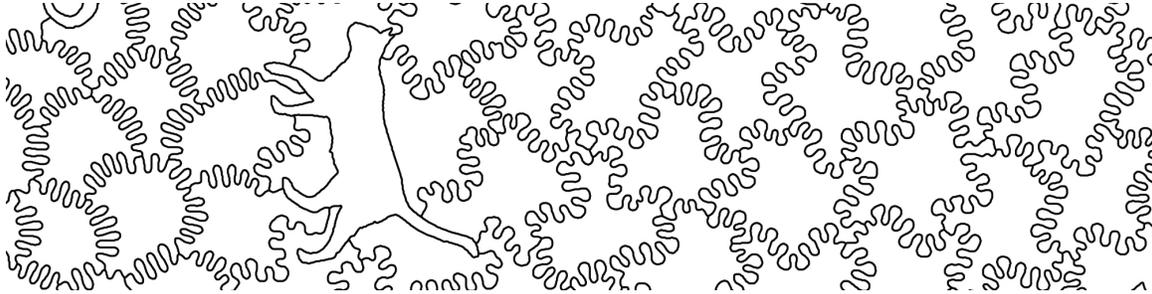


Multiphase Numerical Modeling of Dendritic Solidification for Jigsaw Puzzle Generation

Jesse Louis-Rosenberg *
Nervous System

Alec Resnick †
sprout & co.

Jessica Rosenkrantz ‡
Nervous System



Abstract

We explore an opportunity to expand the stylistic variety of handmade and artisanal jigsaw puzzles through the application of techniques from natural simulation. Typical jigsaw puzzle designs reflect the manufacturing constraints of die-cut, mass-production methods. We generate novel families of puzzle forms by applying a phase field approach to the simulation of dendritic solidification. We extend existing models of solidification with multiphase methods to satisfy aesthetic and geometric considerations specific to jigsaw puzzle design. We present examples of the resulting puzzle forms and discuss possible extensions and refinements.

CR Categories: Computing Methodologies [SIMULATION AND MODELING]: Applications;

Keywords: jigsaw puzzle, phase field methods, dendritic solidification

Links: [DL](#) [PDF](#) [WEB](#) [CODE](#)

1 Introduction

We present a system for generating jigsaw puzzle forms based on a numerical simulation of dendritic solidification using phase field methods. Many software packages exist which generate jigsaw puzzles [Google 2012]. Typically, these systems produce forms resembling mass-produced, die-cut puzzles which have pieces laid out on a rectangular grid with simple, rounded nodules on their boundaries. However, there is a tradition of handmade, artisanal jigsaw

puzzles offering a much richer aesthetic variety [Williams 2004]. Each artist has their own cut style, often incorporating recognizable figures called whimsies. The methods we present here explore the opportunity for for artisanal algorithms as applied to jigsaw puzzles.

The branching forms exhibited by Laplacian growth are a natural fit for the formal requirements of jigsaw puzzles’ intricate, interlocking pieces. Our system implements a phase-field model of dendritic crystal solidification to evolve the boundaries of an initial set of pieces into interlocking shapes. This system enables us to create one-of-a-kind jigsaw puzzles with their own, distinctive cut style.

2 Background

“Laplacian growth” refers to the evolution of a plane interface between two regions driven by the gradient of a harmonic field. These processes typically produce fractal, branching structures which have been studied in fields ranging from pure math [Mineev-Weinstein and Zabrodin 2001], microbial biology [Ben-Jacob et al. 1992], and urban planning [BattyIT 1991], to physics [Krichever et al. 2004][Hastings and Levitov 1998], and computer graphics [Kim et al. 2007].

Some of these systems exhibit discrete growth—like diffusion limited aggregation [Barra et al. 2001]—and can act deterministically or stochastically, as with dielectric breakdown models [Pietronero et al. 1988]. Alternatively, the growth can be continuously defined by a velocity field, as in the viscous fingering of Hele-Shaw cells [Løvoll et al. 2004] or dendritic solidification [Family et al. 1987].

The aesthetic of and extensive literature surrounding dendritic solidification drew us to the phenomena in designing puzzle pieces. Dendritic solidification involves the growth of crystals in a super-cooled environment, a mechanism underlying phenomena including the formation of snowflakes and the engineered microstructure of metal alloys. The simulation of such phenomena has been extensively studied, leading to numerous numerical methods: boundary integral methods [Strain 1989], variational models [Almgren 1993], finite element methods [Zhao and Heinrich 2001], level set methods [Sethian and Strain 1992] and phase-field methods [González-Cinca et al. 2003].

Phase-field methods in particular have been widely applied to the study of dendritic solidification in both 2D [Kupferman et al. 1994] and 3D [Karma and Rappel 1998]. These methods avoid the algo-

* jesse@n-e-r-v-o-u-s.com

† alec@thesprouts.org

‡ jessica@n-e-r-v-o-u-s.com

rhythmic complexity of alternatives which require the explicit representation of a moving, complex, geometric phase boundary. Adaptive grid methods have been used to increase the efficiency of these methods [Provatas et al. 1999]. Phase-field methods have also been used to extend this phenomena, *e.g.* implementing dendritic solidification with fluid flow [Jeong et al. 2001].

3 Methods

Our algorithm is based on Kobayashi’s application of phase-field models to dendritic solidification [Kobayashi 1993] [Kobayashi 1994]. Phase-field models avoid explicitly representing boundaries by replacing a discontinuous phase transition with a smooth transition in an order parameter $p \in [0, 1]$ over a boundary of thickness ϵ from 0 (solid) to 1 (liquid). The numerical results of this model converge to those of the sharp interface representation as $\epsilon \rightarrow 0$. Kobayashi’s model is well-suited to graphical applications because of its incorporation of heuristic parameters which allow one to trade physical correctness for flexibility.

3.1 Phase-field Model

The Allen-Cahn equation (Equation 1) represents the dynamics of a non-conserved order parameter [Allen and Cahn 1979]. The dynamics of the phase field model are governed by an Allen Cahn equation with a Ginzburg-Landau type free energy, which we numerically implement with a finite difference, forward Euler method on a regular grid.

$$\Phi = \int \frac{1}{2} \epsilon^2 |\nabla p|^2 + F(p, m) \quad (1)$$

$$\tau \frac{\partial p}{\partial t} = \frac{\delta \Phi}{\delta p} \quad (2)$$

where ϵ is the diffuse boundary’s thickness, F is the potential driving our system, and τ is a small, positive constant. F is not physically motivated; it is a phenomenological representation of the desired dynamics of the system. We choose a double-well potential parametrized by the phase p and an independent variable m :

$$F(p, m) = \frac{1}{4} p^4 - \left(\frac{1}{2} - \frac{1}{3} m \right) p^3 + \left(\frac{1}{4} - \frac{1}{2} m \right) p^2 \quad (3)$$

$$F_l(p, m) = p(1-p)(p - \frac{1}{2} + m) \quad (4)$$

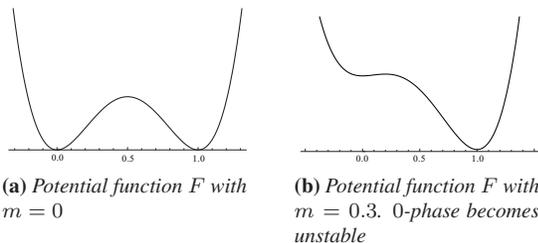


Figure 1: Exemplary dynamics of the double-well potential created by Equation 4 for $F(p, m)$.

As shown in Figure 1, this potential drives the phase to either 0 or 1 when $m = 0$. However, as $m \rightarrow \frac{1}{2}$, the 0-phase becomes unstable, and as $m \rightarrow -1/2$, the 1-phase becomes unstable. The dynamics

of supercooling are modeled by Equation 5, where we introduce a temperature field governing $m(T)$.

$$m(T) = \frac{\alpha}{\pi} \arctan \gamma(T - T_e) \quad (5)$$

Equation 5 uses \arctan to limit $|m| \leq \frac{1}{2}$. T_e is an “equilibrium temperature” at which both phases are stable. For computational convenience, we take $T_e = 0$. γ is a scaling parameter for the temperature, and α limits the instability produced by out-of-equilibrium temperatures.

The temperature dynamics are governed by two elements: diffusion and conservation of enthalpy (Equation 6).

$$\frac{\partial T}{\partial t} = \underbrace{\nabla^2 T}_{\text{diffusion}} + \underbrace{K \frac{\partial p}{\partial t}}_{\text{enthalpy}} \quad (6)$$

Phase changes release and absorb energy. In a closed system, enthalpy is conserved. This conservation is captured by T ’s dependence on the change in phase in Equation 6. K represents the latent heat; in our model, K is dimensionless, allowing us to set $K = 1$ by scaling our temperature appropriately.



Figure 2: As the boundary grows, heat is released. Because the heat diffuses faster than the phase, neighboring areas’ growth is suppressed. As a result, growth of the boundary segment into the supercooled portion is amplified.

The exo- and endothermic phase changes drive boundary instabilities through a local amplification, lateral inhibition process, shown in Figure 2. The boundary’s growth releases heat. The temperature diffuses faster than the phase, limiting the growth of neighboring areas of the boundary. Meanwhile, that portion of the boundary has grown toward the supercooled “open space,” which further encourages the growth of the boundary in that direction.

In the case of isotropic crystal growth, the phase field dynamics simplify to

$$\tau \frac{\partial p}{\partial t} = \epsilon^2 \nabla^2 p + p(1-p)(p - \frac{1}{2} + m) \quad (7)$$

Applying the phase field approach naively to the problem of generating jigsaw puzzles introduces a few problems. Dendritic solidification typically involves a single material in one phase encroaching on another phase of the *same* material, *e.g.* a single, solid piece of ice growing into an expanse of supercooled water (Figure 3). In this setup, one phase grows while another shrinks. We would prefer the simulation to be piece-agnostic and symmetric with respect neighbors’ phases. Rather than one piece acting in a “growing” role and another in a “shrinking” role, the boundary should evolve symmetrically, with identical dynamics in both directions. Otherwise, we introduce an aesthetic asymmetry in the quality of adjacent pieces’ boundaries (Figure 3).

To address the issue of piece symmetry, we treat one piece as supercooled and the other as superheated so both phases grow *into* one another (Figure 3). (Note this setup is physically meaningless—by putting a supercooled and a superheated phase next to each other

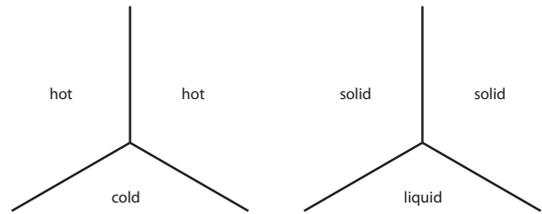


(a) Supercooled everywhere. One phase grows down into the other.

(b) Superheated on the top and supercooled on the bottom. Each phase grows symmetrically.

Figure 3: Demonstration of the aesthetic consequences of [a]symmetry in phase growth dynamics.

we are taking advantage of the free energy functional’s symmetry to create the boundary symmetry we require.)



(a) Impossible to arrange hot/cold with no identical neighbors.

(b) Impossible to arrange solid/liquid with no identical neighbors.

Figure 4: Two phases and temperatures are insufficient to preserve piece symmetry at the intersection of three pieces. Note that phase and temperature fields are *completely independent*.

When we attempt to extend this to three pieces (Figure 4) we discover there is no way to arrange three pieces sharing boundaries without at least one pair of neighbors sharing a boundary *and* a phase. While we could consider the problem of laying out regions to ensure no two neighbors are in the same phase (*e.g.* four regions can be arranged to ensure no two neighbors share a phase), in the general case this is impossible with only two phases. Instead, we assign each piece its own, independent phase. In order to maintain each piece’s superheated/supercooled temperature *relative to its surroundings*, each piece also requires its own, independent temperature. To do this, we need to extend our solidification model to a multiphase simulation. Rather than representing traditional phases of matter (*e.g.* solid, liquid, gas), we now have infinite phases representing a more abstract “state of matter.”

We apply Steinbach *et al.*’s approach to multiphase dynamics [Steinbach *et al.* 1996]. In the two phase example, one phase is represented implicitly, *e.g.* $p_{\text{liquid}} = 1 - p_{\text{solid}}$. In the multiphase approach, each phase is represented explicitly such that $\sum_i p_i = 1$. The one aspect of our dynamics requiring adjustment is the potential function. The most important characteristic of the boundary’s free energy is how it defines the relative stability of the 0- and 1-phases. Because we have numerous, arbitrary, overlapping phases, we need a functional which preserves the double-well behavior but which can handle many phases. We define such a pair-wise phase potential:

$$m_{ij} = \frac{m(T_i) + m(T_j)}{2} \quad (8)$$

$$F'(p_i, p_j) = p_i p_j (p_i - p_j + m_{ij}) \quad (9)$$

$$F'_i(p_i) = \sum_{j \neq i} F'_i(p_i, p_j) \quad (10)$$

Equation 10 expresses the potential of each phase as the sum of the pairwise potential of all phases.

4 Initialization

The multiphase field approach outlined here is very flexible—it only requires a set of boundaries, phase fields, and temperature fields. Many approaches to defining the initial piece boundaries are possible. The boundary definition step is completely separate from their subsequent evolution per the dynamics defined by our free energy functional (Equation 1).

In our system, we’ve included whimsies—well-defined, recognizable shapes artificially placed in the puzzle. These shapes do not change; we treat them as reflective boundary conditions. The remaining pieces’ initial boundaries are formed through a reaction diffusion-process, creating an approximate, generalized Voronoi diagram of a set of starting seed shapes.

In the simplest case, initial seeds are placed via dart throwing [Runions *et al.* 2005]. Points are randomly generated and discarded if they are within ρ of an existing seed. This generates a random set of points of approximately fixed density, a specific case of blue noise. The initial seeds can also be generated in other configurations, *e.g.* uniform grids, phyllotactic patterns, or arbitrary sets of lines and curves. Each initial seed i is assigned a unique chemical A_i which diffuses out and reacts with its neighbors. Each chemical activates itself at a rate μ and inhibits all other chemicals at a rate ν , governed by the dynamics of Equation 11:

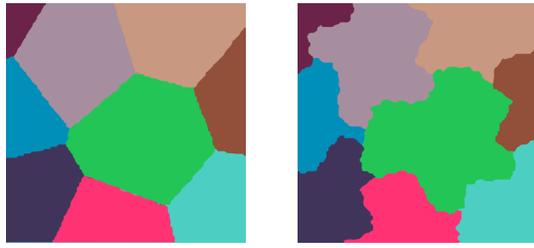
$$\frac{\partial A_i}{\partial t} = \nabla^2 A_i + \mu A_i - \nu \sum_{j \neq i} A_j \quad (11)$$

The process stops when a minimum concentration c_t exists everywhere. The boundaries of each chemical region at this time define the initial piece boundaries.

We then add two families of sinusoidal noise of wavelengths λ_i and U_i , $i \in 1, 2$ to these boundaries. In our system, $\lambda_2 \ll \lambda_1$, meaning that we are effectively adding gross- and fine-scale perturbations to initialize the simulation (Figure 5). Again, this process defines our beginning boundaries but is by no means the only process for initializing boundaries.

After defining the initial boundaries of the pieces, we randomly assign each piece a phase. We designed our framework assuming each piece is assigned a unique phase. Because our simulation’s computational complexity scales with the number of phases, we try to minimize the number of distinct phases we use.

Assigning phases so that no neighbors share a phase—functionally equivalent to assigning each piece a unique phase—is a specific case of a well-known problem in graph coloring. In 1976, Appel demonstrated that four colors is enough to color an arbitrary planar graph with no neighbors sharing a color [Appel and Haken 1989]. Finding that arrangement for four phases is computationally demanding. The “five phase problem” is much more tractable. To



(a) Boundaries defined by the concentration fields after the reaction diffusion process.

(b) Pieces with boundaries perturbed by sinusoidal noise at wavelengths λ_1 and λ_2 .

Figure 5: Initialization of piece boundaries with a combination of dart throwing, reaction diffusion, and boundary perturbation.

find this layout, we take a hill climbing approach, randomly assigning each piece one of five phases, searching for pieces with a color conflict, and toggling their color to resolve the conflict. If we find a situation we cannot resolve by a single color change, we start over.

At this point, we've placed our pieces, defined their initial boundaries, and assigned their phases. The initial phases and temperatures for a piece i are set to $p_i = 1$ and $T_i = -T_0$ and $p_j = 0$ for $j \neq i$ and $T_j = T_0$ for $j \neq i$.

5 Results

To generate the final puzzles, we predefine sets of parameters representing different cut styles. These parameters control the method of seed generation for initial boundaries, the wavelength λ_i and amplitude U of the boundary perturbations, and local parameters of the dynamics such as temperature sensitivity and supercooling. We then generate different regions of the puzzle and assign them different parameter sets to create multiple cut-styles within one puzzle. Figure 6 shows examples of different cut styles

Note that initial shape and boundary perturbations can strongly affect the character of the resulting pattern. This result is consistent with solidification experiments demonstrating the dependence of dendrite form on seed spacing [Diepers et al. 2002].

6 Conclusions

We've made some *really* awesome jigsaw puzzles, but opportunities for improvement still remain. By reducing the number of phases to five to increase efficiency, we create a problem of temperature bleeding. Locally, each piece is roughly in a supercooled environment; but over time, the temperature of pieces with the same phase spreads to one another. Dendrites “melt” into less detailed nodules over the course of long simulations. This could be avoided by increasing the efficiency of the simulation through adaptive grid or domain methods and preserving a unique phase for each piece.

We could also explore extensions to the solidification model and related systems which can be simulated with phase-field methods. More complex solidification models incorporate additional physical characteristics, such as crystal anisotropy, chemical concentrations, surface tension, and fluid flow. Incorporation of these characteristics would allow for more cut style variety.

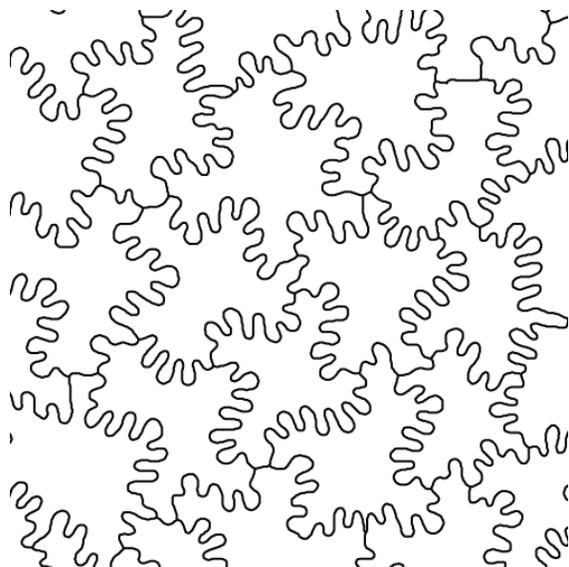
Other Laplacian growth problems have been approached using the phase-field method. Hele-Shaw cells are one such example [Folch et al. 2009] which may be a fruitful area to explore for cut style

inspiration. Unlike solidification simulations, Hele-Shaw cell simulations are phase conservative, making it more conducive to the evolution of puzzle boundaries.

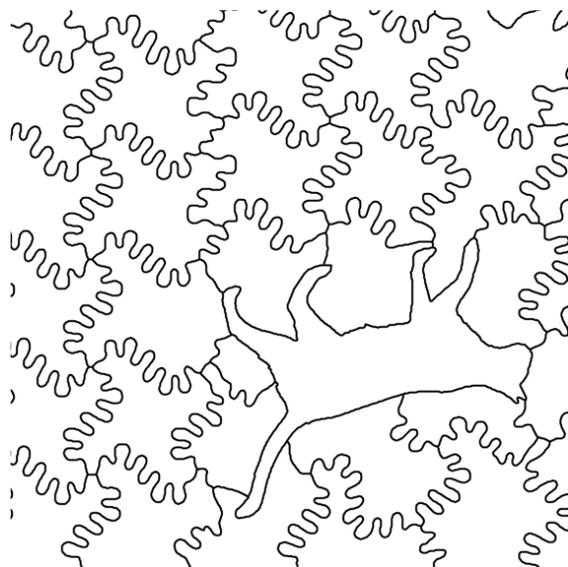
References

- ALLEN, S., AND CAHN, J. 1979. A microscopic theory for antiphase boundary motion and its application to antiphase domain coarsening. *Acta Metallurgica* 27, 6, 1085–1095.
- ALMGREN, R. 1993. Variational algorithms and pattern formation in dendritic solidification. *Journal of computational physics* 106, 2, 337–354.
- APPEL, K., AND HAKEN, W. 1989. *Every planar map is four colorable*, vol. 98. Amer Mathematical Society.
- BARRA, F., DAVIDOVITCH, B., LEVERMANN, A., AND PROCACCIA, I. 2001. Laplacian growth and diffusion limited aggregation: different universality classes. *Physical review letters* 87, 13, 134501.
- BATTYLT, M. 1991. Generating urban forms from diffusive growth. *Environment and Planning A* 23, 511–544.
- BEN-JACOB, E., SHMUELI, H., SHOCHET, O., AND TENENBAUM, A. 1992. Adaptive self-organization during growth of bacterial colonies. *Physica A: Statistical Mechanics and its Applications* 187, 3-4, 378–424.
- DIEPERS, H., MA, D., AND STEINBACH, I. 2002. History effects during the selection of primary dendrite spacing. comparison of phase-field simulations with experimental observations. *Journal of crystal growth* 237, 149–153.
- FAMILY, F., PLATT, D., AND VICSEK, T. 1987. Deterministic growth model of pattern formation in dendritic solidification. *Journal of Physics A: Mathematical and General* 20, L1177.
- FOLCH, R., ALVAREZ-LACALLE, E., ORTÍN, J., AND CASADEMUNT, J. 2009. Pattern formation and interface pinch-off in rotating hele-shaw flows: A phase-field approach. *Physical Review E* 80, 5, 056305.
- GONZÁLEZ-CINCA, R., FOLCH, R., BENITEZ, R., RAMIREZ-PISCINA, L., CASADEMUNT, J., AND HERNÁNDEZ-MACHADO, A. 2003. Phase-field models in interfacial pattern formation out of equilibrium. *Arxiv preprint cond-mat/0305058*.
- GOOGLE, 2012. jigsaw puzzle generating software - google search, Jan.
- HASTINGS, M., AND LEVITOV, L. 1998. Laplacian growth as one-dimensional turbulence. *Physica D: Nonlinear Phenomena* 116, 1-2, 244–252.
- JEONG, J., GOLDENFELD, N., AND DANTZIG, J. 2001. Phase field model for three-dimensional dendritic growth with fluid flow. *Physical Review E* 64, 4, 041602.
- KARMA, A., AND RAPPEL, W. 1998. Quantitative phase-field modeling of dendritic growth in two and three dimensions. *Physical Review E* 57, 4, 4323.
- KIM, T., SEWALL, J., SUD, A., AND LIN, M. 2007. Fast simulation of laplacian growth. *IEEE Computer Graphics and Applications*, 68–76.
- KOBAYASHI, R. 1993. Modeling and numerical simulations of dendritic crystal growth. *Physica D: Nonlinear Phenomena* 63, 3-4, 410–423.

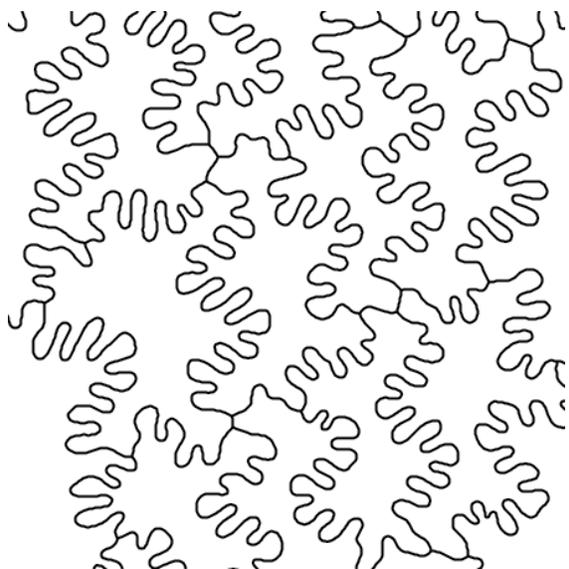
- KOBAYASHI, R. 1994. A numerical approach to three-dimensional dendritic solidification. *Experimental mathematics* 3, 1, 59–81.
- KRICHEVER, I., MINEEV-WEINSTEIN, M., WIEGMANN, P., AND ZABRODIN, A. 2004. Laplacian growth and whitham equations of soliton theory. *Physica D: Nonlinear Phenomena* 198, 1-2, 1–28.
- KUPFERMAN, R., SHOCHET, O., AND BEN-JACOB, E. 1994. Numerical study of a morphology diagram in the large undercooling limit using a phase-field model. *Physical Review E* 50, 2, 1005.
- LØVOLL, G., MÉHEUST, Y., TOUSSAINT, R., SCHMITTBUHL, J., AND MÅLØY, K. 2004. Growth activity during fingering in a porous hele-shaw cell. *Physical Review E* 70, 2, 026301.
- MINEEV-WEINSTEIN, M., AND ZABRODIN, A. 2001. Whitham-toda hierarchy in the laplacian growth problem. *J. Nonlinear Math. Phys* 8, 212–218.
- PIETRONERO, L., ERZAN, A., AND EVERTSZ, C. 1988. Theory of laplacian fractals: diffusion limited aggregation and dielectric breakdown model. *Physica A: Statistical and Theoretical Physics* 151, 2-3, 207–245.
- PROVATAS, N., GOLDENFELD, N., AND DANTZIG, J. 1999. Adaptive mesh refinement computation of solidification microstructures using dynamic data structures. *Journal of Computational Physics* 148, 1, 265–290.
- RUNIONS, A., FUHRER, M., LANE, B., FEDERL, P., ROLLAND-LAGAN, A., AND PRUSINKIEWICZ, P. 2005. Modeling and visualization of leaf venation patterns. In *ACM Transactions on Graphics (TOG)*, vol. 24, ACM, 702–711.
- SETHIAN, J., AND STRAINT, J. 1992. Crystal growth and dendritic solidification. *Journal of Computational Physics* 98, 2, 231–253.
- STEINBACH, I., PEZZOLLA, F., NESTLER, B., SEESSELBERG, M., PRIELER, R., SCHMITZ, G., AND REZENDE, J. 1996. A phase field concept for multiphase systems. *Physica D: Nonlinear Phenomena* 94, 3, 135–147.
- STRAIN, J. 1989. A boundary integral approach to unstable solidification. *Journal of Computational Physics* 85, 2, 342–389.
- WILLIAMS, A. D. 2004. *The Jigsaw Puzzle: Piecing Together a History*. Berkley Hardcover.
- ZHAO, P., AND HEINRICH, J. 2001. Front-tracking finite element method for dendritic solidification. *Journal of Computational Physics* 173, 2, 765–796.



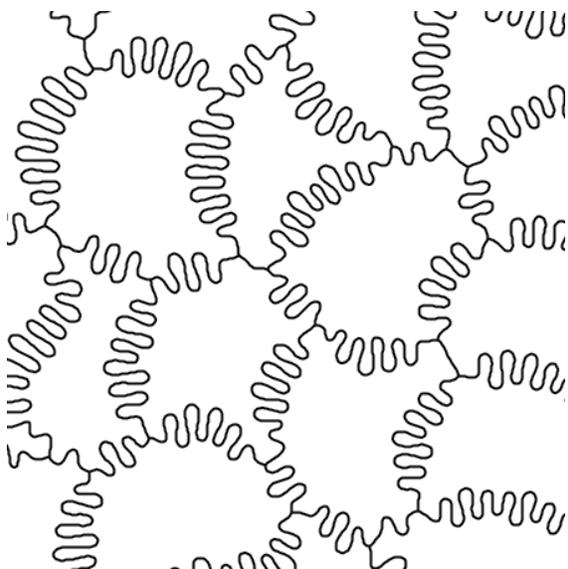
(a) Standard cut: Dart throwing with $\rho = 45$, $\lambda_1 = 20$, $\lambda_2 = 4$, $U_1 = 8$, $U_2 = 1$, $\gamma = 35$



(b) Grid cut with cat whimsy: Regular grid seeds with $\rho = 55$, $\lambda_1 = 20$, $\lambda_2 = 4$, $U_1 = 8$, $U_2 = 1$, $\gamma = 35$



(c) Linear cut: Seed with line segments in the same direction, $\lambda_1 = 20$, $\lambda_2 = 4$, $U_1 = 8$, $U_2 = 1$, $\gamma = 35$



(d) Amoeba cut: Dart throwing with $\rho = 45$, $\lambda_1 = 40$, $\lambda_2 = 4$, $U_1 = 8$, $U_2 = 2$, $\gamma = 35$

Figure 6: In all figures the grid spacing $\Delta x = 0.024$, the time step $\Delta t = 0.00015$, $\epsilon = 0.008$, $\tau = 0.0003$, $\alpha = 0.9$, and $T_0 = 0.3$.