Intro: smooth curves and surfaces in Euclidean space (usually $\mathbb{R}^3$). Usually not as level sets (like $x^2 + y^2 = 1$) as in algebraic geometry, but parametrized (like $(\cos t, \sin t)$).

Euclidean space $\mathbb{R}^n \ni x = (x_1, \ldots, x_n)$ with scalar product (inner prod.) $\langle a, b\rangle = a \cdot b = \sum a_i b_i$ and norm $|a| = \sqrt{\langle a, a\rangle}$).

A. CURVES

$I \subset \mathbb{R}$ interval, continuous $\alpha\colon I \to \mathbb{R}^n$ is a *parametrized curve* in $\mathbb{R}^n$. We write $\alpha(t) = (\alpha_1(t), \ldots, \alpha_n(t))$.

The curve is C^k if it has continuous derivs of order up to k. (C^0=contin, C^1, …, C^∞=smooth). Recall: if I not open, differentiability at an end point means one-sided derivatives exist, or equiv. that there is a diff'ble extension. We deal with smooth curves except when noted.

Examples:

- $\alpha(t) = (a\cos t, a \sin t, bt)$ helix in $\mathbb{R}^3$
- $\alpha(t) = (t^2, t^3)$ curve in $\mathbb{R}^2$ with cusp
- $\alpha(t) = (\sin t, \sin 2t)$ figure-8 curve in $\mathbb{R}^2$
- $\alpha(t) = (t, t^2, \ldots, t^n)$ moment curve in $\mathbb{R}^n$

Define *simple* and *closed* curves (and simple closed curves).

A smooth (or even just C^1) curve α has a *velocity vector* $\dot\alpha(t) \in \mathbb{R}^n$ at each point. The fundamental thm of calculus says $\int_a^b \dot\alpha(t)\,dt = \alpha(b) - \alpha(a)$. The *speed* of α is $|\dot\alpha(t)|$. We say α is *regular* if the speed is positive (never vanishes). Then the speed is a (smooth) positive function of t. (The cusped curve above is not regular at $t = 0$; the others are regular.)

The *length* of α is $\operatorname{len}(\alpha) = \int_I |\dot\alpha(t)|\,dt$. We see

$$\int_a^b |\dot\alpha(t)|\,dt \geq \left|\int_a^b \dot\alpha(t)\,dt\right| = \big|\alpha(b) - \alpha(a)\big|.$$

That is, a straight line is the shortest path. (To avoid using vector version of integral inequality, take scalar product with $\alpha(b) - \alpha(a)$.)

The length of an arbitrary curve can be defined (Jordan) as total variation:

$$\operatorname{len}(\alpha) = \mathrm{TV}(\alpha) = \sup_{t_0 < \cdots < t_n \in I} \sum_{i=1}^n \big|\alpha(t_i) - \alpha(t_{i-1})\big|.$$

This is the supremal length of inscribed polygons. (One can show this length is finite over finite intervals if and only if α has a Lipschitz reparametrization (e.g., by arclength). Lipschitz curves have velocity defined a.e., and our integral formulas for length work fine.)

If J is another interval and $\varphi\colon J \to I$ is an orientation-preserving homeomorphism, i.e., a strictly increasing surjection, then $\alpha\circ\varphi\colon J \to \mathbb{R}^n$ is a parametrized curve with the same image (trace) as α, called a *reparametrization* of α. (Note: reverse curve $\overline{\alpha}\colon -J \to \mathbb{R}^n$, $\overline{\alpha}(t) := \alpha(-t)$ has the same trace in reverse order – orientation reversing reparam.)

For studying continuous curves, it's sometimes helpful to allow reparam's that stop for a while (monotonic but not strictly) – or that remove such a constant interval.

We instead focus on regular smooth curves α. Then if $\varphi\colon J \to I$ is a diffeomorphism (smooth with nonvanishing derivative, so φ^{-1} is also smooth) then $\alpha \circ \varphi$ is again smooth and regular. We are interested in properties invariant under such smooth reparametrization. This is an equivalence relation. An *unparametrized (smooth) curve* can be defined as an equivalence class. We study these, but implicitly.

For a fixed $t_0 \in I$ we define the arclength function $s(t) := \int_{t_0}^t |\dot\alpha(t)|\,dt$. Here s maps I to an interval J of length $\operatorname{len}(\alpha)$. If α is a regular smooth curve, then $s(t)$ is smooth, with positive derivative $\dot s = |\dot\alpha| > 0$ equal to the speed. Thus it has a smooth inverse function $\varphi\colon J \to I$. We say $\beta = \alpha\circ\varphi$ is the *arclength parametrization* (or *unit-speed parametrization*) of α. We have $\beta(s) = \alpha(\varphi(s))$, so $\beta(s(t)) = \alpha(\varphi(s(t))) = \alpha(t)$. It follows that β has constant speed 1, and thus that the arclength of $\beta|_{[a,b]}$ is $b - a$.

The arclength parametetrization is hard to write down explicitly for examples (integrate a square root, then invert the function). But it always exists, and is often the easiest to use for theory. (When considering curves with less smoothness, e.g., C^k, there is a general principle that no regular parametrization is smoother than the arclength parametrization.)

Although for an arbitrary parameter we have used the name t (thinking of time) and written d/dt with a dot, when we use the arclength parametrization, we'll call the parameter s and write d/ds with a prime. Of course, for any function f along the curve, the chain rule says

$$\frac{df}{ds}\frac{ds}{dt} = \frac{df}{dt}, \qquad \text{i.e., } f' = \dot f/\dot s = \dot f/|\dot\alpha|.$$

Suppose now that α is a regular smooth unit-speed curve. Then its velocity α' is everywhere a unit vector, the (unit) *tangent vector* $T(s) := \alpha'(s)$ to the curve. (In terms of an arbitrary regular parametrization, we have of course $T = \dot\alpha/|\dot\alpha|$.)

`End of Lecture 8 Apr 2013`

We should best think of $T(s)$ as a vector based at $p = \alpha(s)$, perhaps as an arrow from p to $p + T(s)$, rather than as a point in $\mathbb{R}^n$. The *tangent line* to α at the point $p = \alpha(s)$ is the line $\{p + tT(s) : t \in \mathbb{R}\}$. (Note about non-simple curves.) (As long as α has a first derivative at s, this line is the limit (as $h \to 0$) of secant lines through p and $\alpha(s + h)$. If α is C^1 near s, then it is the arbitrary limit of secant lines through $\alpha(s + h)$ and $\alpha(s + k)$.) While velocity depends on parametrization, the tangent line and unit tangent vector do not.

We are really most interested in properties that are also independent of rigid motion. It is not hard to show that a Euclidean motion of $\mathbb{R}^n$ is a rotation $A \in \mathrm{SO}(n)$ followed by a translation by some vector $v \in \mathbb{R}^n$: $x \mapsto Ax + v$. Thus α could be considered equivalent to $A\alpha + v\colon I \to \mathbb{R}^n$, $t \mapsto A\alpha(t) + v$. Of course, given any two lines in space, there is a rigid motion carrying one to the other. To find Euclidean invariants of curves, we need to take higher derivatives. We define the *curvature vector* $\vec\kappa := T' = \alpha''$; its length is the *curvature* $\kappa := |\vec\kappa|$.

Recall the Leibniz product rule for the scalar product: if v and w are vector-valued functions, then $(v\cdot w)' = v'\cdot w + v\cdot w'$.

In particular, if $v \perp w$ (i.e., $v \cdot w \equiv 0$) then $v' \cdot w = -w' \cdot v$. And if $|v|$ is constant then $v' \perp v$. (Geometrically, this is just saying that the tangent plane to a sphere is perpendicular to the radius vector.) In particular, we have $\vec{\kappa} \perp T$.

Example: the circle $\alpha(t) = (r\cos t, r\sin t)$ of radius r (parametrized here with constant speed r) has

$$T = (-\sin t, \cos t), \quad \vec{\kappa} = \frac{-1}{r}(\cos t, \sin t), \quad \kappa \equiv 1/r.$$

Given regular smooth parametrization α with speed $\sigma := \dot{s} = |\dot{\alpha}|$, the velocity is σT, so the acceleration vector is

$$\ddot{\alpha} = (\sigma T)^{\cdot} = \dot{\sigma}T + \sigma\dot{T} = \dot{\sigma}T + \sigma^2 T' = \dot{\sigma}T + \sigma^2\vec{\kappa}.$$

Note the second-order Taylor series for a unit-speed curve around the point $p = \alpha(0)$ (we assume without further comment that $0 \in I$):

$$\alpha(s) = p + sT(0) + \frac{s^2}{2}\vec{\kappa}(0) + O(s^3).$$

These first terms parametrize a parabola agreeing with α to second order (i.e., with the same tangent and curvature vector). Geometrically, it is nicer to use the *osculating circle*, the unique circle agreeing to second order (a line if $\vec{\kappa} = 0$). It has radius $1/\kappa$ and center $p + \vec{\kappa}/\kappa^2$. Thus we can also write

$$\alpha(s) = p + \cos(\kappa s)\vec{\kappa}/\kappa^2 + \sin(\kappa s)T/\kappa + O(s^3).$$

(Constant second derivative – parabola – points equivalent by shear. Constant curvature – circle – points equivalent by rotation.)

Any three distinct points in $\mathbb{R}^n$ lie on a unique circle (or line). The osculating circle to α at p is the limit of such circles through three points along α approaching p. (Variants with tangent circles, note on less smoothness, etc.)

Considering $T\colon I \to \mathbb{S}^{n-1} \subset \mathbb{R}^n$, we can think of this as another curve in $\mathbb{R}^n$ – called the *tantrix* (short for tangent indicatrix) of α – which happens to lie on the unit sphere. Assuming α was parametrized by arclength, the curve $s \mapsto T(s)$ has speed κ. Thus it is regular if and only if the curvature of α never vanishes. (Note on curves with nonvanishing curvature – in $\mathbb{R}^2$ versus $\mathbb{R}^3$.)

A1. Plane Curves

Now let's consider in particular plane curves ($n = 2$). We equip $\mathbb{R}^2$ with the standard orientation and let J denote the counterclockwise rotation by 90° so that $J(e_1) = e_2$ and for any vector v, $J(v)$ is the perpendicular vector of equal length such that $\{v, Jv\}$ is an oriented basis.

Given a (regular smooth) plane curve α, its (unit) *normal vector* N is defined as $N(s) := J(T(s))$. Since $\vec{\kappa} = T'$ is perpendicular to T, it is a scalar multiple of N. Thus we can define the (signed) *curvature* $\kappa_\pm$ of α by $\kappa_\pm N := \vec{\kappa}$ (so that $\kappa_\pm = \pm|\vec{\kappa}| = \pm\kappa$).

From $N \perp T$ and $T' = \kappa_\pm N$, we see immediately that $N' = -\kappa_\pm T$. We can combine these equations as

$$\begin{pmatrix} T \\ N \end{pmatrix}' = \begin{pmatrix} 0 & \kappa_\pm \\ -\kappa_\pm & 0 \end{pmatrix}\begin{pmatrix} T \\ N \end{pmatrix}.$$

Rotating orthonormal frame, infinitesimal rotation (speed $\kappa_\pm$) given by skew-symmetric matrix. The curvature tells us how fast the tangent vector T turns as we move along the curve at unit speed.

```
End of Lecture 11 Apr 2013
```

Since $T(s)$ is a unit vector in the plane, it can be expressed as $(\cos\theta, \sin\theta)$ for some $\theta = \theta(s)$. Although θ is not uniquely determined (but only up to a multiple of 2π) we claim that we can make a smooth choice of θ along the whole curve. Indeed, if there is such a θ, its derivative is $\theta' = \kappa_\pm$. Picking any θ_0 such that $T(0) = (\cos\theta_0, \sin\theta_0)$ define $\theta(s) := \theta_0 + \int_0^s \kappa_\pm(s)\,ds$.

This lets us prove what is often called the *fundamental theorem of plane curves* (although it really doesn't seem quite that important): Given a smooth function $\kappa_\pm\colon I \to \mathbb{R}$ there exists a smooth unit-speed curve $\alpha\colon I \to \mathbb{R}$ with signed curvature $\kappa_\pm$; this curve is unique up to rigid motion. First note that integrating $\kappa_\pm$ gives the angle function $\theta\colon I \to \mathbb{R}$ (uniquely up to a constant of integration), or equivalently gives the tangent vector $T = (\cos\theta, \sin\theta)$ (uniquely up to a rotation). Integrating T then gives α (uniquely up to a vector constant of integration, that is, up to a translation).

Now suppose α is a closed plane curve, that is, an L-periodic map $\mathbb{R} \to \mathbb{R}^2$. As above, we get an angle function $\theta\colon \mathbb{R} \to \mathbb{R}$ but this is not necessarily periodic. Instead, $\theta(L) = \theta(0) + 2\pi n$ for some integer n called the *turning number* (or rotation index or …) of α. (It follows that $\theta(s + kL) - \theta(s) = kn$ for any integer k and any s.) Note that the *total signed curvature* $\int_0^L \kappa_\pm\,ds$ of α is $2\pi n$. (If we reverse the orientation of α we negate the turning number.)

We will later prove Fenchel's theorem that $\int \kappa\,ds \geq 2\pi$ for any curve in $\mathbb{R}^n$ (with equality only for convex plane curves). A standard result is the Jordan curve theorem, saying that a simple closed plane curve divides the plane into two regions, one of which (called the *interior*) is bounded. Assuming the curve is oriented so that its interior is on the left, then its turning number is always $+1$. (Hopf gave a nice proof in 1935, which basically proceeds by defining a smooth angle function for all secants of the curve. One can also prove the analogous statement for polygons by induction, and then approximate a smooth curve by a polygon.)

A *vertex* of a plane curve is an extremal point of $\kappa_\pm$, that is a point where $\kappa_\pm$ achieves a local minimum or maximum, so that $\kappa'_\pm = 0$. Since any real-valued function on a compact set achieves a global minimum and maximum, any curve has at least two vertices.

Note that the osculating circle to α at p *crosses* α at p unless p is a vertex. Most people's sketches of osculating cicles are wrong! More generally, suppose α and β are two regular curves with the same tangent at p. Then α stays to the left of β in some neighborhood of p if $\kappa^\alpha_\pm > \kappa^\beta_\pm$ at p. And conversely of course if α stays to the left, then at least the weak inequality $\kappa^\alpha_\pm \geq \kappa^\beta_\pm$ holds at p.

A2. The Four-Vertex Theorem

The *Four-Vertex Theorem* says any simple closed plane curve α has at least four vertices. (Note counterexample $r = 1 + 2\sin\theta$ in polar coords if curve not embedded.) We give a proof due to Bob Osserman (1985).

Lemma: Given a compact set K in the plane (which might be the trace of a curve α) there is a unique smallest circle c enclosing K, called the *circumscibed circle*. (Existence via compactness of an appropriately bounded set of circles; uniqueness by constructing smaller circle containing intersection of two given ones.)

Properties:

1. c must touch K (for otherwise we could shrink c).
2. $c \cap K$ cannot lie in an open semicircle of c (for otherwise we could translate c to contradict (1).
3. thus $c \cap K$ contains at least two points, and if there are only two they are antipodal on c.

For the case of a curve α, by our previous remark, at any point of $c \cap \alpha$ the curvature of α is at least that of c.

Now let's prove the theorem. Let the curvature of the circumcircle c be k. If $c \cap \alpha$ includes an arc, there is nothing to prove. Otherwise suppose $c \cap \alpha$ includes at least $n \geq 2$ points p_i. (At these points $\kappa_\pm \geq i$.) We claim each arc α_i between consecutive p_i and p_{i+1} contains a point with $\kappa_\pm < k$. Then it also contains a vertex q_i (a local minimum of $\kappa_\pm$) with $\kappa_\pm < k$. Since the arc from q_{i-1} to q_i includes the point p_i with $\kappa_\pm \geq k$ is also includes a vertex p_i' (a local maximum of $\kappa_\pm$ with $\kappa_\pm \geq k$. Thus we have found $2n \geq 4$ vertices as desired.

To prove the claim, consider the one-parameter family of circles through p_i and p_{i+1} (with signed curvatures decreasing from k). The last one of these to touch α_i is tanget to α_i at at least one interior point, and since α_i stays to the right of that circle, its signed curvature is even less.

Where did we use the fact that the curve α is simple? (Recall that the theorem fails without this assumption!)

`End of Lecture 15 Apr 2013`

When two curves are tangent at p and don't cross locally, we got an inequality between their signed curvatures. But this assumes their orientations agree at p. By the Jordan curve theorem, a simple curve α bounds a compact region K. Clearly, α and K have the same circumcircle c. If both curves are oriented to have the compact regions to the left, then these orientations agree. Similarly, further application of the Jordan curve theorem ensure that the oriented circular arc from p_i to p_{i+1} used above agrees in orientation with α_i.

A3. Evolutes and the Nesting Theorem

Given a curve $\alpha\colon I \to \mathbb{R}^n$ with nonvanishing curvature, its *evolute* $\beta\colon I \to \mathbb{R}^n$ is the curve of centers of osculating circles: $\beta(t) := \alpha(t) + \vec{\kappa}(t)/\kappa(t)^2$. Let us consider in particular a unit-speed plane curve α with $\kappa = \kappa_\pm > 0$ and write $r = 1/\kappa$ for the radius of curvature. Then the evolute is $\beta(s) = \alpha(s) + r(s)N(s)$. Its velocity is $\beta' = T + r'N + rN' = r'N$, so its speed is $|r'(s)|$. (The evolute is singular where α has a vertex.) The acceleration of the evolute is $r''N + r'N' = r''N - r'T/r$, so its curvature is $\frac{1}{r|r'|}$.

Now consider a planar arc α with strictly monotonic, nonvanishing curvature. By the formula above, its evolute also has nonvanishing curvature, so in particular, the distance $|\beta(s_1) - \beta(s_2)|$ is stricly less than the arclength $\int_{s_1}^{s_2} |r'(s)|\,ds = \left|\int r'\,ds\right| = \left|r(s_1) - r(s_2)\right|$. This simply says the distance between the centers of two osculating circles to α is less than the difference of their radii, that is, the circles are strictly nested. This is the *nesting theorem* of Tait (1896) and Kneser (1914): the osculating circles along a planar arc with stricly monotonic, nonvanishing curvature are strictly nested.

A4. The Isoperimetric Inequality

Another global result about plane curves is the *isoperimetric inequality*. If an embedded curve of length L bounds an area A then $4\pi A \leq L^2$. (Equality holds only for a circle.)

If $R \subset \mathbb{R}^2$ is the region enclosed by the simple closed (C^1) curve $\alpha\colon [a,b] \to \mathbb{R}$, $\alpha(t) = (x(t), y(t))$, then we have by Green's theorem

$$A = \int_R dx\,dy = \int_\alpha x\,dy = \int_a^b x\dot{y}\,dt = -\int_a^b y\dot{x}\,dt.$$

(Actually, the formula gives an appropriately defined algebraic area even if the curve is not simple; no change if parametrization backtracks a bit.)

The trick suggested by Erhard Schmidt (1939) to prove the isoperimetric inequality is to consider an appropriate comparison circle. We deal with a smooth curve α. First find two parallel lines tangent to α such that α lies in the strip between them. Choose coords so make them the vertical lines $x = \pm r$. (Here $2r$ is the *width* of α in the given direction.) Parametrize α by arclength over $[0, L]$ by $(x(s), y(s))$ and parametriz the circle of radius r over $[0, L]$ by $\beta(s) = (x(s), \overline{y}(s))$: same $x(s)$ as for α, and thus $\overline{y}(s) = \pm\sqrt{r^2 - x(s)^2}$. (Note about non-convex curves, etc.) Note that the unit normal vector to α is $N = (-y', x')$, so $\langle N(s), \beta(s)\rangle = -xy' + \overline{y}x'$. We have $A = \int_0^L xy'\,ds$ and $\pi r^2 = -\int_0^L \overline{y}x'\,ds$. Thus

$$A + \pi r^2 = \int_0^L xy' - \overline{y}x'\,ds = \int_0^L \langle -N, \beta\rangle\,ds \leq \int_0^L |N|\,|\beta|\,ds = Lr.$$

Thus by the arithmetic-geometric mean inequality,

$$\sqrt{A\pi r^2} \leq (A + \pi r^2)/2 \leq Lr/2.$$

Squaring and dividing by r^2 gives the isoperimetric inequality.

It is not hard to check that if all these inequalities hold with equality, then α must be a circle.

`End of Lecture 18 Apr 2013`

A5. The Cauchy–Crofton Formula

Given a unit vector $u = u(\theta) = (\cos\theta, \sin\theta) \in \mathbb{S}^1 \subset \mathbb{R}^2$, the orthogonal projection to the line in direction u is $\pi_u\colon \mathbb{R}^2 \to \mathbb{R}^2$, $x \mapsto \langle x, u\rangle\, u$. If $\alpha\colon I \to \mathbb{R}^2$ is a smooth plane curve, then $\pi_u\alpha = \pi_u \circ \alpha$ is its projection (usually not regular!).

The Cauchy–Crofton formula says the length of α is $\pi/2$ time the average length of these projections. By average length we mean

$$⨍_{\mathbb{S}^1} \operatorname{len}(\pi_u\alpha)\, du = ⨍_0^{2\pi} \operatorname{len}(\pi_{u(\theta)}\alpha)\, d\theta := \frac{1}{2\pi}\int_0^{2\pi} \operatorname{len}(\pi_{u(\theta)}\alpha)\, d\theta$$

To prove this, first note that if α is a line segment, the average projected length is independent of its position and orientation and proportional to its length. That is, the theorem holds for line segments with some constant c in place of $\pi/2$. (We could easily compute $c = \pi/2$ by integrating a trig function, but wait!) Next, by summing, it holds for all polygons (with the same c). Finally, it holds for smooth curves (or indeed for all rectifiable curves) by taking a limit of inscribed polygons. (To know we can switch the averaging integral with the limit of ever finer polygons, we can appeal for instance to Lebesgue's monotone convergence theorem.) To compute $c = \pi/2$ it is easiest to consider the unit circle α with length 2π and constant projection length 4.

Note that everything we have said also works for curves in $\mathbb{R}^n$ (projected to lines in different directions) – only the value of c will be different. Similarly, for an appropriate $c = c_{n,k}$ we get that the length of a curve in R^n is c times the average length of projections to all different k-dimensional subspaces.

For any plane curve α, the length of $\pi_u\alpha$ is at least twice the *width* of α in the direction u. If α is a convex plane curve, we have equality, so Cauchy–Crofton says the length is π times the average width. For instance any curve of constant width 1 (like the Reuleaux triangle on an equilateral triangle of side length 1, named after Franz Reuleaux, Rector at TU Berlin in the 1890s) has length π. A unit square has minimum width 1 and maximum width $\sqrt{2}$; since its length is 4, the average width is $4/\pi$.

Writing the various different lines perpendicular to u as $\ell_{u,a} := \{x : \langle x, u\rangle = a\}$ for $a \in \mathbb{R}$, we see that $\operatorname{len}\pi_u\alpha = \int_{\mathbb{R}} \#(\alpha \cap \ell_{u,a})\, da$. Thus Cauchy–Crofton can be formulated as

$$\operatorname{len}\alpha = \frac{1}{4}\int_0^{2\pi}\int_{\mathbb{R}} \#(\alpha \cap \ell_{u(\theta),a})\, da\, d\theta.$$

A6. Fenchel's theorem

Fenchel's theorem says the total curvature of any closed curve in $\mathbb{R}^n$ is at least 2π. (Equality holds only for convex plane curves.) To prove this for C^1 curves α, recall that the tantrix $T(s)$ has speed $\kappa(s)$ and thus its length is the total curvature of α. On the other hand, the tantrix lies in no open hemisphere of $\mathbb{S}^{n-1}$, for if we had $\langle T(s), M\rangle > 0$ for all s then we would get

$$\begin{aligned} 0 < \int_0^L \langle T(s), M\rangle\, ds &= \Big\langle M, \int_0^L T(s)\, ds\Big\rangle \\ &= \big\langle M, \alpha(L) - \alpha(0)\big\rangle = \langle M, 0\rangle = 0, \end{aligned}$$

a contradiction. Fenchel's theorem is thus an immediate corollary of the theorem below on spherical curves.

We will state all results for general n, but on first reading one should probably think of the case $n = 3$ where α lies on the usual unit sphere $\mathbb{S}^2 \subset \mathbb{R}^3$.

To investigate spherical curves in more detail note first that for points $A, A' \in \mathbb{S}^{n-1} \subset \mathbb{R}^n$ the spherical distance (the length of the shortest spherical path, a great circle arc) between them is

$$\rho(A, A') = \arccos\langle A, A'\rangle = 2\arcsin(|A - A'|/2) \le \pi.$$

The points are *antipodal* if $A = -A'$ (i.e., $\rho = \pi, \ldots$). A nonantipodal pair is connected by a unique shortest arc, with midpoint $M = (A + A')/|A + A'|$.

Lemma: suppose A, A' nonantipodal with midpoint M; suppose $\rho(X, M) < \pi/2$. Then $2\rho(X, M) \le \rho(X, A) + \rho(X, A')$. (This says the distance from X to the segment AA' is a convex function.)

To prove this, first note that A, A', X all lie in some three dimensional subspace of $\mathbb{R}^n$, so we work there, and in particular on $\mathbb{S}^2$. Consider a 2-fold rotation around M, taking A to A' and X to some point X'. Using the triangle inequality and the symmetry, we get

$$2\rho(X, M) = \rho(X, X') \le \rho(X, A) + \rho(A, X') = \rho(X, A) + \rho(X, A')$$

as desired.

Theorem: Suppose α is a closed curve on $\mathbb{S}^{n-1}$ of length $L < 2\pi$. Then α is contained in some spherical cap $\{x \in \mathbb{S}^{n-1} : \rho(X, M) \le L/4\}$ of (angular) radius $L/4 < \pi/2$, and in particular in some open hemisphere. (Note as promised that Fenchel is an immediate corollary.)

To prove this, pick two points A, A' on α dividing the arclength in half. Then $\rho(A, A') \le L/2 < \pi$. Let M be the midpoint and let X be any point on α. If $\rho(X, M) < \pi/2$, then by the lemma,

$$\rho(X, M) \le (\rho(X, A) + \rho(X, A'))/2 \le \operatorname{len}(\alpha_{AXA'})/2 = L/4.$$

Thus the distance from M to any point on α is either at most $L/4$ or at least $\pi/2$. By continuity, the same possibility holds for all X; picking $X = A$ we see it is the first possibility.

There are of course other approaches to proving Fenchel's theorem. One goes through an integral geometry formula analogous to our last version of Cauchy–Crofton. (We'll state it just for curves in $\mathbb{S}^2$ but it holds – with the same constant π – in any dimension.) For $u \in \mathbb{S}^2$, the great circle $u^\perp$ is the set of points orthogonal to u. Then the formula says the length of α equals π times the average number of intersections of α with these great circles. (When α itself is a great circle, this is clear, since there are always 2 intersections.)

First note that the length of a spherical curve is the limit of the lengths of spherical inscribed polygons (made of great circle arcs). (Indeed the spherical inscribed polygon always has length larger than the euclidean polygon with the same vertices, which is already approaching the lenth of the curve from below.) Then just as for Cauchy–Crofton, we check this formula first for great circle arcs, then for polygons and then by a (trickier) limiting argument for smooth curves.

With this formula, one can prove Fenchel's theorem for smooth curves by considering height functions $\langle\alpha(s), u\rangle$. Each has at least two critical points (min, max), but critical points satisfy $T(s) \in u^\perp$. That is, the tantrix intersects every great circle at least twice, and thus has length at least 2π.

Without giving precise definitions about knots, we can understand the Fáry–Milnor theorem: a nontrivally knotted curve in $\mathbb{R}^3$ has total curvature at least 4π. For suppose for some height function $\langle\alpha(s), u\rangle$ there was only one min and one max. At each intermediate height, there are exactly two points of α. Joining these pairs by horizontal segments gives an embedded disk spanning α, showing it is unknotted. For a knotted curve, every height function must have at least four critical points, meaning four intersections of the tantrix with every great circle.

`End of Lecture 22 Apr 2013`

A7. Schur's comparison theorem and Chakerian's packing theorem

Schur's theorem is a precise formulation of the intuitive idea that bending an arc more brings its endpoints closer together.

Suppose α is an arc in $\mathbb{R}^n$ of length L, and consider a comparison arc $\tilde\alpha$ in $\mathbb{R}^2 \subset \mathbb{R}^n$ of the same length, such that with respect to a common arclength parameter s, the curvature of $\tilde\alpha$ is positive and everywhere at least that of of $\tilde\alpha$: $\tilde\kappa(s) \ge \kappa(s)$. Assuming that $\tilde\alpha$ with its endpoints joined by a straight segment gives a convex (simple closed) curve, we conclude that its endpoints are closer:

$$\big|\alpha(L) - \alpha(0)\big| \ge \big|\tilde\alpha(L) - \tilde\alpha(0)\big|.$$

Proof: by convexity, we can find s_0 such that the tangent $T_0 := \tilde T(s_0$ to $\tilde\alpha$ is parallel to $\tilde\alpha(L) - \tilde\alpha(0)$. Move α by a rigid motion so that $\alpha(s_0) = \tilde\alpha(s_0)$ and they share the tangent vector T_0 there. We have

$$\big|\alpha(L) - \alpha(0)\big| \le \langle\alpha(L) - \alpha(0), T_0\rangle = \int_0^L \langle T(s), T_0\rangle\, ds,$$

while for $\tilde\alpha$, our choice of T_0 gives equality:

$$\big|\tilde\alpha(L) - \tilde\alpha(0)\big| = \langle\tilde\alpha(L) - \tilde\alpha(0), T_0\rangle = \int_0^L \langle \tilde T(s), T_0\rangle\, ds,$$

Thus it suffices to show $\langle T(s), T_0\rangle \ge \langle\tilde T(s), T_0\rangle$ (for all s).

We start from s_0 (where both sides equal 1) and move out in either direction. While $\tilde T$ moves straight along a great circle with speed $\tilde\kappa$, a total distance less than π, we see that T moves at slower speed κ and perhaps not straight. Thus is geometrically clear that T is always closer to the starting direction. In formulas,

$$\langle\tilde T(s), T_0\rangle = \cos\int_{s_0}^s \tilde\kappa\, ds \le \cos\int_{s_0}^s \kappa\, ds \le \langle T(s), T_0\rangle.$$

Note that this same proof can be made to work for arbitrary curves of finite total curvature. The case of polygonal curves is known as Cauchy's arm lemma and was used in his proof (1813) of the rigidity of convex polyhedra, although his proof of the lemma was not quite correct.

Chakerian proved the following packing result (which again can be generalized to all curves although we consider only smooth curves): A closed curve of length L in the unit ball in $\mathbb{R}^n$ has total curvature at least L. To check this, simply integrate by parts:

$$\operatorname{len}\alpha = \int \langle T, T\rangle\, ds = \int \langle -\alpha, \vec\kappa\rangle\, ds \le \int |\alpha|\,\kappa\, ds \le \int \kappa\, ds.$$

What about nonclosed curves? We just pick up a boundary term in the integration by parts, and find that length is at most total curvature plus 2.

A8. Framed space curves

We now specialize to consider curves in three-dimensional space $\mathbb{R}^3$. Just as for plane curves we used the 4-fold rotation J, in 3-space we will use its analog, the vector cross product. Recall that $v \times w = -w \times v$ is a vector perpendicular to both v and w.

A *framing* along a smooth space curve α is a (smooth) choice of a unit normal vector $U(s)$ at each point $\alpha(s)$. Defining $V(s) := T(s) \times U(s)$ we have an (oriented) orthonormal frame $\{T, U, V\}$ for $\mathbb{R}^3$ at each point of the curve, and the idea is to follow how this frame rotates. As before, expressing the derivatives in the frame itself gives a skew-symmetric matrix:

$$\begin{pmatrix} T \\ U \\ V \end{pmatrix}' = \begin{pmatrix} 0 & \kappa_U & \kappa_V \\ -\kappa_U & 0 & \tau_U \\ -\kappa_V & -\tau_U & 0 \end{pmatrix} \begin{pmatrix} T \\ U \\ V \end{pmatrix}.$$

Here κ_U, κ_V and τ_U are functions along the curve which depend on the choice of framing. We see that $T' = \vec\kappa = \kappa_U U + \kappa_V V$, so these are just the components of the curvature vector in the chosen basis for the normal plane. (And $\kappa^2 = \kappa_U^2 + \kappa_V^2$.) The third function τ_U measures the twisting or torsion of the framing U.

Sometimes in physical problems a framing is given to us by material properties of a bent rod. Mathematically, the curve α might lie on a smooth surface in space; then we often choose U to be the surface normal so that the *conormal* V is (like T) tangent to the surface. (We will explore such *Darboux frames* in detail when we study surfaces.)

`End of Lecture 25 Apr 2013`

But when no external framing is given to us, there are two ways to choose a nice framing such that one of the entries in the matrix above vanishes. The first has no twisting ($\tau_U = 0$), and such a $\{T, U, V\}$ is called a parallel frame or Bishop frame. Given any U_0 at $\alpha(s_0)$ we want U' to be purely tangential, indeed

$$U' = -\kappa_U T = -\langle \vec{\kappa}, U\rangle\, T.$$

But this ODE has a unique solution. Since it prescribes $U' \perp U$ the solution will have constant length, and since $\langle U', T\rangle = -\langle T', U\rangle$, the solution will stay normal to T. If we rotate a parallel framing by a constant angle φ in the normal plane (that is, replace U by $\cos\varphi U + \sin\varphi V$) then we get another parallel framing (corresponding to a different U_0). Indeed any two parallel framings differ by such a rotation. Parallel frames are very useful, for instance in computer graphics when drawing a tube around a curve. One disadvantage is that along a closed curve, a parallel framing will usually not close up.

The second special framing comes from prescribing $\kappa_V = 0$, i.e., $\vec{\kappa} = \kappa_U U$. That is, U should be the unit vector in the direction $\vec{\kappa}$. Here the disadvantage is that things only work nicely for curves of nonvanishing curvature $\kappa \neq 0$. Assuming this condition, we rename U as the *principal normal* N and V as the *binormal* B and call $\{T, N, B\}$ the *Frenet frame*. We have

$$\begin{pmatrix} T \\ N \\ B \end{pmatrix}' = \begin{pmatrix} 0 & \kappa & 0 \\ -\kappa & 0 & \tau \\ 0 & -\tau & 0 \end{pmatrix} \begin{pmatrix} T \\ N \\ B \end{pmatrix},$$

where $\kappa(s)$ is the curvature and $\tau(s)$ is called the *torsion* of α. In terms of a unit-speed parametrization, we have $\alpha' = T$, $\alpha'' = T' = \vec{\kappa} = \kappa N$, so $N = \vec{\kappa}/\kappa$. Finally, $N' = -\kappa T + \tau B$ so $\tau = \langle N', B\rangle = |N' + \kappa T|$. The expansion of the third derivative in the Frenet frame is

$$\alpha''' = (\kappa N)' = \kappa' N + \kappa N' = -\kappa^2 T + \kappa' N + \kappa\tau B.$$

Expressions in terms of an arbitrary parametrization of α with speed $\sigma(t)$ are left as an exercise. Here the nonvanishing curvature condition just says that $\dot\alpha$ and $\ddot\alpha$ are linearly independent, so that $\{\dot\alpha, \ddot\alpha, \dot\alpha\times\ddot\alpha\}$ is an oriented basis. The orthonormal frame $\{T, N, B\}$ is the result of applying the Gram–Schmidt process to this basis.

Of course N and B span the *normal plane* to α at $p = \alpha(s)$. The curve stays to second order in the *osculating plane* spanned by T and N, which contains the osculating circle. The plane spanned by T and B is called the *rectifying plane* (since the projection of α to that plane has curvature vanishing at p).

The Taylor expansion of α to third order around $p = \alpha(0)$ is

$$\alpha(s) \approx p + \Big(s - \tfrac{s^3}{6}\kappa^2\Big)T + \Big(\tfrac{s^2}{2}\kappa + \tfrac{s^3}{6}\kappa'\Big)N + \Big(\tfrac{s^3}{6}\kappa\tau\Big)B$$

where of course T, N, B, κ, τ and κ' are all evaluated at $s = 0$. Exercise: look at the projections to the three planes above, and see which quadratic and cubic plane curves approximate them.

The “fundamental theorem of space curves” says that given functions $\kappa, \tau\colon I \to \mathbb{R}$ with $\kappa > 0$ determine a space curve (uniquely up to rigid motion) with that curvature and torsion. This is basically a standard theorem about existence and uniqueness of solutions to an ODE. For any given $\{T_0, N_0, B_0\}$ the matrix ODE above has a solution, which stays orthonormal and thus gives a framing. (Changing the initial condition just rotates the frames by a constant rotation.) As in the case of plane curves, integrating $T(s)$ recovers the curve α (uniquely up to translation).

Example: a curve with constant curvature and torsion is a helix. Its tantrix traces out a circle on $\mathbb{S}^2$ at constant speed κ. Any curve whose tantrix lies is a circle on $\mathbb{S}^2$ (i.e., makes constant angle with some fixed vector u) is called a *generalized helix*. Exercise: this condition is equivalent to τ/κ being constant.

Suppose now α is a unit-speed curve with $\kappa > 0$. If $\{T, N, B\}$ is the Frenet frame and $\{T, U, V\}$ is a parallel frame, then how are these related? We have of course

$$\begin{pmatrix} N \\ B \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} U \\ V \end{pmatrix}$$

for some $\theta = \theta(s)$. Then $\vec{\kappa} = \kappa N = \kappa\cos\theta\, U + \kappa\sin\theta\, V$ meaning that $\kappa_U = \kappa\cos\theta$ and $\kappa_V = \kappa\sin\theta$. Differentiating $B = -\sin\theta\, U + \cos\theta\, V$ gives

$$-\tau N = B' = -\theta'(\cos\theta\, U + \sin\theta\, V) + 0T = -\theta' N$$

so that $\theta' = \tau$ or $\theta = \int \tau\, ds$. (The constant of integration corresponds to the freedom to rotate the parallel frame.) We see that the twisting or torsion τ of the Frenet frame really does give the rate θ' at which it rotates relative to the twist-free Bishop frame. Sometimes it is useful to use a *complex curvature* $\kappa(s)e^{i\theta(s)} = \kappa_U(s) + i\kappa_V(s)$. Well defined up to global rotation by $e^{i\theta_0}$ in the complex plane (corresponding again to the freedom to rotate the parallel frame).

It is clear that a space curve lies in a plane if and only if $\tau \equiv 0$, if and only if θ is constant, if and only if the complex curvature stays on some fixed line through 0.

As another example, the complex curvature of a helix traces out the circle $|z| = \kappa$ at constant speed.

Bishop (1975) demonstrated the usefulness of the parallel frame by characterizing (C^2 regular) space curves that lie on some sphere. Indeed, α lies on a sphere of radius $1/d$ if and only if its complex curvature lies on a line at distance d from $0 \in \mathbb{C}$. In an appropriately rotated parallel frame, this line will be the line $\kappa_U \equiv d$. (The characterization in terms of the Frenet frame is more awkward, needing special treatment for points where τ and κ' vanish.)

To prove this, note that by translating and rescaling we can treat the case of $\alpha \subset \mathbb{S}^2$, i.e., $\langle\alpha, \alpha\rangle \equiv 1$. It follows that $\alpha \perp T$ so $U := \alpha$ is a framing of itself. From $\alpha' = T$ we see that this framing is parallel. That is, $U = \alpha$, $V = T \times \alpha$ is a Bishop frame. The equation $U' = T$ means $\kappa_U \equiv 1$, as desired. (Note that since the position vector on $\mathbb{S}^2$ is also the normal vector to the spherical surface, $\{T, U, V\}$ is also the Darboux frame for $\alpha \subset \mathbb{S}^2$!) Conversely, suppose α has a parallel frame $\{T, U, V\}$ with $\kappa_U \equiv 1$, i.e., $U' = T$. Then $\alpha - U$ is a constant point P, meaning α lies on the unit sphere around P.

`End of Lecture 2 May 2013`

A9. Framings for curves in higher dimensions

A *framing* along a smooth curve α in $\mathbb{R}^n$ is a choice of oriented orthonormal basis $\{E_1, E_2, \ldots, E_n\}$ at each point of α, where $E_1(s) = T(s)$ is the unit tangent vector, and each $E_i(s)$ is a smooth function. Of course the other E_i (for $i \geq 2$) are normal vectors. The infinitesimal rotation of any framing is given, as in the three-dimensional case, by a skew-symmetric matrix, here determined by the $\binom{n}{2}$ entries above the diagonal. Again it is helpful to choose special framings where only $n-1$ of these entries are nonzero.

In a parallel framing, these are the entries of the top row. That is, the curvature vector T' is an arbitrary combination $\sum \kappa_i E_i$ of the normal vectors E_i, but each of them is parallel with derivative $-\kappa_i T$ only in the tangent direction. Given any framing at an initial point, solving an ODE gives us a parallel frame along the curve.

The generalized Frenet frame exists only under the (somewhat restrictive) assumption that the first $n-1$ derivatives $\dot{\alpha}$, $\ddot{\alpha}$, ..., $\alpha^{(n-1)}$ are linearly independent, and $\{T, E_2, \ldots, E_n\}$ is then the Gram–Schmidt orthonormalization of these vectors. For this frame, it is only the matrix entries just above the diagonal that are nonzero. Thus

$$E_i' := \tau_i E_{i+1} - \tau_{i-1} E_{i-1}$$

In particular $T' = \tau_1 E_2$ so $\tau_1 = \kappa$ is the usual curvature and E_2 is the *principal normal* (the unit vector in the direction of $\vec{\kappa}$). The τ_i are called Frenet curvatures. A “fundamental theorem” says that for any functions $\tau_i(s)$ with $\tau_i > 0$ for $i < n-1$, there is a curve with these Frenet curvatures; it is unique up to rigid motion.