

An Atom is Known by the Company it Keeps: A Constructionist Learning Environment for Materials Science Using Agent-Based Modeling

Paulo Blikstein · Uri Wilensky

Published online: 29 October 2009
© Springer Science+Business Media B.V. 2009

Abstract This article reports on “MaterialSim”, an undergraduate-level computational materials science set of constructionist activities which we have developed and tested in classrooms. We investigate: (a) the cognition of students engaging in scientific inquiry through interacting with simulations; (b) the effects of students programming simulations as opposed to only interacting with ready-made simulations; (c) the characteristics, advantages, and trajectories of scientific content knowledge that is articulated in epistemic forms and representational infrastructures unique to computational materials science, and (d) the principles which govern the design of computational agent-based learning environments in general and for materials science in particular. Data sources for the evaluation of these studies include classroom observations, interviews with students, videotaped sessions of model-building, questionnaires, and analysis of artifacts. Results suggest that by becoming ‘model builders,’ students develop deeper understanding of core concepts in materials science, and learn how to better identify unifying principles and behaviors within the content matter.

Keywords Constructionism · Agent-based modeling · Complexity sciences · Materials science · Engineering education · Modeling · NetLogo · Multi-agent modeling

P. Blikstein (✉)

Learning Sciences and Technology Design Program, School of Education, and Department of Computer Science, School of Engineering (by courtesy), Stanford University, Stanford, CA, USA
e-mail: paulob@stanford.edu

U. Wilensky

Center for Connected Learning and Computer-Based Modeling, Northwestern Institute on Complex Systems, Northwestern University, Evanston, IL, USA
e-mail: uri@northwestern.edu

U. Wilensky

Department of Learning Sciences, School of Education and Social Policy, and Department of Computer Science, McCormick School of Engineering, Northwestern Institute on Complex Systems, Northwestern University, Evanston, IL, USA

1 Introduction

This paper reports on studies of a computer-based learning environment for undergraduate materials science. MaterialSim (Blikstein and Wilensky 2004a, 2005, 2006a, 2008) is an agent-based set of constructionist (Papert 1980) *microworlds* and activities built by the authors within the NetLogo (Wilensky 1999a) modeling environment. MaterialSim enables students to build models and investigate common college-level topics such as crystallization, solidification, crystal growth, and annealing. The design of MaterialSim emerged from extensive classroom observations followed by a literature review of engineering and materials science education, analysis of class materials and interviews with students. Our observations (supported by the literature review) indicated that students' understanding of the subject matter was problematic, and that the conventional teaching strategies and resources were not up to the challenge of the very complex content being taught. Based on this preliminary diagnosis, we created a set of exploratory computer microworlds and learning activities (Blikstein and Wilensky 2004b), and designed a series of studies for evaluating their effectiveness. A total of twenty-one undergraduates who were enrolled in a sophomore-level materials science course participated in three studies in 2004, 2005 and 2006, each comprised of a survey, pre-interview, interaction with the previously programmed computer models, students' construction of new models, and a post-interview.

Our classroom observations suggested that the ever-growing intricacy and extent of college-level content in engineering (and, in particular, materials science) poses a difficult challenge to traditional teaching approaches. One of the reasons is that the important equations and mathematical models taught in undergraduate materials science courses are not only very complex, but are often connected in nontrivial ways to multiple sets of other theories, conceptual units, and equations. Teachers end up resorting to a multitude of equations and models to derive and explain a single canonical phenomenon, and those equations and formulas are oftentimes located in a different area of mathematical modeling (statistical mechanics and geometrical modeling, for example). As a result, a plain linear progression of equations, from simple to complex, is not sufficient. In a typical classroom context, as a new phenomenon is taught to students, a very large number of new connections with previously learned topics will likely arise on multiple levels, generating even more specialized equations to account for those connections. The sheer number of equations generated makes a comprehensive exploration infeasible in the classroom. Our classroom observations revealed that, in a typical 30 minute period, students would be exposed to as many as 12 unique equations with 65 variables in total (not counting intermediate steps in a derivation)—or ~ 2.5 minutes for each equation and 45 seconds for each variable!

This overloading of equations and variables seems a likely candidate for explaining the students' difficulties described above. We decided to investigate this hypothesis and determine: what kind of understanding did this multiplicity of explanation levels and the "overloading" of equations foster in students? In addition to understanding the consequences of the traditional pedagogical approaches, we wanted to explore possibilities of an alternate approach, and examine its consequences for students' understanding of materials science content by using a series of agent-based models (Centola et al. 2000; Collier and Sallach 2001) and microworlds (Edwards 1995; Papert 1980). Our previous research suggested that using a multi-agent modeling platform to represent materials science phenomena might be a better match of content to student cognition.

The agent-based modeling approach, as we will explain in detail, enables modelers to employ simple individual-level rules to generate complex collective behaviors. These simple rules capture fundamental causality structures underlying complex behaviors within

a domain. Wilensky (1999b), Wilensky and Reisman (2006), and Wilensky and Resnick (1999) have pointed out that such rules could be more accessible to students than many of the equations describing the overall, macroscopic behaviors of a system. The agent-based approach is also a better fit with the constructionist pedagogical framework (Papert 1991). In the tradition of constructionist microworlds, agent-based models represent the underlying logic of a system, enabling students to investigate and modify features of that structure and explore the consequences of those changes, and through that exploration and investigation come to understand the domain. In this paper, we will particularly focus on the *building* of agent-based models within a constructionist framework.

The conjecture that using ABM would be a better cognitive match for students is based on research that suggests that using this approach fosters more generative and extensible understanding of the relevant materials science phenomena. In the case of materials science, instead of multiple models or numerous equations, this framework focuses on a small number of elementary behaviors that can be applied to a variety of phenomena. Instead of a *many-to-one* approach (many equations to explain one phenomenon), we attempt here a *one-to-many* approach (one set of local rules to explain many phenomena), through which students would see diverse materials science phenomena not as disconnected one from the other, but rather as closely related emergent properties of the same set of simple atomic or molecular rules. A second major focus of our study was to answer: What kind of understanding do students develop of the materials science content when they study it from this agent-based, *one-to-many* perspective?

In addition to those two driving questions, we wish to explore one further dimension of this pedagogical approach. There have been several recent studies of students using ABM to learn science; in many of these studies the approach taken was to design sequences of models and microworlds for students to explore (e. g., Levy et al. 2004; Stieff and Wilensky 2003). We extend this approach to the domain of materials science but mainly we wish to find out what the effect will be from moving beyond microworlds and enabling students to choose phenomena of interest to them and construct their own models in the domain of material science (for another of such approach, see Wilensky and Reisman 2006).

In this paper, we are focusing on the interviews and laboratory studies prior to the classroom implementation (subsequent design experiments on classroom implementations are reported in Blikstein 2009). We report on a particular pedagogical design and present evidence in the form of excerpts and samples of students' work, which demonstrates that the experience with MaterialSim enabled students to identify and more deeply understand unifying scientific principles in materials science, and use those principles to effectively construct new models.

Materials science is one of the oldest forms of engineering, having its origins in ceramics and metallurgy. In the nineteenth century, the field made a major advance when Gibbs found that the physical properties of a material are related to its thermodynamic properties. In the early twentieth century, the field of materials science concentrated on metals and university departments were often called 'metallurgical engineering departments'. The field has since broadened to include polymers, magnetic materials, semiconductors, and biological materials and since the 1960s has been called materials science. Today, with the explosion of research in nanotechnology, alternative energy, and new materials, it has gained a very significant role in the realm of technological innovation. However, the teaching of materials science has not kept up with the rapid advances in the field. Therefore, before diving in to the study, we step back and contextualize the teaching of materials science within the landscape of engineering education, its recent critique, and calls for reform.

2 A New Scenario in Engineering Education

In 2007, approximately 400,000 students took college-level engineering courses in the US alone (American Society for Engineering Education 2007). As early as the 1960s, education researchers (Brown 1961; Committee on the Education and Utilization of the Engineering 1985; Jerath 1983; MIT Center for Policy Alternatives 1975; Panel on Undergraduate Engineering Education 1986) have pointed out that engineering education lags behind in its adoption of newer approaches to teaching and learning. In recent years there have been numerous calls for reform from the engineering education community and several schools have implemented reform initiatives (Einstein 2002; Haghghi 2005; Russell and Stouffer 2005). The driving force behind engineering education reform programs were both new societal needs (Dym et al. 2005; Committee on the Education and Utilization of the Engineering 1985; Katehi et al. 2004; Tryggvason and Apelian 2006) and technical advances. As basic science and engineering become increasingly intertwined in fields such as nanotechnology, molecular electronics, and microbiological synthesis (Roco 2002), students and professionals have to deal with time scales from the nanosecond to hundreds of years, and sizes from the atomic scale to thousands of kilometers (Kulov and Slin'ko 2004). This wide range of subjects and problems makes it prudent not to try to cover all the relevant knowledge so that students master the knowledge in each domain, but instead to help students develop adaptive expertise (Bransford and Schwartz 1999; Hatano and Oura 2003) that they can apply to new problems and situations.

However, most engineering curricula remain in “coverage mode”—curricula are still so overloaded with transient or excessively detailed knowledge that there is no time for fostering students’ fundamental understanding of content matter (Hurst 1995). This phenomenon of curricular overloading is not exclusive to higher education. Tyack and Cuban (1995) identified the “course adding” phenomenon in most of twentieth century reform initiatives across all levels of education—new courses are regularly added to the curriculum to satisfy new societal needs. However, the situation becomes more problematic as we envision engineering schools in two or three decades from now. At some point the limit is reached and if courses need to be added, others must be removed—but can we afford to exclude anything from the curriculum? A major challenge is in how to go about deciding what courses can be dispensed with (and what knowledge).

A common approach in many universities has been to add hands-on engineering design courses to the curriculum. Design-based courses represented one attempted solution to the overcrowding of courses as they enable multiple content domains to be taught together. Design courses have been highly successful (Colgate et al. 2004; Dym 1999; Dym et al. 2005; Lamley 1996; Martin 1996; Newstetter and McCracken 2000), but they are not the universal answer. First, a significant part of engineering education constitutes of basic science (physics, chemistry), engineering science (fluid mechanics, thermodynamics) and mathematics (calculus, linear algebra). It is challenging for design-based courses to focus on the core conceptual elements of these knowledge domains as the physicality of students’ projects can be an obstacle for learning invisible or microscopic phenomena such as chemical reactions, pure mathematics, or quantum physics. Secondly, the technological tools used in those reform initiatives (such as modeling and design software) are the same employed by professional engineers in their everyday practice and not especially designed for learning. Using professional-based tools might be tempting as they enable students to more rapidly achieve the desired engineering design. In the specific case of materials science, however, this might not be the best choice. Most software tools used in engineering courses do not afford insight into the

computation underlying their design and functioning. For engineering practice, indeed, a tool has to yield reliable and fast results—understanding what’s “under the hood” is not necessarily useful, or possible. But in materials science this could be disadvantageous for learners. The computational procedures might embody an essential, perhaps crucial, aspect of the subject matter—*how* the conventional formulas and representations capture the phenomena they purport to model. Manifestly, no computer modeling environment can uncover all of its computational procedures—it would be impractical to have students wire thousands of transistors so as to understand the underlying logic of the modeling environment. Nevertheless, we believe that most of these environments could be made profitably more transparent to students. But the epistemological issues regarding the tools and knowledge representations in traditional engineering teaching run deeper.

First, in materials science, many of the traditional formulas *themselves* are opaque—they embody so many layers of accumulated scientific knowledge into such a complex and concise set of symbols that they do not afford common-sense insight and grounding of the causal mechanisms underlying the phenomena they purport to capture. Differently from the basic sciences, engineering knowledge is a complex matrix of empirical “engineering laws,” theories derived from fundamental mathematical or physical models, approximations, and rules-of-thumb. Making sense of this complex matrix is challenging for novices. Although using formulas and conventional engineering representations is perhaps conducive to successful *doing* (designing a new alloy, for example) it does not necessarily lead to principled understanding (knowing how each of the chemical elements interact and alter the properties of the alloy¹). Particularly, we are interested in “*extensible*” *understanding*—learning principles from one phenomenon that could be applied to other related phenomena.

Secondly, there is an important distinction to be made in how representations relate to the phenomena they purport to describe. We are not arguing that aggregate equational representations are intrinsically ill-suited for learning engineering or science. Far from it—there are many cases in which equational representations are fruitful for learning. Sherin (2001), for example, showed how the symbolic manipulation of formulas can lead to a gain in conceptual understanding in physics.

We are arguing that in some cases aggregate equations can hide important information needed for learning. In some areas of science, equations are directly postulated at the macro level, i.e., they are not necessarily an aggregation of simpler, local behaviors, or the microscopic behaviors are not relevant to the phenomenon under scrutiny. For example, in simple Newtonian motion, we are interested in predicting the motion of bodies, but looking at the individual atoms of the bodies might not offer additional insight into the phenomenon—the macroscopic and microscopic behaviors could be analogous, i.e., the body and its atoms would be moving in the same fashion. In such areas, aggregate equations reveal most of the needed information. In other domains, however, the opposite is true: equations are an aggregation of microscopic behaviors, and those offer fundamental insights into the phenomenon, and are not analogous to the aggregate equations (for example, statistical mechanics, or diffusion). Therefore, for the latter categories of phenomena, aggregate equational representations might generate an *epistemological gap* (Blikstein 2009)—the mathematical machinery needed to derive macro behaviors from micro behaviors is intricate, and rigorous mathematical frameworks to guide such work are still being developed (see, e.g., Parunak et al. 1998; Yamins 2005; Wilkerson-Jerde and Wilensky 2009). This epistemological gap makes it difficult to keep track of how micro- and macro-

¹ For more on design for learning versus design for use see, for example, Soloway et al. (1994).

level parameters are related and influence each other, or to understand how intuitive, simple micro-behaviors are represented in aggregate analytical forms. Our research, indeed, suggests that an exclusive use of equational representations for those types of phenomena can constitute an obstacle for students in acquiring conceptual understanding in domains of engineering in which the interaction of microscopic entities is at the core of the content matter. For those phenomena, in which equational representations represent an aggregation of micro-behaviors, it seems to be especially beneficial to unpack and deconstruct the traditional aggregate representations, restructuring domains of knowledge around the study of local, individual, “non-aggregated” phenomena (Wilensky et al. 2005; Wilensky and Papert, [in preparation](#); diSessa 2000).

For the most part, however, professional engineering tools whose main goal is arriving at results rather than uncovering processes, emphasize aggregate-level simulation to predict macroscopic variables (Wilensky 1999b, 2003). But the focus on micro-behaviors could make such content intrinsically more learnable and accessible. For example, *temperature* is a macroscopic, aggregate description of a microscopic state of individual molecules (their speed or energy), just as *pressure* is an aggregation of the number of collisions between gas molecules and the walls of the container. At an *aggregate* level, those variables are dependent on a number of different events and phenomena, and thus numerous equations and models have to be employed to predict them, oftentimes “mixing-and-matching” different levels of explanation and mathematical modeling approaches. On the other hand, at the microscopic level, the number of events and phenomena influencing a local interaction is dramatically lower than at an aggregate level, due to the fact that many of the variables observed macroscopically are just emergent properties of the local behaviors.

In this paper, we describe a learning design framework that benefits from this fact, focusing on simple agent-level behaviors (i.e., atomic- and molecular-level interactions) from which complex macroscopic behaviors emerge. We believe that this framework is especially useful in a scenario of increasing technological complexity and specialization. Materials science, for example, has transformed itself considerably over the last decade, with the advent of nano- and bio-materials, as well as the explosion of computational materials science as a core research strand. The number of materials, alloying elements, fabrication techniques, and industrial applications has grown so quickly and vastly that “covering” all the knowledge by simply adding new information to the curriculum would be infeasible. Additionally, the high level of abstraction that the new advances in materials science are bringing makes it increasingly difficult to give students any real world “feel” for the ideas learned in the classroom, as well as clear connections with their previous knowledge. While many archetypal problems in introductory physics would involve one falling body or two colliding objects, typical undergraduate problems in materials science involve simultaneous interactions of billions of atoms. Those interactions generate cascading effects that are hard to predict or understand with conventional mathematical equations, or any real-world intuitions. We posit that the micro-behaviors are easier to understand and model, and could be connected to previous knowledge and intuitions about how individual people or physical bodies behave (Wilensky 1999b). Thus, unifying behaviors embedded in agent-based models are helpful for acquiring solid understanding of these generative principles, which bridge the micro- and macro-levels (Wilensky and Resnick 1999). Consequently, we argue that the new computational tools should not be simple add-ons to the present curriculum, but part of their backbone—eventually restructuring the encoding of the content matter itself. In this, we follow the framework of Wilensky & Papert and their coinage of the word “restructuring,” (Wilensky et al. 2005;

Wilensky 2006; Wilensky and Papert, [in preparation](#)) to refer to the reencoding of knowledge in an alternate representational system.

Our approach is one attempt in this direction. It builds up from previous research on the use of agent-based simulation tools in schools at a variety of educational levels to investigate a wide range of phenomena. Wilensky and Resnick (1999) first noted the need to pay attention to “levels” and possible “slippages” between them, and highlighted the importance of the understanding of emergent behaviors for learning science. Wilensky, Papert, and colleagues have argued that computational representations have reached a point of development where we can embark on a program of radical “restructuring” of the science curriculum using these representations (Wilensky et al. 2005; Wilensky and Papert, [in preparation](#)). Goldstone and Wilensky (2008) have called for such a restructuring of science curricula using common transdisciplinary “patterns” such as energy minimization, positive feedback, and simulated annealing. In terms of implementation in school and universities, over the past decade and a half, educators have successfully employed agent-based modeling in undergraduate chemistry (Stieff and Wilensky 2003), high-school chemistry (Levy et al. 2004, 2006), probability and statistics (Abrahamson and Wilensky 2005; Wilensky 1995), robotics (Berland and Wilensky 2004, 2005), physics (Sengupta and Wilensky 2008; Wilensky 1993, 1999b, 2003), evolution, population dynamics and mathematics (Centola et al. 2000; Wilensky et al. 2000; Wolfram 2002). Ironically, despite the widespread use of multi-agent based modeling in materials science, we have not found significant research investigating the use of such models for learning and teaching it.

We will present and discuss a series of three small studies in laboratory settings of a computer-based learning environment designed within the constructionist framework, which addresses the aforementioned challenges by offering students opportunities to build their knowledge by exploring and designing computer models based on simple computational behaviors. MaterialSim is composed of several exemplar models, effectively being a “library” of simple computational algorithms in materials science. The main model in MaterialSim focus on the topic of *grain growth* in materials. Grain growth is perhaps the most important microscopic characteristic of a material, contributing to, among other properties, strength, toughness, and corrosion resistance.

The user studies were comprised of classroom observations, pre/post interviews, pre-post surveys, and data analysis from individual sessions with students using the curriculum materials.

Before diving into the study, some background information on materials science content and teaching is necessary to illustrate the differences between the traditional and the agent-based representations for grain growth. As the divergences in representation are at the core of this study, the next section will be dedicated to describing these two representations and how they differ. This will prepare the way for the description of our design and data analysis.

3 Equational Versus Agent-based Methods for Grain Growth Modeling

3.1 What is Grain Growth?

Most materials are composed of microscopic “crystals”. Even though we commonly associate the term ‘crystal’ with the material used in glassware manufacturing, its scientific use is different. A crystal is an orderly arrangement of atoms, a regular tridimensional grid in which each site is occupied by an atom. In materials science, scientists use the term “grain” to refer to such an arrangement.

Among other properties, grain size determines how much a material will deform before breaking apart, which is one of the most important issues in engineering design. For example, a car built with steel with a wrong grain size could just break apart during normal use, or be destroyed even in a minor accident. But grain size can change, too—high temperature is the main driving force. This phenomenon, known as *grain growth*, is exhaustively studied in materials science: small grains disappear while bigger ones grow (the overall volume is maintained). Airplanes turbines, for instance, can reach very high temperatures in flight—an incorrectly designed material could undergo grain growth and simply break apart. The following photographs (magnified 850×) show typical results after 20 hours under 900°C.




Fig. 1 Metallic sample before and after grain growth (Blikstein and Tschitschin 1999)

3.2 Equational Representation of Grain Growth

Burke (1949) was one of the first to introduce a law to calculate grain growth and proposed that the growth rate would be inversely proportional to the average curvature radius of the grains:

$$R = kt^n$$

where R is the mean grain size of the grains at a given time, t is time, k is a constant that varies with temperature, and n depends on the purity and composition of the material, as well as other initial conditions. Its theoretical value is 0.5 for 100% pure materials under ideal conditions.

In other words, Burke's law states that large grains (lower curvature radius) grow faster, while small grains (high curvature) have slower growth, or shrink. The mathematical formulation of Burke's law also reveals that, as grains grow, the growth rate decreases. A system composed of numerous small grains (Fig. 1, left) would have a very fast growth rate, while a system with just a few grains (Fig. 1, right) would change very slowly. One of Burke's approximations was to consider grains as spheres with just one parameter to describe their size (the radius). For most practical engineering purposes, this approximation yields acceptable results—however, as we previously discussed, its practical efficacy does not necessarily mean that this approach is the best way to understand the phenomenon.

Due to the applied and integrative aspect of engineering research and practice, oftentimes explanations are drawn from a variety of sources: empirical equations, geometrical

proof, thermodynamics, algebraic deductions, or statistical mechanics. Our classroom observations revealed that, for example, when explaining grain growth and deriving Burke's law, at least three sources were employed during the classes covering the phenomenon:

- The *Laplace-Young equation for pressure* is commonly used in fluid dynamics to calculate surface tension in liquid-gas interfaces (such as a drop of water). It states that the surface tension grows as the pressure difference is increased, and as the radii of curvature decreases. In other words, any small grain (with a low radius of curvature) will have high surface tension, as opposed to large grains. The equation is written as follows:

$$p^\beta = p^\alpha + \frac{2\gamma}{R}$$

$$du = \bar{V}_c dP$$

where α and β are the outside and inside pressures, R is the spherical particle radius, du the change in chemical potential, and \bar{V}_c the partial molar volume.

- The *flux equation* (based on statistical mechanics), which states that the probability of an atom to jump to a neighboring grain increases exponentially with temperature, and therefore the mobility of a grain boundary also grows with temperature. The equation is written as follows:

$$F_{\text{grain1} \rightarrow \text{grain2}} = A_2 n_1 v_1 \exp\left(\frac{\Delta G^\alpha}{RT}\right)$$

where F is the flux, A_2 is the probability of an atom being accommodated in the other grain, n_1 the number of atoms in grain 1 in position to make the jump, and v_1 the vibrational frequency of an atom in grain 1.

- Geometrical approximations are a common technique used to calculate average grain size and also the effect of dispersed particles or impurities in grain growth. Here, the force (P) is applied by a stationary particle to a moving grain boundary. The model assumes that particles are perfect spheres and grain boundaries are lines or “well-behaved” curves. The equation is:

$$P = \pi r \sin(2\theta) \cdot \gamma$$

where P is the force, θ is the angle with the grain boundary, r is the particle radius of the impurity.

We refer to this pedagogical approach as “many-to-one”: many models and equations to describe one phenomenon. Our research suggests that although the *many-to-one* modeling approach is useful in the hands of experienced engineers in real-world situations, this multitude of models can be an obstacle to student understanding.

3.3 Agent-Based Representation of Grain Growth

Apart from equational models, heuristics (engineering “rules of thumb”) are also important instruments for engineering practice. For example, when explaining grain growth, teachers commonly resort to a classic rule of thumb: large grains grow, small grains shrink. However, despite the usefulness of such heuristics to help students gain intuition into particular topics, they are not very generalizable, do not have a formal representation, and are usually domain-specific. The “large grains grow, small grains shrink” rule of thumb,

for example, was shown to be particularly inaccurate when, in the early eighties, with the availability of massive computing power, a new research approach came along: computer simulation of materials, in particular, grain growth. Anderson, Srolovitz and collaborators (Anderson et al. 1984; Srolovitz et al. 1984) proposed the widely known theory for computer modeling of grain growth using the multi-agent based approach (then referred to as “Monte Carlo method”). This kind of simulation not only made predictions faster and more accurate, but also allowed for a completely new range of applications. Researchers were no longer constrained by approximations or general equations, but could make use of more precise mechanisms and realistic geometries. As stated by Anderson et al.:

While it is generally observed that large grains grow and small grains shrink, instances where the opposite is true can be found. [...] The results indicate the validity of a random walk description of grain growth kinetics for large grains, and curvature driven kinetics for small grains. (Anderson et al. 1984, p. 796)

In other words, Anderson et al. state that the classic rule-of-thumb for grain growth (“large grains grow, small grains shrink”) is not always valid, and that randomness plays an important role. Given the microscopic dimensions and small time scale of the phenomenon, practically the only way to visualize this new finding is through computer simulation. In contrast, the *traditional* methods for investigating grain size and growth reflect the tools (and visualization techniques) that were available in the fifties: mathematical abstractions, geometrical modeling, approximations, and empirical data. These traditional methods and techniques, having become the methods of choice to explain the phenomena, made their way to textbooks and classrooms and were solidified as the mainstream path to study grain growth.

However, materials science has grown considerably from its roots in experimental metallurgy and, recently, a significant part of the research breakthroughs have been driven by advances in computational methods. Even though the new computational research tools are enabling researchers to accelerate scientific discovery and explore uncharted territory within their fields, computational methods have not yet reached the classroom. Thornton and Asta (2005) conducted a comprehensive survey about the state of computational materials science in undergraduate and graduate courses at the 20 leading programs in the US. While many universities are creating or planning to initiate computational materials science courses, one striking conclusion is that the prevailing mindset is that students should learn modeling *after* learning the “science”. In other words, computer modeling is regarded as “icing in the cake” to take place after the “real” scientific understanding, and not as an integral part of it. Our work, in contrast, evaluates the usefulness of a different approach: learning the science *by* modeling.

Grain growth is a prototypical example. In the last section, we described how it is common practice to teach students to consider grains as spheres (which they are not), grain boundaries as real entities (whereas they are just imaginary lines between grains), and to make use of numerous metaphors and rules-of-thumb (e.g., “grains swallow others,” “particles hold boundaries,” etc.) to describe and predict changes in the grain.

Agent-based simulation of grain growth offers a different perspective. Its principle is the thermodynamics of atomic interactions—a simple and powerful model which explanatory power covers a wide range of phenomena. The first step is to represent the material as a 2D matrix, in which each site corresponds to an atom and contains a numerical value representing its crystallographic orientation (the angle of orientation of the atomic planes in one particular grain compared to an arbitrary fixed plane).




Fig. 2 Initial and final free-energy calculations. *Black* and *white* arrows denote different or equal neighbors

Contiguous regions (with the same orientation) represent the grains. The grain boundaries are fictitious surfaces that separate volumes with different orientations. MaterialSim’s grain growth algorithm is described below:

- Each element (or agent) of the matrix has its energy² (G_i) calculated based on its present crystallographic orientation (2) and the crystallographic orientation of its neighborhood—the more neighbors of differing orientation, the higher the element’s energy. Figure 2 (left side) shows the central agent with four different neighbors, hence the value of its *initial energy* (G_i) is 4.
- One new random crystallographic orientation is chosen for that agent among the orientations of its neighbors. In this case, as observable in Fig. 2, the current value of the central agent is 2, and the new attempted value is 1.
- The agent’s energy is calculated again (G_f), with the new proposed crystallographic orientation. Figure 2 (right side) shows that there are only two different neighbors in the new situation, thus the *final energy* (G_f) decreases to 2.
- The two states are compared. The value that minimizes the energy is chosen. In this case, $G_i = 4$ and $G_f = 2$, so the latter value is lower and constitutes a state of greater stability.
- A simple, non-technical summary of this model is: the more different neighbors one has, the less stable one is, and thus more inclined to switching to a different orientation.

A crucial feature of the agent-based approach is the capture of the intricacy of the phenomenon in just one model. It is also a simple one which, additionally, could be used to understand other phenomena as well, such as diffusion or recrystallization—in other words, a “one-to-many” modeling framework.

Both traditional methods and computer-based methods of investigating grain growth rely on *modeling*. The scientific enterprise is the process of creating models that are the best approximations to reality we can find. Models from each historical period reflect the tools available at that time. The example of grain growth is illustrative of a common practice in many fields of academic research, in particular engineering. The availability of certain technologies for research shapes how researchers approach a certain problem, and the subsequent “encoding” of the knowledge is heavily influenced by those technologies. As the initially empirical or exploratory hypothesis gradually transitions to becoming a full-blown theory, they transmit much of those influences to the theories themselves, and consequently to the curricula. In this paper, we intend to show that the current “encoding” of the knowledge about grain growth, and materials science in general, is a function of the

² Although the rigorous term would be “free energy,” for simplicity we will use “energy”.

research technology, the state of the field itself, and not an intrinsically superior way of encoding knowledge.

In the following section, we describe the software infrastructure used in the project, and the design of the MaterialSim models.

4 Software Design: NetLogo and MaterialSim

4.1 NetLogo

NetLogo (Wilensky 1999a) is a direct descendant of the Logo language (Papert 1980). It is a freely available, integrated multi-agent modeling environment, designed and developed by the second author at Northwestern University's Center for Connected Learning and Computer-Based Modeling. It includes a graphical user interface for exploring, experimenting with, and visualizing models, as well as a multi-agent modeling language (MAML) used for authoring models (Fig. 3).

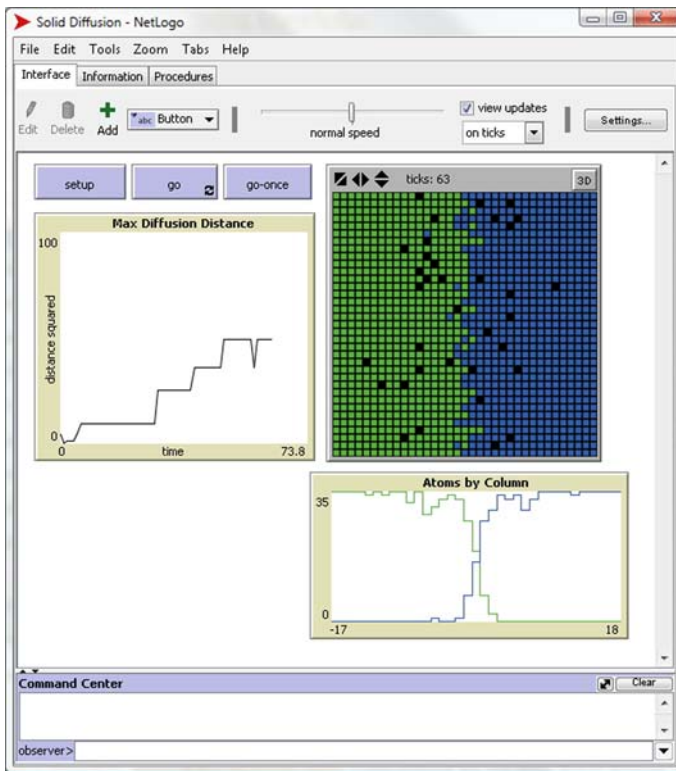


Fig. 3 The NetLogo modeling environment, with the “Solid Diffusion” model

Such languages enable users to easily create and manipulate thousands of computational entities (“agents”) and define simple rules that govern their behavior. For example, to create 100 agents (or ‘turtles’, in NetLogo’s lingo) on the computer screen, the user has to simply type:

```
create-turtles 100
```

To make all of those 100 turtles move 10 units forward, users would type:

```
ask turtles [forward 10]
```

Users can also define simple rules that govern agents' behavior. NetLogo agents can perform simple rule-based behaviors, such as to seek being surrounded by agents with similar properties, or to avoid areas already occupied by other agents. For example, to ask all turtles to check for neighbors (within a one-patch³ radius) and move backwards 10 units in case there are at least four neighbors around, we would use the following command:

```
ask turtles [if (count neighbors in-radius 1) > 4 [back 10]]
```

Such simple agent rules, however, give rise to complex emergent *aggregate* phenomena, many of which are congruent with their traditional macroscopic formula-based descriptions. In addition to the modeling language itself, NetLogo includes a graphical user interface with advanced visualization features, such as multiple topologies and 3D representations. It also includes some specialized tools such as BehaviorSpace (Wilensky and Shargel 2002), which enables users to explore a wide parameter space by running multiple experiments, and automatically logging the data. NetLogo comes with a large library of sample models as well as model-based curricula.

4.2 MaterialSim

We chose the NetLogo modeling-and-simulation environment as a platform as it is well adapted to all of those activities, in particular, NetLogo's "low-threshold, no-ceiling" (Papert 1980; Tisue and Wilensky 2004; Wilensky and Rand 2009) design enables learners to achieve sophisticated results within a relatively short period of time, and its built-in visualization tools allow dynamic, flexible, and customizable views.

MaterialSim is a set of models and activities built by the authors within the NetLogo environment, for investigating materials science phenomena such as crystallization, solidification, casting, grain growth, and annealing. MaterialSim's design is different from many other curriculum projects where instructional designers often prepare a series of models for students. MaterialSim, instead, focuses on students *programming* models. This design choice was based on previous research and on the learning goals of the project. For example, previous studies on students programming their own agent-based models report that participants were able to infer behaviors based on incomplete or sparse information, as well as gain deep understanding of how changes in micro-behaviors can alter a given system (e.g., Berland and Wilensky 2004; Centola et al. 2000). In contrast, scripted curricula (e.g., Gobert et al. 2004; Horwitz et al. 2003) start out with well defined content coverage. While studies of the scripted curricula report positive learning results, in their more 'convergent' approach, they do not necessarily afford insights into other areas of their target domain, nor allot large chunks of time for a deeper examination of the elementary "under-the-hood" behaviors. Since these curricula do not make use of programming, it is understandable that students' familiarity with the behaviors and rules will not be as well developed as in activities which involve model-building. And if one key goal of MaterialSim is to train students to see commonalities across phenomena and domains, having a *strong programming component* became a key design principle.

³ The NetLogo screen is divided into a grid of patches. The size of the patches can be defined by the user.

Creating models is not foreign to undergraduate engineering—it is common for engineering students to have modeling assignments and learn several programming languages while obtaining their degree. However, traditional model-based activities in engineering oftentimes do not afford understanding of microscopic behaviors or elementary physical/chemical principles. Therefore, another key design principle is to *build activities which foreground these micro-behaviors*, and in which students develop a high level of familiarity with the language and the ABM paradigm. In this study, following our framework, the design foci are:

- *Pre-designed programming exemplars* (solidification and grain growth) as to present students with the important algorithms and coding examples which could be useful in the process of building other models
- *Support materials* to help students in learning how to program in NetLogo.
- *Easily transportable code examples*, which students could easily reuse across models.
- *Readily available validation tools*, as to enable students to verify the validity of their models.
- A *persistent library* of student-generated models from previous years, from which students can reuse code and get inspiration for their models.

MaterialSim’s grain growth model, the main exemplar model (Fig. 4), was conceived to enable four kinds of activities:

- *One-dimensional exploration*: users can change variables, draw microstructures using the cursor, and observe their behavior over time.
- *Multi-dimensional exploration*: students can run experiments sweeping entire parameter spaces, as to find out critical points, system rules, mathematical relationships, and patterns.

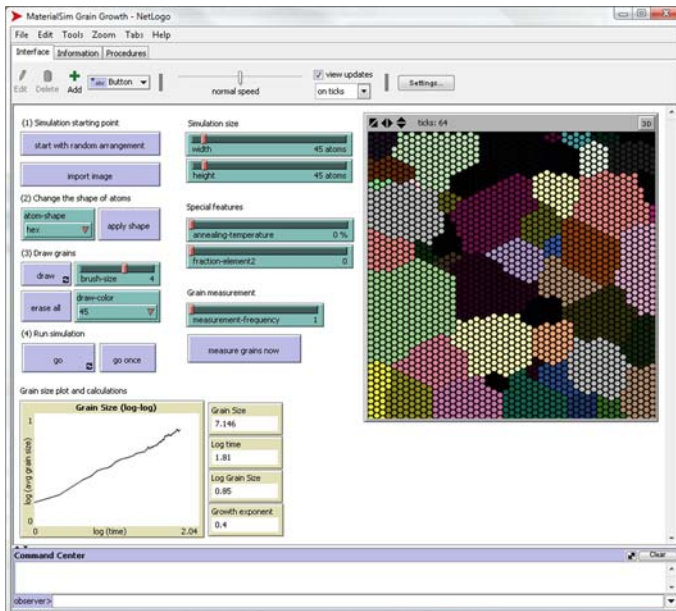
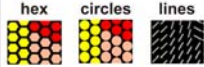


Fig. 4 The interface of MaterialSim’s grain growth model

- *Bifocal exploration*: students can connect real-world and virtual experiments, importing digital pictures from real experiments, and observing their “virtual” evolution. “Bifocal” refers to the simultaneous focus on the model and on the physical phenomenon (Blikstein and Wilensky 2006b).
- *Model building*: students can change, create, or extend the system by coding their own procedures, modifying existing ones, or creating whole new models from scratch, by using the NetLogo modeling language.

Also, the grain growth model offers a number of learning-oriented features, summarized in Table 1:

Table 1

Simulation setup	Atoms' shape	Drawing tools	Additional parameters
<p>(1) Simulation starting point</p> <p>start with random arrangement</p> <p>import image</p> <p>(4) Run simulation</p> <p>go go once</p>	<p>(2) Change the shape of atoms</p> <p>atom-shape hex apply shape</p> <p>hex circles lines</p> 	<p>(3) Draw grains</p> <p>draw brush-size 4</p> <p>erase all draw-color 45</p>	<p>Special features</p> <p>annealing-temperature 0 %</p> <p>fraction-element2 0</p>
<p>Users can start either from a random arrangement of atoms or from a pre-set stage, which can be drawn by the user using the mouse or converted from a digital photo. This enables easy exploration of “what-if” scenarios and the comparison with real-world data.</p>	<p>The appearance of the atoms can be changed for better visualization of particular aspects of the phenomenon. The “lines” mode, for example, was useful to understand the differences in crystallographic orientations.</p>	<p>To investigate particular scenarios, users can “draw” their own microstructures and observe their behaviors, or retouch existing microstructures.</p>	<p>The “annealing-temperature” slider enables experiments with different levels of thermal activity. The “fraction-element2” slider introduces a percentage of dispersed particles of a different material in the sample. Therefore, users can change the temperature and the type of material in order to study their effects on grain growth.</p>

5 Research Design and Methods

The research took place during three spring quarters at the materials science department of a Midwestern research university, with sophomore students enrolled in the “Microstructural Dynamics” undergraduate course. In the first year (2004), six undergraduate students (volunteers) participated in the study. In the second year (2005), eleven students volunteered to participate, and four students participated in the third year (2006), with a total of 21 participants over 3 years. The average class size was 15 students. Each student participated in two individual sessions. The first, 75 min long, was comprised of the following parts:

- Short Likert-scale/open-ended pre-survey to assess students’ familiarity with computers and their attitudes about the course.
- Pre-interview about grain growth and related phenomena, in which students were asked the following questions during a semi-structured interview. These questions were based on exams and assignments used in the course.
 1. What is a grain?
 2. What is a grain boundary?
 3. What is grain growth?
 4. Could you indicate in this picture which grain will grow and which will shrink? (Students were presented with a schematic drawing showing grains of different size in a material)
 5. What is the driving force for grain growth?

6. What is the driving force for recrystallization?
 7. What is the effect on grain growth of dispersed precipitates?
 8. In grain growth, grain boundaries always migrate toward their center of curvature. How does this decrease the free energy?
 9. In recrystallization, the new grains migrate away from their center of curvature. How does this lead to a decrease in the free energy?
- General presentation of the NetLogo programming environment.
 - Demonstration, by the first author, of five canonical agent-based models from the NetLogo models library (fire spread, virus contamination, racial segregation, gas molecules in a container, and a chemical reaction).
 - Hands-on interaction with one MaterialSim model: grain growth (with simultaneous interview). This included running the model with different values for matrix size, temperature, composition, as well as recording and plotting experiments sweeping the entire parameter space of one variable.

As “homework,” the participants were asked to choose a challenging and/or interesting topic from the course and think of a model to build, which would be implemented during the next session. Students also had the option of extending the functionality of one existing model.

The second session (150 minutes) was dedicated to:

- Introduction to the basic commands of the NetLogo modeling language.
- Implementation (i.e., coding) of the new model. Participants were always in front of the computer and in control of the task. The first author would help students as needed with language commands and general programming issues.
- Final interview.

We scheduled the sessions ~1 week after students’ exposure to the topic of grain growth in their regular classes.

All sessions were videotaped, and students’ computer interactions were recorded using real-time continuous screen-capture software. Approximately 65 hours of video were captured, which were selectively transcribed and analyzed. Experiments conducted by students, as well as the models they built, were saved and analyzed. The first author attended the Microstructural Dynamics course 2004, 2005, and 2006, analyzed the class materials and related literature. The classroom observations also generated data about the number of equations, variables, drawings, and plots explained during the class periods (and time spent in each item). Finally, participants were asked to fill up an anonymous web-based post-survey, as to assess their (self-reported) interest and motivation doing the study, as well as usefulness of computer simulation for understanding certain topics in Microstructural Dynamics.

6 Data Analysis

6.1 Pre-Interview Explanations

The pre-interview was semi-structured, and students could also draw pictures to illustrate their explanations. It was an open-book interview, i.e., students could resort to any class

material, book, or website to answer the questions. The first author asked the questions listed in the Sect. 5.

Below is a commented transcription of excerpts of answers to Question 1.⁴ For this part of the analysis, we randomly selected students from the two first studies (2004 and 2005). The goal in this section is to show and discuss students' explanations of some core content in materials science, to which they were exposed during their regular classes about 1 week before the interview.

Interviewer: How would you explain what a grain is?

Bob: A grain is just an orientation of a crystal structure, a group of atoms in a crystal structure, a specific orientation, it is just molecules aligned [pause] in one direction and then you have various grains and each grain has its own direction and when they meet [there is a] grain boundary.

Erika: A grain is just when you have got everything the same, same structure and everything, and then you have a boundary around it and it doesn't line up with the neighboring crystals. [...] One is oriented differently but basically they are the same.

Betty: If you have a metal and you are looking at it under a microscope, at the surface of the metal, where different [pause] what I call grains grow, and these grains are just areas where the atoms collect and the boundaries in between the grains are [pause] sinks and sources for dislocations and [pause] vacancies and so.

Liz: It is so hard to explain... When I think of grain, those kinds of pictures that we see, what comes to [my] mind [...] I feel that it is a word that we use in the English language and you always associate with something small, that is individual, you can see its limits and stuff like that. So when you use it in class, you just associate it with... like... I mean, a grain of rice... it is just one, out of many, that you can see right now.

Ken: A grain is basically a region of materials where there are no dislocations to it.

Ella: A grain? There is the material... A material is just made up of grains... I guess if you... Like, I could say something about the boundaries, grain sizes are different, they can change, with temperature and other variables... I guess I'm not really explaining what a grain is.

These first excerpts illustrate how dissimilar students' explanations were. This may at first seem surprising as the idea and description of a *grain* or *crystal* is a fundamental notion in materials science, a building block for further understanding of a variety of concepts in the field and that students had been exposed to the topic in class just a week prior to the interview. However, based both on the literature on prior conceptions and on our own classroom observations, the diversity of students' explanations is not surprising. The research literature is replete with examples where even simpler ideas are not learned by students in unison, and that their understanding is strongly influenced by previous knowledge, phenomenological primitives (diSessa 1993) and ad-hoc theories. Moreover, our classroom observations showed that instructors overloaded students with a multitude of phenomena and equations without necessarily making a clear distinction between the atomic behaviors themselves and the mathematical descriptions of those micro-behaviors.

Bob and Erika offer a reasonably complete explanation, close to what we would find in a textbook. Most of the other students diverge from that version. This divergence is also

⁴ All names were changed for anonymity.

not coherent: the starting points for the answers are rather different, as well as students' approaches to it. Betty based her explanation on the visual appearance of a grain seen under the microscope. Liz utilizes her previous knowledge about the morphology of a "real-world" grain (a grain of rice). Ken uses the definition of a grain combined with another topic in the course, dislocation theory. Ella apparently understands what a boundary is, but has difficulty explaining what exactly they enclose. Therefore, students resort to a variety of metaphors and explanations for characterizing a grain, at multiple levels: the surface under the microscope, a grain of rice, or atoms with same orientation. Many responses to this first question did not diverge completely from the acceptable concept. However, as the interview progressed and questions started to deal with more complex and dynamic phenomena, the diversity of explanatory models increased, and those initial difficulties unfolded in significant ways. In the following excerpts, students try to explain the phenomenon of grain growth.

Interviewer: How would you explain what grain growth is?

Bob: Grain growth is just when a grain, when more molecules [pause] come into one grain and line up that same direction. The driving force is [long pause]...

Interviewer: Why don't they stay the way they are?

Bob: Well, I mean [pause] I know the method of it, it is diffusion.

Interviewer: But what is the reason they grow?

Bob: Well, grains grow through diffusion, through vacancy diffusion, and atomic diffusion, for one, it is all over the place, temperature increases, molecules move around faster and they just... [pause] but the reason that they would grow [pause] I guess they grow... the driving force is to lower the free energy, overall, there is excess free energy due to dislocations and impurities in the grains, so by growing out, they can get rid of those and thus lower the free energy.

Betty: So when you heat-treat a metal, so when you deform a metal first the grains shrink and become compacted, you get all sorts of dislocations then, like twin boundaries, stuff like that, so if you do a high temperature anneal, then the grains all grow because you increase the energy of the system when you heat it, and so it tends to decrease its internal energy, so the grains become bigger and anneal out the dislocations because [pause] there is a high mobility for the atoms to move, and so they move to the lower energy positions which is into the grains and the *grain decrease*.... ah... the *grain size increases*, and the total area of the grain boundaries decrease, which decreases to decrease the overall energy of the system.

Liz: It is because, it wants to be more energetically stable, or have less energy in the crystal, so it will grow, just to form one big grain, because that's the least energy configuration, and it does this because, by the whole radius of curvature idea, where it starts shrinking.

Chris: Grain growth is... The smaller grains have higher curvatures and higher curvatures is not good, so they will want to shrink and become smaller and smaller, and bigger grains, with a lower radius of curvature will want to expand and so sooner or later will consume the smaller grains.

Peter: Molecules with high energy, which are over here, will jump over to the low energy spot and that's a more desirable position, and that's why grain growth grows.

Interviewer: Ok, when you say high energy or low energy, is it a general kind of energy, or are you talking about a specific kind?

Peter: It's called "free energy," but I don't really know how to explain that.

This question brings about at least three different ways to explain grain growth. The diversity of models and explanation paths is even more apparent than when students were explaining what grains are. Bob, for example, uses the metaphor of free will (“molecules come into the grain and line up”), and employs ideas about diffusion, dislocation and impurities in contradictory ways. He does not resort to the Laplace-Young equation, for instance, to explain the process of decreasing free energy by simply increasing the curvature radius. To him, excess free energy is due to impurities or imperfections in the crystal structure (known as dislocations). “Purity” is taken as a synonym for low energy, and grain growth as a cleansing mechanism by which grain boundaries would “sweep out the dirt”. However, the Laplace-Young equation (studied in class) states a very different idea. Namely, the driving force is the curvature or pressure difference—impurities are not eliminated by grain growth, and growth can exist in 100% pure materials. Here we can again notice that students “mix-and-match” models that appear superficially to be related, such as “grain growing” and “grains boundaries pushing impurities out.” Betty goes even further searching for explanations. The phenomenon she describes (deformation and recrystallization) was taught in a previous section of the course but is, in fact, very different from grain growth. In recrystallization, similarly, crystals grow, but for a different reason and with different kinetics. During the pre-interview, when presented with a printed picture of grains, she incorrectly indicated that the small ones would grow (which would happen in recrystallization). Moreover, she mentions that grains “all grow” to decrease the internal energy of the system, whereas in fact some grow and some shrink (otherwise, evidently, the material would expand). Liz’s explanation, on the other hand, is more coherent, relying on the idea of having “less” energy in the crystal being correlated to the “whole radius of curvature idea,” but without demonstrating how those things connect. Ken, similarly, was more coherent in his explanation, although using vague expressions such as “curvature is not good, so they will want to shrink”.

Bob and Ken provided good additional examples of the model “mix-and-match”.

Interviewer: What is the effect of dispersed particles of impurities?

Bob: I think that... I feel that particles constrict grain growth, and when a grain boundary meets a particle it bends around it, it kind of molds around it, it will slow it down, it won’t make it stop completely, but it will slow it down. It hits the particle, it goes around it, and as it happens, there is a pull on the opposite direction that the grain boundary is moving. They slow... every time the grain boundary is moving out it slows down the growth.

Interviewer: Is it good?

Bob: It depends on how big you want your grain. You know, more particles, the closer they are together, the smaller the grains will end up being, in that case it will be a harder and strong material

Betty: If you have a lattice around the impurity and that increases the energy of the system, and then that is bad, but if you have a lattice and you add particles of similar grain size, or, similar atom size, you can strengthen the material, because this small grains or atoms act to stop dislocation movement, and so it becomes harder and harder to push dislocations through so the plastic deformation becomes harder. The more particles you put [pause] in a system, the harder it is, [pause] the closer the spacing, the harder it will be to put dislocations through, so the harder the material will be.

Liz: Basically, if there is an impurity, the grain boundary will just go around it so it will just do like this, viewing it from a top angle, this will be a grain boundary [showing a drawing]

Interviewer: Will more particles affect grain growth?

Liz: [looks at her class notes for several seconds] As for now, I don't think it does anything to grain size.

Interviewer: What is the effect of dispersed particles?

Ken: if you have two precipitations and if you have a dislocation line, you need to exert a force τ on it, to move the dislocation line, but once it gets to the precipitation, it has to bow out and that will cost more energy, so if you have precipitations it will strengthen the material and that depends on the density of precipitations.

Interviewer: So grain growth slows down or is faster?

Ken: That I am not very sure.

Again, students tried to employ a variety of micro-models: a *force-feedback* model, where particles pull boundaries away, slowing grain growth (Bob); a *classical mechanics* model, in which boundaries can “hit” a particle and slow down (Bob), a *dislocation movement* model (Betty), and a purely *geometrical* one, with no consequences for the behavior of the material (Liz). Betty's explanation has relevant elements of dislocation theory (another topic explored in class weeks before), but does not address completely the question about grain growth; Liz does not see any relationship between grain size and dispersed particles; Bob only sees it as a factor that may decrease speed, but never “stop it completely”. Ken offers a reasonable explanation but fails to say if dispersed particles accelerate or slow down grain growth. His statement is very similar to Peter's, and suggests that they might know a considerable part of the theory, but lack fundamental understanding of the basic atomic behaviors within a material. Ken knows how to recite back the theory about grain growth, but cannot articulate the physical meaning of the formulaic representation.

The pre-interviews, therefore, suggest that students' explanations, sewn together on-the-fly, leverage a variety of models, admix different topics (recrystallization, dislocations, grain growth), and often use the standard vocabulary and rules-of-thumb of the field, but express a weak sense of the interconnectedness, relationships, and possible contradictions amongst all of those components which could be describing different aspects of atomic movement.

6.2 First Session of the User Study: Introduction and Model Exploration

The first session was dedicated to the exploration of the grain growth model. The first activity was simple: observe and reflect on curvature as a driving force for grain growth. Most of the students knew, from their previous classroom instruction, that large grains ‘consume’ small ones, growing toward their center of curvature, and high-curvature boundaries tend to disappear. However, those concepts appeared to be isolated ideas, separate phenomena, and hardly connected to the Laplace-Young equation, which was supposed to be the mathematical evidence for the aforementioned phenomenon.

The goal of this activity was twofold. First, assess students' pre-existing understanding of the phenomenon. Secondly, we carefully observed the cognitive shift as the simulation progressed (Siegler and Crowley 1991). This activity consisted in drawing two grains divided by a curved surface and observing their behavior. The pictures below are snapshots of the dynamic simulation that students observed (Fig. 5).




Fig. 5 The evolution of a curved grain boundary

Before the simulation, most students were unsure of what would happen. Approximately half of the students thought that the larger grain would grow at the expense of the smaller, regardless of the curvature of the boundary separating them, while the other half considered concavity, rather than size, as the main criterion. As they started the simulation, and saw grains growing toward their centers of curvature, they also observed that the movement was not smooth or unidirectional, but that there was intense activity on both grains with random flipping of atoms. The following excerpt suggests that visualizing this progression sparked some changes in Liz's understanding:

Interviewer: Can you describe what you see?

Liz: Just because one grain has a concave side and the other has a convex side, so it comes in towards the concave, because... [pause] does line tension applies in this situation?

Interviewer: Line tension?

Liz: That might be from dislocations... I might be mixing them up. Just because... when you have something... part of the grain is like, curving in, mostly likely other parts of the grain are curving in, so the tension of the grain boundary lines, so the force outside is greater than the force inside, so it will like shrink, it looks like that probably be like straight in the middle, rather than entirely red... just because if the red part also has some concave thing that is off the screen it will just like go together.

Liz is apparently mentioning results of the Laplace-Young equation, which relates surface tension and curvature. However, she cannot yet think at the “micro” level: To visualize what is happening on the computer screen, she has to imagine a large circle going off-screen, which is probably a consequence of what she remembers from class, where grains were always approximated as spheres. She does not yet construe the local interactions along the curved interface as a driving force, but only the “macro,” aggregate level effect of curvature.

The next activity was to draw a microstructure with many grains, but with one of them a lot smaller than the others, as we can see in Fig. 6.

Watching the evolution of this new microstructure was a crucial experience for Liz. She started to transition from memorized rules-of-thumb and a topic-specific model to micro-level reasoning. This excerpt took place when she was observing a triple point—a region where three grains meet and the probability to flip to any of the surrounding grains is the same, as there are two atoms of each grain around the central element (Fig. 6.) While observing this phenomenon, Liz was told to zoom in and out the simulation, so as to also see what was happening at the micro level (following a single atom.)




Fig. 6 Four large grains (yellow, green, light and dark blue) surround a small red grain (left), and a zoomed-in view of the structure showing a triple point (right) (Color figure online)

Liz: Right here there is an equal position for red, yellow and blue, but it just happens to be that blue won, it keeps winning.

Interviewer: How would you explain that?

Liz: Because... if you look at one of those points, either of the three colors, they all have the same number of other colors around it, so it is not favorable to choose one or the other...

Interviewer: What angle is here?

Liz: Oh, so this is the 120 degree angle between the... [pause]

Interviewer: Did you talk about it in class?

Liz: Briefly. He [the professor] said that when you reach a triple junction, it will become 120 degrees.

Interviewer: So are you saying that there is an equal probability?

Liz: Well, I just don't understand why blue is doing so much better, in general. Eventually just one has to become bigger, because this is the most energetically favorable thing, so maybe... blue was bigger, but now yellow is coming back, so maybe next time blue gets bigger again, and they will just keep going. Maybe it will just be like that for a long time.

Interviewer: So what happens to growth speed?

Liz: Eventually they will get like... two big ones... and then it will take forever.

Interviewer: So what could be the law?

Liz: It will eventually taper off... to some point... because if you have a lot of grains then you will... the rate of increase will be faster, but when average grain size increases it gets harder and harder to increase the rest of them, so it just goes...

Interviewer: Why is it harder and harder?

Liz: Just because there isn't a distinct... [pause] being in this orientation is more favorable than this other one so you have to pick and choose... the grains are doing that, but it is not happening quickly just because you know, either one can happen.

In this very short time watching the model, Liz was able to understand and generate hypotheses about two essential ideas: triple points and logarithmic laws (the literature refers to these ideas as particularly hard to understand, e.g., Krause and Tasooji 2007). Rather than trying to assemble statements pulled from regular instruction, Liz departed from what she knew about other phenomena and what she was actually seeing in the model. Even without formally mathematizing the time dependency of grain growth, she

understood the reason for the triple point to be considered a “low-mobility” point in a microstructure. The central atom has two atoms (out of six) of each of the surrounding grains as neighbors, so the switch probability is the same (1/3), and there is no preferred growth direction. She also realized that the time law is not linear: growth speed decreases over time and eventually “tapers off.” Rather than *being told*, Liz arrived at this conclusion *on her own*, by drawing microstructures, changing variables and observing the dynamics of the simulation. Particularly, when asked about the fundamental reason for the “tapering off” of grain growth, she affirmed that “[...] because there isn’t a distinct orientation [which] is more favorable”—in other words, she was getting at the core of the explanation. This same idea could be useful to explain other phenomena in materials science, and we will see how students applied such generative ideas to other phenomena in the next section.

Similarly, Peter and Elise, notwithstanding their initial difficulties in explaining grain growth during the pre-interview, understood exactly the logarithmic nature of the grain growth law:

Interviewer: What could be the rule for grain growth speed?

Peter: As the grains get bigger, each grain is increasingly hard to take away from because it’s bigger, so the interfaces start to be between two large grains, instead of small grains, so an interface between a large grain and a large grain is less likely to have a lot of movement because both grains are large and they are already in a state where they don’t want to shrink.

Interviewer: what will happen to this surface [between two grains]?

Elise: [It’ll] shift up to be vertical. [looking at the simulation] Yes, it’s just getting flatter.

Interviewer: Why do you think it wants to get flat?

Elise: It’s like the nearest-neighbor thing, these want the most nearest green neighbors, the red ones want the most red ones.

Interviewer: [some minutes later, she is looking a triple point] What’s happening here?

Elise: It has the same number around each other, so, the red, the angles are all at equilibrium, they are all a stable formation.

Interviewer: And what’s that angle there?

Elise: It’s a hexagon, I guess it’s 360 divided by three, 120.

Generally, most students knew that the small red grain in Fig. 6 was going to disappear. From their reactions while observing the simulation, they seemed to be expecting a unidirectional animation of the grain being “eaten” by the surrounding ones, and eventually disappearing. This was consistent both with the heuristics and the types of results of aggregate tools, animations, and equations commonly seen in class—processes happen unidirectionally. However, what students observed was different: behaviors emerging from local interactions, which take place with some degree of randomness. At times, the small grain would grow, but most of the times it would shrink. Some of the students wanted to slow down the simulation and use the “zoom” tool to see the process in more detail. But in doing that, students could only see the micro-level phenomenon (atoms flipping to different orientations). By zooming out again, they could observe the *emergent* behavior: curved surfaces disappearing as the Laplace-Young equation would predict. Thus, there is a qualitative difference between traditional learning tools and agent-based modeling: not only are students observing an expected outcome, but *they are able to see the process unfolding* at various levels. Not only is the simulation visually similar to the phenomenon, but also its algorithm loyally emulates the micro-level processes underlying it. This is different from purely numeric simulations in which students are able to compare only outputs, and not the

processes as they unfold. In addition, words commonly used in the classroom, such as “shrink,” “consume,” and “growth” acquired a new meaning. Those metaphorical terms, as our pre-interview data suggested, can mislead students to interpret literally their meaning—working with MaterialSim, students realized that grains were not actually being “consumed” or shrinking: atoms were just switching places, and the metaphors were just describing the net, aggregate effect of such behavior. This was a central element of the whole experience and, as we will observe, deepened as students progressed in the study.

The last activity of the first day was the “BehaviorSpace” experiment. This NetLogo feature enables users to automatically run hundreds of simulations, each under different parameter settings, sweeping entire parameter spaces. Students ran at least one set of experiments, plotted the data, and came up with theories to describe the phenomenon. Most students chose to model the influence of dispersed particles on grain growth. The textbook explanation of this phenomenon takes approximately four pages. It begins with an account of how a force P appears when a boundary attempts to go through a particle, and then calculates the drag force by means of geometrical approximations (Fig. 7).




Fig. 7 The textbook picture explaining how dispersed particles affect boundary migration (Porter and Easterling 1992, p. 141)

Departing from those geometrical approximations, the formula is obtained with a series of derivations (Porter and Easterling 1992, p. 141), which relates the fraction of dispersed particles (f), the mean radius of the particles (r), and the maximum particle size after grain growth (D_{\max}):

$$P = \frac{3f}{2\pi r^2} \cdot \pi r \gamma = \frac{3f\gamma}{2r} \Rightarrow \frac{2\gamma}{D} = \frac{3f\gamma}{2r} \Rightarrow \bar{D}_{\max} = \frac{4r}{3f}$$

However, the NetLogo algorithm is not based on this formula, and Bob was somewhat skeptical that the ‘official’ equation could be matched by the simulation. Thus he programmed NetLogo to run a series of simulations with percentages of particles varying from 0 to 8% (see screenshots and individual plots of grain growth speed in Fig. 8). He also constructed a plot to aggregate the results across all experiments, and subsequently tried to compare their own curve with the theoretical data (dotted line in Fig. 8’s lower plot). To his surprise, the two curves had a very reasonable match. Other students, with slight variations, undertook the same project, or selected different aspects of the phenomenon.

By exploring entire parameters spaces, and having not only the dynamic visualization but also actual numerical data to base their explanations on, these students were able to further generate original hypotheses and find out correlations.

6.3 Second Session: Building Their Own Models

The second session was the model construction part of the study. Students had two and a half hours to learn the basics of the NetLogo language and program a new model of a materials science phenomenon. For this session, which took place 2 or 3 days after the first session, students were asked to bring one idea of their own for a new model of a materials science phenomenon. They pursued questions of their interest or problems that they did not understand during regular instruction. By authoring new models or new features for the

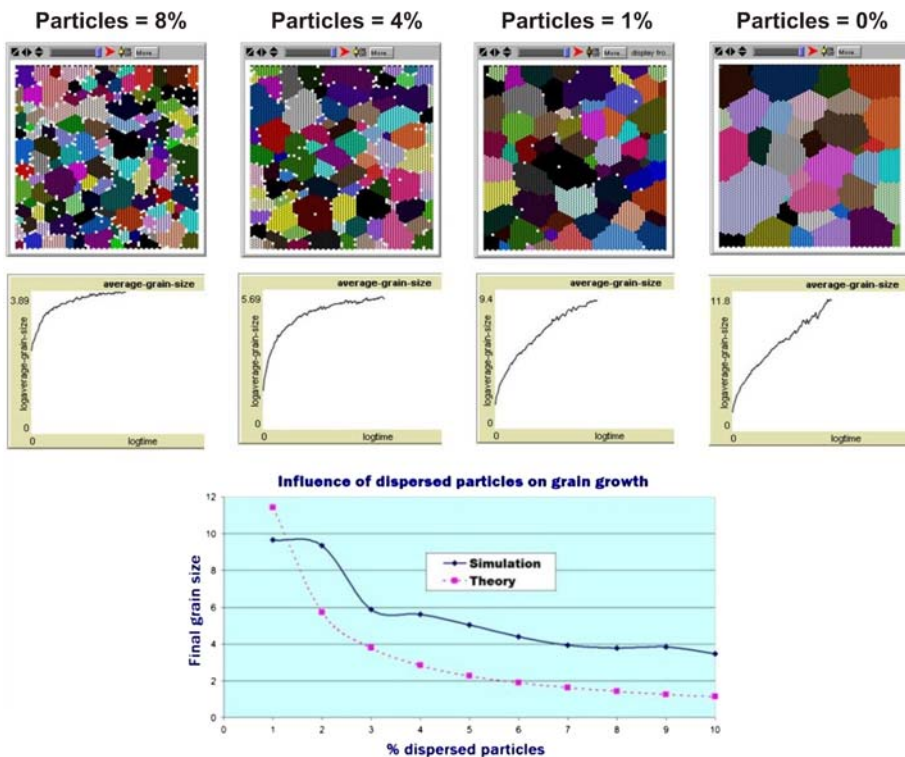


Fig. 8 Sequence of screenshots from Bob’s experiment

existing models, they could elaborate on answers to their research questions. Student achievement was impressive. A comparison between the pre-interview data, when students relied on ready-made statements about the phenomena, and their performance on the last day of the study, when they built their own models relying just on fundamental thermodynamics, suggests that student contact with an agent-based environment effected conceptual gain. Even more than exploring the existing models, constructing their own models was a transformative experience for most. In this section, we will narrate and analyze some of those learning trajectories. The models chosen for this analysis represent typical student work, and the particular choice of which students to include in the data analysis attempted to provide representative examples of the various affordances of ABM employed by students.

6.3.1 Betty's Model

Betty built a model to investigate grain growth with a new and important feature: taking into consideration the misalignment between grains (Fig. 9). In her innovative model, the more misalignment across the boundary, the harder it would be for an atom to jump from one grain to another. The construction of this model presented Betty with many challenges. The first was to convert the grain orientation's angle, which could lie in any of the four quadrants, to a normalized quadrant-independent measure. Betty's solution, after much thinking, sketching and trying out different trigonometric functions, was to use the *arcsine* function. The following picture shows some of her reasoning. From her drawing, we can observe that she was using geometrical analysis from a "micro" level, examining individual atoms, and trying to design an algorithm to account for the trigonometric issues in calculating their misorientation.

She considered that the probability for an atom to jump to the next grain should be dependent not only on the number of different atoms around it, but also on the average misorientation between the two grains. Low misorientation would promote easier migration. Apart from the initial difficulty in figuring out the best trigonometric function for the angle comparison, Betty knew what she needed to do, without resorting to any of the textbook formulas. For her, at the micro-level, adding the misorientation effect was very




Fig. 9 Betty's sketches about angles, sine and arcsine

easy.⁵ She simply added one command to the original grain growth model to implement her change—previously, in the traditional equation-based materials science, making such a change would require long and mathematically demanding derivations. The resulting code of her misorientation calculating function was:

```
;;calculates the absolute value of the arcsin
to-report misorientation [angle]
  report asin (abs (sin (angle)))
end
;;calculates the absolute value of the sum of the two arcsins
to-report calculate-misorientation [angle1 angle2]
  report abs (misorientation (angle1) + misorientation (angle2))
end

;;reports the average misorientation for a given atom
to-report compare-misorientation
  let i 0

  ask neighbors6
  [
    ;;calculates the misorientation between the current atom and each of its 6 neighbors
    set total-misorientation (total-misorientation + calculate-misorientation heading (heading-
of neighbors6))
    set i i + 1 ;update counter
  ]

;;returns the average misorientation
report (total-misorientation / i)

end
```

Then, having her reporters calculate how much each given atom would differ from its neighbors angle-wise, she changed the original grain growth procedure, adding one extra simple “if” statement:

```
;;OLD PROCEDURE
  if future-free-energy <= present-free-energy
    [set heading (future-heading)]

;;BETTY'S NEW PROCEDURE
  if future-free-energy <= present-free-energy
    [
      if (present-heading - ([heading] of one-of neighbors6) < misorientation)
        [set heading (future-heading)]
    ]
```

Yet, aggregate and macroscopic models do not afford such insights as well. The agent-based approach, conversely, provided a “low-threshold” entry point for Betty to implement her ideas by constructing models. Her model was very consistent with known theory, even though she was not cognizant of this theory prior to the interventional study. Betty’s model illustrates one of the main affordances of the agent-based representation: at the micro level, the mathematical machinery required to add new phenomena or parameters to an existing algorithm is much simpler than in traditional representations. Instead of employing numerous equations to add her misorientation effect, just one line of code, at the micro-level, was enough.

6.3.2 Bob’s Model

Bob wanted to include a new parameter in the grain growth model: the size of the dispersed solid particles. The idea was to enable users not only to change the percentage of particles,

⁵ On a more advanced level, similar research was undertaken and published by researchers, such as Ono et al. (1999).




Fig. 10 Bob's attempts to understand how to calculate the particles' areas, first "hard-coded" in NetLogo (*top* section), then figuring out a formula to accomplish it (*middle* section), and finally implementing the code in his model (*bottom* section)

but also their radius. Bob realized that this seemingly simple change was in fact challenging. Given a certain percentage (in mass) of particles, their number had to be adjusted to compensate for the increased mass of each particle. That involved the calculation of the area of each particle (in a hexagonal grid) and the total area of the sample, to determine how many hexagon "seeds" would be necessary for a specific percentage. The first problem involved the conception of an algorithm for calculating the area of polygons placed in the hexagonal grid, which turned out to be an rather elaborate mathematical exercise: particles with different radii, in a hexagonal grid, have different shapes (Fig. 10). Bob realized that a recursive procedure would be adequate; a radius increase by one unit would simply add one layer of near-neighbor atoms to the existing particle.

After completing the model (Fig. 11), he investigated the influence of particle size on grain growth: maintaining the same percentage in mass, how is growth affected by changing each individual precipitate's size? He concluded that the granularity of the particles has a strong influence on the final grain size—in other words, having fine or coarse particles (even with the same overall amount in mass) would have a dramatic effect on the final grain size. He was able to run large batches of simulations in BehaviorSpace, chart the data, and explore possible explanations. By constructing his own model, he went even further than his previous experiment, on the first session, with different percentages of particles (Fig. 8). Having previously addressed "*f*", i.e., the percentage of particles, Bob now fully re-created the $\bar{D}_{\max} = \frac{4r}{3f}$ equation with a new model which enabled him to change the radius "*r*" of the particles.⁶

6.3.3 Anand's Model

Anand did not want to investigate anything related to grain growth, and built a completely different model from scratch. His idea was to explore in detail interfacial energies due to

⁶ On a more advanced level, similar research was undertaken and published by many researchers, such as Gao et al. (1997) and Hassold and Srolovitz (1995).

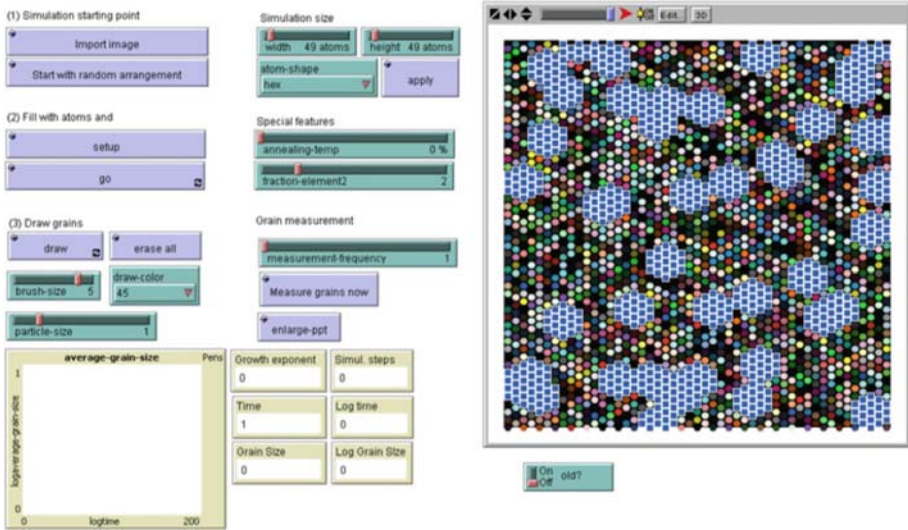


Fig. 11 Bob’s model for particles with varying size

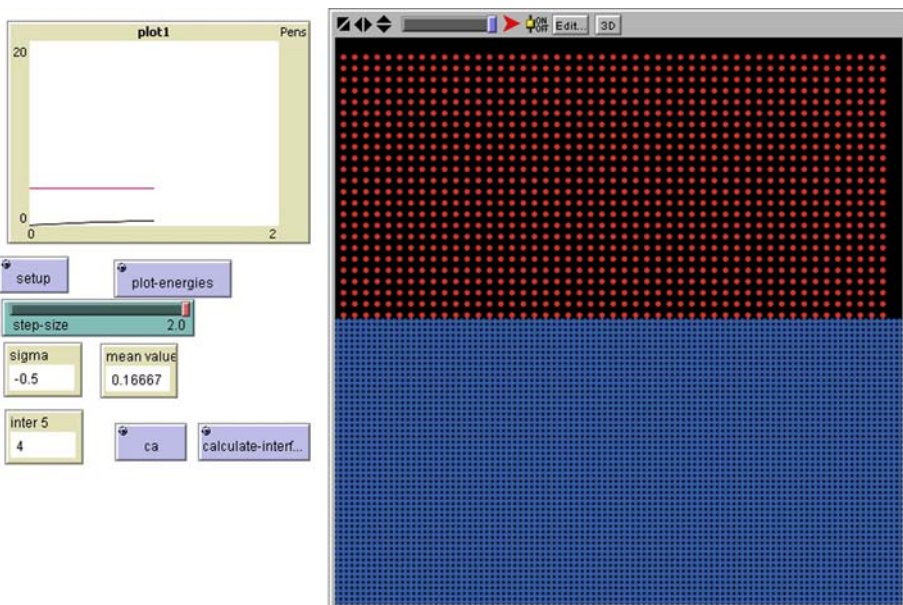


Fig. 12 Anand’s model for detailed study of interfacial energies

atomic misalignment. In other words, his model meant to determine how much energy was stored in an interface between two grains that did not match in terms of atomic spacing. We can observe his model in Fig. 12: one of the grains (in red), has a very large spacing between the atoms compared to the “blue” grain. This causes strain in the interface

between the two. Anand built this model using the same basic algorithm as in the grain growth model: atoms look around and check the state of their neighbors, deciding what to do based on the number of equal and different neighbors. Even though this topic, in the regular curriculum, was completely separate from grain growth, he was able to identify a transferable micro-level model between the two phenomena. He calculated the final energy of each atom with a simple command that counted the number of connected atoms within a certain distance.

```

to calculate-interfacial-energy
  ask element1
  [
    set interfacial-energy (count turtles in-radius 2 with [color = blue])
  ]
end

```

Anand's model is an example of an important affordance of ABM: *one-to-many generativity*. After understanding in detail the principles and algorithm behind the grain growth model, he was able to identify other opportunities to employ the same principles to model and understand other seemingly unrelated phenomena.

6.3.4 Peter's Model

Peter's model was another example of the one-to-many, transferable affordance of the agent-based representation. In the pre-survey, he identified diffusion control as one of the hardest topics in the course. In the second session, that was exactly his choice for building a model. He started it from scratch, and in less than two dozen lines of code and 2 h, had a model the complexity of which is far beyond what is expected from the Microstructural Dynamics course, considering the classroom observations and analysis of class materials. Peter used the agent-based approach and some micro-rules adapted from the grain growth model to create a diffusion model. Even though the two phenomena have differences, he managed to identify the common micro-rules and copy them from one model to the other, with the necessary adaptations, in the same way as Anand and his Interfacial Energy model.

In materials science, it is particularly important to tell apart transformations that are interface-controlled (i.e., the slowest phase happens at the interface of the two materials) from diffusion-controlled (the slowest part is for atoms to "travel" to the interface, where the actual reaction is fast). Knowing the slowest phase of a given process, engineers can greatly optimize it. Peter's purpose was to build a model to investigate this phenomenon. Its textbook explanation is a 5-page narrative with five different equations, which are put together to show that the driving force (referred to as $\Delta\mu_B^i$) is:

$$\Delta\mu_B^i = RT \ln \frac{X_i}{X_e} \cong \frac{RT}{X_e} (X_i - X_e)$$

where X_i and X_e are the chemical compositions of the two materials (Porter and Easterling 1992, p. 177).

Again, Peter ignored the existence of this long sequence of equations. He concentrated on the micro-rules concerning the phenomenon, and realized that the rules were not very different from other models. After all, "liquid" atoms were just moving randomly and "bumping" into a solid, sticking to it according to a certain probability. Just as Bob did, Peter was concerned with the correctness of his work. He generated a series of plots and screenshots so as to match his data with the textbook plots, part of which are shown in Fig. 13.

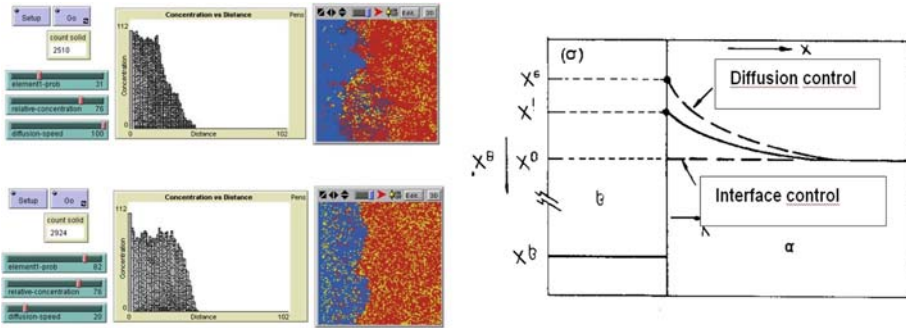


Fig. 13 Results of Peter’s model with diffusion control (*top*, with diffusion speed = 100), interface control (*bottom*, with diffusion speed = 20), and the chart from the textbook, where we can identify a similar same shape for the two concentration curves (note that this last chart was rotated for clarity)

Peter’s algorithm was fairly straightforward: if the atoms are in the liquid, and they bump into a solid, they become solid (with a certain probability, dependent on their chemical properties), hence:

```
if ((breed = element1-liquid) and ;;if you are an atom in the liquid
    (neighbor-breed = solid) and ;; and one neighbor of yours is in the solid
    (element1-probability > random-float 100)) ;; and depending on your diffusion speed
[
    set color neighbor-color ;; switch the atom's color
    set breed neighbor-breed ;; switch the atom's breed
]
```

If the atom is in the liquid (breed different than solid, or “!=solid” in NetLogo language), and it meets an atom different than itself, the atoms switch places—in other words, diffusion is taking place:

```
if ((breed != solid) and ;;if you are in the liquid
    (neighbor-breed != solid) and ;;and one neighbor of yours is also in the liquid
    (diffusion-speed > random-float 100)) ;; and depending on your diffusion speed
[
    set [color] of random-neighbor color ;;switch the neighbor's color
    set [breed] of random-neighbor breed ;;switch the neighbor's breeds
    set color neighbor-color ;; switch the atom's color
    set breed neighbor-breed ;; switch the atom's breed
]
```

In a dozen lines of code, Peter was able to model both diffusion and solidification, manipulating exclusively local rules. At the end of the session, the first author asked Peter about the next steps in his model’s implementation, had he more time to work on it. Again, he demonstrated a solid understanding of how to easily manipulate those basic rules to generate new models:

Peter: I did a liquid to solid model, now I want to be able to invert it, do a solid to liquid algorithm.

Interviewer: And how would you implement it?

Peter: It’s simple: I’ll just invert the probability. It’s just the opposite probability. I don’t have to change much.

[...]

Interviewer: And how would you implement, for example, dispersed particles in a grain growth model?

Peter: I would put in molecules that have zero probability to change to anything else, and zero probability of another molecule to change into them.

Peter's response demonstrated a deep understanding of the process and was in great contrast with his pre-interview data, in which although he correctly identified and explained some phenomena, he failed to see how those principles and knowledge could be put to use to further his own knowledge about a particular topic or other phenomena.

7 Discussion

Computer modeling is posing a serious challenge to extant knowledge encoding schemes in materials science. Researchers have already detected this trend—computer modeling in materials science has more than doubled in the last 10 years (Thornton and Asta 2005). However, materials science students are still obliged to master hundreds of equations and isolated facts. Even if students were to somehow try to connect those equations into a coherent corpus, the mathematical machinery required to accomplish that would be too demanding for most to succeed.

The examples of students' model-building we have described were implemented in less than 2 hours, *including* the time dedicated to learning the basics of the NetLogo language. The relative ease with which students developed their own models, even within such a short timeframe, shows that model building is an approachable task for undergraduate students and supports one of our main claims: agent-based modeling, for some fields of engineering, offers a more principled understanding of the natural phenomena, which, in turn, grants more autonomy for students in learning new content or deriving new theories on their own. Participant students had previous knowledge of the phenomenon from their class work. Nevertheless, during the pre-interview, they demonstrated difficulty in explaining related phenomena in a coherent fashion, resorting to a range of fragmented models and metaphors. The implementation of their own model within an agent-based modeling environment provided students with fewer, simpler rules that were closely related to the physical phenomenon, thus enabling them to better understand and extend the model by adding new proximal rules for the agents.

We compiled evidence suggesting that the agent-based encoding is a good fit to content in materials science. First, the undergraduate courses are overloaded with highly specialized information. Secondly, students demonstrated difficulty in explaining even the most basic concepts in the field, with frequent 'slippage' between levels. And thirdly, throughout the classrooms observations and the interviews, one striking revelation was that the agent-based approach was not a total unknown for textbook authors, teachers and students. The textbook oftentimes makes use of micro-behaviors, simple rules, and agent-based heuristics. For example, when explaining the process of grain growth, the textbook authors use an agent-based approach:

[...] A similar effect occurs in metal grains. In this case the atoms in the shrinking grain detach themselves from the lattice on the high pressure side of the boundary and relocate themselves on a lattice site on the growing grain (Porter and Easterling 1992)

However, the agent-based representation was in this context a mere illustration of the "real" content that would come after, encoded as equations. Arguably, even though the agent-based representations could be easier for students to understand, there was no

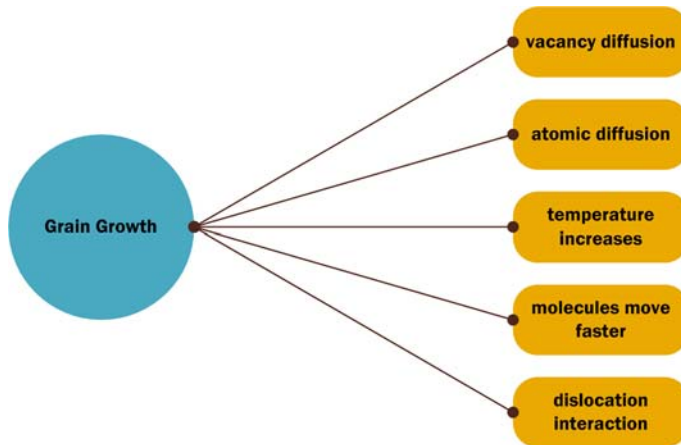


Fig. 14 Bob's one-level explanation

technological infrastructure to “run” those models—the activities and software that we developed could provide this infrastructure. The availability of an expressive tool and an empowering learning environment were crucial elements. As a computational tool, Net-Logo and its agent-based approach was a good fit for capturing students’ intuitions and ideas at the appropriate level. In addition, the constructionist nature of students’ interaction with the tool enabled them to build fluency with this new tool, and perceive themselves as scientists in their own right, transforming seemingly simple ideas and local rules into powerful kernels for scientific modeling. To further understand the cognitive model which the ABM perspective might foster, let us consider again, for example, Bob’s explanation of grain growth:

Bob: Well, grains grow through diffusion, through *vacancy diffusion*, and *atomic diffusion*, for one, it is all over the place, *temperature increases*, *molecules move around faster* [...].

His statement reveals a one-level description of the phenomena, which is compatible with our analysis of the current sparse and linear materials science structuration. Ideas such as “vacancy diffusion” and “increase of temperature” are connected to “grain growth” without a clear hierarchy (Fig. 14).

During the work with MaterialSim, students developed an extra ‘organizing’ layer which grouped some of the surface manifestations on the phenomena under one unifying principle⁷ (Fig. 15). Let us observe Liz’s statement:

Liz: It is because, it wants to be more *energetically stable*, or have *less energy* in the crystal, so it will grow, just to form one big grain, because that’s the least energy configuration [...]

⁷ For elaboration on the idea of organizing layers, see “Papert’s principle” in Minsky’s *Society of Mind* (1986).

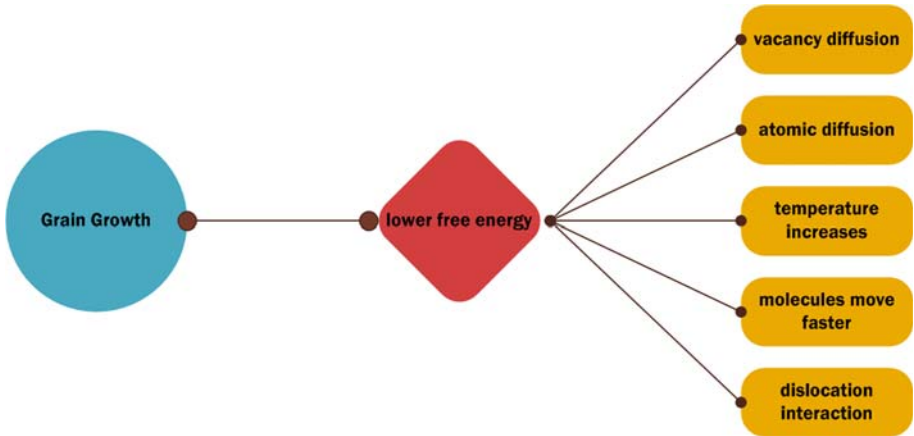


Fig. 15 Liz's two-level structure

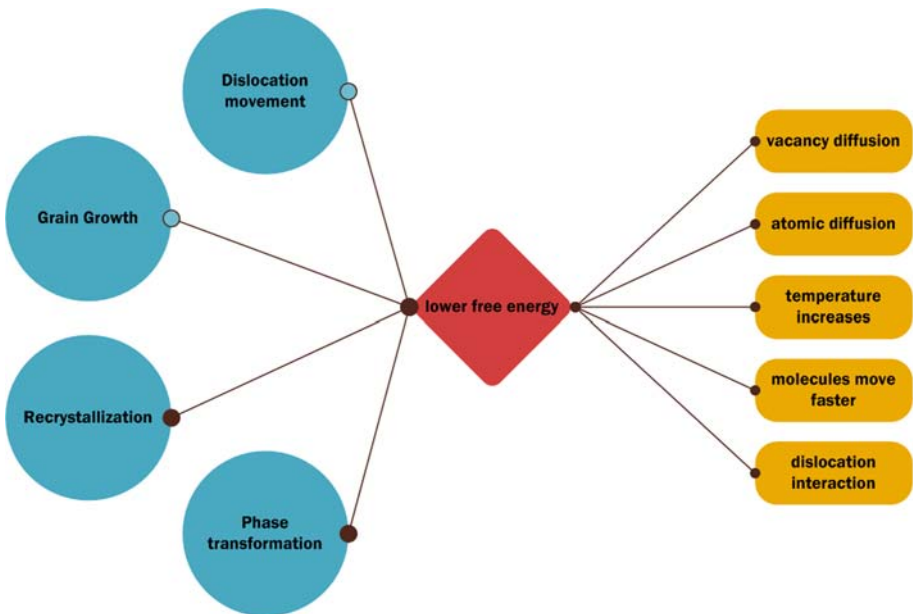


Fig. 16 A two-level structure with multiple phenomena

Liz identified one unifying principle, 'lowering free energy', from which many of those external manifestations derive. An agent-based modeling environment offers low-threshold tools to code and formalize algorithmically this principle, enabling her to 'mobilize' this idea that was previously just a vague mental model. Finally, after the model building, students were able to mobilize these generalizable principles, encoded as computer code, to explain other phenomena that shared the same mechanism (Fig. 16).

8 Conclusion

Design is now fashionable in many engineering schools. Robotics competitions, for instance, are common in various universities. Could we extend the powerful ideas about mechanical construction, engineering design, and projects-based approaches to fields such as materials science, which products are quite different from robots?

Rich, motivating learning is often achieved through an approach of *learning-by-doing*. In areas such as mechanical engineering *doing* and *understanding* could be tightly connected. When students are building a gearing system, all the components are clearly visible. In areas such as chemistry, atmospheric science, biology, and materials science, that is not the case. Learners might observe effects while having little understanding of the underlying causality, as the actual phenomenon is too removed from human size or time scale. Moreover, teaching tools in those disciplines often have relied on “aggregate,” formulaic descriptions. Our studies suggested that the fragmentation and opaqueness of such descriptions could constitute an obstacle to learning. Firstly, the traditional equational descriptions are more context-specific, and do not enable students to make broader inferences about phenomena with similar rules. The mathematical machinery required to manipulate and combine equations is a high threshold. Secondly, these descriptions often lead to heuristics that generate over-generalizations. Students had memorized ideas about phenomena in materials science about which they have no deep understanding. Thirdly, the traditional descriptions often background the actual physical phenomena.

On the other hand, agent-based modeling seems to be a better fit for the content areas discussed in this paper, for three reasons:

1. MaterialSim foregrounded the fundamental physical processes in the material, namely atomic movement and free-energy minimization. Not only the algorithm was exclusively based on those processes, but also the visualization scheme enabled students to *see* them unfolding in real-time. Students observed both favorable and unfavorable atomic flips, grains growing and shrinking, expected and unexpected results. Our data suggests that the observation of those processes was important for student understanding.
2. A core feature of this design is that students can apply a small number of *transferable*, generative models to capture fundamental causal structures underlying behaviors in a range of apparently disparate phenomena within a domain. For example, a free-energy minimization model could enable students to understand not only grain growth, but a wide variety of related phenomena (annealing, interfacial energy, recrystallization, diffusion, phase transformations), which are traditionally taught as *separate* topics with their own models and equations. Most students were able to create their own models by transferring some “kernels” from one model to another.
3. One of the pillars of constructionist theory is the importance of students conducting personally meaningful projects. Our data suggest that coding their *own* models was a particularly valuable learning experience: with a more intense exposure to the tools and methods of agent-based modeling, students had an opportunity to create fluency with the computational representations, by testing and debugging their theories, and reconciling them with previous knowledge—just as a scientist would do.

However, we caution that not all modeling tools are created equal, and thus stress the importance of designing software that foregrounds key transferable, generative models, as to enable students to extend their knowledge into new phenomena. The study, in conclusion, suggests that *less is more*, the knowledge and the exploration of just a few simple

underlying rules of natural phenomena, given the availability of a computational medium to manipulate them, appear to be more generative for students than the aggregate, equation-based representations.

Acknowledgments We thank Prof. Tschiptschin (Engineering School of the University of São Paulo, Brazil), Prof. Srolovitz (Yeshiva University), the members of the Center for Connected Learning and Computer-Based Modeling at Northwestern University, in particular, Spiro Maroulis and Michelle Wilkerson-Jerde, and Prof. Dor Abrahamson (UC Berkeley) for valuable comments on drafts of this paper. We also thank Ann McKenna from the McCormick School of Engineering at Northwestern University. This work was funded under the umbrella of three NSF grants: NSF grants REC-0126227, HCC-0713619 and NSF EEC-0648316.

References

- Abrahamson, D., & Wilensky, U. (2005). ProbLab goes to school: Design, teaching, and learning of probability with multi-agent interactive computer models. *Paper presented at the 4th Conference of the European Society for Research in Mathematics Education*.
- American Society for Engineering Education. (2007). ASEE profiles of engineering and engineering technology colleges. From <http://www.asee.org/publications/profiles/upload/2007ProfileEng.pdf>.
- Anderson, M. P., Srolovitz, D. J., Grest, G. S., & Sahni, P. S. (1984). Computer simulation of grain growth: I. Kinetics. *Acta Metallurgica* (32), 783–792. doi:10.1016/0001-6160(84)90151-2.
- Berland, M., & Wilensky, U. (2004). Virtual robotics in a collaborative constructionist learning environment. *Paper presented at the Annual Meeting of the American Educational Research Association*, San Diego, CA, April 12–16.
- Berland, M., & Wilensky, U. (2005). Complex play systems—Results from a classroom implementation of Vbot. In W. Stroup & U. Wilensky (Chairs) & C. D. Lee (Discussant), *Patterns in group learning with next-generation network technology. Paper presented at the Annual Meeting of the American Educational Research Association*, Montreal, Canada, April 11–15.
- Blikstein, P. (2009). *An atom is known by the company it keeps: Content, representation and pedagogy within the epistemic revolution of the complexity sciences*. Unpublished Doctoral Dissertation, Northwestern University, Evanston, IL.
- Blikstein, P., & Tschiptschin, A. P. (1999). Monte Carlo simulation of grain growth. *Materials Research*, 2(3), 133–137. doi:10.1590/S1516-14391999000300004.
- Blikstein, P., & Wilensky, U. (2004a). MaterialSim: An agent-based simulation toolkit for learning materials science. *Proceedings of the International Conference on Engineering Education*, Gainesville, FL.
- Blikstein, P., & Wilensky, U. (2004b). *MaterialSim curriculum*. Center for Connected Learning and Computer Based Modeling, Northwestern University, Evanston, IL. Retrieved from <http://ccl.northwestern.edu/curriculum/materialsim>.
- Blikstein, P., & Wilensky, U. (2005). Less is more: Agent-based simulation as a powerful learning tool in materials science. *4th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2005)*. Utrecht, Holland.
- Blikstein, P., & Wilensky, U. (2006a). From inert to generative modeling: Two case studies of multiagent-based simulation in undergraduate engineering education. *Paper presented at the Annual Meeting of the American Educational Research Association*, San Francisco, CA.
- Blikstein, P., & Wilensky, U. (2006b). ‘Hybrid modeling’: Advanced scientific investigation linking computer models and real-world sensing. *7th International Conference of the Learning Sciences*. Bloomington, USA.
- Blikstein, P., & Wilensky, U. (2008). Implementing multi-agent modeling in the classroom: Lessons from empirical studies in undergraduate engineering education. In M. Jacobson (Chair) & R. Noss (Discussant), *Complex Systems & Learning: Empirical Research, Issues & ‘Seeing’ Scientific Knowledge With New Eyes*. In G. Kanselaar, J. van Merri’nboer, P. Kirschner, & T. de Jong (Eds.). *Proceedings of the International Conference of the Learning Sciences* (Vol. 3, pp. 266–267) (ICLS).
- Bransford, J. D., & Schwartz, D. L. (1999). Rethinking transfer: A simple proposal with multiple implications. *Review of Research in Education*, 24, 61–100.
- Brown, G. S. (1961). Today’s dilemma in engineering education. *Ire Transactions on Education. E (Norwalk, Conn.)*, 4(2), 48.

- Burke, J. E. (1949). Some factors affecting the rate of grain growth in metals. *Transactions of the American Institute of Mining Engineers* (180), 73–91.
- Centola, D., Wilensky, U., & McKenzie, E. (2000). A hands-on modeling approach to evolution: Learning about the evolution of cooperation and altruism through multi-agent modeling—the EACH project. *Proceedings of the 4th Annual International Conference of the Learning Sciences*. Ann Arbor, MI, June 14–17.
- Colgate, J. E., McKenna, A., & Ankenman, B. (2004). IDEA: Implementing design throughout the curriculum at northwestern. *International Journal of Engineering Education*, 20(3), 405–411.
- Collier, N., & Sallach, D. (2001). *Repast*. University of Chicago. <http://repast.sourceforge.net>.
- Committee on the Education and Utilization of the Engineering. (1985). *Engineering education and practice in the United States: Foundations of our techno-economic future*. Washington, DC: National Research Council.
- diSessa, A. A. (1993). Toward an epistemology of physics. *Cognition and Instruction*, 10(2 & 3), 105–225.
- diSessa, A. A. (2000). *Changing minds: Computers, learning, and literacy*. Cambridge: MIT.
- Dym, C. L. (1999). Learning engineering: Design, languages, and experiences. *Journal of Engineering Education*, 145–148.
- Dym, C. L., Agogino, A. M., Eris, O., Frey, D. D., & Leifer, L. J. (2005). Engineering design thinking, teaching, and learning. *Journal of Engineering Education*, 94(1), 103–120.
- Edwards, L. (1995). Microworlds as Representations. *Proceedings of the 2nd International NATO Symposium on Advanced Technology and Education*.
- Einstein, H. H. (2002). Engineering change at MIT. *Civil Engineering*, 72(10), 62–69.
- Gao, J., Thompson, R. G., & Patterson, B. R. (1997). Computer simulation of grain growth with second phase particle pinning. *Acta Materialia*, 45(9), 3653–3658.
- Gobert, J., Horwitz, P., Tinker, B., Buckley, B., Wilensky, U., Levy, S., et al. (2004). Modeling across the curriculum: Scaling up modeling using technology. *Paper presented at the Proceedings of the Twenty-fifth Annual Meeting of the Cognitive Science Society*. July 31–August 2, Boston, MA.
- Goldstone, R. L., & Wilensky, U. (2008). Promoting transfer through complex systems principles. *Journal of the Learning Sciences*, 15, 35–43.
- Haghighi, K. (2005). Systematic and sustainable reform in engineering education. *Journal of Environmental Engineering—ASCE*, 131(4), 501–502.
- Hassold, G. N., & Srolovitz, D. J. (1995). Computer simulation of grain growth with mobile particles. *Scripta Metallurgica et Materialia*, 32(10), 1541–1547.
- Hatano, G., & Oura, Y. (2003). Commentary: Reconceptualizing school learning using insight from expertise research. *Educational Researcher*, 32(8), 26–29.
- Horwitz, P., Gobert, J., Wilensky, U., & Dede, C. (2003). MAC: A longitudinal study of modeling technology in science classrooms. *Paper presented at the National Educational Computing Conference (NECC)*, Seattle, WA.
- Hurst, K. D. (1995). A new paradigm for engineering education. *Paper presented at the ASEE Frontiers in Education*, Atlanta, GA.
- Jerath, S. (1983). Engineering-education in perspective. *Mechanical Engineering*, 105(2), 92–93.
- Katehi, L., Banks, K., Diefes-Dux, H., Follman, D., Gaunt, J., Haghighi, K., et al. (2004). A new framework for academic reform in engineering education. *Paper presented at the American Society for Engineering Education Conference*, Salt Lake City, UT.
- Krause, S. J., & Tasooji, A. (2007). Diagnosing students' misconceptions on solubility and saturation for understanding of phase diagrams. *American Society for Engineering Education Annual Conference*. Portland, USA.
- Kulov, N. N., & Slin'ko, M. G. (2004). The state of the art in chemical engineering science and education. *Theoretical Foundations of Chemical Engineering*, 38(2), 105–111.
- Lamley, R. A. (1996). A model for an engineering education. *Marine Technology and Sname News*, 33(2), 119–121.
- Levy, S. T., Kim, H., & Wilensky, U. (2004). Connected chemistry—A study of secondary students using agent-based models to learn chemistry. *Paper presented at the Annual meeting of the American Educational Research Association*, San Diego, CA, April 12–16.
- Levy, S. T., Novak, M., & Wilensky, U. (2006). Students' foraging through the complexities of the particulate world: Scaffolding for independent inquiry in the connected chemistry (MAC) curriculum. *Paper presented at the Annual Meeting of the American Educational Research Association*, San Francisco, CA.
- Martin, F. (1996). Ideal and real systems: A study of notions of control in undergraduates who design robots. In Y. Kafai & M. Resnick (Eds.), *Constructionism in practice* (pp. 297–332). Mahwah, NJ: Lawrence Erlbaum.

- Minsky, M. (1986). *The society of mind*. New York: Simon & Schuster.
- MIT Center for Policy Alternatives. (1975). *Future directions in engineering-education—System response to a changing world—a report*. Washington, DC: American Society for Engineering Education.
- Newstetter, W. C., & McCracken, M. (2000). Design learning as conceptual change: A framework for developing a science of design learning. *Conference of the American Society of Engineering Education*. St. Louis, MS: ASEE.
- Ono, N., Kimura, K., & Watanabe, T. (1999). Monte Carlo simulation of grain growth with the full spectra of grain orientation and grain boundary energy. *Acta Materialia*, 47(3), 1007–1017.
- Panel on Undergraduate Engineering Education. (1986). *Engineering undergraduate education*. Washington, DC: National Research Council.
- Papert, S. (1980). *Mindstorms: Children, computers, and powerful ideas*. New York: Basic.
- Papert, S. (1991). Situating constructionism. In S. Papert & I. Harel (Eds.), *Constructionism*. Cambridge, MA: MIT.
- Parunak, H. V. D., Savit, R., & Riolo, R. L. (1998). Agent-based modeling vs equation-based modeling: A case study and users' guide. *Lecture Notes in Computer Science*, 1534, 10–25.
- Porter, D. A., & Easterling, K. E. (1992). *Phase transformations in metals and alloys* (2nd ed.). London: Chapman & Hall.
- Roco, M. C. (2002). Nanotechnology—A frontier for engineering education. *International Journal of Engineering Education*, 18(5), 488–497.
- Russell, J. S., & Stouffer, W. B. (2005). Survey of the national civil engineering curriculum. *Journal of Professional Issues in Engineering Education and Practice*, 131(2), 118–128.
- Sengupta, P., & Wilensky, U. (2008). On the learnability of electricity as a complex system. In M. Jacobson (Chair) & R. Noss (Discussant), *Complex Systems & Learning: Empirical Research, Issues & "Seeing" Scientific Knowledge With New Eyes*. In G. Kanselaar, J. van Merri'nborg, P. Kirschner, & T. de Jong (Eds.). *Proceedings of the International Conference of the Learning Sciences (ICLS)* (Vol. 3, pp. 267–268). Utrecht, The Netherlands.
- Sherin, B. L. (2001). A comparison of programming languages and algebraic notation as expressive languages for physics. *International Journal of Computers for Mathematical Learning*, 6(1), 1–61.
- Siegler, R. S., & Crowley, K. (1991). The microgenetic method: A direct means for studying cognitive development. *American Psychologist*, 46(6), 606–620.
- Soloway, E., Guzdial, M., & Hay, K. E. (1994). Learner-centered design: The challenge for HCI in the 21st century. *Interactions*, 1(2), 36–48.
- Srolovitz, D. J., Anderson, M. P., Sahni, P. S., & Grest, G. S. (1984). Computer simulation of grain growth—II. Grain size distribution, topology and local dynamics. *Acta Metallurgica*, 32(5), 793–801.
- Stieff, M., & Wilensky, U. (2003). Connected chemistry: Incorporating interactive simulations into the chemistry classroom. *Journal of Science Education and Technology*, 12(3), 285–302.
- Thornton, K., & Asta, M. (2005). Current status and outlook of computational materials science education in the US. *Modelling and Simulation in Materials Science and Engineering*, 13(2), R53.
- Tisue, S., & Wilensky, U. (2004). NetLogo: A simple environment for modeling complexity. *Paper presented at the International Conference on Complex Systems (ICCS 2004)*. Boston, MA, May 16–21.
- Tryggvason, G., & Apelian, D. (2006). Re-engineering engineering education for the challenges of the 21st century. *JOM*, 58(10), 14–17.
- Tyack, D. B., & Cuban, L. (1995). *Tinkering toward Utopia: A century of public school reform*. Cambridge, MA: Harvard University Press.
- Wilensky, U. (1993). *Connected mathematics: Building concrete relationships with mathematical knowledge*. Unpublished Doctoral Dissertation, MIT, Cambridge, MA.
- Wilensky, U. (1995). Paradox, programming and learning probability. *Journal of Mathematical Behavior*, 14(2), 231–280.
- Wilensky, U. (1999a). NetLogo. [computer software]. Evanston, IL: Center for Connected Learning and Computer-Based Modeling. <http://ccl.northwestern.edu/netlogo>.
- Wilensky, U. (1999b). GasLab-an extensible modeling toolkit for exploring micro-and-macro-views of gases. In N. Roberts, W. Feurzeig, & B. Hunter (Eds.), *Computer modeling and simulation in science education*. Berlin: Springer.
- Wilensky, U. (2003). Statistical mechanics for secondary school: The GasLab modeling toolkit. *International Journal of Computers for Mathematical Learning*, 8(1), 1–41. (special issue on agent-based modeling).
- Wilensky, U. (2006). Complex systems and reformulating scientific disciplines: Implications for learning, analysis of social systems, and educational policy. *Paper presented at the American Educational Research Association*, San Francisco, CA.

- Wilensky, U., Hazzard, E., & Longenecker, S. (2000). A bale of turtles: A case study of a middle school science class studying complexity using StarLogoT. *Paper presented at the meeting of the Spencer Foundation*. New York, NY, October 11–13.
- Wilensky, U., & Papert, S. (in preparation). *Restructurations: Reformulations of knowledge disciplines through new representational forms*. Working paper. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University.
- Wilensky, U., Papert, S., Sherin, B., diSessa, A., Kay, A., & Turkle, S. (2005). Center for Learning and Computation-Based Knowledge (CLiCK). Proposal to the National Science Foundation Science of Learning Centers.
- Wilensky, U., & Rand, W. (2009). *An introduction to agent-based modeling: Modeling natural, social and engineered complex systems with NetLogo*. Cambridge, MA: MIT.
- Wilensky, U., & Reisman, K. (2006). Thinking like a wolf, a sheep or a firefly: Learning biology through constructing and testing computational theories. *Cognition & Instruction*, 24(2), 171–209.
- Wilensky, U., & Resnick, M. (1999). Thinking in levels: A dynamic systems approach to making sense of the world. *Journal of Science Education and Technology*, 8(1), 3–19.
- Wilensky, U., & Shargel, B. (2002). *Behavior Space [Computer Software]*. Evanston, IL: Center for Connected Learning and Computer Based Modeling, Northwestern University.
- Wilkerson-Jerde, M., & Wilensky, U. (2009). Complementarity in equational and agent-based models: A pedagogical perspective. *Paper presented at the Annual Meeting of the American Research Education Association*, San Diego, CA.
- Wolfram, S. (2002). *A new kind of science*. Champaign, IL: Wolfram Media.
- Yamins, D. (2005). Towards a theory of “local to global” in distributed multi-agent systems (I). *Proceedings of the 4th International Joint Conference on Autonomous Agents and Multiagent Systems*. The Netherlands.