

SMILE[®]: Structural Modeling, Inference, and Learning Engine and GeNIe: A Development Environment for Graphical Decision-Theoretic Models

Marek J. Druzdzel

Decision Systems Laboratory
School of Information Sciences, Intelligent Systems Program,
and Center for Biomedical Informatics, University of Pittsburgh
Pittsburgh, PA 15260

marek@sis.pitt.edu, <http://www.pitt.edu/~druzdzel>

Abstract

SMILE[®] (Structural Modeling, Inference, and Learning Engine) is a fully portable library of C++ classes implementing graphical decision-theoretic methods, such as Bayesian networks and influence diagrams, directly amenable to inclusion in intelligent systems. Its Windows user interface, **GeNIe** is a versatile and user-friendly development environment for graphical decision-theoretic models. Both modules, developed at the Decision Systems Laboratory, University of Pittsburgh, have been made available to the community in July 1998 at <http://www2.sis.pitt.edu/~genie> and have over 1,200 users worldwide (as of April 1999). This document summarizes the basic features of **GeNIe** and **SMILE[®]**.

Decision-theoretic systems

Decision-theoretic systems are increasingly applied in various domains because of their sound foundations, ability to combine existing data with expert knowledge, and intuitive framework of directed graphical models, such as Bayesian networks and influence diagrams. Some of the applications of decision-theoretic systems are: medical diagnosis and therapy planning, machine diagnosis, natural language processing, vision, robotics, planning, fraud detection, processing of military intelligence data in the context of battle damage assessment, and many others (March 1995 issue of *Communications of the ACM* lists several practical applications of Bayesian networks; others can be found in the proceedings of the Annual Conference on Uncertainty in Artificial Intelligence, available on-line in electronic format at <http://www.sis.pitt.edu/~dsl/uai.html>). Given the current interests in the application of decision-theoretic methods and the speed with which they are applied in practice, it can be expected that they will remain to be core modeling tools in intelligent systems.

GeNIe and SMILE[®]

The Decision Systems Laboratory at the University of Pittsburgh (<http://www.sis.pitt.edu/~dsl>) has been working in the domain of decision-theoretic decision support systems for

almost six years. Our theoretical and technical contributions have found their way into a general-purpose decision modeling environment, **SMILE[®]** (Structural Modeling, Inference, and Learning Engine) and its Windows user interface, **GeNIe**. We have made both programs available for non-commercial research, teaching, and personal use since July 1998 at <http://www2.sis.pitt.edu/~genie> and as of February 1999 have more than 1,200 users from over 80 countries. Some of the applications, built using **GeNIe** or **SMILE[®]**, are: battle damage assessment (Rockwell International and U.S. Air Force Rome Laboratory), group decision support models for regional conflict detection (Decision Support Department, U.S. Naval War College) intelligent tutoring systems (Learning and Development Research Center, University of Pittsburgh), medical therapy planning (National University of Singapore), medical diagnosis (Medical Informatics Training Program, University of Pittsburgh; Technical University of Bialystok, Poland). **GeNIe** and **SMILE[®]** have been also used in teaching statistics and decision-theoretic methods at several universities.

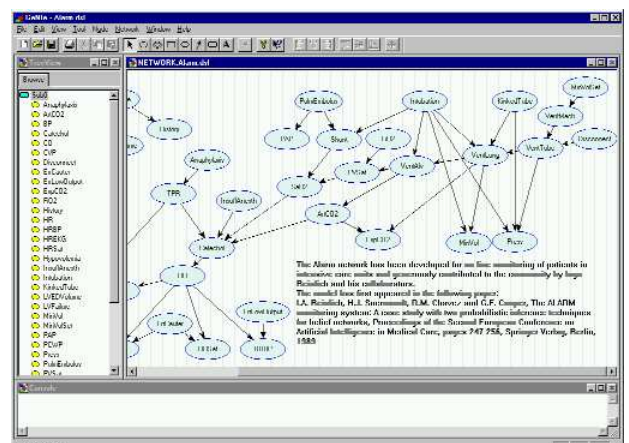


Figure 1: A schematic view of the **GeNIe** window

In the development of **GeNIe**, we stressed accessibility and friendliness of the user interface (see Figure 1 for a snapshot of the program interface). The architecture of the system

(Figure 2) is flexible: **SMILE**[®] is the core reasoning engine that can be embedded in dedicated user interfaces (**GeNIE** is in fact one such interface!). An example of a dedicated user interface to **SMILE**[®], a medical system for diagnosis of liver disorders developed at the Technical University of Bialystok, is shown in Figure 3. The interface shows on the left hand side a list of risk factors and symptoms that a physician can enter and on the right hand side a list of possible disorders according to their likelihood as computed by the program.

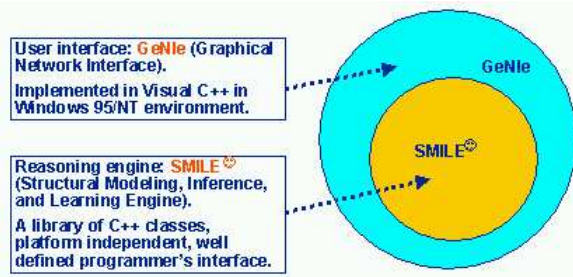


Figure 2: The architecture of **GeNIE** and **SMILE**[®]

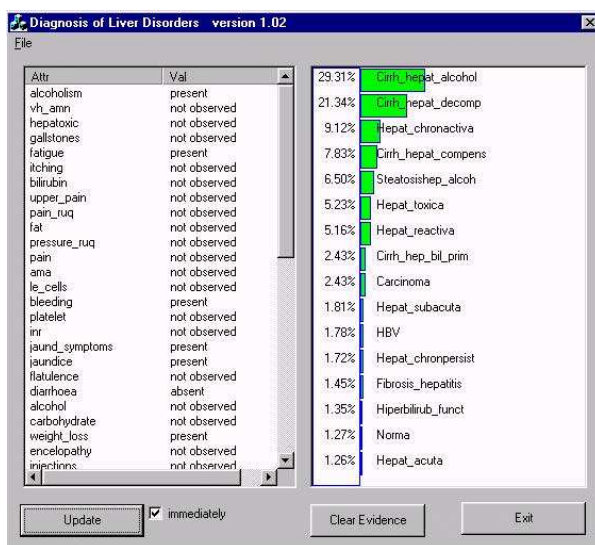


Figure 3: A dedicated interface to a **SMILE**[®] model.

Some features of **GeNIE** and **SMILE**[®]

GeNIE and **SMILE**[®] implement Bayesian networks and influence diagrams with a robust and reliable user interface both in terms of the development environment (**GeNIE**) and the programmer's interface to the library classes (**SMILE**[®]). The modeling language includes hierarchical sub-models, Noisy-OR nodes, deterministic nodes, multiple decision

nodes, multiple utility nodes, linearly additive Multi-Attribute utility nodes (we will have generalized MAU nodes in the future). The development environment includes a pleasant graphical interface with aid in model navigation, such as hierarchical sub-models, a Windows-style tree view, on-screen comments, a comprehensive HTML-based help system, and many other useful features that one would want from a development environment for graphical models. The help system is fully integrated with the World Wide Web and has many useful links to web resources. We use **GeNIE** in teaching and the help system includes many useful documents and tutorials. We believe that it is basically a standalone guide to decision-theoretic modeling. **SMILE**[®] algorithms include relevance reasoning that includes designating nodes as targets, value of information computation, and several Bayesian network algorithms to choose from. **GeNIE** is the Decision Systems Laboratory's research and teaching vehicle, so naturally it will evolve as time goes.

One feature that may contribute to sharing research and development results in the community is that **GeNIE** and **SMILE**[®] are able to read and write several popular Bayesian network and influence diagram file formats, including the most recent version of the Bayesian network interchange format. It can thus be used as a conversion program.

GeNIE runs on Windows 95/98/NT computers. **SMILE**[®] is a fully portable library of classes accessible from C++. We have also developed **SmileX**, an ActiveX Windows component that allows **SMILE**[®] to be accessed from any Windows programming environment, including World Wide Web pages.

Acknowledgments

Support for the development of **GeNIE** and **SMILE**[®] has been provided in part by the Air Force Office of Scientific Research under grant F49620-97-1-0225 and by the National Science Foundation under Faculty Early Career Development (CAREER) Program, grant IRI-9624629. The past and present principal developers **GeNIE** and **SMILE**[®] (listed alphabetically) are: Steve Birnie, Jeroen J.J. Bogers, Jian Cheng, Denver H. Dash, Marek J. Druzzdel, Daniel Garcia Sanchez, Nancy Jackson, Hans van Leijen, Yan Lin, Tsai-Ching Lu, Agnieszka Onisko, Hans Ove Ringstad, Jiwu Tao, Carl P.R. Thijssen, Daniel Tomalesky and Haiqin Wang.