About Shape

Jan Koenderink

DE CLOOTCRANS PRESS

About Shape

Jan Koenderink

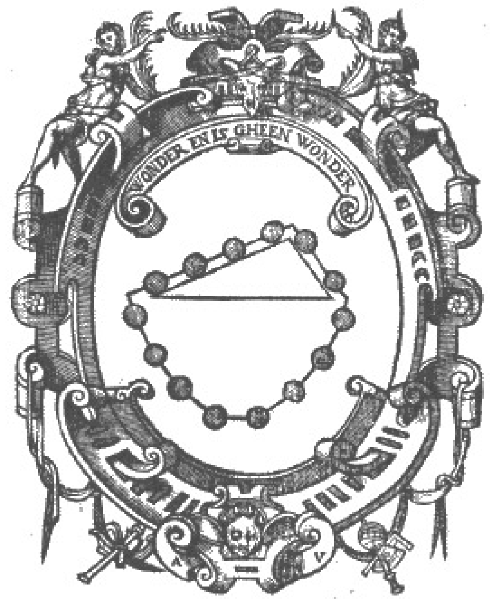

De Clootcrans Press, Traiectum mmxvi

Front cover: Hendrick Goltzius drew the famous Belvedere Torso around 1590. It is believed to be a copy from the 1st c. BCE or CE of an older statue, which probably dated to the early 2nd c. BCE. The Willendorff Venus was found in 1908, it dates from the "Old Stone Age" (28 000–25 000BCE). It is a truly *monumental* sculpture of only 11cm high. Both the Torso and the Venus represent the pinnacle of human shape conception.

De Clootcrans Press
Utrecht The Netherlands
KoenderinkJan@gmail.com

pax / jan koenderink

First edition, 2012

10 9 8 7 6 5 4 3 2 1

Foreword

I wrote this short text as a handout for the 2016 summer course "Visual Neuroscience" at the Rauischholzhausen Castle, organised by Jochen Braun, Wolfgang Einhäuser–Treyer and Karl Gegenfurtner, where I introduced students of Experimental Psychology to the topic of "shape", all in the space of three hours minus half an hour for coffee and biscuits.

Utrecht, july 21, 2016 — Jan Koenderink

Jan Koenderink
University of Leuven (KU Leuven)
Laboratory of Experimental Psychology
Tiensestraat 102 – bus 3711
3000 Leuven
Belgium
jan.koenderink@kuleuven.be

WHAT IS SHAPE?

This is a valid question. Most people feel they know the answer, although they find it hard to express it. This perhaps indicates that a formal understanding is lacking. Not that it matters much, except when you intend to do research on “shape perception”. Then you’d better know what you’re talking about, at least more than the participants in your experiments.

For formal definitions you consult a dictionary. Trying a few different ones is best. You find the usual circular definitions, that is how dictionaries work, they explain words by words. Some give also some “common uses” though, and these are of some interest. Here are a few:

The particular physical form or appearance of something:
— clay can be moulded into almost any SHAPE
— these bricks are all different SHAPES
— our table is oval in SHAPE
— my bicycle wheel is bent out of SHAPE

An arrangement that is formed by joining lines together in a particular way or by the line or lines around its outer edge:
— a round/square/circular/oblong SHAPE
— a triangle is a SHAPE with three sides

Let me start with what may be the simplest example, a planar triangle, that is a configuration of three distinct points in “general position”. A famous illustration is Sengai’s picture of the universe. This is what D.T. Suzuki has to say on it:[1]

> The circle-triangle-square is Sengai’s picture of the universe. The circle represents the infinite, and the infinite is at the basis of all beings. But the infinite in itself is formless. We humans endowed with senses and intellect demand tangible forms. Hence a triangle. The triangle is the beginning of all forms. Out of it first comes the square. A square is the triangle doubled. This doubling process goes on infinitely and we have the multitudinosity of things, which the Chinese philosopher calls “the ten thousand things”, that is, the universe.

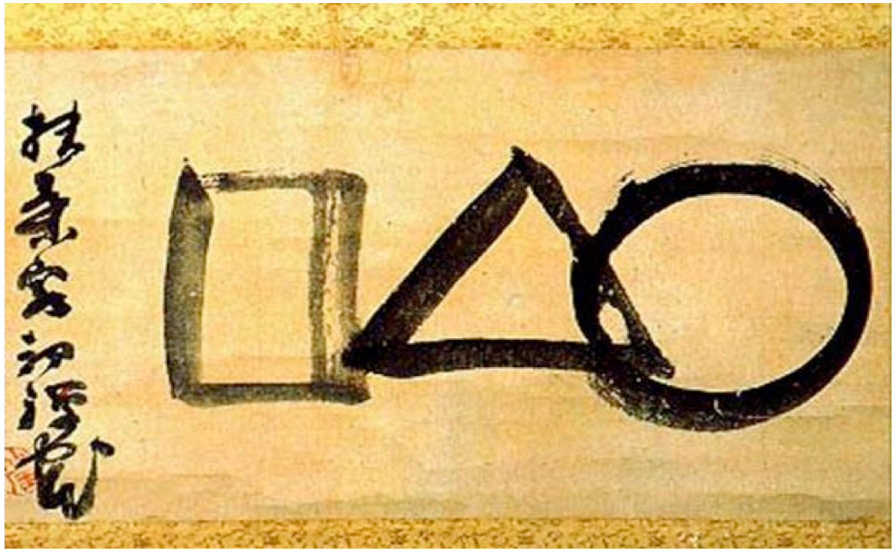

Sengai’s depiction of the Universe.

Suzuki adds that “a more traditional interpretation” is:

> . . . bodily existence is represented by a triangle which symbolizes the human body in its triple aspect, physical, oral, and mental. The quadrangle represents the objective world which is composed of the four great elements, earth, water, fire and air. . . . the ultimate reality, is the circle, that is, the formless form.

This serves to convey that the topic of SHAPE is important and perhaps mysterious. Then the formal definition — not from a dictionary, but from mathematics — may come as a bit bland

[1] Google is your friend in finding information on people you don’t know (like D.T. Suziki or Sengai).

SHAPE is an invariant under the group action of a group of "congruences" or "movements".

However, mathematicians can get really excited over definitions like this, so we should perhaps look a bit further. What does it imply? Well,
— a configuration space (think of "3 points in the plane");
— a group of "movements" (think of "shifts, rotations, scalings");
— an equivalence relation "congruency" that partitions the space of configurations into distinct SHAPES;
— SHAPES may be parameterized in a "SHAPE space".
What you need a shape space for? Well to *name* shapes (through coordinates of some kinds), to measure *differences* between shapes and so forth.

The definition is from the late 19^{th}c. due to people like Felix Klein, who should count among your heroes. Although a first rate mathematician (his students were afraid of him because he outsmarted them), he had a sharp feeling for and appreciation of *intuition*. The mathematical models you see displayed in the older math departments are due to him. Klein stated that the models were not illustrations of the math, but that the math was about the models. For him SHAPES where qualia in the sense of visual and haptic phenomenology.

Klein had no problem to conceive of the distance (in shape space!) between two triangles. Can you? Well, you can cheat! For instance, suppose I define the distance between two triangles as either *one* (when they are congruent, that is to say, the same triangles in different spatial locations and attitudes) or *zero* (when their symmetric difference (or disjunctive union) as planar sets is non-empty in any mutual placement).[2] Why do I say this is a cheat? Because whereas this distance is indeed a good *distance function*,[3] it is not going to help you much. It simply labels two things as same or different. You cannot use it to find a pair of different items that is twice as different as another pair of different items.

We evidently need more useful distance functions. It is evident that this is hard. How different is a horse from a goat? No idea? Is the distance from a horse to an octopus larger than the distance from a horse to a goat? I would say yes! Apparently there "is an answer", at least in some sense, but the problem as stated appears perhaps too hard.[4] It seems better to start with something simpler. As said, I will look at the "shape of triangles'.

You may protest that the shape of triangles is not particularly interesting. Why consider that at all? Well, you may underestimate the power of the triangle. Understand the triangle and you have conquered the (Eucldean) plane, for instance. Here are just a few simple examples:

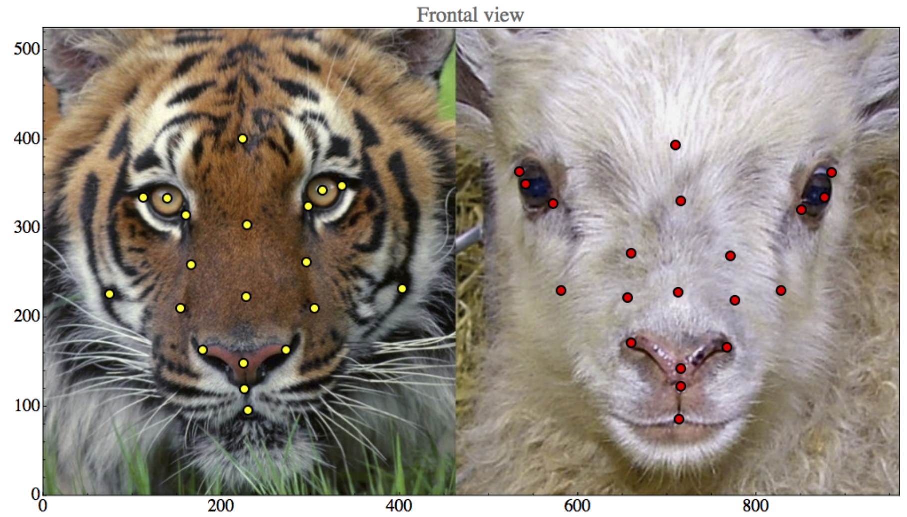

The tiger and the lamb are very different animals. For instance, predators have frontally placed eyes and fixate their prey, whereas their prey animals have laterally placed eyes and hardly need any fixations to see what they're in.

Although the tiger and the lamb have very distinct faces, both have eyes, a nose and so forth. In a formal experiment participants readily provide you

[2] Notice the technical terms "distance", "set", "symmetric difference", "disjunctive union", "planar set" and so forth. If you are not familiar with such a term look it up in Google. In most cases the first hit will do it, but you may have to scout a bit or pursue various trails. As a bonus you're likely to pick up some useful information on the side. Don't be ashamed, many people (including me) do the same.

[3] Look it up!

[4] Cladistics yields one possible answer.

with a dozen or so corresponding locations in the two portraits.

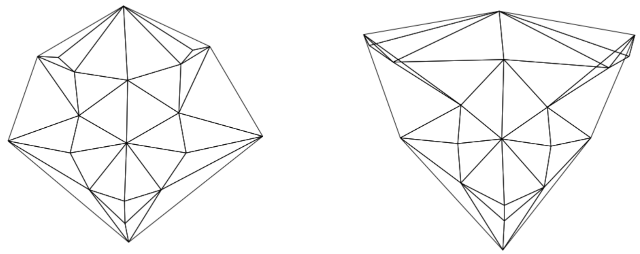

Corresponding triangulations for the tiger and the lamb portraits.

The tiger in the lamb-geometry and lamb in the tiger-geometry.

This is interesting, because you can make a few points go a long way by using *triangulation*. The corresponding points automatically yield corresponding triangulations. Now — using the triangulation — the correspondences set by the participants in the experiment automatically let you relate the full areas with each other.

Here you have plenty of interesting possibilities, for instance, you can both interpolate between and extrapolate either way from the triangulations, allowing you to plot sheepish lambs or unnaturally fierce super-tigers. Using triangles, you really conquered the image plane!

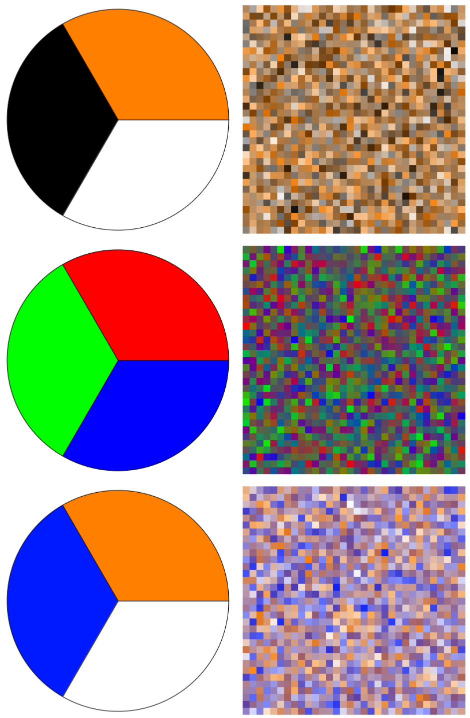

Some triangles in color space. The colors in the textures at right are in the convex hull (that means: are interpolations) of the three colors shown at left. That is what a triangle is, the convex hull of three points. This works out quite differently in different spaces, of course. I might as well have used a "face space" or whatever!

Here is another example of triangle use: Pick three colors and consider a texture filled with all colors obtainable by interpolation. Each such texture is equivalent to a triangle.0 If I have a distance for triangles I can immediately use it to define distances between such chromatic textures.

It would be easy enough (and fun!) to multiply the triangle examples, but I'll leave it at this. By now, you understand that triangles are hardly boring.

It would be appropriate to generate and show a few triangles here. But consider how to make them. Suppose you need to program a "random triangle generator" for your experiments, what would you do? Could you be formally sure all triangles would be "equally likely"?[5] Here is an idea: I generate triples of normally distributed Cartesian coordinates with the same variance. The convex hull of each triple is a "random triangle" for sure!

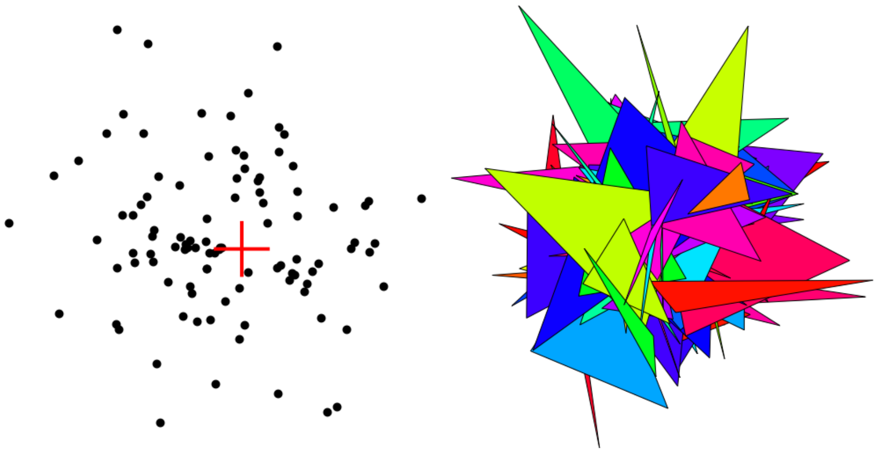

Random triangles generated as the convex hulls of triples of points drawn from a normal distribution.

Indeed, this looks promising. But is it a unique method? No! For instance, I could draw triples from a uniform distribution in a square (an area), or from a circle (a curve!). Are at least the results the same? No! So much is evident from a cursory look. In order to make progress we need formal methods.

[5]Here you should ask "in what sense?", for they trivially *are* in the sense of your algorithm, whatever the algorithm might be like.

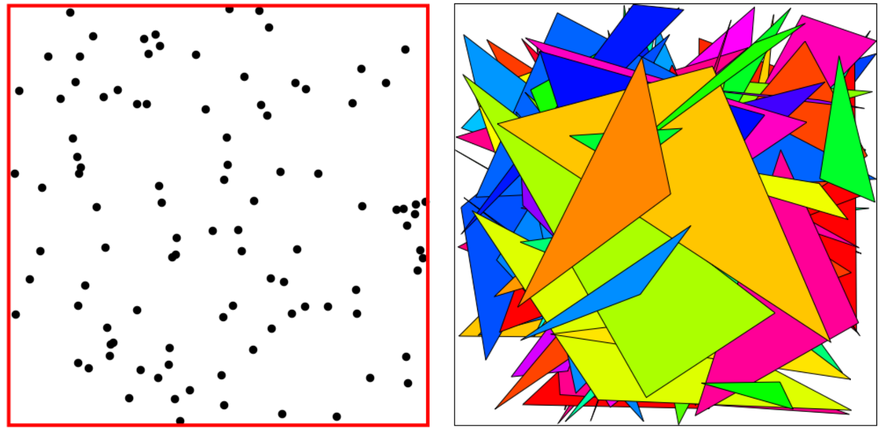

Random triangles generated as the convex hulls of triples of points drawn from a uniform distribution in a square.

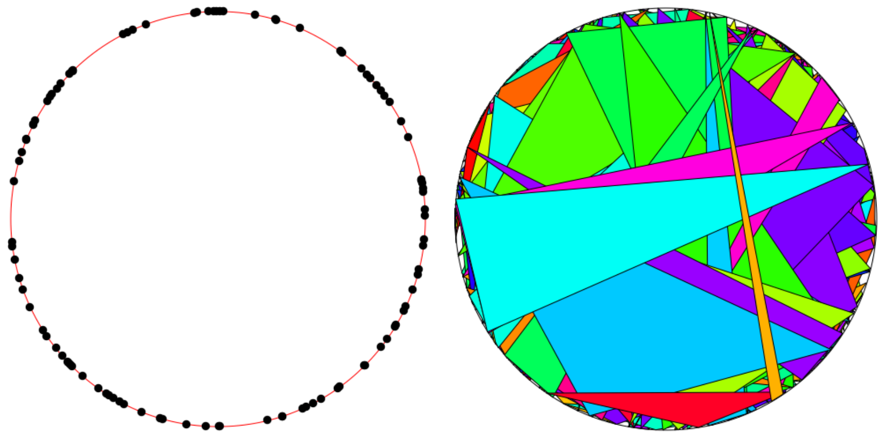

Random triangles generated as the convex hulls of triples of points drawn from a uniform distribution on a circle.

The notion we need at this point is that of the "Procrustes" method. Procrustes was a rather unpleasant person who made sure his guests perfectly fitted the bed he generously offered them by either chopping off the surplus or stretching them forcefully to the right size. The Procrustes method does something similar, hence its name.

Procrustes at his favorite passtime.

For two sets of corresponding point configurations (as in the tiger-lamb experiment), I define a *measure of fit* as the sum of all squared distances of corresponding point pairs. This is a useful number, for it is always positive and only zero when the two point configurations are fully coincident, that is simply *the same*. Useful as it might be, it is hardly a *shape measure* though, for I can increase the number arbitrarily by moving the two configurations away from each other. Here is where Procrustes comes to the rescue.

Although we can't chop pieces off, since that would change a configuration, we may certainly move and rotate them and — perhaps — stretch them. Lets start with the movements. When I let the centers of gravity of the two configurations coincide, I have minimized the measure of fit with respect to translations.[6] Then I simply try all orientations to find a (hopefully unique) minimum. It is the best I can do if I'm not allowed a scaling.

In most cases one would welcome such a scaling, since it appears to fit the common sense notion of SHAPE. So we go Procrustean and use the stretching method.[7] The best scaling can be easily found after factoring out the translation by equating the r.m.s.distance of all points from the center of inertia for the two configurations. After scaling, finding the optimal rotation is perhaps best done using the singular values decomposition method.[8] According to the methods used, finding the best superposition may involve some iterations. The square root of smallest measure of fit then is unique for the pair of point configurations. It is their mutual "Procrustes distance".

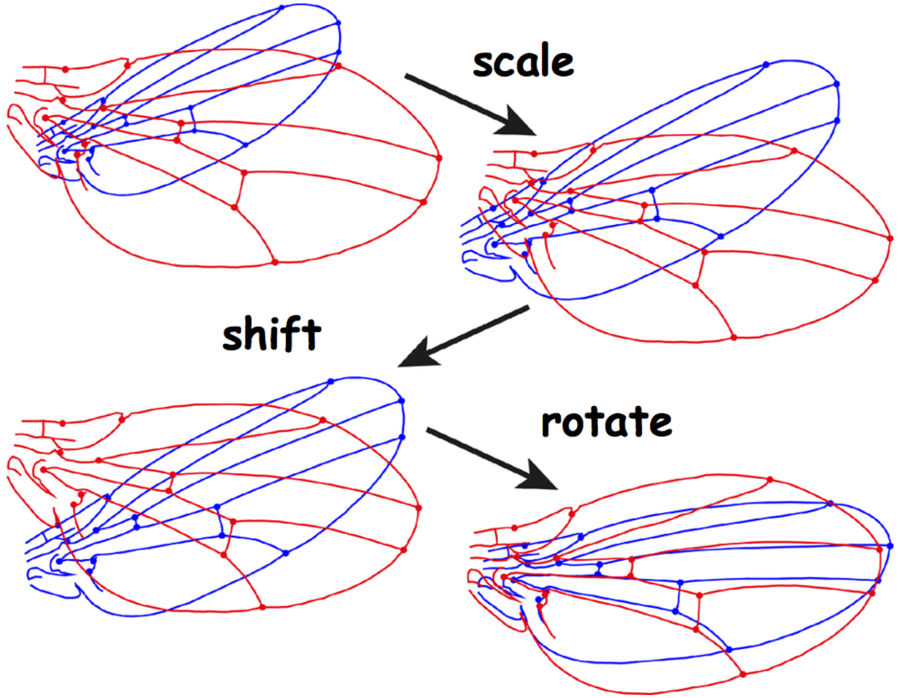

The Procrustes method applied to two insect wings. The method can be applied whenever you have sets of corresponding points. Notice the sequence of operations. Various sequence will do fine.

[6]You may want to prove that to your own satisfaction, it involves only simple algebra.

[7]The chopping method (perhaps preferred by Procrustes the person) is not needed because we can both upscale and downscale.

[8]Notice that the sequence of operations is in itself not too important, but that the particular operations will depend upon the sequence. Not important in principle, but relevant in practice.

Triangles are the simplest objects to which the Procrustes method applies in a non-trivial sense.

There is one additional degree of freedom to consider, namely that of reflection. Two triangles are often considered congruent when they can be brought into full superposition after one of them has been reflected. It is a mere matter of choice, I remark on it a few times below.

I show the analysis of David George Kendall (who should be another one of your heroes), who figured out the structure of the triangle shape space in the mid nineteeneighties. The figure shows what happens graphically, but — of course — Kendall has an algebraic parallel storyline.

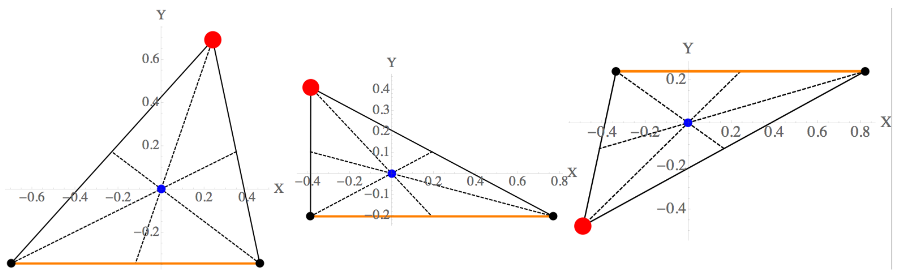

This is Kendall's trick to map triangles to a two-dimensional parameter plane. Move the triangle so its centroid coincides with the origin of the parameter plane, orient the base to be parallel to the x–axis of the parameter place, scale the base-length to $2/\sqrt{3}$ in terms of the parameter plane unit length. Then the top-vertex is the location of the triangle in the parameter plane. (Blue point the origin, orange sides the bases, red points the representations.)

In the example I generated a hundred random triangles using the normal distribution triples and send them through Kendall's Procrustean pipeline, ending up with points in a "triangle plane". Here each point uniquely represents a triangle. Why is this plane useful? This becomes evident in the next step proposed by Kendall.

The parameter plane is nice enough, but—as the algebraic calculations show—the metric of the parameter plane is not the Procrustes metric. If possible, one would like to remap the parameter plane in such a way that the new representation would reflect the Procrustes method. Kendal figured out how to do this, here is how:

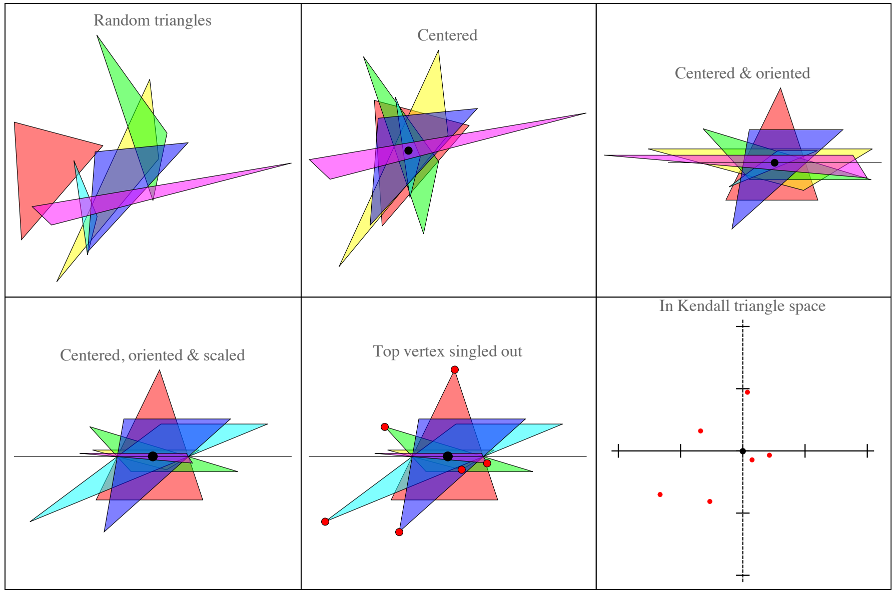

Here I start with six triangles (you can just as easily do it to a million, but it would render the figure hard to read), move them, rotate them and scale them, eventually finding their representation in Kendall's parameter plane.

Define angles $\vartheta = 2\arctan\sqrt{u^2+v^2}$ and $\varphi = \arctan(u, v)$ (this is the "atan2" from C), where $\{u, v\}$ are the Cartesian coordinates of the triangle plane (remember how the axes were defined!) and interpret ϑ as the polar distance (or co-latitude), φ the longitude of the unit sphere.[9] This maps the triangle plane on the unit sphere, which Kendall refers to as the "spherical blackboard" for the occasion. Then Kendall's algebra proves that the spherical distance between the representations of two triangles on the spherical blackboard is their Procrustes distance. It takes a certain "nose" in order to

[9] In case you understand such things: Kendall treats the triangle parameter plane as the stereographical projection of a parameter sphere. Why? Well, as it turns out it makes the sphere metric equivalent to the shape metric.

be able to find such a relation. One needs to grok both the algebra and the geometry, moreover, have a certain empathic feeling for them. That's why I said Kendall should count as one of your heroes. For the rest of us Kendall's spherical blackboard is a wonderful tool. If you ever do experiments involving triangles (remember that you may actually do that without realizing the fact!), you certainly should use the spherical blackboard.

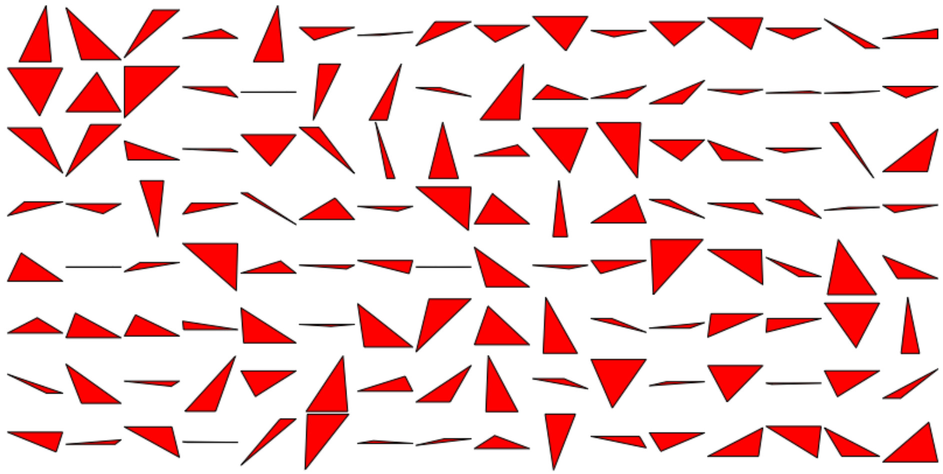

Here is a sample from a uniform distribution on Kendall's spherical blackboard.

Now we can finally draw triangles from a uniform distribution. Of course, defined on the spherical blackboard! Kendall used his method to find the probability densities for the circle and square methods I mentioned above. Not a trivial matter, what would you do? (In this case you can find answers on the Internet.)

Is this the final answer then? No! There are various alternative paths. For instance, you might use another definition of "shape" or use a different geometry than the Euclidean plane.

Here is an idea of a different shape measure. We simply characterize a triangle by its angles. This is a great idea, because it obviates the need for Procrustean methods altogether! As we all know the sum of internal angles in a Euclidean triangle adds up to $180°$. Thus the "shape" may be characterized by two numbers. In practice one uses the three angles and plots this in "barycentric" coordinates. Is this equivalent to the Kendall method? This is easily checked by plotting triangles drawn from a uniform Kendall distribution in the barycentric coordinates. The result is clear: These representations are *different*! So which one is to be preferred? That depends upon your application.

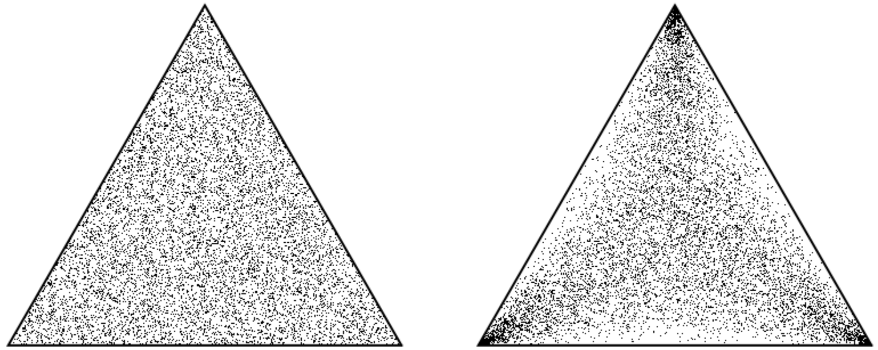

At left the barycentric representation for random triangles generated by drawing three random positions on a circle. At right a similar representation for triangles drawn from a uniform distribution on Kendall's spherical blackboard. In both cases I sampled ten-thousand triangles. This amply suffices to indicate that the shape spaces are different.

The barycentric representation is nice, but it is a nuisance that it has a boundary! This is easily fixed though, because you may glue four copies of the triangle together in such a way that you obtain the topology of the sphere. Since spherical surfaces have no boundary, the problem is solved. The solution comes at a (minor) cost though, because each triangle is now represented through a four-tuple of points.

There is another worry, namely that you may want to consider two triangles to be of the same shape when their angles are the same *irrespective of order*. Thus the triangles ABC, ACB, BAC, BCA, CAB, and CBA would be considered "the same".

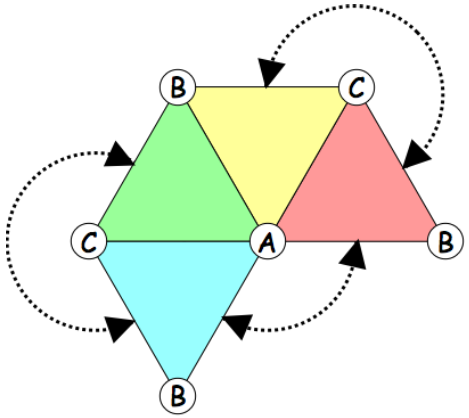

By identifying pairs of edges as indicated I construct a closed surface. It has the topology of the sphere. This has the advantage that there is no boundary, thus a deforming triangle will never run into an obstacle along its path. Each triangle is represented fourfold.

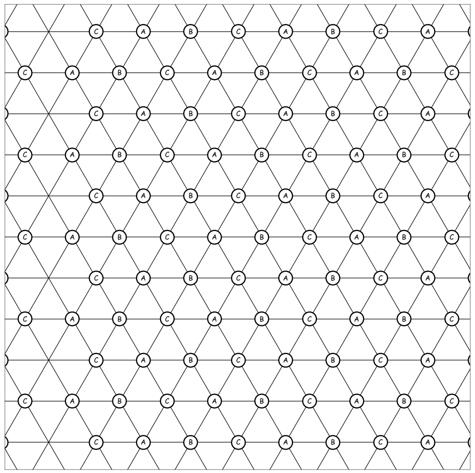

An infinite tiling of the plane works just as well and would be preferred by your formalist brother.

In the topological representation each triangle is now represented as a set of twenty-four points! No big deal though, one simply adjusts the formal definition of what a“point” is. Of course, you need to take care when measuring distances. Simply taking the shortest distance over all contenders does the job.

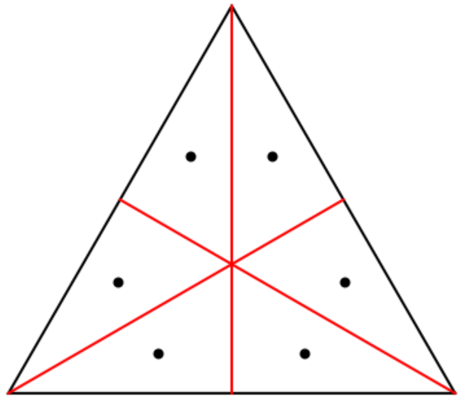

When triangles ABC, ACB, BAC, BCA, CAB, and CBA are considered “the same”, each triangle is represented by six point in the barycentric representation.

The Kendall analysis is specific for the Euclidean plane. What if you leave this trusty territory? Well, when you change the geometry you change the game completely! This is very important, so I’ll give some examples. One of the simplest examples is to consider the shape of triangles in the *affine plane*. In the affine plane two configurations are reckoned “congruent”when they are equivalent under general linear transformations.

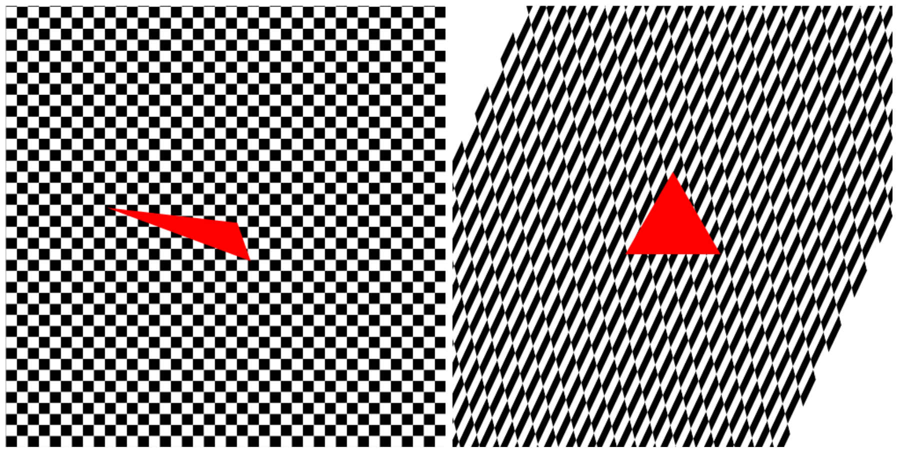

Any triangle can be mapped on any other triangle by some affine transformation. Thus all triangles in the affine plane are congruent!

The figure shows the effect of an *affine transformation*. “Affine transformations” preserve lines and parallelity. They are very common in applications, perhaps even more so than Euclidean movements. The interesting observation is that there exists a unique affine map for any pair of triangles that will map the one to the other. Thus, doing affine Procrustes, any pair of triangles has (affine) Procrustes distance zero. All triangles are affinely identical. *There is only one affine triangle.*

Quadrangles

It is simple enough to generate random quadrangles, although most people tend to be surprised in seeing their initial results.

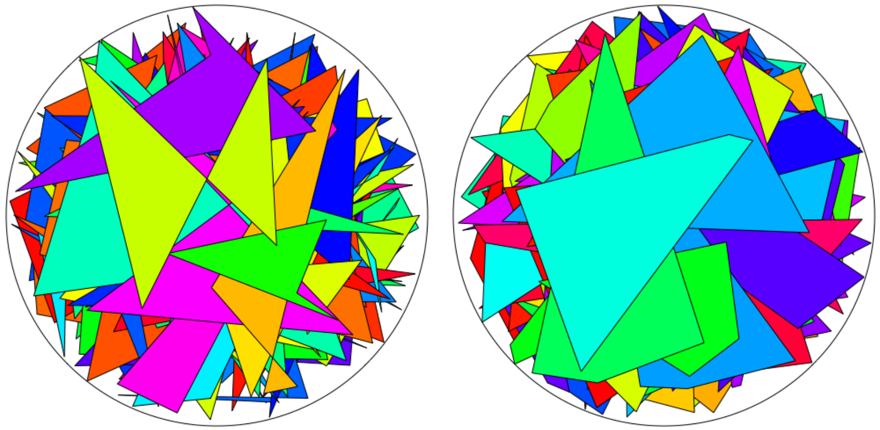

At left random quadrangles generated as sequences of random points. Notice that edges may intersect and that many instances are not convex. At right quadrangles are generated as convex hulls of random tetrads of points. This tends to make better “common sense”.

Although there is nothing wrong with quadrangles that are not convex or even have intersecting sides, this is not what our common sense expects. Quadrangles are perhaps best understood as convex hulls of four points. This tends to make sense to most people.

Although there is only one affine triangle, there are infinitely many affine quadrangles. What is their shape space like? One simple way to obtain an idea is to draw a diagonal, yielding two abutting triangles, and affinely force one triangle to be equilateral. Then the other triangle uniquely parameterizes the quadrilateral. This implicates that the affine quadrangle space is just the Euclidean triangle space.

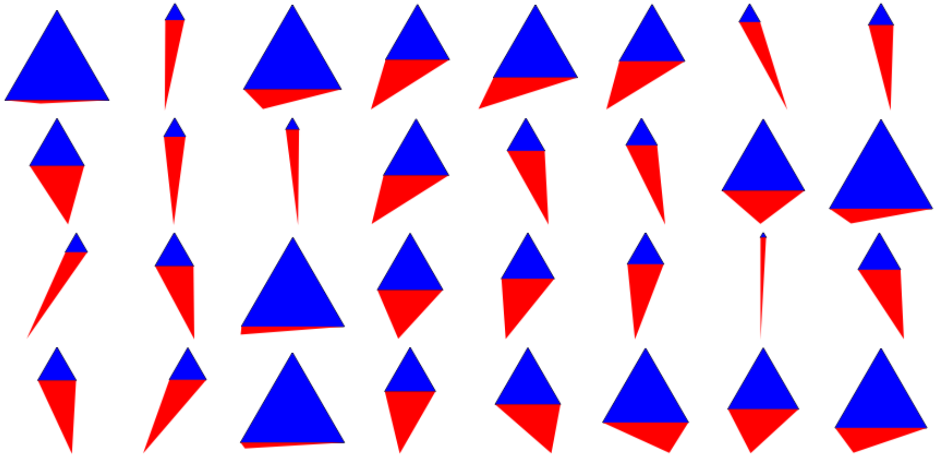

Here a number of random quadrilaterals have been transformed such that one triangle above a diagonal is equilateral. Notice that the quadrilaterals still differ by the other triangle. Thus two quadrilaterals are typically affinely different, their difference being the difference of two Euclidean triangles.

So quadrangles are no big deal.

Of course, we can set another step and allow *projective transformations*. Projective transformations preserve straight lines and their relations. They are very important in applications in (especially computer–)vision. It is a basic fact that any two quadrilaterals are projectively equivalent. Thus there is only one quadrangle in the projective plane. Typically, you wouldn't even recognize it as a quadrangle. However, it can always be represented such that it would look to you like a square.

Most people are familiar with projective geometry as painter's perspective. This is only one application though and one that does not use all the machinery of projective geometry. Nowadays people have apps on their smartphones that do "perspective corrections" for their photograps, If you photograph a building looking upwards, its facade comes out "distorted". The apps do a Procrustes transformation of the projective variety that makes them look right again. They simply change the general convex quadrilateral of the facade (extracted by way of edge detection) into a Euclidian rectangle, usually guessing at the aspect ratio. It makes "the shape come out right" and users are happy.

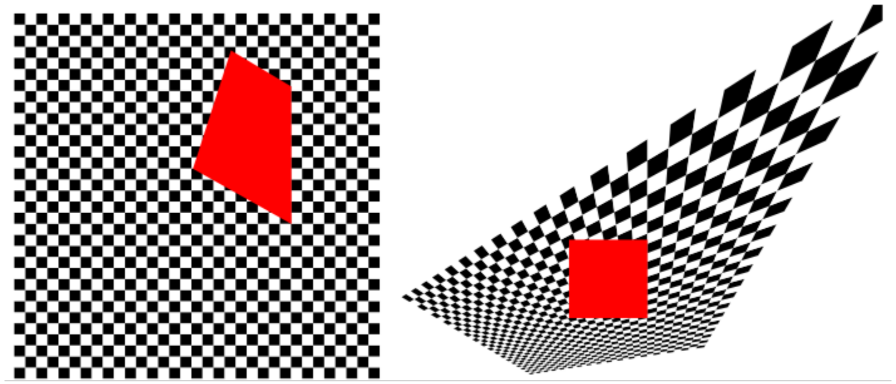

Any two quadrilaterals are projectively equivalent.

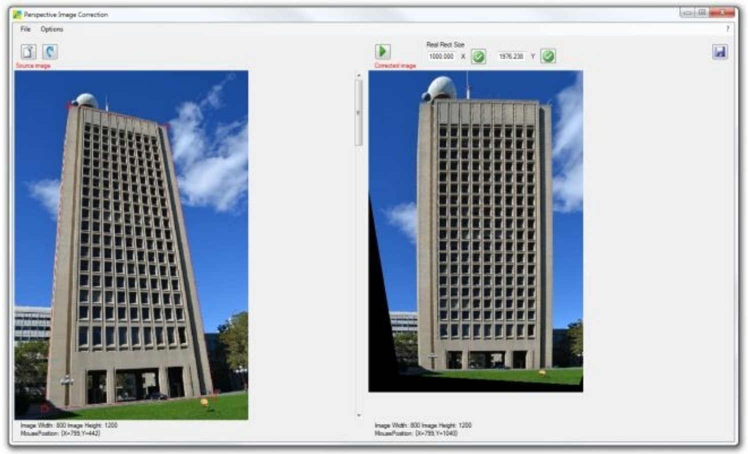

A common application: "perspective image correction". This is easy to use for architectural images, where there are plenty of quadrangles around. If you don't have the necessary calibration data the aspect ratio may be off though. You really get an affine distortion until you get your horizontal/vertical size ratios right. If you happened to know a square in the image that would be easy, of course.

In his Vergleichende Betrachtungen über neuere geometrische Forschungen *Felix Klein explains all we need (and more!) in our present context. Even a superficial understanding of his program will rip you loose from the Euclidean straitjacket you're probably in.*

I could only show the tip of the iceberg here. The point to remember is that selecting the right geometry for your research is a crucial issue. Euclidean geometry is not the holy grail. Physicists have come to understand that, but the other sciences run behind. Don't forget that a change of geometry is a dramatic, qualitative change. Using an inappropriate geometry is likely to require endless fussy patchwork that would be cleared up in one fell swoop when the right choice were made. Don't underestimate this!

Application of affine shape

Here is an elementary application of affine triangular shape. How elementary can it be? There is only *one* affine triangle! The problem is that our common sense eyes often don't see that. The application is a simple case of optical flow that looks complicated to the Euclidean eye, but simple to the affine eye. The upshot is that you need to open your affine eye.[10]

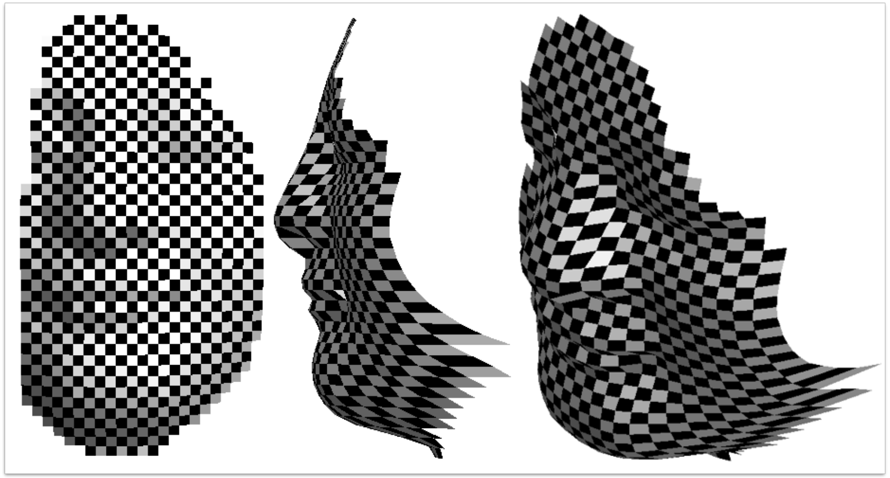

This is a Greek stone sculpture that my friend John Willats (died some years ago, great artist) measured up at the British Museum before the time of 3D scanners. I use it as a test object here. It is convenient because most people know what faces look like.

[10]In case you remember the Flower Power period, it is more important than your Third Eye. Hippies who tried to actually open that generally came to a bad end.

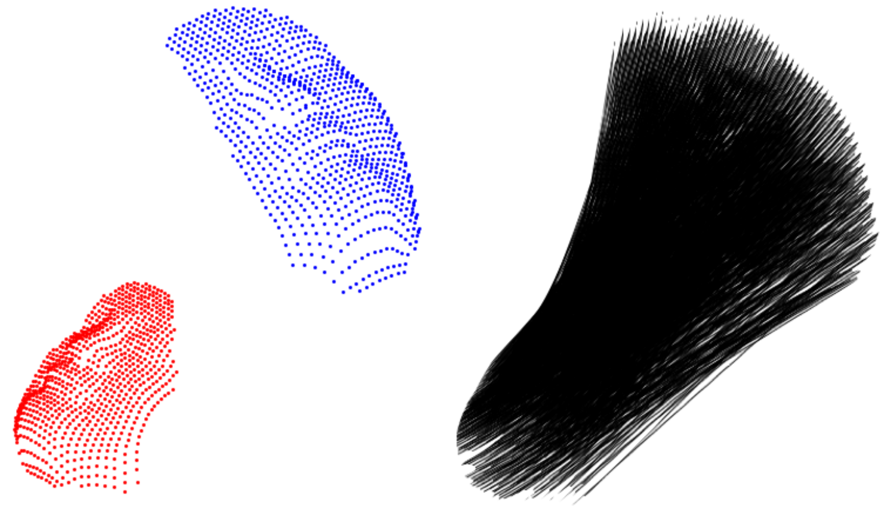

Here are two views of the face from different directions and distances. Can you get the true 3D shape of the face from the flow? Hardly.

Here I consider the "optical flow" (or "parallax") obtained by comparing two views of a single rigid object. I assume that the point-to-point correspondences are known. The flow is obviously somehow modulated by the shape, but there seems to be no obvious way to obtain the shape from the flow. People tried hard before the 1980's, but with no success. The consensus was that you can get shape from *three views*,[11] but not from two. It was supposed a theorem from formal geometry.

Indeed it was, but from Euclidean geometry. Applications oriented people had not yet opened their affine eyes.[12]

[11]Moreover, the algorithms propose for doing that were really UGLY. Sure, you could follow the algebra, but did it tell you anything? No! I'm sure even the heroic authors had no clue, except that it worked.

[12]Computer vision people got wise, they use both affine and projective geometry, big time. Most in experimental psychology have no clue yet, they use the word "affine" often enough, but usually as a sophisticated term for something neither understood nor used. "Projective" is not on the agenda yet.

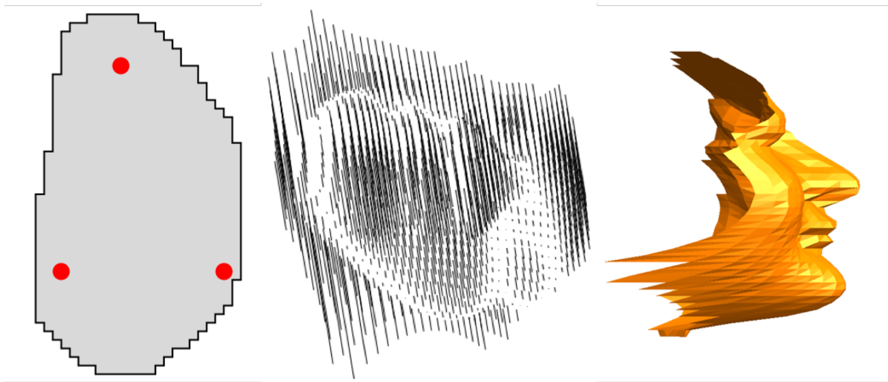

This is the "affine flow". I get it by mapping the triples of indicated points to equal equilateral triangles in the two views, also centering the results so as to annull the average displacement. This is the affine Procrustean trick. Notice that all affine flow vectors are mutually parallel, only their magnitudes count. The magnitudes are proportional to their deviation from the plane of the triangle. This allows me to draw the profile at right. Not a bad reconstruction of the original! All that is lacking is a depth calibration.

All you have to do to see what optical flow captures is to open your affine eye. Perhaps not unexpectedly, it has to do with triangles. Here is what you do: select three points in one view and find their corresponding positions in the other. Now you're ready to do affine Procrustes. You simply map both triangles on the same equilateral triangle. Then the "affine optical flow" becomes zero for them — by construction. But see what happens to the remaining flow vectors, they are all parallel to each other, so you can forget their direction. The direction has nothing to do with the shape, but only with the difference in eye positions for the two views. Their magnitudes are proportional to the distances of points from the plane defined by the three fiducial points. Thus the affine eye sees the 3D shape in the flow without any computation! You can easily compute any possible view from this. Here I have computed a profile view that may be compared to the original. All that is lacking is a depth calibration, the "nose is too long". The important point is that you immediately get the 3D shape modulo a simple affinity involving depth.

This makes for a very nice example of the use of triangles in surprising settings. In this case the problem was solved trivially by moving from a Euclidean to an affine perspective. The solution depended on the fact that three points define a unique plane in *three-dimensional* Euclidean space, whereas projections (of course, planar) triangles in that space yield triangles in the *two-dimensional* Euclidean plane of projection. Since any two triangles are affinely equivalent, such triangles can be mapped upon each other. However, any point not coplanar with the triangle in *three-dimensional* Euclidean space cannot join in the latter transformation. Thus its deviation from the plane becomes manifest after the affine map.

In dealing with images triangles are natural because they cover part of the image plane. A triangulation tiles the image plane with triangular tiles. This is very useful because each tile can be taken as small as desired, with the implication that non-linear deformations can be conveniently linearized. This makes it possible to do complicated transformations by local texture mapping, much cheaper than pixel-based methods. Similar methods apply to higher dimensions. For instance, in three-dimensional space you would "tile" with tetrahedra, and so forth. The space need not be "just space", but could be parameter spaces of arbitrary kind.[13] Such methods generate a need for "good tilings", which generally means tilings having triangles that are close to equilateral. Thus the topic of "the shape of triangles" is bound to pop up in many—often surprising—contexts.

[13] For instance, there may be occasions where it would be advantageous to triangulate Kendall's spherical blackboard itself.

Shape in the Visual Arts

Who knows best what SHAPE *is? Perhaps we should ask visual artists first. For them* SHAPE *is not a formal definition, it is their meal ticket. Once you start to look at the shape-related methods of painters and sculptors you will soon feel to be ready to give up! There are really too many aspects of the matter than that you might hope to formalize. Our formal methods will never be able to reveal all but a few tips of the iceberg and even that with considerable loss. Formalizing intuitive insights necessarily involves hardcore Procrustean methods. I just touch on a few topics, then zoom in on a useful, fairly general, formalism in the next chapter.*

Consider the process of drawing solid shapes for a start. One of its roots is the facility to doodle. Visual awareness turns doodles into solid shapes.

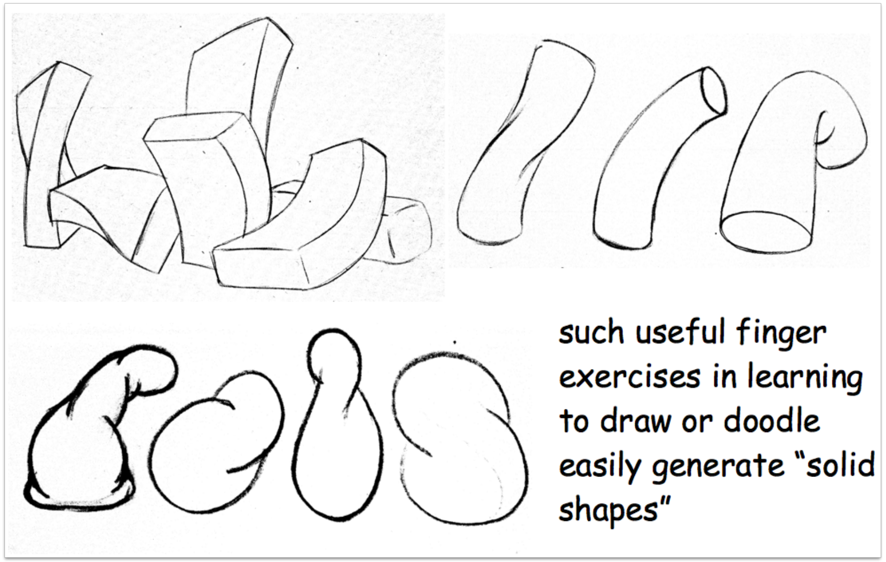

Becoming an expert doodler involves endless finger exercises like these.

There are many ways to doodle, the illustration shows just three common styles. Each work great. It lets you conjure up about any solid shape you want. By combination you easily build more complicated structures. That's really all there is to most effective drawing. Facility in doodling involves endless finger exercises, a bit like keyboard training. Most artists enjoy it. You can bend, twist, bulge and taper such shapes even easier than when you do it in clay. They magically become 3D in visual awareness.

It hardly matters which style you adopt. Most artists stick with a first preference. It does not matter, because it is easy enough to turn such doodles into anything you want. That is called "developing the drawing", which is another facility one needs to acquire. It is of a different kind, you need to dress up the bare doodles with local surface relief.

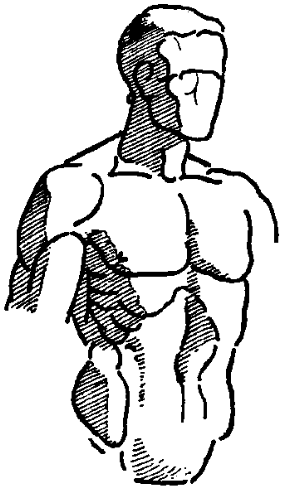

The "cuirasse esthétique" lets you draw the male torso. It has been misused in Western art from the beginnings. It is based on the Roman cuirasses (which were again based on Greek examples like the famous Doryphoros by Polykleitos (ca. 440 BCE*) that should be familiar from Hollywood gladiator movies and comic books.*

The "dressing up" can be applied to many other starting methods, all you need is a good start. The doodles are great for that, because so general. But many people prefer less general but more specific methods.[14] A well known example is the *cuirasse esthétique*. The traditional artist knows it by heart. It allows you to draw a male torso with your eyes closed, it is much less fun than doodling. It is very effective though. The example I show is really a rectangular box (notice the shading!) with a minimum of muscular dressing up. When fully dressed up the artist proudly displays an expert understanding of the male human form.

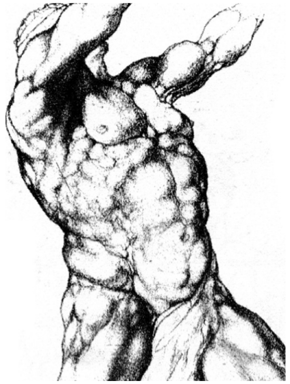

Here is the "cuirasse esthétique" fully dressed up with muscular detail. This drawing is by William Rimmer a mid-nineteenth century family doctor in Massachusetts.

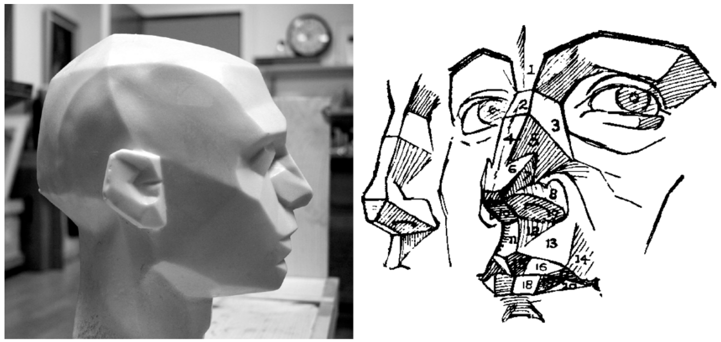

be analytic, then make up for it!

The method of polyhedral approximation. At bottom "ovoid drawing" used in dressing up.

[14] Hence the endles guides on "How to draw cats" and so forth.

Many artists abhor of smooth surfaces because they don't know what to draw within an outline and if you draw only outline you end up with a flat silhouette. They prefer polyhedral approximations and root-learn "the planes of the head" and so forth. This surely works, especially if your dressing up removes or hides your scaffold. The sharp edges will be replaced with soft shading, or "ovoid drawing" is used to add a sprinkle of muscular detail.

This is a "one minute pose" by Gretchen Kelly. It captures the posture and barely suggests volume.

So far the methods were based on volumetric intuitions, with the cuirasse esthéstique perhaps being more like a template that can also be used in an essentially "flat" manner (although that will show). But, of course, it is just as well possible to start with axes, although these are usually "felt" as moving in space, although drawn on paper. The gestural drawing by Gretchen Kelly is a nice example. It perfectly captures the pose of the model in a few brush strokes. Such a method can also be dressed up, although the result is likely to differ from that obtained by other methods. Usually you can see (or feel) how a drawn or painted shape originated.

A somewhat degenerated form of gestural drawing is based on the idea of "stick figures". This usually leads to drawing that indeed look like stick figures, which seriously offends the eye, but in the right hands it is certainly a valid starting method. It is often preferred by amateur artists.

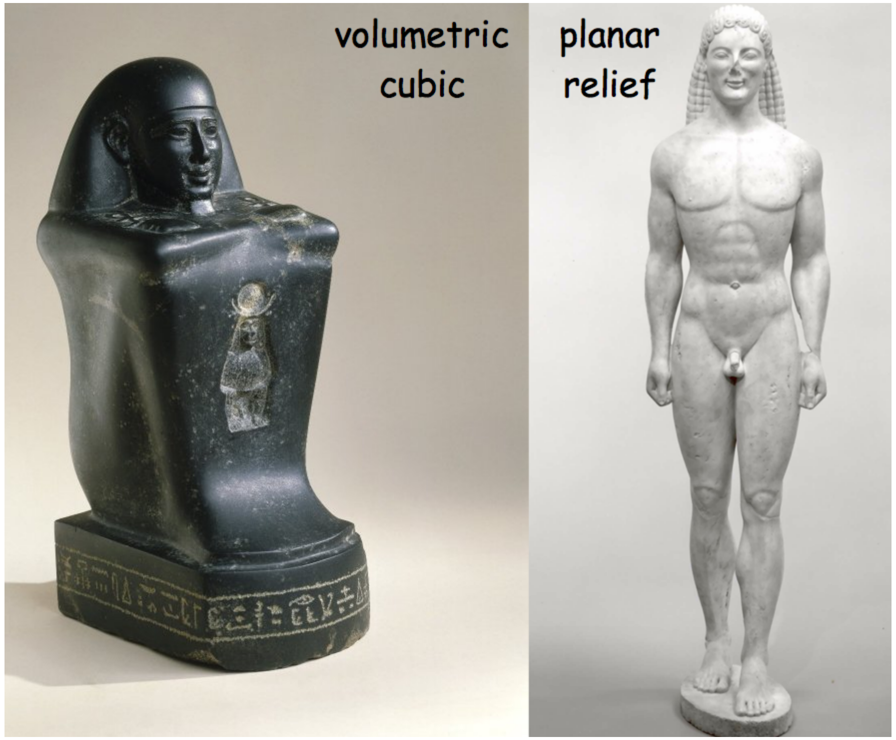

The Egyptian sculpture strikes you as volumetric, whereas the Greek kouros is frontally conceived and almost treated like a relief.

One way to put some order on the chaos is to classify artistic shape conceptions as primarily linear, planar, or volumetric. The Egyptian sculpture first strikes you as a rounded cube, it is clearly conceived as a volumetric block. The Greek kouros is evidently frontally conceived, it might as well have been a relief. It looks like an articulated plane. The gestural drawing by Gretchen Kelly is evidently linearly conceived. This can equally well be done in sculpture, there are numerous examples in contemporary art.

For each of these general families of shapes one might develop formalized treatments. I had to make a choice for this course because of time limits. I decided on articulated planes. It is an apt choice because it leads to a formalism of relative simplicity, but an enormous range of applications. It is a bit like the case of triangles, you might be surprised at the range of — perhaps unexpected — applications.

The Shape of Surfaces

Why the shape of surfaces? Well, because an obvious starting point would be the local shape of surfaces and this leads to a simple but powerful formalism with numerous applications. Because it might not be obvious that such is the case, I start the chapter with an overview of the implications.

In order to know what surfaces are you need to know what space is. So what is space? But what of a question is that? So I will dodge the question and ask "do you know what the surface of the earth is (I mean shape-wise)?" Answers are likely to range from geomorphological descriptions to spherical, oblate ellipsoid and so forth. It apparently depends on the application. There are hills and dales in the landscape, grains of sand and leaves of grass near to you, whereas the globe looks an almost perfect sphere as seen from the moon. Scale is the defining parameter. When I talk of surfaces in this chapter I mean objects like the sphere or perhaps the hills and dales.

Things get a little easier when I regard a "local" environment. This is tricky, because it introduces a second scale parameter. Here is an example:

One used to think that the earth was *flat* like a plane in Euclidian space. Yet we know that hills and dales are seen from many places where we stand and that the earth looks like a sphere as seen from the moon. We might consider a scope large enough to consider the hills and dales as mere local articulations, yet a scope that is too limited to see the sphericity.

Such decisions are made by our visual awareness on the fly and according to present situational awareness. If we say someone has a perfect, smooth skin we do neither imply wrinkles, nor that the skin is extended like a plane. We accept both bulges and pores as not violating the notion of a "smooth skin", albeit for different reasons.

Suppose we accept the flat earth notion for a moment. It probably fits our perceptions here and now fine anyway. Suppose we describe the shapes of hills and dales. The first thing to notice is that the height dimension is qualitatively different from the dimensions of the ground plane. It is measured differently (counting steps works only in the plane) and often expressed in different physical units. For instance, before GPS or radar it was common to measure height in terms of barometric pressure. Thus the "space" we find ourselves in is not 3D, but rather $(2+1)$D. So when I describe the formal theory of articulated planes, I will be using this $(2+1)$D space. I refer to it as "relief space".

Relief space is very important in visual perception. "Visual space" is an articulated plane, namely the "visual field" (2D), articulated by "depth" (1D). A painter paints "pictorial spaces" where the picture plane roughly accounts for the 2D (visual field say) and the 1D is the implicit "pictorial depth".

More generally, any scalar field is such an "articulated plane". Examples are meteorological maps of temperature or barymetric pressure. But scalar fields over two-fold extended continua abound in the sciences.[15] Thus this type of formal shape has numerous applications. Once up to it, you will discover novel applications by the day. Their importance almost beats triangles.[16]

An important final issue is that of smoothness. I will assume that all surfaces of interest are smooth on all scales of interest. It is just that I don't have the time to go into fractal structures. But smooth structures are nice and important in their own right. After all, the earth *does* look like a perfect sphere as seen from the moon, despite the fact that your easily trip if you don't watch your feet.

What types of surface articulations exist?

Alberti was the first person to attempt a complete zoo of surface types. This is a quote from his famous book on painting dating from the early Renaissance (emphasis added by me):

> We have now to treat of other QUALITIES WHICH REST LIKE A SKIN OVER ALL THE SURFACE OF THE PLANE. These are divided into three sorts. Some planes are flat, others are hollowed out, and others are swollen outward and are spherical. To these a fourth may be added which is composed of any two of the above. The flat plane is that which a straight ruler will touch in every part if drawn over it. The SURFACE OF THE WATER is very similar to this. The spherical plane is similar to the EXTERIOR OF A SPHERE. We say the sphere is a round body, continuous in every part; any part on the extremity of that body is equidistant from its centre. The hollowed plane is within and under the outermost extremities of the spherical plane as in the INTERIOR OF AN EGG SHELL. The compound plane is in one part flat and in another hollowed or spherical like those on the INTERIOR OF REEDS or on the EXTERIOR OF COLUMNS.

[15] Yes, including the neurosciences.

[16] Of course, triangulation is the standard way to deal with articulated shapes in practice.

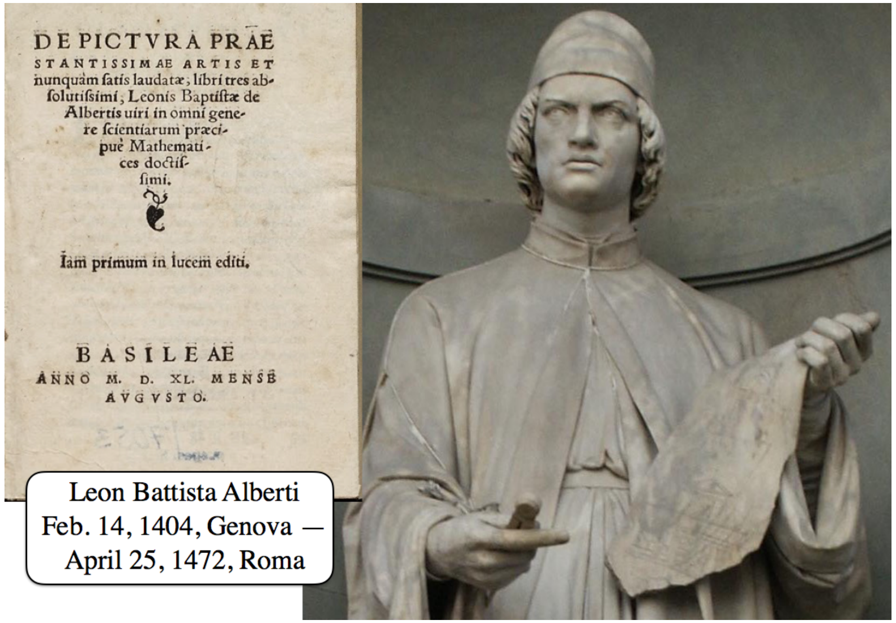

Alberti, a major intellectual force during the early Italian Renaissance.

There are a few exciting insights here. The first one is that Alberti talks about "qualities which rest like a skin over all the surface of the plane". The articulations are like colors. Alberti evidently has relief space in mind.

The other one is the catalogue of shapes. Alberti has these categories, evidently based on the notion of surface *curvature*:

— surface of water

— exterior of a sphere

— interior of an egg shell

— interior of reeds

— exterior of columns

Notice that both "hollowed out" and "swollen outward" admit of degrees. They might be measured on non-negative "more–or–less" scales. The flatness is different though, it does in no way admit of degrees. Intuitively, it is on the edge between convex and concave, which implies that it is not a bona fide category. This extends to the"interior of reeds" and the "exterior of columns". These are evidently borderline cases, but borderlines of what?

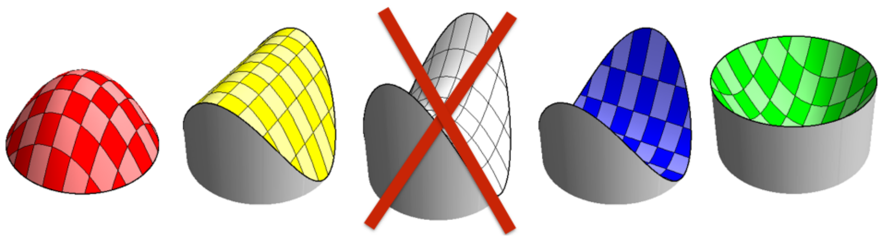

The gist of Alberti's list. "Like a horse's saddle" is sadly lacking.

You can evidently go smoothly from the interior of an egg shell to the interior of reeds, but what if you *move on*? You can evidently go smoothly from the exterior of a sphere to the exterior of columns, but what if you *move on*? In both cases you would arrive at "like a horse's saddle", a surface that is convex and concave at the same time.

Alberti's categories served scientists and artists for centuries. The decisive breakthrough was the paper by Carl Friedrich Gauss of 1827. Gauss achieved a full categorization of local surface shapes and much more. He founded the field of "differential geometry", which is local, smooth geometry, in one fell blow.

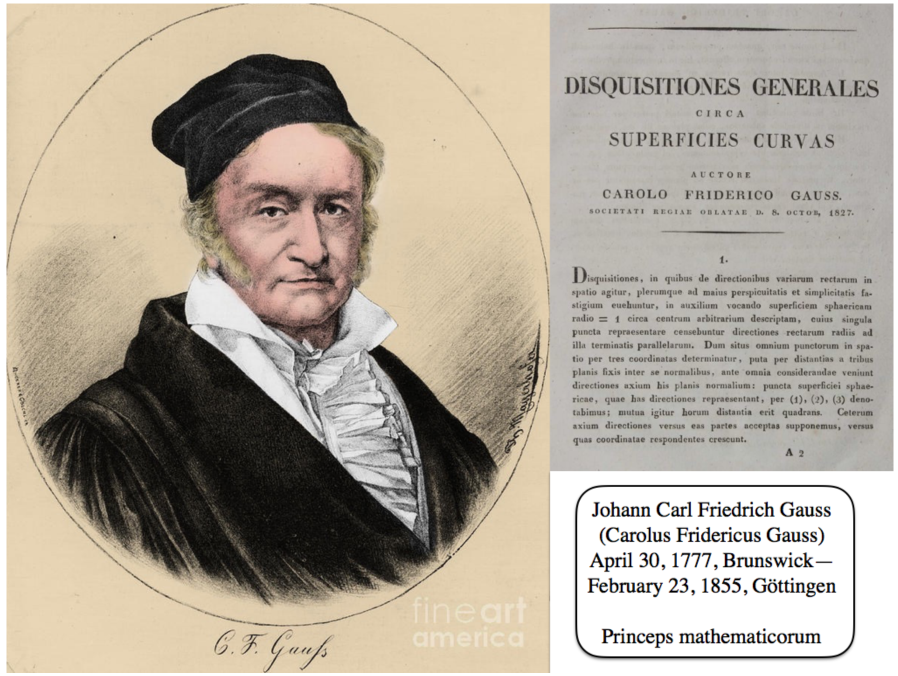

Carl Friedrich Gauss, Prince of Mathematics.

For the present discussion we need to make a few annotations though. The first is that Gauss worked in 3D (Euclidian space) rather than $(2+1)$D (relief space). This introduces numerous complications, but yields hardly any additional insights to my treatment. It merely complicated matters unnecessarily—at least for most applications. It was important because of Gauss's famous *Theorema Egregium* ("Remarkable Theorem"), which throws its shadows all the way to Einstein and beyond, but it is entirely without interest in my present setting.

Gauss did a detailed analysis of how the local surface attitude varies in the environment of a point. As you can image, this is not such an easy task.

As you move in a certain direction the surface tilts in various ways. This is described as a “torsion” and a “curvature”.

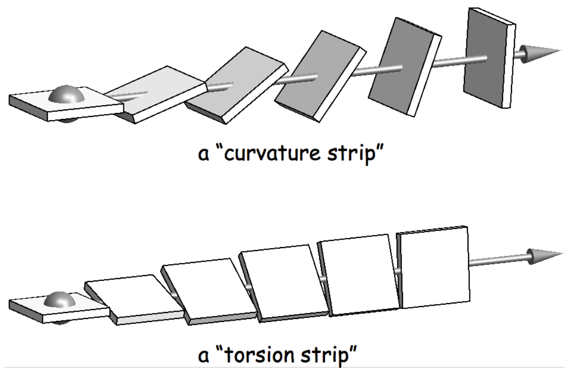

Some examples of surface strips.

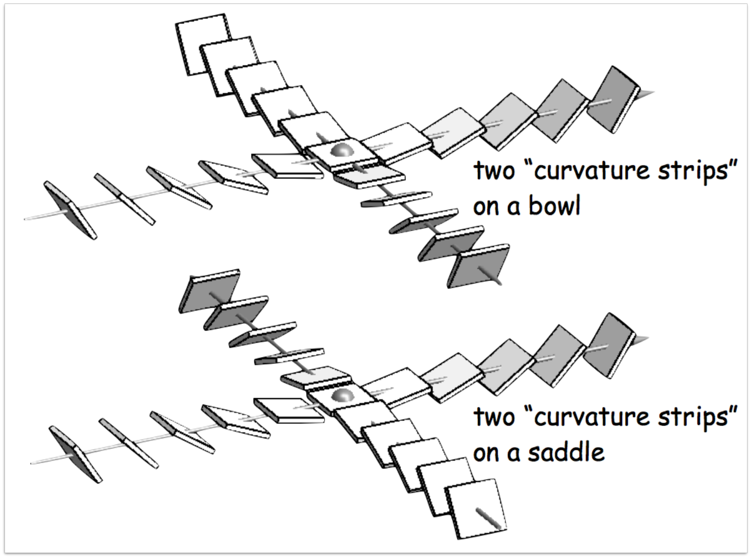

Strips of principal curvature on an elliptic (top) and a hyperbolic (bottom) surface.

Gauss discovered that there are two, mutually perpendicular, directions in which there is no torsion. In these directions the surface is purely curved. They are known as the “principal directions of curvature” of the surface. Of course, they differ from point to point.

Since curvature has a “sense”, there are two kinds of surfaces, those where the principal curvatures have the same sense and those where they are of opposite senses. Thus, in the Gauss categorization, there are exactly two local surface types, he called them “hyperbolical and “elliptical” (I’ll explain why below). The elliptical category may be split into two subcategories “convex” and “concave” if you care to distinguish between up and down (both hill tops and pit bottoms are elliptic).

“Flat” is not a shape at all, it is a singular condition. This makes sense when you think of how hard it is to make a truly flat surface.[17] Formally, they occur with probability zero.

Likewise, cylindrical forms are the transitions between elliptic and hyperbolic (Gauss called them “parabolic”). They occur with probability zero, but—because there are so many points in a surface (severe understatement!)—they occur as points on certain curves, the boundaries between elliptic and hyperbolic regions.

Why the funny names elliptic, parabolic and hyperbolic? There are various ways to explain this. Gauss used algebraic language that derived from the structure of a family of planar curves called “conic sections” by the Greeks that might be said to generalize the circle. Examples are ellipses, hyperbolas and parabolas. However, you probably prefer a more geometric insight. It is easily obtained starting from the classical notion of conic sections.

With “cone” one implies a “right circular cone” (rotationally symmetric), with an infinite axis. That is to say, it looks like two “common” cones (the ones that end in a sharp vertex) in tandem. If you cut it with a plane you either obtain a closed curve (an ellipse), or an open curve of two branches (a hyperbola). In a singular case—dividing the ellipses from the hyperbolas—you obtain an open, single branched curve (a parabola). These curves are the outlines of the “wounds” inflicted by the cutting with a planar blade. The curves display the shape of the cone to who has the eye to appreciate such things.

[17] I used to pay top dollars for flat mirrors when I ran a group at a physics laboratory.

Charles Dupin had the bright idea to extend this technique of revealing shape by cuts to general local surfaces. At a given point he cuts the surface with a plane parallel to the tangent plane at that point. The notion of "tangent plane" may not be familiar, so I'll explain it first. I'll get back to Dupin soon enough.

This is "the earth" in three distinct views. Both scope and scale differ, although the scale–scope ratio is roughly constant. At left is a view of narrow scope and high resolution, here we meet with almost infinite complexity. At center is a far view, at least as judged from the human condition, the horizon distance depends simply on your eye-height. In your awareness the earth is flat, except for some ripples on the water. It takes reflective thought to infer that the horizon might be due to otherwise invisible curvature. At right we (or rather, a NASA *robot) see the earth from a distance. In your visual awareness it appears as a smooth sphere, painted with pretty texture. It is of some interest to estimate how many orders of magnitude separate such views (try!). In fields like the neurosciences you are often confronted with pictures that are easily as far apart (on a log-scale, of course). This should give rise to some thought. All these views depict realities, but these are hardly the same. Your actuality here and now can hardly be at all three simultaneously. The regularities of Nature are also very different in these three worlds.*

Remember that the earth was once supposed to be flat. The "flatness" was a tangent plane at the globe. The flat earth people believed that they were all on the same plane, but we now understand that every geographical location implies a different plane. In practice I'm a flat earth person (the Netherlands, where I live, certainly looks flat to me), except when I interact with people a long distance away, or have to travel there. The only way to *see* that the earth is not flat are NASA photographs taken from outer space.

A formal definition of tangent plane at some fiducial point is as follows: Take any three distinct points in the neighborhood of the fiducial point and construct the unique plane that contains them. Now take the points closer and closer to the fiducial point. When very close indeed, taking them any closer makes no difference any more. The limiting plane—all three points at the fiducial point!—is the tangent plane. Because I assume everything is smooth, such a plane certainly exists.

A physical definition is more fun, but in practice works only for convex elllpitical points. You may use it to construct tangent planes at points on a boiled egg. Simply let the egg touch a (flat) table top. It will touch at one point. The table top is the tangent plane at that point. Got it?

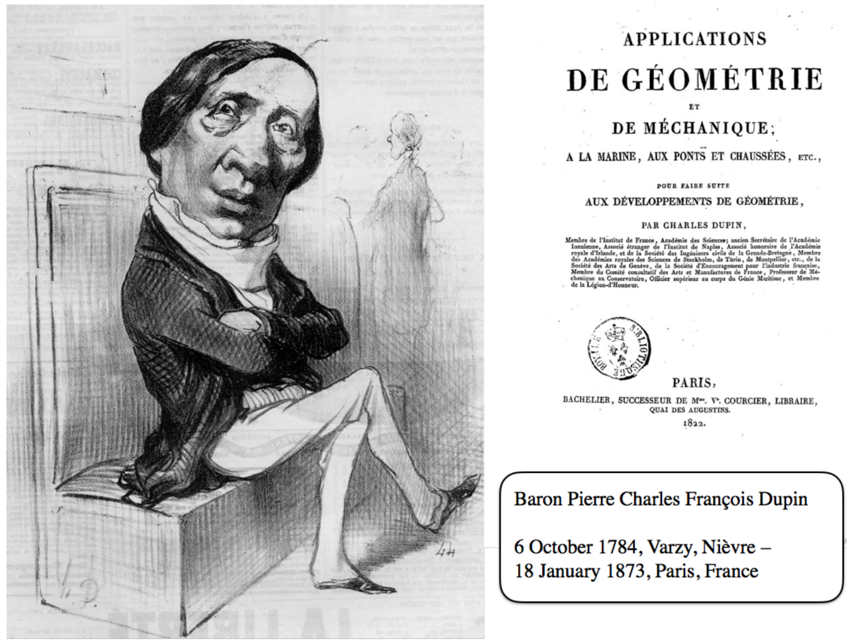

This Dupin caricature of Dupin is by Daumier. Daumier drew many important persons of his time.

Dupin considered cuts very close to the tangent plane. He understood that when you cut close enough you always obtain a conic section. He actually advocated a pair of cuts, one below the surface and one above. That way you get both branches of the hyperbolas and you differentiate between convexities and concavities. The figure you get is called "Dupin's Indicatrix" at the fiducial point. It is a very intuitive, because geometrical, way to arrive at an understanding of local surface shape.

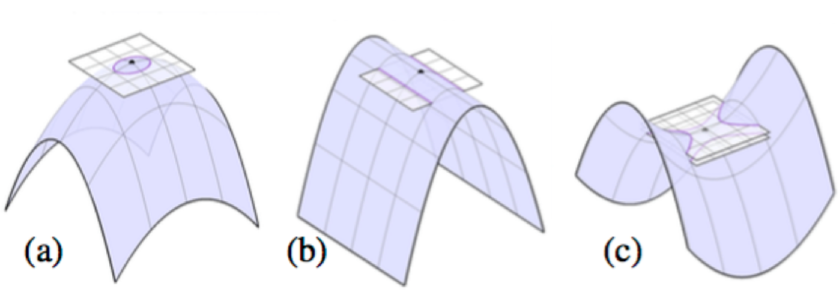

Some instances of Dupin's indicatrix. In (a) the cut is below the tangent plane and it yields an ellipse. In (b) the cut is also below the tangent plane and it yields a degenerate parabola (looks like two parallel lines). In (c) one cut is above the tangent plane and it yields one branch of a hyperbola, whereas the other cut below the tangent plane yields the other branch.

In order to put some structure on it one *assigns a sign* to the surface curvatures. Of course, it is not obvious how to do this! I'm sure this was Alberti's problem: he simply failed to see that a surface may be "hollowed out" and "bulged outwards" at the same time! The point is that a surface is usually curved differently in different directions. This is easily seen in Dupin's indicatrix, because the distance of a point on the indicatrix is inversely proportional to the square root of the curvature.[18] It evidently varies among the curve. Moreover, it makes sense to distinguish between "above the tangent plane" and "below the tangent plane" branches by assigning a sign. Mathematicians conventionally use a plus sign for indicatrix points on the upper cut and negative for points on the lower cut.

It is easy enough to define a measure of "curvature". In (2+1)D-space one uses the value of the second order directional derivative.[19] It has the advantage that the sign is automatically supplied.[20] Since the Dupin's indicatrices are conic sections, the curvature reaches two extremal values if you follow it over all directions. These are known as the "principal directions" and the curvatures as the "principal curvatures". Thus you deal with a pair of curvatures, say $\{\kappa_2, \kappa_2\}$, where I'll let $\kappa_1 > \kappa_2$.

Here the confusion starts! So what is THE curvature of the surface? The so called intrinsic curvature is of major interest in mathematics. Gauss defined it as $K = \kappa_1 \kappa_2$. It is special because it can be found by measurements that are purely within the surface (hence "intrinsic"). It is also most confusing to common sense people because cylindrical and conical surfaces have zero curvature, whereas any fool can see they are not flat. However, Gauss' argument is that they can be "developed into the plane" as they certainly can. Indeed, cylindrical surfaces are often produced by bending flat plates. Moreover, the intrinsic curvature fails to distinguish between "hollowed out" and "swollen outwards" surfaces, which really goes against the grain of common sense—Alberti would cringe!

Another common curvature measure is the so called "mean curvature", not unexpectedly defined as $H = (\kappa_1 + \kappa_2)/2$. This neatly differentiates the Albertian categories. However, the mean curvature of a symmetric saddle is zero, which again conflicts with common sense.

Only in 1890 Felice Casorati, an Italian mathematician, suggested a surface curvature expressly designed to fit common sense, namely $C = \sqrt{(\kappa_2^2 + \kappa_2^2)/2}$. The Casorati curvature is indeed only zero for planes. This is indeed what our intuition understands as "surface curvature". However, it fails to distinguish between the categories of shape. Thus one needs and additional measure for the "quality" of curvature. This is evidently what Alberti intended with "qualities which rest like a skin over all the surface of

[18] Puzzled? Try it in 1D: The curve $\kappa x^2/2$ has curvature κ, the tangent line is the x-axis. The cut at separation d is at distance $a = \sqrt{2d/\kappa} \propto 1/\sqrt{\kappa}$ from the origin, for $\kappa a^2/2 = d$.

[19] In the cortical implementation in V1 this is the activity of a Hubel and Wiesel "line detector".

[20] Yes, the mathematicians have it right, same convention.

the plane". The "quality" is essentially the ratio of principal curvatures. It turns out that $S = \arctan\left((\kappa_2 + \kappa_1)/(\kappa_2 - \kappa_1)\right)$ is a more natural parameter in shape space. This "shape index" takes values in $(-\pi/2, +\pi/2)$.

This is Felice Casorati. He had the guts (not easy for a mathematician) to publish "Mesure de la courbore des surfaces suivant l'idée commune" *in a mathematical journal in 1890. His confrères killed him over that. They couldn't stand the* idée commune, *which is considered unscientific, because subjective. Casorati is better known for his work in complex analysis.*

The Casorati curvature and shape index together define the habitus. The other freedom is the orientation, for which one may use the angle between the axis of maximal principal curvature and some arbitrary reference in the plane. The orientation O has period π, I take it to take values in $[-\pi/2, +\pi/2)$.

The triple $\{C, S, 2O\}$ forms a natural parameterization of shape space.[21] The Casorati curvature measures the distance from the origin, the shape index the latitude and the orientation the longitude. It can be shown that distances in this space represent the natural measure of fit, that is the square root of the average squared deviation.

In measuring the distance between two shapes we can go totally Procrustean.

This is what we do:
— we rotate in the plane, aligning the axes of maximum principal curvature,
— we scale the height, setting both Casorati curvatures to one,
then the remaining r.m.s. measure of fit—that is the Procrustes distance—equals the difference of shape indices. This is a measure of shape *quality*, for instance all convex cylinders[22] ("outsides of columns") have the same shape, they differ only in size (Casorati curvature) and spatial attitude (orientation). This is a result that perfectly fits what I understand as "common sense".

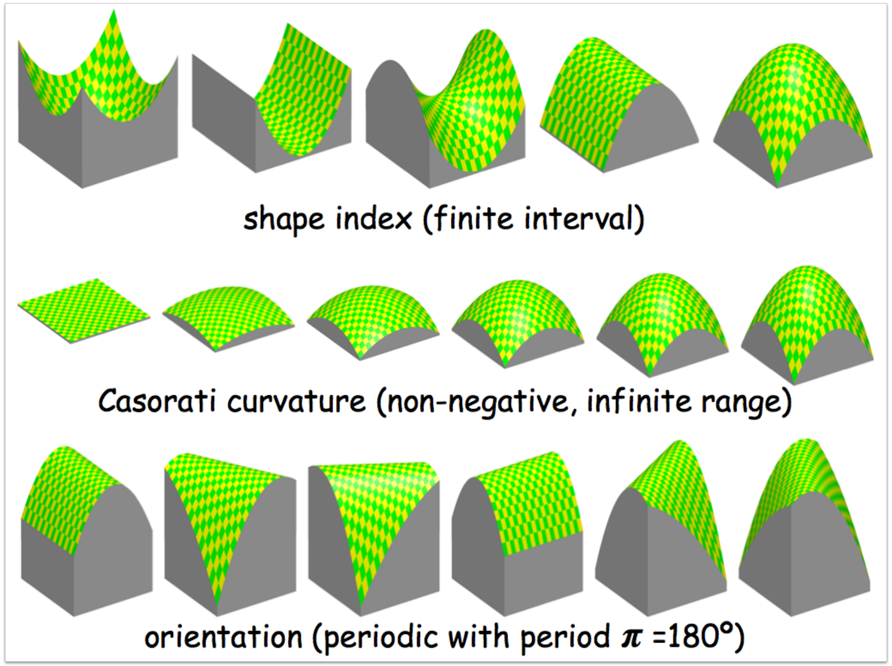

This shows the intuitive parameterization of local surface shape. Shape index, Casorati curvature and orientation together form a natural coordinate system for shape space. Casorati curvature is distance from the origin, shape index latitude and orientation longitude. In this space the metric is Procrustean with respect to a comparison with planarity.

[21] Polar coordinates in Euclidean 3-space. The $2O$ instead of just O is necessary to obtain the 2π periodicity.

[22] I use this example because Gauss does not differentiate cylinders from planes.

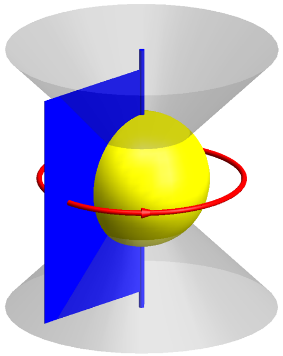

The natural (polar) coordinates of shape space. The yellow sphere is a locus of constant Casorati curvature, the blue plane a locus of constant orientation and the gray cone a locus of constant shape index. The red circle indicates the rotational symmetry (period π, thus each orientation occurs only once!). The blue vertical axis is the "umbilical axis" the locus of shapes like the spherical shell with coincident principal curvatures. For umbilicals the orientation remains indeterminate.

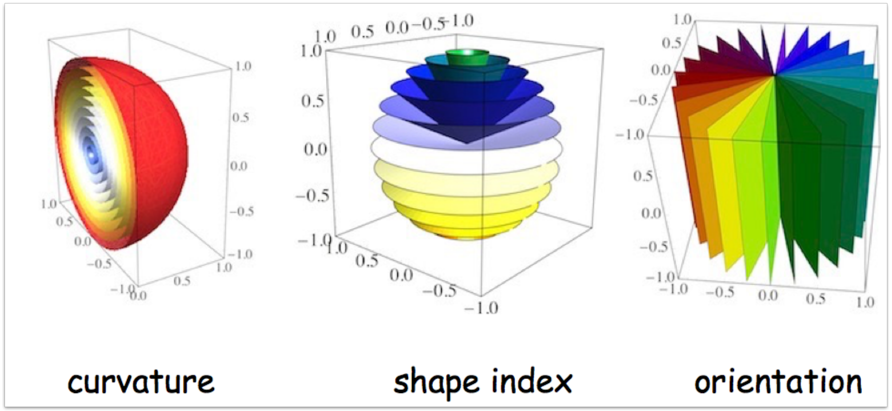

At left loci of fixed Casorati curvature, at center loci of fixed shape index and at right loci of consant orientation (period π, but each orientation occurs only once!).

Given a metrical shape space, you can calculate many things of considerable interest. For instance, the fact that Alberti's oversight was not noticed for centuries gives rise to thought. Is our vision unfit to notice saddle shapes? One fact that appears to fit such a thought is that sculptors generally concentrate on elliptical parts, especially convex ones. A cursory look at the cover picture will suggest as much. If this were true, is it due to the fact that saddles are rare? (That would be the standard Bayesian argument.) Well, we can easily find the probability density functions for a plane that is articulated through isotropic Gaussian noise. The result is interesting: there is actually a higher probability to encounter a hyperbolic than there is to encounter an elliptic shape! So the Bayesian "explanation" does not "save Alberti's face".

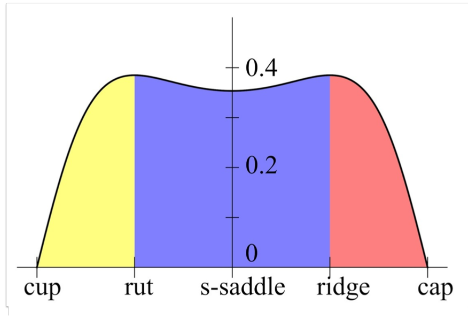

The probability density for the shape index of a plane articulated with isotropic Gaussian noise. The probability to encounter a saddle (the blue range) is about 57%.

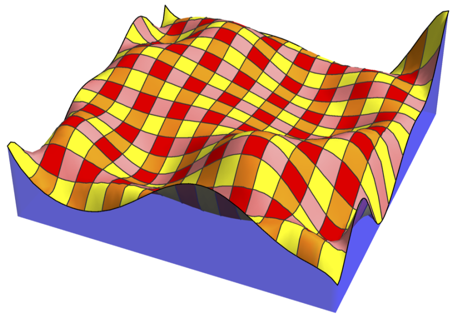

A Gaussian relief is a pretty landscape with flowing hills and dales.

Of course, in practice there will be various kinds of deviations from this. Since a potato has a smooth surface that is curved in itself, it it likely to have a bias towards convexity, like the sphere, which is all convex umbilical. A geographical landscape will lack major concavities due to the inevitable water erosion processes. And so forth. The Gaussian random relief is a mere ideal.

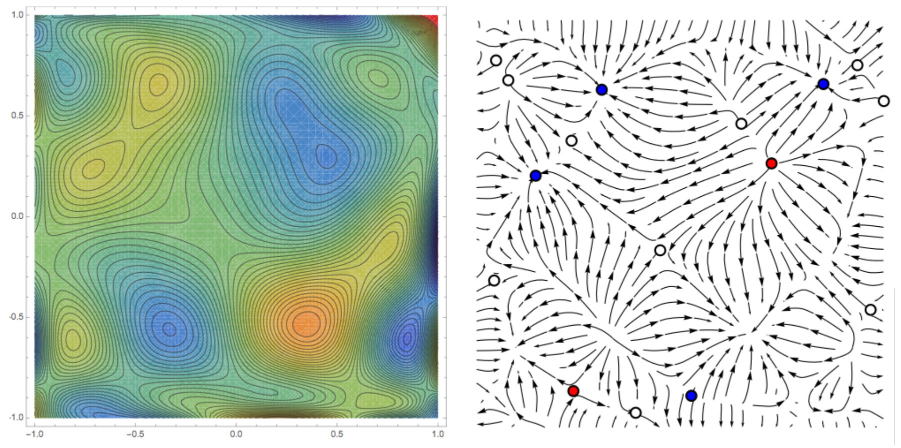

At left a relief map of the Gaussan relief, the curves are loci of equal height. Notice that you get to see the Dupin indicatrices at the pits, peaks and passes of the landscape. At right the flow lines of steepest descent, how water would run down hill (at least initially, I do not consider any dynamics here). This is the "gradient field".

The landscapes are interesting because they show that (one might almost forget after reading Gauss-type differential geometry!) the local concavities, convexities and saddles hardly exhaust the wealth of surface form. Far from it! Of course, it might be a problem that this richesse is really too much for analysis. Here we meet the scale and scope issues again. A hill is a local object, simple to study, a landscape is a multiplicity of hills, complicated to study. In the divide and conquer strategy one studies just one generic hill and is done with. What is left is to study relations between hills and so forth. In rare cases one considers multiple hill objects, think of camel-backs or bikini-tops, for instance. The most useful scope depends upon context.

In order to get a feel for this, consider the case of curves. In the relief view, curves are articulated lines. One description would be to plot height as a function of distance along the line. (This is usually called a"graph", likewise landscapes are often called graphs of functions of two variables.)

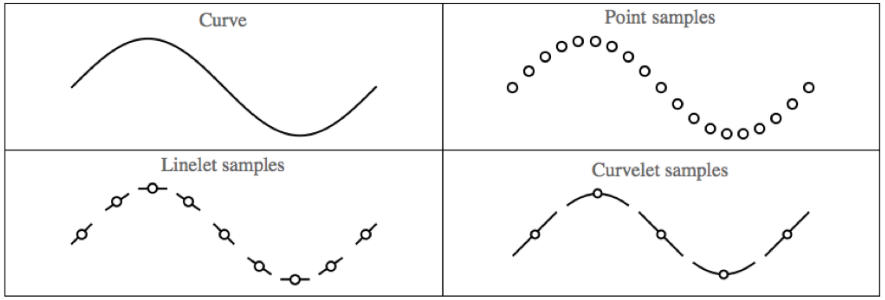

Top left a "graph", which is an articulated line, a relief often known as a "curve". Top right, point samples as a discrete representation of the relief. It is the 0^{th}–order representation. Bottom left, a discrete representation by line samples, the tangent lines at discrete points. It is the 1^{rst}–order representation. Bottom right, a discrete representation by local shape samples. It is the 2^{nd}–order representation.

In applications reliefs are represented through discrete samples. "Samples" are local views. At the rock bottom one has mere point samples. Then the scope is nil, for knowing one point sample tells you nothing about the next one. It is more useful to have sampling of "linelets". At any sample you have a finite slope that tells you something about the local trend and lets you predict the height of the next point, at least if that is close. Even better to have a sampling of curvelet samples. Why better? Well, you can compute several linelets from a single curvelet, but not vice versa. The curvelets imply greater scope. That is why you need more linelets than curvelets (and even more point samples) in order to approximate the relief within a certain tolerance.

Of course, you can use even higher orders. But usually a reasonable sweet spot is order two (curvelets, shapelets). That is because whereas many reliefs are indeed smooth, few are *that* smooth. If the relief is fractal you are better off with point samples even.

Human vision is largely focussed on order two, most likely because the lower orders are not optically specified, or require absolute calibrations not available to physiological mechanisms and the world is too unpredictable over larger scopes. The Hubel and Wiesel "line detectors" are really second order directional derivatives, sustaining local differential geometry of the Gauss type.

Reliefs over large ranges are to be treated as *multitudes*. One way to deal with them is through maps. Maps are convenient because of the possibilities of using various scales (atlases) and degrees of "generalization" or specialization. This is obviously implemented in the various areas of the early visual system.

Useful maps for articulated planes are representations of the height contours and representations of the field of steepest descent curves. Technically one uses the *gradient*, which is often represented as a vector in the direction of steepest descent with a magnitude equal to the slope. A very useful abstraction from this is "gradient space". In gradient space I simply retain the gradient vectors, but ignore their location. I simply let all gradient vectors start at the origin. Why might this be useful? Well, for instance, if you illuminate the relief from a certain direction, with a directional beam, then points that receive the same illumination lie on simple curves[23] in gradient space. If you invert the map from the landscape to gradient space, you immediately get the shading.

Here you meet with an interesting and important phenomenon: inverting the map is not trivial because a single point in gradient space may correspond to several (or many) in the landscape! Thus the gradient plane is a thoroughly wrinkled copy of the articulated plane. Hassler Whitney[24], a mathematician who is another one of my heroes, figured out (in 1955) what kinds of wrinkles to expect. You can have folds and cusps.

The folds and cusps are where the degree of covering of the map changes. The fold and cusps are objects that refer to a wider scope than the shapelets, they are loci where the nature of the local shapelets changes qualitatively. This is the case because folds correspond to parabolic points. Thus the shape index changes between elliptic and hyperbolic at a fold. That is why they are a convenient way to deal with the complexity of reliefs that goes beyond the mere local views. Yet folds and cusps are again "single objects", thus they offer a great opportunity to broaden your view without overloading your bird brain. I really appreciate that.

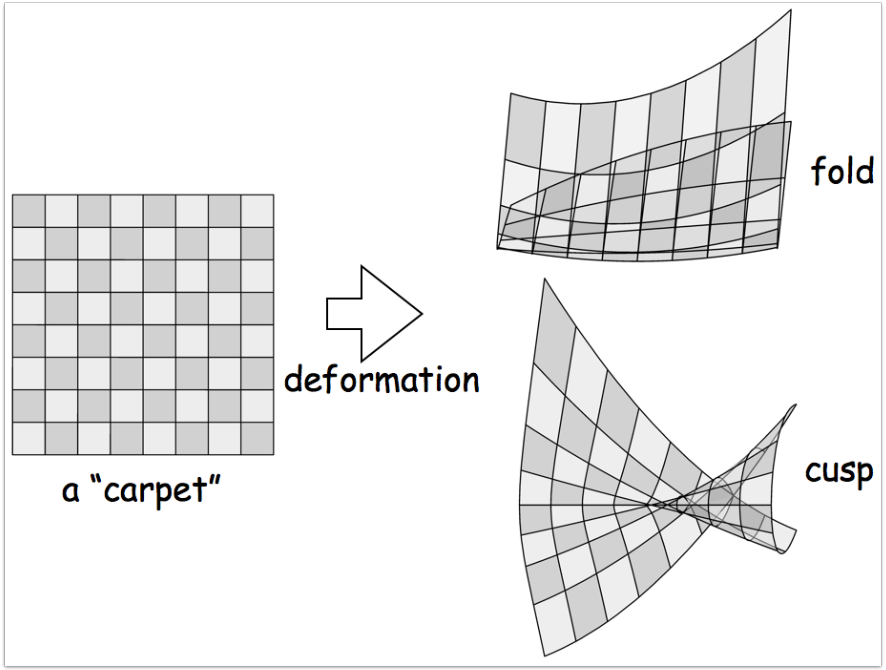

Examples of the Whitney "fold" and "cusp". Notice the degree of covering supplied by the wrinkled carpet. At the fold (top) there is a double covering on one side and no covering on the other side. At the cusp (bottom) the covering is either single or triple.

The study of the wrinkles in gradient space is thus (among more) a cheap way to study higher order properties of relief. Indeed, in some sense it might be said that "nothing happens" between the wrinkles, the folds and cusps are really "where the action is". Even a superficial understanding of the folds allows you to infer important properties of reliefs, even without form derivations. For instance, you can easily show that the stationary points of the shading of reliefs under directional illumination (light spots, dark spots and tonal saddles) is such that they occur *only* on the preimages of the folds. When you vary the direction of the illuminating beam the stationary points travel along

[23] Easy enough to show that these are conic sections.

[24] Remember that Google also helps you to find facts about people, especially famous ones.

the parabolic curves of the relief. You obtain such important, qualitative insights almost for free!

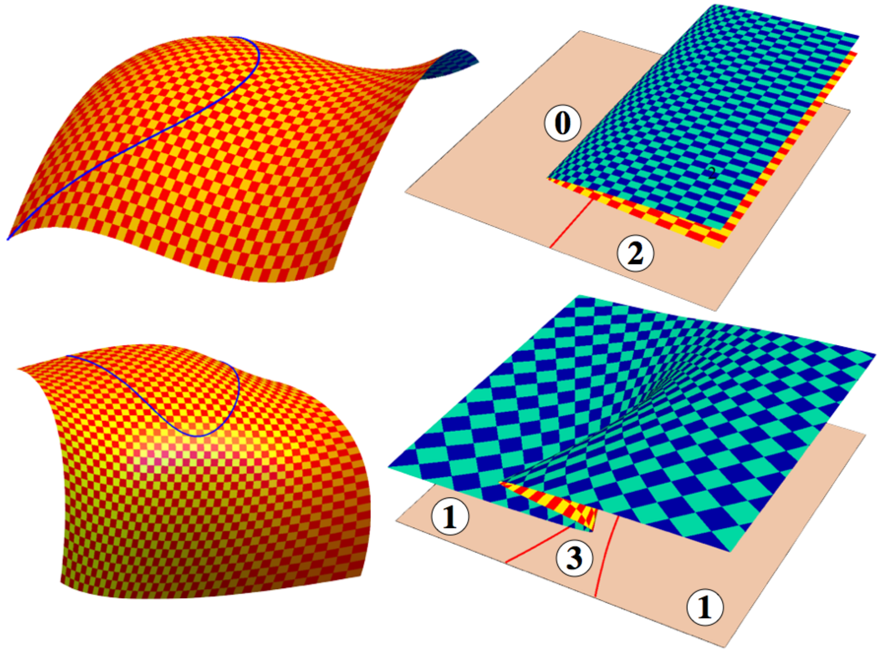

At left examples of landscapes that illustrate the Whitney "fold" and "cusp". The landscapes are smooth throughout, the loci of the folds are the blue curves at left, the locus of the cusp is a point on one of these. At right the corresponding gradient spaces, the pink planes, with the folds indicated in red. The checkered surfaces may help your intuition a bit, they graphically display the "fold" and "cusp".

Maybe that was the reason why Felix Klein proposed the pattern of parabolic curves on a surface as important for the "looks" of a surface. More specifically, he speculated that the beauty of faces might be found in this pattern. In order to check this, he had a student draw the parabolic curves on a copy of a bust of the Apollo Belvedere, at the time considered the pinnacle of male beauty. Hard to say whether he succeeded. Check for yourself! In any case, the idea was hardly crazy. Most painters would agree that the shading patterns reveal the face—they don't paint the face but the shading patterns, then psychogenesis of visual awareness constructs the face for you. But this pattern is essentially determined by that of the parabolic curves.

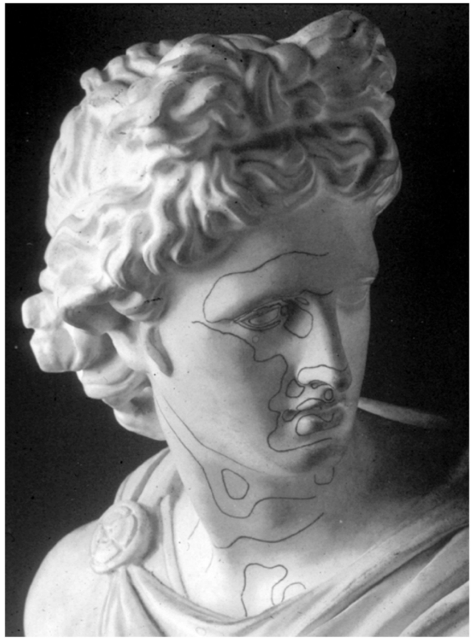

Pattern of parabolic curves drawn on a copy of the face of the Apollo Belvedere by an unidentified student of Felix Klein. This intriguing object is at the Göttingen mathematical institute.

The pattern of parabolic curves (and thus the folds and cusps in gradient space) is of even wider scope than a single piece of fold or a cusp by itself. This is one powerful way to expand your vision of relief.

Apparently the gradient map is a powertool. Might something similar apply to the map into shape space? One would certainly expect so. However, perhaps sadly, there is no extensive formal analysis as to date. I will simply make a few cheap remarks and leave it up to you to become famous!

I have already shown how to map each point of a relief to a point in shape space. By plotting all points, I map the relief plane into shape space. It will be some surface embedded in that space.

Of course, it might be degenerate. A plane is simply mapped on the origin, a mere point. It is simple enough to derive conditions that ensure that a patch of relief will map on a surface patch in shape space. Generally, such conditions are bound to be met with probability one. Exceptions will be the generic singularities of maps from the plane into three-dimensional space. These were again classified by Hassler Whitney.[25]

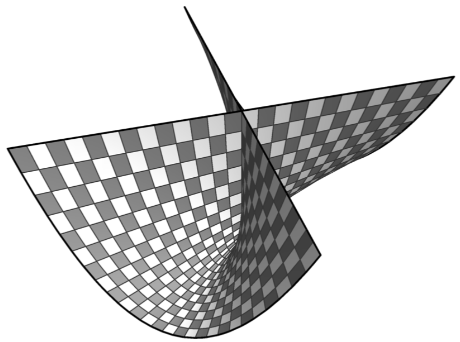

The Whitney umbrella. Notice that the surface passes through itself! The singularity is the "pinch point" where the self intersection ends.

The generic singularity is the "Whitney umbrella".[26] The pattern of pinch points and curves of self-intersection are of obvious relevance, a kind of "shape skeleton" for the relief. It is certain to be important. So far, I have been able to arrive at some interesting relations to the singularities of gradient space, but the real work awaits to be done.

[25] In the nineteenforties.

[26] Why "umbrella"? Well the canonical model is the implicit equation $x^2 = y^2 z$, which also includes the negative z axis—the *handle* of the umbrella.

So *what* is SHAPE?

Now do we finally know "what SHAPE is"? I'll recapitulate the major points. Of course, I'm talking only formal matters here, for accounts on artistic shape you should shop somewhere else. It is not that I feel the formal understanding to be more important than the artistic one, rather the opposite. But a solid understanding of the formal description is a sine qua non for anyone who proposes to undertake an experiment involving "shape".

The first point to understand is that the notion of *shape* presupposes the notion of *space* and that the notion of space involves a notion of *metric*, as well as the notion of a *group of transformations* whose action on configurations leaves their internal metrical properties invariant. Such groups of transformations are often known as the "proper movements"[27] of the space and their actions as *congruences*. Notice that space, metric and movements come as a single package, usually called *a geometry*. This is essentially Felix Klein's famous *Erlanger Programm*.[28]

In most cases we would also consider a wider group of transformations, namely that of "similarities". The group of similarities allows one to *scale* configurations. It includes the group of movements, which are scalings by a factor of one. Euclidean similarities are needed to be able to say such things as that a tennis ball and the earth "have the same shape". One has something similar in other spaces, although the notion of "similarity" may differ widely.[29]

[27] An example of an improper movement would be a reflection.

[28] This is Felix Klein's proposal (*wissenschaftliche Programmschrift*) as he joined the University Erlangen-Nürnberg in 1872, entitled *Vergleichende Betrachtungen über neuere geometrische Forschungen* (printed in Mathematische Annalen **43**, (1893), pp. 63–100). The work had enormous impact upon geometry. (For a complete English Translation get http://arxiv.org/abs/0807.3161.) I take a somewhat narrow view here. For instance, my insistence on a metric is required because I aim to define shape as a quantifiable property.

[29] For instance, in Euclidean space a similarity is parameterized by a single scaling factor, but in (2+1)D space it is natural to distinguish two kinds of similarities (which might be combined in various ways), one pertaining to the 2D part and one pertaining to the 1D part. It is

Given these formal objects, the Procrustes method yields a very general method that allows one to define a "shape space" and put a metric on it. That is typically all that is required for applications in vision.

Notice that this is a very flexible description that allows one to make "shape" a useful concept in mutually very different settings. This is the meaning of the formal definition I suggested at the outset:

> SHAPE is an invariant under the group action of a group of "congruences" or "movements",

which no doubt sounded somewhat mysterious at the time. Throughout this booklet I tried to unpack it for you, step by step.

This formal description is an "ideal one" in the sense that—for instance—neither Euclidian points, lines and planes, nor curves and surfaces exist as material manifestations in physical space. Even the smoothest polished surface is not a surface on the molecular scale. Even the electron is not a point.[30] In cases of real interest the relevant scales and scopes depend on the context and their choice is an important part of the formulation of your research program. Only once these issues have been settled, the formal structures may be applied as suitable "generalizations". Just remember the example of the earth.

Indeed, the earth may be taken to be spherical, flat,[31] or fractal, it is all up to your perspective! Only once you have decided on your perspective can a certain formal shape description be made to apply.

indeed natural, because these parts are not necessarily (or indeed usually) mutually commensurate.

[30] I'm being very selective here. Consider how you would handle the shape of an ocean wave, a swarm of mosquitos, or a galaxy. Problems start on any scale you happen to fancy.

[31] Yes, I wouldn't say the flat earth people are *wrong*, it is just that their perspective ill fits our contemporary framework. Forcefully pushing their views, is merely an ill-conceived cultural mismatch. The "flat earth" fits my personal visual awareness quite well though. Are you ever bothered by the convexity of the floor? Perhaps you are, I'm not. So what?

SOME IDEAS FOR FURTHER READING

For a general introduction to the various geometrical notions I used here, few books can beat Hilbert and Cohn-Vossen's *Anschauliche Geometrie* of 1932 (Berlin: Springer). It has been reprinted many times. An English translation *Geometry and the Imagination* is also available (in the AMS Chelsea Publishing series of the American Mathematical Society, 1999).

A good way to quickly pick up details of Kendall's spherical blackboard is http://www.cantab.net/users/michael.behrend/ley_stats/spher_black/pages/spher_black.html
It also has references to Kendall's original papers (much harder to read).

Gauss's famous treatment *Disquisitiones generales circa superficies curvas* is available on the Internet in various translations. You'll need some true grit to read it. An easier way to pick up the essential ideas with respect to shape is my *Solid Shape* (Cambridge MA: MIT Press, 1990). Hackers download it for free, but the publisher does not encourage me to tell how.

The notion of (2+1)D space is explained in more detail in my *Graph Spaces* (Traiectina: de Clootcrans Press, MMXII). It can be downloaded for free (see following pages).

Details on the shape space are found in my *Shadows of Shape* (Traiectina: de Clootcrans Press, MMXII). Again, this can be downloaded for free.

A tutorial on Scale-Space (the formal background for what I said on scale and scope) can be downloaded from
http://www.gris.informatik.tu-darmstadt.de/~akuijper/course/TUD/vbc96tutorial.pdf
This has references to the most useful sources. I recommend this tutorial (*Introduction to Scale-Space Theory: Multiscale Geometric Image Analysis*, Tutorial VBC '96, Hamburg, Germany, Fourth International Conference on Visualization in Biomedical Computing) because it should appeal to students of the neurosciences.

There are numerous books on artistic shape. My best sources are used bookstores, but that takes time.
I recommend:

Philip Rawson (1969): Drawing (The Appreciation of the Arts/3, ed. Harold Osborne), Oxford University Press.

and especially

L.R.Rogers (1969): Sculpture (The Appreciation of the Arts/3, ed. Harold Osborne), Oxford University Press.

Apart from these you should try to quickly skim through numerous technical treatises on *how to* sculpt, draw, paint, . . . , you name it. Of course, you will have to distill the academic messages yourself.

OTHER EBOOKS FROM THE CLOOTCRANS PRESS:

1. Awareness (2012)
2. MultipleWorlds (2012)
3. ChronoGeometry (2012)
4. Graph Spaces (2012)
5. Pictorial Shape (2012)
6. Shadows of Shape (2012)
7. Through the Looking Glass: on Viewing Aids (2012)
8. Painting to Marble (2012)
9. Experimental Phenomenology: Art & Science (2012)
10. The Spirit of the New Style (2013)
11. World, Environment, Umwelt and Innerworld (2013).

(Available for download here.)

ABOUT THE CLOOTCRANS PRESS

The Clootcrans Press is a selfpublishing initiative of Jan Koenderink. Notice that the publisher takes no responsibility for the contents, except that he gave it an honest try—as he always does. Since the books are free you should have no reason to complain.

THE "CLOOTCRANS" appears on the front page of Simon Stevin's (Brugge, 1548–1620, Den Haag) *De Beghinselen der Weeghconst*, published 1586 at Christoffel Plantijn's Press at Leyden in one volume with *De Weeghdaet, De Beghinselen des Waterwichts*, and a *Anhang*. In 1605 there appeared a supplement *Byvough der Weeghconst* in the *Wisconstige Gedachtenissen*. The text reads "*Wonder en is gheen wonder*". The figure gives an intuitive "eye measure" proof of the parallelogram of forces.

The key argument is

> *de cloten sullen uyt haer selven een eeuwich roersel maken, t'welck valsch is.*

Simon Stevin was a Dutch genius, not only a mathematician, but also an engineer with remarkable horse sense. I consider his "clootcrans bewijs" one of the jewels of sixteenth century science. It is "natural philosophy" at its best.

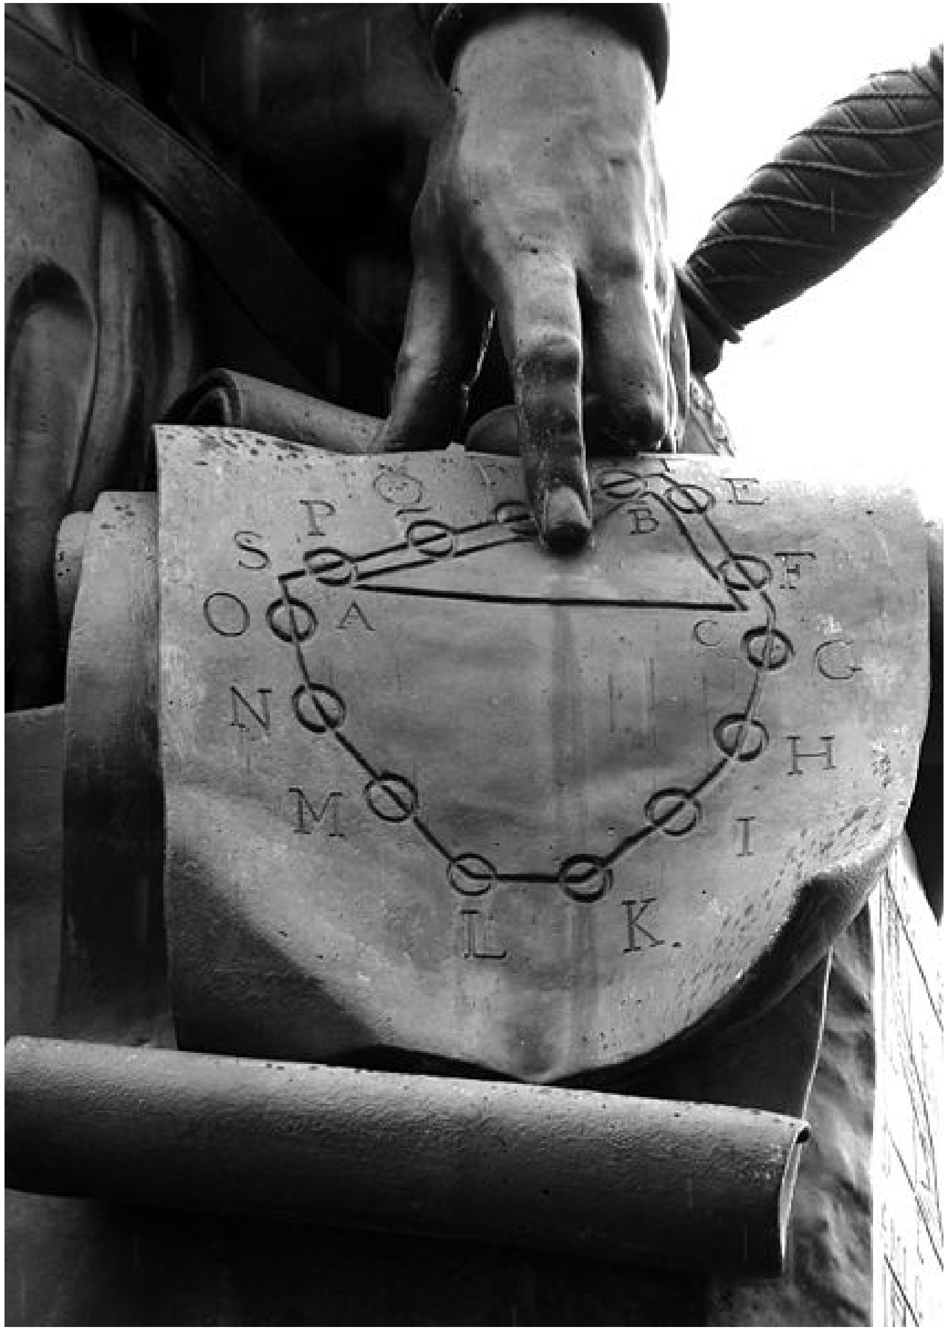
S
P
B
E
F
O
A
C
G
N
H
M
I
L
K