

MODELING DIFFUSION IN A DISCRETE ENVIRONMENT

Adrian Grajdeanu
agrajdea@gmu.edu

Technical Report GMU-CS-TR-2007-1

Abstract

The present report details a model for implementing diffusion in a discrete space. Formulated in 2004 in support of artificial development experiments, the model is mathematically justified starting from the diffusion equation. It has physical plausibility and handles well different shaped and sized diffusion neighborhoods in the sense that it achieves isotropic diffusion in spite of the bias introduced by the discretizing grid. It provides one formulation that encapsulates the diffusion neighborhood details and renders it applicable to linear, planar, spatial and even n-dimensional constructs.

1 Introduction

In the field of artificial development, some researchers choose to endow the environment with diffusion properties as a form of information transfer. This report explains the details of implementing such a feature in the discrete case: the environment consists of a substrate subdivided in a grid of locations. This is not the only way to model an environment, specifically there is the option of considering the space as a continuum, without any discretizing grid. However, this continuous case is not covered in this report.

A significant concern when implementing diffusion in a discrete system is to do it in such way as to avoid artifacts betraying the shape of the grid or the diffusion neighborhood employed. For example, if one considers the space isotropic, any good diffusion algorithm operating on a rectangular discretizing grid will show a circular pattern of diffusion, without betraying the existing rectangular bias.

Back in 2003, when he introduced his Cartesian GP system, Julian Miller [4] uses a rectangular grid and a Moore neighborhood of size 1 (8 neighbors) to simulate diffusion. Based on personal communication dating from 2005 his system treats all neighbors similarly. This introduces an artifact that favors diffusion diagonally: in one step, a substance diffuses a larger distance towards the corner neighbors then towards its lateral ones.

A bit earlier, in 1998, Astor and Adami [1] encountered the same challenge of implementing diffusion in a discrete space. They avoided the issue of variously spaced neighbors by discretizing the space in hexagonal grid and using a neighborhood of 6 immediate neighbors, all of them the same distance away. They leave unanswered, even unasked, the question of how to handle larger neighborhoods on the hexagonal grid, or – even more limiting – how to handle the 3 dimensional space where a hexagonal (or hexaedral) tessellation is non-existent. Further, the update model they used required a limitation on the diffusion coefficient in order to avoid oscillations, thus betraying an implementation that lacks physical plausibility.

Finally, in 2004 I derived a diffusion model for discrete environments to be used in an artificial development experiment. Said model alleviates all the drawbacks mentioned above. Precisely, it handles various shape and size neighborhood, is not limited to two dimensions or even three dimensions for that matter, it has physical plausibility and hides the bias introduced by the discretizing grid. In 2006 I started publishing about this line of research. To date, the developmental model is mentioned in two publications [3] and [2], but none of them dwells into the details of the implementation, or the mathematical justification. It is the purpose of this report to fill in this gap and openly propose to the research community this diffusion model for discrete environments.

The model is formulated from first principles and with mathematical support and validation; the details are presented in section 2. An implementation in the system presented in [3] allows for empirical validation as illustrated in section 3. Finally, no model is without drawbacks, and the current case is no exception. Its own limitations are identified and discussed in section 4.

2 A Discrete Model of Diffusion

Mathematical models of diffusion in the continuous space have been proposed and studied extensively (see http://en.wikipedia.org/wiki/Heat_equation). In deriving a model for the discrete scenario, starting from first principles seems appropriate. The general idea of the diffusion model is that a surface separating two locations will allow molecules to diffuse both ways and the effect can be statistically estimated by a differential equation. In the simplest case (one dimension, homogeneous diffusion) with $f(x)$ denoting the concentration of substance at point x the equation has the following form:

$$\frac{\partial f(x, t)}{\partial t} = k \frac{\partial}{\partial x} \frac{\partial f(x, t)}{\partial x} \quad (1)$$

This equation summarizes the following statement: the overall number of molecules of a substance crossing the surface in a time interval (infinitesimally small) is proportional to the difference in concentrations of that substance on either side of the surface.

Solving this equation, with no boundary condition (unrestricted space) gives the form

$$f(x, t) = \frac{1}{\sqrt{4\pi kt}} \int_{-\infty}^{\infty} f(y, 0) e^{-\frac{(x-y)^2}{4kt}} dy \quad (2)$$

$$-\infty < x < \infty, \quad 0 \leq t < \infty$$

Computational reasons dictate a time and space discretizing, leading to the following simplifications: $t = 1$ (one time step) and the spatial locations being designated as x_i and their neighbors $\dots x_{i-2} x_{i-1} x_i x_{i+1} x_{i+2} \dots$. In these conditions $f(x_i)$ denotes the concentration of substance at location (not a point) x_i . Finally, with the \int becoming a \sum , a one step update is described by the following morphing of equation 2 into:

$$\bar{f}(x_i) = A \sum_{j=-\infty}^{\infty} f(x_{i+j}) e^{-(x_{i+j}-x_i)^2/\mu} \quad (3)$$

where $f(x) = f(x, t = 0)$ and $\bar{f}(x) = f(x, t = 1)$. $\mu > 0$ encapsulates the modulating effects of the diffusion coefficient and A is a normalization factor that insures conservation

$$\sum_{i=-\infty}^{\infty} \bar{f}(x_i) = \sum_{i=-\infty}^{\infty} f(x_i) \quad (4)$$

Substituting 3 in 4

$$\sum_{i=-\infty}^{\infty} \bar{f}(x_i) \quad (5)$$

$$= A \sum_{i=-\infty}^{\infty} \sum_{j=-\infty}^{\infty} f(x_{i+j}) e^{-(x_i-x_{i+j})^2/\mu} \quad (6)$$

$$= \sum_{i=-\infty}^{\infty} f(x_i) A \sum_{j=-\infty}^{\infty} e^{-(x_i-x_{i+j})^2/\mu} \quad (7)$$

$$= \sum_{i=-\infty}^{\infty} f(x_i) \quad (8)$$

The step in 7 is justified by noticing that all combinations of i and j are present in the summations. One can therefore rename x_{i+j} to x_i and factor it out without affecting the form of the exponent. Finally, the equality imposed by 8 (which is in fact 4) is satisfied if

$$A \sum_{j=-\infty}^{\infty} e^{-(x_i-x_{i+j})^2/\mu} = 1 \quad (9)$$

Symbolically, the value of A seems to depend on i , but in fact, under the assumption of isotropic universe (each location has its neighbors arranged in the same pattern), it is the same regardless of i .

In both the above 9, as well as in 3 the effect of neighbor x_{i+j} is modulated by an exponential decay in its distance to x_i , it is practical to limit the computation to a convenient fixed n , or in other words to a local diffusion neighborhood. With that the equation 3 becomes:

$$\bar{f}(x_i) = A \sum_{j=-n}^n f(x_{i+j}) e^{-(x_{i+j}-x_i)^2/\mu} \quad (10)$$

and 9 becomes:

$$A \sum_{j=-n}^n e^{-(x_i-x_{i+j})^2/\mu} = 1, \quad (\forall i) \quad (11)$$

Do equations 10 and 11 describe a process that plausibly resembles diffusion? As a simple validation step, consider the case of $\mu \rightarrow \infty$ (instantaneous diffusion) and $\mu = 0$ (no diffusion).

Equation 10 gives

$$\lim_{\mu \rightarrow \infty} \bar{f}(x_i) \quad (12)$$

$$= \lim_{\mu \rightarrow \infty} \sum_{j=-n}^n f(x_{i+j}) A e^{-(x_{i+j}-x_i)^2/\mu} \quad (13)$$

$$= \sum_{j=-n}^n f(x_{i+j}) \lim_{\mu \rightarrow \infty} \frac{e^{-(x_{i+j}-x_i)^2/\mu}}{\sum_{k=-n}^n e^{-(x_{i+k}-x_i)^2/\mu}} \quad (14)$$

$$= \sum_{j=-n}^n f(x_{i+j}) \lim_{\mu \rightarrow \infty} \lambda_{j,\mu}^i \quad (15)$$

$$= \sum_{j=-n}^n f(x_{i+j}) \frac{1}{2n+1} \quad (16)$$

$$= \frac{1}{2n+1} \sum_{j=-n}^n f(x_{i+j}) \quad (17)$$

where $\lambda_{j,\mu}^i$ introduces a shorthand notation for

$$\lambda_{j,\mu}^i = \frac{e^{-(x_{i+j}-x_i)^2/\mu}}{\sum_{k=-n}^n e^{-(x_{i+k}-x_i)^2/\mu}} \quad (18)$$

The above result in 17 attests that for $\mu \rightarrow \infty$ in one time step the concentrations are instantly averaged to equilibrium.

Similarly, for $\mu \rightarrow 0$ ($\mu > 0$) the same equation 10 gives:

$$\lim_{\mu \rightarrow 0} \bar{f}(x_i) = \sum_{j=-n}^n f(x_{i+j}) \lim_{\mu \rightarrow 0} \lambda_{j,\mu}^i \quad (19)$$

In order to calculate $\lim_{\mu \rightarrow 0} \lambda_{j,\mu}^i$ it is easier to notice that $\lambda_{j,\mu}^i > 0$ and calculate $\lim_{\mu \rightarrow \infty} 1/\lambda_{j,\mu}^i$.

In general

$$1/\lambda_{j,\mu}^i = 1 + \sum_{k=-n, k \neq j}^n \frac{e^{-(x_i-x_{i+k})^2/\mu}}{e^{-(x_i-x_{i+j})^2/\mu}} \quad (20)$$

$$= 1 + \sum_{k=-n, k \neq j}^n \alpha_{j,k,\mu}^i \quad (21)$$

where $\alpha_{j,k,\mu}^i$ introduces another shorthand notation for

$$\alpha_{j,k,\mu}^i = \frac{e^{-(x_i-x_{i+k})^2/\mu}}{e^{-(x_i-x_{i+j})^2/\mu}} \quad (22)$$

For $j = 0$, equation 20 becomes

$$1/\lambda_{0,\mu}^i = 1 + \sum_{k=-n, k \neq 0}^n e^{-(x_i-x_{i+k})^2/\mu} \quad (23)$$

thus,

$$\lim_{\mu \rightarrow 0} 1/\lambda_{0,\mu}^i = 1 \Rightarrow \lim_{\mu \rightarrow 0} \lambda_{0,\mu}^i = 1 \quad (24)$$

under the reasonable assumption that $x_i - x_{i+k} \neq 0, \forall k \neq 0$.

To analyze the case $j \neq 0$, notice from equation 22 that $\lim_{\mu \rightarrow 0} \alpha_{j,k,\mu}^i$ can either be 0, or 1, or ∞ . In particular, for $k = 0$ (and remember $j \neq 0$):

$$\lim_{\mu \rightarrow 0} \alpha_{j,0,\mu}^i = \infty \quad (25)$$

$$\stackrel{(21)}{\Rightarrow} \lim_{\mu \rightarrow 0} 1/\lambda_{j,\mu}^i = \infty \quad (26)$$

$$\Rightarrow \lim_{\mu \rightarrow 0} \lambda_{j,\mu}^i = 0, \forall j \neq 0 \quad (27)$$

Combining 19 with 24 and 27 gives:

$$\lim_{\mu \rightarrow 0} \bar{f}(x_i) = f(x_i) \quad (28)$$

which states that for $\mu = 0$ the diffusion subsides.

Thus, the credibility of equations 10 and 11 having been established, the one step update rule can be further modified as such:

$$\Delta \bar{f}(x_i) = \bar{f}(x_i) - f(x_i) \quad (29)$$

$$\stackrel{(11)}{=} \bar{f}(x_i) - A f(x_i) \sum_{j=-n}^n e^{-(x_i-x_{i+j})^2/\mu} \quad (30)$$

$$= A \sum_{j=-n}^n (f(x_{i+j}) - f(x_i)) e^{-(x_i-x_{i+j})^2/\mu} \quad (31)$$

Notice that just like A , under the assumption of isotropic universe, $d_j = x_i - x_{i+j}$ does not depend on i either. With that, the one step update rule becomes

$$\Delta \bar{f}(x_i) = A \sum_{j=-n}^n (f(x_{i+j}) - f(x_i)) e^{-d_j^2/\mu} \quad (32)$$

Or, in more general terms, if one considers a location x_i and its n neighbors $x_i^1 x_i^2 \dots x_i^n$:

$$\Delta \bar{f}(x_i) = A \sum_{j=1}^n (f(x_i^j) - f(x_i)) e^{-d_j^2/\mu} \quad (33)$$

$$d_j = |x_i - x_i^j| \quad (34)$$

$$A \sum_{k=1}^n e^{-d_k^2/\mu} = 1 \quad (35)$$

where d_j and A are dependent on the size and shape of the diffusion neighborhood.

Finally, even though the model is mathematically justified starting from the one dimensional diffusion equation, this latest incarnation (equation 33) is not limited to any number of dimensions. In effect, it provides a formulation that is oblivious to the dimensionality of the space, such details being encapsulated in the diffusion neighborhood (as in the number and distance of neighbors).

3 Empirical Validation

The model derived in the previous section (2) is implemented in a developmental system presented in [3]. An experiment is configured to use a 107×107 grid with a




Figure 1: Diffusion on a fine grid: 107×107 after 1500 steps

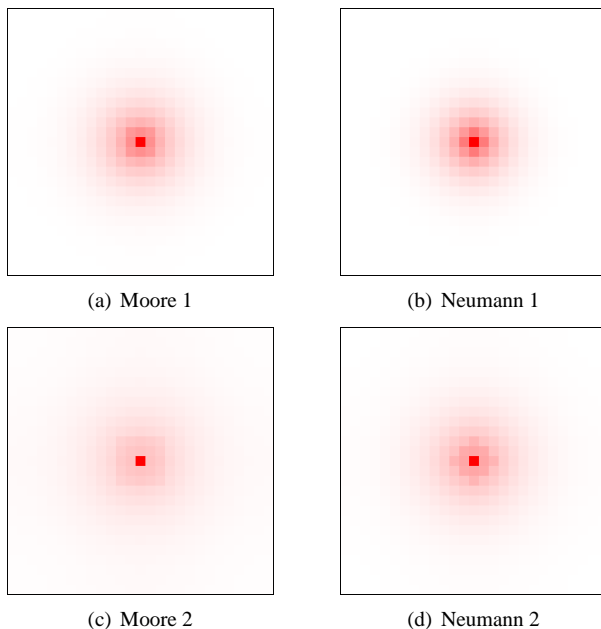


Figure 2: Diffusion on a coarse grid: 27×27 after 50 steps

regenerating well in the center that keeps the concentration of a substance maximal in the center. At each time step the substance is spreading according to various diffusion neighborhoods configured. Figure 1 shows the results of the experiment after 1500 time steps. The panels are labeled according to the diffusion neighborhood used: Moore (square) or Neumann (diamond) and their size. Notice that in all cases the substance spreads from the center in a circular pattern, not betraying the underlying bias of the rectangular grid.

However, the larger the neighborhood, the faster the substance diffuses, even though in all cases the algorithm is run with the same diffusion coefficient. This is one drawback of the model and will be further addressed in section 4.

Figure 2 shows the results of a similar experiment, but on a much coarser grid (27×27). While at macro level diffusion is still isotropic, the bias introduced by the diffusion neighborhood is reflected in visible artifacts. At this scale, this is an unavoidable drawback, also discussed further in section 4. Experimental evidence suggests that if one uses a discretizing grid that is two order of magnitude higher than the diffusion neighborhood size, such artifacts are no longer visible.

4 Drawbacks of the Model

This model, or rather its practical implementation is not without drawbacks.

The first drawback is that limiting the diffusion algorithm to use a certain neighborhood, limits the physical meaning of the diffusion coefficient. In the extreme case, if the diffusion coefficient is maximal (approaching ∞), then diffusion should be instantaneous. According to the model, it is, but only limited to the practical size of the diffusion neighborhood. If it were to be instantaneous to the whole space, then a diffusion neighborhood as large as the whole space would be required. However, in such case, the resulting time complexity of the algorithm would become square in the size of the space.

A second drawback is handling of space edges. The mathematical formulation (equation 33) relies on the assumption of isotropy, or in other words each location has the same configuration of neighbors. If such is not the case, the results of the algorithm are somewhat distorted in the proximity of the simulated space. One possible work-around is to choose a space large enough (or a discretizing grid that is fine enough) such that the practical application of a particular problem either doesn't reach the edges of the space or the effects of the distortion are not meaningful. Another work-around is to use toroidal spaces, such that there are no space edges at all.

A third limitation stems from the very use of a diffusion neighborhood. Even though the algorithm aims at hiding the discretizing grid bias, it is bound to betray some effects of the grid due to the fact that the diffusion neighborhood is embedded in said discretizing grid. However such effects are limited to a scale comparable to the size of the diffusion neighborhood itself. If one analyzes the macro effects, the algorithm is successful in hiding the bias of the grid. Thus, a possible work-around is to choose a discretizing grid that is fine enough for the practical purposes of a particular problem.

Finally, a fourth drawback is the time complexity of the algorithm. In its simplest form (with a small diffusion neighborhood size) the algorithm runs in $O(N)$, where N is the number of locations of the space (the size of the space). A larger space, or a finer discretizing grid both result in increasing the N number of locations, and consequently in a more time consuming diffusion process. Unfortunately, the only remedy to a number of other drawbacks is to increase this size of the space N . The problem is further compounded by the fact that the diffusion algorithm must be run a number of steps comparable to the number of developmental steps in the simulation. As such it turns out that running the diffusion process uses a significant portion of the time required to run a simulation. This fact prompted aggressive algorithmic optimizations such as advance identification and $O(1)$ fetching of neighbors, and pre-calculation of parameters such as the scaling factor A and the exponential factors $e^{-d_j^2/\mu}$ used in equation 33.

Still in one particular experiment when measurements were performed, running with diffusion turned off was one order of magnitude faster than running with it on. This observation dictates the following rule of thumb: if the model minus diffusion is computationally simple and there is no strong indication that diffusion is necessary for solving the problem, one is better off without diffusion. However, if the model is a bit more complex, such that running it dominates the time required to perform diffusion, then it would not hurt to leave diffusion enabled.

5 Conclusions

This diffusion model presented is grounded in mathematical first principles. It has physical plausibility and handles well different size and shape neighborhoods: at macro level (2 order of magnitude larger than diffusion neighborhood size) it hides the bias introduced by the discretizing grid. It is also not limited to a certain dimensionality of the space, providing a formulation that is applicable to linear, planar, spatial and even n-dimensional constructs.

The implementation of this algorithm has some limitations, chiefly among them the time complexity required. However this is a drawback that affects all diffusion algorithms equally. If diffusion were to be implemented in massively parallel manner then it would not be an issue. Barring the availability of such hardware, a better alternative for artificial developmental study is to replace the diffusion mechanism altogether as a means of transferring information, with other systems that explicitly carry it in a more efficient manner.

References

- [1] J. Astor and C. Adami. Development and evolution of neural networks in an artificial chemistry, 1998.
- [2] A. Grajdeanu. Methods for open-box analysis in artificial development. In TBD, editor, *Genetic and Evolutionary Computation Conference*, volume TBD, page TBD. Sheridan Printing, 2007. @ACM, (2007). This is the author's version of work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in proceedings of GECCO'07 <http://doi.acm.org/10.1145/nnnnnn.nnnnnn>.
- [3] A. Grajdeanu and S. Kumar. A novel developmental system for the study of evolutionary design. Technical Report FS-06-03, Developmental Systems; AAAI Fall Symposium, 2006.
- [4] J. F. Miller. Evolving developmental programs for adaptation, morphogenesis and self-repair. In *Proceedings of ECAL*, pages 256–265, 2003.