On divergence-based distance functions for multiply-connected domains

Renjie Chen \cdot Craig Gotsman \cdot Kai Hormann

Abstract

Given a finitely-connected bounded planar domain Ω , it is possible to define a *divergence distance* D(x, y) from $x \in \Omega$ to $y \in \Omega$, which takes into account the complex geometry of the domain. This distance function is based on the concept of *f*-divergence, a distance measure traditionally used to measure the difference between two probability distributions. The relevant probability distributions in our case are the Poisson kernels of the domain at *x* and at *y*. We prove that for the χ^2 -divergence distance, the gradient by *x* of *D* is opposite in direction to the gradient by *x* of G(x, y), the Green's function with pole *y*. Since *G* is harmonic, this implies that *D*, like *G*, has a single extremum in Ω , namely at *y* where *D* vanishes. Thus *D* can be used to trace a gradient-descent path within Ω from *x* to *y* by following $\nabla_x D(x, y)$, which has significant computational advantages over tracing the gradient of *G*. This result can be used for robotic path-planning in complex geometric environments.

Citation Info

Published online January 2018 Preprint arXiv:1801.07099 [math.CV] DOI 10.48550/arXiv.1801.07099

1 Introduction

The *f*-divergence function, first introduced by Kullback and Leibler [6] and later generalized by Csiszár [4], measures the difference between two probability distributions.

Definition 1. Let $f : \mathbb{R} \to \mathbb{R}$ be a strictly convex function with f(1) = 0 and $p, q : E \to [0, 1]$ be two non-negative real functions on some domain *E* such that

$$\int_{E} p(t)dt = \int_{E} q(t)dt = 1.$$

$$d_{f}(p,q) = \int_{E} p(t)f\left(\frac{p(t)}{q(t)}\right)dt$$
(1)

We then call

the f-divergence between p and q.

It is well known [6] that

and

$$d_f(p,q) \ge 0$$

$$l_f(p,q) = 0 \qquad \Longleftrightarrow \qquad p = q.$$

The *f*-divergence is not necessarily a metric, since it is typically not symmetric and does not satisfy the triangle inequality. Many instances of *f* have been proposed over the years, each suitable for some specific application, mostly in probability theory, statistics, and information theory. The interested reader is referred to Liese and Vajda [7] for a survey of the possibilities. We just mention the most popular choices of *f*: Kullback–Leibler: $f(t) = -\log t$, Total Variation: f(t) = |t-1|, Chi Squared (χ^2): $f(t) = t^2 - 1$, Hellinger: $f(t) = (\sqrt{t} - 1)^2$.

In this context, it is useful to remember the concept of the *dual function* of *f*,

$$f^*(t) = t f(1/t).$$

For example, if $f(t) = -\log t$, then $f^*(t) = t \log t$, and if f(t) = |t-1|, then $f^*(t) = f(t)$. It is easy to verify that

- f is (strictly) convex if and only if f^* is (strictly) convex;
- $d_f(p,q) = d_{f^*}(q,p);$
- $f(1) \le d_f(p,q) \le f(0) + f^*(0)$.

Chen et al. [3] took advantage of the concept of f-divergence to define the f-divergence distance $D_f(x, y)$ between pairs of points in a planar domain Ω (see Section 2) and showed that this distance can play an important role in robotic path-planning applications. In that context, given two points x, y in a geometrically complex domain Ω , it is important to be able to generate a path from x to y which stays completely within Ω . Given a non-negative distance function D(x, y) on the domain, which vanishes only when x = y, it is possible to trace such a path by following the gradient of D, as long as this gradient does not vanish at some local minimum of D. Fig. 1 (Left) shows some examples of such paths. Chen et al. [3] show that the f-divergence distance $D_f(x, y)$ has this important property in the case of a simply-connected bounded domain. Their proof uses the concept of conformal invariance and relies heavily on the Riemann conformal mapping theorem [1], namely, that any simply-connected domain is conformally equivalent to the unit disk. They show that D_f is a conformal invariant of the domain, as is the Green's function G, leading to the conclusion that the gradient of the f-divergence distance function is opposite in direction to the gradient of the Green's function, independent of f. Since the harmonic G has the required property, then so does D_f .

Although D_f and G both have the same desirable property and generate identical gradient paths between points, Chen et al. [3, 2] show that there are significant computational advantages to using D_f instead of G in practical path-planning applications, where the domain is discretized.

In this paper we generalize the results of Chen et al. [3] to the more challenging case of multiply-connected domains. Since these are not conformally equivalent to the unit disk, the proof technique used for the simply-connected domain is no longer applicable. Instead, in our central Theorem 3, we provide a direct proof for the χ^2 -divergence distance using methods of complex analysis (see Section 4).

2 The *f*-divergence distance

The *f*-divergence distance between two points in a bounded planar domain Ω is defined using the *Poisson kernel* of Ω at $y \in \Omega$,

$$P(w, y) = -\frac{1}{2\pi} \frac{\partial G}{\partial n(w)}(w, y), \qquad w \in \partial \Omega,$$
⁽²⁾

which is the derivative of the Green's function G(x, y) with pole y at the boundary point w in the direction of the unit outer normal n(w) [5]. Thus P is positive in Ω and for all $y \in \Omega$,

$$\oint_{\partial\Omega} P(w, y) |dw| = 1$$

where |dw| is the usual arc length differential.

Definition 2. Let *f* be a strictly convex function with f(1) = 0 and Ω be a bounded planar domain. The *f*-*divergence distance* from $x \in \Omega$ to $y \in \Omega$ is the contour integral

$$D_f(x,y) := d_f(P(\cdot,x), P(\cdot,y)) = \oint_{\partial\Omega} P(w,y) f\left(\frac{P(w,x)}{P(w,y)}\right) |dw|, \tag{3}$$

The following theorem summarizes the basic properties of the f-divergence distance.

Theorem 1. Let $D_f(x, y)$ be a f-divergence distance on a planar domain Ω . Then:

- D_f is non-negative: $D_f(x, y) \ge 0$.
- D_f is symmetric: $D_f(x, y) = D_f(y, x)$.
- D_f is constant on the boundary $\partial \Omega$ (unless it is infinite there).
- D_f is strictly subharmonic: for any $y \in \Omega$, $\nabla_x^2 D_f(x, y) > 0$.

Proof. See Chen et al. [3].

The second property is a little surprising, since in general the f-divergence of two probability functions (1) is not symmetric. Despite this, the f-divergence distance will typically not be a metric, as it does not satisfy the triangle inequality. A rare exception is the Total Variation divergence distance

$$D_{\mathrm{TV}}(x, y) = \oint_{\partial \Omega} |P(w, x) - P(w, y)| |dw|.$$

Relying on the symmetry of D_f and the properties of the dual function f^* , we have:

Corollary 1. The functions f and f^* generate identical divergence distances,

$$D_f(x, y) = D_{f^*}(x, y).$$

3 Preliminaries

To prove our central Theorem 3, we need a few preliminaries about Green's functions and their consequences on divergence distance functions.

Let Ω be an open, bounded, and finitely-connected domain. Following [5], we consider Ω as a subset of \mathbb{C} and define for any $y \in \Omega$ the *Green's function with pole at y* as

$$G(x, y) = \log \frac{1}{|x - y|} + H(x, y), \qquad x \in \overline{\Omega} \setminus \{y\}, \tag{4}$$

where H(x, y) is the harmonic solution of the Dirichlet problem for the boundary values $f(w) = \log |w - y|$, $w \in \partial \Omega$. Thus *G* vanishes on $\partial \Omega$ and is positive and harmonic on $\Omega \setminus \{y\}$. Furthermore, *G* is symmetric: G(x, y) = G(y, x). We may treat *x* and *y* as complex numbers and, taking advantage of the convenient complex algebra, use the subscripts *x*, *y*, \bar{x} , and \bar{y} to denote the Wirtinger derivatives of G(x, y) with respect to its first and second variable and their conjugates, for example,

$$G_x = \frac{\partial G}{\partial x}, \qquad G_{\bar{x}} = \frac{\partial G}{\partial \bar{x}}, \qquad G_{x\bar{y}} = \frac{\partial^2 G}{\partial x \partial \bar{y}}, \qquad G_{\bar{x}y} = \frac{\partial^2 G}{\partial \bar{x} \partial y},$$

and likewise for *H*. Note that according to this notation, we have, for example,

$$\frac{\partial}{\partial w}G(w,y) = G_x(w,y), \qquad \nabla_x G(w,x) = 2\frac{\partial G}{\partial \bar{x}} \qquad \nabla_y G_x(w,y) = 2\frac{\partial}{\partial \bar{y}}G_x(w,y) = 2G_{x\bar{y}}(w,y).$$

We further say that two complex numbers $a, b \in \mathbb{C}$ are *proportional* and write $a \propto b$, if $\arg a = \arg b$, namely $a/b \in \mathbb{R}$ with a/b > 0.

Observation 1. The Poisson kernel (2) satisfies

$$P(w, y) = \frac{1}{2\pi} |\nabla_x G(w, y)| = \frac{1}{\pi} |G_{\bar{x}}(w, y)| = \frac{1}{\pi} |G_x(w, y)|,$$

because Green's function is real, positive in Ω , and vanishes at $\partial \Omega$, hence its gradient at $w \in \partial \Omega$ points in the opposite direction of n(w), the unit inward normal at w.

Observation 2. For any $w \in \partial \Omega$ and $y, z \in \Omega$ with $y \neq z$, we have

$$\frac{P(w, y)}{P(w, z)} = \left| \frac{G_x(w, y)}{G_x(w, z)} \right| = \frac{G_x(w, y)}{G_x(w, z)}$$

by Observation 1 and because $G_x(w, y) \propto G_x(w, z)$.

Using these observations, we can express the f-divergence distance as a complex contour integral in terms of G_x .

Lemma 1. The *f*-divergence distance (3) satisfies

$$D_f(x, y) = \frac{i}{\pi} \oint_{\partial \Omega} f\left(\frac{G_x(w, x)}{G_x(w, y)}\right) G_x(w, y) dw.$$

Proof. By Observation 1,

$$D_f(x,y) = \frac{1}{\pi} \oint_{\partial\Omega} f\left(\frac{P(w,x)}{P(w,y)}\right) |G_{\bar{x}}(w,y)| |dw| = \frac{1}{\pi} \oint_{\partial\Omega} f\left(\frac{P(w,x)}{P(w,y)}\right) |G_{\bar{x}}(w,x)dw|.$$

Since $G_{\bar{x}}(w, y)$ and dw are orthogonal to each other, with $G_{\bar{x}}(w, y)$ pointing 90° to the left of dw, we have

$$|G_{\bar{x}}(w,y)dw| = i\overline{G_{\bar{x}}(w,y)}dw = iG_{x}(w,y)dw,$$
(5)

resulting in

$$D_f(x, y) = \frac{i}{\pi} \oint_{\partial \Omega} f\left(\frac{P(w, x)}{P(w, y)}\right) G_x(w, y) dw.$$

Note that the complex integral is computed counter-clockwise for the outer boundary loop and clockwise for the internal boundary loops. The Lemma then follows by Observation 2. \Box

Observation 3. Since G(x, y) is harmonic in x and y, G_x is a function of x and not of \bar{x} . For the same reason, $G_{x\bar{y}}$ is a function of only x and \bar{y} . As a result, G_x and $G_{x\bar{y}}$ are both holomorphic in x, except at possible poles. **Lemma 2.** For any $x, y \in \Omega$ with $x \neq y$,

$$G_x(x,y) \propto \frac{1}{G_{\bar{x}}(x,y)}.$$

Proof. Since *G* is real, we have $G_{\bar{x}} = \overline{G_x}$, and the Lemma then follows, because $a \propto 1/\bar{a}$ for any $a \in \mathbb{C} \setminus \{0\}$. \Box **Lemma 3.** For any $y, z \in \Omega$ with $y \neq z$,

$$\lim_{x \to y} (x - y) \frac{G_x(x, y)}{G_x(x, z)} = -\frac{1}{2G_x(y, z)}$$

Proof. It follows from (4) that

$$G_x(x,y) = -\frac{1}{2(x-y)} + H_x(x,y).$$
(6)

Therefore,

$$\lim_{x \to y} (x - y) G_x(x, y) = -\frac{1}{2},$$

 $G_{x\bar{y}}(x,x)\in\mathbb{R}.$

because *H* is twice continuously differentiable and thus H_x is finite. The Lemma follows.

Lemma 4. For any $x \in \Omega$,

Proof. It follows from (6) that

$$G_{x\bar{y}}(x,y) = H_{x\bar{y}}(x,y),$$
 (7)

and since *H* is twice continuously differentiable, $G_{x\bar{y}}(x, x)$ exists, even though G(x, x) does not. By the symmetry of *G*, we further have

$$\overline{G_{x\bar{y}}(x,x)} = G_{\bar{x}y}(x,x) = G_{x\bar{y}}(x,x),$$

implying the Lemma.

4 The χ^2 -divergence distance

Let us now focus on the χ^2 -*divergence distance*, which is given by the function $f(t) = t^2 - 1$, and denote it by $D(x, y) = D_f(x, y)$. Note that D is infinite on the boundary $\partial \Omega$, but this does not pose a problem. We first express D in terms of the derivatives of G and H, mentioned in (4).

Theorem 2. The χ^2 -divergence distance can be written as

$$D(x,y) = 2\frac{H_x(x,x)}{G_x(x,y)} + \frac{1}{2}\frac{G_{xx}(x,y)}{G_x(x,y)^2} - 1.$$
(8)

Proof. By Lemma 1,

$$D(x,y) = \frac{i}{\pi} \oint_{\partial \Omega} \frac{G_x(w,x)^2}{G_x(w,y)} dw - \frac{i}{\pi} \oint_{\partial \Omega} G_x(w,y) dw.$$

According to Observation 3, both integrands are holomorphic. For the second integrand, which has a first-order pole in Ω at w = y, applying Cauchy's Residue Theorem [1] and using Lemma 3, we get

$$\oint_{\partial\Omega} G_x(w,y)dw = 2\pi i \lim_{w\to y} (w-y)G_x(w,y) = -\pi i.$$



Figure 1: (Left) Gradient-descent paths of the χ^2 divergence distance from various source points (green) to a common target point (red). The domain is colour-coded according to the distance. The red curves are contours (level sets) of the distance function. (Right) The integral path γ is a sub-path of the gradient-descent path of the Green's function from target *y* to *x*, continuing further to the boundary point *w*, possibly passing through a saddle point *s*. The gradient of Green's function is visualized as a unit-length direction field, while the magnitude of the gradient is the colour-coded background.

The first integrand has a second-order pole at w = x, and so by Cauchy's Residue Theorem and Lemma 3,

$$\begin{split} \oint_{\partial\Omega} \frac{G_x(w,x)^2}{G_x(w,y)} dw &= 2\pi i \lim_{w \to x} \frac{d}{dw} \left((w-x)^2 \frac{G_x(w,x)^2}{G_x(w,y)} \right) \\ &= 2\pi i \lim_{w \to x} \left(2(w-x)G_x(w,x) \frac{G_x(w,x) + (w-x)G_{xx}(w,x)}{G_x(w,y)} - \left[(w-x)G_x(w,x) \right]^2 \frac{G_{xx}(w,y)}{G_x(w,y)^2} \right) \\ &= -2\pi i \left(\frac{\lim_{w \to x} \left(G_x(w,x) + (w-x)G_{xx}(w,x) \right)}{G_x(x,y)} + \frac{G_{xx}(x,y)}{4G_x(x,y)^2} \right). \end{split}$$

Since it follows from (4) that

$$G_x(x,y) = -\frac{1}{2(x-y)} + H_x(x,y)$$

and

$$G_{xx}(x, y) = \frac{1}{2(x-y)^2} + H_{xx}(x, y),$$

we have

$$\lim_{w \to x} (G_x(w, x) + (w - x)G_{xx}(w, x)) = \lim_{w \to x} (H_x(w, x) + (w - x)H_{xx}(w, x)) = H_x(x, x)$$

and the Theorem follows.

Theorem 3. The gradient of the χ^2 -divergence distance to y is proportional to the negative gradient of the Green's function with pole y, that is,

$$\nabla_x D(x, y) \propto -\nabla_x G(x, y)$$

Proof. Since *G* is harmonic, we have

$$\nabla_x G_x(x, y) = 2G_{x\bar{x}}(x, y) = 0, \quad \text{and} \quad \nabla_x G_{xx}(x, y) = \nabla_{\bar{x}} G_{x\bar{x}}(x, y) = 0$$

The harmonicity of H, together with (7), implies

$$\nabla_{x}H_{x}(x,x) = 2\frac{\partial}{\partial \bar{x}}H_{x}(x,x) = 2H_{x\bar{x}}(x,x) + 2H_{x\bar{y}}(x,x) = 2H_{x\bar{y}}(x,x) = 2G_{x\bar{y}}(x,x).$$

We then conclude from Theorem 2 that

$$\begin{aligned} \nabla_x D(x,y) &= 2 \frac{\nabla_x H_x(x,x)}{G_x(x,y)} - 2 \frac{H_x(x,x) \nabla_x G_x(x,y)}{G_x(x,y)^2} + \frac{1}{2} \frac{\nabla_x G_{xx}(x,y)}{G_x(x,y)^2} - \frac{G_{xx}(x,y) \nabla_x G_x(x,y)}{G_x(x,y)^3} \\ &= 4 \frac{G_{x\bar{y}}(x,x)}{G_x(x,y)} = 2 \frac{G_{x\bar{y}}(x,x)}{|G_x(x,y)^2|} \nabla_x G(x,y) \end{aligned}$$

By Lemma 4, $G_{x\bar{y}}(x, x)$ is real, so it remains to show that the ratio

$$\phi(x) = \phi(x, y) = \frac{\nabla_x G(x, y)}{\nabla_x D(x, y)}$$

is also negative for any $x \in \Omega$ with $x \neq y$.

To keep the notation simple, let us omit the subscript x and tacitly assume that all gradients are taken with respect to x. Following this convention, we then have by the harmonicity of G that

$$0 = \nabla^2 G(x, y) = \nabla \cdot (\phi(x) \nabla D(x, y))$$

= $\nabla \phi(x) \cdot \nabla D(x, y) + \phi(x) \nabla^2 D(x, y)$
= $\frac{\nabla \phi(x)}{\phi(x)} \cdot \nabla G(x, y) + \phi(x) \nabla^2 D(x, y),$

implying

$$\nabla \phi(x) \cdot \nabla G(x, y) = -\phi(x)^2 \nabla^2 D(x, y) \le 0, \tag{9}$$

because *D* is subharmonic. Let us now consider the unique gradient-descent path, defined by ∇G , from *y* to *x* and further on to some boundary point $w \in \partial \Omega$, and denote by γ the reverse sub-path from *w* to *x* with vector differential dz. By definition of γ , we have $dz/|dz| = \nabla G(z, y)/|\nabla G(z, y)|$, and according to the Gradient Theorem, we then have

$$\phi(x) - \phi(w) = \int_{\gamma} \nabla \phi(z) \cdot dz = \int_{\gamma} \frac{\nabla \phi(z) \cdot \nabla G(z, y)}{|\nabla G(z, y)|} |dz| \le 0,$$
(10)

where the last inequality follows from (9).

As $\nabla G(w, y)$ is normal to the boundary of Ω at w, so is $\nabla D(w, y)$ by Theorem 3. But while $\nabla G(w, y)$ points *inwards*, $\nabla D(w, y)$ points *outwards*, because D is subharmonic and hence obtains its maximum on $\partial \Omega$. Consequently, $\phi(w) < 0$, and it follows from (10) that

$$\phi(x) \le \phi(w) < 0.$$

Note that for multiply-connected domains, it may happen that the gradient-descent path of *G* from *y* to *w* encounters a saddle point *s* of *G*. In this case, even though $\nabla G(s) = 0$, there still exists a direction of steepest descent, given by the principal curvature directions (the eigenvectors of the Hessian of *G*) at that point, and we follow this direction to extend the path beyond *s*, see the example in Fig. 1 (Right). Since the number of saddle points is finite, this strategy guarantees that the path eventually terminates at some $w \in \partial \Omega$. Hence, the path γ from *w* to *x* in (10) is well-defined and consequently $\phi(x) < 0$, even in this case.

5 Conclusion

We have shown that the χ^2 -divergence distance of a multiply-connected domain behaves similarly to the Green's function of that domain, namely that it has no local minima, except at the target point, where it has a global minimum of zero. This implies that the divergence distance function may be used to trace paths between pairs of points in the domain by gradient descent. Our numerical experiments have indicated that all other divergence distances (based on other strictly convex f's) also have this desirable property, although we have been unable to prove it. The obstacle is applying Cauchy's Residue Theorem to compute the contour integral in the proof of Theorem 3. For many f's, the resulting integrand may have branches or higher-order poles, significantly complicating the computation of that integral.

References

- [1] L. V. Ahlfors. *Complex Analysis: An Introduction to the Theory of Analytic Functions of One Complex Variable*. International Series in Pure and Applied Mathematics. McGraw-Hill, New York, 3rd edition, 1979. ISBN 978-0-07-000657-7.
- [2] R. Chen, C. Gotsman, and K. Hormann. Efficient path generation with reduced coordinates. *Computer Graphics Forum*, 37(5):37–48, Aug. 2018. Proceedings of SGP. [PDF]
- [3] R. Chen, C. Gotsman, and K. Hormann. Path planning with divergence-based distance functions. *Computer Aided Geometric Design*, 66:52–74, Nov. 2018. [PDF]
- [4] I. Csiszár. Information-type measures of difference of probability distributions and indirect observations. *Studia* Scientiarum Mathematicarum Hungarica, 2:299–318, 1967.
- [5] J. B. Garnett and D. E. Marshall. *Harmonic Measure*, volume 2 of *New Mathematical Monographs*. Cambridge University Press, New York, 2005. ISBN 978-0-521-47018-6.
- [6] S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, Mar. 1951.
- [7] F. Liese and I. Vajda. On divergences and informations in statistics and information theory. *IEEE Transactions on Information Theory*, 52(10):4394–4412, Oct. 2006.