



A Diamond of Geometries:

*The Symmetric Fission of Projective Geometry
into Two Geometries Dual to Each Other,
and Their Fusion into a New Geometry*

*Douglas R. Hofstadter
April–May, 1993*



Introduction: On Symmetry-Breaking and Projective Geometry

It is well known that in projective geometry there is a perfect symmetry between points and lines, known as *duality*. Thus on the projective plane, not only does every pair of points determine a line, but every pair of lines determines a point. This means that there are no parallel lines.

Projective geometry has some unexpected and counterintuitive aspects. For instance, a projective line and a projective point are both “closed” objects. The closure of *points* is a very simple and familiar notion, although “closure” is an unfamiliar name for it. Imagine you have tacked a line to some fixed point on the projective plane; now “ping” the line with your finger so that it spins about the point, like a spinner in a board game. It will of course return to its original position over and over again. This might be called the *rotational* closure of a projective-plane point, and since the same property holds for points on the Euclidean plane as well, it is hardly astounding. In fact, it is hard to imagine how it might have been otherwise. How else could a line tacked to a point behave when pinged?

Were this all there were to the notion of closure, it would not be clear why it deserves mention — it would seem to be a trivial, self-evident, necessary fact. Even granting that this property might deserve an official name, it would still not be clear why the property should be attributed to the *point* in question, rather than to, say, the spinning line, or to point and line taken as a unit.

However, this is *not* all there is to the notion of closure in projective geometry. The essence of the concept emerges far more clearly when one considers the *dual* property — namely, the fact that if one takes any fixed projective-plane line and “pings” a point on it so that it zips along the line, the point will return to its original position over and over again, like a racecar zipping around a closed track. This, despite the fact that projective-plane lines are “straight”. This very counterintuitive property of the projective plane might be called *translational* closure. In this case, it seems natural to attribute the quality of closure to the line itself, rather than to the point zipping along it, or to some point–line unit. After all, it is the track that is closed, not the racecar on it. Similarly, a projective-plane line itself seems to form some kind of loop, and all the point does is simply *reveal* this inherent loopiness. In fact, a projective-plane line does not just *seem* like a loop, it *is* a loop — a completely *straight* loop, to boot (whatever that means).

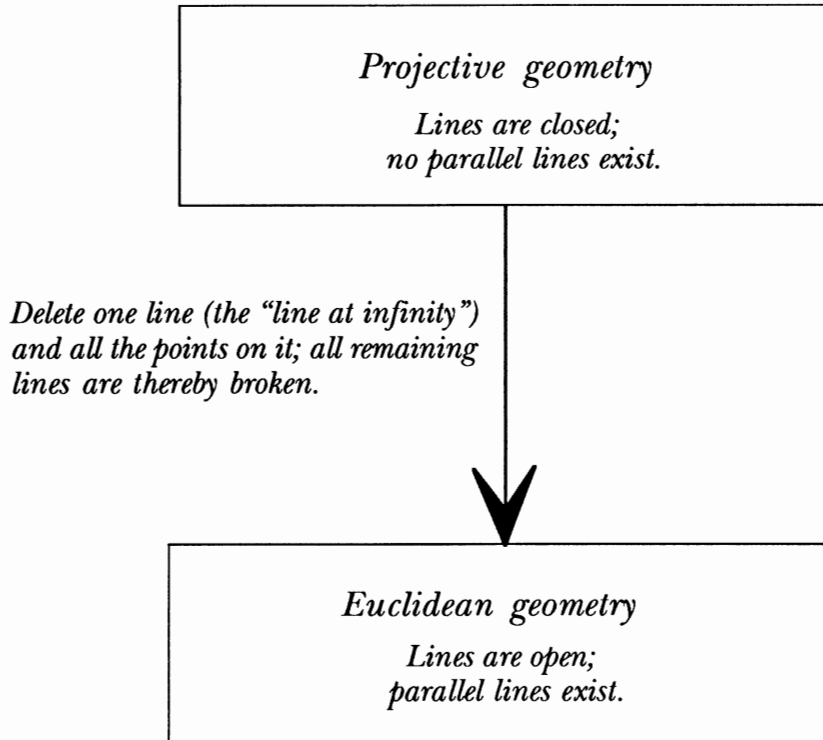
If one thinks long and hard about the duality of projective geometry, one runs into many puzzling questions. For instance, if there is a perfect symmetry between points and lines, then how is it that lines seem to be *made* of points, and not vice versa? How is it that lines seem to be infinite in size, and points infinitesimal (and not the reverse, or something more symmetrical)? When one “pings” a line about a fixed point, the line seems to sweep across the entire projective plane, covering every last bit of space, and yet when one pings a point along a line, the moving point would hardly appear to cover every part of the plane — not even close! How can these seeming asymmetries (and many other similar anomalies) be reconciled with the perfect duality claimed to hold between points and lines?

Such thoughts cannot help but lead one to a reexamination of the intuitions underlying the very notions of “point” and “line”. When one looks carefully, one unearths a host of “Euclidean prejudices” that unconsciously come in and contaminate one’s supposedly neutral imagery. In order to unbury these intuitions, it is very useful to consider the nature of the relationship between Euclidean geometry and projective geometry.

This intimate relationship is usually explained by saying that Euclidean

geometry is what results when a particular line of the projective plane is singled out and called the *line at infinity*. This line, and all the points on it, are considered to be inaccessible or “ideal”. Two lines whose common point lies on the line at infinity are said to be parallel. The upshot of this deletion of one line (and its component points) is a rich new geometry with many features lacking in projective geometry: parallel lines, perpendicular lines, lengths, angles, circles, and so on.

Euclidean geometry could thus be described as arising from projective geometry by destroying the symmetry between points and lines. In fact, removal of the line at infinity somehow “damages” all the remaining lines, in the sense that now all lines are (translationally) *open*. *Points*, however, remain intact — they are still (rotationally) closed, as was described above.



The standard view of the relationship between projective and Euclidean geometry

As one thinks about these matters more and more, one starts to build up a realization that points in the projective plane cannot be imagined in quite the same way as they are imagined in the Euclidean plane. Even the simple act of calling a point “closed” somehow conjures up a swarm of lines passing through the point. Duality eventually forces one into thinking of a point as a complex object, just as complex as a line.

Just as a line consists of an infinite translational set of points, one realizes that a projective point, unlike a Euclidean point, must consist of an *infinite rotational set of lines*. Someone might object and ask, “Why say that a point *consists* of many lines? Why not merely say it is *associated* with an infinite set of lines?” The answer would be that if true conceptual symmetry is to hold, then points must be to lines as lines are to points, and this means that “consists” has to be the proper term.

Of course, this new vision of points as “big” objects leads to a total overthrow of old Euclidean prejudices, and requires deep rebuilding of one’s intuitions. Let us go

back, for instance, to the puzzle of the spinning line covering the full plane whereas the zipping point seems to fail to do so. If one is careful, one realizes that Euclidean prejudices pervade one's sense of *what the projective plane is*. To be concrete, the natural feeling that the spinning line reaches every "part" of space reveals the prejudice that the "parts" of space are *points*. One might counter with the question, "But what else could space consist of?" The answer is: *lines*.

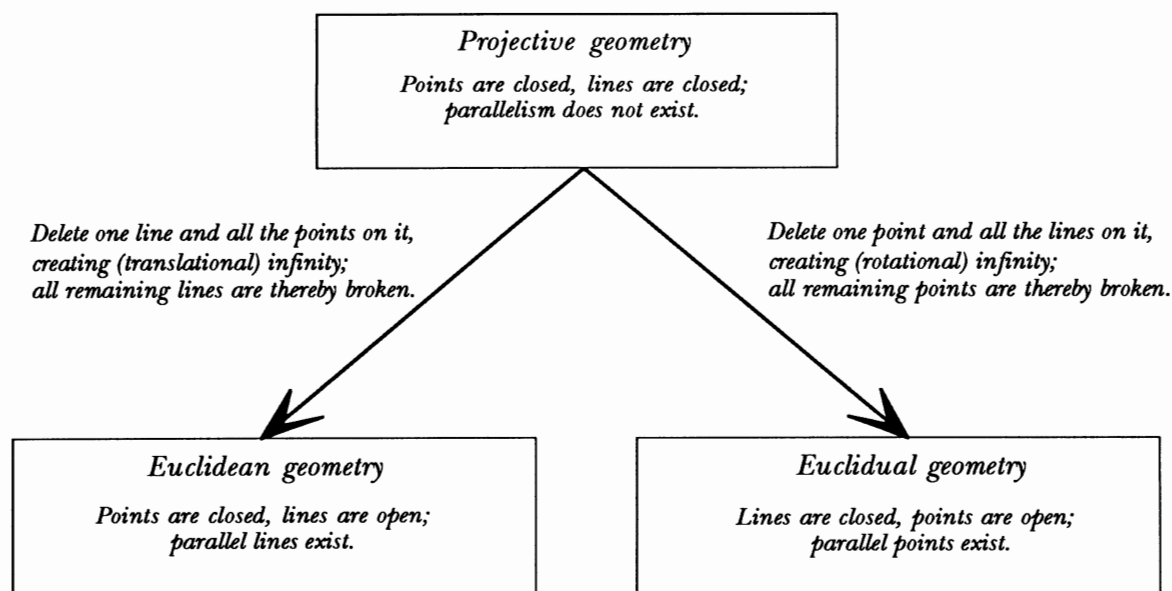
Consider once again a point running along a fixed line. A moment's thought shows that this point, in the course of its travels, is momentarily on *every single line on the plane* (remember, every line in the projective plane crosses any given line — there are no parallels). In that sense, its one-dimensional trip *does* cover the plane entirely. Moreover, consider the spinning line tacked to a fixed point. Although that line undeniably sweeps through every *point* of the plane, it certainly does not come into exact alignment with every *line* of the plane. And so one sees that by enlarging one's perspective on the nature of points and lines, full symmetry is restored, at the price of reconceptualizing the plane itself as consisting of *either points or lines*, depending on circumstances. Full symmetry means that one accepts the idea that, in projective geometry, lines are made of points, and points are likewise made of lines; space is made of points, and space is equally made of lines.

As one gradually gains a deeper and deeper appreciation of the perfect conceptual symmetry of the projective plane, one's sense of what points and lines really are — even *Euclidean* points and lines — undergoes a subtle but genuine change. One is led, for example, to looking for duals, inside Euclidean geometry, of all sorts of phenomena. One begins to think of angles and lengths as dual to each other, for instance. Often such heuristics lead to the invention of new Euclidean concepts, and sometimes to the discovery of new theorems. However, since the Euclidean plane is *not* the projective plane and duality is not a true symmetry in Euclidean geometry, there is no rigor to this process. It is hit-or-miss, which makes it all the more of an adventure, but sometimes it is quite frustrating and disappointing when an idea suggested by this approximate duality simply doesn't pan out at all.

When one jumps back and forth between Euclidean and projective geometry sufficiently often, one becomes very sensitive to these issues of duality and symmetry, and one goes on the lookout both for unexpected appearances of duality within Euclidean geometry and for unexpected appearances of nonduality within projective geometry. The latter, of course, are simply cues that one's intuitions are still insufficiently projective, and can always be fixed up by replacing Euclidean prejudices by symmetric notions.

There is, however, one other place to look for a violation of duality, which is not apparent at first — namely, *at the interface* between the two geometries. In other words, it has to do with the very manner in which Euclidean geometry comes out of projective geometry. Euclidean geometry, as was stated above, is usually cast as resulting from projective geometry by the deletion of a *line*. But this is asymmetric! Why not delete a *point* — would that not be just as good? This simple idea is the premise underlying "Euclidual geometry", which is the principal subject of this article. (Incidentally, the term "Euclidual", coined as a combination of "Euclid" and "dual", is meant to be accented on its second syllable, so that it rhymes with "residual".)

A more symmetric picture




The Alien yet Familiar Face of Euclidal Geometry

Because projective geometry is perfectly symmetric, deleting a line and deleting a point should give exactly the same thing — or rather, exactly isomorphic things. And indeed they do. Euclidal geometry is, as its name implies, the *exact dual* of Euclidean geometry. In some sense, then, there is nothing conceptually new in Euclidal geometry. How could there be? It is formally identical to Euclidean geometry, simply with the concepts “point” and “line” exchanged.

And yet, somehow, it is different — radically different. How and why can this be the case? The answer is, the terms “point” and “line” conjure up *pictures* in our heads, and that is an inevitable part of understanding the meaning of sentences involving those terms. We are not at all like the imaginary mathematicians, posited in so many books on the philosophy of mathematics, who take certain primitive terms as “undefined” and simply let them be whatever the axioms and theorems *define* them as being. That is a formalistic, fairy-tale notion of what the act of doing mathematics is, stemming from a period when mechanizable deductive logic, not ineffable psychological imagery, was considered to be the essence of mathematics. Fortunately, that misguided idea is gradually receding into the past as the roles of intuition, imagery, and so on come to be recognized for what they are.

Even though in principle, points and lines are formally interchangeable in projective geometry, in practice no such symmetry is used in the figures in books on projective geometry. For instance, no author on projective geometry would ever have the gall to label the figure shown below “two points meeting in a line”:



The reason, of course, is that such a label is patently *wrong*. What we see is two

lines meeting in a point, not the reverse. We have strong feelings about what lines and points *are*, and our figures must somehow respect those feelings, at least to the extent that they can do so while also remaining faithful to the phenomena they are intended to represent.

It is for this reason that Euclidual geometry, despite its total formal indistinguishability from Euclidean geometry, will turn out to have an utterly alien, bizarre, counterintuitive, and disorienting feeling to it. Euclidual geometry breaks the symmetry of projective geometry in exactly the opposite way that Euclidean geometry breaks it. In other words, Euclidean and Euclidual geometries represent two modes of symmetry-breaking that together constitute a symmetric duo. Thus, where lines become open and points remain closed in Euclidean geometry, in Euclidual geometry it is the reverse: points become open while lines remain closed. More concretely, this means that when you ping a tacked line, it does *not* return to its initial position, but rather, keeps on twisting forever without ever making it back. And of course, when you ping a point on a line, it *does* return, just as it did in projective geometry. Thus to us Euclideans, Euclidual geometry is even *more* counterintuitive than projective geometry is, because in it, *two* totally taken-for-granted aspects of Euclidean geometry are overturned: the openness of lines and the closure of points.

How can this be understood? How can one build up intuitions for this strange new universe? In the several double-columned pages that follow, an elaborate and systematic comparison between Euclidean and Euclidual geometry will be presented, including a fair number of pictures in order to help intuition along. Despite the fact that it is nothing but an isomorphism and in that sense a triviality, this extended analogy contains many extremely subtle aspects that take time to digest. It is therefore recommended that it be taken quite slowly, point by point (and line by line, of course).

The way in which fresh new light is repeatedly shed on Euclidean geometry as one explores Euclidual geometry is strongly reminiscent of what happens when, in learning a foreign language, one starts to gain glimpses of how one's native language must seem to people "on the outside". Things that as a monolingual one had taken entirely for granted start to seem far from inevitable, and sometimes downright strange. For instance, to a monolingual speaker of English, it is probably inconceivable that a language could get by entirely without definite or indefinite articles — they seem like *necessary* ingredients for expressing what one means, and moreover, they seem utterly natural and obvious. But when one encounters a language like Latin or Russian or Chinese, in which there are no articles at all, one does a double-take and begins to see that words for "the" and "a" are *not* inevitable and necessary features of every language, but simply *possible* ingredients of communicative precision. Likewise, one sees that genders and declensions, which are close to nonexistent in English, can add flavors and flexibilities that English simply lacks, to a language that has them. Of course, there are myriad other cross-language distinctions that could be mentioned, such as the existence of different singular and plural noun forms, which some languages lack, or the existence of verbal modes and tenses that English lacks, and so on and so forth. To come into contact with each one of these cross-language differences is to discover a hidden axis of variability that one would never have suspected existed if one had remained entirely immersed in English.


To be sure, similar disorienting (but in the long run, super-orienting) experiences have happened before in the history of geometry — most notably in the discovery of non-Euclidean geometries, but also in the invention of four-

dimensional and higher-dimensional Euclidean geometries. However, Euclidial geometry is especially powerful in this regard because, unlike the others, it is *entirely isomorphic* to Euclidean geometry, and thus very familiar in one sense, while at the same time being *utterly alien* in a different sense.


Perhaps the act of steeping oneself in Euclidial geometry might best be likened to trying to gain fluency in “hsilgnE” — totally time-reversed English, of the sort that can be heard when a recording of normal English speech is played backwards. Some of the sounds of hsilgnE are very easy to imitate well (“m”, for instance), but others are exceedingly difficult. For instance, the word “two”, when time-reversed, sounds something like “whoosh!”, although normal symbols can’t quite capture the abrupt quenching of the “sh” sound that takes place at the end. The naïve assumption that a time-reversed “t” would sound much like a normal “t” is utterly wrong, because of the explosive attack. This is the kind of clear but subtle revelation about English that one would never obtain, were one not exposed to the alternative language.

Moreover, the grammar of hsilgnE, while obviously perfectly isomorphic to that of English, is nonetheless very alien to our way of thinking. How could anyone feel comfortable putting subjects at the *end* of sentences, or adding pluralization and past-tense markers to the *beginnings* of nouns and verbs, for instance? One could easily spend a lifetime trying to internalize the strange ways in which native hsilgnE speakers form sentences completely effortlessly — and yet by simply recording oneself and playing the tape backwards, one can hear one’s own voice doing it flawlessly!

This kind of half-alien, half-normal feeling is about as close as one can come to the experience of stepping into the Euclidial universe.



*The Euclidean plane is made of infinitely many points,
resembling an infinite case of chicken pox.*




*The Euclidean plane is made of infinitely many lines,
resembling an infinite set of Pick-up Stix.*


Denizens of the Euclidean plane



*Euclidean point
(primitive entity
having no constituents)*




*Euclidean line
(composed of
Euclidean points)*




*Euclidean line segment
(composed of Euclidean points)*


Denizens of the Euclidal plane



*Euclidal line
(primitive entity
having no constituents)*



*Euclidal point
(composed of
Euclidal lines)*



*Euclidal point segment
(composed of Euclidal lines)*

Euclidean Geometry

The Euclidean plane is composed of infinitesimal, indivisible entities called *points*.

A *pencilpoint* is a tool that, when placed on a Euclidean piece of paper, makes a point.

A *line* is a composite entity that contains infinitely many points, whereas a point, being infinitesimal and indivisible, does not contain any lines at all.

The fundamental operation that makes a line out of many points is *sliding*, also known as *translation*.

When a pencilpoint is held against a *ruler* and slid, it traces out many points constituting a finite portion of a line — a *line segment*.

There is a special line, called the *line at infinity*. All points that belong to it are called *points at infinity*. This line and its constituent points do not belong to the proper Euclidean plane.

A line is formed by translating a point from $-\infty$ to $+\infty$. Translationally speaking, therefore, a line is not just composite, but *infinite* in extent.

A line, if translated along itself (*i.e.*, so that each of its constituent points is carried into another), remains exactly the same line.

By contrast, if a line is *rotated*, no matter by how little, it becomes a different line. Rotationally speaking, therefore, a line has *infinitesimal* extent.

Points are *rotationally finite* (*i.e.*, closed) in the sense that if you spin a line about a fixed point, it returns to its starting position.

Lines are *translationally infinite* (*i.e.*, open) in the sense that if you slide a point along a fixed line, it will approach the line at infinity but will never reach it; the point will therefore never return to its starting position.

When a point travels between any two points on a given line, it traces out a *line segment*.

Euclidal Geometry

The Euclidal plane is composed of infinitesimal, indivisible entities called *lines*.

A *rangeline* is a tool that, when placed on a Euclidal piece of paper, makes a line.

A *point* is a composite entity that contains infinitely many lines, whereas a line, being infinitesimal and indivisible, does not contain any points at all.

The fundamental operation that makes a point out of many lines is *twisting*, also known as *rotation*.

When a rangeline is held against a *rotor* and twisted, it traces out many lines constituting a finite portion of a point — a *point segment*.

There is a special point, called the *point at infinity*. All lines that belong to it are called *lines at infinity*. This point and its constituent lines do not belong to the proper Euclidal plane.

A point is formed by rotating a line from $-\infty$ to $+\infty$. Rotationally speaking, therefore, a point is not just composite, but *infinite* in extent.


A point, if rotated about itself (*i.e.*, so that each of its constituent lines is carried into another), remains exactly the same point.

By contrast, if a point is *translated*, no matter by how little, it becomes a different point. Translationally speaking, therefore, a point has *infinitesimal* extent.

Lines are *translationally finite* (*i.e.*, closed) in the sense that if you slide a point along a fixed line, it returns to its starting position.

Points are *rotationally infinite* (*i.e.*, open) in the sense that if you spin a line about a fixed point, it will approach the point at infinity but will never reach it; the line will therefore never return to its starting position.

When a line pivots between any two lines on a given point, it traces out a *point segment*.



A Euclidean writes...

"I am puzzled. To me, a Euclidal point looks so big that it seems to fill all of space by itself!"

A Euclidal replies...

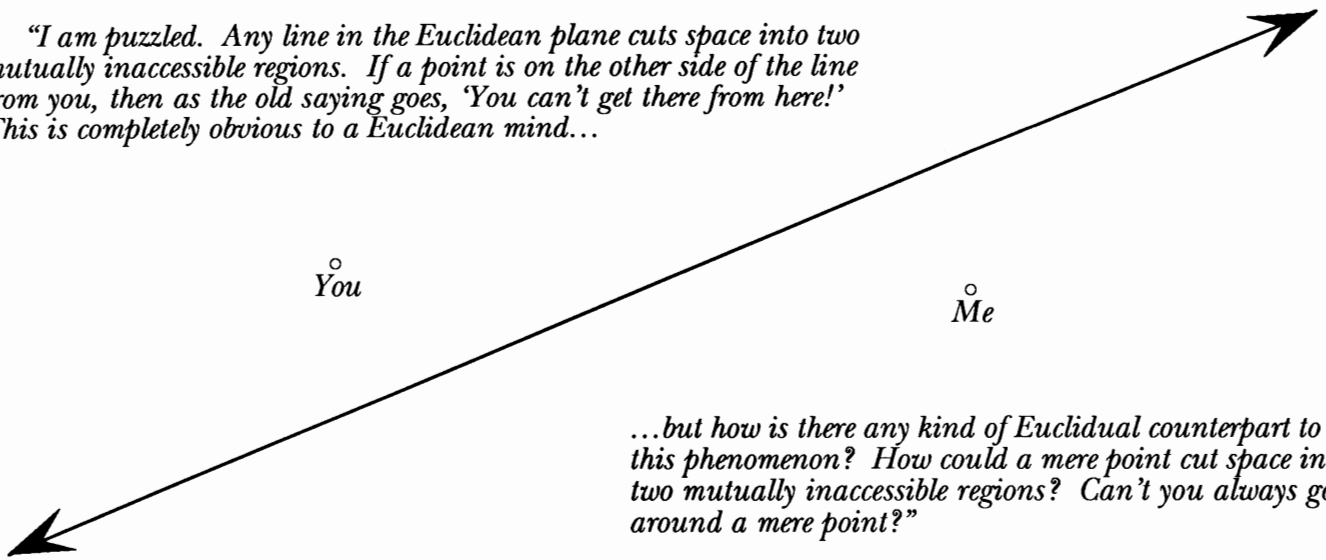
*"Of course it seems that way to you because, like all Euclideans, you unconsciously conceive of space as consisting of all possible points (i.e., dots). If, however, you retooled yourself to think of space as consisting of all possible lines, then you would see that there are lots of 'parts' of space that this Euclidal point misses — for example, line *m*. In fact, this point includes only an infinitesimal fraction of all possible lines! However, if you were to slide the point along any fixed line, it would sooner or later come to include line *m* — or any other line — as one of its constituents, although only for the most fleeting moment.*

And one last remark: we Euclidals, believe it or not, tend to think that one single Euclidean line fills all of space by itself, because when we think of all those infinitely many points on it, what we imagine is lines spraying out in every direction from each one of them, thus including every possible line. What we forget is that that is not what points are, on the Euclidean plane — they are just teeny dots and are not made out of infinitely many lines. What a strange notion of points that is!

Such are the deep-rooted assumptions that one unconsciously brings along when one crosses over from one culture to another. Of course they can be overcome, but it takes time."

A Euclidean writes...

"I am puzzled. Any line in the Euclidean plane cuts space into two mutually inaccessible regions. If a point is on the other side of the line from you, then as the old saying goes, 'You can't get there from here!' This is completely obvious to a Euclidean mind..."





A Euclidual replies...

"You are again forgetting that we Eucliduals are concerned not with points but with lines. If we want one line to coincide with another one, we can always get it to do so by rotating it about the point where those two lines meet (just as you folks would get one point to coincide with another one by sliding the former point along the line that connects the two points). Under normal circumstances, we can always carry out such a rotation (just as under normal circumstances, you folks can always carry out such a sliding operation). Of course you Euclideans could imagine us doing this by using either a clockwise or a counterclockwise rotation, but we don't have that liberty. We have to go whichever way doesn't involve making our line swing through the forbidden point at infinity, whereas you don't see any problem going either way.

But now suppose one point has been deleted from our plane. Can we carry out our rotation? Well, it all depends, as you can see in the picture below. If point P is the deleted one, then You can get to Me, but if point Q is the deleted one, then You cannot get to Me — you run into Q in attempting to do so, and of course can't get around it at all, because going 'the other way', as you Euclideans might put it, is a priori forbidden.

And one last remark: we Eucliduals would be similarly tempted to think you Euclideans could 'go around' a mere line by going 'the other way' — namely, by going out to the line at infinity, crossing it, and coming back on the other side. But that, to you, is inconceivable — just as rotating through the point at infinity is to us."





A gnomonic projection establishes a one-to-one correspondence between points on a plane (including the line at infinity) and antipodal point-pairs on the surface of a sphere that sits on the plane. Through any point P of the plane, draw the line that connects it with the center O of the sphere. This line will always pass through the sphere at two antipodal points P_1 and P_2 , which are then conceptually united and taken to be the image of point P . Under this mapping, the line at infinity goes into the equator. [Figure borrowed from H. S. M. Coxeter, Introduction to Geometry.]

A Euclidean Model and a Spherical Model of the Projective Plane

At this point, we interrupt our two-column presentation to describe how the Euclidual plane can be modeled (*i.e.*, simulated) on the Euclidean plane. Firstly, however, it should be recalled that the *projective* plane can be modeled on the Euclidean plane by adding to the latter the line at infinity, which feels, intuitively, like a huge circle sweeping around the edges of the whole plane (as if it had edges!), despite the fact that that line must be considered just as straight as any of the normal lines on the Euclidean plane. In addition to restoring the line at infinity, one must also drop all metric notions — namely, angles and lengths. As a consequence, not only parallelism but also perpendicularity goes down the drain.

For every cardinal direction (*e.g.*, north, southeast, north-by-northwest, etc.), there is a point on the line at infinity — or rather, for every *opposite pair* of cardinal directions there is a point. Thus, there is a “north–south” point at infinity, an “east–west” point at infinity, and so on. Every normal Euclidean line passes through exactly one point on the line at infinity.

Given that the idea of constant distance has vanished, the concept of “circle” loses its meaning on the projective plane, although, as it turns out, the notion of “conic” does have meaning. In other words, it turns out that conics — ellipses, hyperbolas, parabolas, and circles — can be defined entirely non-metrically, which is a rather startling finding, especially since one’s first exposure to conics in school is almost invariably in terms of constant distances, sums of distances, and differences of distances. (It might well seem contradictory for conics to be definable while circles are undefinable. The explanation is as follows. In projective geometry there is no way of distinguishing one type of conic from any other type. From a projective point of view, a hyperbola and a circle are indistinguishable curves! Thus there is no purely projective way of singling out from the set of all conics just those that are circular, which is why it still makes perfect sense to say that circles are not definable in projective geometry, even though a purely projective construction on the Euclidean plane could, by happenstance, result in a perfect circle.)

A closely related and extremely useful conceptual model of the projective plane is based on the surface of a sphere. The usual way this model is described is to say that a projective point consists in a *pair of antipodal points* on the sphere, and a projective line is a *great circle*. Since any two distinct great circles intersect in a pair of antipodal points, we have every pair of projective lines intersecting in a projective point. The converse, namely every pair of projective points determining a projective line, is also intuitively clear. Consider, for instance, the projective point N–S, consisting of the north and south poles. Any other antipodal point-pair on the globe determines a unique great circle — the meridian, or line of longitude, passing through both itself and N–S.

This model allows one easily to visualize how both points and lines can simultaneously be closed entities, something that seems utterly mystifying if one thinks in purely Euclidean terms. That *points* are closed is exemplified by the image of a meridian twirling about the fixed N–S pole. Obviously it will return to where it started after some time. That *lines* are closed is no harder to see: simply ping any antipodal point-pair on the equator (say), and it will sail right around the globe like Magellan in his ship, coming back in due time to where it started and then starting the same trek right over again. (Unfortunately, Magellan never did quite make it back, showing that history is less reliable than geometry.)

There is a natural mapping of this spherical model of projective geometry onto the extended Euclidean plane. Imagine the sphere to be sitting on the plane with its

south pole smack on the origin. Then imagine an arbitrary line in three-space that passes through the center of the sphere. Each such line of course intersects the sphere in two antipodal points — that is, one projective point. In addition, each such line either intersects the Euclidean plane in one point, or is parallel to it, in which case it intersects the line at infinity at one point. Thus, taken together, the set of all lines through the sphere's center establishes a one-to-one correspondence between point-pairs on the sphere and points on the extended Euclidean plane. Moreover, by imagining planes passing through the center of the sphere and cutting it into two symmetric hemispheres, it is not hard to see that this mapping sends every great circle onto a unique straight line on the plane, and conversely, every planar straight line onto a unique great circle. The equator is of course mapped onto the line at infinity. This elegant mapping of spherical phenomena onto planar phenomena by means of lines and planes that pass through the center of the sphere is known as a *gnomonic projection*.

Note that we have here followed the usual convention of portraying projective points as infinitesimal "point-like" entities, rather than as composite structures made up of many lines. Were we to do the latter, we would describe a projective point not as an antipodal point-pair, but as *the set of all great circles passing through an antipodal point-pair*. This is virtually never done in discussions of projective geometry, which is somewhat strange since it is more faithful to the point-line symmetry that projective geometry is all about. This fact reveals the deep and hard-to-recognize Euclidean prejudices that we all — even professional geometers — carry with us to the projective plane, prejudices that are almost impossible to rid ourselves of, even after years of trying.

By the way, we have said that a point consists of a rotational set of lines tacked down at a fixed spot (note the avoidance of the word "point"!) on the plane. How does this apply to *ideal* points — points at infinity? The gnomonic projection furnishes a quick answer to this question. Take a projective point on the equator (which corresponds to the line at infinity). That point consists of a set of great circles defining planes that cut the flat plane underneath the sphere in a set of parallel lines. In other words, the gnomonic image of any point-pair on the sphere's equator is a set of parallel lines on the plane below, all of which are pointing in the cardinal direction defined by the equatorial point-pair. If you slide the point around the sphere's equator, all the planar parallel lines will obediently rotate in perfect synchrony, like a polite school of infinitely long fish. Amusingly, then, a rectangular grid, consisting of two orthogonally intersecting sets of parallel lines, is nothing more and nothing less than two ideal points located at cardinal directions 90 degrees away from each other!

A Euclidean Model of the Euclidal Plane

We now return to our initial goal — that of modeling the Euclidal plane on the Euclidean plane. First of all, it is clear that we need to retain the line at infinity, since a key property of the Euclidal plane is that there are no parallel lines — that is, that every pair of lines meets somewhere — and this can happen only if we include the line at infinity. Secondly, however, we have to delete some particular *point* of the plane. Which point shall we choose?

Actually, this is an amusing question, because whereas it appears that there are an infinity of possibilities, in fact there are only two. Either you choose a point on the line at infinity, or you don't. That this distinction is all there is to the matter becomes obvious after a moment's thought.

If you *do* choose a point at infinity, it certainly doesn't matter which one, because they are all completely interchangeable. To be more specific, since the Euclidean plane is completely isotropic (a fancy way of saying that it does not single out any cardinal direction as special), once you've chosen your point, you are free to think of it as the N-S ideal point (*i.e.*, you can orient your x - y coordinate system so that the y -axis points in this direction, or if you prefer to think of it more concretely, you can orient your paper so that the point is "straight up, off the top of the page").

On the other hand, if you choose a point *not* on the line at infinity, then your point is a normal Euclidean point, and again, since the Euclidean plane is completely homogeneous (meaning that any point is identical to any other point), what appears to be a vast ocean of choices is reduced to no choice at all: as soon as you've picked a point off the line at infinity, just call it the origin, and center your coordinate system on it, and off you go! In short, your choice is between deleting the N-S ideal point, defined by the direction of the y -axis (and all lines parallel to it), or deleting the normal point (0,0).

In considering this choice, I myself was swayed, probably irrationally, by the connotations of the term "point at infinity". This seemed to me to suggest a point infinitely far away in the sense of *Euclidean distance* (which is simply a Euclidean prejudice rearing its ugly head), so I opted for deletion of the N-S ideal point. The truth is, the term "point at infinity" refers to *rotational* infinity, which has nothing to do with *translational* infinity, but I wasn't quite sure of that when I began. In any case, it's a fine choice.

Now of course merely deleting an infinitesimal dot an infinite number of miles away does nothing serious right *here*, but then again, deleting a point at infinity is a lot more involved than just that. We also have to delete all the *lines* through it, or to speak in a more truly Euclidial fashion, we have deleted nothing at all until we have deleted all the lines that the point *consists* of. Now what lines are those that go through a point infinitely far above and far below the page? To belabor the obvious, they are the N-S lines — the lines parallel to the y -axis, or in other words, the lines whose equation is of the form $x = \text{constant}$. By convention, then, we shall henceforth assume that all these lines have been deleted, or forbidden; they are the ideal lines of the Euclidial plane.

The immediate effect of deleting these lines is to bring into existence *parallel points* : points whose only shared line is a nonexistent line, which is to say, points sharing no line at all. In short, in this Euclidean model of the Euclidial plane, any points that have the same value for their x -coordinate are parallel points.

Moreover, since N-S lines are forbidden or unreachable, one now sees why a line cannot return to its starting state when spinning on a tack — it would have to pass through a nonexistent direction in order to get back! Thus, turn though it might, it will never return, no matter how long one waits. To a Euclidean person, this of course sounds preposterous: what's to keep a spinning line from reaching some particular direction in the plane, and then spinning beyond it? Is there some kind of *force*? Of course not, but that's not the point.

To gain perspective, turn the tables. Consider how silly it would sound if a Euclidial asked us, "Why doesn't a Euclidean point, once pinged, simply move rightwards along the x -axis until it reaches what you Euclideans call 'infinity' and then come trundling on back to its starting position?" We Euclideans would snortingly reply, "You've got the wrong picture in your head. Such a trip out the x -axis is *infinitely long* and would take an infinitely long time, at least if carried out at constant speed. Only if the sliding point accelerated to unlimitedly large velocities could it ultimately 'pass infinity' and come back on the other side."

Well, in perfectly analogous fashion, a Euclidual line pivoting about a fixed point and coming closer and closer to the deleted N–S direction without ever reaching it must be thought of as passing through an infinite sequence of *evenly-spaced rotational stages* precisely isomorphic to the infinite sequence of *evenly-spaced mileposts* along the x -axis. In other words, we must think of the set of lines through a fixed point as being parametrized in such a way that lines arbitrarily close to the forbidden direction have a parameter-setting arbitrarily close to ∞ .

In point of fact, a parametrization with this property is very easy to come by: it simply uses the *slope* of the line. So we are led to imagining the result of “pinging” a Euclidual line to be *not* that the line acquires a fixed *angular* velocity (as would be the case for “Euclidean line-pinging”), but rather that it swivels on a fixed point in such a way that its slope increases by a fixed amount per second. Under these conditions, the line will go on turning forever, gaining slope at a constant rate, yet it will never reach parallelism with the y -axis. If one’s visual system were as innately sensitive to *slope* as human visual systems are to *angle*, then it would seem completely obvious that the pinged line was moving out towards infinity at a constant rate, and of course could never get there. And Euclidials have precisely such visual systems. To them, a N–S line would feel precisely as inaccessible as an infinitely distant point feels to us Euclideans.

And conversely, whereas for us Euclideans the full traversal of a line by a pinged point seems an infinite voyage, no such imagery holds for Euclidials, who consider the trip from the origin out the positive x -axis and the return home (along the negative half of the line, of course) to be as easy as pie. More specifically, they consider it to be as *short* as π , since π is the measure of a full cycle of a Euclidual line, in precise analogy to the fact that in the Euclidean plane, the *angular* measure of the spin undergone by a line that returns to its starting position is π .

Some readers may object that it should be 2π , not π ; this is true in a sense but false in another sense. The proper analysis of this is very subtle, and involves the realization that a point comes back “flipped” after just one tour about a line, and it takes *two* cycles before it returns “unflipped”. This totally absurd-sounding idea is nonetheless the perfect isomorph of the fact that a line, when rotated 180 degrees about a point, comes back into coincidence with itself but *reversed in direction*; it takes another 180 degrees before it is *truly* back where it started. One gains some appreciation for this by recalling how projective points in the spherical model consist of two antipodal points, and imagining such a point-pair going halfway around the globe and seeming to be fully back, whereas in fact the two members of the couple have swapped roles, so that another half-tour of the globe is needed to restore things to how they *really* were. However, although quite helpful, this spherical image is only part of the picture, since it is still mired in the old Euclidean prejudice that points are tiny, indivisible objects. The details of this aspect of the analogy emerge far more clearly in the two-column presentation and the accompanying diagrams.


In the diagrams that follow, when parallel points are represented, they will always be portrayed as having the same x -coordinate. This convention is a consequence of my having chosen the deleted point to be a Euclideanly ideal point. Had I made the opposite choice, the deleted and hence forbidden lines — the lines that define parallelism of Euclidual points — would have all sprayed out radially from the origin. Either choice is fine; I just happened to make the former one.

There is, by the way, a “point at infinity” that in everyday life acts a little like the deleted point giving rise to Euclidual geometry — namely, the sun. Our eyes automatically avoid looking at the sun, and the more closely our gaze approaches


that direction, the harder it is on us. We could pretend that it is *infinitely* hard to look straight at the sun. In this analogy, the sun's rays (which on earth are all parallel, the sun being "essentially" infinitely far away) play the role of the deleted lines, so that any two points along the pathway of a given ray would be parallel. Although the analogy doesn't go too far, it may help to give a certain concreteness to the strange-seeming notions of "forbidden point" and "forbidden lines", not to mention "parallel points".

One last subtlety ought to be mentioned before we resume the two-column comparison, and this is a question that might have occurred to some readers already. Deletion of the N-S ideal point means deleting all the *lines* that that ideal point consists of — and one of those lines is of course the line at infinity. But deleting the line at infinity was what gave rise to Euclidean geometry! So what's going on? The answer is that we should actually have made a distinction between *two types of deletion*. There is *rotational deletion* (which is what we are dealing with in this case) and *translational deletion* (which is what gives rise to Euclidean geometry). They are independent operations, and the effect of rotationally deleting one ideal point does not entail the effects of translationally deleting all the lines that compose it. Otherwise, deleting a single point anywhere at all would mean not only that an infinite set of lines would disappear as well, but also that all the points *on* all those lines would vanish, and all of a sudden, *poof!* — the entire plane would have gone up in a puff of smoke.


Phenomena in the Euclidean plane



Two nonparallel lines share one point




Line segment linking points P and Q, with midpoint M shown




Two rays heading off in opposite directions from point 0 towards (translational) infinity


Phenomena in the Euclidual plane



Two nonparallel points share one line



Point segment linking lines p and q, with midline m shown



Two sprays heading off in opposite directions from line 0 towards (rotational) infinity

A *ray* is a “half-line” consisting of all the points belonging to a given line, starting from a given point and sliding without limit towards the line at infinity.

An arbitrary point chosen from a line divides the line into two rays.

Any two points are infinitesimal parts of a common line; any two lines have at most one point in common.

Parallel lines are lines that have no point in common (alternately, they are lines whose only common point is a point at infinity).

The idea of “parallel points” (points linked by no line) makes no sense.

The degree of rotation of a line about a fixed point is given in terms of *radians* (the natural units of *angle*). A full rotation consists of 2π radians.

If a line is rotated π radians about one of its component points, it comes back into coincidence with itself; however, the two rays defined by the fixed point have swapped places.

Rotating again through π radians brings the line back, with both rays returned to their starting positions.

Two lines, one of which has swiveled through an angle of $\pi/2$ radians relative to the other, are said to be *perpendicular lines*.

Translation of a point along a fixed line is measured in terms of *distance*. Unlike the case with rotation, there is no natural unit of distance, and distances between two points on a line can be arbitrarily large.

The *length* of a line segment is the distance between its two endpoints.

The *midpoint* of a line segment is that point that is equidistant from the segment’s endpoints.

A *curve* is a continuous locus of points.

A curve’s *arc length* is the sum of the lengths of infinitesimal line segments approximating the curve.

A *triangle* consists of three points (or vertices), and the three line segments (or sides) joining pairs of them.

A *spray* is a “half-point” consisting of all the lines belonging to a given point, starting from a given line and twisting without limit towards the point at infinity.

An arbitrary line chosen from a point divides the point into two sprays.

Any two lines are infinitesimal parts of a common point; any two points have at most one line in common.

Parallel points are points that have no line in common (alternately, they are points whose only common line is a line at infinity).

The idea of “parallel lines” (lines linked by no point) makes no sense.

The degree of translation of a point along a fixed line is given in terms of *slidians* (the natural units of *slide*). A full translation consists of 2π slidians.

If a point is translated π slidians along one of its component lines, it comes back into coincidence with itself; however, the two sprays defined by the fixed line have swapped places.

Translating again through π slidians brings the point back, with both sprays returned to their starting positions.

Two points, one of which has traveled through a slide of $\pi/2$ slidians relative to the other, are said to be *perpendicular points*.

Rotation of a line about a fixed point is measured in terms of *twistance*. Unlike the case with translation, there is no natural unit of twistance, and twistances between two lines on a point can be arbitrarily large.

The *swingth* of a point segment is the twistance between its two endlines.


The *midline* of a point segment is that line that is equitwistant from the segment’s endlines.

A *curve* is a continuous envelope of lines.

A curve’s *arc swingth* is the sum of the swingths of infinitesimal point segments approximating the curve.

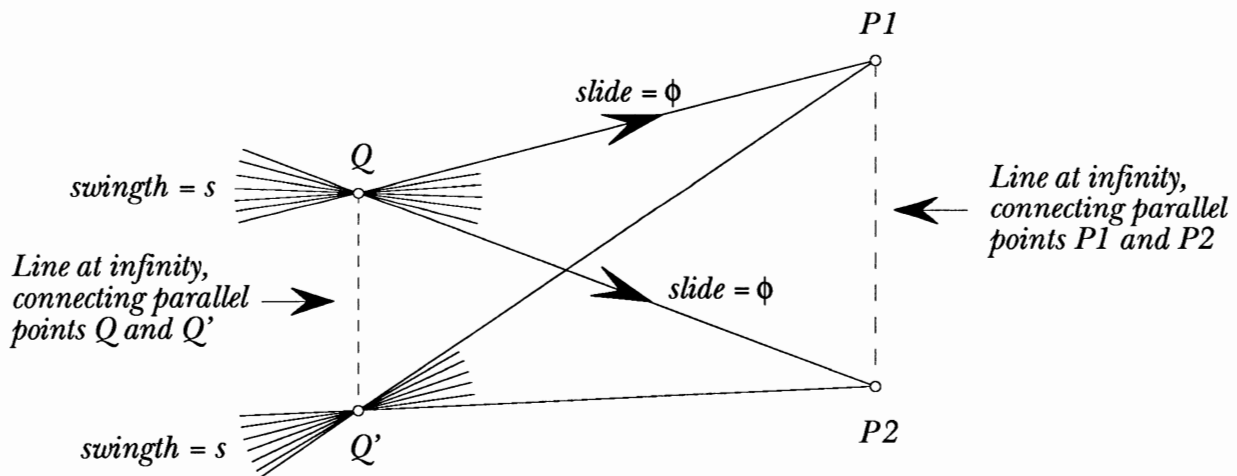
A *trislid* consists of three lines (or sides) and the three point segments (or vertices) joining pairs of them.

The meaning of parallelism in the Euclidean plane



Given two parallel lines ($p1$ and $p2$) and a third nonparallel line (q), the angles from the latter line to each of the parallel lines have the same value. Moreover, if the third line is moved parallel to itself (to q'), then the length of the line segment along it bounded by the two parallel lines stays invariant.

The meaning of parallelism in the Euclidual plane



Given two parallel points ($P1$ and $P2$) and a third nonparallel point (Q), the slides from the latter point to each of the parallel points have the same value. Moreover, if the third point is moved parallel to itself (to Q'), then the swingth of the point segment inside it bounded by the two parallel points stays invariant.

More on the Spherical Model of the Projective Plane

We now take another brief interlude from the two-column presentation to introduce a second model of the Euclidual plane, one that is perhaps the most useful of all.

We begin by recalling that the projective plane can be modeled either on the Euclidean plane (by the addition of a line at infinity and the deletion of metric notions) or on the surface of a sphere (via the gnomonic projection). In terms of enhancing one's intuition, the latter is in many ways more useful, because in it, the absolute equivalence of all lines (and of all points) is self-evident. In particular, the "line at infinity" (played by the equator) is obviously no different in any way whatsoever from any other line. In other words, in this model the "line at infinity" is just a plain old line, identical to every other line, exactly as it should be.

Moreover, the perfect duality of points and lines falls out almost trivially from the spherical model. That this happens is a consequence of the fact that there is a completely natural, simple mapping between *points* on a sphere and *lines* on a sphere. Here's how this "dualization" procedure works. To find the counterpart of any great circle, simply take the antipodal point-pair *furthest away* from the circle. In other words, start anywhere on the great circle and move perpendicularly away from it, until you have gone one quarter of the way around the sphere. Then you are maximally far from the great circle. That's the dual of the great circle, and conversely, the dual of this point is the great circle you began with. For example, under this concrete dualization-operation, the point corresponding to the equator is the N-S pole-pair, and vice versa.

It is not hard to see that if one chooses a point-pair P on a particular great circle q , then the great circle p that is the *image* of P will pass through the point-pair Q that is the image of q . In other words, this type of dualization *respects incidence*. To convince yourself of this, imagine ("with no loss of generality", as mathematicians are so fond of saying) that the great circle q is the equator, so that Q is the N-S pole-pair. Now any random point P on the equator has for its image a great circle p that perpendicularly crosses the equator, a quarter of the way around the globe from P . This means that p is a meridian and thus that it passes through the N-S pole-pair, which is of course Q , and that was what we set out to show — namely, that if a point P lies on a line q , then its image-line p passes through q 's image-point Q .

The upshot of this very simple observation about point-pairs and great circles on the sphere is that any configuration of "points" and "lines" on a sphere, no matter how complex, can be instantly dualized into another different-looking configuration in which the roles of points and lines are totally reversed but in which precisely the same incidence relations hold. This, then, is one of the great powers of the spherical model of the projective plane.


By the way, one of the beauties of this point-line switching operation is that it can be used as a way of turning any point-line picture on the Euclidean plane into a dual picture where lines and points have been reversed but all incidence relations have been retained. All you do is this: take your planar picture and use the gnomonic projection to project it up onto the surface of the sphere. Given this picture on the sphere, perform the point-line reversal just described; you now have a complementary, or dual, picture on the surface of the sphere. Now merely send this new image back down onto the plane by means of the gnomonic projection run in reverse. What results is an utterly different-looking planar picture, and yet it is, in some sense, "the same picture", since it contains all the same information as the original, simply coded differently. In particular, all incidence properties that held for

the original picture remain valid for the dualized picture.

An example of this process is shown below. In short, what you will observe is that a configuration made up of 7 points, 9 lines, and one conic, in which three key lines are concurrent, gets converted into a totally different-looking configuration made up of 9 points, 7 lines, and one conic, in which the corresponding three key points are collinear. Now for the details.

Six line segments — *a* through *f* — define a rather ugly self-intersecting hexagon (shaded). Careful inspection will show that these line segments, if extended to full lines, are all tangent to a conic, which happens to be a circle (the light circle). (The only reason it is a circle is that circles are the easiest conics to draw.) Opposite vertices of this hexagon have been joined to form three more lines, which all intersect in a point, labeled “Brianchon point”. Readers familiar with projective geometry will recognize that this configuration exemplifies Brianchon’s theorem, a brilliant geometrical ruby discovered by Charles Brianchon in 1806:

Consider the hexagon formed by any six lines tangent to any arbitrary conic. Construct the three lines that link opposing vertices of the hexagon (e.g., vertices *ab* and *de*). These three lines are concurrent.



The gnomonic-projection sphere itself sits on the plane of the paper, its south pole clearly indicated, as well as its size (by the heavy circle).

The dual figure to the Brianchon configuration, made by the pole/equator


swapping process on the sphere, is based on six points — A through F — that are the duals of their namesake line segments. These six points define another self-intersecting hexagon of a most different shape. Although it does not quite jump out at the eye, these six points all lie on another conic, namely a hyperbola — the dual conic to the first one. (Typical of the elegance pervading projective geometry is the theorem that states that this type of dualization process operating on any conic yields another conic.) Finally, opposite *sides* of this second hexagon have been joined to form three more points, all of which lie on a single line, labeled “Pascal line”. This configuration illustrates Pascal’s theorem, a gleaming geometrical emerald that is the dual to Brianchon’s theorem. However, it was discovered long before Brianchon’s theorem — in fact, in 1640 by Blaise Pascal, 16 years old:

Consider the hexagon formed by any six points lying on any arbitrary conic. Construct the three points where opposing sides of the hexagon meet (e.g., sides AB and DE). These three points are collinear.

Hopefully, this example makes clear what it means to say that a transformation not only replaces points by lines and lines by points, but also *respects incidence*.

One can add some interesting bells and whistles to the dualization operation by *rotating the sphere* before sending the reversed image on its surface back down to the plane; in fact, one can even *carry or roll the sphere* to any spot on the plane, and, if one wishes, *inflate or deflate it*, before gnomonically projecting the reversed image on its surface back down to the plane. (Note that none of these operations perturbs the internal relationships of points and lines on the sphere’s surface in the slightest.) This shows that there is not just one single way of dualizing a Euclidean picture, but an infinite family of related ways of doing so. Such incidence-respecting point–line reversals are called *polarities* in projective geometry.


The sum of the interior angles of a Euclidean triangle



A Euclidean triangle consists of three points and the line segments they define. It is crucial to keep in mind that a line segment is the finite stretch between two points — i.e., the stretch that bypasses translational infinity — since this fact defines the meaning of “interior angle”. Specifically, the interior angle at any given vertex (vertex A, for example) is that region of the vertex created by joining the vertex to all the points belonging to the opposite line segment (a, in this case). Given this definition, the sum of the interior angles of a triangle is always 180° , or π . Symbolically,

$$\alpha + \beta + \gamma = \pi$$

The sum of the interior slides of a Euclidal trislide



A Euclidal trislide consists of three lines and the point segments they define. It is crucial to keep in mind that a point segment is the finite twist between two lines — i.e., the twist that bypasses rotational infinity — since this fact defines the meaning of “interior slide”. Specifically, the interior slide on any given side (side a, for example) is that part of the side intercepted by lines emanating from the opposite point segment (A, in this case). Given this definition, the sum of the interior slides of a trislide (indicated in the diagram by arrows) is always 180° , or π . Symbolically,

$$a + b + c = \pi$$

The sum of the three interior angles of a triangle is π radians.

The three lines linking a triangle's vertices with the midpoints of opposite sides are concurrent, and their common point is called the *centroid*.

A *circle* is a curve — namely, the set of all points equidistant from a given point (its *center*). Any line segment joining the center to a point on the circle is called a *radius*. By definition, all radii have the same length.

The ratio of the arc length of a circle to the length of any radius is 2π .

The *circumcircle* of a triangle is the unique circle that passes through all three vertices of the triangle, and its center is the *circumcenter*.

An *altitude* of a triangle is a line through a vertex, which is also perpendicular to the opposite side.

The three altitudes of a triangle meet in a single point called the *orthocenter*.

The three angle-bisectors of any triangle meet in a point called its *incenter*, which is the center of its *incircle*, the unique circle that is tangent to all three line segments constituting the triangle.

The circumcenter O, centroid G, and orthocenter H of any triangle all lie on a line known as the *Euler line* of the triangle. The distance from O to G is exactly half that from G to H.

The midpoint of the OH line segment is the center of the *nine-point circle*, a circle that passes through the triangle's three midpoints, its three altitude feet (the points where the altitudes cross the sides), and the midpoints of the line segments joining the orthocenter to the respective vertices of the triangle.

The natural coordinate system for the Euclidean plane is a *biaxial* system, consisting of two perpendicular axes (*i.e.*, lines) that meet in a point called the *origin*.

The coordinates of a point on the plane are given by the perpendicular distances of the point to the two axes.

The sum of the three interior slides of a trislide is π slidians.

The three points linking a trislide's sides with the midlines of opposite vertices are collinear, and their common line is called the *centroidal*.

A *circual* is a curve — namely, the set of all lines equidistant from a given line (its *central*). Any point segment joining the central to a line on the circual is called a *radial*. By definition, all radials have the same swingth.

The ratio of the arc swingth of a circual to the swingth of any radial is 2π .

The *circumcircual* of a trislide is the unique circual that is tangent to all three sides of the trislide, and its central is the *circumcentral*.

An *altitude* of a trislide is a point on a side, which is also perpendicular to the opposite vertex.

The three altitudes of a trislide meet in a single line called the *orthocentral*.

The three slide-bisectors of any trislide meet in a line called its *incentral*, which is the central of its *incircual*, the unique circual that passes through all three point segments constituting the trislide.

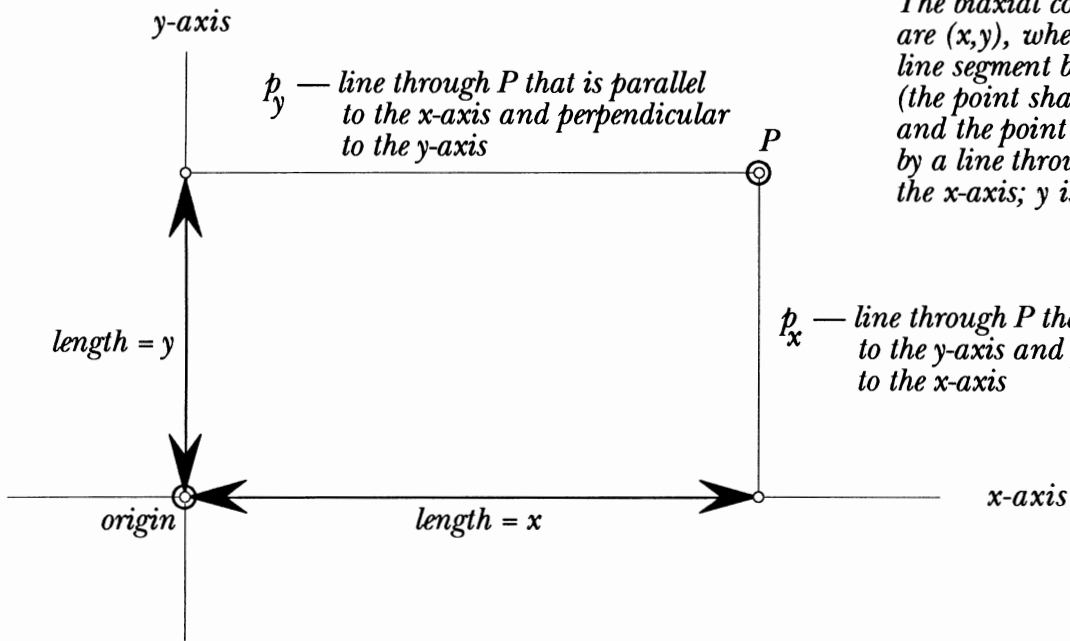
The circumcentral o, centroidal g, and orthocentral h of any trislide all lie on a point known as the *Euler point* of the trislide. The twistance from o to g is exactly half that from g to h.

The midline of the oh point segment is the central of the *nine-line circual*, a circual that is tangent to the trislide's three midlines, its three altitude legs (the lines joining the altitudes to the vertices), and the midlines of the point segments joining the orthocentral to the respective sides of the trislide.

The natural coordinate system for the Euclidual plane is a *bipolar* system, consisting of two perpendicular poles (*i.e.*, points) that meet in a line called the *origin*.

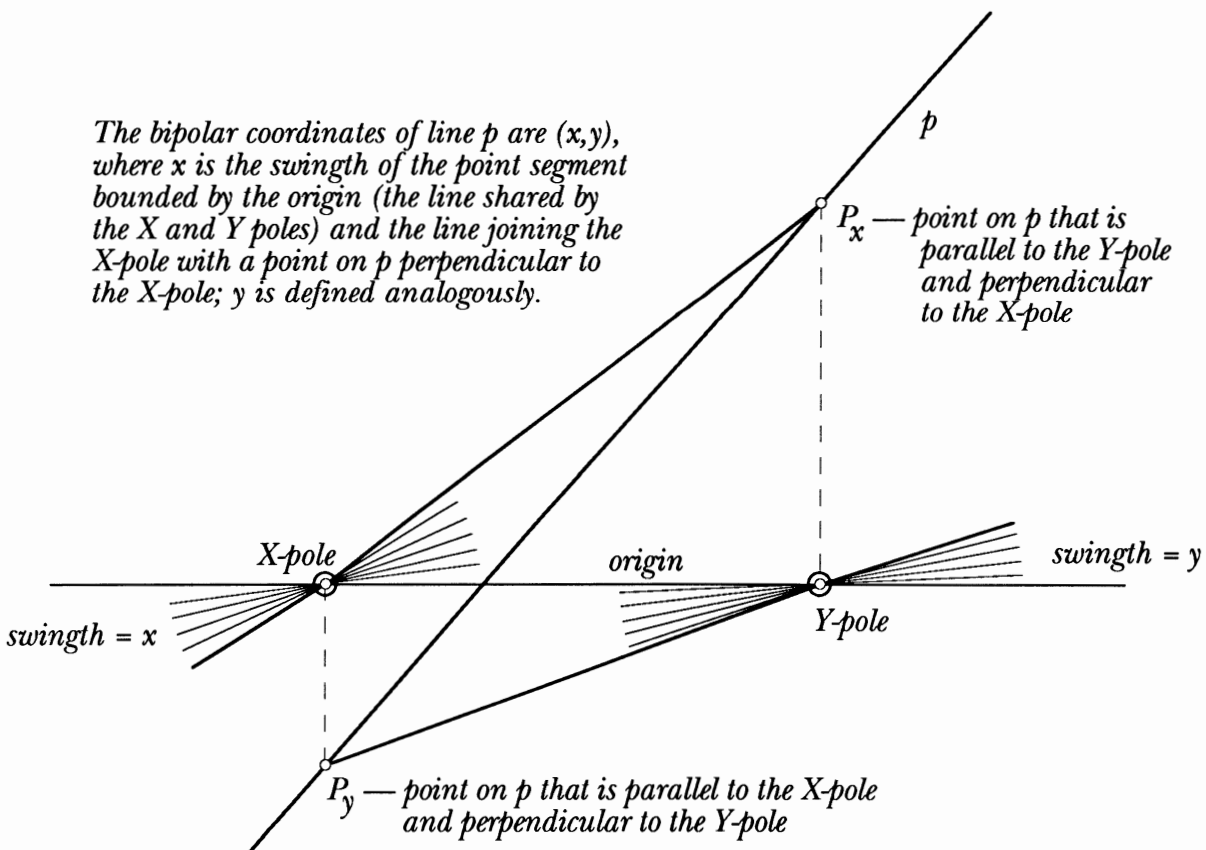
The coordinates of a line on the plane are given by the perpendicular twistances of the line to the two poles.

Biaxial coordinates of a point in the Euclidean plane



The biaxial coordinates of point P are (x,y) , where x is the length of the line segment bounded by the origin (the point shared by the x and y axes) and the point intercepted on the x -axis by a line through P perpendicular to the x -axis; y is defined analogously.

Bipolar coordinates of a line in the Euclidean plane



The bipolar coordinates of line p are (x,y) , where x is the swingth of the point segment bounded by the origin (the line shared by the X and Y poles) and the line joining the X-pole with a point on p perpendicular to the X-pole; y is defined analogously.

P_x — point on p that is parallel to the Y-pole and perpendicular to the X-pole

P_y — point on p that is parallel to the X-pole and perpendicular to the Y-pole

A Tubal Model of the Euclidual Plane

Armed with this deeper understanding of the spherical model of the projective plane, we now turn to a similar three-dimensional model of the Euclidual plane, one with equally remarkable intuition-enhancing powers. To do so, we will describe a mapping that converts our already-existing Euclidean model of the Euclidual plane into a three-dimensional model.

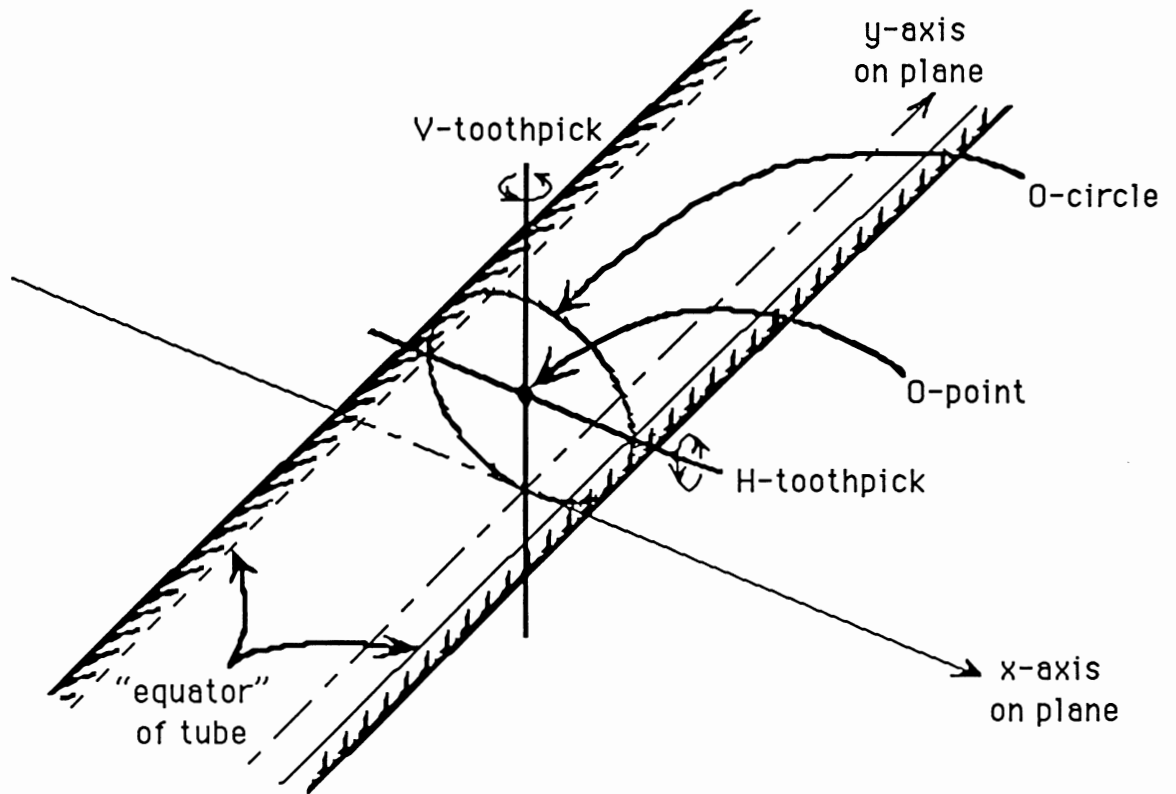
Imagine an infinitely long tube of radius 1 sitting on the Euclidean plane, oriented such that its central axis — an infinitely long line, of course — lies precisely above the y -axis. Let us denote by the term “O-point” that point on the tube’s axis that lies directly above the origin (*i.e.*, where $y = 0$). Now it is very easy to define a “quasi-gnomonic” projection of points on the Euclidean plane up onto the tube’s surface. To get from any point P on the plane to its tubal counterpart P' , we simply construct the straight line that passes through both the O-point and P . This is a unique straight line, of course, and it penetrates the tube in two “opposite” or “antipodal” points. Much as in the spherical model of projective geometry, we conceptually unite these two tubally antipodal points into one single entity P' , which will be the image of P .

Note that as usual, Euclidean prejudices have crept back into the picture, unnoticed: we are again describing points as tiny, indivisible objects. If the tube’s surface is to be considered a model of Euclidual geometry, however, we must rid ourselves of this type of imagery, and replace it with the idea that a point consists of all the *lines* on the tube that pass through the given point-pair P' . But then the question instantly arises: What constitutes a *line* on the surface of a tube? Luckily, that’s very simple: a line on the tube is simply the quasi-gnomonic image of a line on the plane. Saying just this may not make it easy to envision what *shape* a line actually traces out on the surface of the tube. In fact, however, the shape is a perfect ellipse whose center is the O-point.

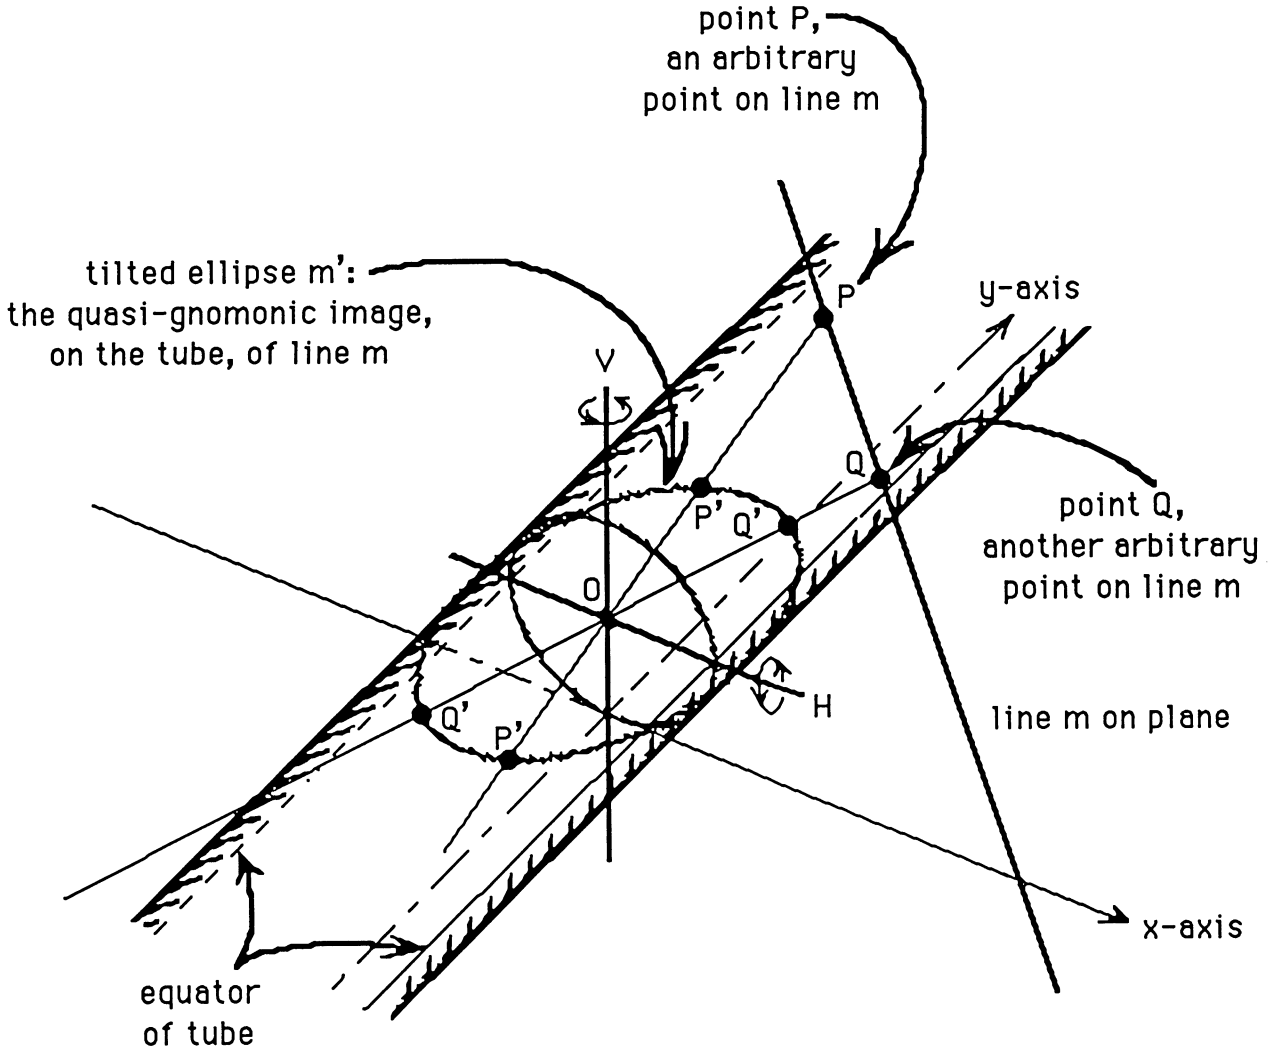
To show this, we consider the following alternate but equivalent way of describing the quasi-gnomonic projection of a planar line m . Imagine the unique plane that cuts the Euclidean plane in the given line and that simultaneously passes through the O-point; in general, this plane will slice the tube obliquely, and in so doing it will define an ellipse m' centered on the O-point. (It is intuitively quite obvious that slicing a tube with a plane can result only in a perfect circle — when the plane is perpendicular to the tube’s axis — or an ellipse.)

So now we know how to conceive of a Euclidual *line* on the tube, and armed with that knowledge, we can figure out how to conceive of a Euclidual *point* as well: a Euclidual point consists in the set of all tubal ellipses centered on the O-point and passing through a fixed point-pair on the tube’s surface. Remember that the two points constituting such a point-pair are connected by a line through the O-point, so that if one of them is 5 units “north” of the O-point, then the other will be 5 units “south” of it (*i.e.*, the points in such a pair have equal but opposite y -coordinates). Similarly, if the *angular* coordinate of either of these tubal points is ϕ (where 0 is the angle assigned to points halfway up the tube — on its “equator”, so to speak — and $\pi/2$ is the angle assigned to points sitting on the bottom of the tube), then the other point will have angular coordinate $-\phi$.

Much as the spherical model of projective geometry greatly facilitates one’s intuition for the simultaneous closure of both points and lines, so the tubal model greatly enhances one’s intuition for how *lines* can be closed while *points* are open — precisely the opposite of what holds in Euclidean geometry. Envisioning the closure of lines is trivial: a tubal line being an ellipse, which is a closed curve, it is



A tube of radius 1 sitting on the plane, tangent to the y-axis. Its circular cross-section directly above the x-axis is called the "O-circle", and the center of the O-circle is called the "O-point". Through the O-point pass two orthogonal "toothpicks": the H-toothpick, which is horizontal, and the V-toothpick, which is vertical. Each of them can twist about its own axis. The two parallel horizontal lines on the tube through which the H-toothpick passes form the "equator" of the tube.



An arbitrary line, line m , has been drawn on the plane. Two arbitrary points on line m — points P and Q — have been connected by straight lines to the O -point. Each such line intersects the tube in a pair of antipodal points (i.e., points symmetrically located with respect to the O -point). If a point were to slide along all of line m , its image points on the tube would trace out a tilted ellipse m' , whose center is the O -point.

Ellipse m' can equally well be described as the intersection of the tube with the plane that passes through both the O -point and the line m . One can imagine obtaining both m and m' by controlling the tilt of a variable plane that passes through the O -point, using the H - and V -toothpicks as controls.

easy to see how a point, when pinged, would zip around the line and come back to its starting point. Of course, what held for the spherical model also holds here — namely, a tubal point, being a point-pair, seems already to be back after just 180 degrees, or π radians, whereas in fact it is only “pseudo-back”; true restoration of the point’s original state requires one more cycle through π radians. (We shall henceforth refer to trips of points along lines as involving a certain number of *slidians* — intended to rhyme with “Lydians” — rather than *radians*, since we wish to emphasize that we are dealing with a measure of linear separation between points, and not with angles between lines.)


The non-closure of points is a little subtler than the closure of lines is. This requires us to imagine the set of all tubal lines passing through a given tubal point. Some of these lines — *i.e.*, ellipses — are rather small, being nearly perpendicular to the axis of the tube. However, others are very large, being nearly parallel to the axis of the tube. In the limit, such lines would be arbitrarily elongated ellipses, stretching as far along the tube as the eye could see. Clearly, *something* is going to infinity as these “lines” approach parallelism with the y -axis. This notion of a parameter increasing at a constant rate, and thus able to head for but never to actually reach infinity, is just what is needed to make non-closure understandable. Note that if one of these lines actually *did* reach the infinite status, that would mean that it had become parallel with the tube’s axis, and thus that it ran all the way from $y = -\infty$ to $y = +\infty$. Such an “ellipse” on the tube’s surface would in actuality consist of two straight N–S lines, diametrically across from each other on the tube. This antipodal line-pair is of course one of the infinite set of lines that were deleted when the point at infinity was deleted.

Slide and Twistance Made (Relatively) Easy


The tubal model allows us easily to visualize many counterintuitive phenomena of the Euclidual plane. For instance, the notion of *slide* between two points is reduced to a very simple and clear idea on the tube: the amount of slide between two points is simply the *difference in their angular coordinates* ϕ_1 and ϕ_2 . It makes no difference where the points happen to be located longitudinally along the tube (*i.e.*, their y -coordinates). Two points are said to be *perpendicular* if the slide between them has value $\pi/2$, meaning that they are located exactly one-quarter of the way around the tube from each other, no matter where they are along its axis.

If one projects this down onto the Euclidean plane, one instantly sees that any point on the y -axis is perpendicular to any ideal point of the (extended) Euclidean plane, since the latter all have an angular coordinate of 0, whereas the former all have angular coordinate $\pi/2$. Similarly, any two points having angular coordinates ϕ and $\phi + \pi/2$ will be mutually perpendicular, whatever their y -coordinates. Forbidden lines are thus *lines of constant slide* relative to any given point. (Note that the Euclidean dual of this statement is the proposition that ideal points are “points of constant angle” relative to any given line. Though this sounds a little obscure, what it means is simply that the angle between a fixed line in the plane and any variable line through a fixed ideal point is invariant. It is just a fancy way of saying, then, that for people scattered across any flat region, no matter how large, the angle between the ground and an infinitely distant star is the same.)

One of the most surprising and elegant properties of the tubal model is what a *circual* turns out to look like in it. To see how this works, one must first think about the most natural way to impose a Euclidual coordinate system on the tube. One wants a *bipolar* system — the dual of a *biaxial* system — in which lines rather than



Perpendicular points on the E-plane are such that the difference between the azimuthal angles of their tubo-gnomonic projections is 90 degrees. More intuitively put, their tubal images are 90 degrees around the tube from each other. To a Euclidual eye, perpendicular points look as far apart as any pair of points can get from each other.



The slide between any two points on the E-plane depends solely on their x-coordinates. To find its value, project the two points up onto the tube by tubo-gnomonic projection, then take the difference between their azimuthal angles. This is the value of the slide. Note that this does not imply that the slide between two points is a function solely of the difference between their x-coordinates; for instance, segment PQ is clearly much shorter than segment RS , yet $\text{slide}(P,Q)$ is exactly equal to $\text{slide}(R,S)$.

points are the basic entities to which one assigns coordinate pairs. The dual to a Cartesian coordinate system is one in which the X and Y poles, like the x and y Cartesian axes, are perpendicular to each other. We can freely choose the poles to both lie on the x -axis, which makes the x -axis itself now become the “origin” — that is, the *line* relative to which all twistances are to be measured, just as in Cartesian coordinates, the origin is the *point* relative to which all distances are to be measured. And what is a twistance? A twistance is a measure of the rotational separation of two lines, superficially reminiscent of an angle but ranging between $-\infty$ and $+\infty$, and thus more like a distance.

One might think that the twistance between two lines would just be a function of the angle between them, but it's nothing of the sort; twistances are trickier than that. Imagine, for instance, two lines that are parallel in the Euclidean plane — for simplicity's sake, the east–west lines $y = 0$ and $y = 1$. One would expect that since they are parallel, the twistance between them would be zero. Wrong! There is a non-zero twistance between *any* pair of nonidentical lines, parallel or not, just as in Euclidean geometry, there is a non-zero *distance* between any pair of nonidentical points (whether or not a Euclid might think they were parallel points).

Quasi-gnomonic projection up onto the tube makes this initially troubling idea much more understandable. The line $y = 0$ becomes a circle running vertically around the tube — in fact, we shall henceforth refer to this special circle as the *O-circle* — and the line $y = 1$ becomes an ellipse tilted with respect to the tube's axis, and intersecting the O-circle in two antipodal points on the tube's “equator” (the pair of infinitely long horizontal lines facing each other halfway up the tube). Despite the fact that $y = 0$ and $y = 1$ are parallel lines on the Euclidean plane, they have tubal images (the O-circle and the tilted ellipse) that not only intersect but do so at a non-zero angle. Although this does not tell us the twistance between them, it gives a good intuitive sense for why the twistance is not zero.

But let us return to the coordinatization of the tube and the plane. As was said a moment ago, we shall take as our poles two perpendicular points on the Euclidean plane's x -axis (or, on the tube, the quasi-gnomonic images of those points). For simplicity's sake, we might as well make them the Euclidean origin ($x = 0, y = 0$) and the “east–west” ideal point defined by the direction of the x -axis. The tubal image of that ideal point is the point-pair on the tube's equator having y -coordinate 0. So now imagine a pair of toothpicks (I could have said “poles”, but that might have been too confusing!) stuck through the tube, one vertically and one horizontally, at these points' images. We shall call these the *H-toothpick* and the *V-toothpick*. The H and V toothpicks pass through each other precisely in the middle of the tube, at the O-point. We will define our bipolar coordinates in terms of rotations of these two toothpicks.

First, however, we shall introduce one more useful piece of terminology. Imagine a variable plane passing through the O-point. We shall call this plane the *O-plane*. The O-plane always cuts the tube in some ellipse centered on the O-point; when the O-plane is perpendicular to the tube's axis, that ellipse is of course the O-circle. These various “O-entities” will play key roles in our development of bipolar tubal coordinates. Note that the two toothpicks constitute two perpendicular diameters of the O-circle.

Imagine now that the O-plane can be made to rotate about the O-point by twisting either toothpick. (A very useful image, for those who know it, is the delightful Scandinavian marble maze called “Labyrinth”, in which a wooden board is made to tilt at different angles by two independent twists controlled manually by knobs located on adjacent sides of a wooden box.) What we are interested in is not

the rotated plane itself, but the *ellipse* made by its intersection with the tube — in other words, in what the O-circle becomes as the toothpicks are twisted.

If one looks down from infinitely far above the plane and twists only the V-toothpick (the one sticking straight up at us out of the Euclidean plane), the resultant ellipse will of course appear to be a line segment, because one is seeing it edge-on. As the V-toothpick is twisted, we see the line segment crossing the tube at an increasing angle, growing longer and longer until it eventually is parallel with the tube. In a similar manner, if one moves out the x -axis to the east–west ideal point and twists only the H toothpick (the one floating one unit above the x -axis and parallel to it), one again sees a rotating ellipse edge-on, which is to say, a line segment crossing the width of the tube at various angles, growing unboundedly long as the angle of rotation of the H -toothpick approaches 90° .

If (as in the Labyrinth game) we now allow ourselves a combination of *two* twists, one of the V-toothpick and one of the H-toothpick, we can turn the original circle into any desired “line” on the tube (ellipse centered on the O-point). Although that claim might sound surprising, it is just the Euclidal version of the pedestrian Euclidean fact that any possible point on the Euclidean plane can be reached by translating the origin twice — once parallel to the x -axis, and once parallel to the y -axis.

To finish our coordinatization task, we merely need to assign *numbers* to each of the two toothpick-twists. That, luckily, is quite easy. We imagine making either one of the two twists by itself; then we observe how far the resultant ellipse extends up or down the y -axis — in other words, we take that ellipse’s extremal y -coordinate. That number, which can of course range between $-\infty$ and $+\infty$, will be the bipolar coordinate associated with the given toothpick (*i.e.*, pole). We do this for both poles, and the *pair* of such y -coordinates will constitute the bipolar coordinates of any given line (*i.e.*, ellipse).

Guided by a sense of duality, we have now figured out how to assign two independent bipolar coordinates to any of the allowed ellipses on the tube (and by projection, these same coordinates will of course apply to the corresponding line on the Euclidal plane). Now in Euclidean geometry, given a coordinate system, one of the most natural questions that arises is: “How can we use it to calculate *distances*?” The analogous question here is, of course, “How can we calculate *twistances*?” In Euclidean geometry, the answer is that if we *combine* the two coordinates (x, y) of any point using the Pythagorean formula $\sqrt{(x^2 + y^2)}$, we get the distance of the given point from the origin. More generally, the formula $\sqrt{[(x_2 - x_1)^2 + (y_2 - y_1)^2]}$ gives the distance between an arbitrary pair of points. The analogous answer is therefore that the Pythagorean formula applied to the bipolar coordinates of any two lines will give the *twistance* between those lines. Now, of course, we can easily calculate the *twistance* between the parallel Euclidean lines $y = 0$ and $y = 1$; it turns out to be 1.


What is the Shape of a Circual?

Having now a formula for *twistance*, we are in a position to define the notion of a “*circual*” quite rigorously, as the set of all lines having a fixed *twistance* t from a given line. In other words, a *circual* is the envelope of all lines whose bipolar coordinates (x, y) satisfy the familiar-looking equation

$$(x - a)^2 + (y - b)^2 = t^2.$$

The line (a, b) is the central of the *circual*, and t is (the swingth of) its radial.

Although it is far from self-evident, it turns out, when one calculates it out,



On the Euclidual plane, a circual is defined to be the envelope of all lines equitwistant from a given line (the central of the circual). Since we are here looking at a tubal realization of the Euclidual plane, a line is realized as an arbitrary ellipse centered on the O-point. We wish to find all such lines that are equitwistant from a particular such line. As our central, we shall choose the O-circle itself — the circular cross-section of the tube directly above the x-axis.

A variable line is determined by twisting the H- and V-toothpicks. The twistance of the line from the central is given by the Pythagorean formula $\sqrt{h^2 + v^2}$, where h and v are the twists of the two toothpicks (both h and v can range between $+\infty$ and $-\infty$). The equation of this circual is therefore $\sqrt{h^2 + v^2} = t$, where t , a constant, is the radial of the circual. This equation is more commonly written

$$h^2 + v^2 = t^2.$$

It turns out that the envelope of all such lines (i.e., tubal ellipses centered on the O-point) is a pair of circles running around the tube, symmetrically located with respect to the O-point. If $t = 0$, then the two circles coincide and equal the O-circle. As t increases in magnitude, they separate.

If an arbitrary tubal line rather than the O-circle is chosen for the central of the circual, then the equation of the circual becomes

$$(h - h_0)^2 + (v - v_0)^2 = t^2,$$

where h_0 and v_0 are the bipolar coordinates of the line chosen to be the central. In this case, a circual turns out to be a pair of ellipses on the tube, again located symmetrically with respect to the O-point, and identically tilted.

that circuals on the tube are, amusingly enough, shaped like circles and ellipses! However, unlike the ellipses that correspond to lines on the plane, these circles and ellipses are *not* constrained to be centered on the O-point (or in other words, the planes containing them are not constrained to pass through the O-point). Thus a simple example of a circual is any circular cross-section of the tube located at any y -coordinate other than 0. This is in fact what circuals whose central is the origin look like — that is, circuals all of whose lines are equidistant from the O-circle. As the “radial” — that is, the twistance from the central, given by t — is varied, the circle-that-is-a-circual slides along the axis of the tube. When the radial is zero, then the circual coincides with the O-circle $y = 0$. This is the smallest possible circual whose central is the origin.

From my statement that circuals can be located at *negative* as well as positive y -coordinates, readers might be wondering whether such circuals would not then belong to negative radials. Actually, the fault is mine: I slightly oversimplified the truth about circuals. A circual on the tube consists not of *one* circle or ellipse, but of a *pair* of them, symmetrically constructed with respect to the O-point. In other words, for every point on a circual, its antipodal point (defined as its exact reflection through the O-point) is also on the circual. So in fact, there is no difference between a negative and a positive radial, as one might hope.

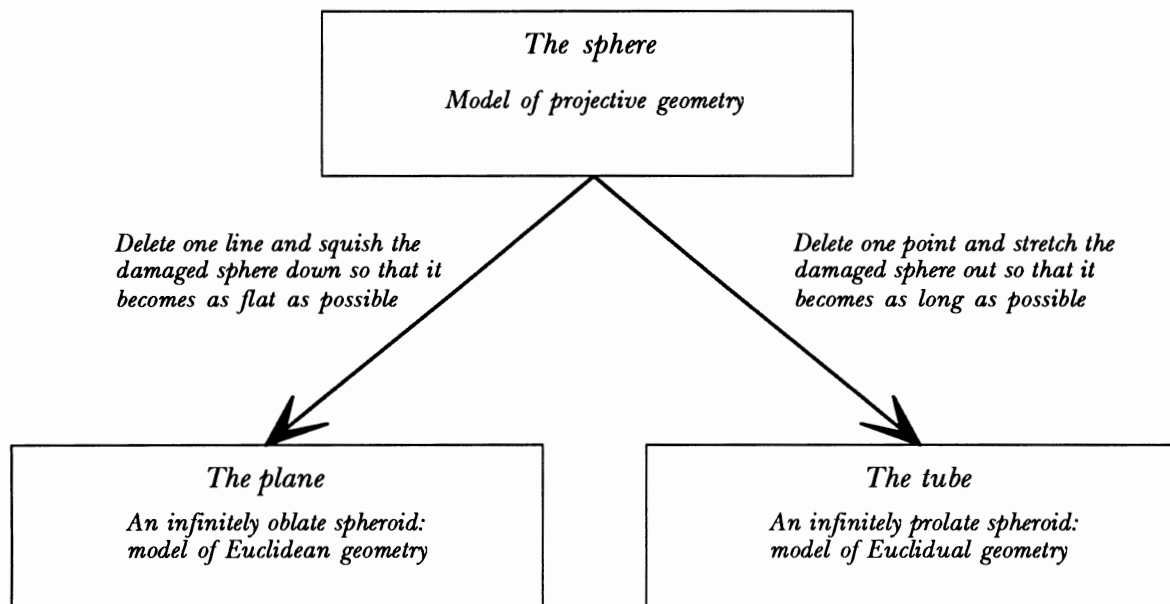
Now that we know how circuals on the tube look, you might well ask, “What does a circual look like *in the Euclidean-plane model*?” To answer this, all we need to do is to use our quasi-gnomonic projection to get from the tube down onto the plane. And here, our knowledge of conic sections comes in handy. Recall that what you get when you slice any cone by a plane parallel to its axis is a hyperbola. Well, it turns out that this is what is going on in our quasi-gnomonic projection, and so *hyperbolas* are the representatives of circuals in the Euclidean-plane model.

To see that the quasi-gnomonic projection of a pair of antipodally symmetric ellipses on a tube down onto a plane amounts to slicing a cone by a plane parallel to its axis is not nearly as hard as it sounds — in fact, it is trivial. First of all, the plane involved is of course the Euclidean plane on which the tube is sitting. What is the cone? Well, its apex is the O-point, and the cone itself is the three-dimensional shape traced out by all the lines connecting the O-point to points on the ellipses. Why is this a cone? Well, in the special case when the two ellipses are circles, this shape is obviously a cone. In the general case, it is almost as easy to see; it suffices to recall that a cone sliced obliquely by a plane gives an ellipse. Now just run that statement backwards: given an ellipse in a plane, if you connect all of its points to some point *not* in that plane, you get a cone. Put the Euclidean plane back into the picture, and we now have a *bona fide* cone and a plane — not just any plane, but a plane parallel to the cone’s axis — cutting it. *Voilà* — a hyperbola!

There are of course many, many further intriguing questions that might pop to mind. We have but scratched the surface of Euclidal geometry, but this gives a good flavor for the subject. In what follows, we shall make some concluding philosophical musings.

On Abstract Geometries and Their Concrete Realizations

Intuition suggests that there must be some close relationship between the sphere as a model of projective geometry and the tube as a model of Euclidal geometry. This is in fact the case, but to see things with maximal clarity, we need to bring in a third element — namely, the (Euclidean) plane as a model of Euclidean geometry. The three-way relationship can now be portrayed concisely, as follows.



Not surprisingly, there is a kind of duality between the natural *planar* model of Euclidean geometry and the natural *tubal* model of Euclidal geometry. More specifically, a plane and a tube can be looked upon as maximally opposite types of distortions of the sphere. In particular, a plane is what you get when you remove a sphere's equator and then "squish" the remainder of the sphere completely flat — in other words, an infinite plane is a *maximally oblate spheroid*. By contrast, a tube is what you get when you remove one pair of diametrically opposite points from a sphere and then "stretch" the remainder of the sphere so that it is infinitely long — in other words, an infinite tube is a *maximally prolate spheroid*. This, then, is the connection between the sphere, the plane, and the tube.

We have made continual reference to various objects and spaces as *models* of one or another variety of geometry. The implication would seem to be that there is something less than genuine about these "models", in the sense that a model train is not a real train. However, it is time to discredit this view. Probably the best way to put it is that a term like "projective geometry" or "Euclidal geometry" stands for an *abstraction* rather than for an actual space. As such, projective geometry can be *realized* in various different ways. So it would be more accurate to say that the sphere and the extended Euclidean plane are two different concrete *realizations* of projective geometry. These are to projective geometry what two actual physical trains are to the abstract *concept* "train". They are completely genuine, full-fledged instances of projective geometry rather than half-gauge or quarter-gauge "models" that aren't quite the real thing.

This is no less true of the famous "models" of hyperbolic and elliptic geometry (the two most famous non-Euclidean geometries), such as the spherical model of elliptic geometry, or Poincaré's and Beltrami's ways of embedding hyperbolic geometry in the Euclidean plane. Geometry books always use the term "Poincaré

model", as if it were good for giving you an *idea* of what hyperbolic geometry *really* is, but implying that it is certainly not the real McCoy itself, even though every single theorem of hyperbolic geometry holds in the alleged model. (Sometimes a three-dimensional model using a so-called "pseudosphere" is even trotted out as the *genuine* hyperbolic geometry, in contrast to mere *models*.) Once again, the implication is wrong. Henri Poincaré, Eugenio Beltrami, and Felix Klein found full realizations of the abstraction known as "hyperbolic geometry". Neither is imperfect or less than complete. The fact that the Poincaré model is embedded in the Euclidean plane and that it requires calling arcs of circles "lines" might make some people think it is some kind of "cheat" or "trick", but it certainly is not.

By the way, just as one can tamper with the postulates of Euclidean geometry and come up with non-Euclidean geometries, so one can tamper with the postulates of Euclidual geometry and come up with *non-Euclidual* geometries. For instance, whereas in Euclidual geometry, exactly *one* point parallel to a given point can be found on any given line, there is a non-Euclidual geometry in which *no* such parallel point can be found (known as "elliptual" geometry, of course), and another non-Euclidual geometry in which *many* such parallel points can be found ("hyperbual" geometry). The Euclidual mathematicians Henriette Ligneronde, Eucretina Bellatrima, and Felicia Gross even discovered ways to simulate hyperbual geometry on the Euclidual plane! The fact that the Ligneronde model is embedded in the Euclidual plane and that it requires calling arcs of circuals "points" might make some people think it is some kind of "cheat" or "trick", but it certainly is not.

We have so far seen two models — that is, realizations — of projective geometry, and two of Euclidual geometry. What about models of good old Euclidean geometry? It might seem strange to suggest that the Euclidean plane might be merely *one possible way* of realizing the abstraction called "Euclidean geometry" — yet if non-Euclidean geometries admit of multiple models, why should the same not hold for Euclidean geometry?

Indeed, it turns out that both of the famous non-Euclidean geometries contain models of Euclidean geometry, just as Euclidean geometry contains models of both of them. Tit for tat! Such models allow the denizens of non-Euclidean worlds to develop an intuition for what, to them, is the highly counterintuitive subject of Euclidean geometry! Of course, to us, such models would seem doubly weird, since they give us an inkling of how our world looks to "people" whose world looks very strange to us.

In this paper, it turns out that we have inadvertently come up with yet another model of Euclidean geometry — namely, Euclidual geometry itself. All we need to do is call Euclidual lines "points" and Euclidual points "lines", and then we can interpret any of our pictures of phenomena in the Euclidual plane as illustrating — admittedly, in an extremely bizarre way — truths of Euclidean geometry. We also have a bizarre model of Euclidean geometry on the Euclidual tube as well as on the plane. Thus a tubal ellipse centered on the O-point could be called a "point" instead of a "line", and the set of all such ellipses that pass through a given antipodal point-pair could be called a "line" instead of a "point". Then, thanks to our careful analysis of Euclidual geometry, we would have an extremely hard-to-follow but genuine realization of Euclidean geometry on the surface of the tube.

This type of conceptual double-reversal is delightful to contemplate but a bit too mind-boggling to be worth following out any further. The main point is simply that Euclidean and Euclidual geometries, being each other's duals, are *ipso facto* models of each other. In that sense, they are indistinguishable.


This brings us back to the difference discussed earlier between, on the one

hand, taking the terms of geometry as *undefined*, and on the other hand, bringing to them prior images and connotations. If we really bring no imagery whatsoever to the field of geometry, then “point” and “line”, when discussed in projective geometry, should be as interchangeable in our minds and our speech as “left” and “right” — in fact, even more so! As a consequence, the ideas of deleting a point and deleting a line should feel no more distinct from each other than driving on the right side of the road and driving on the left side. Someone who felt this way would think, “Big deal — deleting a point and deleting a line are as alike as Tweedledum and Tweedledee!” To such high-powered but imagery-free thinkers, Euclidean and Euclidal geometries would have to be completely boring variants of each other, conveying no new imagery and no new results whatsoever.

To us lesser beings, however, there seems to be a vast difference. Lines, to us, really *are* infinitely long non-closed structures made of points, and points really *are* infinitesimal spots about which lines can be freely rotated as often as one wants. We feel as if the world we live in is genuinely Euclidean, not Euclidal. Indeed, we find it hard if not impossible to imagine what it would be like to live in a Euclidal world. How could it ever happen that lines could *not* freely rotate around points? How could it ever happen that running along a straight line would eventually bring you back home? Even if cosmology reveals that space itself is curved and closed, that will not make things any better; it will simply tell us that the universe we live in violates our deep and primordial intuitions about space.


My point is this: our thoughts and imagery about space are fundamentally and deeply Euclidean, and trying to get used to any other system requires great mental effort on our part, and even after expending such effort, we remain contaminated by many unconscious Euclidean prejudices. Thus, despite the formal symmetry, even interchangeability, of Euclidean and Euclidal geometries, something in us tells us that *we are Euclideans, not Euclidals*.

Similar triangles on the Euclidean plane



The two triangles shown above are similar because their sides are parallel in pairs (i.e., each pair of corresponding sides shares an ideal point).

Similar trislices on the Euclidean plane



The two trislices shown above are similar because their vertices are parallel in pairs (i.e., each pair of corresponding vertices shares an ideal line).

Two triangles are *similar* if they have the same interior angles.

Two triangles are *congruent* if they have the same interior angles and their sides have the same lengths.

The operations of *translation* and *rotation* are both shape- and size-preserving — that is, each of them preserves both relative angles and relative distances.

Translation of any planar figure is carried out as follows: Add a pair of constants (x_0, y_0) to the biaxial coordinates (x, y) of each point in the figure.

Rotation of any planar figure is carried out as follows: Given a center of rotation O (any point) and any component point P of the figure, shift the line OP by a fixed angle (*i.e.*, add a constant value θ_0 to the angle θ that line OP makes with some fixed reference line in the plane); then on the new line, find a point P' having the same distance from O as P did.

The operation of *dilation* is shape-preserving — that is, it preserves both relative angles and ratios of distances.

Dilation of any planar figure is carried out as follows: Uniformly scale all distances as measured from a fixed center point while keeping all angles (measured relative to an arbitrary reference line through that point) fixed.

In Euclidean geometry, exactly one line parallel to a given line can be found through any given point.

In *elliptic* geometry (one variety of non-Euclidean geometry), no such parallel line exists.

In *hyperbolic* geometry (another variety of non-Euclidean geometry), an infinite number of such parallel lines exist.

Two trislices are *similar* if they have the same interior slides.

Two trislices are *concordant* if they have the same interior slides and their vertices have the same swings.

The operations of *twistation* and *relocation* are both shape- and size-preserving — that is, each of them preserves both relative slides and relative twistances.

Twistation of any planar figure is carried out as follows: Add a pair of constants (x_0, y_0) to the bipolar coordinates (x, y) of each line in the figure.

Relocation of any planar figure is carried out as follows: Given a central of relocation o (any line) and any component line p of the figure, shift the point op by a fixed slide (*i.e.*, add a constant value ϕ_0 to the slide-value ϕ that point op has with reference to some fixed point in the plane); then on the new point, find a line p' having the same twistance from o as p did.

The operation of *oblation* is shape-preserving — that is, it preserves both relative slides and ratios of twistances.

Oblation of any planar figure is carried out as follows: Uniformly scale all twistances as measured from a fixed central line while keeping all slides (measured relative to an arbitrary reference point on that line) fixed.

In Euclidal geometry, exactly one point parallel to a given point can be found on any given line.

In *elliptual* geometry (one variety of non-Euclidal geometry), no such parallel point exists.

In *hyperbual* geometry (another variety of non-Euclidal geometry), an infinite number of such parallel points exist.

The View from Euclidualia

What would it be like to be a Euclidual? This is a question one cannot help asking oneself as one grapples with Euclidual geometry. One thing that comes to mind very quickly is that a Euclidual would have to wonder just as earnestly, "What would it be like to be a Euclidean?" In fact, we clearly understand that for each aspect of Euclidual geometry that is boggling to us, there is a corresponding and isomorphic aspect of Euclidean geometry that is precisely as boggling to our Euclidual counterparts. For instance:

"What sense does it make to say that a point consists of many lines?", we wonder. At the same time, Eucliduals wonder, "What sense does it make to say that a line consists of many points?"

We wonder, "How can there be a non-zero twistance between parallel lines?" At the same time, Eucliduals wonder, "How can there be a non-zero distance between parallel points?"

One of the stranger things about Eucliduals is how different their visual systems must be from ours. Eucliduals and we bring very different eyes to one and the same figure. What simpler Euclidean concept is there than that of *congruent triangles*? To our eyes, two congruent triangles are *identical*. Yet to Euclidual eyes, they have nothing in common. To come to grips with this anomaly, we must carefully analyze why it is that we Euclideans think two congruent triangles are identical. The answer is quite simple: we can *shift* either triangle and make it coincide perfectly with the other. And this is where Euclidean prejudices creep in — namely, prejudices about what types of "shifting" are allowed in bringing two figures into coincidence. For us, operations that preserve *angles* and *distances* are the right kinds of shifting-operations — to our eyes, they are "shape-preserving". Thus, rotation and translation preserve *angles*, but of course Eucliduals don't care a fig about angles — they are concerned with *twistances*. Similarly, translation and rotation preserve *distances*, but unfortunately, Eucliduals find distances a completely silly way of looking at separation between points — what they care about is *slides*.

This means that to get a feeling for Euclidual vision, we must find the Euclidual analogues to our operations of translating and rotating figures in the plane. Let us first consider translation. Translation of a figure consists in:

Adding a pair of constants (x_0, y_0) to all the biaxial coordinates (x, y) of points constituting the figure.

The Euclidual version of this operation would thus be:

Adding a pair of constants (x_0, y_0) to all the bipolar coordinates (x, y) of lines constituting the figure.

Note that this operation is designed specifically to *preserve all twistances between lines*, just as translation preserves all distances between points. (Think about how the x_0 and y_0 will get canceled out in the Pythagorean formulas for distance and twistance.) We shall therefore baptize this operation with the name *twistation* (the phonetic resemblance to "translation" is of course not accidental).

Now what about the dual to rotation? Well, rotation is a slightly more intricate operation, but basically, it is the following.

Given a center of rotation O (any point) and any component point P of the figure, shift the line OP by a fixed angle (*i.e.*, add a constant value θ_0 to the

angle θ that line OP makes with some fixed reference line in the plane); then on the new line, find a point P' having the same distance from O as P did.

We can dualize this pretty straightforwardly, to make a new operation that we shall call *relocation* (again, the phonetic resemblance to "rotation" is intended).

Given a central of relocation o (any line) and any component line p of the figure, shift the point op by a fixed slide (*i.e.*, add a constant value ϕ_0 to the slide-value ϕ that point op has with reference to some fixed point in the plane); then on the new point, find a line p' having the same twistance from o as p did.

For Eucliduals, relocation and twistation play analogous roles to our rotation and translation, respectively. And therefore, the figures that Eucliduals effortlessly see as having "the same shape" look nothing like each other, to us. Two trislides that "look alike" to Eucliduals — that is, that can be brought into exact coincidence with each other by relocation and twistation — are called *concordant*.

What about the Euclidual counterpart of *similar* triangles? Two triangles are similar if they have exactly the same set of three angles. That's easy to dualize: two trislides are *simular* if they have exactly the same set of three slides. (Note that sharing just two slides is enough, since the size of the third slide is determined by the condition that the sum of all three slides must equal π .) Just as two similar triangles look, well, *similar* to us, although not identical, so to Eucliduals, two similar trislides appear, well, *simular*, although not identical.

Similarity in Euclidean geometry can also be defined by means of the operation known as *dilation* — the uniform scaling-up or scaling-down (*i.e.*, multiplication by a constant factor) of all distances from a fixed center point. This operation too has a Euclidual counterpart that we shall name *oblation*. Oblation is the uniform scaling-up or scaling-down (*i.e.*, multiplication by a constant factor) of all twistances from a fixed central line.

We have now analyzed, to some extent, what makes two figures look exactly the same to possessors of one type of visual system and utterly different to possessors of the dual visual system. Despite all of this intellectual analysis, we still cannot help but think to ourselves, "Euclidual eyes must be very strange indeed! How could anyone possibly look at two perfectly congruent triangles and not see them as having the same shape, yet look at two utterly dissimilar 'concordant' triangles and see them as indistinguishable?" But we know that at the very same time, our Euclidual counterparts are thinking to themselves, "Euclidean eyes must be very strange indeed! How could anyone possibly look at two perfectly concordant trislides and not see them as having the same shape, yet look at two utterly dissimilar 'congruent' trislides and see them as indistinguishable?"

These are truly irreconcilable worldviews; still, we have tried our best to plunge ourselves into the Euclidual universe and to "see it as the Eucliduals do". Yet in some sense, the very best way of understanding their universe would be to realize that *their universe looks to them exactly as our universe looks to us*. If this is the case, however, then they must feel about lines exactly as we feel about points, and they must feel about points exactly as we feel about lines. So once again it sounds as if the only difference between our universes is a notational one — what words one uses for what things.

If that were all there were to it, though, then we certainly wouldn't be so convinced that our universe is Euclidean and *not* Euclidual. Could it be that we are fooling ourselves? Does our intuitive conviction that lines are *translationally*

infinite actually come from dealing with points that are *rotationally* infinite, except that we code it wrongly in our brains? Are we all born wearing invisible “inverting spectacles” that convert physically rotatory experiences into linear ones in the mind, and vice versa (a somewhat extreme variation on the famous theme of inverting spectacles that reverse right and left)? Are we humans *actually* living in a Euclidal universe that we erroneously *perceive* as Euclidean? Are we Euclidals who just *think* they are Euclideans?

It would seem that this is an unanswerable question. It certainly cannot be answered via experimentation, the way that Gauss sought to determine whether the space we live in is Euclidean or non-Euclidean by measuring the sum of the angles of a very large triangle. It seems that we simply innately think we live in a Euclidean world, and a Euclidal world makes no sense to us. Why this psychological asymmetry? Is it a property of the human brain, or a property of the universe, or what?

A Few Speculative Remarks on Physics and Euclidal Geometry

One cannot help but wonder if this new perspective on what space might be like could conceivably have anything to do with the physics of our universe. One thing is certain: the laws of physics often treat rotation and translation in very similar yet not identical ways. For example, there is the translational concept of *linear momentum*, and the analogous rotational concept of *angular momentum*. Both are conserved under all circumstances, as far as we know. Yet angular momentum is *quantized* (restricted to certain specific values) whereas linear momentum is not — it can assume any value whatsoever. (This holds for free particles; for bound particles, linear momentum is also quantized.) What does this asymmetry mean? Where does it come from? What would a world be like in which the reverse were the case — namely, a world where linear momentum was always quantized, but angular momentum was not? Would that kind of world perchance be Euclidal?

A second potential link to physics is suggested by a disturbing point/line asymmetry in electromagnetism — the fact that the electric field admits of *point* sources (*i.e.*, charged particles) but no *line* sources, while the magnetic field admits of *line* sources (*i.e.*, charged currents) but no *point* sources. The lack of point sources for the magnetic field deeply bothered the great English physicist P. A. M. Dirac, and he worked out a theory of what he called *magnetic monopoles*, which are isolated magnetic north and south poles — something that so far has never been seen in our universe. In Dirac’s monopole theory, lines and points act very counterintuitively, in a manner somewhat reminiscent of the reversals in Euclidal geometry. Could it be, then, that the asymmetry of electromagnetic theory as embodied in Maxwell’s equations is the result of an asymmetric way of “breaking” a more “projective” (*i.e.*, fully symmetric) sort of physics, in which there is a complete and perfect symmetry between electric and magnetic phenomena?

This is a very speculative idea, but perhaps someday someone will follow it out and discover some new physical meaning to the concepts of projective geometry.

Geometries of “A Somewhat Bizarre Nature”

My explorations into the disorienting Euclidal world were conducted in utter solitude over a several-week period, and of course I was deeply curious the whole time about whether I was blazing new pathways or just trodding down old ones. I

looked with great care at all my books on geometry, which number several dozen, to see if I could find any references to such things as “parallel points” or other tell-tale phenomena of the Euclidual world. But I came up with nothing.

Finally, though, on a day when I was putting the very finishing touches on this article, I opened up a rather old book on non-Euclidean geometry by D. M. Y. Sommerville (published in 1914), merely seeking the name of the three-dimensional shape that is often presented as a model for hyperbolic geometry. I found it — “pseudosphere” — but while riffling through Sommerville’s pages, I also found something else. I saw a picture that looked exactly like what I have herein called a “trislid”, and its caption said something about the perimeter of a triangle being constant and equaling a multiple of π . That certainly caught my eye! I eagerly looked for more discussion and for references, but there was very little given. Indeed, the entire passage, not even one page long, was set in smaller type than the rest of the text, indicating that it was something of a digression and not very important, and began as follows:

The other geometries, in which the measure of angle is either hyperbolic or parabolic, are of a somewhat bizarre nature.

For example, if the absolute degenerates to two imaginary lines ω, ω' , and two coincident points Ω , the case is just the reciprocal of the Euclidean case; linear measurement is elliptic, K being imaginary, and angular measurement is parabolic, k being infinite. In this geometry the straight line is of finite length = πKi .

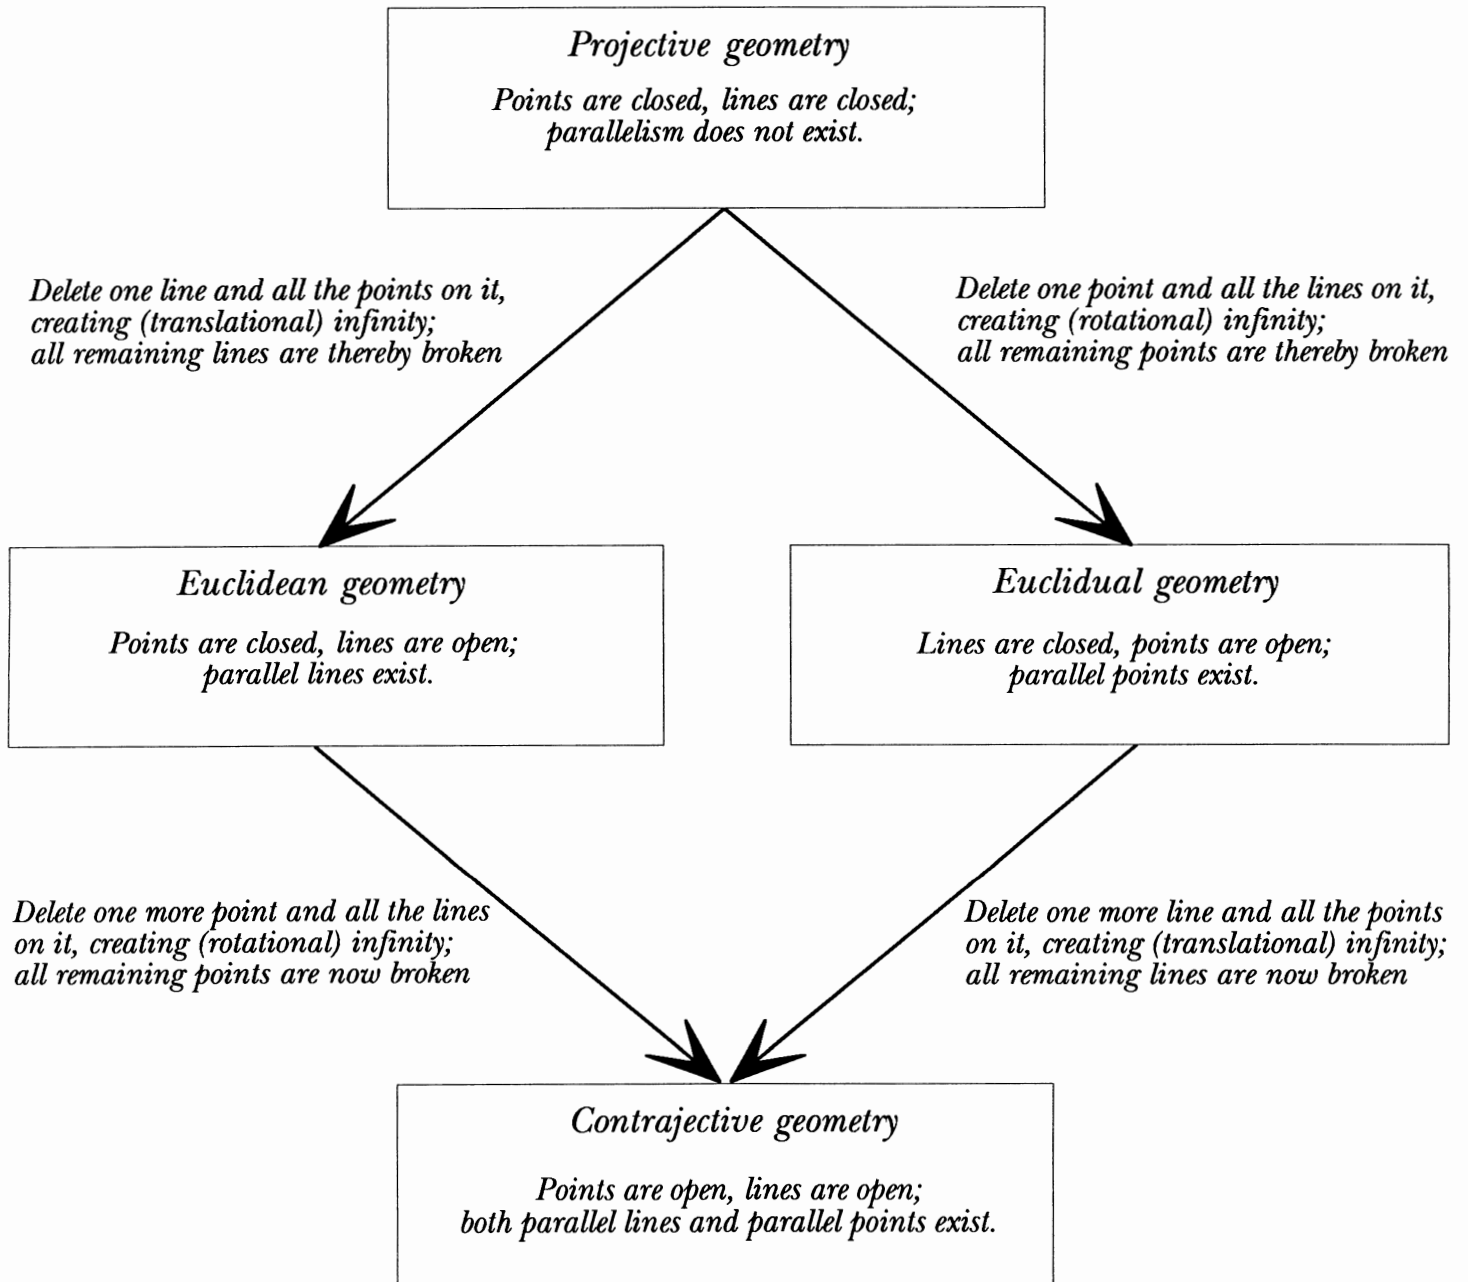
Sommerville really said very little more than that, other than describing the appearance of a trislid and stating the fact about its constant perimeter. No mention was made of the existence of parallel points, which I would have thought would be one of the most curious and striking aspects of this “somewhat bizarre” geometry. And there were no references at all! Clearly, however, people almost a century ago had certainly sniffed at, if not fully inhaled, some of these ideas. What I read also seemed to imply that projective, Euclidean, Euclidual, elliptic, hyperbolic, and other geometries known at the time were merely a few isolated points (or lines!) in an infinite space of geometries, defined by several continuous parameters. This was a little humbling and a little disappointing to me, but not seriously so.

Conclusion: Making a Diamond of Geometries

In any case, I had one last idea up my sleeve, and my intuition somehow told me it was genuinely new, perhaps even rich in implications. That idea involved the restoration of full symmetry to the Euclidean and Euclidual universes. Of course my idea was not to put back the missing line or point, since that would just result in projective geometry once again. Instead, the idea was to go in the opposite direction. That is to say, if Euclidean geometry results from “damaging” the projective plane by deleting one line, and if Euclidual geometry results from damaging it by deleting one point, why not *doubly* damage the projective plane, by deleting from it both one line *and* one point?

The result of such an operation, which I have dubbed *contrajective geometry*, would be totally symmetric — in other words, self-dual, just like projective geometry — but it would feature lines that are open and points that are open, in contrast to projective geometry’s closed lines and points. It would also feature both parallel lines and parallel points, all in one universe. Moreover, the natural pair of shape-preserving shifting-operations in contrajective geometry would not be the

The fully symmetric Diamond of Geometries



Euclidean pair (translation and rotation) nor the Euclidual pair (twistation and relocation), but rather a hybrid pair — translation and twistation — which draws one member from each of the “broken” geometries.

Contrajective geometry fits in a symmetric position below both Euclidean and Euclidual geometry, and thus completes a seemingly inevitable quartet — the “Diamond of Geometries” alluded to in the title of this article.

What is contrajective geometry like? What are some of its theorems? One thing for sure is that neither the sum of the interior angles nor the sum of the interior slides of all three-sided, three-vertexed entities is invariant, because the notions of “angle” and “slide” are irrelevant in this geometry. Only distances and twistances need apply here! But what exactly does this mean? How to visualize it? Is there some kind of natural model for contrajective geometry, playing the role that the sphere, the plane, and the tube do for the other three geometries? I must admit that at this stage I do not have the foggiest idea; these remain beckoning questions for the future.