UCLA
COMPUTATIONAL AND APPLIED MATHEMATICS

Diffusion Generated Motion by Mean Curvature

Barry Merriman
James Bence
Stanley Osher

April 1992

CAM Report 92-18

Department of Mathematics
University of California, Los Angeles
Los Angeles, CA. 90024-1555

Diffusion Generated Motion by Mean Curvature

BARRY MERRIMAN, JAMES BENCE, AND STANLEY OSHER

1. Introduction

Motion by mean curvature is the art of moving a planar curve normal to itself with velocity equal to the curvature. More generally, the velocity might be some function of curvature, or we might want to move a surface in three dimensions.

There are several well developed numerical methods for curvature dependent motion, but all have weaknesses. To understand what is lacking, and to set the stage for our new approach, we briefly consider the strengths and weaknesses of common methods: front tracking, reaction diffusion, and level set.

In front tracking methods, the curve is replaced by a string of discrete points, the curvature is computed by finite differences, and the points move in time with the computed velocity. This has the advantages of high accuracy, computational efficiency, and generalizes well to arbitrary motion laws. The main disadvantage is the complicated algorithm needed to track fronts that merge, break, or otherwise change topology. In three dimensions these complications increase severely.

In reaction diffusion methods, the curve is represented as a sharp front separating two stable equilibrium states of a strong reaction, and a small diffusion is added, which results in motion by mean curvature of the front. This has the advantage of requiring no special treatment for topological changes in the curve, and the algorithm is simple, since all that is required is finite differences to solve the reaction diffusion equation on a uniform grid. It also generalizes easily to three dimensions. The main disadvantage is that the reaction front width is much smaller than the scales of interest in the problem—the length scales of the curve itself—yet this front width must be well resolved on the grid in order to compute the correct time evolution. In fact we have rigorously proven elsewhere that if the front is not well resolved then incorrect speeds will

The first author was supported in part by an NSF Postdoctoral Fellowship in the Mathematical Sciences and the remaining authors were supported in part by Grant #NSF DMS91-03104 and Department of Commerce Grant #DOC GONANBID1125.

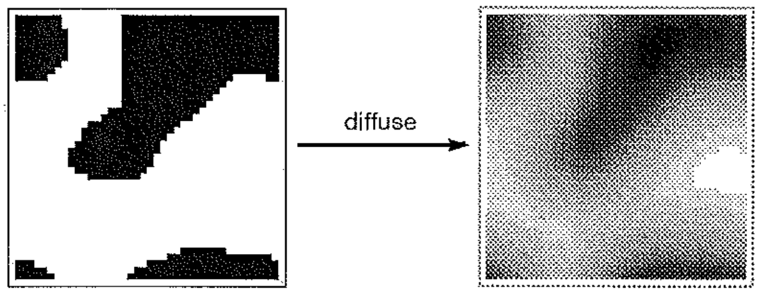

FIGURE 1. Diffusion of a set.

be observed on the computer. Thus the method is computationally inefficient due to the need for an artificially fine grid. Adaptive gridding can greatly improve efficiency, but at the cost of complicating the algorithm. Also, this method has difficulty generalizing to handle arbitrary curves that have self intersections, triple junctions, or other juncture points.

In level set methods the curve is represented as a level set of a function on the plane, and this function in turn evolves according to a partial differential equation of Hamilton-Jacobi type, chosen precisely to produce the proper level set motion. This has the advantages of requiring no special treatment of topological changes and a simple algorithm, since only finite differences on a uniform grid are needed solve the Hamilton-Jacobi equation. Further, the grid is not artificially fine as in the reaction diffusion methods, so it is reasonably efficient. It is sufficiently flexible to handle arbitrary motion laws, and extends easily to three dimensions. The main disadvantage is that the method does not extend to handle curves with self intersections, triple junctions or other junctions.

From this brief overview, we see there is no method that has a simple algorithm, accommodates curves with complicated structure (triple junctions, intersections, topological changes, higher dimensions), and is computationally efficient. Our effort to find such a method lead us to the present work.

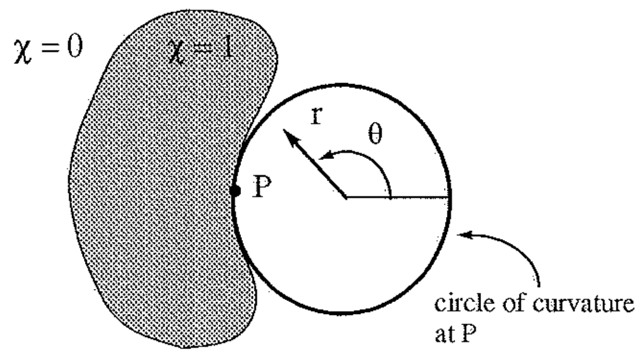

FIGURE 2. Analysis of diffusion of the characteristic function.

2. Intuition

Before describing the new method in detail, let us present the intuition that underlies its discovery.

Imagine that a set of points in the plane is allowed to "diffuse", as shown in figure 1. The diffusion rapidly blunts the sharp points on the boundary, but has little impact on the flatter parts. It seems that diffusion creates some sort of curvature dependent boundary motion—assuming we can give meaning to a precise boundary within the fuzzy boundary of the diffused set. This suggests there will be an algorithm in which we repeatedly diffuse a set, "recover" a new set, and thereby generate curvature dependent motion of the set boundary.

3. Analysis of Diffusion Generated Motion

We apply a formal analysis to clarify the intuition from the previous section. Our goal is to determine how the set boundary evolves when we diffuse the set. By "diffuse the set", we mean apply diffusion to the characteristic function of the set, χ. So we formally analyze the diffusion equation:

$$\frac{\partial \chi}{\partial t} = D \nabla^2 \chi.$$

Consider a point, P, of interest on the set boundary (see figure 2), and construct a local polar coordinate system with origin at the center of curvature of P. Write out the diffusion equation in these (r, θ) coordinates:

$$\frac{\partial \chi}{\partial t} = \frac{D}{r}\frac{\partial \chi}{\partial r} + D\frac{\partial^2 \chi}{\partial^2 r} + \frac{D}{r^2}\frac{\partial^2 \chi}{\partial^2 \theta}.$$

Because of the local circular symmetry, we have $\frac{\partial \chi}{\partial \theta} = \frac{\partial^2 \chi}{\partial^2 \theta} = 0$ near P, and the equation reduces to

$$\frac{\partial \chi}{\partial t} = \frac{D}{r}\frac{\partial \chi}{\partial r} + D\frac{\partial^2 \chi}{\partial^2 r}.$$

This is an *advection*-diffusion equation in the radial direction, with advective velocity $V = D/r$. At P, $r = \rho$, the radius of curvature, and so $V = D\kappa$ there, where $\kappa = 1/\rho$ is the curvature. Thus the total effect of diffusing the set is to initially move the boundary radially at speed $D\kappa$, while simultaneously diffusing χ radially about this position, as indicated in figure 3. Note that the diffusive contribution does not affect the motion of the $\chi = 1/2$ level set—this level moves *exactly* with the advective velocity. Thus the diffusion generates motion by mean curvature of the boundary, as long as we identify this boundary with the $\chi = 1/2$ level of the diffusing characteristic function. Before the local analysis breaks down, we need to redefine a new set, and repeat the diffusive process to generate further motion. The recovered set is the one with this level as boundary, namely $\{\chi \geq 1/2\}$.

If instead we use a variable diffusivity diffusion equation,

$$\frac{\partial \chi}{\partial t} = \nabla D \nabla \chi,$$

the same formal analysis yields the same result. By allowing the diffusivity to have the functional form $D = D(x, \chi, \nabla\chi, \nabla\nabla\chi)$, we can formally obtain a wide class of curvature dependent motion, anisotropic motion, and other motion laws.

4. Numerical Method

The formal analysis suggests numerical methods for evolving an initial curve σ can be based on the following algorithm:

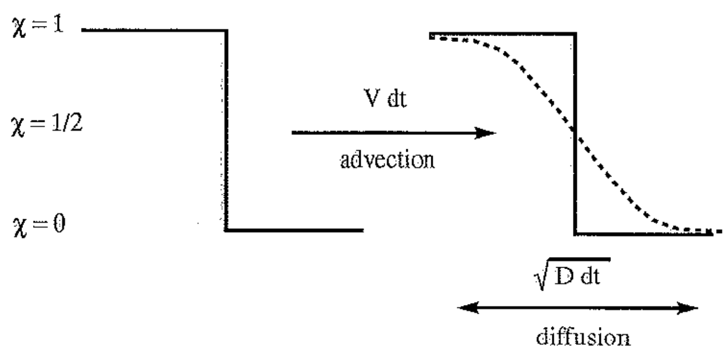

FIGURE 3. Profile view of the initial motion.

Given: a set with boundary σ, characteristic function χ;
For: χ characteristic function of a set;
construct $\chi(\tau) =$ diffusion of χ for time τ;
define χ as characteristic function of the set $\{\chi(\tau) \geq 1/2\}$;
define σ as the boundary of this set;
Repeat;

The only parameter in the method is the time step, τ. The basic requirement is that it be short enough so that the local analysis of the previous section is valid, but also long enough so that the boundary curve moves by at least one grid cell on the spatial grid (otherwise the curve would be stuck). Since the local analysis breaks down when diffusive information travels on the order of the local radius of curvature—at that point the local analysis is polluted by external information—the first requirement is $\sqrt{D\tau} \ll \rho$. The second is $D\kappa\tau \gg dx$. These combine to give an upper and lower bound on τ,

$$(\frac{\rho}{\delta x})^2 \gg \frac{\tau D}{\delta x^2} \gg \frac{\rho}{\delta x}.$$

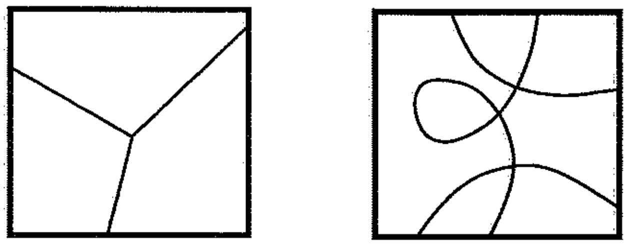

FIGURE 4. General curves of interest.

As long as $\rho/\delta x$ is large—i.e., as long as the grid resolves the radius of curvature—there are τ in this allowable range.

5. Extension to Arbitrary Curves

The basic method we have developed applies only to curves that bound a set. This does not include curves with self-intersection, triple junctions, or other junctions, such as in figure 4. To be of practical interest, we need to generalize the method to these "multiple region" cases. The minimum demands we place on a generalization are that it reduce to the previous case away from junctions and intersections, and that it provide some simple way to allow for different possible stable triple junctions (i.e. arbitrary angles in a stable triple junction).

We can achieve the generalization simply by symmetrizing the algorithm we already have. That is, the case we treated was really a two region case: the curve of interest divides the plane into two regions, and we could choose χ to be the characteristic function of either region to represent the curve. This asymmetry suggests the algorithm should be recast to use both these characteristic functions in a symmetric way.

To obtain this symmetry, suppose the curve divides the plane into 2 regions, with characteristic χ_1 and χ_2. As per the algorithm , we could move the curve by diffusing either χ_i. The updated sets used in the algorithm would then be either $\{\chi_1 \geq 1/2\}$ or $\{\chi_2 \geq 1/2\}$. These can be symmetrically characterized as $\{\chi_1 > \chi_2\}$ and $\{\chi_2 > \chi_1\}$. Thus, working with *both* sets, we can diffuse both independently, and recover new sets $i = 1, 2$ as the sets on which $\chi_i, i = 1, 2$ is biggest. In this way, the algorithm can be rewritten using both sets symmetrically, instead of one or the other.

This version of the algorithm extends immediately to the case of moving a curve σ that divides the plane into N regions:

Given: N sets that partition the plane, with the union of the boundaries equal to curve σ, and characteristic functions $\chi_1, \ldots, \chi_N$;
For: $\chi_1, \ldots, \chi_N$ characteristic functions of sets that partition the plane;
construct $\chi_i(\tau)$ = diffusion of χ_i for time τ, $i = 1, \ldots, N$;
define χ_i as characteristic function of the set $\{\chi_i(\tau) \geq \chi_j(\tau), j = 1, \ldots, N\}$;
define σ as the union of the boundaries of the sets represented by $\chi_i, i = 1, \ldots, N$;
Repeat;

This generalization does reduce to the original algorithm away from any junctions or intersections. As for allowing choice of stable triple junction, note we have presented here the maximally symmetric generalization. This, not surprisingly, results in a maximally symmetric stable triple junction: all angles equal. However, other stable triple points can be enforced by altering the procedure for recovering sets from fuzzy sets.

To understand this, its best to visualize the general set reconstruction process as follows: note that because $\chi_i \geq 0$ and $\sum_i^N \chi_i = 1$ initially, and because the diffusion evolution operator is linear with $\chi \equiv 1$ as a steady solution, these properties hold for all time. Thus, for any given point x in the plane, its "state vector" $(\chi_1(x), \ldots, \chi_N(x)) \in R^N$ stays on the unit tetrahedron $\sum_i^N \chi_i = 1, 1 \geq \chi_i \geq 0$ (see figure 5 for the $N = 3$).When the χ correspond to sets, this state vector is at a vertex, since all entries are 0, except for one, which is 1. During the diffusion, the state moves into the interior of the tetrahedron, since generally all components are positive. To recover actual sets—i.e. determine which

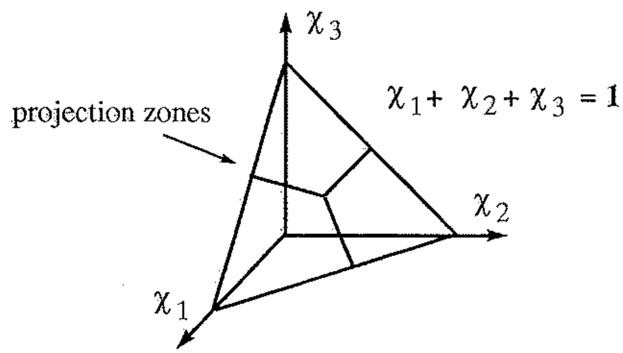

FIGURE 5. The tetrahedron in state space.

set contains x—we need to project back onto a vertex. We have chosen here the maximally symmetric projection, namely moving the point back to the closest vertex. If, instead, we choose some other way to push the state back to a vertex, we find empirically that this results in a different stable triple junction. The precise relation between the choice of "projection zones" on the tetrahedron in state space and the resulting stable triple junction in real space is not yet clear, but it seems to be a simple one. It appears that all stable triple junctions are obtainable by a suitable choice of projection.

When starting from an initial triple junction that is not the stable one, we find empirically that it rapidly evolves to the stable form, as is desired.

6. Generalization to Three Dimensions

The algorithms generalize directly to higher dimensions: that is, diffusion generates motion by mean curvature of set boundary surfaces in three dimensions. There is no need to make any modifications to the algorithm—simply solve the three dimensional diffusion equation.

7. A Generalized Huygens Principle

Diffusion generated motion by mean curvature has a geometric interpretation as a generalized Huygens principle. To see this, we will first convert the standard geometric Huygens principle to an algebraic form, compare this to the algebraic form of diffusion generated motion, and from this comparison deduce its geometric expression.

The standard Huygens principle is used to move a curve with a constant normal velocity (see figure 6). We proceed by drawing identical little circles centered so that they are entirely on one side of the curve and tangent to it, and take the locus of these circle centers as the new curve position. If we represent the curve as the boundary of a set with characteristic function χ, and convolve this with a circular kernel K (the characteristic function of a disk will do),$\chi * K$, the 0 level set of this smoothed out χ is the same updated curve as in the geometric statement. The corresponding updated set is $\{\chi * K > 0\}$. Thus the standard geometric Huygens principle for motion with constant velocity is equivalent to the algebraic procedure of convolving with a kernel and advancing to a 0 level set of the result.

Since diffusion generated motion advances the set as $\{\chi * K > 1/2\}$, (here K is the diffusion kernel), this suggests the geometric formulation would consist of drawing identical little circles *centered so that the curve cuts them in half* (by area), and taking the locus of these circle centers as the new curve. A simple geometric analysis shows this principle does indeed advance the boundary with velocity proportional to curvature. This also shows that K could be taken to be any circularly symmetric kernel, rather than just the diffusion equation Gaussian kernel. (However, the diffusion kernel is convenient to compute with because the convolution can be accomplished indirectly by solving the finite difference equations.)

The algebraic point of view clearly shows there is a continuum of Huygens principles, defined by advancing the set as $\{\chi * K > \alpha\}$. $\alpha = 0$ corresponds to the standard Huygens principle and motion with constant velocity, and $\alpha = 1/2$ corresponds to the new Huygens principle and motion by mean curvature. Geometrically, α represents the fraction of the circle area that lies on one side of the curve.

Note that all circularly symmetric, positive smoothing kernels will lead to motion by mean curvature if the set is advanced as $\{\chi * K > 1/2\}$. By choosing nonsymmetric kernels, we can get anisotropic motions. The corresponding geometric statement is that if we draw some other shape besides little circles, we get a Huygens principle for anisotropic, curvature dependent motion.

By generalizing Huygens principle, we have achieved a symmetry between the algebraic and geometric points of view: we now have a general class of geometric principles based on drawing certain little shapes and taking a locus of points as the new curve, and algebraic principles based on convolving characteristic functions with kernels and using a level set to update the set boundary.

This algebraic-geometric duality extends to the three dimensional case as well.

8. Experimental Status

The results from several experiments can be seen in the video that accompanies the paper.

We have used diffusion generated motion to compute the evolution of a circle collapsing by mean curvature—a problem where the analytic solution is known—and verified that it accurately computes the motion for this benchmark problem.

We have also done a variety of experiments on simple closed curves, triple junctions and curves that have self intersections, and found the algorithm is robust, behaves reasonably, and converges as the grid is refined.

In general, the experimental results are promising, and the method is quite simple to implement, since it essentially is just solving the diffusion equation. The initial experiments were computationally intensive, due to using explicit time stepping in the diffusion equation, and required workstation-hours to run on a 100×100 grid. We now use implicit time stepping, and the method runs in less than a workstation-minute.

We have not yet tried any experiments in three dimensions, or attempted more general curvature dependent laws of motion.

9. Summary

We have developed a new method for mean curvature dependent motion. It works in any number of dimensions and easily accommodates curves with triple junctions, arbitrary junctions and intersections. Computationally, the algorithm consists almost entirely of solving the diffusion equation on a domain containing the curve of interest, using a grid fine enough to resolve the length scales of interest. Thus the algorithm is extremely simple to implement, and reasonably efficient.

The method is based on the observation that diffusion of a set initially moves its boundary with a velocity proportional to mean curvature. This can be understood algebraically based on a simple analysis of the diffusion equation. It can also can be understood geometrically in terms of a generalized Huygens principle.

Experiments based on the method are promising, and there remain many things to test. We have developed a useful and intuitive formal theory to guide the development, but the rigorous theory is not yet in place.

Department of Mathematics, UCLA, Los Angeles, Calif. 90024
E-mail: barry@math.ucla.edu

Department of Mathematics, UCLA, Los Angeles, Calif. 90024
E-mail: jbence@math.ucla.edu

Department of Mathematics, UCLA, Los Angeles, Calif. 90024
E-mail: sjo@math.ucla.edu