Serious Games:

Improving Public Policy through Game-based Learning and Simulation

PUBLICATION 2003-1

Subjectivity and Simulation

Chris Crawford

Foresight and Governance Project

Woodrow Wilson International Center for Scholars

http://wwics.si.edu/foresight

ABOUT THE **AUTHOR:**

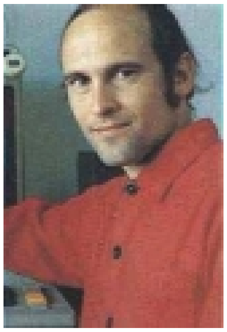

Chris Crawford started out writing games for the Atari Home Computer System in the early 1980's. After writing such classics as Eastern Front (1941) and Legionnaire, he founded the Games Research Group. While there, he wrote The Art of Computer Game Design and created Excalibur, a game about the Arthurian legends. He continued to design and create games throughout the 1980s and early 1990s (Balance of Power, Patton versus Rommel, Trust & Betrayal, Guns & Butter, Balance of the Planet, and Patton Strikes Back). Along the way he also wrote two more books about game design (most recently Understanding Interactivity), founded The Journal of Computer Game Design, and started the Computer Game Developer's Conference. His current project, now nine years in the making, is a technology for interactive storytelling and a development environment that permits non-technical artists to control the technology.

This paper is one of a series of articles supported by the Foresight and Governance Project at the Woodrow Wilson International Center for Scholars in Washington, DC, as part of the Serious Games Initiative. Serious Games explores the application of computer-based games to public policy and administration challenges. An electronic version of this paper can be found at www.seriousgames.org. We appreciate comments on this article that were provided by Ben Sawyer at Digitalmill.

Subjectivity and simulation

Simulation has, historically, been confined to simple problems, most often of a physical nature. Early simulations did little more than provide numerical results for well-understood physical systems such as first- and second-order linear differential equations. It didn't take long before simulation was applied to more mathematically complex systems, especially those that defied analytic solution. Perhaps the most ambitious applications of simulation in the early years were the attempts by the Club of Rome to model the future supply and consumption of natural resources. The alarming results of these simulations generated considerable concern and discussion, but by the 70s these simulations had lost credibility. Simulations of complex, ill-defined systems came to be regarded as unavoidably subjective and therefore unreliable.

Current attitudes towards simulation remain colored by early experiences. My impression is that users of simulations expect objective accuracy, and are either naively accepting or cynically dismissive of the results of a simulation. Most people seem to judge simulation by the Boolean truth value of its results; it is either correct or incorrect. This attitude hampers progress in the application of simulation. Such simulations, be they for training, or other relevant (albeit not predictive) purposes have great use if properly presented. If one considers perception to be reality then one can not deny the fact, that a "realistic seeming" simulation offers some sense of reality that can be useful. This is significant, because most game developers strive to create products (fantastical or real-world based) that encourage suspension of disbelief in their players.

Current attitudes towards simulation remain colored by early experiences.

In the process of putting more substance into their products, game developers have refined techniques that are applicable to simulation software for training, policy exploration, education, and more. Indeed, the incorporation of entertaining elements to a simulation enhances its grip on the user and accordingly increases its educational impact.

Game-based simulations (such as those that I've spent a good deal of my lifetime building) have pioneered many aspects of modern day simulation development – especially of simulations that require some level of subjectivity, multiple viewpoints, player based input, and player-based customization.

Game-based simulation developers have a lot to share about how to tackle the specifics of designing simulations based on imperfect and often subjective information. In this paper I will explore this issue, and present specific arguments and ideas in support of creating "realistic simulations based on imperfect data".

Other representations of reality

To appreciate the core of the problem, we must step back from simulation and think in wider terms of any representation of reality. We represent reality in a wide variety of media: paintings, drawings, sculpture, text, cinema, and so forth. Such representations are not strictly accurate. The Mona Lisa’s hands are impossibly smooth, and she is presented against a landscape that cannot exist anywhere on this planet. Yet no one objects to these patent inaccuracies; they are accepted as part of the overall representation. In war movies, grenades explode with the power of artillery shells. Hyperbole is the stuff of verbal description; we are not surprised when something is described as “incredibly” big or fast or expensive – yet the term means literally that the description should not be believed. Indeed, we not only accept such subjectivity, we factor it into our interpretation of the representation. When we see the unrealistically smooth skin in a Renaissance portrait, we do not reject the portrait as a falsehood; we recognize that the artist is attempting to communicate the beauty of his subject. When the Hollywood movie shows overly large explosions, we do not complain that the movie is too dramatic; we expect the movie representation to be visually spectacular.

There is purpose and utility in the exaggeration used in other media. The artist using such techniques can focus our attention on the true content of the expression. Leonardo had no intention of presenting a simple photographically accurate depiction of his subject; he strove to show something of the woman’s personality, her mystery and aloofness. Michelangelo’s statue of David has unnaturally large hands and feet, but this is no error – Michelangelo deliberately highlighted his hands and feet to accentuate the representation of the manly, virile qualities of his subject. The absurdly energetic explosions in war movies serve to accentuate the sense of danger of the hero’s predicament. These qualities (aloofness, manliness, and danger) cannot be directly depicted in the media chosen by their artists, but they can be represented by focusing the audience’s attention on some particular aspect of the representation.

Yet we do not extend these broad-minded considerations to simulation; simulation is expected above all to be accurate. The aura of digital precision that saturates our image of the computer spills over into our appreciation of simulation, blinding us to the possibilities of the medium.

Algorithms as a medium of expression

In this paper I argue that simulation’s greatest utility can be tapped only when we recognize it as a medium of expression rather than a form of calculation. Our problem here is a variation on C.P.Snow’s “Two Cultures” thesis. Fifty years ago, Dr. Snow observed that Western intellectual culture was bifurcating into two mutually antagonistic camps: science and engineering on one hand, and arts and humanities on the other. The Two Cultures divide has cleaved deep into our mentalities, leading us to falsely isolate mathematics from the arts and humanities. We unconsciously treat mathematics as confined to the world of science and engineering, where it is

…simulation’s greatest utility can be tapped only when we recognize it as a medium of expression rather than a form of calculation.

used to give exact answers. Since simulation relies on mathematical algorithms, our unconscious biases prevent us from seeing mathematics as a medium of expression.

An example of my point might help the reader. Consider the following formula:

Anger = Temper * Provocation

This formula says that the Anger a person will feel is equal to the magnitude of the person's Temper multiplied by the magnitude of the Provocation applied to the person. It is a simple, straightforward, and completely understandable statement. The quantification of personality traits such as Temper might seem unfamiliar, but surely you can concede this once you agree that the personality trait Temper comes in varying amounts. Some people have a bigger temper than others; this implies that quantification can be useful.

The mere fact that some personality traits can be quantified does not mean that my formula is correct, or that any of the values I might use for Temper and Provocation are accurate. But accuracy is not the desideratum in cases where we seek to express subjective ideas. If I attempted to use the above formula to predict the behavior of real individuals, then I would surely fail. But if instead I use this formula to control the behavior of fictional characters in a fictional representation of a story, then this formula is justifiable as my own subjective interpretation of human behavior.

Of course, another artist might represent human behavior in other terms and with other formulae, and none of us would denigrate such a representation merely because it is different from my own. We recognize the complexity of human affairs and grant wide latitude in such representations. However, if said artist decided upon this formula:

Anger = Magnanimity * Affection

then we would have a basis for objecting, because no reasonable assessment of human behavior supports such a claim. This is nonsense. We can therefore object to mathematical expressions if they are patently false, but not merely because they are mathematical.

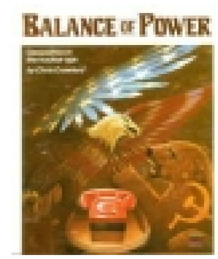

I have used an egregiously subjective example; I shall now tackle a problem a bit closer to home. What about political simulation? In designing my game of geopolitical conflict, **Balance of Power**, I used a great many formulae that reflected subjective assessments on my part. For example, I used the following formula for the popularity rating of governments:

Change in popularity = % change in consumer spending per capita – 3%

(The actual formula was a bit more complex but I shall shed that additional messiness.) The idea behind this formula is simple: a government's popularity is based on economic performance. So long as the populace enjoys steadily improving economic conditions, they will be happy. Note that this formula sets 3% as the base expectation level of the people. They expect an overall improvement of 3% per year; if their lot does not better by at least that much, then they grow disaffected with the government.

The Unexpungeability of Subjectivity

Is this correct? Of course not; it is an approximation. Thousands of subtle factors enter into the overall popularity of a government. As I write this essay, the Bush administration is enjoying phenomenal popularity despite a weak economy. My formula fails badly in this case. A more complex formula would be required to obtain results that accurately reflect the state of the American polity in 2002. Nevertheless, I maintain that the formula I used is a good first approximation of reality. My own interpretation of history suggests to me that economic performance is the single most important factor in the popularity of governments. Other factors can outweigh it, but if we must have a single-term formula, economic performance is probably the best term to start with.

The important point here is that absolute correctness is impossible; approximation is unavoidable and all approximation requires some judgment. Thus, simulation is inevitably a subjective exercise requiring the exercise of judgment on the part of the designer. We can always increase the correctness of our simulations – my formula could be expanded with more terms to obtain more accuracy. But the absence of those terms does not invalidate my simulation; it merely makes it less accurate in the matter of government popularity. If government popularity is central to the intent of the simulation, then perhaps more accuracy in this formula would be desirable. But if government popularity is of secondary importance, then the formula functions adequately.

This brings us to an important guideline for designing simulations: we must tune the content to the message. A simulation that concentrates on economic factors must have detailed economic algorithms, whereas a simulation concentrating on military factors can use less complicated economic calculations. Every simulation has a message, an intent. This concept has long been clear to game designers, who have learned through thousands of trials and errors that content delivered is not the same as content received.

You may be uncomfortable with this assertion. Should we not strive for objective truth, minimizing the role of subjectivity? While absolute objectivity is certainly desirable, the sad truth is that it is forever unattainable. Truth always outstrips the capacity of the human mind; therefore any human representation of reality must be a considered subset of reality: that is, a subjective digestion of reality. Thus, subjectivity is inevitable.

While absolute objectivity is certainly desirable, the sad truth is that it is forever unattainable.

I first realized this when studying that most objective of sciences, physics. In my lower division years I learned about the full range of physics, taking courses in all the important fields. In my upper division years, I took upper division versions of my lower division courses, in which many of the truths I had learned were revealed to be simplistic approximations of physical truth. Then in graduate school I repeated the process, studying the same material in even greater detail. My lower division teachers deliberately lied to me: they presented me with a distorted, subjectively chosen approximation of the full truth. So did my upper division teachers. I now realize that even the courses in grad school were still just approximations of truth – which I shall never learn to infinite accuracy.

The point of this long ramble is this: subjectivity is inevitable. We cannot cleanse our simulations of subjectivity. But while we cannot eliminate subjectivity, we can manage it through three techniques: 1) declare the limitations of the simulation; 2) identify the most questionable subjective elements of the simulation; and 3) make subjective elements subject to modification by the user.

…game developers have amassed practical experience about how to bridge the gaps in information, data, and perfect accuracy, that inevitably exist in simulations like complex socio-economic, geo-political, or human behavior based topics.

Game-developers have spent a myriad of hours developing simulations using these three techniques. Even more importantly, they've refined such techniques in order to create simulations that appear to be more and more realistic to their users. Thus, game developers have amassed practical experience about how to bridge the gaps in information that inevitably exist in simulations like complex socio-economic, geo-political, or human behavior based topics.

Declaring the limitations

Simulations fall short of perfection in two ways: range limitations and intrinsic inaccuracies. The first problem arises when the user of the simulation takes it into behavioral territory outside the range contemplated by the designers. Suppose that the user of an economic simulation abuses the economy so egregiously that the simulation generates negative growth exceeding –100%, generating a negative gross domestic product. Oops. A simulation of population problems might be mismanaged to yield a total world population of –30 million. This kind of absurdity arises only from incompetent design. Every simulation designer must take care to place limitations on such excursions. This is most often managed by placing asymptoting factors into the formulas used in the simulation.

The interactive storytelling technology I have developed provides a good example of how to handle this kind of problem. The software system uses hundreds of variables that can be recalculated at runtime in very complex ways, so there is no simple or direct safety mechanism. Therefore, the software requires that all such variables fall in the range 0 – 100, and that these variables be modified through a special function that I call AdjustXXX.

The AdjustXXX function specifies the variable that is to be modified, and the amount that is to be added to it. The software, however, does not simply add the incremental value (or decremental, if the passed value be negative) to the variable. Instead, it treats the passed value as a percentage of the distance between the current value of the variable and the extreme being approached. Thus, if the function is asked to increment a variable with a value of 0 by 10, then it will add 10% of the distance between 0 and 100 –

which just happens to be 10. In this case, the result will be 10. But if the variable has a value of 90, and the increment is again 10, then the function will increase it by only 10% of the distance between 90 and 100, yielding a final result of 91 for the variable. This insures that the variable always remains within the required range.

The second failure of limitations arises when the simulation is operating within its normal design range, but generates unnecessary inaccuracies. Back in the days when we did integer arithmetic with 8-bit values, this kind of problem was everywhere, but nowadays we don't have to worry about inaccuracies arising from large numbers (except in certain extreme cases that most programmers already understand). These problems most often arise from poorly designed repetitive algorithms. While a single straightforward calculation can normally be relied upon to produce acceptable results, it is possible with complex systems of equations to generate undesirable (i.e., possibly perceived as unrealistic) behavior.

One example of this is oscillatory behavior in a complex system of equations. A minor fluctuation in one value can be amplified in its effects as it propagates through the system, yielding substantial excursions in other variables. Designers often respond to this problem with negative feedback. In systems using negative feedback, some part of the output of one segment of the system is inverted and fed back into the input of a preceding segment. But negative feedback does not solve the problem, it merely converts excursions to infinity into oscillations. My advice in this situation is to equip each formula with shock absorbers that insure that the variable in question moves only, say, half the distance to the computed value. Such algorithmic shock absorbers make for a better ride than simple negative feedback springs. Game designers have developed a wide range of similar techniques (and player testing processes) to preserve the stability of the behavior of their simulations.

Identifying subjective elements

This is an easy and suboptimal solution, amounting to nothing more than a cover-your-ass disclaimer. The identification of subjective elements takes expository rather than interactive form, and thus lacks the communicative power of the simulation. A user deeply enmeshed in the complexities of a good simulation will hardly notice, and take little heed of, warnings buried in the documentation about controversial algorithms, yet the designers will all too readily feel that they have discharged their obligation to truth by burying such warnings in the documentation. I recommend against this approach.

Making subjective elements modifiable by the user

This, I believe, is the most effective solution to the problem of subjectivity in simulation. By sharing the subjectivity with users, we invite them to internalize it and consider its significance, and ultimately to better understand the complexities of the problem. Games have long sported a variety of customization schemes in the form of alternate scenarios, checklists of rules to apply or disable, or player attributes that

By sharing the subjectivity with users, we invite them to internalize it and consider its significance, and ultimately to better understand the complexities of the problem.

can be modified by the player. The popularity of such schemes demonstrates the willingness of players to consider the implications of such variability.

Rather than present a theoretical discussion of how to accomplish this, I shall describe the solution I developed for my simulation, **Balance of the Planet**. This simulation ideally exemplifies solutions to the problems that many simulation designers face. It addresses global environmental issues, with a strong emphasis on economic factors. It encompasses more than 120 public variables in a complex mesh of inter-relationships. Clearly, a simulation of this scale on this topic presents many challenges. My strategy for tackling these problems was as follows:

First, I broke the simulation up into four levels. In the first level, the user merely navigates through the hyperlinks of cause and effect in the simulation, becoming familiar with the scope of the simulation and the many complicated inter-relationships in environmental issues. In the second level, the user plays the simulation, attempting to maximize a point score based on net global human happiness, a variable derived from a great many factors in the simulation. The user can influence the system by levying global taxes on various commercial activities directly affecting the environment, using the proceeds to subsidize various environmental research efforts. In the third level, the subjectivity of the simulation is addressed monolithically, and the user is given the opportunity to play the game with a predefined set of biases, such as a pro-nuclear bias, and environmentalist bias, a pro-growth bias, and so forth. In each of these biases, the algorithms are adjusted to reflect the beliefs intrinsic to that bias.

These three levels are intended to be preparation for the fourth and highest level, in which the user directly controls all the algorithms used in the simulation. This is achieved by the expedient of casting all formulae in the form:

Variable = Coefficient #1 * Input Variable #1 + Coefficient #2 * Input Variable #2…

Confining myself to linear equations imposes serious behavioral constraints on the formulae; a great many minor inaccuracies are generated as a result. Nevertheless, I believe that linear equations are necessary to make the formulae accessible to users. While the intricacies of systems of linear equations may evade the grasp of most users, the simple notion of a weighting factor is easily grasped.

Here is a sample window showing how this is implemented:

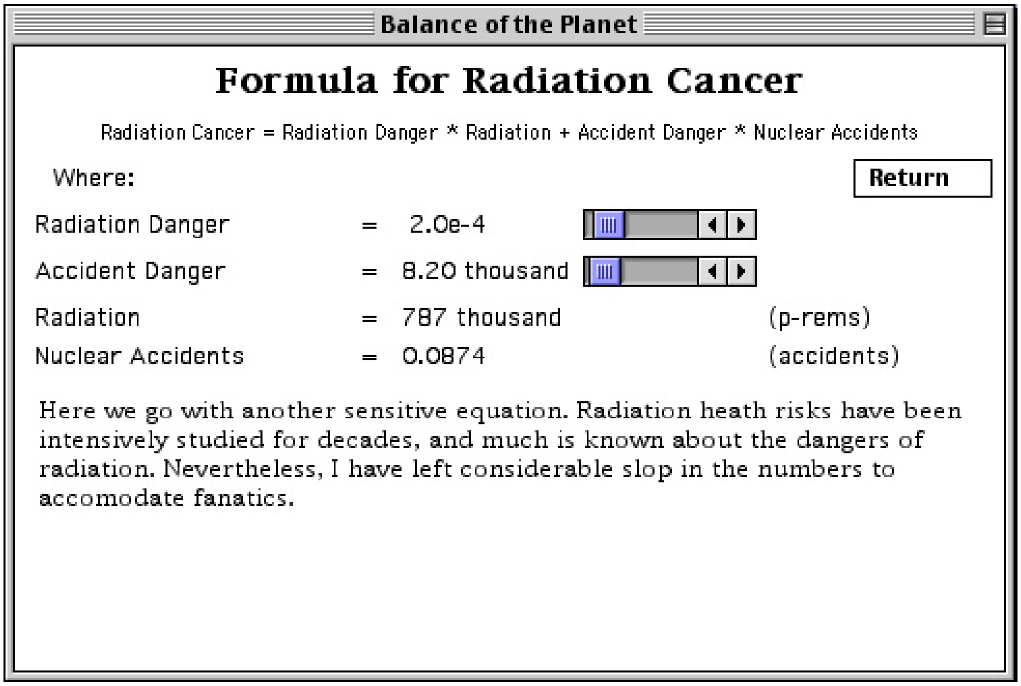

The two weighting coefficients can be adjusted by the user with a scrollbar. This is important because scrollbars can be set with predefined upper and lower limits to insure that the user does not take the simulation outside of its intended behavioral range.

The value of making subjectivity manipulable by the user is made clear with this window:

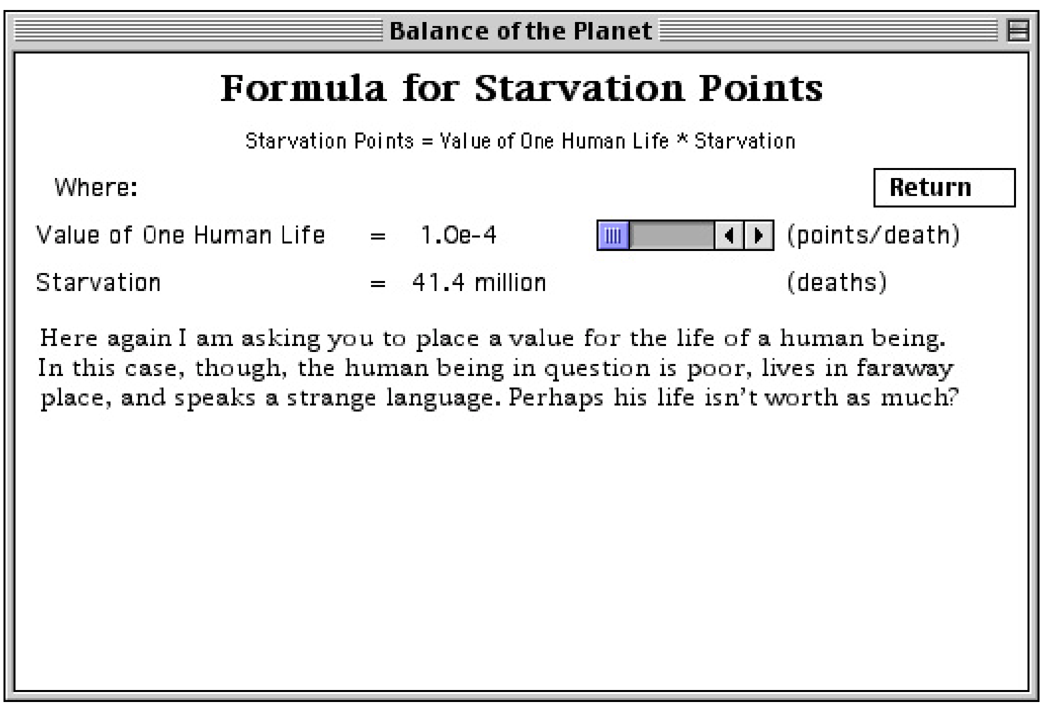

In this window we see the greatest value of this technique. All policy simulations must ultimately deal with matters of personal value. Just how valuable is a human life? How important is freedom, or security? No designer can impose such value judgments on users with a clear conscience.

All policy simulations must ultimately deal with matters of personal value…. No designer can impose such value judgments on users with a clear conscience.

Conclusions

The inclusion of subjectivity in a simulation is a delicate business. Often, simulation designers cop out by designing values and subjective factors right out of the simulation. I believe that the methods demonstrated here are far the better, for instead of avoiding values or sweeping them under the rug, it brings the question of personal values to the forefront. In games our goal is to create simulations and worlds that force players to confront their own values and make explicit their assumptions; this is the most powerful way to bring home the real lessons of simulation. For ultimately, simulation is not a mechanical exercise nor is it a means of bottling truth inside a computer – it is a way to bounce our ideas and values against reality and see how they bounce back.

Woodrow Wilson International Center for Scholars
Foresight and Governance Project

For more information contact:
Foresight and Governance Project
Woodrow Wilson International Center for Scholars
1300 Pennsylvania Avenue, N.W.
Washington, DC 20004-3027
Phone: 202-691-4255
Fax: 202-691-4001
Email: rejeskidw@wwic.si.edu

http://wwics.si.edu/foresight