Computer Aided Geometric Design 66 (2018) 52–74

Contents lists available at ScienceDirect

Computer Aided Geometric Design

www.elsevier.com/locate/cagd

Path planning with divergence-based distance functions

Renjie Chen [a,*], Craig Gotsman [b], Kai Hormann [c]

[a] *Max Planck Institute for Informatics, Saarbrücken, Germany*
[b] *Jacobs Technion-Cornell Institute, Cornell Tech., New York, USA*
[c] *Università della Svizzera italiana, Lugano, Switzerland*

ARTICLE INFO

Article history:
Available online 7 September 2018

Keywords:
Path planning
Divergence
Potential function
Distance function

ABSTRACT

Distance functions between points in a domain can be used to automatically plan a gradient-descent path towards a given target point in the domain, avoiding obstacles that may be present. A key requirement from such distance functions is the absence of spurious local minima, and this has led to the common use of harmonic potential functions. This choice guarantees the absence of spurious minima, but is well known to be slow to numerically compute and prone to numerical precision issues. To alleviate the first of these problems, we propose a family of novel *divergence distances*. These are based on *f-divergence* of the *Poisson kernel* of the domain. Using the concept of *conformal invariance*, we show that divergence distances are *equivalent* to the harmonic potential function on simply-connected domains, namely generate paths which are *identical* to those generated by the potential function. We then discuss how to compute two special cases of divergence distances, one based on the *Kullback–Leibler*, the other on the *total variation* divergence, in practice by discretizing the domain with a triangle mesh and using Finite Elements (FEM) computation. We show that using divergence distances instead of the potential function and other distances has a significant computational advantage, as, following a pre-processing stage, they may be computed online in a multi-query scenario up to an order of magnitude faster than the others when taking advantage of certain sparsity properties of the Poisson kernel. Furthermore, the computation is "embarrassingly parallel", so may be implemented on a GPU with up to three orders of magnitude speedup.

1. Introduction

Path planning in a planar domain containing obstacles is an important problem in robotic navigation. The main challenge is for an autonomous agent to move from one point (the *source*) in the domain to another (the *target*) along a realistic path which avoids the obstacles, where the path is determined automatically and efficiently based only on knowledge of the domain and local information related to the current position of the agent. This important problem has attracted much attention in the robotics community and is the topic of ongoing research, some of the most important techniques being the classical Dijkstra algorithm (Dijkstra, 1959), the A^* and D^* search algorithms, configuration space sampling algorithms and potential functions. The interested reader is referred to the recent survey by Goerzen et al. (2010) and Souissi et al. (2013) for more details. The family of path planning algorithms most relevant to our work is that based on so-called *potential functions*, inspired by the physics of electrical force fields, first proposed in the late 1980s by Khatib (1986) and developed

* Corresponding author.
E-mail addresses: renjie.c@gmail.com (R. Chen), gotsman@cs.technion.ac.il (C. Gotsman), kai.hormann@usi.ch (K. Hormann).

https://doi.org/10.1016/j.cagd.2018.09.002
0167-8396/© 2018 Published by Elsevier B.V.

by Kim and Khosla (1992), Rimon and Koditschek (1992), and Connolly and Grupen (1993) soon after. The idea is, given the target point, to construct a scalar function on the domain, such that a path to the target point from any other source point may be obtained by following the negative gradient of the function. While elegant, Koren and Borenstein (1991) have identified a number of significant pitfalls that these methods may encounter, the most important being the presence of so-called "trap" situations – the presence of local minima in the potential function. To avoid this, the scalar function must have a global minimum (typically zero-valued) at the target, and be void of local minima elsewhere in the domain. The presence of "spurious" local minima could be fatal, since the gradient vanishes and the agent becomes "stuck" there. Other critical points, such as saddles, are undesirable but not fatal, since a negative gradient can still be detected by "probing" around the point.

Designing and computing potential functions for planar domains containing obstacles has been a research topic for decades. Perhaps the most elegant type of potential function is the harmonic function (Kim and Khosla, 1992; Connolly and Grupen, 1993), which has very useful mathematical properties, most notably the guaranteed absence of spurious local minima. Alas, the main problems preventing widespread use of these types of potential functions are the high complexity of computing the function, essentially the solution of a very large system of linear equations, and the fact that very high precision numerical methods are required, as the functions are almost constant, especially in regions distant from the target. This paper addresses the first of these issues. We describe a family of new functions, which, while quite distinct from the harmonic potential function, generate *exactly* the same gradient-descent paths. However, they do this at a tiny fraction of the computational cost.

In practical path planning scenarios, the planar domain is described by a set of polygons representing the domain boundary and the obstacles within, which can be quite complicated. A good potential function should be "shape-aware", in the sense that it should produce paths which naturally circumvent the obstacles. The agent is armed with an automatic algorithm relying on auxiliary data structures which, given its current position in the domain, can efficiently compute the direction in which it should proceed towards the target. As we shall see later, there is a tradeoff between space and time complexity in achieving this goal.

Although the classical term is "potential function", in this paper we use the more generic term "distance function" for the guiding scalar function. We believe this is more appropriate, as in a sense, the function measures a scalar distance value between the source and the target, which the path-planner tries to decrease as it advances towards the target. Although not identical to the classical shortest-path distance (also known as "geodesic" distance), this distance also takes into account the geometry of the domain and the obstacles.

The rest of this paper is organized as follows. We start with a mathematical analysis of a number of distance functions on continuous planar domains: In Section 2 we consider distance functions used for path-planning which are common in the literature, based on various forms of the Laplace operator, including the classical harmonic potential function. In Section 3 we introduce our new family of divergence distance functions, and show (in Appendix A) that they are all equivalent to the Green's function and the Poincaré metric on the disk. In Section 4 we direct our attention to the more practical case of a discretized planar domain and provide explicit algebraic expressions and computation methods for the distance functions. There we show how divergence distance functions may be computed much faster than any of the traditional distances. In Section 5 we provide more experimental results and insights. We conclude in Section 6 with a summary and open questions.

2. Laplacian-based distance functions

The Green's function

Classical potential functions are based on harmonic functions, which satisfy the second-order linear differential Laplace equation

$$\nabla^2 f = 0 \tag{1}$$

where $\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$ is the Laplace operator, also called the *Laplacian*. These are particularly attractive since a harmonic function satisfies a "minimum/maximum principle" – it obtains its minimum and maximum on the domain boundary, implying that the domain interior contains no local extrema. Beyond this, harmonic functions have many other "nice" properties and have been studied extensively for decades. Rather than providing a detailed exposition here, we refer the interested reader to the book by Axler et al. (2001) and the related text by Garnett and Marshall (2005), which contain a wealth of information, including all the classical results we use here. Denote by Ω an open bounded planar domain, by $\partial\Omega$ its boundary, by q the source point, and by p the target point. Note that the mathematical translation of "obstacles" in the domain, is to "holes" in Ω. If there are no obstacles, Ω is simply connected and has a single exterior boundary loop. If there are obstacles, Ω is multiply connected having a single exterior boundary loop and multiple interior boundary loops. For convenience, we identify the plane $\mathbb{R}^2$ with the complex field $\mathbb{C}$, and much of our notation and formula will use complex number algebra. For example, a point $(x, y) \in \mathbb{R}^2$ is identified with the point $z = x + iy \in \mathbb{C}$, its conjugate is $\bar{z} = x - iy$, its