Projectible Predicates in Analogue and Simulated Systems

James Mattingly (Georgetown University)
Walter Warwick (Micro Analysis & Design)

Abstract: We investigate the relationship between two approaches to modeling physical systems. On the first approach, simplifying assumptions are made about the level of detail we choose to represent in a computational simulation with an eye toward tractability. On the second approach simpler, analogue physical systems are considered that have more or less well-defined connections to systems of interest that are themselves too difficult to probe experimentally. Our interest here is in the connections between the artifacts of modeling that appear in these two approaches. We begin by outlining an important respect in which the two are essentially dissimilar and then propose a method whereby overcoming that dissimilarity by hand results in usefully analogous behavior.

We claim that progress can be made if we think of artifacts as clues to the projectible predicates proper to the models themselves. Our degree of control over the connection between interesting analogue physical systems and their targets arises from determining the projectible predicates in the analogue system through a combination of theory and experiment. To obtain a similar degree of control over the connection between large scale, distributed simulations of complex systems and their targets we must similarly determine the projectible predicates of the simulations themselves. In general theory will be too intractable to be of use, and so we advocate an experimental program for determining these predicates.

---

the object of the natural history I propose is . . . to give light to the discovery of causes and supply a suckling philosophy with its first food.

Francis Bacon, *The Great Instauration*

1. Introduction. Many years ago now Nelson Goodman attempted to explain, in part, what accounts for our choice of predicates in our descriptions of the natural world. He was animated by the realization that our explanations of the nature and legitimacy of causal relations, laws of nature, and counterfactuals all depend strongly on each other. The solution, as he saw it, was to investigate why certain predicates like green and blue were widely considered to be appropriate and adequate to our attempts to characterize natural happenings, while others like grue and bleen were not. This problem, as he presented it, was not one of logical definability and it could not be solved by identifying those predicates that are, in the case of grue/bleen versus green/blue for example, temporally indexed. The point is this: As a matter of mere description of the features of the world, there is very little constraint on the legitimate, completely general properties we can dream up while the true causal processes in nature, the true laws of nature, the true counterfactual dependencies in nature, all have to do with natural kinds; these kinds are picked out by what Goodman called projectible predicates.

Goodman's particular solution to the question of how to identify proper projectible predicates for new domains of inquiry need not concern us. Indeed there seems to be general consensus that his solution is wanting. It is enough for our purposes however that we keep in mind the general lesson: finding adequate adjectives to describe possible predicates is trivial; finding proper kinds is hard. This lesson applies equally well to computer simulations of human performance. Such simulations comprise multiple levels of resolution and are intended to help practitioners understand the complex and sometimes emergent behaviors that result when simulated agents interact with each other and their simulated environment. Unfortunately, despite the considerable skill practitioners' bring to bear in hand-tuning the underlying models, these simulations often produce spurious, even inscrutable behavior.

We argue below that much of the problem lies in the failure to identify and exploit principled abstractions that can structure and constrain the interactions of the component models. These abstractions are the reflection of what we choose to represent (and not represent) in our models. In the spirit of Goodman, we might think of these abstractions as the "projectible" predicates of the simulation. And here, too, not just any abstraction will do; it is not enough, for instance, simply to defer to the object-oriented abstractions imposed by the software architecture of the simulation. What is needed is serious consideration of the "nomological" features of the world we wish to simulate.

Our aim for what follows is simply to begin to think about some ways of finding projectible predicates for computer simulations that parallel those used in the construction and analysis of analogue systems in the physical sciences. We consider these because they show clearly what we desire in much of our simulation practices: the ability to mimic in a relatively controllable situation behavior of interest that arises in relatively uncontrollable situations.

We begin by discussing examples of some of the general kinds of problems that plague the simulation of human performance. The examples reveal breakdowns between the abstractions imposed at different levels of resolution within the simulation. In turn, these general problems point to two underlying questions: What are the laws of nature of these simulations (i.e., what makes behavior at one level of abstraction independent from behavior at other levels)? And how do we change only some of these laws in a way that stops short of encoding each and every feature of the simulation by hand? Let us pause here to emphasize and clarify this our orientation. We treat simulations *as* natural systems displaying lawlike phenomenal behavior. We bracket for now the question of the underlying source of this behavior and focus instead on attempting to understand the actual regularities displayed by the system. We treat these regularities as the laws of nature of the simulation. Then the answer to the above two questions, we believe, involves a shift in focus from the development of individual behavior models to a more deliberate exploration of the general abstractions on which those models depend. Again, speaking in the spirit of Goodman, this is a matter of finding and learning how to manipulate the projectible predicates of the simulation itself.

Next, as to how we might conduct such an exploration, we turn to a discussion of a common technique in the physical sciences: analogue modeling. Analogue modeling bears important resemblances to other kinds of modeling in physics, but has a unique flavor that may offer some insight for difficult conceptual problems in the simulation of human performance. In particular, we take analogue models to be themselves a type of simulation. We focus on cosmological analogues in Bose Einstein condensates. These are interesting analogue systems, but are also nifty because they occupy about as extreme a separation of scale and resolution as possible. We note that the artifacts of the one system (the Bose Einstein condensate) are features of the other (the universe as a whole). This leads naturally into a more general discussion of ontology: the ontology of the target system; the ontology of the analogue system. We begin to ask here about the laws of nature of the analogue system itself, and the laws that the system is meant to represent. In analogue systems the *fundamental* laws are still the real laws of nature, but the utility of a given analogue comes from seeing it also as a different natural system embodying different natural laws.

We then return to questions about the simulation of human performance. In the analogue systems we employ a blend of mathematical analysis and experiment to fix the abstractions of the simulation. But we have seen in the history of large scale distributed computer simulations of human performance that they remain intractable to analysis through theory alone. We therefore call for a general program of experimental computer simulation. That is we call for a program that takes as given certain exemplar simulations and asks what regularities they actually display and what counterfactuals they actually support. But two major problems remain: How do we connect the projectible predicates of the simulation to those that are of interest to us? Is it really possible to manipulate these predicates without changing the basic underlying code, and thus vitiating the whole project? We conclude pessimistically. We think the general approach we advocate, an experimental program for computer science, is worth pursuing, but we see little hope for immediate payoff. The situation seems now like that confronting Bacon when he advocated simply performing all of the experiments there are, and thereby learning all of nature's laws. If we knew which kinds of experiment were really worth doing, it would be because we had a better handle on the plausible projectible properties and abstractions. Here we are content to clarify our proposal and to make plausible the utility of pursuing it.

2. Breakdowns in the computer simulation of human performance. Like any other large software application, computer simulations often produce unexpected behaviors. Sometimes these behaviors are the result of software bugs that are bad enough to cause spurious behavior—e.g., a simulated agent stops moving or interacting with its environment—but not so bad to cause the entire simulation program to crash. Although unexpected, such behaviors are not so much artifacts of the simulation as they are products of the inevitably imperfect software development process. But the same behavior, from a phenomenological perspective, might also be caused by software functioning exactly as it was designed to. Whether we consider a given behavior to be a bug or an artifact depends as much on the cause of the behavior as it does on the nature of the behavior itself. Indeed, it is one thing for an entity to freeze in a simulation because, say, of a missing semi-colon or because an array index was improperly initialized and quite another when that behavior follows from the interaction of software components that are functioning exactly according to specification.

It might be tempting here to think about the difference between a "programming error" and a "logic error;" a bug results from a programming error while an artifact is the result of a logic error. But that distinction exists only because we all agree where the programming stops and the logic starts, so to speak. To return to a more literal invocation of Goodman, both programming error and logic error are projectible predicates in the context of computer simulation—they are robust classes of problems that support counterfactual reasoning (e.g., "If I initialize the array index to zero, I'll avoid the overflow error"; "Incrementing the resource count before advancing the simulation clock led to some entities having extra resources during the simulation"). The more interesting problems are a subset of the logic errors, those that arise when we're forced to rethink our abstractions after the fact, when it turns out that what we took to be a "projectible" predicate in our simulation isn't, or worse, can't be given our representational assumptions. These are the problems that occur because we have no principled way of deciding what to represent in our simulations and, hence, no means for ensuring that abstractions at any one level of resolution are appropriately independent from the abstractions at other levels.

Each of the following three examples reveals a breakdown in the layered abstractions we use in the simulation of human performance—each more problematic than the last. In every case we're confronted with a simulation artifact—an interaction gone awry between a simulated entity and its simulated environment—with no easy solutions in sight. The examples all follow from both our direct and indirect experiences working with in the US Army's OneSAF Test Bed, v.2 simulation environment (OTB). OTB is a large-scale, distributed simulation environment designed to address a wide range of simulation needs (e.g., simulation-based analysis, simulation-based training, simulation-based acquisition, etc.). The code base comprises literally hundreds of libraries and over a million lines of software code. OTB simulates entity-level interactions (e.g., two opposing soldiers shooting at each other) using a number of behavior models together with a wide variety of physics-based models to represent features of the environment (e.g., weather and terrain features, vehicle characteristics, ballistics etc). Suffice it to say, the simulated world of OTB is both rich and complex. At the same time, behaviors in OTB can be quite brittle.

2.1. Target detection. One of the most primitive, yet important, behaviors for a simulated agent is the ability to "perceive" its environment. Of course simulated agents don't see anything directly, rather various target detection algorithms are used to determine the probability that one agent will see another. Darken and Jones (2007) describe their attempts to improve the most common target detection algorithms. Along the way, they also describe some interesting artifacts.

For example, they point out that the most basic target detection algorithms depend on a line-of-sight calculation (i.e., ray tracing) from the simulated agent to the *top* of the target. As a rough and ready solution, such an approach ensures that a gross representation of inter-visibility is generated—that agents are not visible behind mountains, buildings or other large objects—but it also means it's possible for an agent to hide while standing under an umbrella or, conversely, to betray his position by exposing a millimeter of his simulated scalp. Obviously, this can lead to seemingly strange engagement behaviors as relatively well obscured targets are quickly identified while far more obvious targets are ignored.

More sophisticated algorithms have been implemented that include consideration of target brightness, background brightness and size of the visible portion of the target. While these more sophisticated approaches avoid some of the obvious pitfalls of the line-of-sight algorithms, Darken and Jones point out that a whole new class of artifacts arises. For example, relative contrast between the target and the background can be confounded by the averaging routines used to represent color and brightness in the simulation. Thus, a grey target might become invisible when placed before a black and white checkerboard background. Likewise, the color (independent of brightness), shape, and texture of a target do not figure into these more sophisticated detection algorithms, even though all of these qualities are known to have a significant effect on human target detection. A hunter safety vest makes fine camouflage while visual clutter does nothing to distract simulated agents from their targets. (Darken and Jones, p102)

While it is possible to imagine yet better algorithms that could take these factors into account, the salient point here is that this process of improvement is and will continue to be ad-hoc. A priori, it is very hard to judge whether ray tracing to the top of a target should be preferred to ray tracing to, say, center-of- mass. It is harder still to anticipate how standard methods for computing contrast might confound agent behaviors. In short, there is very little to guide our efforts in determining how we should represent one of the most basic behaviors of a simulated agent. Moreover, how can we avoid having to worry about the lower-level implementation of, say, contrast computations when our focus is on the ability of a simulated agent to detect a target. Our choices have profound consequence for agent behavior, but more often than not these become clear only after the fact.

2.2 Basic movement As an entity-level simulation, each simulated soldier is imbued with its own (i.e., a local instance of) basic route-finding and movement routines. Given a goal-location to move to, these routines operate over a representation of both the terrain and the objects in it to plan a route through space and time to the location. Given the constraints present in the simulated world—speed of movement, distance to target, possibility of terrain—not every target location will be reachable, which one would expect. But there are other instances when locations turn out to be unreachable because the soldier gets stuck within a wall, or the ground, and simply cannot move or because upon arriving at the target location the solider would essentially be left hovering in the air. The problem here isn't one of avoiding simulated obstructions but rather of representing the obvious physics—two objects can't occupy the same space, and soldiers can't fly. Of course, some of the relevant physics have been represented, but they come into play at exactly the wrong time, only after the entity has moved into but not out of an impossible position.

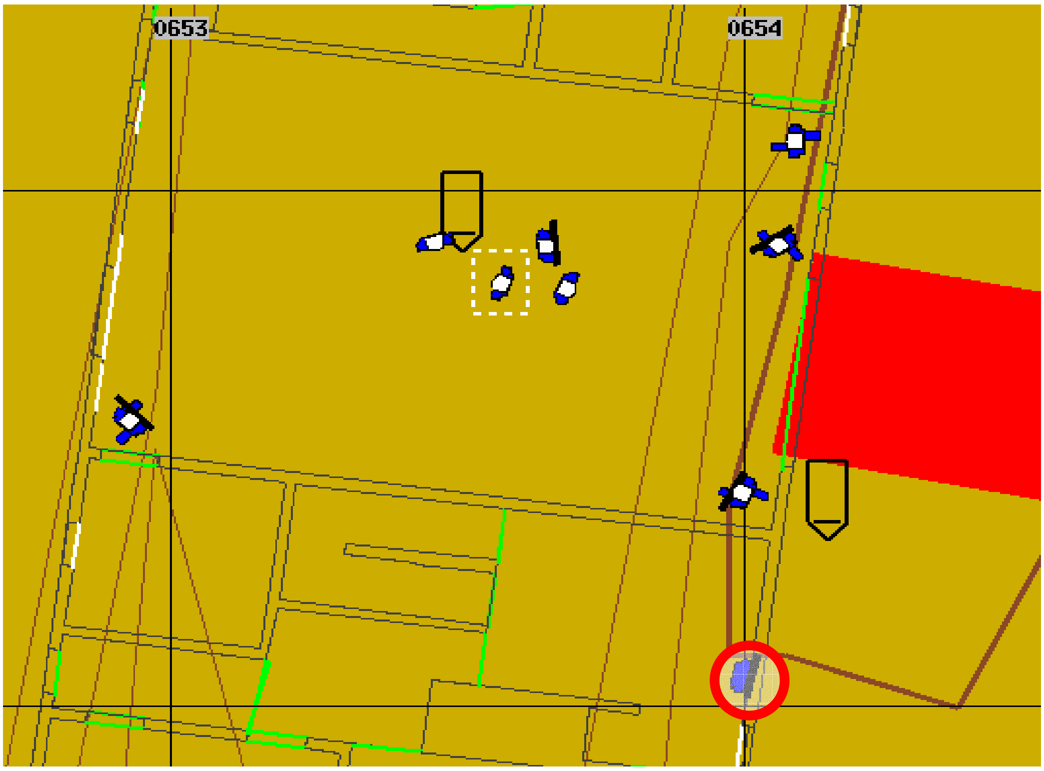

Figure 1. A screen shot of pathological entity behavior in OTB. Here the highlighted entity has moved into the room only to become "stuck" in the wall. The simulation continues despite the fact that one entity has moved into a physically impossible location.

Although such breakdowns are treated within the OTB development community as software "defects" and remedied with various workarounds (e.g., by opening up an entity editor by hand and adjusting the x-y-z location of the soldier), the fundamental problem is one of a mixed fidelity representation where some, but not all of the relevant detail has been represented. So, again, we ask what's the heuristic we should apply for what gets represented and what doesn't? Obviously, we don't want different objects occupying the same space, unless, of course, it's the simulated gas cloud through which a simulated solider must pass. We might treat the cloud as an object extended in space, but not the sort that prevents other objects from occupying the same space. And even for other hard bodies, we might want some of them to be squeezeable—like clowns into a car—and others to be not-so-squeezable—like 70-ton M1A1 tanks crossing between concrete barricades. A good deal of detail starts to stand out as we think more and more about the seemingly obvious physics of object interactions. At the same time, we cannot simply add more and more detail to our simulation model without running the risk of making the whole enterprise intractable.

The conventional wisdom in the simulation community holds that the fidelity of the simulation should be driven by the questions the simulation is intended to address. For example, there is no need to represent the color of a driver's socks if the simulation is intended to answer questions about traffic flow on an interstate highway. As an admonition to the over-eager model developer, there is something right about the conventional wisdom; there is no reason to add extraneous detail to models. But how is one to know *a-prior* which detail will be extraneous? The examples above demonstrate how a seemingly level headed approach to the representation of bodies interacting in a simulated world goes wrong. In fact, the examples above betray a pair of competing orientations. The view from the top holds that we should model only enough detail to get the phenomenology of the simulation right—that is, we simulate just enough of the "physics" to engender only those interactions that reflect the salient features of the situations we're simulating. Thus we model humans as hard objects. The view from the bottom, however, starts with a small handful of computational techniques that can be used to manifest various simulated behaviors and hopes that these will in fact produce all and only those interactions needed to get the phenomenology right. Hard objects can't occupy the same spaces, unless they happen to, and then they are stuck. What is needed here are some principled connections between the two views—how do we ensure that our simulated interactions will behave in the right way without recapitulating all the actual physics in our simulated world?

2.3. Tactical movement. But even if we could identify and supplement the OTB environment with the right amount of physical detail, problems still remain at the behavioral level—that is, the level at which the simulated humans are intended to think and act like people. The problem here isn't behavior qua physics but, rather, a question of balancing the behavior of simulated people as both objects and agents. Consider for example what happens when a group of simulated soldiers must move in a coordinated manner.

Such behaviors are especially brittle in and around buildings. For instance, a group of simulated soldiers that march into an obstructed hallway will often get stuck—permanently—because the first soldier in the line can't clear the obstruction and can't back out and the soldier behind him is still trying to move forward, and so on down the line. In this case the "physics" are fine insofar as the soldiers can't pass around each other in the constrained interior space, but agent behavior is lacking. There is no representation of "excuse me" built into the simulated soldier. Similarly, there are cases when simulated soldiers should move in formation and for whatever reason one soldier in the formation will find his route obstructed. This immediately spawns a route re-planning task for that agent, which will often get the agent back to his original place in the formation, but only after a long, seemingly inexplicable detour through enemy territory. Sometimes a single will soldier will fail to move at all if it happens that constraints imposed by moving in formation cannot be satisfied given a particular configuration of geography and inter soldier spacing. Rather than tighten the spacing to his buddy the soldiers simply stays put.

In these cases, the richness of the physical representation actually comes at the cost of the agent-level representation. Indeed, the route-planning behaviors of the agents are quite sophisticated, using ground-truth information (i.e., variables used in the simulation that represent the physics of the simulated world—locations, sizes, velocities etc) about stationary and moving obstacles to plan and dynamically adjust routes through both time and space. And yet, despite this sophistication, or perhaps because of it, there's little code written in to determine when a plan needs to be adjusted (e.g., tighten formation rather halt movement) or how it might be communicated (e.g., "excuse me").

The point here is that getting the physics right enough, as hard as that is, still isn't enough to capture the right kinds of agent interactions. And in the case of military simulations, these so-called tactical movements are a highly salient feature of the actual performance of human soldiers and, hence, should be a central feature of soldier simulations. Thus, the problem of identifying useful abstraction recurs at a new level and our efforts must be redoubled. Just as there is a question of how much physics to include, there is also a question of how much agent-based behavior to include, keeping in mind, again, that we can't nor would we even want to include every imaginable agent behavior.

The problems we've just described are all the result of breakdowns in the layered abstractions we impose. None of the individual problems are insurmountable; extra detail can always be included in the underlying models just as one-off patches to higher level behaviors can always implemented. But, in general, such fixes are symptomatic of a deeper issue: namely, we expect the assumptions we make at one level of abstraction to be independent from those we make at other layers of abstraction, and yet we have very little in the way of principled guidance to inform those assumptions. In fact, there's something of an embarrassment of riches at play here. Computers are nothing but tools for supporting layered abstractions—thus physical circuits become logical circuits, logical circuits become hardware implementations which, in turn support an ever growing hierarchy of increasingly abstract software languages. But because so many abstractions are possible, it makes the job of identifying those most appropriate to a human performance simulation all the harder.

3. The idea of an analogue system. So how do we go about deciding what to include and what not to include in our simulations? How do we ensure that representation we impose at one level of abstraction don't end up confounding representations we make at other levels of abstraction? To answer these questions we turn to an unlikely example.

Cosmologists are hampered by a significant obstacle: They cannot conduct experiments to test their models. To overcome this difficulty they have had recourse to prolonged observation, and intensive theoretical analyses. But these do not completely overcome the necessity for actual experimental feedback in ruling out theories and suggesting new classes of theory.

It has lately been realized that some classes of quasi-experiments, observing and manipulating systems that are analogous in appropriate ways to the universe as a whole, would, if they could be performed, provide important experimental data to cosmologists. Unruh (1995) has shown, for example, that one can model black holes by sinks in classical fluids---the so-called dumb-holes. Moreover some features of Hawking radiation can be modeled---waves traveling out of the hole even though the fluid flow is faster than the speed of water waves. But many such classes of quasi-experiment themselves suffer by being composed mostly of experiments that are themselves too difficult to perform---perhaps impossible even in principle involving sensors that may themselves be unphysical, or some other insurmountable obstacle.[2] However there are some that are clearly performable in principle and of those some appear to be performable with present levels of technology.

As a particular example we consider a Bose-Einstein condensate, which we describe

---

[2] For analysis of a similar problem in the context of semiclassical gravity see (Mattingly 2006).

shortly, as an analogue of the universe as a whole. The point to the analogue is to test the predictions of a semiclassical theory of quantum gravity indirectly by giving experimental access to various parameters that are not fixed in the general theory. Seeing how the analogue system changes in response to varying these parameters, coupled with observation of the cosmos constitutes, effectively, a cosmological experiment. Semiclassical gravity is the current de facto theory of quantum gravity and is widely used to guide theory construction in the quest for a more principled future quantum gravity. For example, the behavior of black holes predicted by semiclassical gravity is a minimum standard for any candidate theory of quantum gravity and quantum cosmology. If that candidate's predictions differ in the wrong way from those of the semiclassical theory, then it's off the table. Thus an experimental test of semiclassical gravity theory will give empirical input into quantum gravity itself---input that is sorely lacking to date.

3.1 Bose-Einstein condensates. Bose Einstein condensates are predicted by quantum mechanics. In quantum mechanics the statistical distribution of matter is governed by two distinct theories of counting for two distinct types of matter. Every material system possesses, according to the quantum theory, an intrinsic angular momentum. That is, every material system possesses angular momentum that arises not from any mechanical movement of the system, but merely due to its composition. This angular momentum can take on values that are either half-integer multiples of Planck's constant or whole-integer multiples. Systems with half-integer intrinsic momentum (fermions) are governed by Fermi-Dirac statistics; those with whole-integer intrinsic momentum (bosons) are governed by Bose-Einstein statistics. These two different statistics turn out to have significant consequences for the behavior of large collections of the various type of entity. The basic idea of the two classes of statistics are well-known. Fermions cannot all be in the same quantum state; bosons may all be in the same state. A Bose-Einstein condensate is the state of a collection of bosons that are all in the same state together. Since they all share their quantum state, there is no difference between the elements composing the condensate---the condensate behaves as though it were a single object.

Since 1995 and the production of a Bose-Einstein condensate in the gaseous state by Cornell and Wiemann, many physicists have become interested in these systems as possible experimental test-beds for studying quantum cosmology. This is extraordinary on its face. What could be less like the universe with its distribution of objects on every length scale and its curved spacetime geometry than a small container of gas (on the order of $10^{9—10}$ atoms) with fluctuations in the phase velocity of sound propagating through it? And yet one can find analogous behaviors in these systems that make the one an appropriate experimental system for probing features of the other. One feature of interest in cosmological models governed by semiclassical theories is pair-production caused by the expansion of the universe[3]. Barceló, Liberati, and Visser (2003) have shown how to manipulate a Bose Einstein condensate in such a way that it will mimic certain features of an expanding universe exhibiting semiclassical particle production. That is, they show how to mimic in a Bose Einstein condensate a semiclassical scalar field propagating in spacetime that produces particle pairs as the universe expands.

[3] This is discussed in Birrell and Davies (1982) for example. Many interesting features of semiclassical models have to do with particle production under various circumstances. One reason for their interest is that these are features we can imagine actually observing in the cosmos.

It is well known to theorists of Bose-Einstein condensates that all of their important features can be captured in the Gross-Pitaewskii equation:

\begin{equation}
i\hbar \frac{\partial}{\partial t}\psi(t,\bm{x})= \left(-\frac{\hbar^2}{2m}\nabla^2 + V_{ext}(\bm{x}) + \lambda |\psi(t, \bm{x})|^2\right) \psi(t,\bm{x})
\end{equation}

This is a non-linear approximation to the Schrödinger equation with the self-interaction term given by a function of the modulo square of the wave function. In their proposed setup, Barceló, Liberati, and Visser propose a series of generalizations to this equation. By allowing arbitrary orders of the modulo square of the wave function, by allowing the non-linearity to be space and time dependent, by allowing the mass to be a tensor of third rank, by allowing that to be space and time dependent as well, and finally by allowing the external potential to be time dependent, they arrive at a new Schrödinger equation:

\begin{equation}
i\hbar \frac{\partial}{\partial t} \psi(t,\mathbf{x}) =-\frac{\hbar^2}{2\mu}\Delta_\hbar \psi(t, \mathbf{x})-\frac{\xi\hbar^2}{2\mu} ~^{(3)} R(h) \psi(t,\bm{x}) + V_{ext}(t,\bm{x}) + \pi' (\psi^\ast\psi)\psi(t,\bm{x})
\end{equation}

And this equation has characteristics that allow it to be cast into the form that describes perturbations in the wave function propagating through an effective, dynamical Lorentzian metric. With a suitable form for the potentials one can use this equation to replicate a general-relativistic spacetime geometry.

It is also possible to show that, in the regimes of the experimental setup they identify, the Bose-Einstein condensate mimics very well the behavior as a whole of the expanding universe, and especially the behavior of scalar fields propagating in that universe. As the interaction between the components of the condensate is modified, the effective scattering length changes, and these changes are equivalent in their effect to the expanding universe. Under that "expansion" these scalar fields will exhibit pair production. And Barceló, Liberati, Visser give good reason to suppose that actual experimental tests can be conducted, in the near future, in these regimes. Thus the Bose-Einstein condensates are appropriate analogue models for the experimental study of important aspects of semiclassical cosmology. We can therefore use the condensate to probe the details of cosmological features of the universe, even though the analogue system has very little qualitative similarity to the universe as a whole. (For example the condensate isn't really exapanding.)

We now pull back for a moment and try to get a clearer picture of analogue systems. The general idea of these systems is this. We use actual physical systems to investigate the behavior of other physical systems. Stated in this way, the point appears trivial. Isn't this no more than just plain old experimental physics? What of significance is added when we call the experimental situation an analogue? Aren't all experiments analogues in this sense? We can answer the question in the negative by being more precise about the nature of analogue models. In a completely broad sense it is true that all experimental systems are themselves analogue systems---unless all we are interested in probing is the actual token system on which we are experimenting. When we experiment we allow one system to stand in for another system that differs from the first in various ways. If these two systems are not token identical then they merely analogous, being related by something other than strict identity.

That is correct as far as it goes, but in the vast majority of cases, the experimental system is related to the target system by something like a similarity transformation. That is to say that generally we have to do with changes of scale, or with approximations, or suppressing certain parameters in constructing the experimental system. So for example in precision tests of Newtonian particle physics we will attempt to find experimental systems for which the inevitable finite size of the bodies will not relevantly change the results of the test. We see that taking the limit of smaller and smaller particles does not change the results to the precision under test. In this case we have a system that approximates the target system by the continuous change in the value of a parameter as that value approaches zero. This kind of thing is quite standard. We attempt to suppress effects due to the idiosyncratic character of the actual systems with which we have to deal, character that tends to deviate from that of the target system in more or less regular ways.

Analogue systems in their full generality are not like that. These systems are not necessarily similar to the target systems they are analogues for. In the general case analogue systems are neither subsystems of the systems of interest, nor are they in any clear sense approximations to such subsystems (as billiard balls might be to Newtonian particles). The laws that operate in these systems are not the laws operative in the target systems. That last remark is too fast, of course. Rather we should say that even though the laws of physics are the same for all physical systems, the phenomenal features in the analogue that are being taken as analogous to those of the target system arise from very different effects of the laws of nature than they do in the target system.

In cases of interest the proximate physical causes differ markedly. Consider the following example: When speech causes a human body to perform some action---say the kicking of a leg under a doctor's orders for example---the relevant explanation is of a radically different character than when the cause is a direct physical manipulation---say when the doctor strikes the patellar tendon with a hammer for example. In both cases of course the ultimate causal agency is (overwhelmingly) the electromagnetic fields. But the salient causes are quite different.

The appropriate description of the causes that are operative in an analogue system, even though merely artifactual features of that system, are what we mean by projectible predicates in this context. Even though we use suggestive terminology, the fact is that our normal predicates do not obviously apply in these cases. We have merely identified sub-systems (that is, isolated regimes) of the analogue systems that support counterfactual, causal descriptions appropriate to our interests as modelers. These sub-systems can provide useful insight into their target systems only if their behavior is stable in the right way. And the right way is that they are independent of the finer and grosser details of the underlying systems of which they are parts; the sub-systems, *as* proper analogues, must be protected from the effects of significant changes of the super-systems.

Look again at the Bose-Einstein condensate. That analogue system is a strange one. The target of the simulation is a continuous classical spacetime metric that is coupled to the expectation value of a quantum field. This is being simulated by the analogue system of a single, unified quantum state supporting classical sound waves. As we saw, Barcelo, Liberati, and Visser generalize the governing equations for the Bose Einstein condensate by proposing new potentials, and then show that the new system is governed by equations of motion that are essentially non-relativistic but which encode a Lorentzian spacetime geometry. Their formal analysis allows one to consider the metric encoded by these equations of motion to be dynamical.

However the metric of the actual space they consider is non-dynamical across the span of the system. The processes operative there are radically, qualitatively unlike those of the semiclassical Einstein equation. Instead the "metric" is really a feature of the tensorial mass distribution. So the similarity is neither by approximation nor by suppression of parameters; instead it is something other. This is more like a simulating ideal particle mechanics using standing waves in a river, than with billiard balls say. And we could see "particle" creation in such a simulation too---some dip and some hump might emerge from the same location and move off scene. Here the connection between the simulation and its target is as indirect as that of the leg kicking case. The behavior is being caused in the one case by peculiar features of the condensate, and in the other by the interaction of the spacetime geometry with a quantum field[4]. We have a system with new "physical laws" that are merely artifacts of the analogue system. And it is those artifactual laws that we hope will shed light on the target system, the universe as a whole.

To emphasize the point: Even the descriptive terminology we use to apply to the Bose-Einstein condensate is merely artifactual. We have a "mass" term, and we talk about "particles", and the "metric" but these are no more than names that we apply to certain projectible features of the condensate to indicate the analogical role they will play in our later analysis. The key work is in finding the stable features of the condensate, identifying them with features of interest in the cosmos, and showing that the sub-system in which these features are the operant causes are independent of the vagaries of the super-system.

Before closing this section we point out one further, important conclusion to be drawn about this class of experiment: The members of this class are much closer to what is normally called a simulation, than to more common model systems. That is, we try to mimic the behavior of the target system by features of the simulation that arise from qualitatively dissimilar causes.

[4] The simulation is even farther removed from the target system of course because we need also to identify the caused behavior by analogy as well. We don't address that issue here.

4. Natural, analogue, and simulated laws. Our basic orientation is that it is fruitful to view computer simulation as a kind of analogue experiment. Returning to Goodman's distinction between projectible and non-projectible predicates, we can make the following observation. In general, in computer simulations, and in analogue simulation more generally, we do not know the natural kinds and very often we do not have a good sense for what predicates of the simulation are projectible. More seriously we do not know what the relation is between the projectible predicates of the one system and those of the other---especially since we are often trying to use the simulation to find those predicates that are the most fruitful for framing our hypotheses.

4.1 Theories of Experiment. To get clear on the issues that face us, it will be worthwhile to introduce some categories that will allow us to relate the general aspects of analogue physical systems to those of computer simulations.[5] What we are trying to do is relate one system to another by some mathematical relation between them. Generally we have a theory of some experimental set-up, including a method for turning our finite data into mathematical models involving perhaps continuous values of various parameters. To investigate some system type by conducting an experiment we abstract away irrelevant features of the experimental setup using our theory of the experiment and construct this data model. We can then compare that model to the various mathematical models that arise in our theoretical descriptions of the target system.

In the more general case of analogue systems there is the added complication of restricting the

[5] The reader may notice here some similarity to the relation between the physical analogue systems and the target systems in Suppes' framework from "Models of Data" (1962).

parts of the experimental system we allow to be relevant despite being just as significant in magnitude and scope as other parts. For example in the Bose-Einstein condensate we pay attention only to certain vibrational modes even though other modes are present and impact the state of the system---we do not identify *those* modes with particle anti-particle pairs, we do identify *these* modes with them.

Even with these complications however the basic framework is the same. We have two systems, the target and the experiment. To bring the results of our experiment to bear on the target we need a more or less explicit theory of that experiment---less in cases where the similarity between the experiment and the target is obvious, more in cases where the connection is more tenuous. An experiment in Newtonian particle mechanics using a billiard table needs a much less explicit theory of the experiment than an experiment in quantum cosmology using a Bose-Einstein condensate.

We have seen that even in the latter case it is possible to provide a very general analysis of the causal artifacts of the system and their complicated interactions with each other. We did not discuss the breakdown of the model, but it turns out that the system can be made stable and the analogy itself can be made stable across a wide range of parameter values---and most importantly, that range is known.

Our theory of the experiment in this case functions as it should despite the fact that the underlying features of the experimental system are so far removed from the salient, artifactual, causal agents of the experiment *as* a model of the target system

We have seen however that things are not so clear in the case of interacting systems modeled by large scale computer simulations of human performance. While there is a tight qualitative similarity between the classes of artifact that arise there, we have very little control over the higher level artifacts of the simulation, and what control we do have requires tedious case by case attention. Looking to independence results in analogue modeling suggests how this might be achieved for the simulated physics. And in a similar vein, Warwick and Napravnik (2004) have outlined a method for standardizing the agent behaviors in OTB by fixing a "catalog" of perceptions and actions that a simulated agent has access to. This effort circumscribes agent behaviors more explicitly than is currently done within OTB and thereby makes it easier to localize shortcomings in behaviors when they occur.

Still, what is needed is a principled method for determining what should, and should not be represented in a simulation to achieve the appropriate level of fidelity. A theory of the computer simulation as an experiment would not only define, say, a particular catalog of perceptions and actions, but would also provide, among other things, the justification for using one catalog rather than another. Unfortunately such considerations are far from the minds of most human performance modelers.

This is most clearly seen by switching gears from discussion of a large-scale, distributed simulation human performance, to that of the small self-contained representations of cognitive modeling. In this context, questions about the physics of a simulated world are largely absent, while questions about agent interactions with that world or other agents are often tightly constrained (cf. Anderson and Lebiere (1998)). In most cases, the goal of a cognitive model is to develop a computational representation of whatever it is the researcher believes is going on "inside the head" of human agents as they perform various tasks—often those drawn from the experimental psychology literature—and to compare the performance of that model against human performance. While there are a host of serious debates within the field about the nature of cognition (e.g., symbolic versus connectionist views), there is a surprising unanimity when it comes to the decidedly hypothetico deductive face cognitive modelers present when discussing their enterprise in general: computational models implement theories of cognition, simulations generate predictions which either confirm or disconfirm theories and thereby support explanations of cognition. Unfortunately, for all its homespun appeal, this view of the cognitive modeling enterprise is undercut to the extent that decisions about what gets represented in a cognitive model are largely up for grabs.

Warwick and Fleetwood (2006) make exactly this point by showing how three qualitatively different approaches to the modeling of a classic paradigm in cognitive psychology can yield quantitatively similar predictions about human performance. In short, the predictions severely underdetermine the choice between the simulated cognitive mechanisms that generate them. One might respond along the lines drawn by Smith and Minda (2000) who questioned whether this particular task is sufficiently rich to discriminate among models. Or, conversely, one might revisit Roberts and Pashler's (2000) concerns that human performance data is so squirrelly that drawing any conclusion about cognition is unwarranted. One might take a pragmatic angle and argue that having different models producing equally good predictions is a boon to simulation; so long as the phenomenology is right within the simulation, who cares how the behaviors were generated.

For our purposes none of these responses hit the mark. The problem here is not necessarily with the task being modeled or the data used to evaluate the model's fit or whether the human behavior modeler has enough tools at his disposal. Rather, the problem here is rooted in the flexible levels of representation that are so readily supported within computational models of cognition. Once again, without some principled means for identifying what should and shouldn't be represented in the model, we cannot use the models for generating predictions or explanations. Nor do we have any guarantee that a particular model that happens to produce reasonable predictions in one case will generalize to a new case. The ability to predict, explain and generalize all follow from our ability to identify what's doing the work in generating the model's behavior. As far as we're concerned, artifactual behavior might as well be accidental behavior if we can't point to the specific mechanisms that generate the behavior. In fact, this points to one of the dirty secrets of cognitive modeling, namely, that the goodness of the model is often not so much a reflection of the theory being implemented but rather of the skill of the modeler in representing the task given a particular modeling architecture. Worse, this problem is not limited to the examples we cite here, but has been identified in other cases where model "bake-offs" had been pursued (cf., Gluck and Pew 2001) and have led to similarly confounding results.

Cognitive modeling and human performance modeling are all about experiments, but the wrong kind. Rather than trying to identify correct theories of human cognition, greater emphasis should be placed on understanding how representational choices impact computer simulations. Such simulations are sufficiently complex that such understanding has yet and is unlikely to yield to pure analysis. Instead we advocate an experimental program for finding appropriate theories of experiments employing the artifacts arising out of the underlying code in computer simulations.

We suggest that only in this way can we insulate our experimental setup from peculiarities of the interaction of the bits of code themselves, and isolate the interactions of the artifacts that comprise the experiment. And only by so doing can we find the projectible predicates in terms of which we can develop the operant laws of the simulation.

This marks a complete reversal from the current methodology in human performance simulations. Rather than view the computer as a virtual sand box in which we can advance our theories, we see the simulation as a world unto itself that must understood on its own terms before experiments can be properly conducted. In the case of analogue modeling in physics, we are forced to confront and reconcile the manifest disconnects between the target system and the simulated system. We can span this disconnect only because we lean our hard-earned understanding of the operative physics. Too often the very same disconnect in computer simulation is ignored merely because is doesn't occur to us that there could be undiscovered physics in the world we created.

5. Conclusion.

Computer simulation is *essentially* analogical in a strong sense. The laws of the simulation are never those of the simulated system. They are always the laws of the actual computer we use. Moving from Newtonian mechanics and the process of human cognition, through our equations of motion for the mechanics conjoined with a going theory of the mind, and on to the 2-d image being observed and manipulated by a user of one of these large modular platforms requires covering a lot of very rough conceptual terrain.

We saw in the case of simulation by means of analogue physical systems how important it is to have a theory of the experiment, and concomitantly, a theory showing that and how the relevant parts of that analogue operate independently of the other parts of the system within well-specified limits. All of this amounts to having a good theory of what predicates ("mass" say, or "metric") we may successfully project in the analogue system itself.

We are far from having such theories in the case of many of the large-scale distributed environments used for the simulation of human/environment interaction. But these theories and their attendant battery of projectible predicates are sorely needed for these environments. We have suggested that a first step in that direction is to explicitly contrast the work done in computer simulation with that done in physics simulation by means of analogue systems. In particular, since generally we lack the theoretical tools needed to analyze the connection between the user level output simulation and the input at the level of code, we advocate an experimental program that will try to chart and catalogue some of these connections empirically.

Partly the trouble arises from how far removed computer simulation is from the world it is intended to represent. Even in the case of mathematical modeling we are closer to the physical world than we are here. Every bit of "physics" in a simulation is the product of some computational artifact we have harnessed for this purpose. But if this is so, then rather than trying to force our simulations to manifest only those "physical" laws that interest us, why not begin trying to harvest the "physical" laws they produce naturally?

We are trying to develop a method of finding out what the invariants, and objects, and laws are for a given class of simulation. In the case of "catalog" interface to OTB alluded to above, for example, by circumscribing the inputs to and outputs from a human performance model, we can start to understand how a fixed set of "resources" affects our ability to represent behavior. Supporting an abstraction of medium-sized, incompressible objects moving over a terrain database with 1-meter resolution might be adequate for the simulation of tanks in the open field, but might lead to completely unreal behavior in the simulation of dismounted soldiers moving through interior spaces. The physics of our simulated world are not invariant with respect to resolution—they are fine at some scales and incoherent at others. Contrast this situation with that of climate modeling where the geometry is discrete, the temperature differences are rational numbers, and time is also discrete. All of these are unreal assumptions, but even so we have good reason to expect in those cases that our simulated world is good enough---the physics *of* the model, and the connection between the model and the world are sufficiently independent of these unphysical features that we can ignore the differences. However in the case of weather simulation in particular, much of our confidence seems to come from empirical studies of the necessary fineness needed in the simulation so that we *can* have such independence. But we don't have that kind of confidence with our current simulation platforms for agents---the agents get stuck, they walk on air, etc. The physics of the simulation is opaque, and worse the relation between small changes at the code level and changes in the physics of the simulation is similarly opaque and beyond our control. At this point, empirical, i.e., *experimental* study would come in very handy.

On the other hand, we do know that much of the experimental study in cognitive modeling does not help us uncover invariant structure. Quite the contrary, such studies leave questions about mechanism underdetermined. We have in this class of human decision-making models no tools for and no insight into how to extend our successful simulations. Even if we appeal to the invariance of the underlying "cognitive architecture" as a reason to believe that our results could generalize, the truth of the matter is that the skill of the modeler has far more to do with successful modeling of new domains than whatever robustness we'd like to ascribe to the architecture itself. More often than not, all we can do is capture known empirical features of a situation in a multitude of ways. Until we develop the tools needed to confidently extend our predictions that derive from these various models we will gain insight neither into predicting human behavior nor into understanding the real process by which such decisions are made.

The problem here, as elsewhere in computer simulations, is that we simply cannot theorize from the ground up how our simulations will behave under various small modifications of their underlying code, and neither can we see from the descriptive accuracy of given simulations how much confirmation is conveyed downward from the simulation to the proposed theory. This, in some sense, is necessarily the case. We turn to computer simulation when the route from the theory to prediction is too messy and complicated to traverse with theoretical tools alone.

Our conclusions are necessarily modest. We do not have strong results to point to that justify our experimental approach to the analysis of computer simulation, nor indeed do we have much more than the Baconian injunction "find the causes". Instead we have given some sketchy and preliminary, but, we think, substantial indications that such an approach is likely to bear fruit. By analogy with the analysis of analogue models in the physical sciences, themselves a kind of simulation, we see potential in experiment for gaining understanding of and control over those projectible predicates that arise out of the artifacts of large-scale computer simulations.

We end with a slogan: A good law of nature is where you find it . . . and so is nature itself.

REFERENCES

Anderson, J. R. and Lebiere, C. (1998). *The Atomic Components of Thought.* Mahwah, NJ: Erlbaum.

Anderson, M. and J. Ensher, M. Matthews, C. Wieman, and E. Cornell, "Observation of Bose–Einstein condensation in a dilute atomic vapor",*Science* **269**, 198 (1995).

Barceló, C., S. Liberati, and Visser, M. (2000). "Analog gravity from Bose-Einstein condensates." arXiv:gr-qc/00110262 v1.

Barceló, C., S. Liberati, and Visser, M. (2003). "Probing semiclassical analogue gravity in Bos3-Einstein condensates with widely tunable interactions." arXiv:cond-matt/0307491 v2.

Best, B. J. (2006). "Using the EPAM Theory to Guide Cognitive Model Rule Induction." In Proceedings of the 2006 Behavior Representation in Modeling and Simulation Conference. Baltimore, MD. SISO.

Darken, C. J. and B. Jones. (2007) Computer Graphics-Based Target Detection for Synthetic Soldiers. In Proceedings of the Sixteenth Conference on Behavioral Representations In Modeling and Simulation. Norfolk, VA. SISO.

N.D. Birrell, P.C.W. Davies (1982). *Quantum Fields in Curved Space*. Cambridge: Cambridge University Press.

L. J. Garay, J. R. Anglin, J. I. Cirac, and P. Zoller (2000). "Sonic analog of gravitational black holes in Bose-Einstein condensates." *Physical Review Letters* **85**, 4643. Version cited here: arXiv:gr-qc/0002015.

Gluck, K. A., J. J. Staszewski, H. Richman, H. A. Simon, and P. Delahanty (2001). "The right tool for the job: Information processing analysis in categorization." Proceedings of the 23rdAnnual Meeting of the Cognitive Science Society. London: Erlbaum.

Gluck, K. A., and R. W. Pew (2001) "Overview of the Agent-based Modeling and Behavior Representation (AMBR) Model Comparison Project." Proceedings for the 10th Conference on Computer Generated Forces and Behavioral Representation. Norflok, VA. SISO. 3-6.

Goodman, N. (1983) *Fact, Fiction, and Forecast* 4th Edition. Cambridge, MA: Harvard University Press.

Mattingly, J. (2006) "Why Eppley and Hannah's thought experiment fails." *Physical Review D **73**, 064025.*

Roberts, S. and H. Pashler (2000). "How Persuasive Is a Good Fit? A Comment on Theory testing." *Psychological Review* **107**(2): pp358-367.

Smith, J. D., and J. P. Minda (2000). "Thirty categorization results in search of a model." *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **26**, 3-27.

Unruh, W. G. (1995). "Sonic analogue of black holes and the effects of high frequencies on black hole evaporation." *Physical Review D **51***, 2827-2838.

Warwick, W., and M. Fleetwood (2006). " A Bad Hempel Day: The Decoupling of Prediction and Explanation in Computational Cognitive Modeling." Proceeding for the 2006 Fall Simulation Interoperability Workshop. Orlando, FL. SISO.

Warwick, W. and L. Napravnik (2005). "SAFBots: A Uniform Interface for Embedding Human Behavior Representations in Computer Generated Forces." Proceedings for the Fourteenth Conference on Behavior Representation in Modeling and Simulation, Universal City, CA, SISO.