14. Deep structure II. catastrophe theory

Erik Dam and Bart M. ter Haar Romeny

Out[100]=

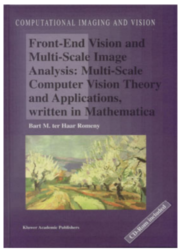

From the book:
ter Haar Romeny, B. M.:
Front–end vision and multi–scale image analysis:
Multi–scale computer vision theory
and applications, written in Mathematica.
ISBN: 978–1402015038
Publisher: Springer Dordrecht (2003)
Download: **www.romeny.info/FEV**

Initialization

In[31]:=

```
url = "https://www.romeny.info/FEV-CD/images/";
SetOptions[ArrayPlot, ColorFunction → GrayLevel];
$HistoryLength = 0;
gD[im_, nx_, ny_, σ_] := DerivativeFilter[im, {ny, nx}, σ];
Off[DerivativeFilter::arg1];
gDf[im_List, nx_, ny_, σ_] :=
 Module[{xres, yres, gdkernel},
        {yres, xres} = Dimensions[im];

  gdkernel = N[Table[Evaluate[D[1/(2 π σ^2) Exp[-(x^2 + y^2)/(2 σ^2)], {x, nx}, {y, ny}]],
      {y, -(yres - 1) / 2, (yres - 1) / 2},
      {x, -(xres - 1) / 2, (xres - 1) / 2}]];
        Chop[N[√(xres yres) InverseFourier[Fourier[im] ×
        Fourier[RotateLeft[gdkernel, {yres / 2, xres / 2}]]]]]]
AutoCollapse[] := (If[$FrontEnd =!= $Failed,
    SelectionMove[EvaluationNotebook[], All, GeneratedCell];
    FrontEndTokenExecute["SelectionCloseUnselectedCells"]]);
```

14.1 Catastrophes and singularities

The previous chapter illustrates a number of approaches that explore the deep structure. However, there are a number of caveats. The edge focusing technique

implicitly assumes that the edges for the signal can be located at the adjacent lower scale level in a small neighborhood around the location at the current scale. As mentioned, no formal scheme for defining the size and shape of the neighborhood is presented. Furthermore, this method ignores the problems encountered when edge points merge or split with increasing scale.
Analogously, the multi-scale watershed segmentation depends on the behaviour of the dissimilarity measure singularities (the notion of dissimilarity is defined in chapter 13, section 6.1). Even though the linking of the watershed catchment basins is quite robust due to the matching of regions (opposed to tracking of points as in the edge focusing paradigm), the linking in the presence of merges and splits of regions is not explicitly established above. Without knowledge of how the dissimilarity singularities can behave in scale-space, we can only hope that the method will work on other images than the ones used for the illustrations.
The finding of explicit schemes for linking in the neighborhood of changes in the singularity structures requires explicit knowledge of the changes. A change in the singularity structure in denoted a *catastrophe*. Catastrophe theory (or with a broader term: singularity theory) is the theory that analyses and describes these changes. Catastrophe theory allows prediction of which changes in the singularity structure can be expected. Thereby the schemes that involve the singularities can be designed to detect and handle these events as special cases.

Actually, this analysis was done for the multi-scale watershed segmentation method in [Olsen1997].

The field of catastrophe theory is vast and rather complicated. The focus of this introduction is to give a condensed introduction of the central definitions, with an intuitive understanding of the effects we often observe in 'deep scale-space'.

14.2 Evolution of image singularities in scale-space

An important property of linear scale-space is the overall simplifying effect of increasing scale.
In general, this implies that the number of appearances of a given image feature decreases as scale increases. In particular, this is the qualitative behavior for the image singularities - blurred versions of an image will in general contain less singularities than the original one.
As mentioned in chapter 2 section 2.8, this notion is formalized by Florack [Florack-1992a], which leads to a prediction on the number of singularities in n-dimensional signals/images. Specifically, the number of singularities can be expected to decrease with a slope of -1 for 1D signals and -2 for 2D images (in general: $-n$ for n-

D signals; see chapter 1 for the reasoning to derive this).
More precisely, when the scale levels are generated with the usual exponential sampling $\sigma = \epsilon\, e^{\tau}$ the logarithm of the number of singularities decrease with these slopes as a function of the scale parameter τ.
This is illustrated for the "blobs" image from the scale selection section in the previous chapter. For simplicity we count the number of maxima instead of the number of singularities (then we can use the function **nMaxima**). Following the argument of Florack, the relative decrease in the number of maxima is equivalent to the decrease in the number of singularities.

In[48]:=
```
countMaxima[im_] := Module[{p, d = Depth[im] - 1},
   p = Times @@ Table[(Sign[im - Map[RotateLeft, im, {i}]] + 1)
        (Sign[im - Map[RotateRight, im, {i}]] + 1), {i, 0, d - 1}] / 4^d;
   Count[Flatten[p], 1]];
noisyblobs =
  255 ImageData[ColorConvert[Import["blobs.gif"], "Grayscale"]] +
   10 Table[Random[], {128}, {128}];
Off[General::munfl]
levels = 15;
step = 0.15;
data = Table[{step t, countMaxima[nb[t] = gDf[noisyblobs, 0, 0, E^(step t)]]},
  {t, levels}];
GraphicsRow[Append[ArrayPlot /@ {nb[1], nb[8], nb[15]}, ListLogPlot[
   data, Joined → True, AxesLabel → {"t", "N"}]], ImageSize → 500]
Print["Slope = ", Coefficient[Fit[Log[data], {1, t}, t], t]];
```

Out[52]=

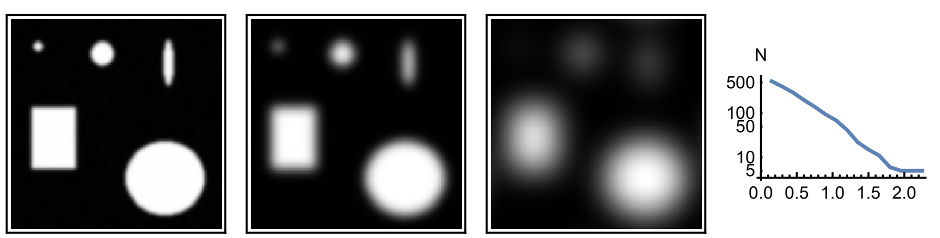


```
Slope = -2.04188
```

Figure 14.1 The evolution of the number of singularities for a set of noisy 2D blobs. Images blurred with $\sigma = e^{0.15} = 1.16$, $\sigma = 3.32$ and $\sigma = 9.49$ pixels. The observed slope is close to the predicted value of -2.

The blob image from figure 13.3 is used with noise added (S/N ratio = 1/10). First we define the scale levels used (such that $\sigma = \epsilon\, e^{\text{step}\,\tau}$ where the **step** ensures sufficiently small scale steps, and $\epsilon = 1$). We display the lowest, middle and highest levels of the selected range of 15 scales and plot the logarithm of the number of maxima as a function of **step t**. The slope is calculated with a linear least square **Fit**.

The overall effect of blurring signals and images is simplification. This is exemplified by the decrease of maxima above. However, this general notion reveals nothing about *how* the singularities disappear. More specifically, it gives no insight into the local structure of the signals and images at the specific point in scale-space where a singularity is annihilated. Furthermore, we get no information about whether singularities are *created* as well. In order to investigate these matters we need a bit of mathematics - introduced in the following through some central concepts from *Catastrophe Theory*.

Task 14.1. Check the expected decrease in the number of singularities for a 1D signal as it was done for a 2D image above. Use the 1D noisystep signal.

14.3 Catastrophe theory basics

The field of catastrophe theory is quite extensive and complicated. This introduction focuses on giving an intuitive understanding of the concepts most related to computer vision and image processing. Therefore, the presentation is also somewhat less strict than possible. For a comprehensive introduction see [Gilmore1981].

The singularities are central feature points of an image - or more general for a function. The singularities alone offer a good qualitative description of the structure of a function. When a function undergoes an evolution it is therefore central to capture where the set of singularities change, in order to analyze the evolution. These points are denoted *catastrophes* since this is where the qualitative structure changes.
In order to describe these events properly, a few definitions are needed. They are presented quite briefly - the concepts are then illustrated by a number of examples.

14.3.1 Functions

For a smooth (≡ infinitely differentiable, indicated with C^∞) function *f* the parameters are divided into *n state* and *m control* parameters:
$f(x_1, .., x_n, c_1, .., c_m) \in C^\infty(\mathbb{R}^{n+m}, \mathbb{R})$. For the intuitive understanding, think for the concept of scale-space of the *state* parameters as *spatial* coordinates, and think of a single *control* parameter, namely *scale*.

14.3.2 Characterization of points

For a smooth function $f(x_1, .., x_n, c_1, .., c_m) \in C^\infty(\mathbb{R}^{n+m}, \mathbb{R})$ a given point $p \in \mathbb{R}^{n+m}$ is either:

Regular : $\exists\, l \in [1..n]$ such that $\frac{\partial f}{\partial x_l} \neq 0$

$$\text{Morse Singularity}: \quad \frac{\partial f}{\partial x_i} = 0 \quad \text{and} \quad \text{Det}\left(\frac{\partial f}{\partial x_i x_j}\right) \neq 0$$

$$\text{Catastrophe}: \quad \frac{\partial f}{\partial x_i} = 0 \quad \text{and} \quad \text{Det}\left(\frac{\partial f}{\partial x_i x_j}\right) = 0$$

In the equations above, tensor notation is used for the spatial parameters with subscripts *i* and *j* (but not for *l*).

Note that a *regular* point is only required to be a singularity with respect to the state (or spatial) parameters. Singularities and catastrophes are found at locations where the gradient is zero, so at horizontal locations in the image intensity landscape. A catastrophe differs from a singularity in that the second order structure has a degenerate Hessian matrix, i.e. the determinant of the Hessian matrix vanishes, and the Hessian matrix is thus singular in catastrophes.
Morse singularities are non-degenerate singularities in the state (or spatial) parameters. When the determinant of the Hessian matrix for the state parameters is vanishing the singularity becomes degenerate.
A non-degenerate singularity is *stable.* This means that a slight perturbation of the function will not change the local qualitative structure of the function - there will still be a singularity of the same type near the original with a value close to the original.
Degenerate singularities are *not stable* - a slight perturbation can cause a change in the local structure of the function. Such a perturbation could be caused by a slight change in the *control* parameters. Specifically, in scale-space the singularity structure changes at degenerate singularities when the scale is changed. These are the *catastrophe points* where *creations* or *annihilations* of singularities occur.

14.3.3 Structural equivalence

Two functions are *locally structurally equivalent* at a point if a diffeomorphism (a smooth invertible function with smooth inverse) exists such that a change of the coordinate system for one function with this diffeomorphism will make the functions equal in a neighborhood around the point.
Two functions are *globally structurally equivalent* if they are *locally structurally equivalent* at all points.

We will not formalize this definition in mathematical notation. The key point is that these definitions imply that two functions are structurally equivalent if their singularity structures are corresponding - the topological ordering and the types of the singularities are equivalent.

Slight perturbations of a function will in general leave it structurally equivalent with itself. The singularities move a bit and change value, but the topological structure remains the same. However, in the presence of catastrophe points, a slight perturbation will change the singularity structure and the function no longer remains structurally equivalent to itself.

14.3.4 Local characterization of functions

Analogous to the characterization of points into three classes, the local structure of a function is characterized by the following theorems. Here, f is a given smooth function $f(x_1, .., x_n, c_1, .., c_m) \in C^\infty(\mathbb{R}^{n+m}, \mathbb{R})$.

Implicit Function Theorem:
At a given regular point the function f is locally structurally equivalent with the function g where

$$g(x_1, .., x_n, c_1, .., c_m) = x_i$$

In other words, the Implicit Function Theorem states that at a regular point the function is locally equivalent with its tangent plane.

The Morse Lemma:
At a given Morse singularity point the function f is locally structurally equivalent with the function g where

$$g(x_1, .., x_n, c_1, .., c_m) = \frac{1}{2!} x_i x_j$$

The Morse Lemma states that at Morse singularity points the local structure is defined by the second order terms.

The Splitting Lemma:
At a given catastrophe point for the function f, the Eigenvalues for the Hessian matrix can be ordered by absolute value with the first d being zero - corresponding to the degree of degeneracy. The function f is then locally structurally equivalent with the function g where

$$g(x_1, .., x_n, c_1, .., c_m) = g_{nm}(x_1, .., x_d) + \sum_{i=d+1}^{n} \sum_{j=d+1}^{n} \frac{1}{2!} x_i x_j$$

The *Splitting Lemma* states that the function can be split into two parts: A non-Morse part and a Morse part. Accordingly, the parameters are split into "bad" and "good" parameters. The bad parameters correspond to the degenerate directions of the Hessian matrix. However, the *Splitting Lemma* does not characterize the local structure of the non-Morse part of the function. This is done by a theorem by the French mathematician René Thom (1923-).

14.3.5 Thom's theorem

Let f be a given smooth function $f(x_1, .., x_n, c_1, .., c_m) \in C^\infty(\mathbb{R}^{n+m}, \mathbb{R})$. *At a catastrophe point the eigenvalues for the Hessian matrix can be ordered by absolute value* *with the first d being zero. The function is then locally structurally equivalent with a function g where*

$$g(x_1, .., x_n, c_1, .., c_m) = CatGerm(d) + Perturb(d, m) + \alpha\, x_i\, x_j$$

If m≤5 then CatGerm(d) is one of the Catastrophe Germs and Pert(d,m) is the corresponding Pertubation listed in the table below:

Name	Nickname	*m*	*d*	CatGerm (*d*)	Perturb (*d*, *m*)
A_2	Fold	1	1	x^3	$c_1 x$
$A_{\pm 3}$	Cusp	2	1	$\pm x^4$	$c_1 x + c_2 x^2$
A_4	Swallowtail	3	1	x^5	$c_1 x + c_2 x^2 + c_3 x^3$
$A_{\pm 5}$	Butterfly	4	1	$\pm x^6$	$c_1 x + c_2 x^2 + c_3 x^3 + c_4 x^4$
A_6	□	5	1	x^7	$c_1 x + c_2 x^2 + c_3 x^3 + c_4 x^4 + c_5 x^5$
D_{-4}	Elliptic Umbilic	3	2	$x^2 y - y^3$	$c_1 x + c_1 y + c_3 y^2$
D_{+4}	Hyperbolic Umbilic	3	2	$x^2 y + y^3$	$c_1 x + c_1 y + c_3 y^2$
D_5	Parabolic Umbilic	4	2	$x^2 y + y^4$	$c_1 x + c_1 y + c_3 x^2 + c_4 y^2$
D_{-6}	□	5	2	$x^2 y - y^5$	$c_1 x + c_2 y + c_3 x^2 + c_4 y^2 + c_5 y^3$
D_{+6}	□	5	2	$x^2 y + y^5$	$c_1 x + c_2 y + c_3 x^2 + c_4 y^2 + c_5 y^3$
$E_{\pm 6}$	□	5	2	$x^3 \pm y^4$	$c_1 x + c_2 y + c_3 xy + c_4 y^2 + c_5 xy^2$

Figure 14.2 Table of elementary catastrophes for $m \leq 5$. The names in the first column were originally proposed by Thom [Thom1975]. The nicknames come from their visual appearance (see MathWorld [http://mathworld.wolfram.com/Catastrophe.html] with interactive plots, Gray [Gray1993], Bruce & Giblin [Bruce1984] and Scanns [Scanns2000]). The c_1 to c_5 factors are also called the control factors.

The *Catastrophe Germ* is the local structure of the function for the specific set of control parameters at the catastrophe point.

The *Pertubation* terms determine how the function behaves when the control parameters vary (in a neighborhood around the catastrophe point).

14.3.6 Generic property

A property for a system is *generic* if an open, dense subset of the system possesses the property. In probabilistic terms, a property is *generic* if it is possessed with probability one.

In this context, the term generic is used to characterize which catastrophes are generic for images or for differential operators on images - these are the so-called *generic events*.

14.3.7 Dimensionality

Analysis of the dimension of the involved spaces can often determine whether a property is generic. As an example, lets look at the set of singularities in an *n*-dimensional image. A singularity point is determined by all *n* first order partial derivatives equaling zero. This means that we have *n* conditions in a *n*-dimensional space. Under the assumption that these conditions are independent, the space where the conditions are met is a $n - n = 0$ dimensional space. This means that the set of singularities contains only isolated points - or more precisely, a singularity point is generically isolated. We know that for a 2D image the singularities are the maxima, minima and saddle points of the intensity landscape, which are easily recognized as isolated points.

What about catastrophes? The potential catastrophe point is required to be a singularity (n conditions) and then the Hessian is required to have a degenerate direction (an extra condition). This means that $n + 1$ conditions are to be met resulting in an n dimensional space. Or in other words: generically, a given image contains no catastrophe points.

In scale-space we have an extra parameter - the scale.

This means that in scale-space the set of singularities is generically a 1-dimensional space (the 'path' of the singularity over scale) and that catastrophes do generically occur in isolated points.

This very short and informal treatment of the dimensionalities of the singularity and the catastrophe sets is only meant as an appetizer. In order to present the above considerations properly mathematically, the image and the set of conditions should be represented as manifolds in jet-space where the independency of the terms can be investigated through the concept of *transversality*. However, this is far beyond the scope of this introduction - for a richly illustrated interactive tutorial in *Mathematica* see [Sanns2000], for a more comprehensive treatment see [Gilmore1981, Florack1994b, Olsen2000, Dam2000, Bruce1984].

14.3.8 Illustration of the concepts

The example below shows a function with two state parameters x and y and one control parameter a. The local structure differs for three choices of the control parameter a: To the left the function has a saddle and a local minimum, in the middle a single saddle, and to the right no singularities at all. In general, the function has two singularities for negative values of a, one singularity for $a = 0$, and no singularities for positive a.

In[54]:=
```
f[x_, y_, a_] := x^3 - 24 y^2 + a x;
GraphicsRow[Plot3D[f[x, y, #], {x, -30, 30}, {y, -30, 30}] & /@
  {-300, 0, 300}, ImageSize → 500]
```
Out[54]=

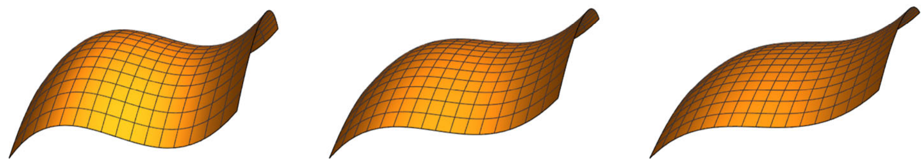

Figure 14.3 The evolution of the Fold catastrophe as a function of the control parameter a. As the control parameter changes, the structure of the function changes. For negative values of the control parameter, the function has a local maximum and a saddle point. For positive values of the control parameter, the function has no singularities in the neighborhood of the observed point. The catastrophe occurs for control parameter equal to zero.

Since the singularity structure is different for positive and negative values of a, there must be a catastrophe point for $a = 0$.

And reassuringly, for $a = 0$ the singularity at $(x, y) = (0, 0)$ is a catastrophe point as well. This is easily verified mathematically: for $a = 0$ the determinant of the Hessian is $6\,x\,(-48)$, which is zero at the singularity point $(x, y) = (0, 0)$.

The *Splitting Lemma* states that we can split the state parameters into "good" and "bad" parameters (locally around the catastrophe point $(x, y, a) = (0, 0, 0)$ where the function changes singularity structure for increasing a).

Actually, it need not be the original state parameters that are split - the degenerate direction need not be aligned with the original axes. However, in this case the "bad" parameter is the x. As stated by the splitting lemma, the function can be split into a part with the bad parameters, and a part with the good parameters in the shape of a sum of second order monomials (a monomial is a polynomial consisting of a product of powers of variables, e.g., x, $x\,y^3$, $x^4\,y\,z^2$, etc). The non-Morse part of the function can be recognized from the list of *catastrophe germs* in *Thom's Theorem*. It is known as the *fold catastrophe*. It should be noted that - since there is only one control parameter - this is the only generic catastrophe.

In[55]:=
```
fold[x_, a_] := x^3 + a x;
GraphicsRow[
 Plot[fold[x, #], {x, -30, 30}, PlotLabel → "a=" <> ToString[#],
    AxesLabel → {"x", ""}, Ticks → None] & /@
  {-300, 0, 300}, ImageSize → 500]
```

Out[56]=

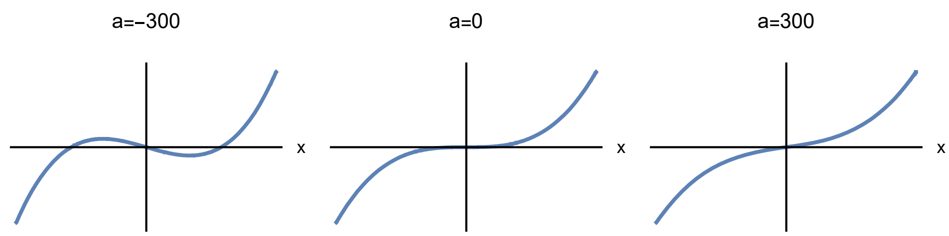

Figure 14.4 The "bad" parameter x from the previous example. The original function can be split into the Morse and the non-Morse part. The non-Morse part is structurally equivalent to the *fold* catastrophe.

The canonical fold catastrophe is further analysed below. We derive the singularity and catastrophe sets for the function fold(x, a):

In[57]:=
```
Solve[∂x fold[x, a] == 0, {a}]
Solve[{∂x fold[x, a] == 0, ∂x,x fold[x, a] == 0}, {a, x}]
```

Out[57]=

$\{\{a \to -3x^2\}\}$

Out[58]=

$\{\{a \to 0, x \to 0\}\}$

The above description of the singularity set is used to plot the *fingerprint* or the *signature* of the function - the position of the singularities against the value of control parameter:

In[59]:=
```
Plot[-3 x^2, {x, -1, 1}, ImageSize → 250]
```

Out[59]=

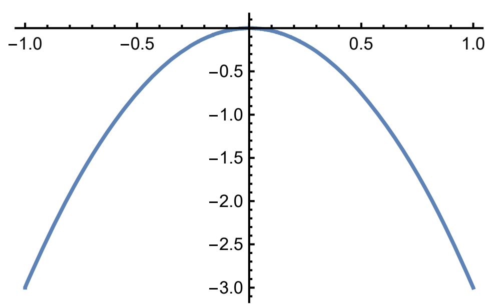

Figure 14.5 The *fold* catastrophe. The catastrophe at a=0 is where the two singularity strings meet and annihilate. A point where several singularity strings meet is also denoted a *bifurcation*. This is the *fingerprint* of the fold catastrophe. Compare with the catastrophes in the fingerprint in figure 14.11.

The *cusp* catastrophe from Thom's theorem is also commonly encountered. The canonical germ and perturbation that give rise to the cusp catastrophe is

$x^4 + c_1 x^2 + c_2 x$. Since there are two control parameters, this is slightly more complicated than the fold catastrophe.

Task 14.2 Illustrate the fingerprint for the cusp catastrophe (like the illustration in figure 14.5 for the fold catastrophe).

```
In[60]:= Clear[cusp]; cusp[x_, c1_, c2_] := x^4 + c2 x^2 + c1 x;
GraphicsGrid[
 Table[Plot[cusp[x, c1, c2], {x, -10, 10}, Axes → None],
  {c2, 0, -30, -15}, {c1, -80, 80, 40}], ImageSize → 450]
```

Out[61]=

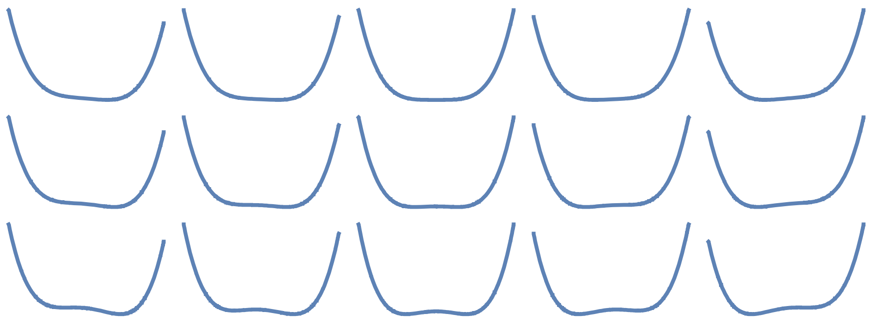

Figure 14.6 Various perturbed shapes for the cusp catastrophe germ. It has two stable states: one with a single minimum and one with a minimum,maximum,minimum triple (or the same states with maximum and minimum switched). Conceptually, one control parameter allows transition between the states by "tilting" the two minima until one minima is annihilated with the central maxima (or the reverse process) - these events are fold catastrophes. The other control parameter allows transition between the states by letting the two minima approach each other until they are merged together with the maxima into one single minimum. The cusp catastrophe point is where both control parameters are zero.

Here are a few other illustrations and some plot commands to study them:

```
In[63]:= GraphicsRow[{
    cusp = ContourPlot[x^3 - y^2, {x, -.5, 4},
      {y, -6, 6}, Contours → {0}, PlotLabel → "Cusp"],
    cusp3D = ContourPlot3D[4 x^3 + 2 u x + v,
      {u, -2.5, 2}, {v, -2, 2}, {x, -1, 1}, PlotPoints → 6,
      PlotLabel → "Cusp3D", ViewPoint → {-4.000, 1.647, 2.524}],
    swallowtail = ParametricPlot3D[{u v^2 + 3 v^4, -2 u v - 4 v^3, u}, {u, -2, 2},
      {v, -.8, .8}, BoxRatios → {1, 1, 1}, PlotLabel → "Swallowtail",
      ViewPoint → {4.000, -1.390, -3.520}]}, ImageSize → 500]
Out[63]=
```

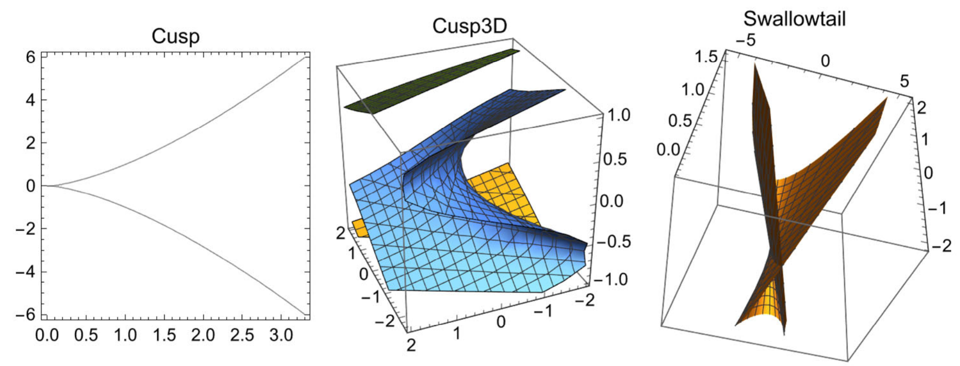

Figure 14.7 The *cusp* (2D and 3D) and the *swallowtail* catastrophe. From the wonderful booklet by [Sanns2000]. See also:
MathWorld [http://mathworld.wolfram.com/Catastrophe.html].

14.4 Catastrophe theory in scale-space

As stated earlier, the natural application of catastrophe theory towards scale-space theory is to view the spatial parameters as state parameters and the scale as the single control parameter. The differentiability of the scale-space even ensures that the functions are smooth. However, it is slightly more complicated than that. In scale-space theory, there is a fixed connection between the spatial parameters and the scale parameter given by the diffusion equation: $L_t = L_{xx} + L_{yy}$. This gives a severe restriction compared to the general space of functions with one control parameter.

Thereby, the canonical catastrophe germs listed in Thom's theorem might not apply to scale-space images. Fortunately, the work of James Damon (UNC) reveals that we can in fact still apply similar results to Thom's theorem for all practical purposes [Damon1995, Damon1997].

Another caveat in scale-space singularity analysis is the image. Images from natural scenes behave nicely, but artificial test images often possess nasty properties. Two typical examples are symmetry and areas with constant intensity. Across a symmetry axis singularities appear in pairs. This means that the expected catastrophes appear in variants where two (or more) symmetric catastrophes occur simultaneously at the same place (an example of this is shown in the next section). This apparently causes non-generic catastrophes to appear. Actually, it is not the catastrophes that are non-generic - symmetry in images is non-generic.

Areas with constant intensity in artificial test images can cause unexpected results. In theory this should cause no problems since the areas no longer have constants intensity at any given scale > 0. In practice, implementations with blurring kernels of limited size will however leave areas with constant intensity (gradually smaller with increasing scale). A simple consequence is apparent areas of singularities - as opposed to the expected isolated points.

14.4.1 Generic events for differential operators

Thom's theorem states that the only generic catastrophe with only one control parameter is the fold.
At first hand, we would therefore not expect to encounter any other catastrophes in scale-space, where we only have scale as control parameter. However most applications, like edge detection, examine singularities for differential operators and not singularities for the raw image. Depending on the differential operator this induces other catastrophes as well.

An example of higher order singularities induced by a differential operator is illustrated in the following. The operator is the square of the first derivative of the original image (the gradient squared). In order to understand why this simple operator can induce other generic catastrophes we look at the Taylor series expansion of a one-dimensional function *f* (around 0 for simplicity) and the derivative of this expansion squared.

```
In[64]:= Clear[f]; Series[f[x], {x, 0, 4}]
```

Out[64]=

$$f[0] + f'[0]\,x + \frac{1}{2} f''[0]\,x^2 + \frac{1}{6} f^{(3)}[0]\,x^3 + \frac{1}{24} f^{(4)}[0]\,x^4 + O[x]^5$$

When we look for singularities the zeroth order term is not interesting. Intuitively we can use the spatial parameter x as a free parameter that allows us to find points p where $f_x(p)$ is zero. Therefore we find singularities in generic signals (and images). When we have a control parameter, we can turn this extra "knob" until we find points where $f_{xx}(p)$ is zero as well. Then we find catastrophe points where the first and second order structures are vanishing - the canonical fold catastrophe germ. But since we have no more knobs to turn, we cannot get rid of the higher order structure and therefore higher order catastrophes are non-generic.

In[65]:=
```
D[Series[f[x], {x, 0, 5}], x]^2
```

Out[65]=

$$f'[0]^2 + 2 f'[0] f''[0] x + \left(f''[0]^2 + f'[0] f^{(3)}[0]\right) x^2 + \left(f''[0] f^{(3)}[0] + \frac{1}{3} f'[0] f^{(4)}[0]\right) x^3 + \left(\frac{1}{4} f^{(3)}[0]^2 + \frac{1}{3} f''[0] f^{(4)}[0] + \frac{1}{12} f'[0] f^{(5)}[0]\right) x^4 + O[x]^5$$

The situation is somewhat different for the derivative signal squared. Again, we can use the two free parameters (x and the control parameter) to find points p where $f_x(p) = f_{x\,x}(p) = 0$. The remaining part of the derivative of the signal squared is then $\frac{1}{4} f_{xxx}(p)^2 x^4 + O(x^5)$. We see that the third order structure completely disappears and we therefore generically have the cusp catastrophe present in the derivative of the signal squared. It is also worth noticing that this implies that for each fold in the original signal, there is a cusp in the derivative of the signal squared.

We first demonstrate these findings in scale-space by looking at a simple signal composed of a few sines and cosines. First the definition of the signal and the derivative squared:

In[66]:=
```
from = -100; to = 400; resolution = 25;
f[t_] :=
  Sin[t/resolution] + Cos[1.2 t/resolution + 0.9] + Sin[1.3 t/resolution + 0.6];
GraphicsRow[Plot[#, {t, from, to}] & /@ {f[t], f'[t]^2}, ImageSize → 500]
```

Out[68]=

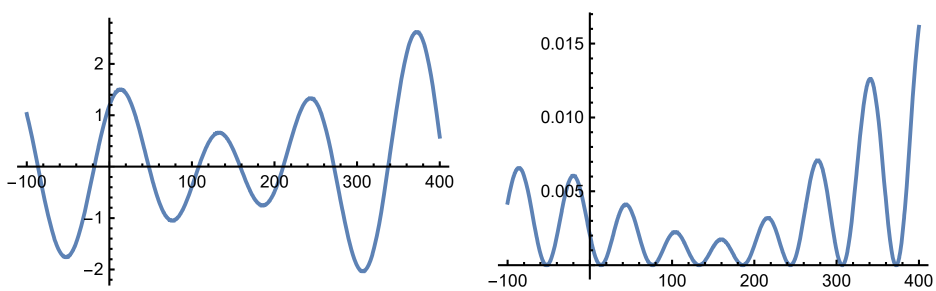

Figure 14.8 A simple signal and the corresponding derivative squared.

We investigate how this signal evolves in scale-space by convolving with a Gaussian using an exact Fourier domain implementation:

In[77]:=
```
gausskernel[x_, σ_] := 1/(σ Sqrt[2 π]) Exp[-x^2/(2 σ^2)];
signal[σ_, x_] = Simplify[
   InverseFourierTransform[FourierTransform[gausskernel[x, σ], x, ω] ×
     FourierTransform[f[x], x, ω], ω, x], σ > 0];
dsignalsquared[σ_, x_] = (∂x signal[σ, x])^2 // Chop;
GraphicsGrid[
 {Plot[signal[#, x], {x, from, to}, Ticks → {Automatic, None},
     PlotLabel → "σ = " <> ToString[#]] & /@ {1, 25, 40, 50} ,
  Plot[dsignalsquared[#, x], {x, from, to}, Ticks → {Automatic, None}] & /@
   {1, 25, 40, 50}} , ImageSize → 500]
```

Out[80]=

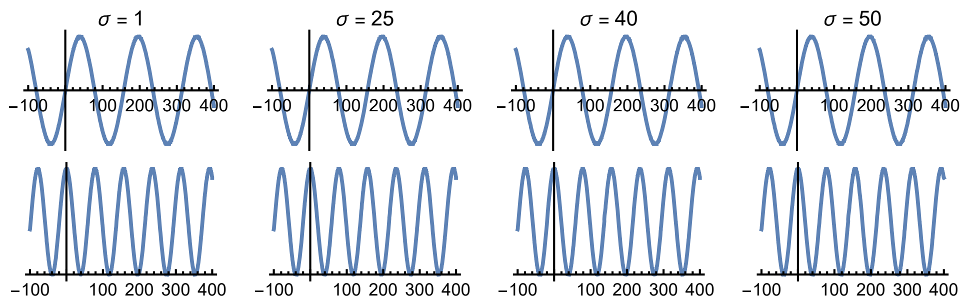

Figure 14.9 Top row: the original signal at scales $\sigma = 1, 25, 40, 50$. At the center of the signal a maximum and a minimum melt together as scale increases and are annihilated in a fold catastrophe. Bottom row: the derivative squared at the same scales. A {minimum, maximum, minimum}-triple is annihilated into a single minimum in a cusp catastrophe located where the fold is in the original signal.

In order to see that these nice derivations actually hold for real discrete images as well, we inspect a random signal. First the signal and the derivative squared is constructed.

In[132]:=
```
noise = RandomReal[{-1, 1}, {512}];
```

We then illustrate the evolution as scale increases:

In[133]:=

```
p[order_, σ_, n_] := ListPlot[DerivativeFilter[noise, {order}, σ]^n,
  Joined → True, PlotLabel → "σ = " <> ToString[σ]]
GraphicsGrid[
 {p[0, #, 1] & /@ {1, 25, 40, 50}, p[1, #, 2] & /@ {1, 20, 30, 40}}]
```

Out[134]=

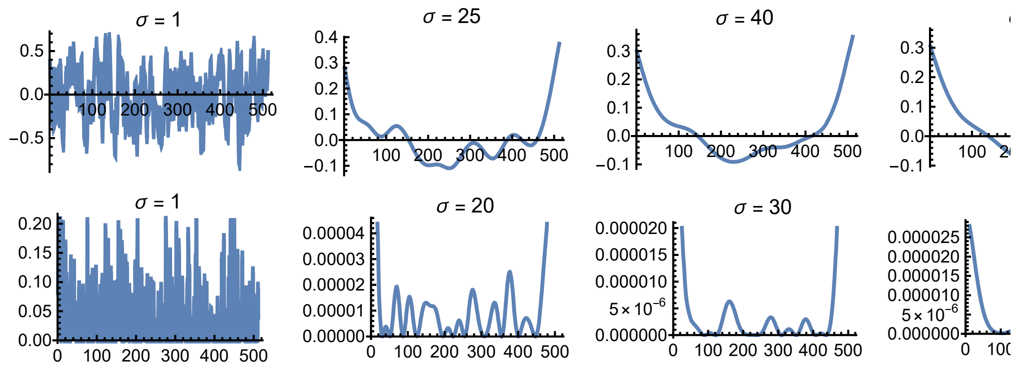

Figure 14.10 The scale-space evolution for a random signal and the derivative signal squared. The signals are displayed at scales $\sigma = 1, 25, 40, 50$ as in figure 14.9.

Finally we display the fingerprints for the random signal and the derivative of the signal squared. The fingerprints are slightly cluttered at low scale due to the inherent randomness of the signal but the fold/cusp pairs are obvious.

In[137]:=

```
signal[σ_] := DerivativeFilter[noise, {0}, σ]
fingerprint[signal_, maxscale_] := Module[{scsig, scLx, sig},
  scsig = Table[signal[E^(t Log[maxscale]/200)], {t, 50, 200}];
  scLx = scsig - Map[RotateLeft, scsig];
  sig = Reverse@Map[(RotateLeft[#] - #) &, Sign[scLx]];
  ArrayPlot[sig, ImageSize → 440]]
fingerprint[signal, 60]
```

Out[139]=

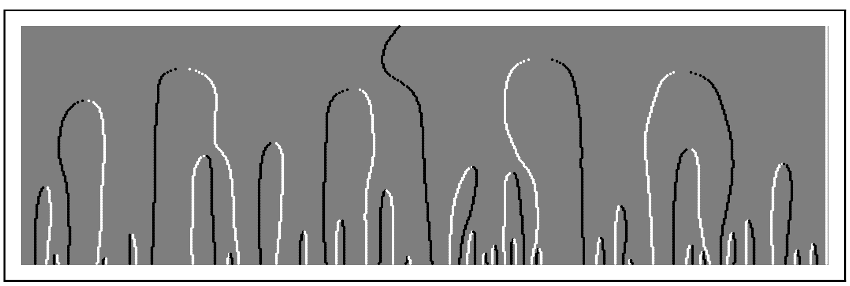

Figure 14.11 Fingerprints for the random signal and the derivative signal squared. For each fold catastrophe in the signal there is a corresponding cusp catastrophe in the derivative signal squared.

Thom's theorem states that the only generic catastrophe for a function with only one control parameter is the *fold*. However, as we have seen in the previous section, when we construct *differential* expressions from the original generic function we can induce higher order catastrophes as well.

A simple example is the derivative squared where the cusp catastrophe is generic.

In[140]:=
```
signal2[σ_] := DerivativeFilter[noise, {0}, σ]^2;
fingerprint[signal2, 60]
```
Out[141]=

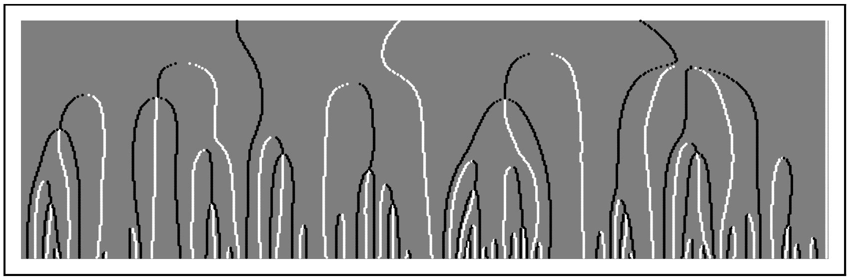

Other differential expressions can induce catastrophes of even higher order. Therefore, it is necessary to analyze each differential expression individually in order to reveal the generic events for the singularities as scale is increased.

14.4.2 Generic events for other differential operators

Other corners measures are studied in [Sporring1998a].

Among the investigated differential operators are the gradient magnitude used above as dissimilarity measure for the watershed segmentation. The generic events for this operator are the fold and the cusp - for both annihilations as well as creations are generic.

In[142]:=
```
Show[Import[url <> "nongeneric_isocat.jpg"], ImageSize → 300]
```
Out[142]=

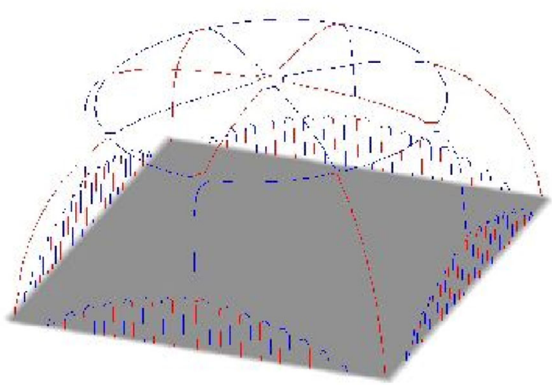

Figure 14.12 The singularities for the isophote curvature located on a specific isophote. The red/blue curves are the maxima/minima. These singularities can be perceived as corners. The

singularity strings are followed up through scale-space revealing non-generic catastrophes. Both the blue ring and the catastrophe at the top, where eight singularities are annihilated, display highly non-generic behavior. This is due to the non-generic test image - a perfect square. Illustration from [Dam1999].

This is derived in [Olsen1997]. Note, however, that creations only are generic in 2D and higher dimensions. Creations in 1D can never occur.

The generic events for the isophote curvature are studied in [Dam1999]. Again, the generic events are the annihilation and creation fold and cusp catastrophes. From this work we also have an ensemble of apparently non-generic catastrophes due to symmetry in the test images. An example is displayed in figure 14.12.

14.4.3 Annihilations and creations

In traditional catastrophe theory there is no preferred orientation for the control parameters. When the singularity structure for a function changes new singularities are as likely to appear as old ones are to disappear. Annihilations and creations are simply reverse events between different states for the local structure of a function. This is not the case for linear scale-space functions where scale is the control parameter. We only study the evolution of the functions for *increasing* scale. As mentioned earlier, increasing the scale results in a general simplification of the image functions. This implies that singularities are annihilated much more often that they are created.

As an example, see the fingerprint for the noisystep signal in figure 14.11. The figure shows fold annihilations - but no creations at all. This is in fact no coincidence. For a 1D signal (with no special properties or symmetries), creations are non-generic in scale-space. For images of dimension 2 (or higher) creations are generic. The famous example from the literature, where this was first discovered by Lifshitz and Pizer [Lifshitz1990], is the 'dumb-bell example' (see figure 14.13).

In[145]:=

```
db = Table[Chop[Exp[-(x - π)^2/2] Exp[-y^2/(2 (Sin[x] + .1)^2)]],
   {y, -4, 4, .2}, {x, 0, 2 π, .04}];
GraphicsRow[
ListPlot3D[#, ViewPoint → {1.489, -2.605, 1.968}, Mesh → False] & /@
  {db, gD[db, 0, 0, 3]}, ImageSize → 400]
```

Out[146]=

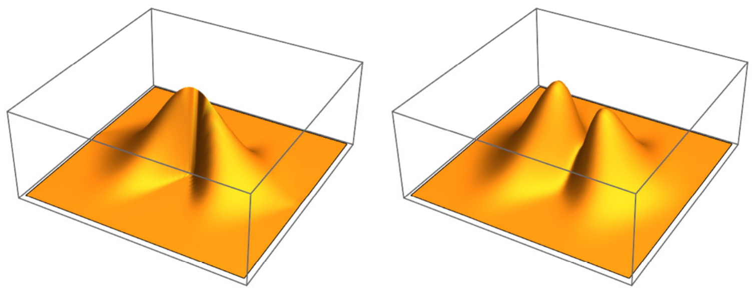

Figure 14.13 The dumb-bell image (left) consists of two blobs with a narrow bridge in between with a maximum. Blurring (right, $\sigma = 3$) has much more effect on the bridge, so two new maxima and a saddle are created under blurring. Illustration from [Lifshitz1990].

In general for catastrophes from images or differential operators on images, annihilations are much more common than creations [Kuijper2002a, Florack2000b, Kuiper2002b]. Furthermore, the scale-range where singularities can be observed resulting from a creation in scale-space is generally relatively short.

Another interesting aspect about singularities resulting from creations is their "validity". In a sense, the singularities that arise from creations have no originating structure in the original image. The feature corresponding to such a singularity can not be linked down to its "cause" in the image. Therefore, the singularities arising from creations are discarded in certain applications. An example of this is actually the linking scheme in the multi-scale watershed segmentation. Since the regions are linked from fine to coarse scale, regions resulting from creations never enter the linking tree. They are simply ignored on purpose.

14.5 Summary of this chapter

Catastrophe theory is the theoretical foundation that allows us to analyze the evolution of the singularities in scale-space. The theory allows prediction of the behavior of differential operators in scale-space. Ideally, any deep structure application should therefore include a comprehensive analysis of the generic behavior.

However, catastrophe theory is quite complicated. Therefore a strict theoretical approach has little appeal for the general image analysis community. Fortunately, a number of central results can be applied without more than a certain intuitive understanding of catastrophe theory. The purpose of this section was to provide a first step towards such a basic intuitive understanding.