$$

or, in dual form (omitting the overall minus),

$$F[\frac{d}{dt}\phi] = F[\nabla \cdot (\mathbf{v}\,\phi)] + F[\frac{\partial}{\partial t}\phi] = 0\,.$$

An algorithm for solving the latter has appeared in the International Journal of Computer Vision [9], which also discusses the case of density sources.

## References


[1] R. van den Boomgaard. *Mathematical Morphology: Extensions towards Computer Vision*. PhD thesis, University of Amsterdam, March 23 1992.

[2] R. van den Boomgaard. The morphological equivalent of the Gauss convolution. *Nieuw Archief voor Wiskunde*, 10(3):219–236, November 1992.

[3] R. van den Boomgaard and L. Dorst. The morphological equivalent of the Gaussian scale-space. In Sporring et al. [20], chapter 15, pages 203–220.

[4] R. van den Boomgaard and A. W. M. Smeulders. The morphological structure of images, the differential equations of morphological scale-space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(11):1101–1113, November 1994.

[5] L. Dorst and R. van den Boomgaard. Morphological signal processing and the slope transform. *Signal Processing*, 38:79–98, 1994.

[6] L. M. J. Florack. *Image Structure*, volume 10 of *Computational Imaging and Vision Series*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1997.

[7] L. M. J. Florack, B. M. ter Haar Romeny, J. J. Koenderink, and M. A. Viergever. The Gaussian scale-space paradigm and the multiscale local jet. *International Journal of Computer Vision*, 18(1):61–75, April 1996.

[8] L. M. J. Florack and R. Maas. Pseudo-linear scale-space theory: Towards the integration of linear and morphological scale-space paradigms. Technical Report UU-CS-1997-25, Department of Computer Science, Utrecht University, 1997.

[9] L. M. J. Florack, W. J. Niessen, and M. Nielsen. The intrinsic structure of optic flow incorporating measurement duality. *International Journal of Computer Vision*, 27(3):263–286, 1998.

[10] B. M. ter Haar Romeny, L. M. J. Florack, J. J. Koenderink, and M. A. Viergever, editors. *Scale-Space Theory in Computer Vision: Proceedings of the First International Conference, Scale-Space'97, Utrecht, The Netherlands*, volume 1252 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, July 1997.

[11] B. K. P. Horn. *Robot Vision*. MIT Press, Cambridge, 1986.

[12] T. Iijima. Basic theory on normalization of a pattern (in case of typical one-dimensional pattern). *Bulletin of Electrical Laboratory*, 26:368–388, 1962. (In Japanese).

[13] J. J. Koenderink. The structure of images. *Biological Cybernetics*, 50:363–370, 1984.

[14] J. J. Koenderink. Scale-time. *Biological Cybernetics*, 58:159–162, 1988.

[15] J. J. Koenderink. Local image structure. In P. Johansen and S. Olsen, editors, *Theory & Applications of Image Analysis*, volume 2 of *Series in Machine Perception and Artificial Intelligence*, pages 15–21. World Scientific, Singapore, 1992. Selected papers from the 7th Scandinavian Conference on Image Analysis.

[16] J. J. Koenderink. What is a “feature”? *Journal of Intelligent Systems*, 3(1):49–82, 1993.

[17] J. J. Koenderink and A. J. van Doorn. Receptive field families. *Biological Cybernetics*, 63:291–298, 1990.

[18] T. Lindeberg. *Scale-Space Theory in Computer Vision*. The Kluwer International Series in Engineering and Computer Science. Kluwer Academic Publishers, 1994.

[19] L. Schwartz. *Théorie des Distributions.* Hermann, Paris, second edition, 1966.

[20] J. Sporring, M. Nielsen, L. M. J. Florack, and P. Johansen, editors. *Gaussian Scale-Space Theory*, volume 8 of *Computational Imaging and Vision Series*. Kluwer Academic Publishers, Dordrecht, 1997.

[21] A. P. Witkin. Scale-space filtering. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 1019–1022, Karlsruhe, Germany, 1983.