



in smaller, more understandable chunks than does Soar's native language. Herbal (Cohen, Ritter, and Haynes 2005; Morgan, et al. 2005) is a more recent work that provides a more compact and tractable development environment complete with an Eclipse IDE. Like Taql, Herbal compiles to code that can be executed in Soar. Unlike Taql, Herbal has a relatively shallow learning curve and a simpler grammar than native Soar.

Taql, Herbal, and Maslow are similar in that they are all working towards reducing the cost and complexity of developing models of human behavior. Herbal and Taql both operate in the context of Soar and are languages that are primarily used for developing models of one or more aspects of human cognition. Contrary to Taql and Herbal, Maslow is intended to capture and communicate the essence of human motivation, not human cognition. Maslow is expected to find use as an instrument of communication between very abstract models produced by social scientists and the software developers that implement those models.

In this regard, Maslow bears more similarity to the Unified Modeling Language (UML) (Rumbaugh, Jacobson, and Booch, 1999) than it does to either Taql or Herbal. UML is a very general language that can model a variety of systems, although its typical use is in software engineering and business process modeling. The generality of UML is often regarded as a significant feature, however, in the present context a modeling language that is more specific to the needs of social scientists is likely the superior alternative.

## 2.2 Rational choice theory

Rational choice theory (see Allingham, 2002) is a framework for analyzing social and economic behavior. In rational choice theory, the world is ever evolving from the current state to one of many possible future states. Over the set of possible future world states, an agent has some preferential ordering. A *rational* agent selects actions in a manner which biases the evolution of the world towards one of the agent's preferred states. Rational choice theory has sometimes been criticized for oversimplifying assumptions such as perfect information and ideal games. The subfield of bounded rationality (e.g. Simon, 1957) addresses many of those concerns that focus on the challenges of games where agents have limits in both information and cognitive capability.

## 2.3 Influence diagrams and situated decisions

Influence diagrams (Howard and Matheson, 2005) are graphical structures that represent and visualize decisions, non-deterministic elements, and outcomes of decisions. Node types are shown in Figure 2.1.

Decision nodes (Figure 2.1-a) represent a decision (choice) made by a single agent. Chance nodes (Figure 2.1-b) represent non-deterministic elements such as the outcome of a gamble and are not directly influenced by the decision-making agent. General variables (Figure 2.1-c) are deterministic functions of their inputs and used to aggregate and better visualize the decision situation. Objective nodes (Figure 2.1-d), a special case of the general variable, represent the decision making agent's net gain/loss ("payoff") in utility as a consequence of the decision and the non-deterministic elements.

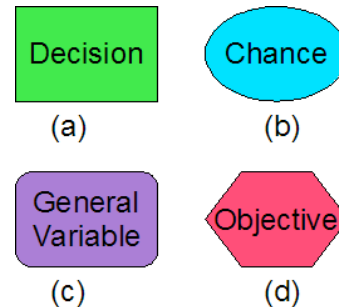


Figure 2.1 Influence diagram nodes

An arc between nodes is directional and is typed by the nodes at the head and tail of the arc. An arc from chance node to an objective node or general variable is a functional arc and indicates that the variable or objective is influenced by the out-come of the chance node. Arcs between decision nodes are sequential arcs and represent a sequence of decisions that occur in the topological order of the sequential arcs. Arcs from a chance node to a decision node are informational arcs and represent the potential state of information that an agent has available when making the decision.

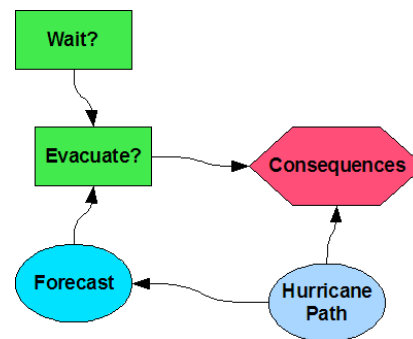


Figure 2.2 Hurricane evacuation example

A compact, canonical example of an influence diagram is the hurricane evacuation decision as depicted in Figure 2.2. The first decision to be made is whether or not the decision maker should wait for an updated forecast. The next decision, whether to evacuate or not, will be influenced by the forecast, if available. The

consequences of the decision will be the product of the decisions made and the path which the hurricane takes.

### 2.4 General theories of human motivation

There are certain elements of the human experience which seem to be common. For instance, at the most basic level, all humans need air, water, and food. However, the common aspects of human experience seem to extend far beyond individual subsistence. Many psychological theories have been advanced which aim to capture common human values, ambitions, and actions. Maslow's Hierarchy (Figure 2.3) (Maslow, 1943) is a classic example of such theories (and the inspiration for the name of the language presented here). Alderfer's Existence, Relatedness, and Growth (ERG) (Alderfer, 1972) builds on Maslow's earlier work and replaces the original hierarchy with a parallel relationship between the three dimensions he identifies.




Figure 2.3 Maslow's Hierarchy

Whereas Maslow and Alderfer have advanced psychological models, sociology has also attempted to advance theories of human motivation. For instance, the Fundamental Human Needs identified by Max-Neef, et al (1989) propose that human motivation is described across nine dimensions: subsistence, protection, affection, understanding, participation, leisure, creation, identity, and freedom. Similar in some respects is the work of Nussbaum and Sen (1993) where human welfare (and motivation) is described in terms of capabilities and the ability to move from capability towards actuality. Recent work by The World Bank (Alkire, 2002) considers the possibility of unifying the sociologically inspired theories into a usable metric of human welfare which has the potential of creating universal human utility function.

### 3 Maslow Graphical Language

Maslow is composed of four elements (Figure 3.1) which are called welfare, aspect, stimulus, and action. Each model must have one and only one welfare (Figure 3.1-a) node. This node represents the overall utility state of the agent. Welfare nodes are a special

case of the more general aspect nodes. An aspect node (Figure 3.1-b) represents some component of the overall welfare and can be arbitrarily decomposed. Stimulus nodes (Figure 3.1-c) embody conditions and procedures that influence an aspect of an agent's welfare. Action nodes (Figure 3.1-d) represent alternative courses of action that will positively affect the associated aspect. In building a model, each instantiated element is given a short name and a sufficient description to convey the function of the instantiated node.




Figure 3.1 Maslow elements

In general, stimulus nodes decrease utility and executed actions increase utility. Note that a planning arc represents a belief on the part of the agent that executing the associated action will in some way improve the condition of the associated aspects. Maslow makes no assumptions about the actual outcome of the action and implementations of the action are not constrained to producing positive results.

The grammar of directed influential connections is straightforward. Decomposition arcs denote aggregation or subsumption and can connect an aspect node to one or more aspect nodes or to the root welfare node. Affecting arcs connect a stimulus node to one or more aspect nodes. Planning arcs are placed in order to denote an association between an action node and one or more welfare nodes.

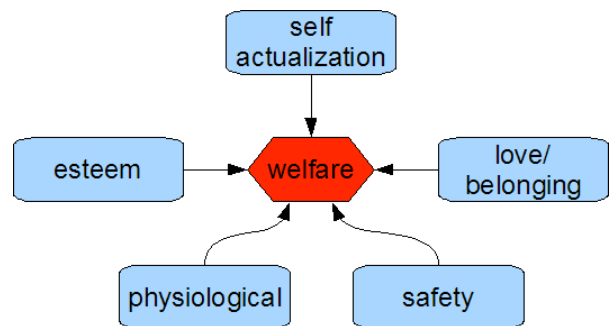
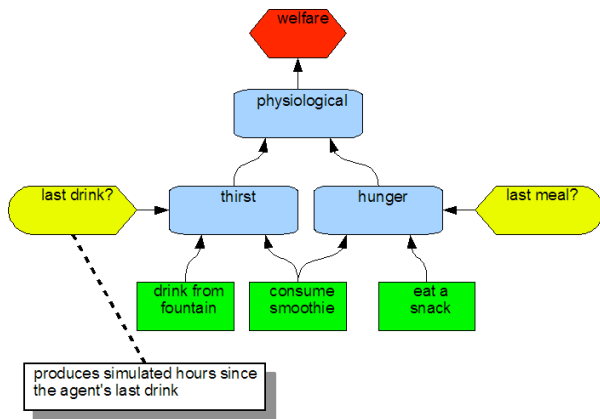


Figure 3.2 Aspects of Maslow's Hierarchy

Figure 3.2 shows the top level welfare and aspect nodes used to represent the components of Maslow's hierarchy. There is an explicit ordering in the hierarchy which implies a utility function over the satisfaction of the components of the hierarchy. *Maslow* does not explicitly represent such a utility function but utility functions are implicit to the specific fusion algorithms used in the aspect nodes and the heuristics implemented to select potential actions. An agent's overall utility function is an emergent phenomenon resulting from the interactions between the states of stimulus processes and aspect fusion.



**Figure 3.3 Usage example**

For the purpose of demonstration, this paper will focus in on the physiological component of Maslow's hierarchy. Figure 3.3 illustrates the use of each element in defining a simple agent. There are two stimulus nodes which represent a rule or process that influences the “thirst” and “hunger” aspects. The “last drink?” stimulus node is also shown with a pop-out text box that contains the longer description associated to the node. The “thirst” and “hunger” nodes are both aspects of the “physiological” aspect. The “physiological” aspect directly influences overall welfare. There are three actions shown. The planning arc that associates “drink from fountain” to the “thirst” aspect depicts the agent's expectations that executing the “drink from fountain” action will positively impact the “thirst” aspect. The “eat a snack” action has a similar relationship with the “hunger” aspect. A smoothie is a kind of drink typically composed of yogurt, crushed ice, fruit, etc. The “consume smoothie” action is associated to both the “thirst” and “hunger” aspects, accordingly.

## 4 Implementing a Maslow Model

*Maslow* is intended to be relatively neutral of the underlying simulation framework in which the agent models are implemented. However, there are some assumptions in the language with the result that some

frameworks will be better implementation targets than others. Following are some considerations for implementing a Maslow model.

### 4.1 Fusion in aspect nodes

Currently, *Maslow* does not communicate a specific algorithm for fusing the welfare of subordinate aspects into higher level aspects. This is left up to the framework to implement a general policy for fusion. Alternatively, supplementary documentation can specify fusion algorithms for individual agents or even individual aspect nodes.

### 4.2 Action selection

*Maslow* does not communicate an action selection heuristic. As with aspect fusion, this is left to either the framework or to specific agents. According to rational choice theory, an agent should execute the script that it believes will have the greatest positive impact on welfare, but how this impact is estimated is not specified.

Of note is the lack of prescriptive state for selecting actions. Agents in many ABS applications are designed with a set of discrete states. As the simulation progresses, the agent moves into and out of these states. Each state constrains the set of actions which the agent can execute. *Maslow* does not define such a finite set, allowing for an agent to be, for instance, “happy” and “sad” simultaneously. *Maslow* can approximate the action selection strategy of the state-based agents by implementing code that reduces the state of all aspects to a single, descriptive value.

### 4.3 Action execution

Maslow's actions are assumed to be coarsely grained. That is, rather than a single small action, *Maslow* assumes that actions contain scripts or plans.

These scripts are a sequence of atomic statements. Atomic statements have an associated time cost and cannot be preempted. Between actions, it is acceptable for *Maslow* to re-evaluate its aspects and possibly preempt the currently executing script in favor of another script.

It is also acceptable, although not necessary, for *Maslow* agents to execute more than one script simultaneously. In this case, the script should contain statements for the acquisition and release of finite resources. Deadlock conditions may occur as a result, however delay will likely result in one or more of the deadlocked scripts to be preempted, effectively breaking the deadlock.

#### 4.4 Simulation time

Maslow does not communicate a model of time. Managing the advancement of simulation time is left to the implementation framework. Maslow can operate in a discrete-time framework where actions are scheduled and the simulation clock is advanced to the next action. In the case where stimuli or actions are continuous with respect to time, a hybrid discrete-continuous model can be used or by advancing the simulation clock in small increments

### 5 Features

*Maslow* has several significant features of which several are briefly discussed below.

#### 5.1 Vehicle for communication

*Maslow* is used to build intermediate models between actual human behavior and the executable simulation code. A model expressed in this language explicitly depicts the relationship between stimuli, actions, and welfare and is easier to comprehend than a collection of if-then rules or program source code. Social science modelers are free to concentrate on the structure and rationale of the agent and can then use this intermediate model as a vehicle for communicating with programmers in much the same way that the Unified Modeling Language (UML) is used to convey program design to programmers.

#### 5.2 Encourages reusable models

Another feature of *Maslow* is its ability to promote reuse. Elements from an intermediate model can be copied from one model to another, importing the structure and semantics from one model to the other. Furthermore, researchers can also develop templates for intermediate models that are, for instance, specialized for various cultures that can span from national and ethnic cultures down to the micro-cultures of communities-of-practice.

Canonical intermediate models can be defined and easily shared. Canonical models can facilitate research by porting behavior across simulations. For instance, a particular market simulation could be run using different cultural models to explore how the differences in macro behavior between, for instance, Germans and Canadians, affect market outcomes.

#### 5.3 Potential in a web 2.0 world

Web 2.0 technologies have dramatically altered the potential for collaboration in research and development. Maslow has the potential to be an excellent means for sharing with other researchers

using wiki technology. Using Uniform Resource Identifiers (URIs) or Digital Object Identifiers (DOIs), canonical models can be uniquely identified as network accessible resources. A standard wiki can be used to hold the underlying textual description of a Maslow intermediate model (serialized using, e.g. GraphML) and provide additional services such as version tracking, indexing, and cross-linking with other constructed models or components. With more-advanced wiki technology such as OntoWiki (Auer, Dietzold, and Riechert, 2006), specific implementations of model components can be associated to the intermediate models. As a system, these technologies will enable a researcher to develop social science simulations by first "shopping" for already developed models and then copying or modifying relevant components. Collaboration along such lines will establish intellectual capital which will continue to pay dividends in future research as ever more complex behaviors will be easier to construct by recombination of already existing models.

### 6 Conclusion

The purpose of this paper is to introduce *Maslow*, a simple, graphical language for creating intermediate agent models of human behavior. Maslow is not meant to be a complete model of human motivation, action, and cognition. Such technology is well beyond the scope of this paper. Rather, *Maslow* is intended to serve as a concept prototype and to spur further conversation on technologies that reduce the cost and barrier to entry for agent-based simulation for the social sciences.

Even as a concept prototype, Maslow has the potential to convey several advantages to modelers. These advantages are primarily in creating clear designs which are easily visualized and understood in much the same way that UML is used in software engineering. Such designs also facilitate communication between those social scientists that will use the simulation for research and the programmers that must implement the simulations in executable code. Furthermore, Maslow promotes reuse between simulations and within the broader social science community and is amenable to publication, distribution, and collaboration using Web 2.0 technologies.

### 7 References

- Alderfer, C. (1972) *Existence, Relatedness and Growth*. The Free Press: New York.
- Alkire, S. (2002) Dimensions of Human Development, *World Development* 30(2) pp. 181-205, 2002.

Allingham, M. (2002) *Choice Theory: A Very Short Introduction*, Oxford University Press.

Auer, S., Dietzold, S. and Riechert, T. (2006) OntoWiki - A Tool for Social, Semantic Collaboration. In *Proceedings of the 5th International Semantic Web Conference (ISWC 2006)*.

Cohen, M.A., Ritter, F.E., and Haynes, S.R. (2005) Herbal: A High-Level Language and Development Environment for Developing Cognitive Models in Soar. In *Proceedings of the 14<sup>th</sup> Conference on Behavior Representation in Modeling and Simulation (BRIMS)*.

Howard, R.A. and Matheson, J.E. (2005) Influence Diagrams, *Decision Analysis* 2(3) pp. 127-143.

Iba, T. (2004) A Framework and Tools for Modeling and Simulating Societies as Evolutionary Complex Systems, In *Proceedings of the 2nd. International Conference of the European Social Simulation Association*.

Maslow, A.H. (1943) A Theory of Human Motivation, *Psychological Review* 50 pp.370-396.

Max-Neef, M. et al. (1989) Human Scale Development: An Option for the Future, *Development Dialogue*, 5-81.

Morgan, G.P., et al. (2005) Increasing Efficiency of the Development of User Models. In *Proceedings of the 2005 Systems and Information Engineering Design Symposium*.

Newell, A. (1990) *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

Nussbaum, M. and Sen, A. (1993) *The Quality of Life*, Oxford: Clarendon Press.

Rixon, A., Moglia, M., and Burn, S. (2005) Bottom-up Approaches to Building Agent-based Models: Discussing the Need for a Platform, In *Proceedings of the Joint Conference on Multi-agent Modeling for Environmental Management*.

Rumbaugh, J., Jacobson, I., and Booch, G. (1999) *The Unified Modeling Language Reference Manual*. Addison-Wesley.

Simon, H. (1957) A Behavioral Model of Rational Choice, in *Models of Man, Social and Rational:*

*Mathematical Essays on Rational Human Behavior in a Social Setting*. New York: Wiley.

da Silva, P.S. and De Melo, A.C.V. (2008) Reusing Models in Multi-Agent Simulation with Software Components, In *Proceedings of the 7<sup>th</sup> International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Yost, G.R. (1993) Acquiring Knowledge in Soar. *IEEE Expert* 8, pp. 26-34.

## Author Biography

**NATHAN T. DENNY** (*Scientist, 21<sup>st</sup> Century Systems, Inc.*) received his Master of Science in Computer Science from Southern Illinois University in 1998 and his Bachelors of Science in Computer Science and Mathematics from Southern Illinois University in 1997. Mr. Denny is also pursuing a PhD in Computer Engineering from the University of Arizona. He has published in diverse fields such as design for testability (DFT) of very large scale integrated circuits (VLSI), information re-use in case-based reasoning, Internet spam control, automation in agricultural irrigation, peer-to-peer networking and knowledge management in global software development. His current research interests include distributed computing, artificial intelligence, human-machine interfaces, cognitive science, knowledge management, and agile software development in the 24-Hour Knowledge Factory. Mr. Denny assists with research and development at the Omaha office of 21<sup>st</sup> Century Systems, Inc.

