Subjective Logic for Composing Utility Functions from Maslow Models

Nathan T. Denny
21st Century Systems, Inc.
6825 Pine St.
Omaha, NE 68105
402-502-8439
nathan.denny@21csi.com

Keywords:
social science, utility functions, graphical models

ABSTRACT: *Researchers in the social sciences often collaborate with software developers to create agent-based simulations that are increasingly used in the study of sociology, political science, economics, etc. Maslow is a nascent, graphical (network or connectionist) modeling language that aims to make the modeling of motivation more intuitive to the social scientist and facilitate the translation of simulation specifications into executable code. This paper builds upon the Maslow language, illustrating how subjective logic can be used as a means to represent influence between elements in a Maslow model. So constructed, an acyclic Maslow model can be expressed as a subjective logic expression which in turn can be compiled into executable code. The result is a model that can represent motivations with arbitrary detail that is also computationally efficient. The detail and scalability of this approach may be of particular interest in multi-agent simulations of large groups, where a good degree of modeling fidelity can be achieved with relatively little impact in computing performance.*

1 Introduction

Agent-based simulations (ABS) have broad applicability and can be applied to modeling teams of robots, the spread of infectious diseases, and even entire ecosystems. ABS has found increasing use in the study of sociology and economics where researchers can simulate organizational behavior, market exchanges, and other social interactions to study the emergence of macro characteristics from micro entities. In the present context, these micro entities are behavioral models that are proxies for real human behavior.

As has been noted (see. Iba, 2004; da Silva and de Melo, 2008; Rixon, Moglia, and Burn, 2005), ABS simulations are not always easy to develop. Available simulation platforms typically require some degree of technical ability in order to implement simulations using what is often (e.g. Java) a general purpose programming language. Social scientists must either acquire the necessary technical skills themselves or collaborate with software developers that already possess the technical know-how. Both options can be prohibitive and costly.

For those social scientists that do their own software development, re-use of previous models is enticing (Newell, 1990). Indeed, the software engineering community seems to be able to deliver, to some degree, on its long promise of object and component reuse. However, this has only come about after many years of incremental accumulation of intellectual capital, accreting into software libraries and frameworks. By comparison, ABS simulations are too new and too few to have built up enough intellectual property and most ABS studies build their models from scratch with highly-domain specific agents.

The division of effort between social scientist and software developer is an efficient use of resources, but is not without difficulties. In particular, describing a behavioral model at a granularity that is easily understood by both the social scientist and the software developer may not be trivial. Furthermore, the description should outlive the lifetime of the study, thereby promoting model re-use in later studies.

Maslow is a nascent, graphical (network or connectionist) modeling language that aims to make the modeling of motivation more intuitive to the social scientist and facilitate the translation of simulation specifications into executable code. This paper builds upon the Maslow language, illustrating how subjective logic can be used as a means to represent influence between elements in a Maslow model. So constructed, an acyclic Maslow model can be expressed as a subjective logic expression which in turn can be compiled into executable code. The result is a model that can represent motivations with arbitrary detail that is also computationally efficient. The detail and scalability of this approach may be of particular interest in multi-agent simulations of large groups, where a good degree of modeling fidelity can be achieved with relatively little impact in computing performance.

2 Maslow

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 2.1) (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.

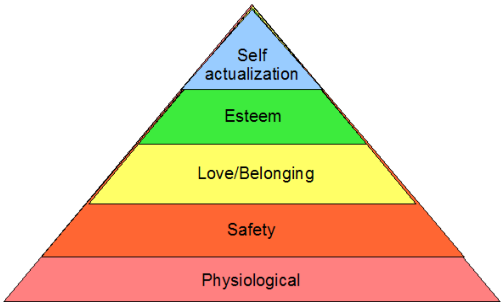

Figure 2.1 Maslow's Hierarchy

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

The human brain has a nearly universal structure, with the location of specialized functions found in more-or-less the same relative locations across individuals. This lends support to the concept of a universal cognitive architecture that can model human cognition. A consequence of both universal structure and universal cognitive architecture is the existence of a universal architecture of human utility functions. Although Abraham Maslow did not describe his work as such, his eponymous hierarchy reflects such a universal architecture of human utility.

Maslow (Denny, 2009) is a simple, graphical language which is intended to model human motivation in much the same way that the Unified Modeling Language (Rumbaugh, Jacobson, and Booch, 1999) describes software architecture. The *Maslow* graphical language is composed of four elements (Figure 2.2) which are called welfare, aspect, stimulus, and action. Each model must have one and only one welfare (Figure 2.2-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 2.2-b) represents some component of the overall welfare and can be arbitrarily decomposed. Stimulus nodes (Figure 2.2-c) embody conditions and procedures that influence an aspect of an agent's welfare. Action nodes (Figure 2.2-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.

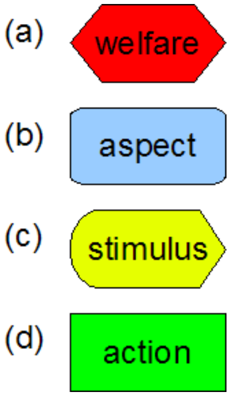

Figure 2.2 *Maslow* elements

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.

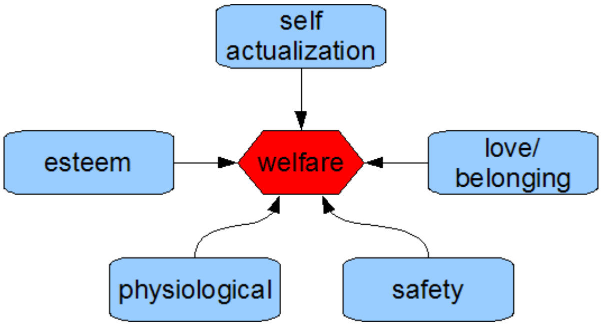

Figure 2.3 Aspects of Maslow's Hierarchy

Figure 2.3 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.

3 Overview of Subjective Logic

Subjective Logic (Josang, 1997, 2009) is a type of probabilistic logic that is often used in evidential reasoning (e.g. Lindahl and Petrov, 2007 and Lindahl and Zhu, 2007) where belief, disbelief, and uncertainty must be explicitly and simultaneously accounted. Before discussing the method by which Subjective Logic can be used to compose utility functions, a brief introduction to Subjective Logic and a summary of the relevant algebraic operations is in order.

In contrast to systems described by Boolean Logic, for those systems described by Subjective Logic the basic object is an opinion rather than a fact. An *opinion* $\omega^A(x)$ about some proposition "x" held by source "A" is a 4-tuple of the belief (b_x^A), disbelief (d_x^A), uncertainty (u_x^A), and relative atomicity (a_x^A). (Atomicity is the base-rate of the proposition.) Note that $b_x + d_x + u_x = 1$, so while it is not necessary to specify all three of the values, it is convenient when performing certain calculations.

The Subjective Logic algebra provides an array of operations that manipulate opinions. These operators have many applications in evidential reasoning and data fusion. For the present purpose, only the consensus and discount operators are of interest.

The consensus operator (written as ⊕) is used for belief fusion, providing the capability to fuse possibly conflicting opinions while still forming coherent, summary judgments. The underlying calculations on the belief tuple elements are given in Figure 3.1.

$$
\begin{aligned}
K &= u_x^A + u_x^B - u_x^A u_x^B \\
b_x^{A,B} &= \frac{b_x^A u_x^B + b_x^B u_x^A}{K} \\
d_x^{A,B} &= \frac{d_x^A u_x^B + d_x^B u_x^A}{K} \\
u_x^{A,B} &= \frac{u_x^A u_x^B}{K} \\
a_x^{A,B} &= \frac{a_x^A u_x^B + a_x^B u_x^A - (a_x^A + a_x^B) u_x^A u_x^B}{K - u_x^A u_x^B}
\end{aligned}
$$

Figure 3.1: Subjective Logic Consensus Operation

Subjective logic also provides a well developed "discount" operation (written as ⊗) that can be used for modifying the contribution of evidence based upon a subjective measure of confidence in the source of the evidence. The discount operator thus provides a rather general means of describing degrees influence and can be used to represent semantic similarity, relevance, trust, etc. The calculations for implementing a discount operator over belief tuples is shown in Figure 3.2.

$$
\begin{aligned}
b_x^{A,B} &= b_B^A b_x^B \\
d_x^{A,B} &= b_B^A d_x^B \\
u_x^{A,B} &= d_B^A + u_B^A + b_B^A u_x^B \\
a_x^{A,B} &= a_x^B
\end{aligned}
$$

Figure 3.2: Subjective Logic Discount Operation

4 Composing Utility Functions

As a modeling tool, Maslow is predicated upon Rational Choice Theory (see Allingham, 2002). That is, agents have a utility function and reason and act to maximize the utility function. Although Rational Choice Theory is sometimes derided as too simple a model of human behavior, most of the criticisms of simplicity are well addressed by Bounded Rationality (e.g. Simon, 1957).

The welfare, aspect, and stimulus nodes of an executing *Maslow* model are essentially the component variables of a utility function. The welfare node is the ultimate dependent variable and contains the present, summarized utility state of the agent. Stimulus nodes contain the state of external stimuli. Aspect nodes are intermediate variables that are calculated as a function of other aspects and affecting stimuli. Both decomposition arcs (between aspect nodes) and affecting arcs (from stimulus to aspect) carry a measure of influence that is defined over the range [0,1.0].

An example model is shown in Figure 4.1 where a subgraph of a Maslow model focuses on the influence of claustrophobia on welfare. The color fill in the boxes next to each arc represent the degree of influence that propagates along the arc. The agent in Figure X is highly sensitive to claustrophobia. The same structure is re-used in Figure 4.2, with another agent that is relatively insensitive to claustrophobia.

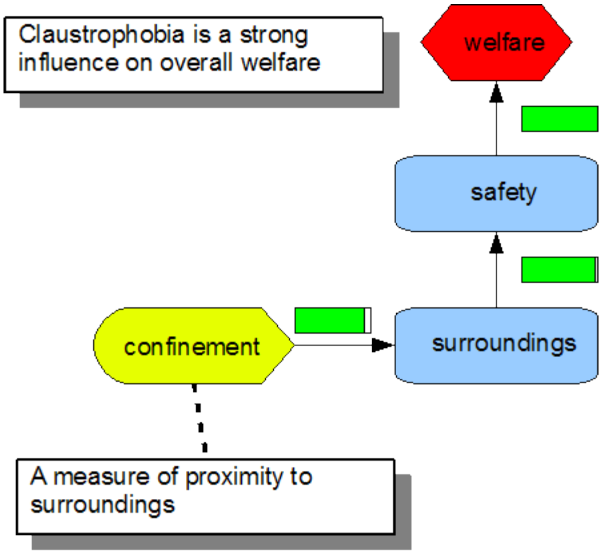

Figure 4.1. Claustrophobic agent

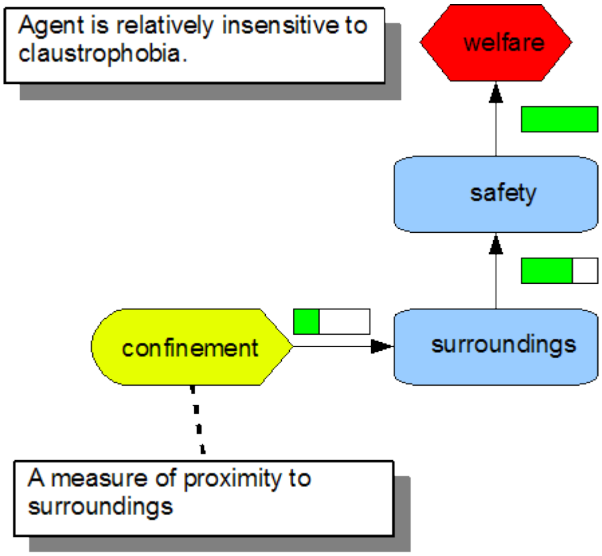

Figure 4.2. Agent is little affected by claustrophobia

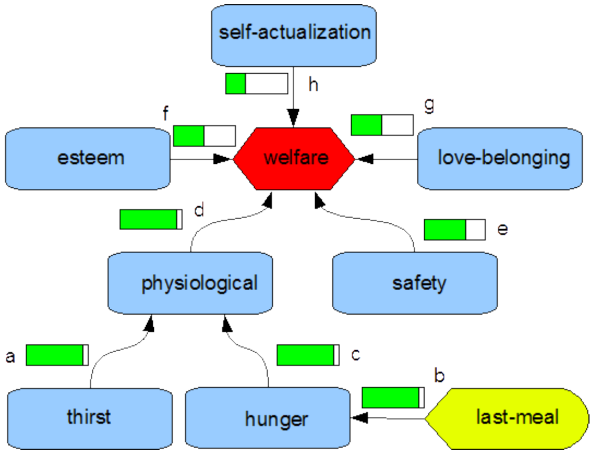

Figure 4.3. Propagating influence

The process of generating a computable utility function from a Maslow model is relatively straightforward: aspect and stimulus nodes are treated as opinions while decomposition and affecting arcs act as discounts on propagated influence. The compiler would then traverse the model in topological order (working from the exterior nodes to the interior nodes) and generate an infix expression of the graph. For example, the physiological contribution of the model shown in Figure 4.3 can be represented algebraically as:

$$(\oplus\ (\otimes\ (\oplus\ (\otimes\ \text{last-meal}\ b)\ \text{hunger})\ c)\ (\otimes\ \text{thirst}\ a))$$

Before executing the model, the infix expression would first be compiled to byte-code or machine code for efficient evaluation. This latter characteristic is of particular importance when running simulations of large groups where demands on computing resources can be severe.

When the model is executed at run-time, aspect and stimulus nodes are stateful and hold the default vacuous opinion where all belief mass is allocated to uncertainty. As stimuli act on the model, the influence from the stimuli propagates through the network of aspect nodes, changing their state and ultimately influencing whatever reasoning engine is employed for the agent. As Subjective Logic is not yet widely supported in reasoning engines, the Subjective Logic expectation function is a simple and convenient function for mapping from a 4-tuple belief vectors to the more common representation of belief as a scalar in the range of [0, 1.0]. (The expectation function loses information and should only be used on the result taken from the welfare node.)

5 Conclusions and Future Work

To date, Maslow has remained ambiguous on how influence was to be propagated from stimulus through aspects to the overall welfare of the agent. Although Subjective Logic was developed for evidential reasoning, there is an intuitive similarity between evidence and influence and the algebra of Subjective Logic lends itself for use in composing functions from relatively distinct components. Given an acyclic *Maslow* model, the model can be assembled into an infix expression in subjective logic and can then be further compiled into byte-code or machine-code that can be efficiently executed at run-time.

Maslow is still in its infancy and undergoing gradual improvement. Maslow remains agnostic to the reasoning mechanism, but this may need to be changed given commitments that the model is now assuming. Furthermore, the method of composing utility functions that has been described here represents only the instantaneous utility. For a higher-fidelity model, the language and framework must be amended to include something akin to the inertia that individuals often have in their emotional (the surface manifestation of welfare) states.

6 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.

Denny, N.T. (2009) Maslow: A Graphical Framework for Communicating Models of Human Behavior. In *Proceedings of the 19th Conference on Behavior Representation in Modeling and Simulation (BRIMS)*.

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*.

Josang, A. (1997) Artificial Reasoning with Subjective Logic. In *Proceedings of the 2nd Australian Workshop on Commonsense Reasoning*.

Josang, A. (2009) Conditional Reasoning with Subjective Logic. *Journal of Multiple-Valued Logic and Soft Computing*, 15(1).

Lindahl, E.R. and Petrov, P. (2007) Situation Awareness and Dynamic Group Formation of Agents and Humans Using Jabber. In *Proceedings of the 12th International Command and Control Research and Technology Symposium (12th ICCRTS)*, 2007.

Lindahl, E.R. and Zhu, Q. Epistemic Belief Frames in Distributed Effects-based Reasoning. In *Proceedings of the International Conference on Integration of Knowledge Intensive Multi-Agent Systems (KIMAS 2007)*, 2007.

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.

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 7th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Author Biography

NATHAN T. DENNY ***(Scientist, 21st 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 21st Century Systems, Inc.