CONVEX REGIONS, SHADOWS, AND THE GAUSS MAP

CHARLES L. EPSTEIN

Introduction

We consider a simple model problem for medical image reconstruction. Let $D \subset \mathbb{R}^2$ be a bounded convex region in the plane. Supposing that D is not accessible to direct measurements, we consider how to reconstruct D by observing its shadows from all possible directions. This leads us to consider the Gauss map as a way to represent curves in the plane. Next we attempt to reconstruct the boundary of a region by measuring its width in all directions. Exercises are provided to help the reader test his knowledge of the material.

1. Basic concepts

In this note we are concerned with convex regions in the plane. A open subset $D \subset \mathbb{R}^2$ is convex if, for every pair of points, $p, q \in D$ the line segment l_{pq} joining p to q lies entirely within D. By letting p and q tend (separately) to the boundary of D we see that the same condition holds for points p and q in the closure of D. This implies that if $l \subset \mathbb{R}^2$ is any line then $l \cap D$ is either the empty set or a connected segment contained in D.

Convex and non-convex regions.

Let $D \subset \mathbb{R}^2$ and suppose that x is a point on the boundary of D. Because D is an open set we can choose two sequences $< p_n >, < q_n >$ contained in the boundary of D so that

((1)) $p_n \neq q_n$ for any n,
((2)) $\lim_{n\to\infty} p_n = x = \lim_{n\to\infty} q_n$.

We would like to make an assertion about the behavior of the lines containing the segments $l_{p_n q_n}$ as $n \to \infty$ but, we have no way to discuss the convergence of a sequence of lines. What is required is a parametrization for the set of lines in the plane. A line is the set of solutions to a linear equation

$$l_{a,b,c} = \{(x, y) \,:\, ax + by = c\} \text{ where } a^2 + b^2 \neq 0.$$

The coefficients (a, b, c) could be used as parameters for the set of lines. For many purposes it is not a very good choice: if $\alpha \neq 0$ then (a, b, c) and $(\alpha a, \alpha b, \alpha c)$ define the same line. A better parametrization is obtained by normalizing the pair (a, b) to have unit length. In this way pairs $(t, \omega) \in \mathbb{R} \times S^1$ parametrize lines with

$$l_{t,\omega} = \{(x, y) \,:\, \langle (x, y), \omega \rangle = t\}.$$

As $l_{t,\omega} = l_{-t,-\omega}$, each line appears twice in this parametrization. In fact the pair (t, ω) specifies a unique *oriented line.* A line is oriented by selecting a direction along the line. This direction is often called the *positive direction.* To see that $\mathbb{R} \times S^1$ parametrizes the oriented lines it is convenient to let

$$\omega(\theta) = (\cos\theta, \sin\theta) \text{ and } \hat{\omega}(\theta) = (-\sin\theta, \cos\theta). \tag{1}$$

The second vector is determined by the conditions

$$\langle \omega, \hat{\omega} \rangle = 0 \quad \det(\omega\hat{\omega}) > 0.$$

In our definition ω is a direction orthogonal to the line $l_{t,\omega}$, and therefore $\hat{\omega}$ is a vector parallel to it; this vector defines the positive direction along the line. The line $l_{t,\omega}$ has the parametric representation

$$l_{t,\omega} = \{t\omega + s\hat{\omega} \,:\, s \in \mathbb{R}\}.$$

Exercise 1.1. Show that $|t|$ is the distance from $(0,0)$ to $l_{t,\omega}$.

We have establised the following proposition.

Proposition 1.1. *The oriented lines in* $\mathbb{R}^2$ *are in one-to-one correspondance with the points of* $\mathbb{R} \times S^1$.

Definition 1.1. A sequence of oriented lines $< l_{t_n,\omega_n} >$ converges to the oriented line $l_{t,\omega}$ if and only if

$$\lim_{n\to\infty} \omega_n = \omega \text{ and } \lim_{n\to\infty} t_n = t.$$

Remark 1.1. The set of (un-oriented) lines is the quotient space

$$\mathbb{R} \times S^1/\sim \text{ where } (t,\omega) \sim (-t,-\omega).$$

For our applications the set of oriented lines is more useful.

We now consider the convergence of the sequence of lines containing the seqments $< l_{p_nq_n} >$. To each line we can associate a direction

$$\hat{\omega}_n = \frac{p_n - q_n}{|p_n - q_n|},$$

a normal direction ω_n and an affine parameter t_n. The segment $l_{p_nq_n}$ lies in l_{t_n,ω_n}. Since the unit circle is compact the sequence $< \omega_n >$ has a convergent subsequence, which for simplicity we continue to denote by $< \omega_n >$. In light of exercise 1.1, the affine parameters also lie in a bounded set and therefore have a convergent subsequence. Hence there is a subsequence $< n_j >$ so that the lines $l_{t_{n_j},\omega_{n_j}}$ converge to a line l_{t^*,ω^*}. The common limit point, x of the sequences $< p_{n_j} >, < q_{n_j} >$ must lie on l_{t^*,ω^*}. On the other hand

$$l_{p_{n_j}q_{n_j}} = \overline{D} \cap l_{t_{n_j},\omega_{n_j}}$$

and therefore the limiting line does not contain any points in D. This gives a different characterization of a convex set: For each point x on the boundary of D there is a line l_x which contains x but such that $l_x \cap D = \emptyset$. Such a line is called a *support line at* x.

The support line as a limit of chords.

Proposition 1.2. *An open subset $D \subset \mathbb{R}^2$ is convex if and only if each point $x \in bD$ has a support line.*

Exercise 1.2. Prove the converse part of this Proposition.

Remark 1.2. A point on the boundary of convex region may have more than one support line and a support line can contain more than one boundary point. The former case arises if the boundary of D has "corners" and the latter case arises if the boundary of D contains line segments. If each support line meets the boundary in exactly one point then we say that D is *strictly convex.*

Exercise 1.3. Show that a point on bD has a tangent line if and only if it has a *unique* support line.

Some pathologies of convex sets.

A line divides the plane into two half spaces. Given a pair (t, ω) define

$$\begin{aligned} H^+_{t,\omega} &= \{(x,y) \,:\, \langle (x,y), \omega\rangle \geq t\} \\ H^-_{t,\omega} &= \{(x,y) \,:\, \langle (x,y), \omega\rangle \leq t\}. \end{aligned} \tag{2}$$

Exercise 1.4. Show that

$$H^{\pm}_{t,\omega} = H^{\mp}_{-t,-\omega}.$$

If $l_{t,\omega}$ is a supporting line for a convex region D then D is contained in either $H^+_{t,\omega}$ or $H^-_{t,\omega}$. An orientation is fixed on the boundary of D by choosing (t, ω) so that $D \subset H^-_{t,\omega}$. We call a support line with this property an *oriented support line.* At points of bD which have a tangent line, the orientation of the boundary is defined to coincide with $\hat{\omega}$ where $l_{t,\omega}$ is the oriented support line.

Exercise 1.5. Show that the orientation of bD defined by the oriented support lines is the counterclockwise direction.

Exercise 1.6. Let D be a bounded convex domain. For each direction $\omega \in S^1$ show that there is a unique $t \in \mathbb{R}$ so that $l_{t,\omega}$ is an oriented, support line for D.

Animation showing family of support lines.

2. Shadows

We would now like to determine a bounded convex region from a knowledge of its "shadows." The first step is to give a precise description of a shadow. For this purpose we imagine a light source placed very far away from D so that the rays of light emanating from the source are very close to parallel, with common direction $\hat{\omega}$. The shadow, in dierction ω is the projection of D on a screen, placed beyond D, at right angles to the direction $\hat{\omega}$. The portion of the screen which is dark is the shadow. To describe the shadow we need to find the maximal interval $(t_{\min}, t_{\max})$ so that

$$l_{t,\omega} \cap D \neq \emptyset \text{ if } t_{\min} < t < t_{\max}.$$

The lines $l_{t_{\min},\omega}$ and $l_{t_{\max},\omega}$ are support lines for bD. For a direction ω we define *the shadow function* of D to be

$$H_D(\omega) = t_{\max}.$$

With this choice, the region lies in the half space $H^-_{H_D(\omega),\omega}$ for every $\omega \in S^1$.

Exercise 2.1. Show that $H_D(-\omega) = t_{\min}$.

The shadow projected onto a screen.

To continue, it is useful to use the parametrization of S^1 given in (1). In this representation θ is only defined up to multiples of 2π. Note that

$$\partial_\theta \omega = \hat{\omega} \text{ and } \partial_\theta \hat{\omega} = -\omega. \tag{3}$$

To simplify notation we set $h_D(\theta) = H_D(\omega(\theta))$. The line $l_{h_D(\theta),\omega(\theta)}$ meets bD. If D is strictly convex then each support line touches bD at exactly one point, hence there is unique number $s(\theta)$ so that the point

$$Z_D(\theta) = h_D(\theta)\omega(\theta) + s(\theta)\hat{\omega}(\theta) \in bD.$$

If we can determine the function $s(\theta)$ then we can determine the boundary of D. To do that one only needs to observe that the tangent line to bD at the point $Z_D(\theta)$ is parallel to $\hat{\omega}(\theta)$. Assuming that $s(\theta)$ is differentiable and using (3) gives the direction of the tangent line at $Z_D(\theta)$

$$\partial_\theta Z_D(\theta) = (h'_D - s)\omega(\theta) + (h_D + s')\hat{\omega}(\theta).$$

In order for $\partial_\theta Z_D(\theta)$ to be parallel to $\hat{\omega}(\theta)$ it is necessary that

$$s(\theta) = h'_D(\theta).$$

If the boundary of D is differentiable and strictly convex then it has a parametric representation in terms of its shadow function

$$Z_D(\theta) = h_D(\theta)\omega(\theta) + h'_D(\theta)\hat{\omega}(\theta). \tag{4}$$

If bD is not strictly convex or non-differentiable then this argument does not apply. Instead we use a more geometric construction. Since, for every $\omega \in S^1$, $D \subset H^-_{H_D(\omega),\omega}$ it follows that

$$D \subset \bigcap_{\omega \in S^1} H^-_{H_D(\omega),\omega}.$$

This interesection is exactly the closure of D. To verify this claim we need to show that

$$\overline{D} \supset \bigcap_{\omega \in S^1} H^-_{H_D(\omega),\omega}.$$

which is equivalent to the assertion that

$$\overline{D}^c \subset \bigcup_{\omega \in S^1} \left[H^-_{H_D(\omega),\omega} \right]^c.$$

Let p be a point in $\overline{D}^c$ and let q be the point in $\overline{D}$ nearest to p.

Exercise 2.2. Show that the point q is unique.

If $\widetilde{q}$ is any other point in $\overline{D}$ then the line segment $\{(1-t)q+t\widetilde{q} : t \in [0,1]\}$ is contained in $\overline{D}$. As q is the point in $\overline{D}$ closest to p we have the inequality

$$F(t) = \langle p - ((1-t)q + t\widetilde{q}), p - ((1-t)q + t\widetilde{q})\rangle > F(0) \text{ for } t \in (0,1].$$

Differentiating F at $t = 0$ gives

$$0 \leq F'(0) = 2\langle p - q, q - \widetilde{q}\rangle$$

which implies that

$$\langle p - q, \widetilde{q}\rangle \leq \langle p - q, q\rangle. \tag{5}$$

Set $\omega = (p-q)/\|p-q\|$ and $t = \langle \omega, q\rangle$, then (5) shows that $l_{t,\omega}$ is a support line for D at q. Indeed D lies in $H^{-}_{t,\omega}$. In light of exercise 1.3, $t = H_D(\omega)$. As $p \notin \overline{D}$ the distance $\sqrt{\langle p - q, p - q\rangle} > 0$ and therefore

$$\langle \omega, q\rangle < \langle \omega, p\rangle.$$

This shows that $p \in \text{int}\, H^{+}_{t,\omega}$ and proves that

$$\overline{D} = \bigcap_{\omega \in S^1} H^{-}_{H_D(\omega),\omega}. \tag{6}$$

The reconstruction problem has several different solutions. The first method only works under restrictive hypotheses whereas the second method is generally applicable. The parametric solution, (4) is somewhat neater but, in truth, less useful in applications. To use this formula we need to know $h_D(\theta)$, exactly for all values of θ and this is rarely the case for real measurements. Usually we can expect to have approximate values for $h_D(\theta)$ for θ belonging to a finite set $\{\theta_1, \dots, \theta_m\}$. For a finite set the interesection

$$\bigcup_{j \in \{,\dots,m\}} H^{-}_{h_D(\theta_j),\omega(\theta_j)}$$

is a convex polygon containing D. Such a polygon is an *outer* approximation to D. If instead of the exact values $\{h_D(\theta_j)\}$ we have instead approximate values $\{h^m_D(\theta_j)\}$ for the shadow function, then

$$\widetilde{D}_m = \bigcup_{j \in \{1,\dots,m\}} H^{-}_{h^m_D(\theta_j),\omega(\theta_j)}$$

defines a convex polygon which again provides a reasonable *outer* approximation to $\overline{D}$.

To use (4) to obtain an approximate reconstruction requires more accurate measurements as we need to approximate $h'_D(\theta_j)$. Using a finite difference approximation we can obtain approximate values for a finite set of points,

$$\{h^m_D(\theta_j)\omega(\theta_j) + \frac{h^m_D(\theta_j) - h^m_D(\theta_{j-1})}{\theta_j - \theta_{j-1}}\hat{\omega}(\theta_j)\},$$

which lie on bD. Connecting adjacent points with straight lines gives an *inner* approximation to D, that is a polygon which lies inside of D. In order for the difference quotient,

$$\frac{h^m_D(\theta_j) - h^m_D(\theta_{j-1})}{\theta_j - \theta_{j-1}}$$

to be close to $h'_D(\theta_j)$ is it generally necessary to take $|\theta_j - \theta_{j-1}|$ small. Moreover the measurement errors

$$\{|h^m_D(\theta_j) - h_D(\theta_j)|, |h^m_D(\theta_{j-1}) - h_D(\theta_{j-1})|\}$$

must be small *compared to* $|\theta_j - \theta_{j-1}|$. This explains the difficulty of using this formalism in a practical problem.

An outer approximation to a convex region.

3. The range of the shadow function map

Suppose that $h(\theta)$ is a function defined on the unit circle. Is h the shadow function of a convex domain in the plane? We restrict attention to the somewhat easier case that h is twice differentiable and ask if it is the shadow function for a smooth, *strictly* convex domain. From formula (4) it follows that the boundary of the domain must have the parametric representation

$$Z_h(\theta) = h(\theta)\omega(\theta) + h'(\theta)\hat{\omega}(\theta).$$

In addition we know that $\hat{\omega}(\theta)$ should be the *positively* oriented tangent direction at $Z_h(\theta)$. Computing the derivative we find that

$$\partial_\theta Z_h(\theta) = (h + h'')\hat{\omega}(\theta).$$

A necessary condition for h to be the shadow function of a strictly convex domain is that

$$h''(\theta) + h(\theta) > 0 \text{ for all } \theta \in [0, 2\pi). \tag{7}$$

For a smooth function which satisfies (7), $\partial_\theta Z_h$ is non-vanishing and therefore the image, $Z_h(S^1)$ is a smooth, oriented curve Γ_h, immersed in the plane.

The sum $h'' + h$ has a nice geomeric interpretation. If Γ is a curve in the plane, with arclength parametrization $\gamma(s)$ then

$$\tau(s) = \gamma'(s)$$

a unit vector, tangent to Γ. Its direction defines an orientation on the curve. Let $\nu(s)$ be the unit vector orthogonal to $\tau(s)$ such that $\det(\tau(s)\nu(s)) > 0$. The derivative of ν is a multiple of τ,

$$\nu'(s) = \kappa(s)\tau(s).$$

The function $\kappa(s)$ is called the curvature of the curve. The sign of the curvature of a curve depends on the choice of orientation, changing the orientation reverses the sign of the curvature. For example, the counterclockwise oriented unit circle has curvature $+1$.

Proposition 3.1. *If h is a smooth 2π-periodic function which satisfies* (7) *then the curvature of Γ_h at $Z_h(\theta)$ is given by*

$$\kappa(\theta) = \frac{1}{h(\theta) + h''(\theta)}. \tag{8}$$

The proof is a calculation using the definition of curvature and is left to the reader. The condition on the shadow function is simply that, with the orientation given by Z_h, the curve has positive curvature at each point.

To show that h is the shadow function of the strictly convex domain we need to show that Γ_h has no self intersections. It follows from the implicit function theorem that for each $\theta \in S^1$ there is an $\epsilon > 0$ so that the arc

$$\gamma_{\theta,\epsilon} = Z_h((\theta - \epsilon, \theta + \epsilon))$$

is embedded with positive curvature at each of its points. From the discussion above it is clear that, for small enough $\epsilon > 0$,

$$\gamma_{\theta,\epsilon} \subset H^-_{h(\theta),\omega(\theta)} \tag{9}$$

We use the Gauss map to show that the curve Γ_h is embedded, that is

$$Z_h(\theta_0) = Z_h(\theta_1) \text{ iff } \theta_0 = \theta_1 \mod 2\pi.$$

Definition 3.1. Let Γ be an oriented, immersed curve in the plane given by a $\mathcal{C}^1$-map $\gamma : [0, 1] \to \mathbb{R}^2$ with $\gamma'(t) \neq 0$ for any $t \in [0, 1]$. The map which assigns to a point on Γ its unit tangent vector is called the Gauss map,

$$g_\gamma(t) = \frac{\gamma'(t)}{\|\gamma'(t)\|}.$$

As Z_h parametrizes Γ_h in terms of its tangent vector the Gauss map is given by

$$g_h(\theta) = \hat{\omega}(\theta).$$

The Gauss map is a strictly monotone map from Γ_h to the unit circle, which goes around the circle exactly once. To show that Γ_h must be embedded we choose a direction ω and consider the family of lines normal to this direction $\{l_{t,\omega}\}$. For $t >> 0$ Γ_h is disjoint from the line $l_{t,\omega}$. Letting t decrease we find the largest $t = t_{\max}$ so that $l_{t,\omega} \cap \Gamma_h \neq \emptyset$. Clearly $l_{t_{\max},\omega}$ is tangent to Γ_h at a point $Z_h(\theta_0)$. In fact as the curvature of Γ_h is positive, and Γ_h lies in $H^-_{t_{\max},\omega}$ we conclude that $\hat{\omega}(\theta_0)$ is the oriented tangent vector at $Z_h(\theta_0)$.

This in turns shows that Γ_h must be embedded in a neighborhood of $Z_h(\theta_0)$. Otherwise there would exist $\theta_1 \neq \theta_0$ with $Z_h(\theta_0) = Z_h(\theta_1)$. The tangent vectors at these two points must be parallel for otherwise Γ_h could not lie in a half space. This leaves only the possibility that $\theta_0 = \theta_1 + \pi$ mod 2π. Because the Gauss map is one-to-one this implies that

$$\hat{\omega}(\theta_0) = -\hat{\omega}(\theta_1) \text{ and } h(\theta_1) = -h(\theta_0).$$

From (9) it follows that a neighborhood of $Z_h(\theta_0) = Z_h(\theta_1)$ lies in both $H^-_{h(\theta_0),\omega(\theta_0)}$ and $H^-_{-h(\theta_0),-\omega(\theta_0)}$ and hence, in their intersection,

$$H^-_{h(\theta_0),\omega(\theta_0)} \cap H^+_{h(\theta_0),\omega(\theta_0)} = l_{h(\theta_0,\omega(\theta_0)}.$$

This is impossible because Γ_h has positive curvature and so cannot agree with a straight line over an interval!

Each point on Γ_h which is obtained as a "first point of contact" with a family of parallel lines $\{l_{t,\omega} : t > t_{\max}\}$ has a neighborhood which is embedded. As the choice of direction $\omega \in S^1$ is arbitrary and the Gauss map from Γ_h to S^1 is one-to-one this shows that every point on Γ_h arises as such a first point of contact. Thus the condition (7) is the necessary and sufficient for h to be the shadow function of region with a smooth, strictly convex boundary. We summarize this discussion as a theorem.

Theorem 3.1. *Let h be a twice differentiable function on S^1 then h is the shadow function of a strictly convex region with differentiable boundary if and only if*

$$h''(\theta) + h(\theta) > 0 \text{ for all } \theta \in [0, 2\pi).$$

We close this section with some exercises showcasing the marvelous properties of the Gauss map and shadow function.

Exercise 3.1. Suppose that h is a function satisfying (7). Show that another formula for $Z_h(\theta)$ is

$$\theta \mapsto (-\int_0^\theta (h(s) + h''(s)) \sin(s) ds, \int_0^\theta (h(s) + h''(s)) \cos(s) ds).$$

Exercise 3.2. Suppose that h is a function satisfying (7). Show that the area enclosed by Γ_h is given by the

$$\text{Area}(D_h) = \frac{1}{2} \int_0^{2\pi} [(h(\theta))^2 - (h'(\theta))^2] d\theta.$$

Explain why this implies that a function satisfying (7) also satisfies the estimate

$$\int_0^{2\pi} (h'(\theta))^2 d\theta < \int_0^{2\pi} (h(\theta))^2 d\theta.$$

Exercise 3.3. Which positive functions $\kappa(\theta)$ defined on S^1 are the curvatures of closed convex curves? Prove the following result: A positive function $\kappa(\theta)$ on S^1 is the curvature of a closed, strictly convex curve (parametrized by its tangent direction) if and only if

$$\int_0^\infty \frac{\sin(s) ds}{\kappa(s)} = 0 = \int_0^\infty \frac{\cos(s) ds}{\kappa(s)}.$$

Exercise 3.4. Let D be a convex region with shadow function h_D. For a vector $v \in \mathbb{R}^2$ define the translated region

$$D^v = \{(x, y) + v : (x, y) \in D\}.$$

Find the relation between h_D and h_{D^v}. Explain why this answer is inevitable in light of the formula for the curvature.

Exercise 3.5. Let D be a convex region with shadow function h_D. For a rotation $A \in SO(2)$ define the rotated region

$$D^A = \{A(x,y) \, : \, (x,y) \in D\}.$$

Find the relation between h_D and h_{D^A}.

4. Determining a region from its width

In the previous section we showed that a convex region can be reconstructed from its shadow function. This suggests the following question: Can we reconstruct a region knowing only the widths of its shadows in all directions? This question can easily be rephrased in terms of the shadow function. Supppose that $h_D(\theta)$ is the shadow function of a convex domain then the width of the shadow in the direction θ is simply

$$W_D(\theta) = h_D(\theta) + h_D(\theta + \pi).$$

Exercise 4.1. Prove this relation.

From this formula it follows that a region *cannot* be determined from the width of its shadows. The question is simply whether there are two domains D and D' so that $W_D \equiv W_{D'}$. Using the results of the previous section we can easily construct such pairs.

The width of the shadow in a given direction.

Let h be a smooth function on S^1 which satisfies (7). As shown above, there is a strictly convex domain D with $h_D = h$. Let k be another smooth function on S^1 which satisfies

$$k(\theta) = -k(\theta + \pi) \text{ for all } \theta. \tag{10}$$

Since h satisfies (7) and k is smooth, for sufficiently small ϵ the sum $h + \epsilon k$ also satisfies (7). This means that there is a domain D_ϵ so that $h_{D_\epsilon} = h + \epsilon k$. On the other hand

$$[h(\theta) + \epsilon k(\theta)] + [h(\theta + \pi) + \epsilon k(\theta + \pi)] = h(\theta) + h(\theta + \pi).$$

That is D and D_ϵ have the same width in every direction.

There are many functions which satisfy (10). This condition is equivalent to a condition on the Fourier series of k.

Proposition 4.1. *A function k on S^1 satisfies* (10) *if and only if*

$$k(\theta) = \sum_{j=1}^{\infty} [a_j \sin(2j+1)\theta + b_j \cos(2j+1)\theta].$$

The proof is left as an exercise. From the proposition we deduce that the set of strictly convex domains with a given width function is infinite dimensional. In particular this means that there are domains with the same width functions which are not translates of one another. A very interesting example is given below.

Animation of a non-circular region with constant width.

Department of Mathematics, University of Pennsylvania
E-mail address: cle@math.upenn.edu