Submitted to the IV International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2005), Utrecht, Holland.

Less is More: Agent-Based Simulation as a Powerful Learning Tool in Materials Science

Paulo Blikstein
Northwestern University[1]
paulo@northwestern.edu

Uri Wilensky
Northwestern University[1]
uri@northwestern.edu

ABSTRACT

This paper reports on a user study of a computer-aided learning environment for Materials Science. "MaterialSim" is an agent-based set of microworlds built in the NetLogo modeling environment, for investigating crystallization, solidification, metallic grain growth and annealing. Six undergraduate students enrolled in an introductory Materials Science course participated in the study, in which they could run experiments and build models. The rationale for the design is that the agent-based perspective may foster deeper understanding of the relevant scientific phenomena. A core feature is that students can apply a small number of local rules to capture fundamental causality structures underlying complex behaviors within a domain. We present evidence in the form of excerpts and samples of students' work, which demonstrates that experience with MaterialSim enabled them to identify and understand some of the unifying principles in Materials Science and build sophisticated new models based on those principles.

Categories and Subject Descriptors

K.3.1 [**Computer Uses in Education**]: *Computer-assisted instruction*

General Terms

Algorithms, Design, Human Factors, Theory.

Keywords

Agent-based modeling, Engineering education, Materials Science.

1. INTRODUCTION

Engineers, who frequently work at the interfaces of different fields of knowledge, are in a privileged position to integrate knowledge, be creative, and reinvent the world. In spite of this, engineering teaching and learning have remained remarkably traditional. Some universities have been proposing changes and introducing design-based curricula which include hands-on engineering design, intensively using modeling and simulation software packages [7, 8]. However, most of the tools used in such new initiatives are the same employed by professional engineers in their everyday practice and not especially designed for learning. For instance, these tools do not afford insight into the computation underlying their design and functioning. 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, traditional Material Science formulas *themselves* are opaque—they do not afford common-sense insight and grounding of the causal mechanisms underlying the phenomena they purport to capture. Thus, using formulas is perhaps conducive to successful *doing* but not to *understanding.*

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). To accomplish this goal, numerous equations, models, tools and heuristics have to be employed, oftentimes "mixing-and-matching" different levels of explanation to deliver the expected results. Our *modeling-for-understanding* framework, conversely, focuses on simple agent-level behaviors (atomic-level interactions, free energy minimization).

Throughout this work, we refer to one of the two common uses of the term agent. Our *agents* are employed in the context of *Multi-Agent Modeling Languages*. They are elementary computational constructs, with very simple rules governing their behavior, from which complex macroscopic behaviors emerge.

In this paper, we will present and discuss a user study of a system designed within the *modeling-for-understanding* framework: **MaterialSim**, a multi-agent Materials Science modeling environment built with the NetLogo language [13]. The study was comprised of classroom observations, pre/post interviews and data analysis from individual sessions with students using the system.

2. RESEARCH DESIGN & METHODS

Six sophomore year undergraduate students (volunteers) participated in the research, all enrolled in the "Microstructural Dynamics" undergraduate course in the engineering school of a liberal arts university in the spring of 2004. They were all computer literate and had had some contact with computer programming, typically MatLab. The MaterialSim's models explored by students during the study dealt with one particular phenomenon in Materials Science: *grain growth in metals*. We scheduled the sessions to happen between one and two weeks after students' exposure to this topic in their regular classes. Each student participated in two individual sessions. The first, one hour long, was comprised of a general presentation about NetLogo and MaterialSim, pre-test/interview about grain growth, and hands-on interaction with MaterialSim's environment, in which students performed experiments of their choice. As "homework", they

[1] Center for Connected Learning and Computer-Based Modeling
2120 Campus Drive – Evanston, IL, USA. +1 (847) 467 3859

were asked to devise a new feature for MaterialSim, which would be implemented on the next day. The second session (2 hours) was dedicated to a brief introduction to the NetLogo modeling language, the implementation (i.e., coding) of a new feature to MaterialSim, and a final interview.

All sessions were videotaped, and students' computer interactions were recorded using real-time continuous screen-capture software. Approximately 20 hours of video and 10 hours of screen capture were analyzed and selectively transcribed. Experiments done by students, as well as the models they built, were saved and analyzed. The first author attended three weeks of classes with students and analyzed the recommended materials and literature.

3. 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.

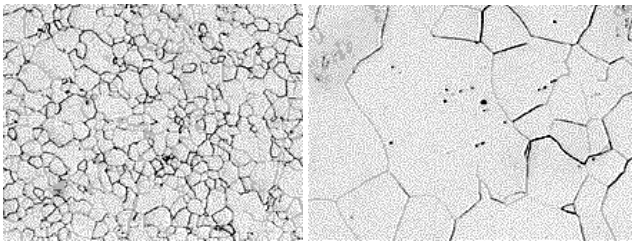

Figure 1: Metallic sample before and after grain growth [3]

Burke [6] 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 [3, 4]. Anderson, Srolovitz et al. [1, 2] 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." [2, p. 794]

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. The increasing importance of disciplines such as Systems Engineering and Complex Systems [9] are examples of the worth of this skill. Moreover, the introduction of digital tools is radically transforming engineering practice, but engineering education still does not benefit as much as it could from the new possibilities brought about by these new tools.

4. SOFTWARE DESIGN

4.1 The NetLogo environment

NetLogo [3] is an integrated multi-agent modeling environment. 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. 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.

4.2 MaterialSim design

MaterialSim is a set of exploratory models built in the NetLogo environment [13]. 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.

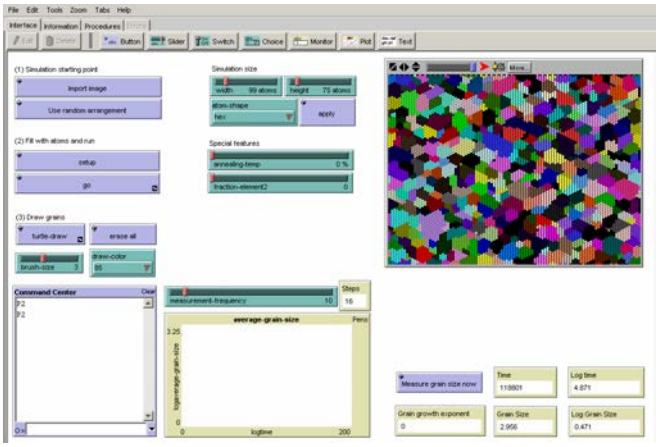

Figure 2. MaterialSim's grain growth model

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 [10] combined with a multi-agent approach, NetLogo's "low-threshold, no-ceiling" design enables learners quickly to achieve sophisticated results within a relatively short period of time. Moreover, its built-in visualization allows dynamic, flexible, and customizable views of model results. 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 physics [11, 12, 14-16].

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).

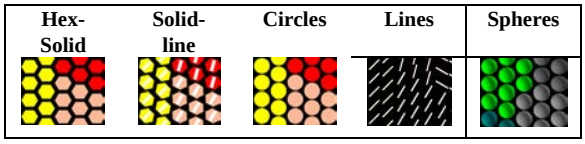

Figure 3. Different visualization modes

- 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.

5. APPROACHES TO GRAIN GROWTH LEARNING

5.1 Observed Classroom Practice

Grain growth is a good example of an area in which the teaching tools have not kept pace with the existing research tools. Even though the new research tools make a deeper understanding of the phenomenon 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 classes covering grain growth: the Laplace-Young equation for pressure, the flux equation based on statistical mechanics, and various geometrical approximations, as we have described in detail in previous work [5]. Later in this paper, we suggest that this multitude of models is an obstacle to student understanding.

5.2 Agent-based grain growth simulation

Agent-based simulation of grain evolution, on the other hand, offers a very different perspective. Its principle is both simple and fascinating: the thermodynamics of atomic interactions. 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:

a) 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_f**) is **4**.

b) 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".

c) 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**.

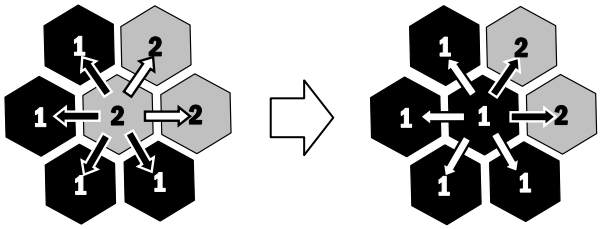

Figure 4: Initial and final free-energy calculations. Black and white arrow denote different or equal neighbors.

d) The two states are compared. The value that minimizes the free energy is chosen. In this case, $\mathbf{G_i=4}$ and $\mathbf{G_f=2}$, so the latter value is lower and constitutes a state of greater stability [3, 4].

6. DATA ANALYSIS

6.1 Pre-test/interview explanations

The pre-test consisted of two activities: first, students were asked to identify, in a printed picture of a metallic microstructure, grains that would grow or shrink, and then asked five questions.

Table 1: Pre-test/interview questions

A friend missed two weeks of classes. Now she/he is back and wants to catch up. How would you explain to her/him these ideas?
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 effect on grain growth of dispersed particles? Why?

The hypothetical situation ("a friend who missed classes") was an effective "ice-breaker" and enabled students not to feel that they were being tested. Below we have a commented transcription of some excerpts (the names were changed). The findings were consistent with the hypothesis that the multiple descriptive models used in their classroom instruction resulted in the students generating a diversity of explanations for the phenomena.

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, and [pause] you know, and 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 [pause] and if for some reason if they were the same direction they would be considered one grain, and since they are lined up there would be no boundary anymore.

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, you know, those kinds of pictures that we see, what comes to [my] mind... just one of those grains [pause] 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. So I suppose that is how I associate it with things.

These first excerpts already illustrate how dissimilar students' understanding was, which might be surprising as they are all Materials Science majors and have attended classes about the same topic a week earlier. Bob offers a reasonably complete explanation, based on crystallography, but says that when two grains had "the same direction they would be considered one grain", and mixes atoms with molecules. Betty, on the other hand, bases the explanation on a visual model, and integrates models from other topics in a vague way. Liz uses yet another explanation path, utilizing her previous knowledge about the morphology of a "real-world" grain to understand the Materials Science meaning of the word. Thus, 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). The responses to the third question further reinforce this diversity:

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 [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: What is the reason they grow?

Bob: Well, I mean, 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.

This third question brings about at least three different ways to explain grain growth. Bob uses the metaphor of willpower ("molecules come into the grain and line up"), and employs ideas about diffusion, dislocation and impurities in a 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 materials, "cleaning" it of impurities and dislocations. 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 not completely related to grain growth. During the pre-test, when presented with a printed picture of grains, she incorrectly indicated that the small ones would grow (which is what happens in recrystallization but not in grain growth). 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. 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 another consequence of thermal agitation, but unrelated to grain growth).

Liz's explanation is more coherent, relying on the idea of having "less" energy in the crystal being correlated to the "whole radius of curvature idea" (which is correct) but without demonstrating how those things connect. This analysis suggests that, in fact, students are mixing and matching different explanation fragments from a multitude of models from different sections of the course, with different granularities, levels of analysis, and mathematical descriptions. They make *ad-hoc* attempts to stitch together an explanation, and some achieve a reasonably coherent one – however, reading between the lines of those statements, we can glimpse their difficulty in identifying the fundamental driving force behind the phenomenon. The fifth question elicited students' "model mix-and-match" even clearer.

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.

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".

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. Constructivist theory would describe students as trying to accommodate the new knowledge into existing structures and representations. Indeed, as most knowledge in Materials Science is removed from everyday experience, students might be very receptive to the first "rules-of-thumb", mental models and explanations they come across in class, as a way to make such abstract ideas more acceptable.

6.2 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.

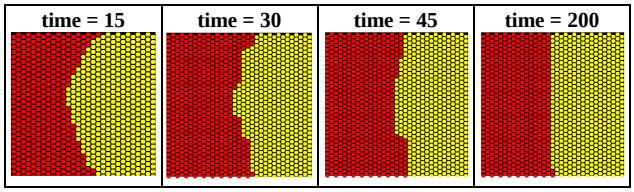

Figure 5. The evolution of a curved grain interface

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.

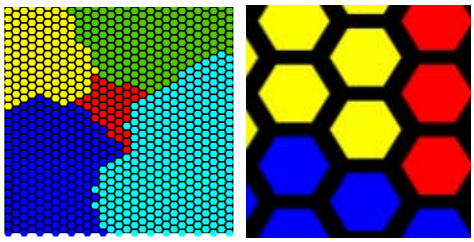

Figure 6. Four large grains (yellow, green, light and dark blue) surround a small red grain (left), and a zoomed-in view of the structure showing a triple point (right)

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 (see Figure 6, right side), 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 or 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.

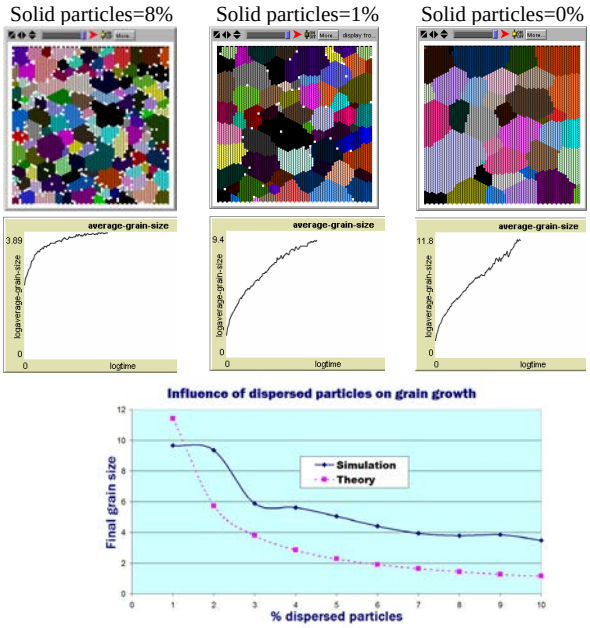

Figure 7. Sequence of screenshots from students' experiments

6.3 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.

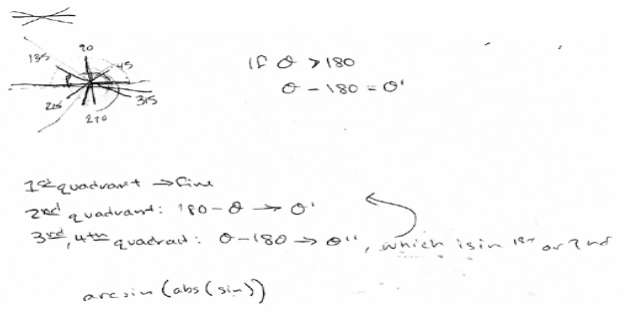

Figure 8: Betty's reflection about angles, sine and arcsine.

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.

7. 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 it 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, within a Constructionist learning environment, seems to be a better fit for the content areas discussed in this paper, for three reasons:

1) MaterialSim foregrounded the fundamental physical processes transpiring 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 *anchor 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, phase transformations), which are traditionally taught as *separate* topics with their own models and equations.

3) One of the pillars of Constructionist theory [10] is the importance of students conducting personally-meaningful projects. Coding their *own* models, students had an opportunity to test and debug their theories, as well as reconcile them with previous knowledge. Each student had a different learning trajectory, but that did not translate into misconceptions – our data, in fact, appears to indicate the opposite. 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 of just a few simple underlying rules of natural phenomena appears to be more generative for students than the more encapsulated aggregate, equation-based knowledge.

8. Acknowledgements

Prof. André Paulo Tschiptschin (University of São Paulo), Dor Abrahamson, Prof. Peter Voorhees (Northwestern University) and the students who participated in the study.

9. References

[1] Anderson, M.P., Srolovitz, D.J., Grest, G.S. and Sahni, P.S. Computer simulation of grain growth - II. Grain size distribution, topology and local dynamics. *Acta Metallurgica* (32). 793-801.

[2] Anderson, M.P., Srolovitz, D.J., Grest, G.S. and Sahni, P.S. Computer simulation of grain growth: I. Kinetics. *Acta Metallurgica* (32). 783-792.

[3] Blikstein, P. *Monte Carlo or the meaning of life.* Dept. of Materials Science and Metallurgical Engineering (USP), São Paulo, 1998. http://www.pmt.usp.br/paulob/montecarlo

[4] Blikstein, P. and Tschiptschin, A.P. Monte Carlo Simulation of Grain Growth. *Materials Research, 2* (3). 133-137.

[5] Blikstein, P. and Wilensky, U., MaterialSim: An Agent-Based Simulation Toolkit for Learning Materials Science. in *International Conference on Engineering Education*, (University of Florida, USA, 2004).

[6] Burke, J.E. *Trans. Am. Inst. Min. Engrs* (180). 73.

[7] Colgate, J.E., McKenna, A. and Ankenman, B. IDEA: implementing design throughout the curriculum at Northwestern. *International Journal of Engineering Education, 20* (2).

[8] Martin, F. Ideal and Real Systems: A study of notions of control in undergraduates who design robots. in Kafai, Y. and Resnick, M. eds. *Constructionism in Practice*, Lawrence Erlbaum Associates Inc, Mahwah, NJ, 1996, 297-332.

[9] Ottino, J.M. Engineering complex systems *Nature*, 2004, 399.

[10] Papert, S. *Mindstorms: children, computers, and powerful ideas*. Basic Books, New York, 1980.

[11] Stieff, M. and Wilensky, U. Connected Chemistry: Incorporating Interactive Simulations into the Chemistry Classroom. *Journal of Science Education and Technology, 12* (3). 285-302.

[12] Wilensky, U. GasLab-an Extensible Modeling Toolkit for Exploring Micro-and-Macro-Views of Gases. in Roberts, N., Feurzeig, W. and Hunter, B. eds. *Computer Modeling and Simulation in Science Education*, Springer Verlag, Berlin, 1999.

[13] Wilensky, U. NetLogo, Center for Connected Learning and Computer-Based Modeling. http://ccl.northwestern.edu/netlogo, Evanston, IL, 1999.

[14] Wilensky, U. and Reisman, K. Thinking like a Wolf, a Sheep or a Firefly: Learning Biology through Constructing and Testing Computational Theories. *Cognition & Instruction*.

[15] Wilensky, U. and Resnick, M. Thinking in levels: a dynamic systems perspective to making sense of the world. *Journal of Science Education and Technology, 1* (8). 3-18.

[16] Wolfram, S. *A new kind of science*. Wolfram Media, Champaign, IL, 2002.