

Blikstein, P. and Wilensky, U. (2006, April). From inert to generative modeling: two case studies of Multi-Agent-Based Simulation in Undergraduate Engineering education. In D. Abrahamson (Org.), U. Wilensky (Chair), and M. Eisenberg (Discussant), *Small steps for agents... giant steps for students?: Learning with agent-based models*. Symposium conducted at the annual meeting of the American Educational Research Association, San Francisco, CA.

From inert to generative modeling: two case studies of Multi-Agent-Based Simulation in Undergraduate Engineering education

Paulo Blikstein (paulo@northwestern.edu)

Uri Wilensky (uri@northwestern.edu)

Center for Connected Learning and Computer-Based Modeling

School of Education and Social Policy

Northwestern University

1. Introduction

This paper reports on a user study of a computer-based learning environment for college-level Materials Science. MaterialSimTM (Blikstein & Wilensky, 2004, 2005) is an agent-based set of *microworlds* built by the authors within the NetLogo (Wilensky, 1999b) modeling environment. MaterialSim was created for investigating phenomena such as crystallization, solidification, grain growth and annealing. The design of MaterialSim emerged from extensive classroom observations followed by a literature review on engineering and Materials Science education, extensive classroom observations, 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 designed a user study for evaluating its effectiveness. A total of seventeen undergraduate students enrolled in an sophomore-level Materials Science course participated in the study in 2004 and 2005, which was comprised of a survey, pre-interview, interaction with the previously-programmed computer models, and students' construction of new models.

Our classrooms observations suggested that the ever-growing intricacy of college-level content in Materials Science is such that there is urgent need for new teaching and learning approaches. The reason is that the important equations and mathematical models in Materials Science, at this level of education, are often connected (in nontrivial ways) to multiple sets of other theories and other equations. As a result, a plain linear progression of equations, from simple to complex, is not sufficient. The result is that teachers have to resort to a multitude of models to explain a single phenomenon, each of them requiring a different set of equations, some even situated in a different locus of mathematical modeling (statistical mechanics and geometrical modeling, for example). Our classroom observations revealed that, in a typical 30-

minute period, students would be exposed to as many as 18 unique equations with 95 variables in total (not counting intermediate steps in a derivation) – which means approximately 1.5 minutes for each equation and 20 seconds for each variable. Our first research question was to investigate the consequences of this particular teaching strategy: what kind of understanding did this multiplicity of explanation levels and the “overloading” of equations foster in students?

As the data suggested that student understanding was problematic, we started to investigate different design frameworks to address the issue. We chose to utilize the *multi-agent-based* computer modeling approach (Collier, 2001; Wilensky, McKenzie, & Centola, 2000) within a Constructionist (Papert, 1980) framework. The multi-agent modeling approach, as we will explain in detail, enables modelers to depart from simple individual-level rules to generate complex collective behaviors. These simple rules capture fundamental causality structures underlying complex behaviors within a domain. Wilensky, Resnick and colleagues (Centola, 2000; Wilensky & Reisman, 2006; Wilensky & 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. We embed the multi-agent modeling approach with the Constructionist framework by designing materials and activities that enable students to explore multi-agent models and microworlds and then to choose an area of their interest and construct a model of a phenomenon in that area.

The rationale for using agent-based modeling is that this perspective may foster different, and perhaps more generative and extensible understanding of the relevant scientific phenomena. Instead of multiple models or numerous equations, this framework focuses on a small number of elementary behaviors which can be applied to a variety of scientific phenomena. Instead of a *many-to-one* relationship (many equations to one phenomenon), we attempt here a *one-to-many* (one set of local rules to many phenomena) relationship. In this approach, scientific phenomena are not “stand-alone” entities, disconnected one from the other, but emergent properties of the same set of local rules. Our second research question was: What kind of understanding do students develop of the Materials Science content when they study it from this agent-based, *one-to-many* perspective?

A third research question was to investigate the learning outcomes of students who build their own computer models. In other words, does coding (i.e., programming) multi-agent-based models generate deeper understanding of the scientific phenomena, as opposed to just interacting with ready-made models?

To answer our research questions, we present evidence in the form of excerpts and samples of students’ work, which demonstrates that the experience with MaterialSim enabled them to identify and understand some of the unifying principles in Materials Science and to use those principles to effectively construct new models.

2. A new scenario in engineering education

Engineering teaching and learning have not changed much in the past few decades. This conservatism, which has been pinpointed as early as the sixties (Brown, 1961; Jerath, 1983; MIT Center for Policy Alternatives, 1975) is still prevalent in numerous engineering schools. In recent years there have been calls for reform from the engineering education community and several

schools have proposed a number of reform initiatives (Einstein, 2002; Haghghi, 2005; Russell & Stouffer, 2005). A popular approach has been to introduce design-based curricula which include hands-on engineering design, oftentimes using modeling and simulation software (Colgate, McKenna, & Ankenman, 2004; Lamley, 1996; Martin, 1996; Newstetter & McCracken, 2000). Some of the reasons mentioned for this push for curricular reform in Engineering are the declining interest of American high-school students for an engineering career (down 18% since 1991), the changing workforce demographics, the new industrial dynamics brought about by “mass customization” and other new manufacturing needs and techniques (Katehi et al., 2004).

Apart from the broader societal changes, technical advances have also been an important driving force for engineering education reform programs. As basic science and engineering become increasingly intertwined in fields such as nanotechnology, molecular self-organization, 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 & Slin'ko, 2004). The 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 (Hatano & Oura, 2003) which they can apply to new problems and situations.

Many researchers and industrial leaders in the field have been pointing out that reform initiatives are falling short of those new challenges. Bazzo (1998) and Blikstein (2001) criticized the “banking” (Freire, 1974) approach to curricular reform – simply adding new courses to the curriculum but not making structural changes to it. Hurst (1995) pointed out that syllabi and curricula were so overloaded with transient or excessively detailed knowledge that there was no time for fostering students’ fundamental understanding of content matter. In fact, most of those reform initiatives were based on grand views about “blocks” of content or general skills needed for future engineers, or even socio-economic needs of whole countries (Munjal, 2004), rather than detailed studies on how engineering students learn.

In addition, 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. This might be due to the belief that “doing”, most of the times, leads to “learning” (see, in contrast, Soloway, Guzdial & Hay, 1993). For instance, 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. But in some particular areas, such as Materials Science, this could be disadvantageous, especially 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. Evidently, no computer modeling environment can uncover all of its computational procedures – it would be impractical, for example, to have students wire hundreds of transistors so as to understand the underlying logic of the modeling environment. Nevertheless, we believe that the level of opaqueness of the environment, in many cases, should be made more transparent to students, after careful analysis of the content matter being taught.

Moreover, in Materials Science, many of the traditional formulas *themselves* are opaque—they embody so many layers of accumulated scientific discovery 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. Thus, 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 *conceptual understanding* (knowing how each of the chemical elements interact and alter the properties of the alloy). Particularly, we are interested in *generative, extensible understanding* – learning principles from one phenomenon that could be transferable to other related phenomena.

This is not to say that traditional formulaic representations are intrinsically negative - Sherin (2001), for example, showed how the manipulation of formulas can lead to conceptual understanding in Physics. Our work so far has suggested, nevertheless, that the exclusive use of formulaic representations could constitute an obstacle for conceptual understanding in some domains of engineering, especially in areas in which microscopic interactions of millions of elements (such as atoms or molecules) happen simultaneously. In those areas, it seems to be especially beneficial to unpack, and deconstruct the traditional representations, restructuring domains of knowledge around the study of local phenomena (Wilensky et al., 2006).

For the most part, professional engineering tools target *modeling-for-doing*, which emphasizes “aggregate”-level (Wilensky & Reisman, 1998; Wilensky & Stroup, 2001) simulations to predict macroscopic variables. 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, we posit, in the multi-agent perspective, the number of events and phenomena influencing a local interaction is dramatically lower than at an aggregate level, due to the fact that the many of the variables observed macroscopically are just emergent properties of the local rules. In this paper we propose a learning design framework which benefits from this fact, focusing on simple agent-level behaviors (i.e. atomic-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, fabrications techniques, and industrial applications has grown so quickly and vastly that “covering” all the knowledge by simply adding new courses 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 which are hard to predict or understand with conventional mathematical equations. We posit that unifying, “anchor” multi-agent models are useful for generating solid understanding of generative principles, in order to bridge the micro- and macro-levels (Wilensky & Resnick, 1999).

Therefore, the new computational tools should not be simple add-ons to the present curriculum, but part of their backbone.

Our approach is one attempt in this direction. It builds up from previous research on the use of multi-agent simulation tools in school and research environments, to investigate a variety of phenomena in chemistry and physics (Stieff & Wilensky, 2003; Wilensky, 1999a; Wilensky & Reisman, 2006; Wilensky & Resnick, 1999; Wolfram, 2002). From our literature review, however, it appears that no research has been done in using multi-agent computer-based modeling for teaching topics in Materials Science, even though this technique is widely used in Materials Science research.

Wilensky & Resnick (1999) first noted the need to pay attention to “levels” and highlighted the importance of the understanding of emergent behaviors for Science learning. 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., 2006). Goldstone and Wilensky (2005) have called for such a restructuring of science curricula using common transdisciplinary “patterns” such as energy minimization, positive feedback and simulated annealing.

We will present and discuss a user study of a computer-based learning environment designed within a multi-agent modeling framework and a Constructionist pedagogy, which tries to address the aforementioned challenges by offering students opportunities to build their knowledge by designing and understanding simple computational behaviors that generate complex collective behaviors. The MaterialSim materials we developed focus on the topic of *grain growth* in materials. The user study was comprised of classroom observations, pre/post interviews and surveys, and data analysis from individual sessions with students using the materials.

3. Research Design & Methods

The research took place during two spring quarters at the Materials Science Department of a Midwestern liberal arts university. In the first year (2004), six sophomore year undergraduate students (volunteers) participated in the study, all enrolled in the “Microstructural Dynamics” undergraduate course. In the second year (2005), eleven students volunteered to participate. The average class size in both years was 15 students, although the average class attendance was around 13 students. Each student participated in two individual sessions. The first, 75 minutes 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:
 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? Why?
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 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).

As homework, they were asked to pick 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 just extending the functionality of the existing grain growth model.

The second session (150 minutes) was dedicated to:

- Introduction to the primitive commands of the NetLogo modeling language.
- Implementation (i.e., coding) of the new model.
- Final interview.

We scheduled the sessions to happen approximately one week after students' exposure to this topic in their regular classes.

All sessions were videotaped, and students' computer interactions were recorded using real-time continuous screen-capture software. Approximately 50 hours of video were captured, 40% of that total was selectively transcribed and analyzed. Experiments done by students, as well as the models they built, were saved and analyzed. The first author attended the Microstructural Dynamics course both in 2004 and 2005, 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 exact 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.

4. Computer-based methods vs. traditional methods for Grain Growth modeling

Most materials are composed of microscopic “crystals”. Even though we commonly associate the term ‘crystal’ with the material used in glassware manufacturing, its scientific meaning is different. A crystal is just 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 significantly increase the risk of serious injury for the passengers. 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 850x) show typical results.




Figure 1: Metallic sample before and after grain growth (Blikstein & Tschiptschin, 1999)

Burke (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:

$$R = kt^n$$

where R is the mean grain size at a given time, t is time, k is a constant that varies with temperature, and the theoretical value of n is 0.5.

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 (see Figure 1, left) would have a very fast growth rate, while a system with just a few grains (see Figure 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.

Massive computing power, in the early eighties, has made a new and promising approach possible: computer simulation of grain growth. Anderson, Srolovitz et al. (Anderson, Srolovitz,

Grest, & Sahni, 1984a, 1984b) proposed the widely known theory for computer modeling of grain growth using the Monte Carlo method and a cellular-automata approach. 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, Srolovitz 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, Srolovitz, Grest, & Sahni, 1984b)

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.

Both traditional methods and computer-based methods of investigating grain growth rely on *modeling*. It behooves us to remember a “secret” of science. The scientific enterprise is the process of creating models that are the best approximations to reality we can find. The models of each time period reflect the tools available at that time. Traditional scientific models employed the best representational tools available – mathematical equations. But the recent availability of computational representations and tools enables new kinds of models. Given this rapid growth of tools, we argue that engineering schools should prepare students not only to use existing tools, but also to have the adaptive expertise needed to redesign tools in light of technological innovations or in the presence of new engineering challenges.

5. Software Design: NetLogo and MaterialSim

NetLogo (Wilensky, 1999b) is a freely-available, integrated multi-agent modeling environment, developed at the Center for Connected Learning and Computer-Based Modeling at Northwestern University, under the direction of Prof. Uri Wilensky. 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. Such languages enable users easily to create and manipulate thousands of graphical agents and define simple rules that govern the agents’ behavior. The NetLogo agents can perform simple rule-based behaviors, such as to seek or to avoid being surrounded by other agents. 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. It also includes some specialized tools such as BehaviorSpace (Wilensky, 2000), which enables users to explore a wide parameter space by running multiple experiments and automatically logging the data.

MaterialSim is a set of exploratory models built by the authors of this paper within the NetLogo environment. Currently there are models for investigating crystallization, solidification, casting, grain growth and annealing.

The system was conceived to enable four kinds of activities:

- One-dimensional exploration: users can change variables, draw microstructures and observe their behavior over time.
- Multi-dimensional exploration: students can run experiments altering multiple parameters to find out rules, mathematical relationships, and patterns.
- “Multi-world” exploration: Students can connect real-world and virtual experiments, importing digital photos from real experiments and observing their “virtual” evolution.
- Model building: students can change, create or extend the system by coding their own procedures or modifying existing ones, using the NetLogo modeling language.

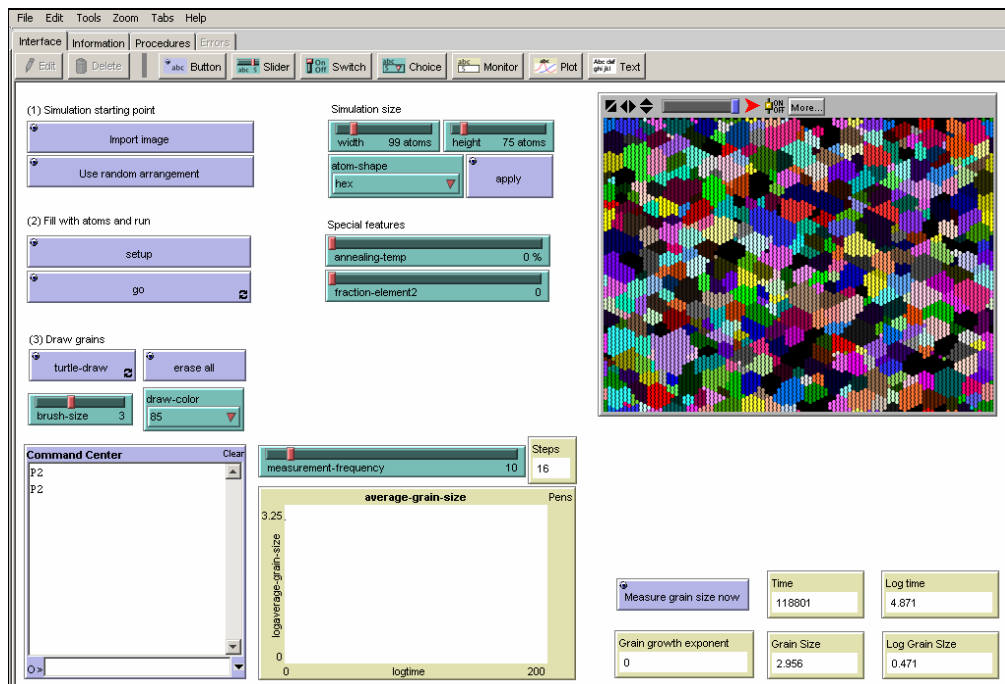


Figure 2. MaterialSim’s grain growth model

We chose the NetLogo modeling-and-simulation environment as a platform as it is well adapted to all of those activities. Using a constructionist design framework combined with a multi-agent approach, NetLogo’s “low-threshold, no-ceiling” (Papert, 1980; Tisue & Wilensky, 2005) design enables learners to achieve sophisticated results within a relatively short period of time. Moreover, its built-in visualization tools allow dynamic, flexible, and customizable views. In addition, NetLogo and other multi-agent simulation tools have been used in many school and

research environments to investigate a variety of phenomena in chemistry and physics (Stieff & Wilensky, 2003; Wilensky, 1999a; Wilensky & Reisman, 2006; Wilensky & Resnick, 1999; Wolfram, 2002).

MaterialSim’s grain growth module has a number of features that were specifically designed as learning tools:

- Users can start either from a random arrangement of atoms or from a pre-set stage, which can be drawn by the user or converted from a digital picture.
- The appearance of the “atoms” can be changed for better visualization of particular aspect, such as number of grains, alignment of atoms etc. (see Figure 3).




Figure 3. Different visualization modes

- Users can change the temperature, which increases the likelihood of non-energetically-favorable orientation flips. In addition, a user-determined percentage of second-phase particles can be introduced in the sample, which alters the rate of grain growth.

In addition to the Grain Growth model, which was the main focus of this study, MaterialSim includes models for diffusion, solidification and crystallography.

6. Approaches to Materials Science learning

Materials science and engineering 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. Thornton and Asta (2005) recently conducted a comprehensive survey about the state of computational Materials Science in undergraduate and graduate courses at the 20 leading programs in the United States. While many universities are creating or planning to create Computational Materials Science courses, one striking conclusion from their work is that the prevailing mindset in most of those institutions is that one should learn modeling *after* learning the science. In other words, computer modeling is regarded as “icing in the cake” for scientific understanding, and not an integral part of science learning. Our work purports to evaluate the usefulness of a different approach: learning the science **by** modeling.

Even though the new computational research tools are making a different understanding of the various phenomena in Materials Science more accessible, they have not yet reached the classroom. Common practice in teaching grain growth is to think of grains as spheres (which they are not), “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., “dislocations climb”, “grains swallow others”, “particles hold boundaries” etc.) to describe and predict changes in the grain.

This teaching practice is widespread in engineering education: traditional engineering teaching relies on a large number of ad-hoc models, approximations, metaphors, and shortcuts, which have been accumulating over decades. They constitute a web of ideas of very different natures, granularities, levels of analysis, and mathematical descriptions. Due to the applied and integrative aspect of engineering research and practice, oftentimes explanations are drawn from a variety of sources: geometrical proof, thermodynamics, algebraic deductions, and statistical mechanics. Grain growth is one example: our classroom observations revealed that at least three of sources were employed during the class 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 basically 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 being accommodated in the other grain, n_1 the number of atoms in grain 1 in position to make the jump, v_1 the vibrational frequency of an atom in grain 1.

- Geometrical approximations is a common technique used to calculate grain size and also the effect of second-phase particles in grain growth. Here, the force (P) is applied by a particle to grain boundaries. The model assumes the particles as spheres and boundaries as lines or surfaces.

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

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

Later in this paper, we suggest that although this *many-to-one* modeling approach might be efficient for predicting properties of materials in real world setting, this multitude of models could be an obstacle to student understanding.

Multi-agent simulation of grain growth offers a different perspective. Its principle is the thermodynamics of atomic interactions – one of the *extensible*, *transferable*, *anchor* models, which we mentioned in section 1 (Introduction). The first step is to represent the material as a hexagonal 2D matrix, in which each site corresponds to an atom and contains a numerical value representing its crystallographic orientation. Contiguous regions (containing 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 free energy (G_i) calculated based on its present crystallographic orientation (Q_i , represented by an integer) and its neighborhood (the more neighbors of differing orientation, the higher its free energy). Figure 4 (left side) shows the central agent with four different neighbors, hence the value of its *initial free energy* (G_i) is 4.
- One new random crystallographic orientation is chosen for that agent (Q_f), among the orientations of its neighbors. In this case, as observable in Figure 4, the current value of the central agent is “2”, and the new transition value is “1”.
- The agent’s free energy is calculated again (G_f), with the new proposed crystallographic orientation ($Q_f=1$). Figure 4 (right side) shows that there are only two different neighbors in the new situation, thus the *final free energy* (G_f) decreases to 2.




Figure 4: Initial and final free-energy calculations. Black and white arrows denote different or equal neighbors.

- The two states are compared. The value that minimizes the free 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.

Pre-test/interview explanations

The pre-test was conducted as a semi-structured interview, and students could also do free drawing to illustrate their thinking. 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 section 3 (Research Design & Methods).

Below we have a commented transcription of some excerpts of Question 1 (all names were changed for anonymity). The goal in this section is to show and discuss student understanding of core ideas in Materials Science.

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 directions 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 (and, oftentimes, incorrect) students' explanations were, which might be surprising as they are all Materials Science majors and have attended classes about the same topic just a week before the interview. What is more, However, as diSessa (1993) extensively discussed on his work with Physics learning, students' understanding is strongly influenced by previous knowledge, phenomenological primitives (*p-prims*) and *ad-hoc* theories. The idea and description of a *grain* is a reasonably basic notion in Materials Science, from which a large number of important concepts derive. Bob and Erika offer a reasonably complete explanation, close to what we would find in a textbook. Most of the other students, on the other hand, diverge from that version. This divergence is 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. Ken mixes the definition of a grain with another topic in the course, dislocation theory. Ella apparently understands what a boundary is, but has a difficult time explaining what they enclose. Students resort to a variety of metaphors and explanations for characterizing a grain: the surface under the microscope, the grain of rice or the atoms with same orientation. Many responses to this first question did not diverge completely from the acceptable concept. However, the questions were structured in a growing level of complexity and "dynamicity". The first one dealt with a definition of a static entity (grain), whereas the third, for example, refers to a phenomenon which evolves with time (grain growth). As students went on in the interview and questions started to deal with more

complex and dynamic processes, their diversity of explanation models increased, as we will see in the following commented excerpts.

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*.... ahn... 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. 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 way. 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, whereas the Laplace-Young equation (studied in class) stated a very different idea. Impurities are

not eliminated by grain growth, and growth can exist in 100% pure materials. Apparently, he imagines that grain growth drives impurities to the outside of the material, “cleaning” it. Here we can notice that students “mix-and-match” models that appear superficially to be related, such as “grain growing” and “grains 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 different reasons, and with different kinetics. During the pre-test, 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 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 expressions such as “curvature is not good, so they will want to shrink”.

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

Interviewer: What is the effect of dispersed particles?

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 moulds 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 increase 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, the students tried to employ a variety of models: a *force-feedback* model, where particles pull boundaries away, slowing grain growth (Bob); a *dislocation movement* model (Betty), and a purely *geometrical* one, with no consequences for the behavior of the material (Liz). Betty's explanation draws from dislocation theory (another topic explored in class weeks before), but does not address 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". Interestingly, 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 good part of the theory, but lack fundamental understanding of it. After an almost accurate explanation about grain growth, Peter says that the driving force is "called free energy, but I don't really know how to explain that."

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 contradictions of all those components which, in fact, are describing different aspects of atomic movement. This relates directly to our first research question.

7. First session: introduction and model exploration

As described in the Research Design and Methods section, the first session was dedicated to the exploration of existing models. The first activity was simple: observe and reflect on of curvature as a driving force for grain growth. Most of the students knew that large grains consume small ones, grains grow 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. 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.

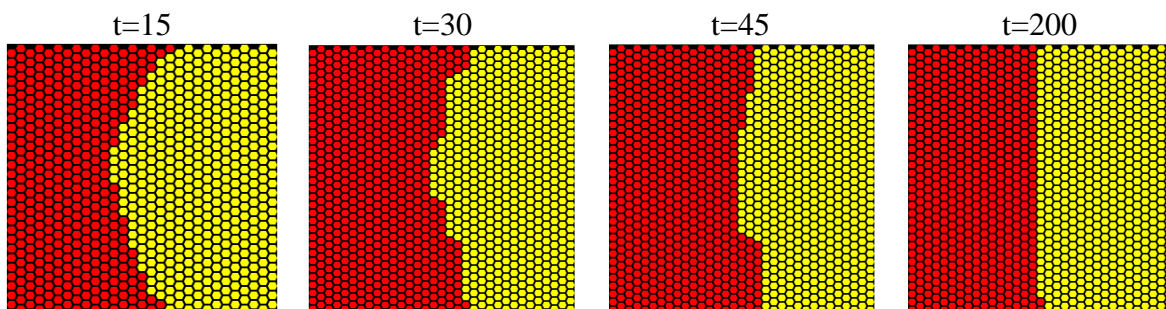


Figure 5. The evolution of a curved grain interface

Before the simulation, most students were unsure of what would happen. As they saw grains growing toward their centers of curvature, they also observed random flipping of atoms.

The following excerpt suggests that visualizing this evolution 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 the results of the Laplace-Young equation, which relates surface tension and curvature. However, she cannot yet think in 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 one of them a lot smaller than the others, as we can see in Figure 6. Liz continued with her line of reasoning.




Figure 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)

Watching the evolution of this microstructure was a key experience for Liz. She started to move from rote memorization and a topic-specific model to more general, principled knowledge. This excerpt took place when she was observing a triple point, a region where three grains meet and the flipping probability is the same for all (as there are two atoms of each grain around the central element).

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... it 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, working with MaterialSim, Liz was able to understand and generate hypotheses about two essential (and complex) ideas: triple points and the time dependency of grain growth. Without realizing it, she understood the reason for the triple point to be considered a “low-mobility” point in a microstructure. The central atom has two atoms 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, so growth speed decreases over time and eventually “tapers off”. The additional importance of this discovery is that, rather than *being told*, Liz arrived at this conclusion *on her own*, by drawing microstructures, changing variables and observing the dynamics of the simulation.

Generally, most students knew that the small grain was going to disappear. From their reactions while observing the simulation, they seemed to be expecting a unidirectional animation grains being “eaten” by the surrounding ones. This was consistent both with the heuristics and the types of results of aggregate tools, animations, and equations. However, what students observed was different: Behaviors emerge 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 jumping to different positions). 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*. Therefore, not only is the simulation visually similar to the phenomenon, but also its algorithm loyally emulates the micro-level of the phenomenon's underlying process. This

is different from purely numeric simulations in which what students are able to compare are 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-test data suggested, can lead to misconceptions. Working in MaterialSim, students realized that grains were not being “consumed” or shrinking: atoms were just switching places.

The last activity of the first day was the “BehaviorSpace” experiment. This NetLogo feature allows users to automatically run hundreds of simulations each under different parameter settings. Students ran at least one set of experiments, charted the data, and came up with theories to describe the phenomenon. Most students chose to model the influence of dispersed particles. Figure 7 has a sequence of images at the same time step for different percentages of dispersed particles, as well as the individual NetLogo plots and a chart showing a very good fit with the theoretical data (dotted line). In this activity, students could even further generate their own hypotheses and equations, having not only the dynamic visualization but also actual numerical data on the evolution of the microstructure.

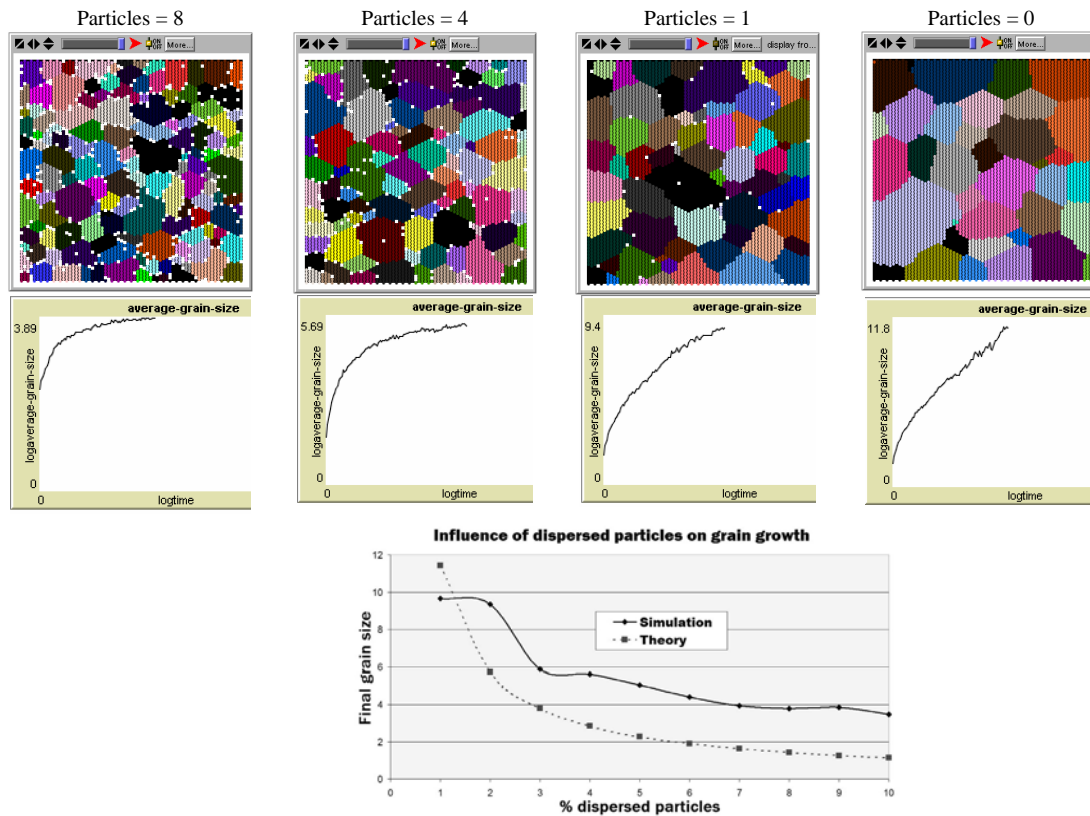


Figure 7. Sequence of screenshots from students' experiments

8. Second session: building their own models

The tasks of the second session were learning the basics of the NetLogo language, and programming new additions to the models. Students were asked to come up with their own ideas for extending MaterialSim's models. They pursued questions of their own and authored novel features for the models, which helped them elaborate on answers to their research questions.

Student achievement was impressive. A comparison between the pre-test data, when students relied on ready-made statements about the phenomenon, 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 students.

Betty built a model that incorporates misalignment between grains: In her innovative model, the more misaligned, 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 useful measure independent of the quadrant. Betty's solution, after much thinking/drawing, was to use the arcsine function. The following picture shows some of her reasoning. From her drawing, we can observe that she was using ideas from geometry, but now in a "micro" level, taking into consideration the orientation of individual atoms.

The probability for an atom to jump to the next grain was no longer just a function of the number of different atoms around it, but also of the average misorientation among them. The higher the misorientation, the harder it will be for that atom to move to another grain. Very low misorientation, on the other hand, would promote easier growth. This all now seemed intuitive to Betty.

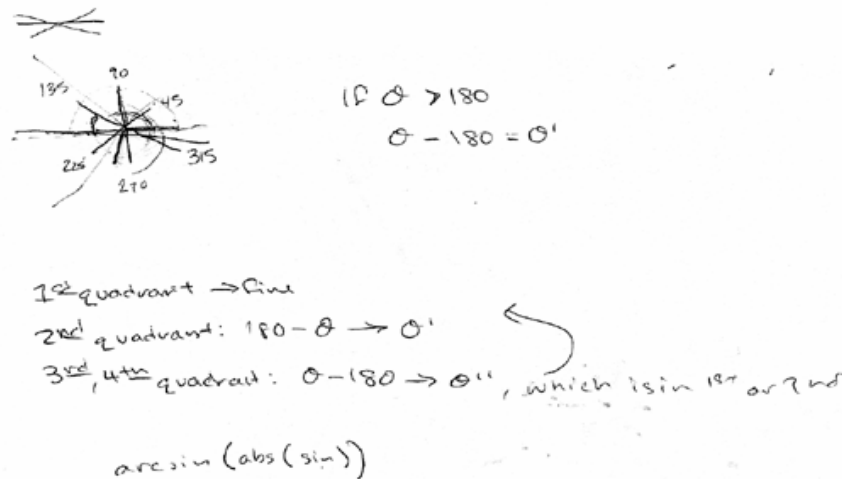


Figure 8: Betty's reflection about angles, sine and arcsine.

Yet, aggregate and macroscopic models do not afford such insight. 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.

Bob had a different idea: he wanted to include a new parameter: the size of the dispersed solid particles. The idea was to allow users not only to change the percentage of particles, but also the radius, which was a complex challenge. Bob realized that given a certain percentage (in mass) of particles, their number had to be adjusted to compensate for the increased mass of each. 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 a formula for calculating the area of polygons placed in the hexagonal grid, which turned out to be an interesting mathematical exercise. Bob realized that a recursive procedure would be adequate, as new layers were being added after each iteration. After completing the model, Bob investigated the influence of particle size on grain growth: maintaining the same percentage in mass, how is growth affected by changing an individual precipitate's volume? Bob was able to run large batches of simulations in BehaviorSpace, chart the data, and explore possible explanations.

Anand did not want to investigate grain growth, and built a model from scratch. His idea was to explore in detail interfacial energies due to atomic misalignment. In other words, his model meant to determine how much energy was stored in an interface between two grains which did not match in terms of atomic spacing. We can observe his model in Figure 9: one of the grains (in red), has a very large spacing between the atoms compared to the “blue” grain. This causes the interface between the two to store energy, as the atomic structure is somehow stretched. Anand built this model using the same “kernel” as the grain growth model: atoms look around and check the state of their neighbors, deciding what to do based on the proportion of equal and different neighbors. Even though this topic, in the regular curriculum, was separate from grain growth, he was able to identify a “transferable” model between the two phenomena.

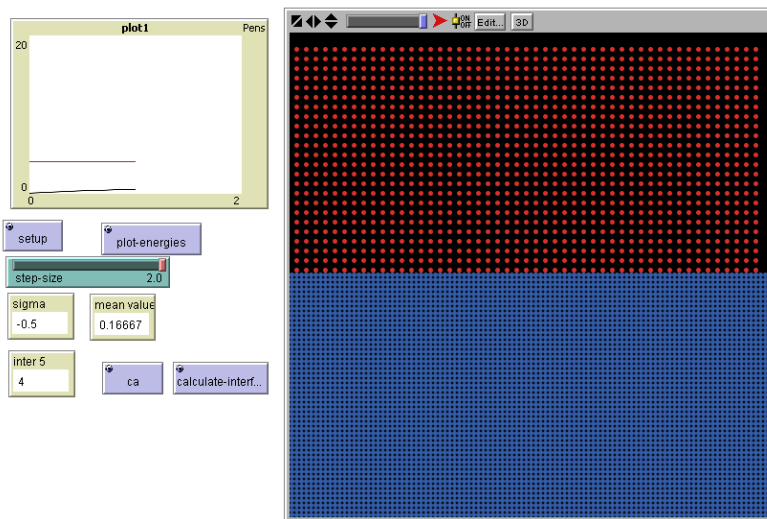


Figure 9 – Anand’s model for detailed study of interfacial energies

These examples of student model building were implemented in less than two 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 support 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 pretest, they

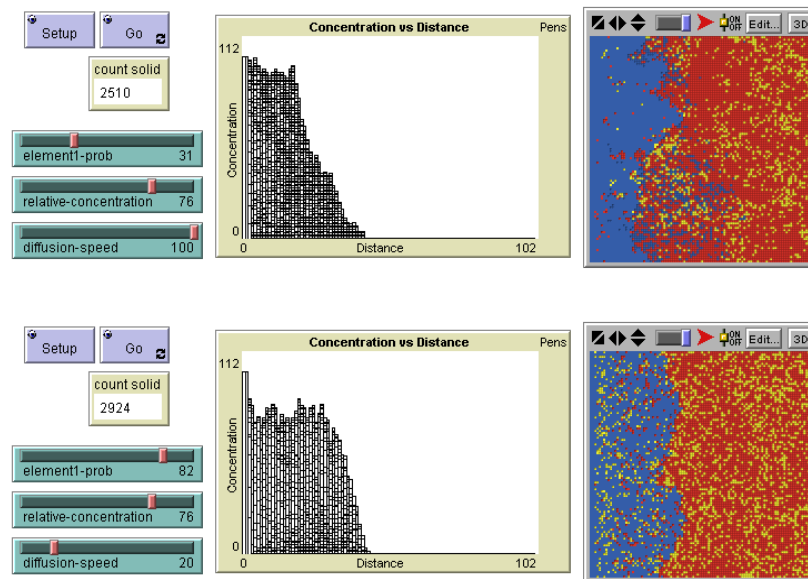
demonstrated difficulty in explaining related phenomena in a coherent fashion, resorting to a range of models and metaphors in a fragmented fashion. The implementation of their own model within an agent-based simulation environment, however, provided students with fewer, simpler rules that were closely related to the physical phenomenon, thus enabling students to better understand and extend the model by adding new proximal rules for the agents.

Peter’s model was an excellent example of the generative, transferable aspect of the agent-based approach. 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 two hours, coded a model which complexity is far beyond what is expected from the Microstructural Dynamics course, considering my classroom observations and analysis of class materials. Peter used the agent-based approach and some “kernels” from grain growth to create a diffusion model. Even though the two phenomena have differences, he managed to identify the common kernels and copy them from one model to the other, with the necessary adaptations, in the same way as Anand and his Interfacial Energy model. When Peter told me his future plans for the model, at the end of the session, he explored another common agent-based heuristics:

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.






Figure 10. 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 the exact same shape for the two concentration curves. Note that this last chart was rotated for clarity purposes.

9. Conclusion

Design is now fashionable in many engineering schools. Robotics competitions, for instance, are common in various universities. However, not all of engineering is mechanical. Could we extend the powerful ideas about mechanical construction to fields such as Materials Science, with products 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 visible and clearly laid out. 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”, formula-based descriptions. Our user study suggested that the fragmentation and opaqueness of such descriptions could constitute an obstacle to learning. Firstly, the traditional descriptions are more context-specific, and do not enable students to make broader inferences about phenomena with similar rules. Secondly, the descriptions often lead to heuristics that foster misconceptions. Students had memorized ideas about grain growth for which they have no “feel” or intuition. 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 a ready-made model.

3) One of the pillars of Constructionist theory is the importance of students conducting personally-meaningful projects. We have some preliminary indication, to be confirmed in future work, that coding their *own* models was a valuable learning experience. Students had an opportunity to test and debug their theories, as well as reconcile them with previous knowledge. In the post survey, students self-reported motivation in the model-building session to be one of the most positive aspects of the experience.

However, we caution that not all modeling tools are created equal, and thus stress the importance of designing software that foregrounds key generative models, as to enable students to extend their knowledge into new phenomena by making use of them... The study, in conclusion, suggests that the knowledge and the exploration of just a few simple underlying rules of natural phenomena appear to be more generative for students than the more encapsulated aggregate, equation-based knowledge.

10. References

- Anderson, M. P., Srolovitz, D. J., Grest, G. S., & Sahni, P. S. (1984a). Computer simulation of grain growth - II. Grain size distribution, topology and local dynamics. *Acta Metallurgica*(32), 793-801.
- Anderson, M. P., Srolovitz, D. J., Grest, G. S., & Sahni, P. S. (1984b). Computer simulation of grain growth: I. Kinetics. *Acta Metallurgica*(32), 783-792.
- Bazzo, W. A. (1998). *Ciência, Tecnologia e Sociedade: e o contexto da educação tecnológica (Science, technology and society: the context of technological education)*. Florianópolis: Editora da UFSC.
- Blikstein, P. (2001). *Transdisciplinary Science and Technology Ateliers*. Unpublished MSc. thesis, University of São Paulo, São Paulo.
- Blikstein, P., & Tschiptschin, A. P. (1999). Monte Carlo Simulation of Grain Growth. *Materials Research*, 2(3), 133-137.
- Blikstein, P., & Wilensky, U. (2004). *MaterialSim: An Agent-Based Simulation Toolkit for Learning Materials Science*. Paper presented at the International Conference on Engineering Education, Gainesville, Florida, USA.
- Blikstein, P., & Wilensky, U. (2005). *Less is More: Agent-Based Simulation as a Powerful Learning Tool in Materials Science*. Paper presented at the IV International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2005), Utrecht, Holland.
- Brown, G. S. (1961). Today's Dilemma in Engineering Education. *Ire Transactions on Education*, E 4(2), 48-&.
- Burke, J. E. (1949). *Trans. Am. Inst. Min. Engrs*(180), 73.
- Centola, D., Wilensky, Uri & McKenzie, E. (2000). Survival of the groupiest: Facilitating students' understanding of the multiple levels of fitness through multi-agent modeling- The EACH Project. *International Journal of Complex Systems*.
- Colgate, J. E., McKenna, A., & Ankenman, B. (2004). IDEA: implementing design throughout the curriculum at Northwestern. *International Journal of Engineering Education*, 20(2).
- Collier, N. S., D. (2001). Repast. University of Chicago (2001). <http://repast.sourceforge.net>.

- diSessa, A. A. (1993). Toward an epistemology of physics. *Cognition and Instruction*, 10(2 & 3), 105-225.
- Einstein, H. H. (2002). Engineering change at MIT. *Civil Engineering*, 72(10), 62-69.
- Freire, P. (1974). *Pedagogy of the oppressed*. New York,: Seabury Press.
- Goldstone, R. L., & Wilensky, U. (2005). Promoting Transfer through Complex Systems Principles (Manuscript in preparation).
- Haghighi, K. (2005). Systematic and sustainable reform in engineering education. *Journal of Environmental Engineering-Asce*, 131(4), 501-502.
- Hatano, G., & Oura, Y. (2003). Commentary: Reconceptualizing School Learning Using Insight From Expertise Research. *Educational Researcher*, 32(8), 26-29.
- Hurst, K. D. (1995). *A new paradigm for Engineering Education*. Paper presented at the ASEE Frontiers in Education.
- Jerath, S. (1983). Engineering-Education in Perspective. *Mechanical Engineering*, 105(2), 92-93.
- Katehi, L. P. B., Banks, K., Diefes-Dux, H. A., Follman, D. K., Gaunt, J., Haghighi, K., et al. (2004). *A New Framework for Academic Reform in Engineering Education*. Paper presented at the 2004 American Society for Engineering Education, Nashville, TN.
- 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.
- 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 Associates Inc.
- MIT Center for Policy Alternatives. (1975). Future Directions in Engineering-Education - System Response to a Changing World - Report by Center for Policy Alternatives for MIT School of Engineering. *Engineering Education*, 65(5), 382-382.
- Munjal, M. L. (2004). Slow pace of engineering education reforms (vol 84, pg 1175, 2003). *Current Science*, 86(7), 900-900.
- Newstetter, W. C., & McCracken, M. (2000). *Design Learning as Conceptual Change: A framework for developing a science of design learning*. Paper presented at the Conference of the American Society of Engineering Education, St. Louis, MS.
- Papert, S. (1980). *Mindstorms : children, computers, and powerful ideas*. New York: Basic Books.
- 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.
- Sherin, B. (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.
- 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., & Mark, A. (2005). Current status and outlook of computational materials science education in the US. *Modelling and Simulation in Materials Science and Engineering*, 13(2), R53.
- Wilensky, U. (1999a). 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 Verlag.
- Wilensky, U. (1999b). NetLogo. Evanston, IL: Center for Connected Learning and Computer-Based Modeling. <http://ccl.northwestern.edu/netlogo>.
- Wilensky, U., McKenzie, E., & Centola, D. (2000, June 14-17). *A Hands-on Modeling Approach to Evolution: Learning about the Evolution of Cooperation and Altruism through Multi-Agent Modeling-The EACH Project*. Paper presented at the The Fourth Annual International Conference of the Learning Sciences, Ann Arbor, MI.
- Wilensky, U., Papert, S., Sherin, B., diSessa, A. A., Kay, A., & Turkle, S. (2006). *Center for Learning and Computation-Based Knowledge (CLiCK). Proposal to the National Science Foundation - Science of Learning Center*. Unpublished manuscript.
- 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 perspective to making sense of the world. *Journal of Science Education and Technology*, 1(8), 3-18.

Wolfram, S. (2002). *A new kind of science*. Champaign, IL: Wolfram Media.