

2006-2496: A CASE STUDY OF MULTI-AGENT-BASED SIMULATION IN UNDERGRADUATE MATERIALS SCIENCE EDUCATION

Paulo Blikstein, Northwestern University

Doctoral student at the Center for Connected Learning and Computer-Based Modeling at the School of Education and Social Policy (Northwestern University).

Uri Wilensky, Northwestern University

Associate Professor at the Department of Electrical Engineering and Computer Science and at the School of Education and Social Policy at Northwestern University. Director of the Center for Connected Learning and Computer-Based Modeling (CCL).

A Case Study of Multi-Agent-Based Simulation in Undergraduate Materials Science Education

1. Introduction

This paper reports on a user study of a computer-aided learning environment for college-level Materials Science. “MaterialSim” is an agent-based set of *microworlds* built by the authors within the NetLogo¹ modeling environment, for investigating phenomena such as crystallization, solidification, grain growth and annealing. This design started out from a literature review on Materials Science education, extensive classroom observations and analysis of class materials. This preliminary phase suggested that students’ understanding of the subject matter was problematic, and that the teaching strategies and technological learning 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 introductory Materials Science course participated in the study in 2004 and 2005, which was comprised of a survey, pre-interview, interaction with the computer models, and model building.

The complexity of college-level content in the field of Materials Science is such that teachers oftentimes end up resorting to a multitude of aggregate-level models to explain a particular phenomenon, each of them requiring a different set of equations. Our first research question was to investigate the consequences of this common teaching strategy: what kind of understanding did this multitude of explanation models foster in students?

MaterialSim’s *multi-agent-based* computer models were built within a different framework. Instead of resorting to a variety of aggregate-level description, the core feature of the multi-agent-based perspective is the application of a small number of local rules to capture fundamental causality structures underlying complex behaviors within a domain. The rationale for the design is that this agent-based perspective may foster different, and perhaps more generative understanding of the relevant scientific phenomena. Instead of multiple models or numerous equations, this framework focuses on a small number of elementary models which can be applied to a variety of scientific phenomena. This is in opposition to common teaching strategies in Materials Science. Our second research question, thus, was to study students’ understanding of the Materials Science content departing from this agent-based perspective.

A third research question, about which we will present only preliminary results, was to investigate the learning outcomes of building 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 build new models based on those principles.

2. Engineering education is flux?

Engineers, who frequently work at the interfaces of different fields of knowledge, are in a privileged position to integrate knowledge and be creative. Nevertheless, as in most domains, its teaching and learning have remained remarkably traditional for decades. This phenomenon, which has been pinpointed as early as the sixties²⁻⁴ is nowadays prevalent in numerous engineering schools, which, in recent years, have been trying to overcome it by proposing a number of reform initiatives⁵⁻⁷. Particularly, a popular approach has been to introduce design-based curricula which include hands-on engineering design, oftentimes using modeling and simulation software⁸⁻¹¹. Some of the reasons mentioned for this urge for curricula reform are the declining interest of American high-school students for the career (down 18% since 1991), the changing workforce demographics, the new industrial dynamics brought about by “mass customization” and other new manufacturing needs and techniques¹².

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 synthesis¹³, students and professionals have to deal with time scales from the nanosecond to tenths of years, and sizes from the single atoms to machines many thousands of meters long¹⁴.

Many researchers and industrial leaders in the field have been pointing out that reform initiatives are falling short of those new challenges. Bazzo¹⁵ and Blikstein¹⁶ criticized the “banking” approach of simply adding new courses to the curriculum, and Hurst¹⁷ 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 to future engineers, or even socio-economic needs of whole countries¹⁸, 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”. For instance, most software tools used in engineering courses tools 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. For learning purposes, however, this is disadvantageous, because the computational procedures embody an essential, perhaps crucial, aspect of the subject matter—*how* the conventional formulas and representations capture the phenomena they purport to model. Moreover, especially in fields such as Chemical engineering and Materials Science, 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, using formulas and conventional engineering representations is perhaps conducive to successful *doing* (designing a new alloy, for example) but not necessarily to *understanding* (knowing how each of the chemical elements interact and alter the properties of the alloy).

Therefore, professional tools of the domain are targeted for *modeling-for-doing*, rather than for *modeling-for-understanding*. The former often emphasizes “aggregate”-level simulations to predict macroscopic variables (e.g., time, temperature, concentration, current, pressure). Those variables are oftentimes dependent on a number of events and phenomena, and thus numerous equations, models, tools and heuristics have to be employed, oftentimes “mixing-and-matching” different levels of explanation to achieve the expected results. In this paper, conversely, we propose a *modeling-for-understanding* framework, which focuses on simple agent-level behaviors (atomic-level interactions, free energy minimization) from which complex macroscopic behaviors emerge.

We believe that this framework is especially useful in a technological scenario of increasing 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 the “banking” approach of simply adding new courses to the curriculum would be impossible. 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 a typical problem 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 models are necessary for bridging the micro- and macro-levels¹⁹ – therefore, the new computational tools should not be simple add-ons to the present curriculum, but part of their backbone.

Our *modeling-for-understanding* approach is one attempt in this direction. It builds up from previous research on the use of multi-agent-based simulation tools in school and research environments, to investigate a variety of phenomena in chemistry and physics¹⁹⁻²⁴. Wilensky & Resnick¹⁹ pioneered the discussion of “levels” and highlighted the importance of the understanding of emergent behaviors for Science learning. More recently, Goldstone and Wilensky²⁵ called for a radical restructuring of Science curricula around 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 the modeling-for-understanding framework, which tries to address the aforementioned challenges by offering students opportunities to build their knowledge of complex natural phenomena from the ground up, by designing and understanding simple computational behaviors with growing levels of complexity. The user study was comprised of classroom observations, pre/post interviews and surveys, and data analysis from individual sessions with students using the system.

3. Research Design & Methods

The research took place during two spring quarters at the Materials Science Department of a liberal arts university. On the first year (2004), six sophomore year undergraduate students

(volunteers) participated in the study, all enrolled in the “Microstructural Dynamics” undergraduate course. On the second year (2005), eleven students volunteered to participate. The average class size on 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 views on the course.
- Pre-interview about grain growth and related phenomena, in which students were asked the following questions:
 1. What is a grain?
 2. What is a grain boundary?
 3. What is grain growth?
 4. What is the driving force for grain growth?
 5. What is the driving force for recrystallization?
 6. What is the effect on grain growth of dispersed precipitates? Why?
 7. In grain growth, grain boundaries always migrate toward their center of curvature. How does this decrease the free energy?
 8. 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 (fire spread, virus contamination, racial segregation, gas molecules in a container, and a chemical reaction).
- Hands-on interaction with one of MaterialSim’s 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 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 using a microgenetic framework²⁶. Experiments done by students, as well as the models they built, were saved and analyzed. The first author attended almost entirely the Microstructural Dynamics course both in 2004 and 2005, analyzed the class materials and related literature. 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. Grain Growth: A Primer

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. In Materials Science, scientists use the term “grain” to refer to them.

Among other properties, grain (or crystal) size determines how much a material will deform before breaking apart, which is one of the most important issues in engineering design. High temperatures can cause grain size to change. 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 ²⁷

Burke²⁸ 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, t is time, k is a constant that varies with temperature, and the theoretical value of n is 0.5.

Massive computing power has made a new and promising approach possible: computer simulation of grain growth. Anderson, Srolovitz et al.^{29,30} 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.”²⁹

In other words, Anderson et al. state that the classic rule-of-thumb for grain growth (“large grains grow, small grains shrink”) is not necessarily valid and that randomness plays a much more important role. Given the microscopic dimensions and small time scale of the phenomenon, one of the only ways to visualize this new finding is through computer simulation. The traditional methods for investigating grain size and growth, nevertheless, reflect the tools (and visualization techniques) that were available in the fifties: mathematical abstractions, geometrical modeling, approximations, and empirical data. We are reminded of a “secret” of science: models are just approximations of reality and reflect the tools that were available when they were created. Engineering schools should prepare students not only to use existing tools, but also to redesign them in face of innovative technologies or in the presence of new engineering challenges.

5. Software Design: NetLogo and MaterialSim

NetLogo¹ 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, a co-author of this paper. 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. There are also some specialized tools such as BehaviorSpace, 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 first author 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.

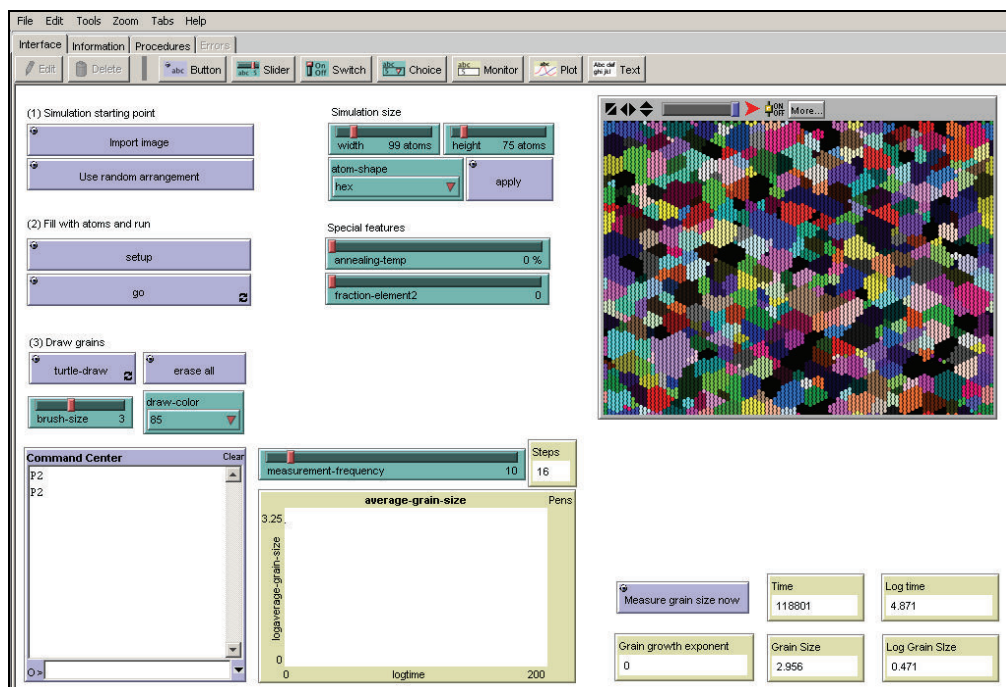


Figure 2. MaterialSim’s grain growth model

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 framework³¹ combined with a multi-agent approach, NetLogo’s “low-threshold, no-ceiling” design enables learners to achieve sophisticated results within a relatively short period of time. Moreover, its built-in visualization tools allows dynamic, flexible, and customizable views. In addition, NetLogo and other multi-agent-based simulation tools have been used in many school and research environments to investigate a variety of phenomena in chemistry and physics¹⁹⁻²³.

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 imported from a bitmapped image.
- The appearance of the “atoms” can be changed for better visualization of particular phenomena (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 changes the rate of grain growth.

6. Approaches to Materials Science learning

Materials science and engineering has grown considerably from its roots in experimental metallurgy and, recently, the main research breakthroughs have been driven by advances in computational methods³².

Thornton and Asta³² 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. Specifically, our classroom observations revealed that at least three of these sources were employed during the class covering grain growth:

- The *Laplace-Young equation for pressure* is commonly used to 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 with temperature, and therefore the mobility of a grain boundary also grows with temperature. The equation is written as follows:

$$A_2 n_1 v_1 \exp\left(\frac{-\Delta G^\alpha}{RT}\right)$$

where 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 the 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 this multitude of models is 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 “anchor models”, or kernels, which we mentioned in the Introduction. The first step is to represent the material as a square or 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, but they have no physical existence in the simulation. 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 attempted 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-structure interview, and students could also draw 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 seven questions, listed in the Research Design & Methods section.

Below we have a commented transcription of some excerpts of Question 1 (students’ names were changed). The goal in this section is to show and discuss student understanding of core ideas materials’ Microstructural Dynamics.

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.

Erin: 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 an 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.

Elise: 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. 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 Erin 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. Interestingly, 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. Elise 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: *visual* (surface under the microscope), *analogical* ("grain of rice") or *geometrical* (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 refers to a dynamic phenomenon (growth). As students went on in the interview, and questions started to deal with more complex and dynamic processes, their diversity of explanation models and heuristics increased, as we will see in the following commented excerpts.

Interviewer: How would you explain to a friend of yours 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. Interestingly, not even one of the students utilized any of the equations or formal explanations seen in class, or tried to explain analytically (with symbols and equations) what grain growth is. Bob, for example, uses the metaphor of willpower ("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. Interestingly, he also states that temperature causes molecules to "move faster", which would be a "driving force", confusing causality and mechanism. 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 very different from grain growth. In recrystallization, similarly, something "grows", but for completely different reasons. 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). She also had a hard time making the distinction between driving force and mechanism. Interestingly, when explaining the phenomenon in her own words, she mentions that

“atoms move to the lower energy positions, which is into the grains, and the grain decreases”, correcting her own conclusion just seconds later, which might suggest that she is “assembling” an explanation on-the-fly, resorting to what she remembers from class. Also, 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, which is not the case).

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, sewed 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... 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, 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 dramatically different from purely numeric simulations in which what students are able to compare are only outputs, and not the processes. 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 experiment, 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.

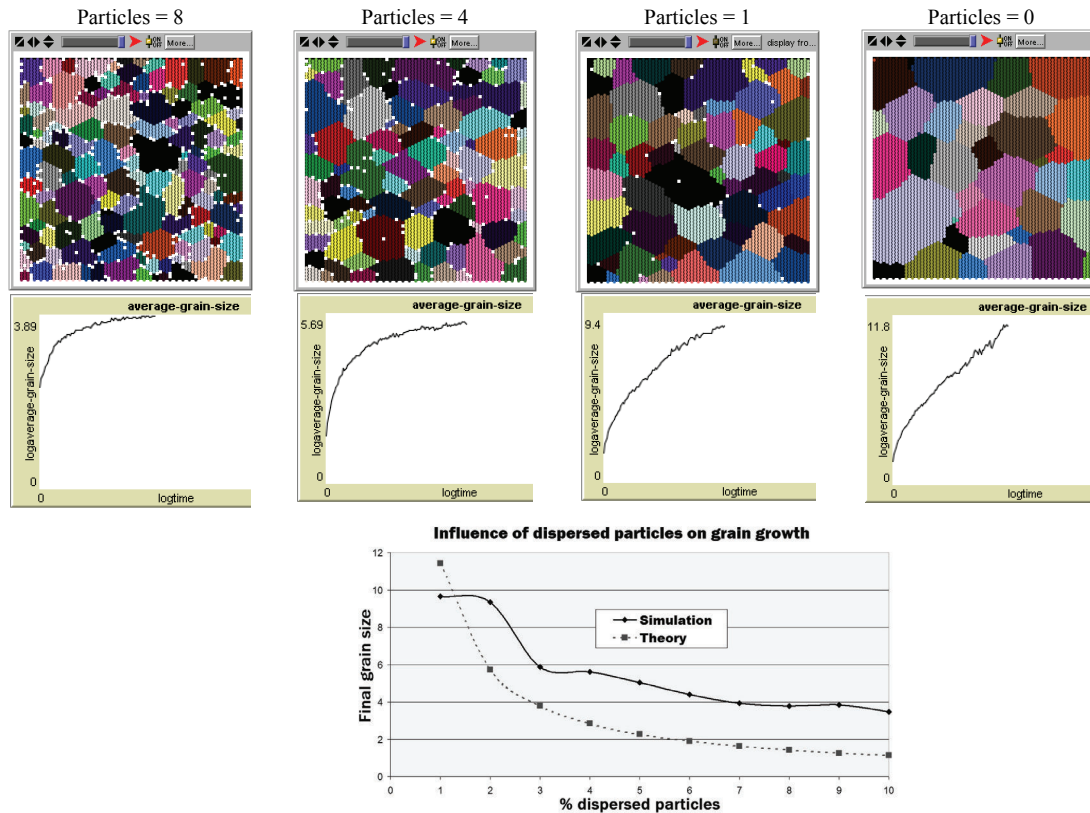


Figure 7. Sequence of screenshots from students' experiments

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.

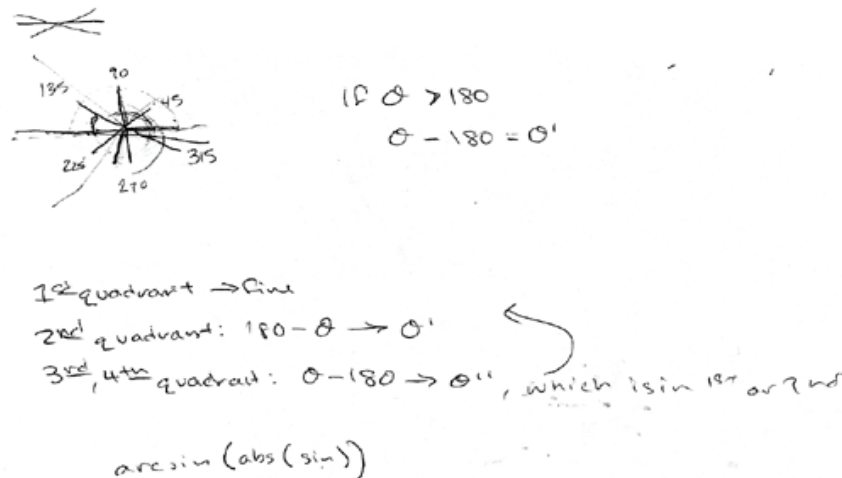


Figure 8: Betty's reflection about angles, sine and arcsine.

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.

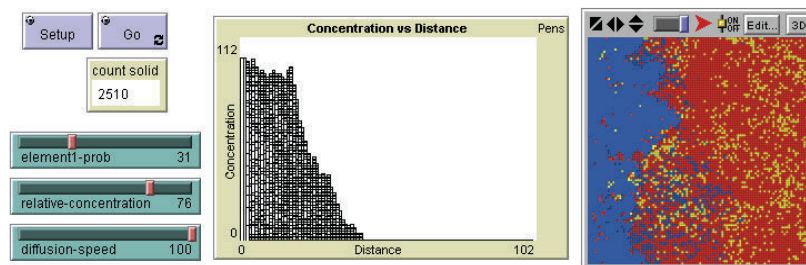
Both examples of student model building were implemented in less than two hours, *including* the time dedicated to learning the basic of the NetLogo language. The relative ease with which students developed their own models, even within such a short timeframe, shows that model building is an approachable 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 aspect of the agent-based approach. In the pre-survey, he identified diffusion control and 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. When he 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.



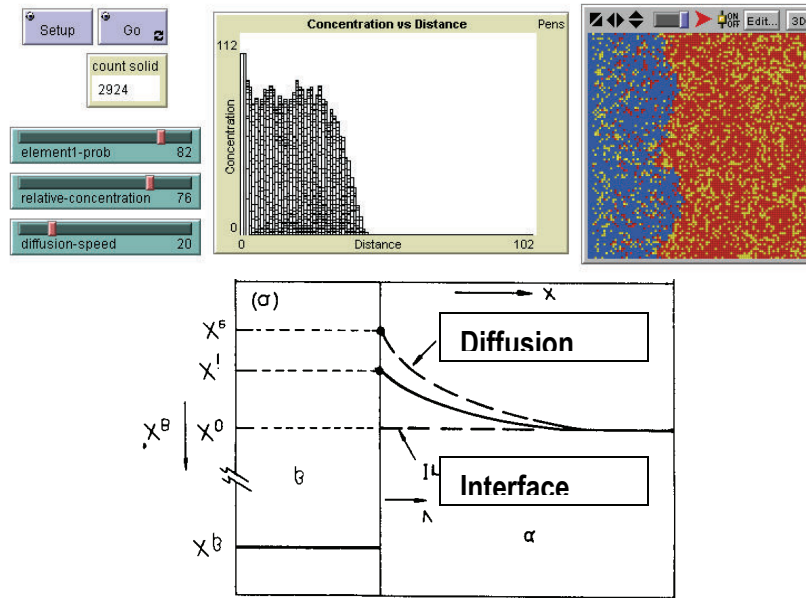


Figure 9. Results of Peter's model with diffusion control (top, with diffusion speed = 100), interface control (bottom, with diffusion speed = 20), and the chart from the textbook, where we can identify the exact same shape for the two concentration curves. Note that this last chart was rotated for clarity 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 for learning. First, they are more context-specific, and did not enable students to make broader inferences about phenomena with similar rules. Secondly, they might generate "rules-of-thumb" and heuristics which could lead to misconceptions. Students had memorized ideas about grain growth for which they had no "feel" or intuition. Thirdly, they 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 essential for student understanding.

2) A core feature of this design is that students can apply a small number of “kernel”, *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, recrystallization, diffusion, phase transformations), which are traditionally taught as *separate* topics with their own models and equations. Some students were even able to create their own models by transferring some “kernels” from a ready-made model.

3) One of the pillars of Constructionist theory³¹ 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 enables modeling-for-understanding and not solely modeling-for-doing. The study, in conclusion, suggests that *less is more*: 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

1. Wilensky, U., *NetLogo*. 1999, Center for Connected Learning and Computer-Based Modeling. <http://ccl.northwestern.edu/netlogo>: Evanston, IL.
2. Brown, G.S., *Today's Dilemma in Engineering Education*. *Ire Transactions on Education*, 1961. E 4(2): p. 48-&.
3. MIT Center for Policy Alternatives, *Future Directions in Engineering-Education - System Response to a Changing World - Report by Center for Policy Alternatives for MIT School of Engineering*. *Engineering Education*, 1975. 65(5): p. 382-382.
4. Jerath, S., *Engineering-Education in Perspective*. *Mechanical Engineering*, 1983. 105(2): p. 92-93.
5. Haghghi, K., *Systematic and sustainable reform in engineering education*. *Journal of Environmental Engineering-Asce*, 2005. 131(4): p. 501-502.
6. Russell, J.S. and W.B. Stouffer, *Survey of the national civil engineering curriculum*. *Journal of Professional Issues in Engineering Education and Practice*, 2005. 131(2): p. 118-128.

7. Einstein, H.H., *Engineering change at MIT*. Civil Engineering, 2002. **72**(10): p. 62-69.
8. Martin, F., *Ideal and Real Systems: A study of notions of control in undergraduates who design robots*, in *Constructionism in Practice*, Y. Kafai and M. Resnick, Editors. 1996, Lawrence Erlbaum Associates Inc: Mahwah, NJ. p. 297-332.
9. Colgate, J.E., A. McKenna, and B. Ankenman, *IDEA: implementing design throughout the curriculum at Northwestern*. International Journal of Engineering Education, 2004. **20**(2).
10. Newstetter, W.C. and M. McCracken. *Design Learning as Conceptual Change: A framework for developing a science of design learning*. in *Conference of the American Society of Engineering Education*. 2000. St. Louis, MS: ASEE.
11. Lamley, R.A., *A model for an engineering education*. Marine Technology and Sname News, 1996. **33**(2): p. 119-121.
12. Linda P.B. Katehi, K.B., Heidi A. Diefes-Dux, Deborah K. Follman,, K.H. John Gaunt, P.K. Imbrie, Leah H. Jamieson,, and W.C.O. Robert E. Montgomery, and Phillip Wankat. *A New Framework for Academic Reform in Engineering Education*. in *2004 American Society for Engineering Education*. 2004. Nashville, TN: American Society for Engineering Education.
13. Roco, M.C., *Nanotechnology - A frontier for engineering education*. International Journal of Engineering Education, 2002. **18**(5): p. 488-497.
14. Kulov, N.N. and M.G. Slin'ko, *The state of the art in chemical engineering science and education*. Theoretical Foundations of Chemical Engineering, 2004. **38**(2): p. 105-111.
15. Bazzo, W.A., *Ciência, Tecnologia e Sociedade: e o contexto da educação tecnológica (Science, technology and society: the context of technological education)*. 1998, Florianópolis: Editora da UFSC. 319.
16. Blikstein, P., *Transdisciplinary Science and Technology Ateliers*, Unpublished MSc. thesis. *Integrated Systems Department*. 2001, University of São Paulo: São Paulo.
17. Hurst, K.D. *A new paradigm for Engineering Education*. in *ASEE Frontiers in Education*. 1995: American Society for Engineering Education.
18. Munjal, M.L., *Slow pace of engineering education reforms (vol 84, pg 1175, 2003)*. Current Science, 2004. **86**(7): p. 900-900.
19. Wilensky, U. and M. Resnick, *Thinking in levels: a dynamic systems perspective to making sense of the world*. Journal of Science Education and Technology, 1999. **1**(8): p. 3-18.
20. Steiff, M. and U. Wilensky, *Connected Chemistry: Incorporating Interactive Simulations into the Chemistry Classroom*. Journal of Science Education and Technology, 2003. **12**(3): p. 285-302.
21. Wilensky, U., *GasLab-an Extensible Modeling Toolkit for Exploring Micro-and-Macro-Views of Gases*, in *Computer Modeling and Simulation in Science Education*, N. Roberts, W. Feurzeig, and B. Hunter, Editors. 1999, Springer Verlag: Berlin.
22. Wolfram, S., *A new kind of science*. 2002, Champaign, IL: Wolfram Media. xiv, 1197.
23. Wilensky, U. and K. Reisman, *Thinking like a Wolf, a Sheep or a Firefly: Learning Biology through Constructing and Testing Computational Theories*. Cognition & Instruction, in press.
24. Sengupta, P. and U. Wilensky. *N.I.E.L.S: An Emergent Multi-Agent Based Modeling Environment for learning Physics*. in *AAMAS 2005 (ABSHL session)*. 2005. Utrecht, Netherlands.
25. Goldstone, R.L. and U. Wilensky, *Promoting Transfer through Complex Systems Principles*. 2005.
26. Siegler, R.S. and K. Crowley, *The microgenetic method: A direct means for studying cognitive development*. American Psychologist, 1991. **46**(6): p. 606-620.
27. Blikstein, P. and A.P. Tschipschin, *Monte Carlo Simulation of Grain Growth*. Materials Research, 1999. **2**(3): p. 133-137.
28. Burke, J.E., *Trans. Am. Inst. Min. Engrs*, 1949(180): p. 73.
29. Anderson, M.P., et al., *Computer simulation of grain growth: I. Kinetics*. Acta Metallurgica, 1984(32): p. 783-792.
30. Anderson, M.P., et al., *Computer simulation of grain growth - II. Grain size distribution, topology and local dynamics*. Acta Metallurgica, 1984(32): p. 793-801.
31. Papert, S., *Mindstorms : children, computers, and powerful ideas*. 1980, New York: Basic Books. 230.
32. Thornton, K. and A. Mark, *Current status and outlook of computational materials science education in the US*. Modelling and Simulation in Materials Science and Engineering, 2005. **13**(2): p. R53.