STAG: Smart Tools and Applications in Graphics (2021)
P. Frosini, D. Giorgi, S. Melzi, and E. Rodolà (Editors)

Straightedge and Compass Constructions on Surfaces

C. Mancinelli[1] and E. Puppo[1]

[1]DIBRIS - University of Genoa, Italy

Abstract

We discuss how classical straightedge and compass constructions can be ported to manifold surfaces under the geodesic metric. After defining the equivalent tools in the manifold domain, we analyze the most common constructions and show what happens when trying to port them to surfaces. Most such constructions fail, because the geometric properties on which they rely no longer hold under the geodesic metric. We devise some alternative constructions that guarantee at least some of the properties of their Euclidean counterpart; while we show that it is usually impossible to guarantee all properties together. Some constructions remain still unsolved, unless additional tools are used, which violate the constraints of the straightedge and compass framework since they take explicit distance measures. We integrate our constructions in the context of a prototype system that supports the interactive drawing of vector primitives on a surface represented with a high-resolution mesh.

1. Introduction

The ancient Greek mathematicians developed a set of geometric techniques, which go under the name of *straightedge and compass constructions*, to draw a number of planar geometric figures and arrangements, involving straight lines, circles, and angles. The peculiar aspect of such constructions is that they do not require taking any explicit measure. More formally, the only permissible constructions are those granted by Euclid's first three postulates, and they are based on two ideal tools: the straightedge, which can be positioned between any two points and extended indefinitely in both directions; and the compass, which can trace circles of any radius by starting with its needle and pencil points at two given points in the plane. Besides, all points of intersection between straight lines and circles drawn with such tools can be found.

The straightedge and compass constructions can be used to define vector graphics in the plane. In fact, several graphics primitives and constructions made available in the GUI of drawing systems can be addressed with such basic tools. This work is part of our effort to bring vector graphics to the manifold domain, i.e., by assuming a surface as a canvas [MNPP21, NPP21]. Here, we investigate to which extent the straightedge and compass constructions can be ported to the manifold setting, by using equivalent tools.

On a surface, a geodesic line is the counterpart of a straight line in the plane; and a geodesic circle is defined as an isoline of the distance field from a given point. Similarly to the Euclidean case, we will take for granted that the shortest geodesic connecting two points on a surface can be found and extended indefinitely on both sides; that the isoline of the distance field from a point and passing through another point can be traced; and that the intersections between any two such curves can be found.

The main challenge here is that constructions in the Euclidean context rely on geometric properties that no longer hold in the geodesic metric. In fact, even the basic properties of straight lines and circles do not hold on a surface without additional conditions. For instance, a long-enough geodesic line may self-intersect; there might exist infinitely many geodesic segments of different lengths joining two distinct points; and even the shortest geodesic segment between a pair of points might not be unique. Similar issues arise for circles: a generic isoline of the distance field from a point is guaranteed neither to be homeomorphic to a standard circle, nor to be smooth at all points; and equal angles at the center do not intercept equal chords or arcs on the circle.

Nonetheless, geodesic line segments and circles are somehow well-behaved as long as they the are "small enough". In particular, we will restrict our study to constructions that occur inside a *strongly convex region* (see Sec. 3 for a formal definition). Note that, differently from the Euclidean case, convex sets on a manifold cannot be too big; in particular, they cannot cross the cut locus of anyone of their points. Inside convex regions, geodesic lines and circles are smooth and behave like in the Euclidean plane from a topological point of view, i.e., in the way they intersect.

Even in this favorable case, most metric properties, which are crucial for the straightedge and compass constructions, are lost. In Sec. 4, a quick review of the main constructions reveals that most of them fail, or produce limited results, once they are lifted to the manifold setting. In order to overcome such limitations, we resort to additional tools, such as the square set to trace perpendicular lines, and the composition of distance fields to trace further isolines, which are neither geodesics, nor arcs of circles, but have some of their properties in the Euclidean case.

Our extended set of tools is sufficient to support efficiently most

DOI: 10.2312/stag.20211469

constructions necessary to create vector graphics on surfaces. On this basis, we have implemented a first prototype system, which we demonstrate in Sec. 5.

2. Related work

The straightedge and compass constructions rely on basic theorems of the Euclidean geometry, which relate lengths and angles. If one tries to find similar relations on a manifold surface, then curvature must be taken into account. This subject was thoroughly investigated during the last two centuries, in the context of the theory of intrinsic geometry of surfaces. To the best of our knowledge, Rodrigues was the first mathematician to address the problem [Rod16], using a spherical mapping to study the ratio of the areas of corresponding surfaces. His work provides the first definition of intrinsic curvature, as formalized later by Gauss in his *Theorema Egregium* in 1827. Many results followed, which compare the geometry of a general manifold with that of a model space with constant curvature. See [CE75, Cha06] for a comprehensive account on the subject.

Referring just to cases addressed in this paper, a corollary of the Gauss-Bonnet theorem relates the internal angles of a geodesic polygon to the curvature of the region it encloses. Such result explains the challenge in defining an isosceles geodesic triangle (see Section 4.7) and, more in general, in addressing constructions that require geodesic lines of given lengths *and* forming given angles.

Alexandrov investigated thoroughly the relations between quantities measured on a surface with their counterpart on surfaces with constant curvature (a.k.a. CAT – Cartan-Alexandrov-Topogonov – spaces) [Ale48]. In a nutshell, it turns out that geodesic lines, which are cast from a common source along different directions, tend to converge if the curvature of the space is positive, and to diverge if it is negative. Based upon these facts, many comparison theorems involving Alexandrov and CAT spaces have been proposed in the literature. See [AKP19] for a recent account on this subject; interestingly enough, the title of the chapter addressing geodesic triangles is *The ghost of Euclid*.

The literature concerning tools for geodesic computations is vast and has been recently surveyed in [CLPQ20]. The specific methods we rely on are summarized in Sec. 5.

3. The geodesic arsenal

Definitions. Let S be a smooth surface embedded in $\mathbb{R}^3$. The embedding induces a Riemannian metric, defining the length $L(\gamma)$ of any curve γ on S. The *geodesic distance* $d(p,q)$ between two points $p,q \in S$ is the infimum of length L over all curves γ having their endpoints at p and q; one such curve γ_{pq} satisfying $L(\gamma_{pq}) = d(p,q)$ is called a *shortest geodesic path* between p and q.

A shortest geodesic path may not be unique. The *cut locus* of a point p is defined as the closure of the set of points that can be connected with p with more than one minimal geodesic; the injectivity radius r_p is in fact the distance of p from its cut locus. A *normal ball* centered at p is a ball that does not intersect the cut locus of p. A set $U \subset S$ is said to be *strongly convex* if for each pair of points

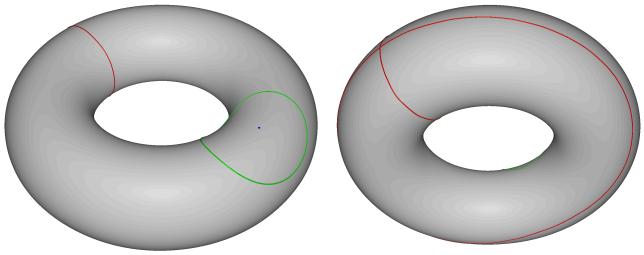

Figure 1: *The cut locus (in red) of a point p (in blue) on a Torus (front and back view. The convex ball centered at p (in green) cannot cross the outer equator, since otherwise it would contain pairs of points that are connected by a geodesic that do not belong to the ball. The cut locus has been computed with the method proposed in [MLP21].*

p and q in U there exists a unique shortest geodesic path γ_{pq} connecting p to q in S which is entirely contained in U, and moreover this property holds also for every open ball contained in U. By definition, a convex set U cannot extend beyond a maximal normal ball centered at any of its points. Figure 1 shows an example of cut locus of a point and a strongly convex ball centered at that point.

Geodesic curves can be also characterized by their *straightness*. In order to assess the curvature of lines in the intrinsic geometry of S, one needs to introduce the covariant derivative, which we omit here for brevity. Intuitively, from an extrinsic point of view, a geodesic curve γ does not make any further turn except the strictly necessary to follow the curvature of S: it turns *with* S, but it does not turn *on* S. Thus, geodesics play the role of straight lines on S.

Similarly to straight lines in Euclidean space, a geodesic curve may extend indefinitely and is completely defined by a point p and its tangent vector in the tangent plane T_pS at p. The *exponential map* $\exp_p : T_pS \longrightarrow S$ maps vectors of the tangent plane to points on the surface, where point $\exp_p(t)$ is defined as the other endpoint of the geodesic path traced from p in the direction of t for length $|t|$. In general, the exponential map is not injective; the *injectivity radius* of p is the maximum radius r_p such that $\exp_p(t)$ is invertible at all $t \in T_pS$ such that $|t| < r_p$.

Let γ and γ' be two geodesics having a common endpoint at p; the angle between them at p is defined from their tangent directions in the tangent plane T_pS. See Fig. 2 for an example.

From now on, we will assume that all our constructions will be contained in a strongly convex ball.

Geodesic tools. We will assume that the following primitives for geodesic computations are available. Their implementation is briefly addressed in Sec.5.

- *Shortest-path:* given points $p,q \in S$, return the shortest geodesic path γ_{pq} connecting them;
- *Tangent:* given a curve γ on S and one of its points p, return the direction $t \in T_pS$ tangent to γ at p;
- *Geodesic-tracing:* given point $p \in S$ and a tangent direction $t \in T_pS$, trace a geodesic through p and tangent to t at p;

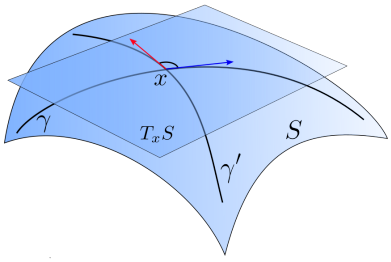

Figure 2: *Two geodesic lines γ and γ' intersecting at point $x \in S$ form an angle defined by their tangents at x on the tangent plane T_xS (red and blue arrows).*

- *Distance-field:* given $x \in S$, compute the field $d_x : S \longrightarrow \mathbb{R}$ where $d_x(y) := d(x,y)$;
- *Isoline:* given field d_p and a point $q \in S$ return the isoline of d_p that goes through q;
- *Intersect:* given any two lines on S, not necessarily geodesic, return their intersections.

Note that *Shortest-path* only allows us to trace segments between two endpoints. However, once we have one such segment, the joint use of *Tangent* and *Geodesic-tracing* allows us to extend it indefinitely from both sides. We thus define the derived operation *Geodesic-line* that traces an arbitrarily long line through a and b. The *Geodesic-line* generalizes the straightedge to the manifold setting. Likewise, the joint use of the *Distance-field* and *Isoline* primitives allows us to obtain the derived operation *Geodesic-compass*, which generalizes the compass.

Note that, the *Tangent* and *Geodesic-tracing* operators allow us to bring directions from the surface to the tangent plane, and, in the opposite direction, straight lines from the tangent plane to the surface; this will allow us to directly exploit local constructions about angles in the Euclidean setting.

Conversely, the *Distance-field* alone does not belong to the straightedge and compass framework. On the other hand, since the *Distance-field* is anyhow necessary to implement the *Geodesic-compass*, we will use it directly to address constructions where the basic tools fail, without taking explicit measures though.

4. Most used straightedge and compass constructions

Straightedge and compass constructions consist of iteratively applying the following five basic constructions:

1. creating the line through two existing points;
2. creating the circle through one point with center another point;
3. finding the intersection point of two existing, non-parallel lines;
4. finding the intersection points of a line and a circle;
5. finding the intersection points of two circles.

We address the above constructions in the manifold setting by means of *Geodesic-line* (1); *Geodesic compass* (2); and *Intersect* (3, 4, 5). Note that, within a strongly convex set U, it is guaranteed that: any geodesic line will not self-intersect and will be the unique shortest geodesic between its endpoints; any geodesic circle will be homeomorphic to the standard circle; two geodesic lines either do not intersect or intersect at a single point; a line and a circle, or two circles, either do not intersect, or are tangent at one point, or intersect at exactly two points. Fig. 3 show the five basic constructions in the manifold setting.

There exist many straightedge and compass constructions in the Euclidean setting. We review just a few of them, analyzing the problems that hinder their direct extension to the manifold setting.

4.1. Copying, adding and subtracting segments

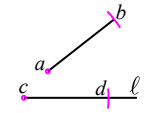

Given a line segment ab in the plane and a line ℓ through another point c, find a point d on ℓ such that ab and cd have the same length. In the plane, the aperture of the compass is taken at ab, then the needle point is placed at c and a circle is traced; point d is taken as any of the two intersections of the circle with line ℓ. Notice that we are assuming a non collapsible compass here; it can be shown that the same result can be achieved with a collapsible compass, through a more involved procedure though.

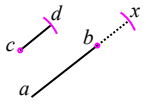

Given two line segments ab and cd in the plane, extend ab at b for a length equal to cd. In the plane, segment ab is extended to a line with the straightedge; the aperture of the compass is taken at cd and a circle is traced by placing the needle point at b; the intersection x of this circle with the line is taken, which lies on the opposite side of a wrt b; line segment ax is the result.

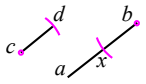

Given two line segments ab and cd in the plane, assuming that ab is longer than cd, shorten ab at b by the length of cd. In the plane, the aperture of the compass is taken at cd and a circle is traced by placing the needle point at b; the intersection x of this circle with ab is taken; line segment ax is the result.

All three constructions can be successfully ported to the manifold setting in a straightforward way, by means of the basic constructions, as described in the previous section. In this case, everything works fine because we are addressing just distances and collinearity, whose properties are preserved in the manifold setting.

4.2. Operations with angles

In the plane, an angle is defined by two half-lines ℓ_a and ℓ_b incident at a point x, which can be built with the straightedge, given two points a and b lying on them, respectively.

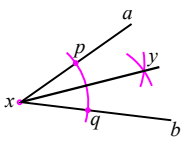

In the plane, an angle can be bisected as follows. Place the needle point of the compass at x, trace any circle and let p and q be its intersections with ℓ_a and ℓ_b. Place the needle point at p, and next at q, with aperture pq trace other two circles; let y be any of their two intersection points. The line ℓ_y through x and y bisects the angle at x. An additional property of the bisector is that all its points are equidistant from ℓ_a and ℓ_b.

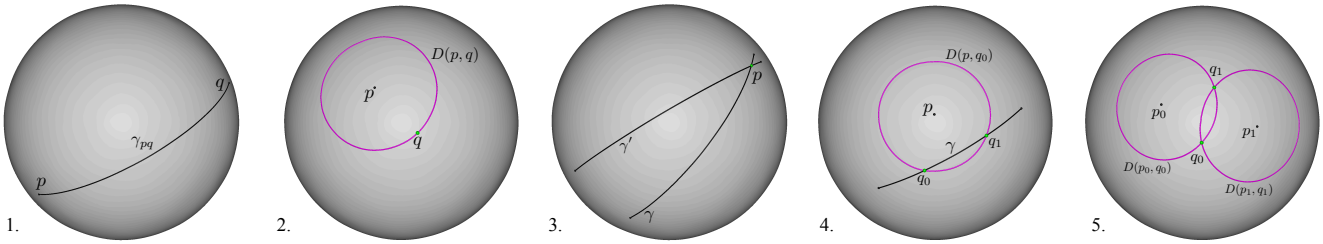

Figure 3: *The five basic constructions on a sphere. The black curves are geodesic lines, while the curves in magenta are geodesic circles. We denote with $D(p,q)$ the geodesic circle centered at point p and passing through point q.*

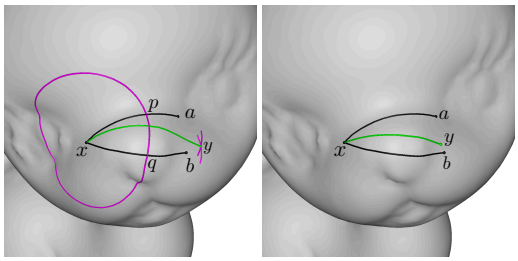

Figure 4: *The bisection of an angle axb by reproducing the Euclidean construction (left); and by tracing a geodesic with initial tangent vector mathbft_x obtained with the Euclidean construction in the tangent plane T_xS (right). In the first case, the green line γ_{xy} does not have any of the properties of a Euclidean bisector. In the second case, the angle at x is truly bisected, but the points of γ_{xy} are not equidistant from lines γ_{xa} and γ_{xb}, and γ_{xy} may eventually intersect one of them.*

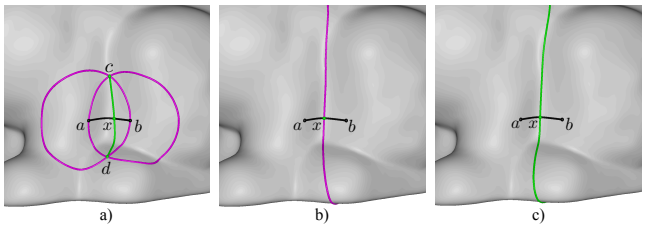

Figure 5: *The bisector of a geodesic segment computed by reproducing the Euclidean construction (a); by the zero isoline of the difference of distance fields from a and b (b); and by tracing a geodesic from the midpoint of the segment (obtained as in b) along the orthogonal direction computed in the tangent plane T_xS (c). Only the third construction gives a geodesic line orthogonal to ab and through its midpoint.*

This construction fails in the manifold setting: neither the geodesic line through x and y bisects the angle at x, nor its points are equidistant from the input lines. We can resort to operators *Tangent* and *Geodesic-tracing* to find a geodesic line that bisects the angle. Given lines γ_{xa} and γ_{xb} intersecting at x, find their tangent vectors t_a and t_b at x. Extend such vectors to lines in the tangent plane, and use the Euclidean construction to find line ℓ_y as above; let t_y be its direction at x. Trace geodesic line γ_{xy} through x and along direction t_y. The resulting line bisects the angle in the sense that the angles formed by its tangent at x and the tangents of the two input lines at x are equal, by construction. However, similarly to the previous case, its points will not have the equidistance property from the input lines. In fact, the locus of equidistant points from the two lines is not a geodesic line in general, and finding it is beyond the scope of this paper, as it requires using the distance fields from γ_{xa} and γ_{xb}, while we limit our distance fields to have their sources at single points. Fig. 4 shows results obtained with the two methods.

A number of other constructions deal with operations on angles, such as copying an angle to another place, adding or subtracting angles, or creating angles of a few specified amplitudes, namely $\pi/2$, $\pi/3$, $\pi/4$, $\pi/6$. These problems are somehow *local* to the point x at the tip of the angle, and can be addressed in the manifold setting in a straightforward way, by finding the tangents of the geodesic lines that define the angles at play, resolving the Euclidean construction in the tangent plane, and using the resulting directions to trace geodesics back to the surface S. For this reason, we do not analyze such constructions in detail.

4.3. Line segment bisector and midpoint

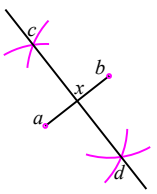

In the plane, the bisector is constructed as follows. Given points $a,b \in \mathbb{R}^2$, first use the straightedge to trace the straight-line segment joining them. Then place the needle point of the compass at a and the pencil point at b and trace a circle; repeat the same operation with needle at b and pencil at a. Let c,d be the intersection points of the two circles; use the straightedge to trace segment cd. In the Euclidean setting, the *straight line* line through c,d intersects segment ab *orthogonally* and *at its midpoint* x; this is also the *locus of points that have equal distance* from a and b.

The construction above fails miserably in the manifold setting. If we use *Geodesic-line* and *Geodesic-compass* to obtain points c,d and connect them as above, and we use *Intersect* to compute the intersection x of the two geodesic paths γ_{ab} and γ_{cd}, then in general x will *not* be the midpoint of γ_{ab} and the two paths will *not* be orthogonal at x. Concerning distances, we only know that c and d are equidistant from a and b, but distances can be different at all other points of γ_{cd}.

In order to overcome this limitation, we resort to our additional

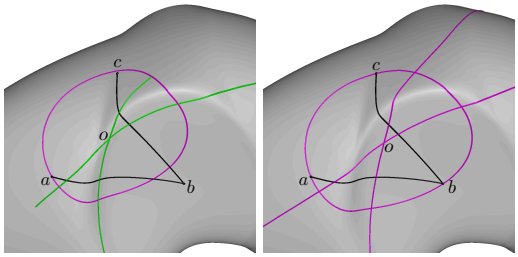

Figure 6: *Circle through three points a,b,c. Left: A straightforward reproduction of the Euclidean construction fails because the intersection of the two green lines is not equidistant from a,b,c; a similar failure would occur when intersecting lines from the construction of Fig.5(c). Right: The intersection of curves obtained with the construction of 5(b) gives the correct center of the geodesic circle.*

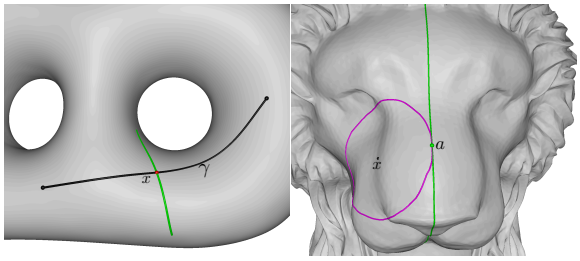

Figure 7: *Left: Perpendicular to line γ through point x computed by placing the square set at x: the tangent direction $\mathbf{t}_x$ in T_xS is obtained with the Euclidean construction starting at the tangent of γ in T_xS. Right: Line tangent to a circle centered at x and through a: the square set is placed at a and oriented according to the tangent of γ_{xa} at a.*

tools. Let d_a, d_b be the two distance fields with sources at a and b, respectively. Compute the difference field $d_{ab} = d_a - d_b$; the two points c and d belong to the zero isoline of this field. If we extract the *Isoline* through c of d_{ab}, the resulting line $\gamma_{ab}^{\perp}$ will intersect γ_{ab} at its midpoint x, which can be found with *Intersect*. It can be easily seen that $\gamma_{ab}^{\perp}$ and γ_{ab} intersect orthogonally at x. Besides, all points of $\gamma_{ab}^{\perp}$ are equidistant from a and b, by construction. However, $\gamma_{ab}^{\perp}$ is *not* a geodesic line, hence not *straight* in the manifold sense.

If we want to obtain a geodesic line that intersects orthogonally γ_{ab} at its midpoint, we have to give up equidistance from a and b at all points, and we need a further construction. We first compute the *Tangent* t_x of γ_{ab} at x; this tangent defines a line ℓ in the tangent plane T_xS. We use the Euclidean construction to find the line $\ell^{\perp}$ through x and orthogonal to ℓ in T_xS, as follows: first place the needle of the compass at x and trace a circle; let p and q be the intersections of such circle with ℓ; next find $\ell^{\perp}$ as the perpendicular bisector of segment pq. Let $t^{\perp}$ be the direction of $\ell^{\perp}$ in T_xS, use *Geodesic-tracing* to draw a geodesic line from x in direction $t^{\perp}$ on S. The resulting line is the desired result.

Fig. 5 shows results obtained with the three different methods.

The Square-set as a derived operator. Given any curve γ and a point x on it, the following construction can be used to obtain a geodesic line intersecting γ at x orthogonally: find the tangent of γ at x; find its orthogonal direction in the tangent plane as above; and finally trace the geodesic through x along such direction. This procedure implements a further operation that we call *Square-set*, which will be used as an atomic operation in the following.

4.4. Circle through three non-collinear points

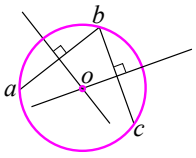

In the plane, given three non-collinear points a,b,c, this construction can be done by first computing the perpendicular bisectors of segment ab and bc; then intersecting such two bisectors at point o; and finally tracing the circle centered at o and through a (and, consequently, through b and c). The same procedure trivially gives the circle circumscribed to a triangle abc.

This construction relies on the fact that all points on a bisector are equidistant from the endpoints of the input segment, a property which is not fulfilled in the manifold case when the bisector is a straight line. However, if the two bisectors are obtained as isolines of the difference distance field, as described above, then their intersection will indeed be equidistant from the three points, hence we can use it as the center for a geodesic circle through them. See Fig. 6 for a comparison of the two approaches.

4.5. Perpendicular to a line at a point

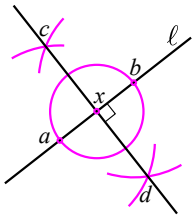

In the plane, let ℓ be a line and x a point on it, we want to find a line through x and orthogonal to ℓ. To this aim, it is sufficient to trace any circle centered at x, finding its intersections a,b with ℓ, and then finding the bisector of line segment ab.

While this construction can be used in the manifold setting too, by resorting to the construction described in Sec. 4.3, it is easier here to use the *Square-set*, placing it at x and orienting it according to the tangent of the support line at x.

The same construction can be used to find the tangent at a point a to a circle centered at x and through a. This is in fact the perpendicular to segment ax and passing through a. Fig. 7 shows both such constructions.

4.6. Perpendicular to a line through a point not on the line

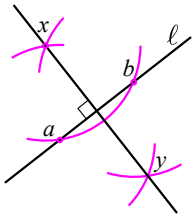

In this case, point x lies outside ℓ and we are again asked to find the perpendicular to ℓ through x. In the plane, we also trace a circle centered at x, with an aperture larger than its distance from ℓ; we find the intersection points a,b of this circle with ℓ and we trace other two circles centered at a and b with the same aperture; by construction, such circles intersect at x and at another point y; segment xy is orthogonal to ℓ.

Unfortunately, this construction cannot be used in the geodesic setting, because the orthogonal projection of a point x onto a geodesic γ will not be the midpoint of a segment intercepted on γ with a circle centered at x. We rather have to define the problem in terms of distances, as finding the point z on γ that minimizes the distance from x. If we could find z, there is guarantee that the geodesic path γ_{xz} meets γ orthogonally at z, because it is a radial path of the circle centered at x and tangent to γ. But finding z seems not possible without taking any measure. One possible workaround, which however violates the rules of the straightedge and compass framework, is to restrict the distance field d_x to γ and finding its minimum along it.

A similar problem, which poses the same challenge, is mirroring a point x about a line not containing it. Once we have found the perpendicular from x to the line, it is sufficient to trace a circle centered at the projection z of z with aperture xz and then find the intersection between such circle and the geodesic line extending xz.

4.7. Triangles

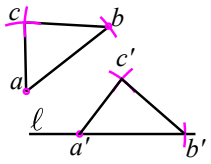

A triangle can be copied to another place with the same construction, both in the planar and in the manifold setting. Let abc be a triangle, ℓ a line and a' a point on ℓ. We want to copy the triangle in such a way that a goes to a', b goes to a point b' on ℓ, and c is placed at a point c' accordingly. We first draw a circle with amplitude ab centered at a' and we select a point b' as one of the two intersections of the circle with line ℓ. Next we trace two more circles, one with amplitude ac centered at a' and another with amplitude bc centered at b'; we select point c' as one of the intersections of such two circles. In the manifold setting, the result is a triangle with edges of the same length of abc, but nothing can be said about its angles. Moving a triangle while preserving the amplitude of its angles is inherently impossible in general, for consequences of the Gauss-Bonnet theorem.

Creating an equilateral triangle is among the simplest constructions: given an edge ab, intersect the two circles with radius ab and centered at a and b, respectively. Any of their two intersections can be chosen as the third vertex c of the triangle. The same procedure works in the manifold setting too, if we aim at obtaining a triangle with three edges of the same length. This does not guarantee any other of the properties of the equilateral triangles, e.g., having three equal angles, having three equal heights that bisect the angles and bisect the edges, etc. Constructions fulfilling even one of such requirements seem not easy to obtain in the manifold setting.

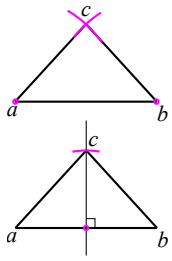

Likewise, it is easy to build an isosceles triangle on a basis ab with the diagonal edges of a given length (transferred with the compass from some given segment). Alternatively, one can build an isosceles triangle of a given height, by first constructing the perpendicular bisector of ab and then transferring the height on it with the compass. Both such constructions work to some extent in the manifold setting, too. However, the first construction will not warrant anything about either equality of the angles at the basis, or the height from c to bisect ab. While the second construction will just warrant the latter property, but neither that the diagonal edges, nor that the angles at the basis are equal. In our GUI, we implement a more practical, yet equivalent, variant of the first construction: we consider the *Isoline* of points equidistant from a and b, as in Section 4.3, and we let the user choose the length of the sides by dragging point c along such bisector.

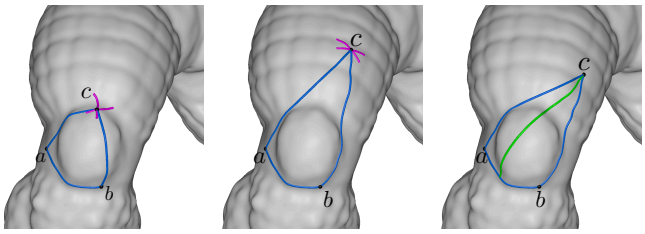

Figure 8: *Equilateral triangle (a) and isosceles triangle obtained by reproducing the Euclidean construction have three/two edges of the same length (b). A triangle with vertex c lying on the perpendicular bisector of the basis ab does not have any property in terms of either edge or angle equalities (c).*

Figure 8 shows examples of triangles on a surface, obtained with the constructions described above.

The examples above demonstrate that when we deal with some regular figure in the manifold setting, we cannot ask it to fulfill all properties such regular figure has in the planar setting together. We usually can ask it to have at most one property at a time, and the construction and the result will be different depending on which property we aim at fulfilling. More than that, it is not always clear if and how some such properties can be obtained with a geometric construction.

4.8. Squares and rectangles

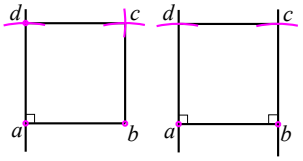

A square can be built from one of its edges ab as follows. A line perpendicular to ab and through a is built first, as described in Sec. 4.5. Then the the length of ab is transferred to segment ad on such a line by placing the needle point of the compass at a. Finally, the needle point of the compass is placed at b and next at d with the same aperture ab, and the intersection c of the two circles gives the last vertex of square $abcd$.

This same construction works in the manifold setting too. However, the resulting polygon will have four edges of equal length, but only angle dab is guaranteed to be a square angle. An alternative construction consists of tracing perpendicular lines at both a and b, transferring the length of ab on both of them, and connecting the points c and d obtained in this way. In this case, in the manifold setting we obtain a quadrilateral with three edges of the same length, namely ab, ad and bc, and two right angles dab and abc; but noth-

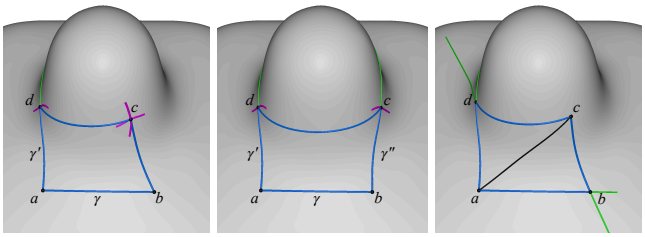

Figure 9: *Rectangles obtained with different constructions: by tracing two perpendicular lines γ and γ' intersecting at a and tracing opposite sides of the same length (a); by tracing two lines γ' and γ'' perpendicular to γ at a and b and setting points d and c on γ' and γ'' at equal distance from a and b, respectively (b); by tracing the diagonal ac, transferring angle bac to acd and tracing two lines perpendicular to ab and cd at a and c, respectively (c).*

ing can be said about the length of edge *cd* and the amplitude of angles at *c* and *d*.

The same constructions apply to draw a rectangle, except that the aperture of the compass to obtain the vertical edges can be different than the length of *ab*. The outcome in the manifold setting has the analogous (lack of) properties.

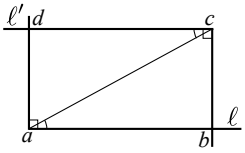

We describe a third construction, which is more appropriate to the GUI of drawing systems. Given a basis line ℓ and a point a lying on it, a diagonal segment ac is traced first. Then the angle between such segment and line ℓ is transferred at c, as described in Sec. 4.2, to obtain a line ℓ' parallel to ℓ. Finally, two lines are traced through a and c, which are perpendicular to ℓ and ℓ', respectively. The intersections of such lines with the first two lines give the other two vertices b and d of the rectangle. This construction applies to the manifold setting, too, by copying the angle in the tangent planes and using the square set to trace perpendicular lines. However, the resulting quadrilateral has two square angles at a and c, but nothing can be said on the amplitude of the other two angles, and opposite edges are not congruent in general. A number of other constructions can be devised, which are all equivalent in the Euclidean setting, while none of them can warrant congruent opposite edges and four right angles. Each such construction privileges some of the properties of rectangles, at the expense of others.

Fig. 9 shows examples of rectangles on a surface, obtained with the constructions described above.

4.9. Regular polygons

The limitations about regularity are even more evident for polygons with more edges. In the Euclidean setting, straightedge and compass constructions of a k-gon are known for several values of k: given a circle, the length of the edge of a k-gon is obtained first, which is then transferred repeatedly to the circle, and the dots are eventually connected. All such constructions are based on the fact that angles at the circle's center of equal width define arcs of equal length and chords of equal length.

All these constructions fail in the manifold setting, because the relations between the amplitude of angles and the lengths of arcs and chords no longer hold. More specifically, for no value of k it is clear how to find a radius of the geodesic compass, such that if that distance is transferred to a geodesic circle k times then the polygon closes correctly. The following approximations is possible with the tools in our arsenal, though. Define θ_k to be $\frac{2\pi}{k}$: this angle can be found with straightedge and compass construction in the plane for several values of, k, then transferred to the surface as in Sec. 4.2; next transport such angle $k-$ times to the tangent plane of the center c of a given circle; trace the geodesics along the directions defining the transported angles, until intersecting the circle; join the consecutive intersections with geodesic lines. The resulting polygon has equal angles at the center, but in general it neither has equal edges, nor equal internal angles; and his vertices do not define equal arcs on the circle.

More in general, when defining the counterpart of a regular polygon on a manifold, we need to clarify which properties such polygon should have. It is not clear whether a k-gon with equal sides and equal angles is even possible in the general case, due to the curvature of the internal region. While a polygon with equal sides and all its vertices of a geodesic circle is well defined, its construction seems not trivial and not possible with our arsenal. Even a simple construction that splits a circle into k arcs of equal length and then connects the dots requires measuring distances in the manifold case, thus breaking the constraints of the straightedge and compass framework.

4.10. Parallel lines

There are a number of constructions in the plane that deal with parallel lines. As a matter of fact, the concept itself of parallel lines is ill-defined in the manifold setting. Given a point on $x \in S$ and a direction t_x on its tangent plane, e.g., the tangent of a geodesic line γ_x through x, this direction can be transported to a "parallel" direction t_y lying in the tangent plane of any other point $y \in S$. This is done via *parallel transport*, a fairly complex differential geometry operation that we have not considered in our preliminaries. Once t_y is given, we could trace the geodesic through y tangent to t_y and considering it a "parallel" to γ_x. The trouble here is that the direction t_y will be different depending both on the starting point x that we select on γ_x, and on the trajectory that we choose to transport t_x. Therefore, the result is not unique and it is somehow arbitrary.

A straightforward possibility it to take a reference line γ and define a bundle of "parallel" lines as all those lines that intersect γ with a given angle. Given points $x_0, \ldots, x_n$ along γ and the reference angle, it is possible to use the construction in Sec.4.2 to trace such parallel lines through the points x_i.

Because of all such considerations, we leave a more thorough study of parallel lines to the manifold setting as future work.

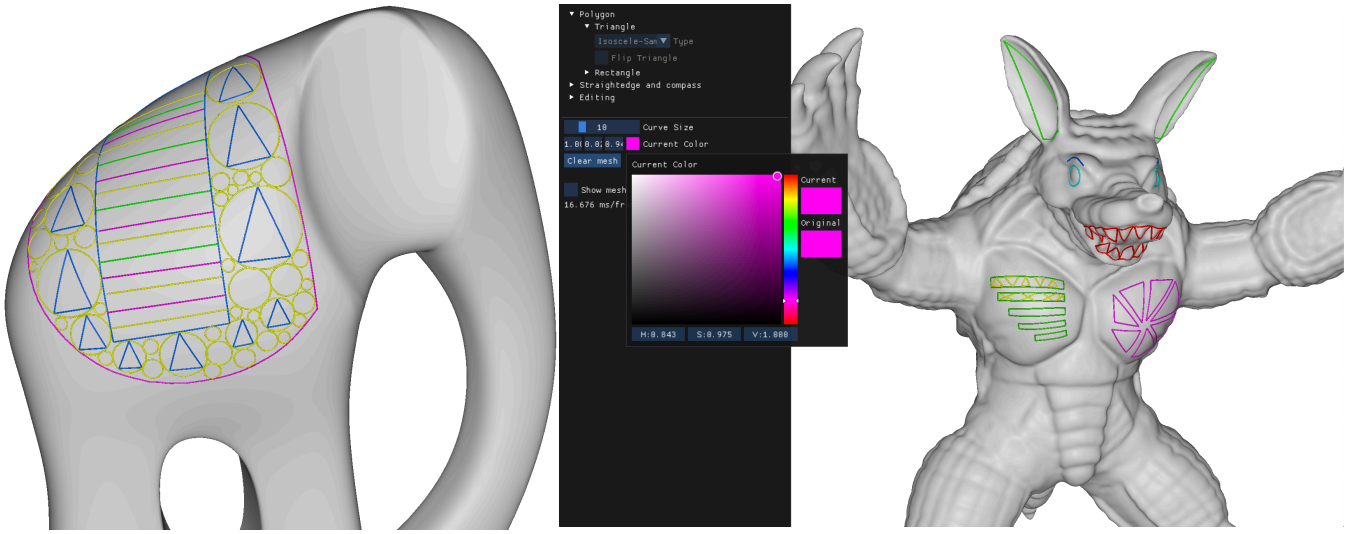

Figure 10: *The graphical user interface of our prototype system, with some geometric constructions obtained interactively on a meshes consisting of 1M triangles (left) and about 350K triangles (right).*

4.11. More constructions

There are several other constructions that are easy in the plane, while it is not clear how they can be done in the manifold case:

- Tangents to a circle through an external point: the construction in the plane is based on the fact that an angle at the circumference in a half circle measures $\pi/2$. As all constructions that relate angles and lengths, this cannot be used in the manifold setting.
- Circle inscribed in a triangle: the construction in the plane is based on the fact that all points in the bisectors of angles are equidistant from the edges. This is no longer true in the manifold case. On the other hand, it is not clear how the locus of points that are equidistant from two edges – which is not a geodesic line in general – can be constructed, unless the distance fields from the edges are available (see also the discussion in Sec. 4.2).

5. Implementation

All the constructions described in the previous section, which apply to the manifold setting, have been implemented by means of the primitives defined in Sec. 3 and included in a prototype system, which supports their interactive usage on high resolution meshes. Note that, on a mesh M, any line is approximated with a polyline containing one segment per crossed triangle. The primitives have been implemented on top of a light data structure encoding meshes [PNC19], as follows:

- *Shortest-path:* We use a variation of the algorithm proposed in [XW07]. The main feature of our method lies in its efficiency in computing the initial guess: we extract a strip of triangles connecting the endpoints by navigating a dual graph on the mesh. The navigation of such graph is optimized using well known heuristics such as SLF and LLL [Ber98]. All the details about the implementation and comparisons with state-of-the-art methods can be found in [MNPP21].
- *Tangent:* This is trivial if a point p lies inside a triangle, since the plane containing the triangle coincides with the tangent plane. If p lies on an edge, it is sufficient to unfold the incident triangles, and compute a discrete approximation of the tangent at p of a polyline consisting of two consecutive segments. If p is a vertex, then the total angle about p on M is rescaled to 2π to map the 1-ring of p to the tangent plane; then the segments crossing the triangles incident at p are mapped accordingly, and the same discrete approximation is applied; the resulting direction is pulled back to the mesh by an inverse rescaling of angles.
- *Geodesic-tracing:* We apply a straightforward implementation of the *straightest geodesics* described in [PS98]. In practice, we propagate a geodesic line across an edge e by flattening the two triangles incident at e; and across a vertex v by rescaling its 1-ring to the tangent plane, as above, and extending the incoming direction on a straight trajectory with respect to such plane.
- *Distance-field:* We rely on a very efficient graph-based solver described in [NPP21], which computes the distance field at all vertices of M, or on the vertices of a desired region of interest.
- *Isoline:* We linearly interpolate a field inside each triangle of M. For each triangle t, which crosses a given isovalue, the segment of isoline crossing t is computed independently.
- *Intersect:* Lines on M are encoded as paths, each consisting of a strip of triangles containing the line, together with parametric coordinates on each edge crossing the strip, which are used to encode the intersection points. Intersections between a pair of lines can be found in linear time in the total number of triangles in the corresponding paths. Each triangle intersecting one of the paths is assigned a unique tag; next the triangles forming the other path are scanned, and intersections are computed at tagged triangles.

Our implementation can support interactive times on meshes up to millions of triangles. All primitive operations and constructions are supported via intuitive click-and-drag, which mimic the behav-

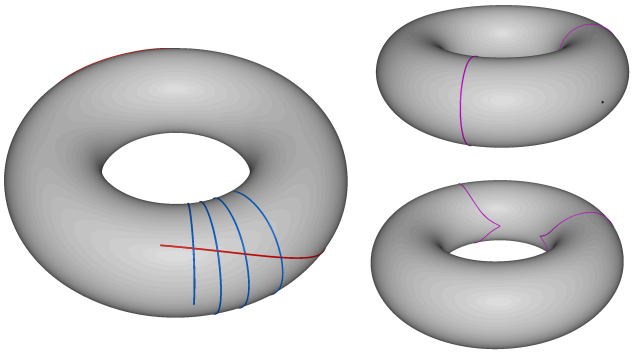

Figure 11: *Left: two straightest geodesics intersecting in more than one point. Right: a geodesic circle with a radius greater than the injectivity radius of its center is not diffeomorphic to its Euclidean counterpart. The cusps (bottom) arise at the intersections of the circle with the cut locus of its center.*

ior of standard 2D drawing systems. Fig. 10 shows examples of mesh drawing through the GUI of our system. A live session is shown in the accompanying video.

6. Discussion

It follows from our analysis that the *Geodesic-line* and *Geodesic-compass* alone are sufficient to port only few of the existing straightedge and compass constructions to the manifold setting. We extend the scope of constructions by allowing for a more extended usage of the *Distance-field* operator, which can support constructions that need the computation of midpoints and orthogonal lines. These constructions already support several operations in the context of interactive vector graphics on surfaces. A few relevant constructions, however, such as the projection of a point to a line and the computation of a tangent to a circle through an external point, are still not supported. Such operations may require explicit measures, e.g., finding the minimum of a distance field along a given line, which are forbidden in the straightedge and compass framework. This is not difficult to implement and integrate into our prototype system, though.

A relevant limitation stems from the impossibility to warrant the congruence of both lengths and angles together in tracing regular polygons. This limitation is intrinsic to the geodesic setting and does not depend on the type of construction applied. This fact makes regular tilings hardly applicable to manifold surfaces, without introducing some modifications. Tilings can be addressed by relaxing some conditions on the preservation of angles and/or lengths, but it remains a challenging problem how to extend them to large regions. This problem is tightly related to the design of N-RoSy fields [VCD*17] and, in particular to the presence of field singularities, which is a consequence of the Gauss-Bonnet theorem.

A possible avenue is to relax the constraint of lines to be straight, in geodesic terms, trading some straightness for other properties. This leads to the a concept of *as-straight-as-possible* lines under given constraints, e.g., joining their endpoints with a prescribed length or with given tangent directions. Computing such lines entails investigating Jacobi fields [PHD*10, Le 19] and related optimization problems.

A further challenge is extending our primitives to work on large non-convex regions of the surface. In fact, even basic properties of lines and circles are lost outside strongly convex regions, thus bringing further issues. For example, geodesic lines could intersect in more than one point, as shown in Figure 11, left. Similarly, since circles has been defined as isolines of geodesic distance fields, beyond the cut loci of their centers they will no longer be diffeomorphic to their Euclidean counterparts (Figure 11,right). Some operations, including all primitives that can be computed with distance fields, as well as an extension of Bézier surfaces, have been addressed already in [MNPP21, NPP21]. We plan to address the remaining issues in our future work.

References

[AKP19] ALEXANDER S., KAPOVITCH V., PETRUNIN A.: Alexandrov geometry: preliminary version no. 1, 2019. arXiv:1903.08539. 2

[Ale48] ALEXANDROV A. D.: Intrinsic geometry of convex surfaces. *OGIZ, Moscow-Leningrad* (1948). 2

[Ber98] BERTSEKAS D.: *Network optimization: continuous and discrete models*. Athena Scientific, Belmont, Massachusetts, 1998. 8

[CE75] CHEEGER J., EBIN D.: *Comparison theorems in riemannian geometry*. North-Holland mathematical library. North-Holland Pub. Co, 1975. Includes bibliographical references (pages 169-172) and index. 2

[Cha06] CHAVEL I.: *Riemannian Geometry: A Modern Introduction*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2006. URL: https://books.google.it/books?id=3Gjp4vQ_mPkC. 2

[CLPQ20] CRANE K., LIVESU M., PUPPO E., QIN Y.: A survey of algorithms for geodesic paths and distances, 2020. arXiv:2007.10430. 2

[Le 19] LE BRIGANT A.: A Discrete Framework to Find the Optimal Matching Between Manifold-Valued Curves. *Journal of Mathematical Imaging and Vision 61*, 1 (2019), 40–70. arXiv:1703.05107, doi:10.1007/s10851-018-0820-2. 9

[MLP21] MANCINELLI C., LIVESU M., PUPPO E.: Practical computation of the cut locus on discrete surfaces. *Computer Graphics Forum 40*, 5 (2021), 261–273. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.14372, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.14372, doi:https://doi.org/10.1111/cgf.14372. 2

[MNPP21] MANCINELLI C., NAZZARO G., PELLACINI F., PUPPO E.: b/surf: Interactive bézier splines on surfaces, 2021. arXiv:2102.05921v2. 1, 8, 9

[NPP21] NAZZARO G., PUPPO E., PELLACINI F.: geoTangle: interactive design of geodesic tangle patterns on surfaces. *ACM Trans. Graph.* (2021). 1, 8, 9

[PHD*10] POTTMANN H., HUANG Q., DENG B., SCHIFTNER A., KILIAN M., GUIBAS L., WALLNER J.: Geodesic patterns. *ACM SIGGRAPH 2010 Papers, SIGGRAPH 2010* (2010). doi:10.1145/1778765.1778780. 9

[PNC19] PELLACINI F., NAZZARO G., CARRA E.: Yocto/GL: A Data-Oriented Library For Physically-Based Graphics. In *EG Ital. Chap. Conf.* (2019). https://github.com/xelatihy/yocto-gl. 8

[PS98] POLTHIER K., SCHMIES M.: Straightest geodesics on polyhedral surfaces. In *Mathematical Visualization*. Springer-Verlag, New York, 1998, pp. 135–150. 8

[Rod16] RODRIGUES O.: Recherches sur la théorie analytique des lignes et des rayons de courbure des surfaces, et sur la transformation d'une classe d'integrales doubles, qui ont un rapport direct avec les formules de cette théorie. In *École Polytechnique Corresp. III*. 1816, pp. 162,182. 2

[VCD*17] VAXMAN A., CAMPEN M., DIAMANTI O., BOMMES D., HILDEBRANDT K., TECHNION M. B.-C., PANOZZO D.: Directional field synthesis, design, and processing. In *ACM SIGGRAPH 2017 Courses* (New York, NY, USA, 2017), SIGGRAPH '17, Association for Computing Machinery. 9

[XW07] XIN S.-Q., WANG G.-J.: Efficiently determining a locally exact shortest path on polyhedral surfaces. *Computer Aided Design 39*, 12 (2007), 1081–1090. 8