A Bouquet of Exotic Geometries

Douglas R. Hofstadter
Center for Research on Concepts & Cognition
Indiana University • 510 North Fess Street
Bloomington, Indiana 47408

July, 1993

The purpose of this article is to relate a number of geometrical explorations I have recently made as a consequence of a kind of personal obsession with the notion of point/line duality, as well as a fascination with gnomonic projection. These explorations have yielded new (for me) ways of understanding hyperbolic and related geometries via the sphere, and several new "hybrid" types of geometry, described below in some detail.

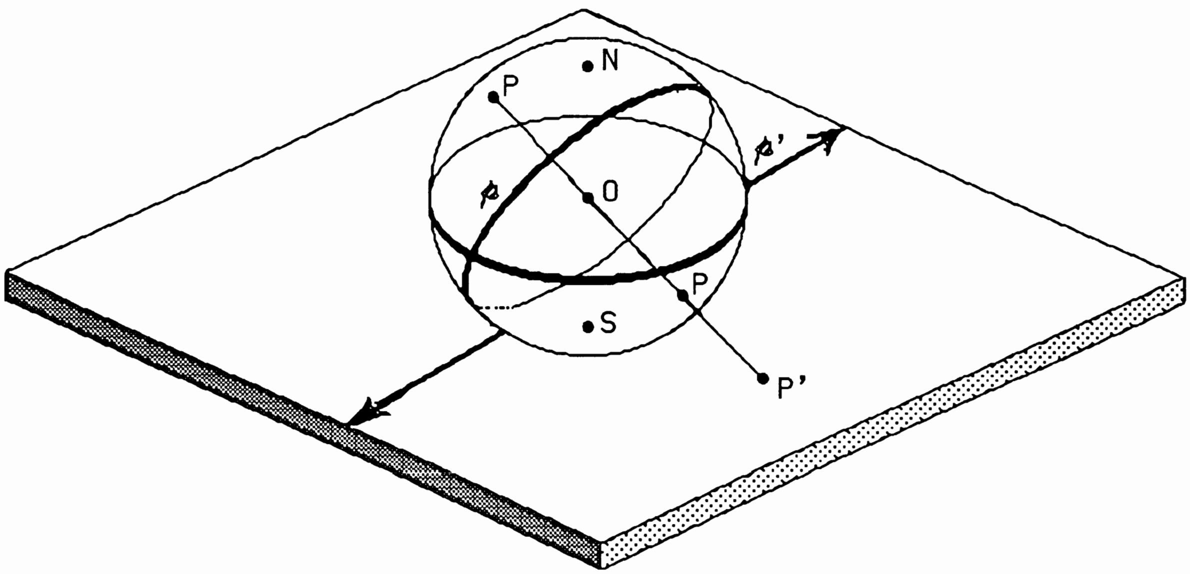

Gnomonic projection and the absolute polarity. If a point on the plane is joined to the center of the sphere by a line, the line hits the sphere in two antipodal points, called the "gnomonic image" of the point. Thus here, the gnomonic image of point P' is the antipodal point-pair P. If you take the gnomonic image of all the points on a straight line, you get a great circle. Thus great circle p and planar line p' are each other's images.

There is a natural duality between great circles and antipodal point-pairs, exemplified by the equator and the north/south poles, which are each other's duals. In general, the great circle located midway between two antipodal spherical points is the dual of that point-pair, and vice versa. This relationship is called the "absolute polarity". In the picture, great circle p and antipodal point-pair P are each other's duals; their gnomonic images on the extended Euclidean plane, straight line p' and point P', are each other's duals.

For me, gnomonic projection enormously clarifies the interrelationships among the geometries that I find most interesting. This means essentially *two-dimensional* geometries, including all of the following (plus some new ones, to be described in a few pages): projective, elliptic, Euclidean, and hyperbolic, as well as their duals: projective (self-dual, obviously), elliptic (self-dual), Euclidual, and "hyperbual"; and finally, contrajective (yet another self-dual geometry).

I was first led in this direction, a bit over a year ago, by my fascination with the *absolute polarity* on the surface of a sphere (with antipodal points equated). (I didn't know its name, since I had never heard of polarities at the time.) It was thrilling to me to realize that I could begin with any diagram on the extended Euclidean plane, gnomonically project it upwards onto the sphere, then *dualize* it in the obvious way (namely, by converting great circles into their absolute poles and point-pairs into their absolute polars), and finally project the dualized image back down onto the extended Euclidean plane, thereby creating *a specific transformation of the plane that concretely realized the abstract notion of duality*.

When I first thought of this transformation of the extended Euclidean plane, a bit over a year ago, I supposed it might be original, but of course it turned out simply to be a polarity, carried out via the intermediary of a three-dimensional sphere. In fact, if you recast it purely in terms of planar operations (without any reference to the sphere), it is just "antireciprocation" — that is, reciprocation in some circle, followed (or preceded) by reflection across the circle's center. I was disappointed to find out how very old this idea was (dating from about 1820), when I read the chapter on projective geometry in Coxeter & Greitzer's book *Geometry Revisited,* about a year ago. That was the first time I had ever heard of the notion "polarity".

As I continued pondering this image of what one might call *sphere-mediated polarities,* it occurred to me that you could get a large family of related polarities by *rotating* the sphere either before or after carrying out the dualization step (via the absolute polarity), and then gnomonically reprojecting the rotated-and-dualized image back down onto the plane. (Since any rotation of the sphere preserves great circles as well as all incidence relations on the sphere's surface, throwing in a rotation in this fashion will not disturb any projective properties.) My hope was that this might in fact yield *all* the polarities on the extended Euclidean plane. However, I soon realized that since there are five degrees of freedom in an arbitrary polarity and only three in an arbitrary rotation of the sphere, I had fallen far short of getting all the possible polarities, and consequently I lost most of my interest in the idea.

Some months later, I realized I had overlooked the fact that there are other things (besides rotations) that you can do to the sphere before reprojecting the image down onto the plane. Anything that leaves intact the pattern on the sphere's surface will still yield a perfectly valid polarity. Thus you can *translate the sphere as a whole,* carrying its south pole to any desired spot on the plane; you can also *inflate or deflate the sphere,* and finally you can *reflect the sphere in a mirror,* thus reversing parity. Altogether, these operations, together with rotations, appear to yield *six* degrees of freedom, which are enough — in fact, more than enough — to yield all polarities on the plane. Since I knew that every combination of these operations definitely leads to some polarity or other, I concluded that there must be some redundancy in my counting, so that there are probably in truth just five, not six, degrees of freedom in the family.

Thus I was led to making the following conjecture, which is almost surely correct:

> Gnomonic Conjecture I: *An arbitrary polarity of the projective plane can be made by carrying out a gnomonic projection G onto a sphere, followed by an arbitrary combination of rotations, translations, inflations, and reflections of the sphere, then followed by a "dualization" of the sphere, and finally followed by the inverse gnomonic projection* G^{-1} *back down onto the plane.*

Soon after making this conjecture, I realized that if you omit the dualization step, you produce *collinearities* instead of *correlations* of the plane. I figured that in all likelihood you could get an arbitrary collinearity in exactly the same way. So I made an analogous conjecture:

> Gnomonic Conjecture II: *An arbitrary collinearity of the projective plane can be made by carrying out a gnomonic projection G onto a sphere S, followed by an arbitrary combination of rotations, translations, inflations, and reflections of S, and finally followed by the inverse gnomonic projection* G^{-1}*, carrying S back down onto the plane.*

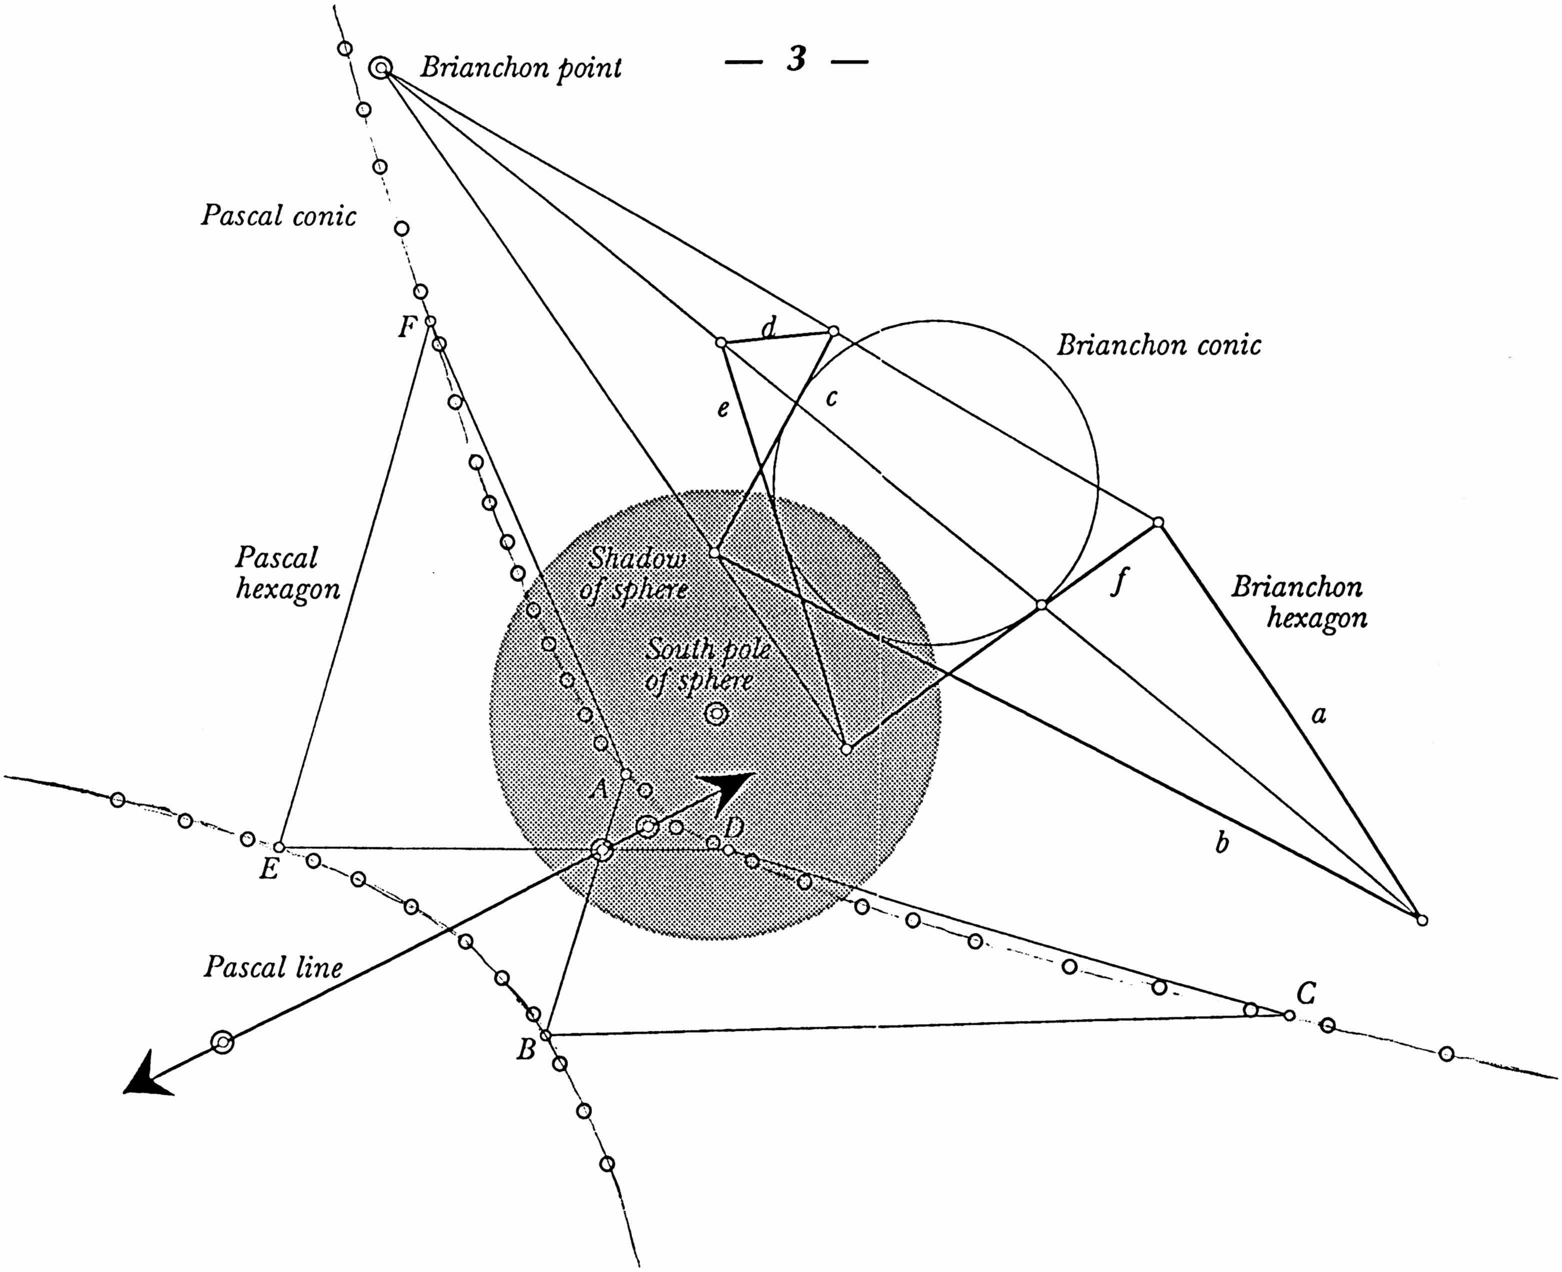

Using the gnomonic projection and the absolute polarity in tandem to realize, in a concrete manner, projective geometry's point–line duality. In the upper right, ***Brianchon's theorem*** *is illustrated:*

> *If a hexagon is circumscribed about a conic, the three lines that link opposing vertices of the hexagon are concurrent.*

The "Brianchon conic" is a perfect circle, but that was done merely for ease of drawing. The hexagon, with sides labeled "a" through "f", looks strange because it is self-intersecting, but it is no less a hexagon for that. The special point of concurrency is labeled "Brianchon point".

This entire configuration is then "dualized" by sending it up onto a sphere (whose shadow is shown) by a gnomonic projection and then using the absolute polarity on the sphere to map antipodal point-pairs to great circles and vice versa. This done, the dualized image on the sphere is sent back down onto the plane via the gnomonic projection. The resulting "gnomonic–dual image" of line a is thus point A, and so on. The Pascal conic (the gnomonic dual of the Brianchon conic) is a hyperbola, and the new planar configuration in its entirety illustrates ***Pascal's theorem:***

> *If a hexagon is inscribed in a conic, the three points where opposing sides of the hexagon meet are collinear.*

And finally, to complete the perfect symmetry, the image of the Brianchon point is the Pascal line.

What would seem to be a reasonable special case of this would be a one-dimensional version, namely this:

Gnomonic Conjecture III: *An arbitrary one-dimensional projectivity of any line on the projective plane can be made by carrying out a gnomonic projection G onto a sphere S, followed by an arbitrary combination of rotations, translations, inflations, and reflections of S, so long as each one of them preserves the image of the line, and finally followed by the inverse gnomonic projection G^{-1}, carrying S back down onto the plane.*

* * *

All this was most exciting and quite beautiful, and kept alive my fascination with the "gnomonic approach" to visualization of projective geometry. The next thing that happened was that I was musing about the following two facts at about the same time:

Fact I: *Suppose we have been given a line and a point not on it. Then:*

- *in **elliptic** geometry, there are **zero** parallel lines through the point;*
- *in **Euclidean** geometry, there is **one** parallel line through the point;*
- *in **hyperbolic** geometry, there are **infinitely many** parallel (i.e., non-intersecting) lines through the point.*

Fact II: *When you make a gnomonic projection from a sphere onto the extended Euclidean plane or vice versa, then:*

- *in order to model the **elliptic** plane, you delete **zero** lines;*
- *in order to model the **Euclidean** plane, you delete **one** line.*

The suggestive resonance between these two facts led me to wondering if there wasn't a line missing from Fact II, whose addition would make it fully analogous to Fact I. That third line, of course, would have to look this way:

- *in order to model the **hyperbolic** plane, you delete **infinitely many** lines.*

Ironically, I was adding a line that spoke of deleting infinitely many lines! This suggestion seemed extremely appealing and eminently reasonable, but there was one question it left unanswered: *Which infinite set of lines to delete?*

Since the line usually deleted to give Euclidean geometry is the equator, the image that instinctively leapt to my mind was that of a "fat equator" — a symmetrical *band* around the equator, delimited by two horizontal lines of latitude at equal angles above and below the equator. If we imagine this "fat equator" as being black and the rest of the sphere as being transparent, then when you project the sphere downwards gnomonically onto a white plane, what you get is a *white circle surrounded by blackness extending to infinity.*

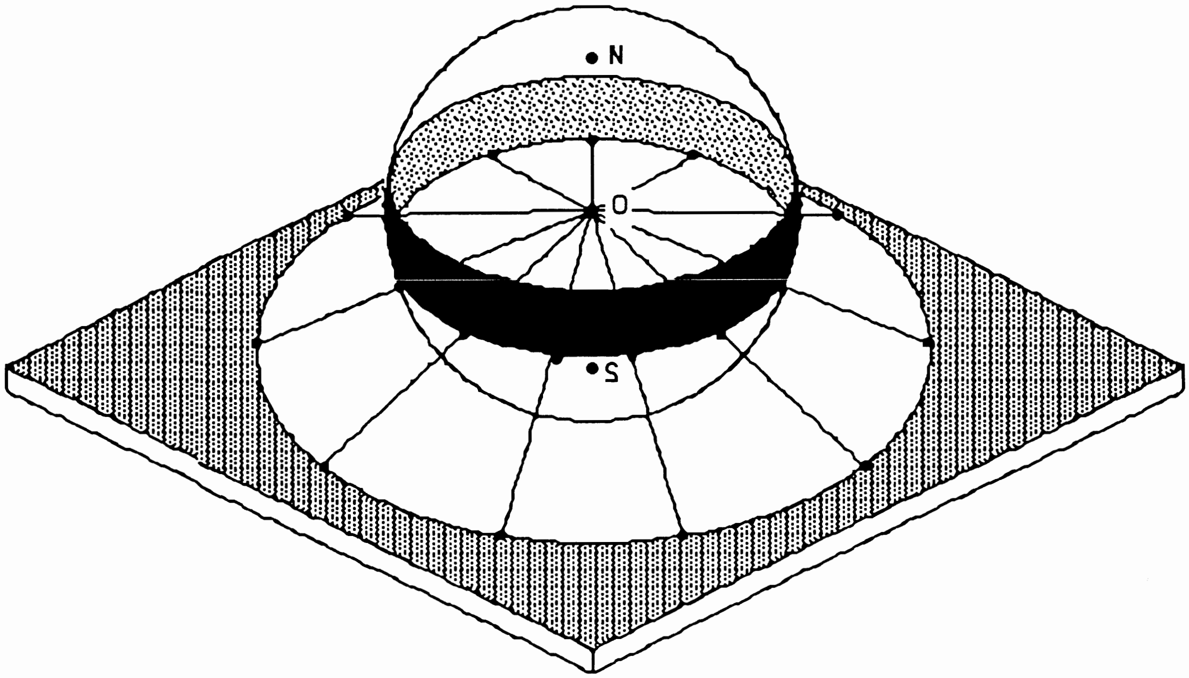

A sphere with a fat equator has been projected gnomonically onto the extended Euclidean plane. The portion of the sphere above and below the fat equator maps onto a white disk centered on the sphere's south pole. The planar image of the fat equator itself maps onto the black exterior region, stretching in all directions out to infinity.

This at first made me think of the Poincaré model of hyperbolic geometry, but on reflection I concluded I was probably dealing with the *Beltrami–Klein* model, since gnomonically-projected great circles become straight lines (not circles) on the plane.

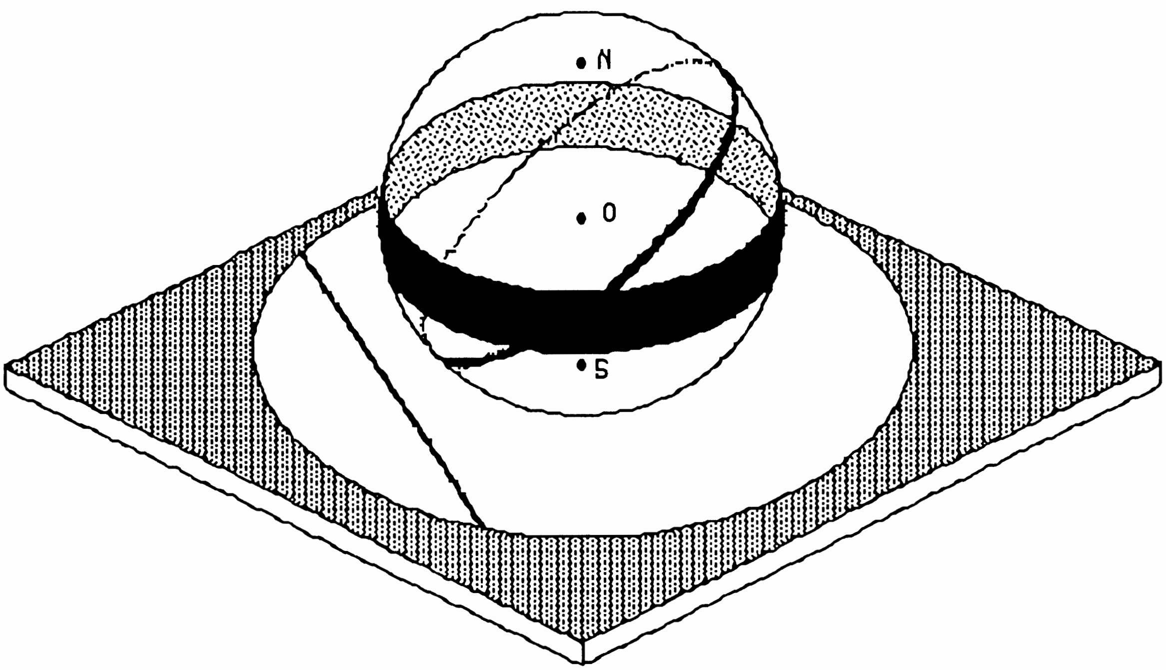

The nondeleted portion of a great circle projects gnomonically downwards onto a chord of the circle delimiting the "legal" portion of the plane.

A little more thought added clarity concerning the nature of the "fat equator".

Although at first one might tend to conceive of it as consisting of an infinite set of *lines of latitude,* that would be a bad way of looking at it. Rather, the fat equator should be conceived of as *an infinite set of great circles*: namely, those that you get by *tilting the equator* through all angles less than or equal to some particular fixed angle. Of course, these deleted lines (or great circles, depending on whether you're looking at the plane or the sphere) are the lines that define parallelism, in the sense that any two lines that meet on any deleted line are said to be parallel. And the bottom and top edges of the fat equator play a special role, in that any two lines that intersect on them are *asymptotic* or *bounding* parallels. From now on, think of these upper and lower edges of the fat equator as colored *red.* Thus, two red circles divide the black (forbidden) zone from the white (permitted) zone on the sphere. When projected down onto the plane, they yield a single red circle separating the white disk from the infinite black zone.

Then I realized that the two red-colored edges of the fat equator on the sphere are really the gnomonic images of the conic at infinity — that is, the conic usually referred to as *the Absolute.* In this case, the Absolute is a circle on the plane — the limit circle of the Klein model, of course. But then I remembered that in its full generality, the Absolute is an *arbitrary conic,* not just the special case of a circle. No problem: you can get an arbitrary conic by gnomonic projection of the fat equator. All you need to do is, as before, first *rotate* the sphere, *translate* it, *inflate* it, *reflect* it if you like. By a suitable choice of such operations followed by a downward gnomonic projection, you can make the fat equator's two red edges map onto any desired conic on the plane.

To summarize what I have said so far, then, my "fat equator" idea really seemed to be just a new way to think about the concept of the Absolute, mediated once again by the sphere, along with gnomonic projection and the set of all operators that act on the sphere as a whole but leave its surface-pattern undisturbed. I had never seen this idea spelled out before, but it didn't seem as if it could be entirely novel. Still, it made much clearer for me a whole lot of things about the relationships among projective, elliptic, Euclidean, and hyperbolic geometries, leading me to wonder why it is so ignored.

One small thing, for instance, that I had read in books but never understood before was the claim that in Euclidean geometry, the Absolute is not *one* line, but two *coincident* lines. It is often stated that this is a degenerate case of a conic; this would be easier to visualize if it were located somewhere in the middle of the plane, but for me it was hard to conceive of the line at infinity as *two* lines. But now, all I needed to do to see this was to imagine making the fat equator thinner and thinner. As its thickness approaches zero, its two red-colored edges — lines of latitude in opposite hemispheres located symmetrically with respect to the equator — eventually come to coincide on the equator itself. For me, this provided a very satisfying intuition for the statement that the line at infinity is really two coincident lines.

* * *

Having gotten this far, I then naturally turned to my idea of *dual* geometries. Thus, I thought, hyperbual geometry — the dual of hyperbolic geometry — must surely come from deleting the dual of the fat equator on the sphere. But what *is* the fat equator's dual? The absolute polarity makes the answer quite obvious: it is the interior of a circular zone surrounding the North and South poles, a zone one might call the *polar caps,* or the *ice caps.* So now imagine deleting the polar caps (colored blue, with their edges being green circles, say) from the sphere (transparent) and then making a gnomonic projection downwards onto the (white) plane. What you get is a green circle (the Absolute) whose interior is blue (this is the forbidden zone, corresponding to the blue polar caps), and which is surrounded by an infinite expanse of white. In short, what this produces is the planar complement of the Beltrami–Klein model.

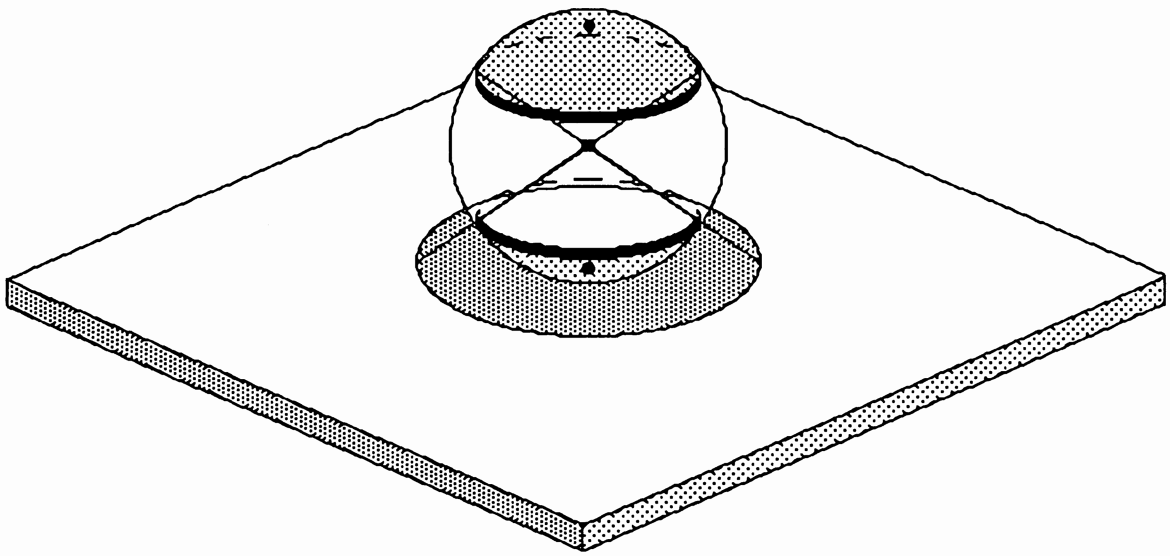

A sphere with deleted polar caps has been projected gnomonically onto the plane. The deleted caps map onto a small shaded disk centered on the sphere's south pole. The image of what remains is the white exterior region, stretching in all directions out to infinity.

Now since we are dealing with a *dual* geometry — the hyperbolic analogue of Euclidual geometry— this curious plane, just like the Euclidual plane, should be conceived of as consisting of *lines,* not of points, and the blue disk as the set of *deleted points at infinity.* Just as in Euclidual geometry, deletion of a point really means deletion of all *lines* that pass through it, so in fact this hyperbual plane is quite impoverished, in that it contains no line passing through the blue disk. Fortunately, however, there are lots of lines left over. Obviously, if you carry this whole procedure through, you will construct all of hyperbual geometry, exactly in parallel fashion to the way I constructed Euclidual geometry. The only difference is that you have deleted a circular *neighborhood* of points rather than a *single* point.

Needless to say, you can now rotate and translate (etc.) the sphere, and the deleted blue area on the plane (with its green-colored Absolute boundary) will change its shape, location, and size accordingly, thus giving you all the different possible realizations on the plane of hyperbual geometry. The only difference between these planar models is that they have differently-shaped Absolutes, but of course the geometry that takes place within the confines of the Absolute is an invariant.

*   *   *

Here is a summary of the geometries I have mentioned so far, all characterized as resulting from the projective plane as realized on the surface of a sphere, with antipodal points equated, and various portions of the sphere possibly deleted.

1. *Projective* geometry, gotten by deleting *no* lines and *no* points from the sphere. This nonmetric geometry is its own dual.
   Size of a line: —. Number of parallel lines thru a point: 0.
   Size of a point: —. Number of parallel points on a line: 0.

1a. *Elliptic* geometry, gotten by imposing the natural projective metric onto projective geometry, in which a line and a point both acquire size 2π. This metric geometry is its own dual.
   Size of a line: 2π. Number of parallel lines thru a point: 0.
   Size of a point: 2π. Number of parallel points on a line: 0.

2. *Affine* geometry, gotten by deleting the equator (one line) from the sphere. The dual of this semi-metric geometry is iffine geometry.
   Size of a line: ∞. Number of parallel lines thru a point: 1.
   Size of a point: —. Number of parallel points on a line: 0.

2a. *Euclidean* geometry, gotten from affine geometry by imposing an absolute involution of points on the equator, defining perpendicular lines. The dual of this metric geometry is Euclidual geometry.
   Size of a line: ∞. Number of parallel lines thru a point: 1.
   Size of a point: 2π. Number of parallel points on a line: 0.

2′. *Iffine* geometry, gotten by deleting the north/south pole (one point) from the sphere. The dual of this semi-metric geometry is affine geometry.
   Size of a line: —. Number of parallel lines thru a point: 0.
   Size of a point: ∞. Number of parallel points on a line: 1.

2a′. *Euclidual* geometry, gotten from iffine geometry by imposing an absolute involution of lines on the north/south pole, defining perpendicular points. The dual of this metric geometry is Euclidean geometry.
   Size of a line: 2π. Number of parallel lines thru a point: 0.
   Size of a point: ∞. Number of parallel points on a line: 1.

3. *Hyperbolic* geometry, gotten by deleting a "fat equator" (infinitely many lines) from the sphere. The dual of this metric geometry is hyperbual geometry.
   Size of a line: ∞. Number of parallel lines thru a point: ∞.
   Size of a point: 2π. Number of parallel points on a line: 0.

3′. *Hyperbual* geometry, gotten by deleting "polar caps" (infinitely many points) from the sphere. The dual of this metric geometry is hyperbolic geometry.
   Size of a line: 2π. Number of parallel lines thru a point: 0.
   Size of a point: ∞. Number of parallel points on a line: ∞.

4. *Contrajective* geometry, gotten by deleting both the equator (one line) and the north/south pole (one point) from the sphere. This nonmetric geometry is its own dual.
   Size of a line: —. Number of parallel lines thru a point: 1.
   Size of a point: —. Number of parallel points on a line: 1.

4a. *Ecliptic* geometry, gotten by imposing the natural projective metric onto contrajective geometry, in which a line and a point both acquire size ∞. This metric geometry is its own dual.
   Size of a line: ∞. Number of parallel lines thru a point: 1.
   Size of a point: ∞. Number of parallel points on a line: 1.

Note that the latter two are *hybrid* geometries, in the sense that they involve deleting two different types of things at once. The resulting geometry in some sense inherits different traits from its "mother" and "father" geometries — specifically, it inherits its *points* from its mother and its *lines* from its father.

In the case of contrajective geometry, the mother and father geometries are iffine and affine geometry, respectively, which means that *contrajective geometry has iffine points and affine lines.* (And projective geometry has affine points and iffine lines.)

Similarly, the mother and father of ecliptic geometry are Euclidual and Euclidean geometry, respectively, which means that *ecliptic geometry has Euclidual points and Euclidean lines.* (And elliptic geometry has Euclidean points and Euclidual lines.)

All this is pretty crazy, or at least mind-boggling — but as if this weren't enough, I then realized that many other hybrid possibilities were opened up by the idea of deleting fat equators and polar caps. Here is a list:

5. *Hyperb-Euclidean* geometry, gotten by deleting the equator (one line) and polar caps (infinitely many points) from the sphere. This hybrid geometry's mother is hyperbual geometry, and its father is Euclidean geometry. It therefore has *hyperbual points* and *Euclidean lines*. The dual of this metric geometry is hyperb-Euclidual geometry.
   Size of a line: ∞. Number of parallel lines thru a point: 1.
   Size of a point: ∞. Number of parallel points on a line: ∞.

5′. *Hyperb-Euclidual* geometry, gotten by deleting a fat equator (infinitely many lines) and the north/south pole (one point) from the sphere. This hybrid geometry's mother is Euclidual geometry, and its father is hyperbolic geometry. It therefore has *Euclidual points* and *hyperbolic lines*. The dual of this metric geometry is hyperb-Euclidean geometry.
   Size of a line: ∞. Number of parallel lines thru a point: ∞.
   Size of a point: ∞. Number of parallel points on a line: 1.

6. *Subjective* geometry, gotten by deleting a fat equator (infinitely many lines) and polar caps (infinitely many points) from the sphere. This hybrid geometry's mother is hyperbual geometry, and its father is hyperbolic geometry. It therefore has *hyperbual points* and *hyperbolic lines*, but it has no projective metric. This nonmetric geometry is its own dual.
   Size of a line: —. Number of parallel lines thru a point: ∞.
   Size of a point: —. Number of parallel points on a line: ∞.

6a. *Subtropical* geometry, gotten by imposing the natural projective metric onto subjective geometry, in which a line and a point both acquire size ∞. This hybrid geometry's mother is hyperbual geometry, and its father is hyperbolic geometry. It therefore has *hyperbual points* and *hyperbolic lines*. This metric geometry is its own dual.
   Size of a line: ∞. Number of parallel lines thru a point: ∞.
   Size of a point: ∞. Number of parallel points on a line: ∞.

At some stage, I started realizing that the affine/iffine family, whose members are characterized by having blank entries in their descriptive tables (somewhat reminiscent of defective verbs, such as the verb *falloir* in French, whose conjugations lack one or more entries), was far from complete. The way I calculated it, there are twelve of these semi-metric geometries *in toto.*

A couple more of them are described below:

7. *Hyperaffine* geometry, gotten by deleting a fat equator (infinitely many lines) from the sphere, but leaving out the concept of perpendicularity. The dual of this semi-metric geometry is hyperiffine geometry.
   Size of a line: ∞. Number of parallel lines thru a point: ∞.
   Size of a point: —. Number of parallel points on a line: 0.

7′. *Hyperiffine* geometry, gotten by deleting polar caps (infinitely many points) from the sphere, but leaving out the concept of perpendicularity. The dual of this semi-metric geometry is hyperaffine geometry.
   Size of a line: —. Number of parallel lines thru a point: 0.
   Size of a point: ∞. Number of parallel points on a line: ∞.

9 fully metric geometries

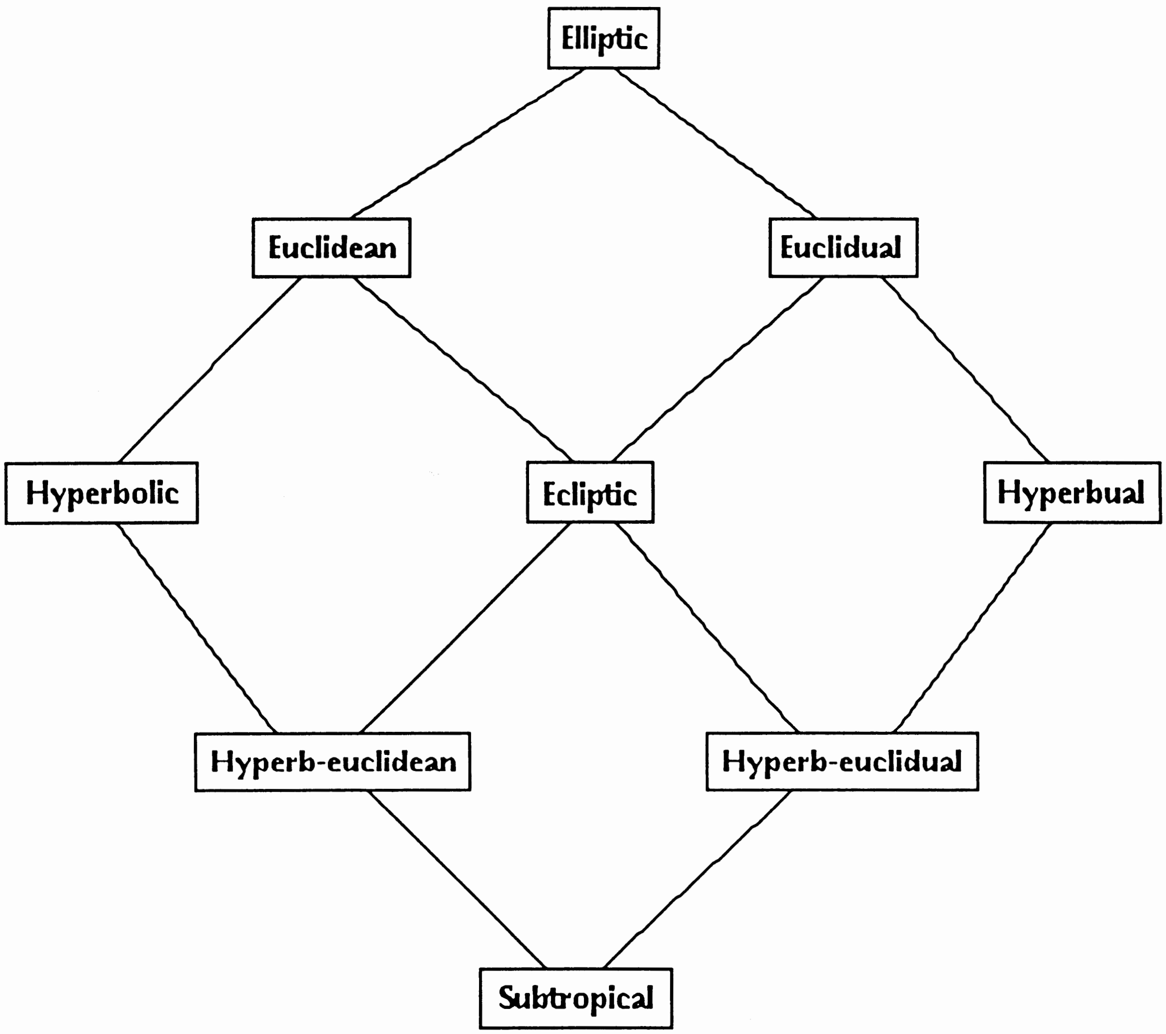

9 fully metric geometries,
12 semimetric geometries,
and 3 nonmetric geometries

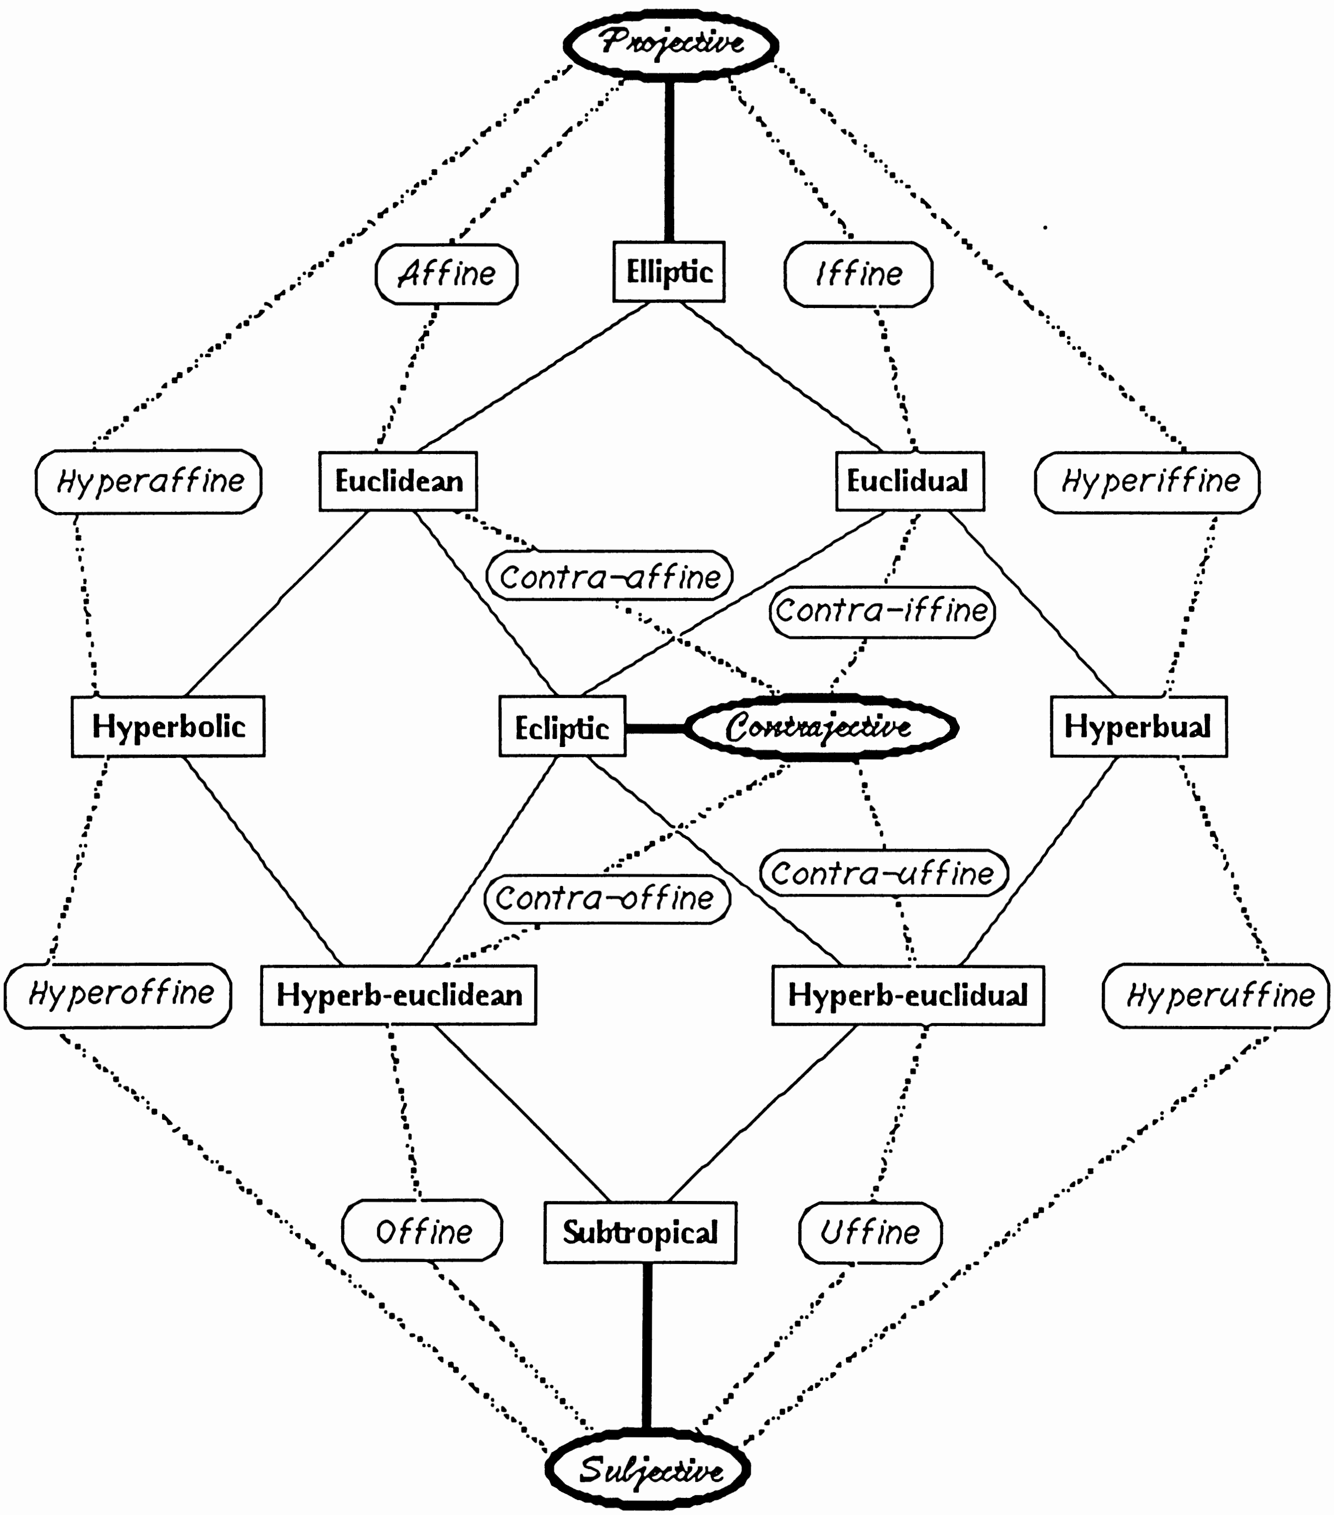

There are also *offine* and *uffine* geometries, which are to subtropical geometry what affine and iffine are to projective geometry. Then there are also *hyperoffine* and *hyperuffine* geometries, and four more. All twelve of them are shown on the accompanying chart.

All told, including the twelve semi-metric geometries, there seem to be 24 natural geometries that come out of canonical ways of deleting points, lines, "fat points", or "fat lines", either alone or in simple combinations. If you leave out the semi-metric geometries, there are still twelve geometries — too many to even *name,* let alone *understand*!

* * *

There is no compelling reason, other than elegance, to make the two deleted regions in a hybrid geometry be centered on locations that are each other's absolute duals. For example, in making contrajective geometry (which involves deleting one line and one point), if you delete the equator, you are not thereby compelled to delete the north/south pole in particular — you can delete any point whatsoever on the sphere. Of course, there are only two "essentially different" types of points to delete — ones off the equator, and ones on it. Therefore, contrajective geometry and its metric partner, ecliptic geometry, divide into two cases: deleted point and line *incident,* versus deleted point and line *non-incident.*

Because the deleted line functions on the plane essentially as a *horizon,* I choose to use that name for it. And, following the convention used in perspective drawing, I call any point on the horizon a *vanishing point.* (Of course, vanishing points are not real points because they lie on the deleted line.) Similarly, because the deleted point functions on the plane essentially like the *sun* (you cannot look directly at it, and the closer you come to doing so, the harder it gets), I choose to use that name for it. And not following any prior convention, at least to my knowledge, I call any line that emanates from the sun a *sunbeam.* Therefore, contrajective geometry's two cases are where *the sun is above the horizon,* and when *the sun is exactly on the horizon.* For short, therefore, I refer to them as the "midday" and "sunset" versions of contrajective (or ecliptic) geometry.

I know almost nothing about contrajective or ecliptic geometry, but the one basic fact I do know is shown in the accompanying figure. It is simply the fact that to each pencil of parallel lines, there is a naturally associated range of parallel points, and vice versa. More simply put, there is a *one-to-one correspondence between vanishing points and sunbeams.* This seems to be a fundamental, central fact about such geometries, but what it really means or implies, I don't yet know. I feel like an ignoramus, in fact.

* * *

Just as contrajective geometry subdivides into two cases depending on the relative locations of the two deleted entities, so *all* hybrid geometries subdivide. Some of them split into three cases, and some of them into a whole bunch. Take hyperb-Euclidual, for instance, which involves deleting a fat line (which we might as well take to be centered on the equator) and one point (which we will consider mobile). There are three essentially different positions for the point relative to the fat equator: *outside* the fat equator (*i.e.* in the white zone), *on either edge* of the fat equator (*i.e.,* on either red curve), and *inside* the fat equator (*i.e.,* in the black zone). More conventionally put, the deleted point can fall either in the *real* region of the plane, or on the *Absolute,* or in the *ultra-ideal* region.

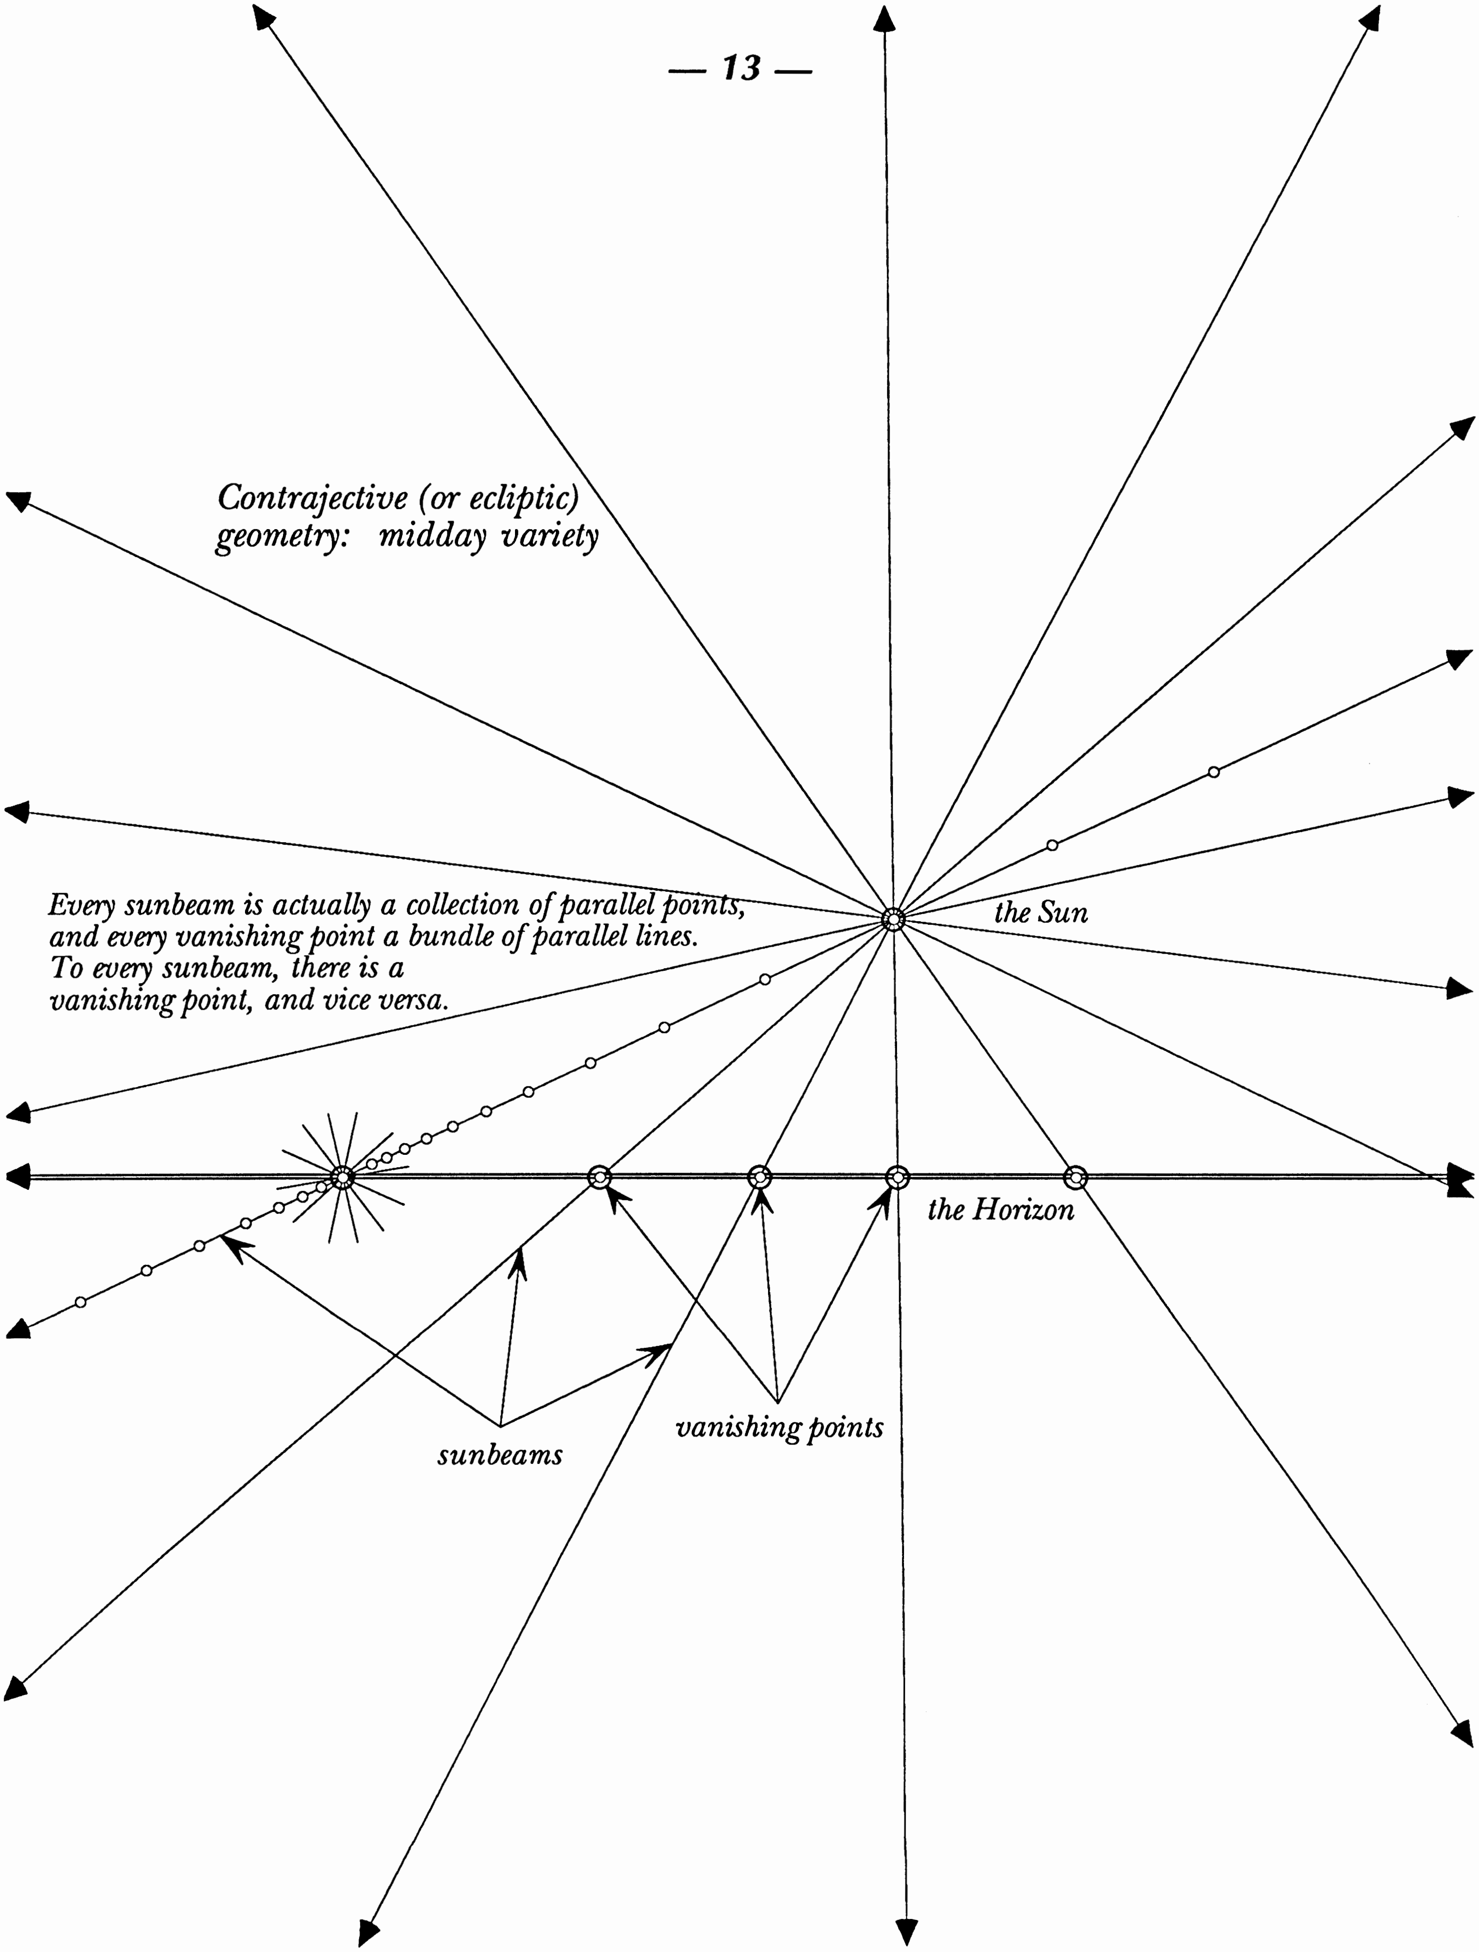
Contrajective (or ecliptic) geometry: midday variety
Every sunbeam is actually a collection of parallel points, and every vanishing point a bundle of parallel lines.
To every sunbeam, there is a vanishing point, and vice versa.
the Sun
the Horizon
sunbeams
vanishing points

Actually, mentioning the Absolute brings up a subtle point: in any hybrid geometry, there really are *two* Absolutes: a *Point Absolute,* which is a *locus,* and a *Line Absolute,* which is an *envelope.* The Point Absolute is the boundary of the deleted (simple or fat) line, curiously, and the Line Absolute is the boundary of the deleted (simple or fat) point. The reason for this reversal is that deleting a point really means declaring all the lines that belong to that point to be ideal — that is, to consist of parallel points. So the operation of deleting a point is really more about lines, and conversely, deleting a line is really more about points.

Let's jump to the messiest case by far, which is subtropical (or subjective) geometry. Here we are deleting both a pair of antipodal caps *and* the fat equator. Where can the caps lie, relative to the fat equator? The easiest case is when there is no overlap at all. Then you have the most opposite case, which is when one is contained entirely in the other. There are actually two such cases, dual to each other. It's easy enough to envision two antipodal caps sliding around in perfect synchrony on the surface of the globe and sometimes being contained entirely inside the fat equator, assuming they are small enough, but what about the reverse case? How could the fat equator ever be contained inside a pair of caps, no matter how big they are?

The trick is, one has to remember that in a self-dual geometry like this, both the fat equator and the caps can be thought of either as a collection of ultra-ideal *points* or as a collection of ultra-ideal *lines.* The intuitive image of caps contained inside the fat equator, given above, is basically a *point-oriented* image. To see the fat equator as contained inside the caps, however, we have to switch over to a *line-oriented* point of view (so to speak). In this view, the caps, rather than being merely a set of points, consist in the set of all *lines* that pass through those points (if the caps are small, these lines are nearly meridians, for instance). Thus we imagine the caps sliding around the sphere and, if they are wide enough, there will be certain positions when the lines they define include all the lines comprising the fat equator — when this happens, then the fat equator is inside the caps. From a point-oriented stance, this happens exactly when the fat equator is thin enough, and oriented in such a way, that it passes through the caps, like a multilane highway passing through two antipodal cities.

From now on, since every case has a dual, I'll just describe the far-easier-to-visualize point-oriented cases. And since the caps are antipodal, I'll just talk about one of them. So, to begin with, we can easily imagine a cap as being *externally tangent* to the fat equator, or as being *internally tangent.* Internal tangency subdivides again, depending on whether the cap is wider or narrower than the fat equator. We can even imagine a cap as being *doubly internally tangent* to the equator— namely, when it is completely inside the fat equator but has the same width as it does. By the way, this is a self-dual case, of course.

I think there is another type of double internal tangency. Imagine the fat equator reaching out to, say, the Tropics of Cancer and Capricorn, and the cap being tilted so that it is tangent to both Tropics on their *north* sides. I think this is a valid case.

Another very weird case is where one cap is centered on the north pole, and is very fat — so fat that it descends nearly to the equator itself — and the fat equator is also so fat that it actually *overlaps* this very fat polar cap. Then there are measure-zero cases such as when the caps reach to the Tropics of Cancer and Capricorn and where the fat equator does likewise, so that the Tropics serve simultaneously as the Point Absolute and the Line Absolute. There are probably more cases that I've missed, but this certainly gives the idea. It's a huge, ridiculous zoo of geometries!

* * *

This act of cataloguing geometries *ad nauseam* is kind of fun but makes me also feel very weird. On the one hand, I am at long last doing "high-level mathematics" — a kind of abstract level of mathematics that feels far more sophisticated, in some sense, than the mere act of discovering various curious concurrences of lines in triangle geometry on the Euclidean plane. I feel I'm almost being "modern" — almost doing "twentieth-century" mathematics. On the other hand, I feel ridiculously out of touch with all the wonderful, specific, idiosyncratic geometrical phenomena that must be going on in *every single one* of these individual geometries.

I at least got a *taste* of Euclidual geometry, and I have at least tried a tiny bit to think about contrajective geometry, and hope to continue — but what about all the rest? I don't know Word One about a single one of them, and probably, as far as most of them are concerned, never will. It's sad. It really reminds me of those linguists who study African languages and who make grand classificatory schemes, grouping them into large families and subfamilies and plotting interrelationships and historical tendencies and so on, all while knowing not one word of any of them. It may seem grand, but on another level it is very remote from the phenomenon of language itself.

Somehow, in mathematics as in linguistics, there must be a happy medium. I don't want to just manipulate geometries as if they were opaque little capsules — I want to peer *inside* those capsules, and learn *specifics.* Of course, I can't hope to learn a great deal about each of these dozens of geometries, any more than I can hope to learn dozens of foreign languages, but I would certainly like a *glimpse* inside each capsule, in the same way as I can recognize Vietnamese and Cantonese, have a smattering of Russian and Hindi, and know a *few* languages well enough to converse in them.

Plus, it strikes me that this kind of geometry-manipulation, while fun for a short time, becomes a nearly empty shell-game quite soon. There is nothing really *gripping,* nothing really strange or exciting or beautiful going on at these ethereal levels of abstraction. The really beautiful phenomena are lower down, at a more visualizable level. And the really deep *ideas,* too, all spring out of lower-level phenomena, or at least so I deeply believe.

Maybe finding a deep idea quickly sends you soaring up into the stratosphere of abstraction, but whatever you are doing up there remains rooted way down on some very low level, and that's where its power comes from. If you forget about that origin and remain stuck (and possibly stuck-up) up in the oxygen-poor, image-poor stratosphere of abstraction, you will probably wind up eventually just counting angels on the head of a pin, metaphorically speaking. Only if you occcasionally dip back down to the low, oxygen-rich layers of the atmosphere will you refresh yourself and renew your excitement with down-to-earth, genuinely unpredictable, weird, exotic phenomena that you could never have dreamt of from your remote perch in the abstract sky. Thus renewed, you may well shoot back up to the stratosphere, but at least you will be in touch with the doings on earth, and your inspiration will be coming, in some sense, from the right place.

* * *

Several weeks ago, Donald Coxeter wrote to me that my earlier article "A Diamond of Geometries" reminded him of something in Boris Rosenfeld's *History of Non-Euclidean Geometry,* so I immediately rushed out to the library, checked it out, and checked it out. On first perusal, I failed to spot anything that seemed related, but a clever friend of mine decided to look up "Sommerville" in the index, and that led him to Rosenfeld's discussion of Klein's, Cayley's, and Sommerville's *projective metrics.* I quickly recognized that this was indeed related to my investigations of Euclidual geometry and the rest of the "bouquet of geometries" described herein, but the language was so different from mine that I had a rough time figuring out what the exact overlap is.

Whereas I characterize a geometry by specifying what gets deleted from a projective plane (one line or one point, or possibly one of each), Rosenfeld talks about the values of the coefficients K and k that figure into the cross-ratio-based formulas for determining what I call *line-sep* (the generalization of *angle* that most naturally fits the geometry in question) and *point-sep* (the generalization of *distance* that most naturally fits the geometry in question). Each of these coefficients K and k can take on only three "essentially different" values: real, imaginary, or infinite (when K or k is finite, its precise value makes no difference), and since there are two such coefficients, there are nine total combinations, leading to nine "essentially different" geometries. Actually, Rosenfeld has *three* constants, K, k, and κ, because he is dealing with three-dimensional rather than two-dimensional geometries, so there is a similar cross-ratio-based formula for *plane-sep* as well. As a consequence, in this scheme, there are 27 different 3-D geometries.

It was quickly apparent to me that Sommerville's "co-Euclidean" space $^*R_3$ (the nomenclature and notation are Rosenfeld's, not Sommerville's, I would guess) is a 3-D version of my "Euclidual plane". On the other hand, it was not at all clear whether my "contrajective geometry" (what you get when you simultaneously delete *both* a point *and* a line from the projective plane) fits into this projective-metric scheme, in the sense of being the geometry yielded by some particular choice of Sommerville's parameters K and k.

Obviously the projective metrics that Klein, Cayley, and Sommerville studied and that Rosenfeld discusses include all sorts of weird geometries, but I couldn't make head or tail of most of them. I had no idea what *quasielliptic* or *quasihyperbolic* spaces are, or what *pseudoisotropic* spaces or *co-Galilean* spaces are, or what *flag space* is. (The definition of flags on Rosenfeld's p. 260 — "manifolds of planes whose intersections with a nested system of fixed planes have prescribed dimensions" — left me completely in the dark.)

It's hard for me to imagine that Sommerville, when he came up with his "Big 27", wouldn't have experienced an *embarras du choix* similar to the one I experienced. I mean, who has time to fully explore 27 different geometries? Life is finite, after all! That's why I wonder if he — or anyone else — ever really looked deeply into, say, Euclidual geometry, in the sense of *finding Euclidean models* of it (as I do in my "Diamond of Geometries" article) or carefully *describing* it. Has anyone besides myself ever graphically discussed the weird, disorienting properties that such a space has? By this, I mean explicitly bringing up the existence of parallel points and perpendicular points, the lack of parallel lines, the fact that all lines have finite length, the fact that points are "made out of" lines, so to speak, rather than the reverse, and so on — all the things that I spelled out in such detail in my article "A Diamond of Geometries".

I guess that in some rough sense, that article attempts to do for Euclidual geometry (or as Rosenfeld would have it, "co-Euclidean" geometry) what *Flatland* did for four-dimensional geometry, in the sense that it tries to spell out, in the most basic and vivid terms, the truly disorienting nature of this "alternate world".

Post Scriptum

Writing this article was a distinct help to me in analyzing and classifying my own bouquet of geometries. After finishing the first draft of it (it started life out, in fact, as a letter to Donald Coxeter!), I went back to Rosenfeld's discussion of Sommerville's 27, and went very carefully through it. It took quite a while, but eventually I was able to make a mapping between Sommerville's and my geometries. The mapping was certainly complicated by the fact that Sommerville was operating in 3-space while I was working in 2-space, but I eventually came to see that in a certain sense, everything I had devised had already been thought up by Sommerville some 80 years earlier — in roughly 1913. This was a cause for sadness, or at least a sense of letdown.

However, I have a certain kind of consolation. The discussion of Sommerville's stuff took place toward the very end of Rosenfeld's hefty tome (over 400 pages) in a chapter on group theory and non-Euclidean geometry, in a section called "Quasisimple and *k*-Quasisimple Lie Groups". Typical of the lucidity of the prose in that section is this gem of a sentence:

> Quasielliptic and quasihyperbolic spaces $S(m,n)$ and $S(m,n;\ l_0 l_1)$ are special cases of *quasi-Riemannian* and *quasipseudo-Riemannian spaces,* and *k*-quasielliptic and *k*-quasihyperbolic spaces $S(m_0\ m_1 \ldots m_{k-1})$ and $S(m_0\ m_1 \ldots m_{k-1};\ l_0 l_1 \ldots l_k)$ are special cases of *k-quasi-Riemannian* and *k-quasipseudo-Riemannian spaces,* which are spaces with a homogeneous connection in whose tangent spaces there are defined corresponding projective metrics.

If you can make head or tail of this, more power to you!

Ironically, the very opacity of this section provides me with a certain kind of consolation prize. Even if I am not the first person to get here, I feel that I am a bit like Reinhold Messner, the Italian–Tyrolean mountaineer who was the first person to carry out the remarkable feat of climbing Everest without the aid of any oxygen. Not that I consider the ascent of these particular slopes of geometry to be the equivalent of climbing Everest, admittedly, but my analogy is meant to stress that *my* way of getting there was a horsies-and-doggies method — I didn't use any sophisticated concepts whatsoever. Certainly no equations, no fancy group-theoretical ideas, no "quasi-pseudo" fizgigs with homogeneous (or inhomogeneous) lolligaggles — just plain old visual imagery involving "fat equators", "ice caps", the "sun", the "horizon", and so forth. It's about as down to earth as you can get! And yet with my low-falutin' no-oxygen techniques, I still got to the very same place as those high-falutin' mathematicians did, with their piles of subscripts and superscripts. (I didn't mention earlier that Rosenfeld's jungle of superscripts with subscripts was too complex for my poor little word-processing program to reproduce, so I just used parentheses.)

So my consolation prize is my hunch that my achievement is very different in kind from Sommerville's; it comes from such an intuitive, imagistic source that surely, I must have some different kinds of insights into it all from those he had. At the very least, there must surely be *some* fresh insights to be had from this new, highly intuitive angle of approach, even if I myself haven't yet found them.

In any case, I feel that I'm making a kind of slow catch-up with respect to modern mathematicians. My first rediscoveries turned out to be 170 years old or so; then this new wave of rediscoveries, made about a year later, is about 80 years old — so at this rate, by a bit under a year from now, I ought to be fully caught up with modern mathematicians and making my own discoveries. If this rate continues, then by gosh, in just *two* years from now I will be making discoveries that will only be made 100 years from now! Needless to say, I am very much looking forward to that.

*   *   *
* *
* *
**
*
*