Applying Software Engineering to Agent Development

Mark A. Cohen, Frank E. Ritter, and Steven R. Haynes

■ *Developing intelligent agents and cognitive models is a complex software-engineering activity. This article shows how tools to create intelligent agents can be improved by taking advantage of established software-engineering principles such as high-level languages, maintenance-oriented development environments, and software reuse. We describe how these principles have been realized in the Herbal integrated development environment, a collection of tools that allows agent developers to exploit modern software-engineering principles.*

T his article is about how to make it easier to create intelligent agents by applying established software-engineering principles. We present an example integrated agent development environment that realizes these principles and provides lessons for other agent architectures, both existing and in development.

Creating complex software is not a new problem, and the software-engineering community has developed principles to guide solving complex problems with software. Developing intelligent agents is a complex software-engineering activity but the benefits of applying software-engineering principles such as high-level languages, maintenance-oriented development environments, and software reuse to intelligent agent development have not yet fully migrated to the agent-development community. We demonstrate how these principles have been realized in the Herbal Toolset, a collection of tools that embody these three core principles. In this article, we introduce the Herbal Toolset and the principles that informed its design.

High-Level Languages for AI and Cognitive Modeling

Intelligent agents and cognitive models are useful in many applications, but for different reasons. The processing speed and accuracy of an intelligent agent can help people dial cell phones more quickly and accurately when driving a car, while a cognitive model can help predict common human errors and their causes, which can lead to better cell phone design for novice users. Agents can be used as colleagues and adversaries for education and training, especially in domains where actual human participation could be dangerous (Jones et al. 1999). In these areas, opponents that follow predictable scripts can be made more interesting using cognitive models of human adversaries (Laird 2001). Prediction and psychological insight are two important outcomes of a cognitive model that separate it from other types of intelligent agents. In addition, computer interfaces can be tested efficiently using models of human users (for example, Ivory and Hearst [2001]).

Unfortunately, development of intelligent agents and cognitive models is challenging. In more conventional software-engineering domains, high-level languages are often used by developers because they simplify the creation of complex systems. However, many of the intelligent agent and cognitive-modeling architectures in use today rely on low-level, rule-based programming languages (for example, Soar, ACT-R). While programming in these languages is not as primitive as using assembler, these low-level production systems do exhibit a similar problem: the concepts represented in low-level rules sometimes bear little resemblance to the concepts used by the programmer to solve the problem. Higher-level representations are needed to close the gap between the theory and the behavior-representation languages, and this need is increasingly recognized and evident in the literature (for example, Cooper and Fox [1998], Jones et al. [2006], Salvucci and Lee [2003]).

The absence of higher-level languages that incorporate cognitive theory as an explicit object in the language (instead of using rules) has not gone entirely unnoticed (Ritter et al. 2006). In response, researchers have begun developing higher-level languages that simplify the encoding of behavior by creating representations that map more directly to a theory of how behavior arises in humans.

For example, G2A (St. Amant, Freed, and Ritter 2005) is a high-level representation language that allows for the creation of ACT-R models using the goals, operators, methods, and selection rules (GOMS) model of human activity (John and Kieras 1996). One naturalistic experiment has shown that G2A reduced the amount of effort required to produce ACT-R models by more than 90 percent (St. Amant, Freed, and Ritter 2005).

ACT-Simple (Salvucci and Lee 2003) and Cog-Tool (John et al. 2004) are additional examples of high-level languages designed to simplify cognitive modeling and have been shown to be useful for quickly building models that predict expert performance. They are similar to G2A in that they provide a GOMS-based higher-level language that compiles into low-level ACT-R rules.

ACT-Simple, CogTool, and G2A are examples of how combining the simplicity of an abstract behavior-representation language with the complexity of a lower-level cognitive architecture can simplify the modeling task. Unfortunately, these tools support only one underlying architecture (ACT-R). In addition, the high-level language supported by these tools is based on GOMS, which is limited to modeling expert behavior in simple tasks. Providing support for multiple architectures and a high-level language based on a richer psychological theory would open up new opportunities for these development platforms.

The high-level symbolic representation (HLSR) project is another example of a high-level cognitive-modeling language. The HLSR project aims at creating a formal language that encompasses a wide variety of modeling tasks using a variety of cognitive architectures (Jones et al. 2006). Currently, HLSR creates Soar and ACT-R models. HLSR supports multiple architectures using microtheories, which describe how an HLSR architectural construct will compile into a specific architecture. The ability of HLSR to use microtheories to remove the architectural dependencies from the code that represents the cognitive model is an important accomplishment. This allows modelers to implement a model once, yet execute in different architectures.

However, the architectural neutrality of HLSR results in the lack of explicit support for widely used unified theories of cognition (for example, the PSCM of Soar, and ACT-R). Instead, the microtheories hide this theory. Because modelers often use one of these theories of cognition to understand how to perform a task, they must take an extra step to translate their understanding of the task into a description using HLSR. This gap between the modeler's conceptualization of behavior, and its realization in the high-level language, is exactly what a high-level language is intended to address.

The lessons described above are important for builders of agent and cognitive-modeling development environments, and these lessons informed the high-level agent-construction language adopted by the Herbal Toolset introduced in this article.

Maintenance-Oriented Development Environments

For complex systems, the process of maintenance is the most expensive phase of a system's development life cycle (Boehm 1987, Brooks 1995). A recent study by the National Institute of Standards and Technology (Tassey 2002) showed that U.S. programmers spend more than 70 percent of their time testing and debugging. For programmers of intelligent agents and cognitive modelers this problem may be even more acute, given the complexity of the software they typically develop.

Fortunately, the use of high-level languages can help with maintenance (Brooks 1995). A review by Hordijk and Wieringa (2005) categorized the factors that influence software maintainability. Included in these factors were code-level properties such as code complexity and duplication. High-level languages help here because they reduce code complexity and duplication.

In addition to high-level languages, a survey done by Hordijk and Wieringa (2005) also identified development environments as a factor that influences maintainability. This implies that creating maintenance-oriented environments (environments that explicitly support software maintenance) can help reduce the cost of software development.

Researchers have taken notice. For example, Ko, Aung, and Myers (2005) looked at how Java programmers perform software maintenance. Their study suggested that programmers form a working set of task-relevant code fragments that is typically built by using a find and replace dialogue or by visually searching the source code. Ko and colleagues believe that support for working sets can help simplify maintenance by bundling related source code fragments into sets that can be easily and quickly retrieved, examined, and modified.

Another way to streamline the maintenance process is to make it easier for programmers to access the design rationale underlying a system (La-Toza, Venolia, and DeLine 2006). Easy access to this design rationale is crucial, but developers currently spend much of their time reconstructing design rationale that is implicitly embedded within a program's source code (LaToza, Venolia, and De-Line 2006). It has long been recognized that one of the key barriers to comprehending software is the invisibility of its structure relative to the functions it performs (Brooks 1987). For intelligent agents and cognitive models, this problem is especially acute as these systems are often intended to carry out complex and consequential operations similar to, and sometimes in place of, people. Because of this, intelligent agents are expected to be accountable for the actions they perform, to explain how and why their reasoning led to a particular action.

An early implementation of this idea was the XPLAIN architecture for creating intelligent systems (Swartout 1983). XPLAIN was designed expressly to provide explanations of an expert system using the design rationale underlying its structure and behaviors. The XPLAIN knowledge base includes justifications for system structure, behavior, and general problem-solving strategies, as well as a mapping between terms and definitions used in the design to those in its domain of intended use. Other examples of how design rationale has been used to support explanation in intelligent and other complex systems are available (Haynes 2005).

More recently, Haynes, Cohen, and Ritter (2009) have proposed a unified approach using design patterns to support explanation in agent-development environments. Their work presents a set of guidelines for developing agents that explain themselves based on a study of the questions that users asked while working with an intelligent system. Three general design patterns for creating explainable agents emerged from the analysis of this study:

Ontological explanations are designed to provide answers to questions about the static structure of an agent's design. Mechanistic explanations provide insight into how the components within an agent interact to produce behavior. Finally, the operational explanations describe how a modeler can access and utilize an agent's functionality.

These explanation types help provide the rationale underlying the design of the agent. The advantage of access to these explanations is that developers and users can understand the intent of the program designers. Building support for these types of explanations into a high-level language and agent-development environment can simplify the creation of intelligent software by providing reusable and extendable models of explanation delivery.

Software Component Reuse

Reuse of source code, even within a single program, can reduce development and maintenance costs (Boehm 1987, Brooks 1995, Krueger 1992). One of the earliest discussions of the importance of software reuse was presented by McIlroy (1968). He described a view of the future of software reuse that included the availability of safe, well-tested software components tailored to each user's needs.

A useful framework for understanding reuse was developed by Krueger (1992). Krueger breaks software reuse into four dimensions: (1) abstraction, (2) selection, (3) specialization, and (4) integration. Abstraction allows programmers to consider a programming task separate from the concrete realities of the modeling language and is the reason why

high-level languages provide more effective support for reuse. Although often taken for granted, abstraction in high-level languages is one of the most successful vehicles of software reuse (Brooks 1987, Krueger 1992). Selection allows programmers to locate and select appropriate reusable components. Good maintenance-oriented environments make it easy for developers to search for reusable components based on criteria that are tied directly to the design rationale. Specialization allows programmers to tailor reusable components to their specific needs. Without support for specialization, developers are unable to configure a component for their specific use. Integration, the fourth of Krueger's dimensions, allows developers to combine reusable components into a working program. A maintenance-oriented environment should play a major role in simplifying the integration piece of reusable software.

Design patterns provide an effective way to implement Krueger's dimensions. Design patterns are reusable templates that provide solutions to recurring problems (Gamma et al. 1995). A design pattern consists of four central elements: The pattern name, which makes it possible for developers to identify and communicate about a pattern; a problem, which helps developers recognize when a particular pattern is useful; a pattern solution, which provides an abstract description of the pattern and how it can be used to solve the problem; and the consequences, which discuss the trade-offs related to the pattern's use.

The Herbal Toolset

The software-engineering community has developed strong theories and principles about how to solve complex problems with software solutions. High-level languages have been shown to simplify software development (Beck and Perkins 1983; Daly 1977; Maxwell, Wassenhove, and Dutta 1996), yet many popular agent-development environments do not yet support high-level programming languages (Jones et al. 2006). The advantages of software maintenance and accessible design rationale are also clear (Boehm 1987; Brooks 1987; LaToza, Venolia, and DeLine 2006; Tassey 2002), yet many popular agent-development architectures (for example, Soar and ACT-R) lack development environments that incorporate these features (Cooper and Fox 1998, Pew and Mavor 1998, Ritter et al. 2003). Finally, code reuse has been shown to reduce development and maintenance costs (Boehm 1987, Brooks 1995, Krueger 1992), yet reuse is difficult to achieve in most intelligent agent development environments because they are built on lower-level languages.

The Herbal Toolset is an attempt to improve agent development by providing a high-level language and maintenance-oriented agent-development environment that offers first-class support for design rationale and software reuse. Software-engineering research suggests that this toolset, or any toolset designed with these principles in mind, will lead to more productive agent developers and useful agents. The next section explains how the Herbal Toolset realizes these three principals.

Herbal: A High-Level Behavior-Representation Language

To simplify agent programming and cognitive modeling, a high-level behavior-representation language and associated parser and compiler were designed and implemented as the Herbal Toolset. This high-level language is based on the Problem Space Computational Model (PSCM) and is represented using the Extensible Markup Language (XML).[1] This language is currently compiled into productions that execute within two popular agent architectures: Soar[2] and Jess.[3]

The Problem Space Computational Model as a High-Level Language

The high-level language supported by the Herbal Toolset is based on PSCM (Lehman, Laird, and Rosenbloom 1996; Newell et al. 1991). PSCM defines behavior as movement through a problem space (see figure 1), which is a high-level organizational representation to explain high-level behavior. PSCM provides a proven conceptual model for implementing a high-level language for agent development.

A problem space is defined by a set of states (for example, $S0$, $S1$) and a set of operators (that is, $O0$, $O1$). A task is formulated when a problem space (P) is adopted, a desired goal (D) is set, and the state of the problem space ($S0$) is initialized. The task is attempted as operators are selected and applied to the current state, transforming the problem space into a new state. Finally, the task terminates when the current state matches the goal (Newell et al. 1991).

PSCM was first proposed by Allen Newell, and this theory was implemented in the Soar Cognitive Architecture (1990). The successful use of PSCM by Soar for the creation of cognitively plausible agents provides evidence of the utility of PSCM as a unified theory of cognition (Jones et al. 1999, Tambe et al. 1995). This prior success, along with our own experience using Soar, has provided motivation for our use of PSCM as the foundation for the high-level language used by the Herbal Toolset.

PSCM can also serve as a general organizational structure for AI-oriented intelligent agents. Clancey (1981) argued that rule-based agent devel-

Figure 1. PSCM Defines Behavior as Movement through a Problem Space.

Based on Exhibit 11.7 in Newell et al. (1991).

opment is complex because the problem-solving strategy is often implicitly hidden within the rules. A language explicitly providing PSCM constructs can alleviate this problem by partitioning behavior into a hierarchy of problem spaces, operators, states, and desired goals.

XML and XML Schema

The PSCM-based high-level language supported by the Herbal Toolset takes the form of an XML application, which implements PSCM and is translated into a low-level rule-based representation for execution within an agent environment (see figure 2).

Using XML as an intermediate representation provides many benefits. For example, XML allows for the creation of structured documents that can directly represent the hierarchical structure of PSCM as a higher-level agent-development language. In addition, the portable text format used by XML is easily readable by both people and computers. In fact, there are a large number of robust XML editors that parse XML and provide a graphical environment for quickly and safely editing the XML (for example, XMLSpy, oXygen, XMetaL). XML can also be transformed into other formats using the Extensible Stylesheet Language (XSL).[4] This makes it possible to transform Herbal agent code into other formats such as HTML documentation or scalable vector graphics (SVG).[5] Finally, the popularity of XML reduces the developer's learning curve that might otherwise form a barrier to its adoption.

The Herbal high-level language specification is defined using XML schema.[6] The use of XML schema for our language was advantageous for many reasons. XML schema allowed us to provide clear documentation of the structure and content of the Herbal high-level language. In addition, the Herbal XML schema is used directly by XML parsers to validate the content of an Herbal program. Lastly, most commercial and open source XML editors utilize XML schema to provide features such as syntax highlighting and autocompletion.

An Herbal program is made up of six different types of XML documents, each defining a set of reusable components: types, conditions, actions, operators, problem spaces, and agents. These documents are referred to as libraries in the Herbal language. In most cases, the components in these libraries mirror the elements of PSCM. However, there are components in the Herbal high-level language (such as types, conditions, and actions) that supplement PSCM by providing additional levels of abstraction and support reuse in ways the rules
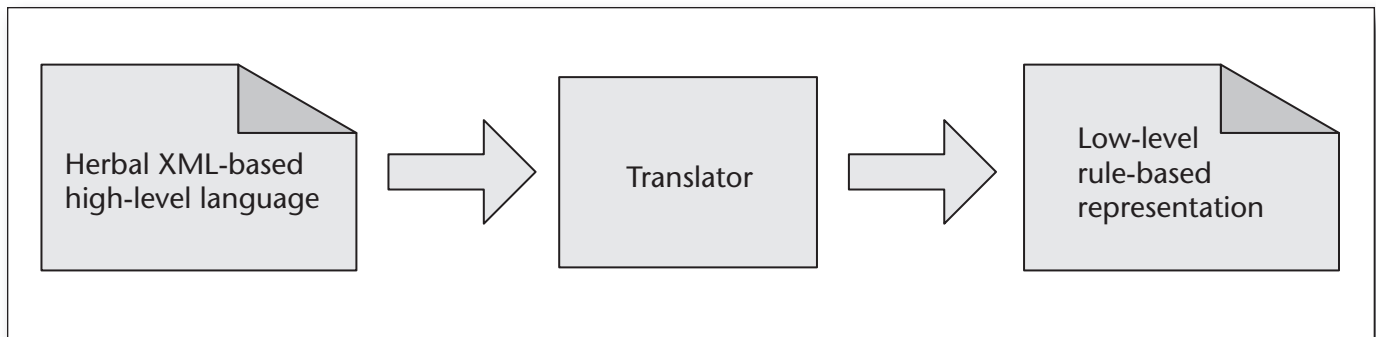
Figure 2. High-Level XML Representation Is Translated into Low-Level Rule-Based Representations.

XSchema Specification for an Operator   **Instance of an Operator**

```
<xs:complexType name="operatorType">        <operator name='driveRight'>
  <xs:sequence>                               <if>
   <xs:element name="if"                       <conditionref condition='okRight'/>
             type="ifType"                    </if>
             minOccurs="1"                    <then>
             maxOccurs="1"/>                   <actionref action='moveRight'/>
   <xs:element name="then"                    </then>
             type="thenType"               </operator>
             minOccurs="1"
             maxOccurs="1"/>
  </xs:sequence>
  <xs:attribute name="name"
type="xs:ID" use="required"/>
</xs:complexType>

<xs:complexType name="ifType">
  <xs:sequence>
   <xs:element name="conditionref"
             type="conditionRefType"
             minOccurs="0"
             maxOccurs="unbounded"/>
  </xs:sequence>
  </xs:complexType>

<xs:complexType name="thenType">
  <xs:sequence>
   <xs:element name="actionref"
             type="actionRefType"
             minOccurs="0"
             maxOccurs="unbounded"/>
  </xs:sequence>
</xs:complexType>
```

Table 1. The Operator XML Schema Specification and an Operator Instance of This Specification.

cannot, ironically, because the rules are too coarse.

The left side of table 1 shows a section of an XML schema file that defines an operator element. An operator is a major component of PSCM and is a first-class object in the Herbal high-level language. According to the specification shown in table 1, an operator element has a unique name and child elements of ifType and thenType. The ifType element contains references to conditions, and the thenType element contains references to actions. These references point to conditions and actions that are defined in a separate XML document (library) and whose syntax is specified in a separate XML schema.

The right side of table 1 lists a typical section of Herbal source code. The XML shown here declares an instance of an operator called driveRight and obeys the schema given in the left side of table 1. The driveRight operator will be proposed when the condition okRight is true, and when the operator is applied an action called moveRight will move the agent to the right. The details of the okRight condition and the moveRight action are given in additional libraries.

The XML schema and associated XML code shown in table 1 can be edited graphically using any of the commercial or open source XML editors. For example, figure 3 shows an XML document, containing instantiations of several operators, being edited in XML Notepad. In figure 3, XML Notepad indicates that an operator is missing the required name attribute and can help the programmer add this. For additional information about the Herbal high-level language specification, see the *Herbal XML Programmer's Guide* (Friedrich, Cohen, and Ritter 2007) available from the Herbal website.[7]

The Herbal Parser and Compiler

Code written in the Herbal high-level language can be transformed into executable productions for either the Soar or Jess agent architectures. The first phase in this transformation (shown in figure 4) consists of parsing the XML code and creating a document object model (DOM) of PSCM. The Herbal parser is written in Java, and the DOM consists of a hierarchical collection of Java objects.

A standard XML parser is used to parse the XML libraries. This parser validates the XML based on the associated XML schema and is extended with custom logic that checks for semantic errors.

The DOM is used for the creation of useful visualizations or the creation of executable productions. The Herbal compiler is responsible for the transformation of the DOM into executable code.

The Herbal application programming interface consists of a set of Java interfaces and abstract classes that support creating compilers for multiple architectures. The compiler API can be implemented and extended to transform the DOM into different types of executable code. There are currently two concrete compilers in the API: one that produces Soar productions and another that produces Jess productions and facts.

The main challenge in creating an agent compiler is deciding how to transform the PSCM DOM into a semantically equivalent set of productions for a specific architecture. The degree of difficulty of this transformation is related to how explicitly the underlying language supports PSCM. For example, the Soar architecture is a direct instantiation of PSCM, while Jess provides no explicit support for PSCM.
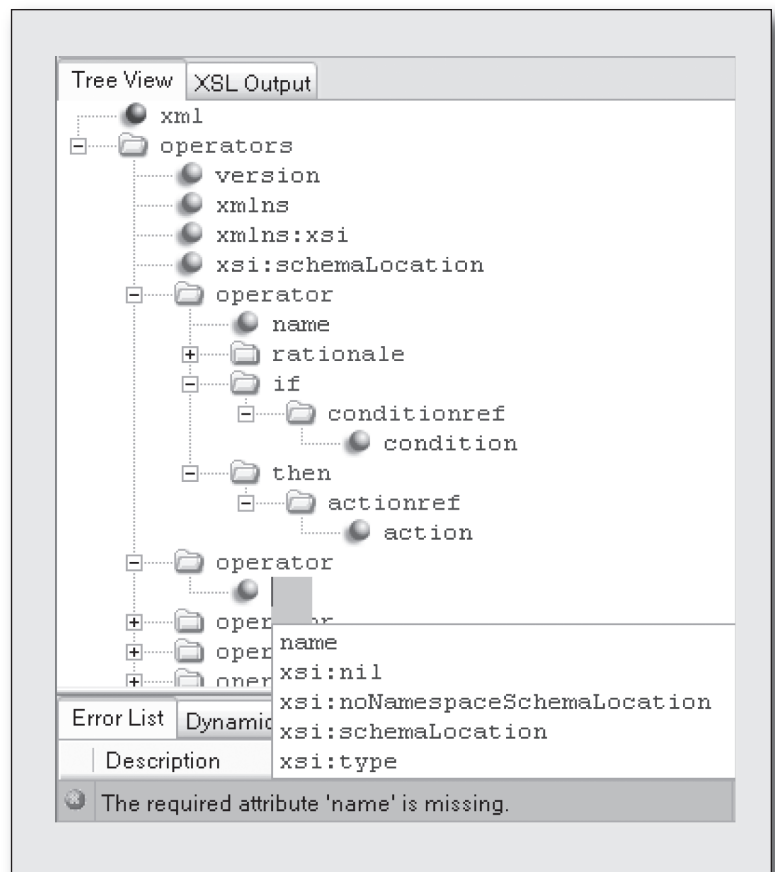

Figure 3. Herbal Programming Using XML Notepad.

A few examples are provided to illustrate how the Soar and Jess compilers transform the PSCM DOM. Table 2 shows Herbal XML code and the resulting Soar and Jess code for a condition called *dirty*, which tests whether a vacuum cleaner agent is on a dirty square (Cohen 2005). This translation is straightforward because both Soar and Jess have clear support for the concept of a condition. However, this example shows one of the main benefits of the high-level language: unlike the resulting Soar and Jess code, the Herbal XML language makes sets of conditions explicit. In other words, in Herbal, sets of conditions are named and unambiguous, making it easy for developers to recognize them and, importantly, reuse them.

A second example, given in table 3, illustrates how the Soar and Jess compilers transform Herbal XML code for an action called *clean*. This translation is less straightforward because Soar and Jess have different support for the interaction with the environment. Again, this example shows the advantage of the high-level language. The Herbal actions are easy to identify and reuse.

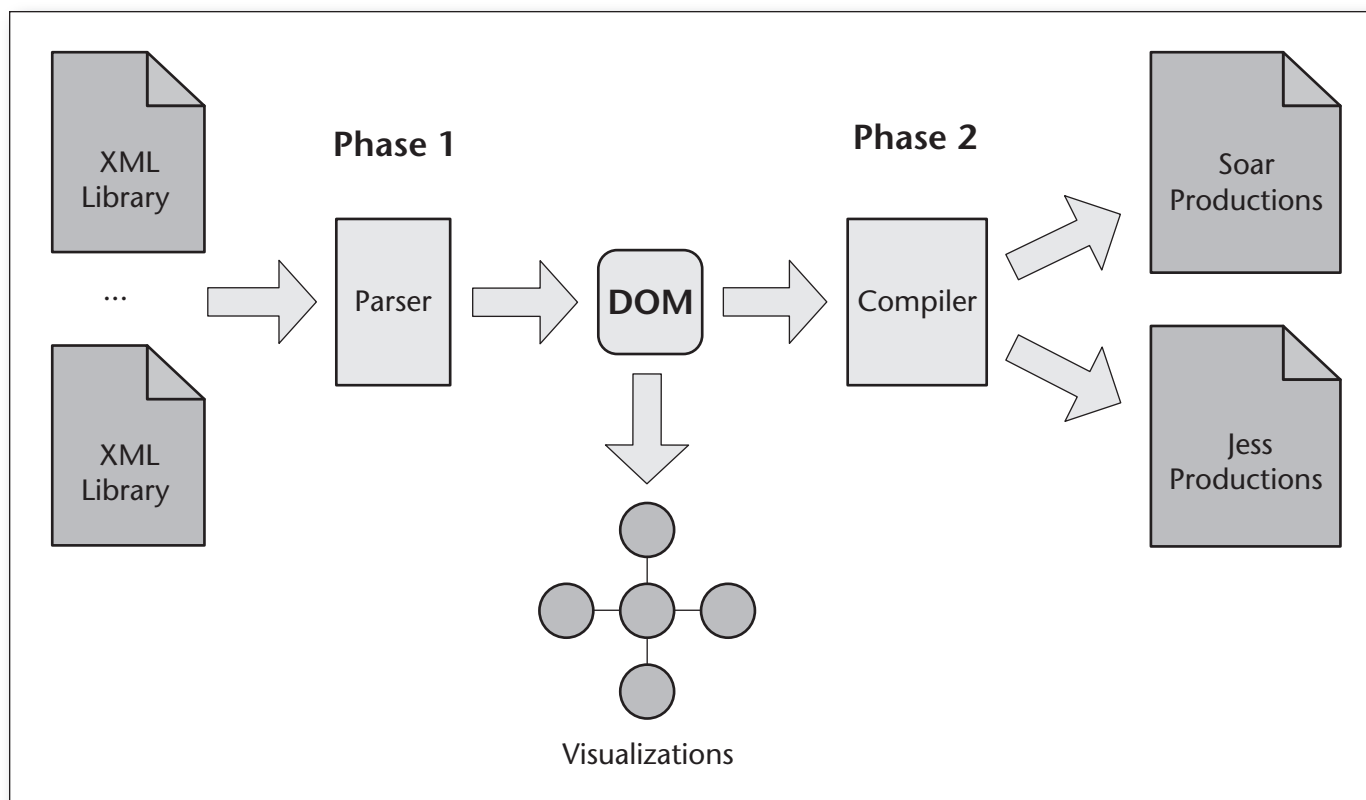Soar defines explicit structures to support an agent's communication with its environment.

Figure 4. Parsing and Compiling Herbal XML Source Code.

Architecture	Source Code

Herbal XML Language	```<condition name='dirty'>```
	``` <match type='vacuum.types.spot'>```
	``` <restrict field='status'>```
	``` <eq>dirty</eq>```
	``` </restrict>```
	``` </match>```
	```</condition>```

Compiled Soar Code	```(<vacuum-types-spot2> ^status <status2> |dirty|)```
Compiled Jess Code	```(topspace::vacuum.types.spot (status ?status1&:(eq* ?status1 "dirty")))```

Table 2. A Translation from an Herbal Condition to Soar and Jess Source Code.

These structures take the form of an input link and an output link. As a result, the Herbal compiler adds the Clean working memory element directly to the output link (labeled <i2> in table 3). Jess, on the other hand, has no special language constructs that deal with agent/environment interaction, so the Clean command is treated like any other fact in working memory.

The third example, shown in table 4, demonstrates how the Herbal compiler transforms an Herbal operator. Recall that the operator is an important component of PSCM. Surprisingly, Soar

Architecture	Source Code
Herbal XML Language	```<action name=clean'>
<add type='vacuum.types.action'>
<set field='move'><value>clean</value></set>
</add>
</action>```
Compiled Soar Code	```(<i1> ^output-link <i2>)
-->
(<i2> ^|vacuum.types.action| <vacuum-types-action20>)
(<vacuum-types-action20> ^move |clean|)```
Compiled Jess Code	```(assert (topspace::vacuum.types.action (move "clean")))```

Table 3. A Translation from an Herbal Action to Soar and Jess Source Code.

does not have a single explicit syntax for declaring operators. Instead, operators arise implicitly with productions. In addition, unlike Soar, the Jess language has no concept of operators at all. As a result, the operator concept must be simulated in Jess using a basic production.

Table 4 shows how the Herbal compiler produces operators in Soar and Jess. This example clearly illustrates the advantages of using a high-level language. In table 4, the Herbal high-level language declaration of an operator is far more obvious, and in the case of Soar, far more concise. The first-class status of conditions and actions in the Herbal high-level language made this possible because the detailed specifications of conditions and actions are referenced rather than duplicated. This is a clear example of how high-level languages can reduce errors and development time by eliminating code duplication and making code easier to read and understand.

For the Herbal to Jess rendering, the Herbal operator is translated directly into a single production. However, in Soar an operator consists of a proposal rule and an application rule (Lehman, Laird, and Rosenbloom 1996). The proposal rule fires when the operator is appropriate for the current situation. The application rule contains knowledge about how the operator changes working memory. The distinction between operator proposal and operator application allows for interruptability, which is an important part of the psychological plausibility of Soar agents and is also necessary to support learning in Soar. The Soar productions shown in table 4 show how the Herbal compiler produces both the proposal and the application rules for the Herbal operator, reducing the time and chance for error inherent in coding these operators by hand.

The examples given in tables 2, 3, and 4 illustrate how the Herbal high-level language represents and implements PSCM. In some cases (for example, the addition of conditions and actions as first-class objects), these modifications have added greater granularity, which allows for easier reuse, while in other cases (for example, simulated operators in Jess), sacrifices were made in the richness of the problem-solving abilities and psychological plausibility of PSCM. These sacrifices are apparent when creating models in architectures that do not provide direct support for PSCM. For example, interruptability, which was described previously as an important part of the psychological plausibility of Soar agents, is not accounted for in agents compiled to Jess. While minimized, these trade-offs are common throughout the design of the Herbal Toolset. In all cases, however, the high-level code is more explicit, making it easier for the developer to recognize and reuse the key components of PSCM.

Herbal: A Tool for Supporting Maintenance

The Herbal Toolset includes an integrated development environment (IDE) that provides a graphical environment for creating and maintaining agents by leveraging the popular Eclipse extensible platform[8] and by providing integral support for design rationale and working sets.

The Herbal IDE

The Herbal IDE is implemented as an Eclipse plug-in. Eclipse is a universal platform providing an open and extensible IDE that provides many advantages. First, Eclipse provides a framework for the creation of powerful development tools. This framework consists of the modern IDE features ex-

Architecture	Source Code
| Herbal XML Language | ```xml
<operator name='clean'>
   <if>
      <conditionref condition=dirty'/>
   </if>
   <then>
      <actionref action='suck'/>
   </then>
</operator>
``` |
| Compiled Soar Code | ```
sp {propose*clean
   (state <local> ^top <top> ^parent <parent> ^name cleanps)
   (<top> ^io <i1>)
   (<i1> ^input-link <i2>)
   (<top> ^|origvac.types.status| <vacuum-types-spot2>)
   (<vacuum-types-spot2> ^status <status2> |dirty| )
-->
   (<local> ^operator <o> + =)
   (<o> ^name clean)
   (<o> ^count <count>) }
sp {apply*clean
   (state <local> ^top <top> ^name cleanps ^operator <o>)
   (<o> ^name clean)
   (<top> ^io <i1>)
   (<i1> ^output-link <i2>)
-->
   (<i2> ^|vacuum.types.action| <vacuum-types-action20>)
   (<vacuum-types-action20> ^move |suck| )}
``` |
| Compiled Jess Code | ```
(defrule clean
   (topspace::vacuum.types.spot
      (status ?status18&:(eq* ?status18 "dirty")))
=>
   (assert (topspace::vacuum.types.action (move "suck") )) )
``` |

Table 4. A Translation from an Herbal Operator to Soar and Jess Source Code.

pected by developers, including project management, multiple project and code views, and real-time compilation. In addition, as the popularity of the Eclipse IDE has grown, the learning curve for using the Herbal IDE is reduced for developers already familiar with Eclipse. Finally, Eclipse is free and executes on a variety of platforms.

Graphical editing in Herbal is accomplished with the Herbal GUI editor (shown in figure 5). Like the Herbal high-level language, the editor is library centered. Using the editor, programmers can create or modify existing library components (that is, types, conditions, actions, operators, problem spaces, and agents) without having to write code in the Herbal high-level language—the Herbal XML code is created as the developer interacts with the GUI editor.

The GUI editor can also be used as a means for teaching the Herbal XML language. Developers can create a PSCM component using the GUI editor and then inspect the generated XML code. By switching between the editor and the generated code, programmers can learn the syntax of the Herbal high-level language.

While the editor simplifies the creation of PSCM components, some developers may prefer to work directly with the Herbal high-level language. At any time during development, programmers can edit the Herbal XML code directly, and these changes are immediately reflected in the GUI editor (see figure 6).

Typical of most Eclipse plug-ins, the Herbal compiler is automatically invoked as the programmer works. With each change, the Herbal IDE compiles the Herbal XML code into both Soar and Jess productions. This feature serves as an important

mechanism to support novice developers learning the underlying Soar or Jess programming languages: Herbal programmers can create PSCM constructs using either the Herbal GUI editor or the Herbal high-level language and then inspect the generated Soar and Jess code to learn how these constructs can be implemented in the underlying architectures. In classroom evaluations of Herbal, this strategy proved to be very useful, especially in computer science classes in which learning how to program the underlying architecture (in this case Jess) was a course objective (Cohen, Ritter, and Haynes 2009).

Figure 7 shows the Herbal IDE displaying multiple views of an Herbal library. The top left view shows the Herbal GUI editor. To the right of the GUI editor is a snapshot of the Herbal high-level XML code. The bottom two views in figure 7 show the generated Jess and Soar code. Finally, along the very bottom of figure 7 is a list of current warnings and errors. In this case, a typo made by the developer has generated a warning. Double-clicking this warning will open an editor to the appropriate location in the model so the warning can be resolved. By implementing Herbal within Eclipse, developers gain the advantage of its modern code view and debugging capabilities.

Design Rationale

The Herbal IDE incorporates the three general explanation design patterns introduced by Haynes, Cohen, and Ritter (2009). These design patterns help with the creation of explainable agents, which can be easier to understand, debug, and modify.

Ontological explanations are designed to provide answers to questions about the static structure of an agent's design. In support of ontological explanations, the Herbal IDE provides the model browser view shown in figure 8. The model browser view makes it easy to browse the static PSCM structure of an Herbal agent and simplify the maintenance of these structures.

Mechanistic explanations provide insight into how the components within an agent interact to produce behavior. Mechanistic explanations are typically generated while an agent is executing. As a result, the Herbal IDE provides these types of explanations in the form of a run-time debugger and trace tool. The Herbal debug view (see figure 9) shows a trace of the agent interacting with its environment. For each event, the agent and its current problem space is shown, including the currently satisfied conditions and operators and any actions that have been executed. Finally, relevant working memory is also given.

Finally, the operational explanations describe how a modeler can access and utilize an agent's functionality. When making changes to an agent's

Figure 5. The Herbal GUI Editor.

behavior, operational explanations help the programmer decide what components are available and how these components can be used. The operational explanations are available in the model browser view as answers to questions such as "How do I use it?" and "How does it work?" (see figure 8). This information must be provided by developers as they build library components or else these explanations must be reconstructed by interpreting source code functionality.

These three explanation patterns, and the design rationale they provide, support software maintenance and reuse by providing developers access to the intent and design decision making carried out by a component's original developers. This information should make it easier to identify, select, and specialize components developed for analogous applications and fully to comprehend why a piece of software has the structure and behavior that it does. This understanding contributes to maintenance by avoiding code changes that conflict with designers' intent and consideration of implementation constraints. In addition, because the explanations are given in the context of PSCM, they can be more psychologically plausible.

Working Sets

As discussed earlier, studies done by Ko, Aung, and Myers (2005) suggest that better support for working sets can help simplify the maintenance task. As a result, support for working sets is included in the Herbal IDE.

As shown in figure 10, the Herbal IDE makes it possible for developers to build a working set of task-relevant code fragments. Working sets can be built manually by the developer or by searching the libraries using keywords related to the current main-

Figure 6. Developing Agents by Both Using the GUI Editor and Editing the Herbal XML by Hand.

Figure 7. The Herbal IDE Showing Multiple Views of an Herbal Library.

tenance task. The collection of code fragments can then be saved as a named working set and shared among developers. Finally, double-clicking items in the working set will open the code fragment in the Herbal GUI editor for inspection.

Herbal: A Tool for Supporting Reuse

The Herbal toolset was designed to support several different forms of reuse from the creation of libraries, instantiation of behavior design patterns,

Figure 8. Ontological Explanations Supported Using the Model Browser View.

Figure 9. Mechanistic Explanations Supported Using the Debugger View.

and the support for the reuse of low-level PSCM components. The Herbal high-level language allows for the creation of libraries of reusable components that are uniquely defined using namespaces. In addition, the language allows for the rationale behind the design of each component to be captured as part of the component's definition. This rationale makes it easier for developers to understand how to reuse existing components.

Libraries

Herbal is library centered, in that Herbal projects consist of XML documents that define several libraries of reusable components. There are six different types of Herbal libraries: types libraries, condition libraries, action libraries, operator libraries, problem-space libraries, and agent libraries.

The dependencies between the contents of these libraries are shown in figure 11. The foundation of all the Herbal libraries is the types library. This library contains the set of data types available to the agent programmer. From these types, the programmer can define conditions and actions that can add, edit, remove, or test for the existence of instances of the defined types. Operators are then built from these conditions and actions, and problem spaces are built from a set of conditions and operators that activate the problem space. Finally, agent behavior can be defined by a hierarchy of problem spaces. This layered approach allows developers to specify behavior at the most appropriate level of abstraction for a given problem.

Herbal libraries are uniquely qualified using a namespace. This allows developers to create any

Figure 10. Support for Working Sets in the Herbal IDE.

Figure 11. The Dependencies between the Six Different Types of Libraries in Herbal.

number of libraries and share them across models, a fundamental benefit of reusable code implemented in most modern, higher-level programming languages. The Herbal IDE supports library sharing graphically using wizards for the importing and exporting of libraries across projects. This feature automatically detects library dependencies, ensuring that the required library components are included in the export.

As an example of how library reuse is supported

Figure 12. Building Custom Agents by Reusing Libraries.

in Herbal, libraries of reusable components for the vacuum cleaner agent environment were created and prefixed with the namespace "Vacuum" (Cohen 2005). These libraries were then combined with other libraries and reused to build new vacuum cleaner agents. For example, figure 12 shows how more aggressive vacuum cleaner agents can be created by reusing the existing vacuum cleaner libraries. In figure 12, more aggressive operators (aggressive operators could be operators that cause vacuums to clean more aggressively and do less wandering) might be created based on conditions and actions contained in the stock vacuum cleaner libraries. These operators are then used to build new problem spaces and aggressive agents based on these problem spaces. This type of reuse has been used in the classroom environment so that students can assemble agents from reusable components, thus allowing them to spend more time focusing on modeling more complex behavior and less time developing components to implement fundamental behaviors.

Design Rationale

Operational explanations help with Krueger's (1992) selection and integration dimensions and therefore are an important key to supporting reuse. The Herbal IDE supports Krueger's concept of selection and integration by allowing developers to filter components based on the component's design rationale.

For example, suppose that a developer wishes to find the set of model components that are responsible for a vacuum cleaner agent cleaning dirty squares. Using the Herbal IDE, model components can be found that are related to the keyword *clean*, and these components can be placed into a working set that can be saved and reused. Herbal's selection mechanism takes full advantage of design rationale, which makes it possible for developers to browse library components based on their operational explanations.

Behavior Design Patterns

In evaluations of Herbal in classroom settings, it became clear that there are certain common metabehaviors that are frequently required to implement basic, recurring agent functionality. For example, many of the agents created by students for the vacuum cleaner environment and the dTank environment (Ritter et al. 2007) implemented looping constructs. For the vacuum cleaner agents, behaviors like "while the vacuum is on a clean square search for dirt using this pattern of movement" were common. For the dTank agents, behaviors like "while no enemy tank is spotted search for an enemy using this search strategy" were repeatedly implemented.

Figure 13. The Behavior Design Pattern Wizard.

Structured programming paradigms like looping constructs are useful in agent programming but can be a challenge to program in a typical rule-based language. This challenge presented a barrier to the students that limited what they could accomplish in their projects. In addition, similar looping constructs are often repeated throughout an agent program. High-level support for these constructs can allow modelers to reuse the behavior they generate, as opposed to duplicating it or creating it from scratch.

To address this problem, and to promote the reuse of metabehaviors such as looping, the behavior design pattern wizard (see figure 13) was incorporated into the Herbal development environment. This wizard makes it possible for the agent developer to generate instantiations of useful metabehaviors using existing PSCM components.

Looping Patterns Common in Agent Behavior

Currently, the behavior design pattern wizard allows developers to instantiate three different patterns of looping: the fixed-order loop, the implied-order loop, and the unordered loop. These loops are currently novel to Herbal, but would be a useful design pattern for most rule-based modeling languages.

The fixed-order loop mimics the classic structured while loop. This loop creates behavior in which a collection of actions is performed in a specific order, but only as long as a set of entry conditions evaluates to true. While a simple construct, this type of loop can be a challenge to create using a rule-based language. Table 5 shows pseudocode for a fixed loop implemented by a vacuum cleaner agent.

The implied-order loop is similar to the fixed-order loop except that some of the items within the loop might not execute during a loop iteration. Table 5 shows pseudocode for an implied-order loop implemented for a vacuum cleaner agent.

The unordered loop has no counterpart in the classical structured programming paradigm because the order of the contents within the loop is not specified. As long as the loop conditions are true, the operators inside the loop become candidates for application, and actions that are possible might not be performed. This provides a source of variability in the behavior. Pseudocode for an unordered loop is also given in table 5.

Additional metabehaviors can be included in the Herbal behavior design pattern wizard by editing a configuration file and providing a custom Java class. This allows developers to increase the library of metabehaviors as more recurring behaviors are identified.

Discussion and Conclusion

The Herbal toolset is an example of applying mod-

Pattern Type	Pseudo code

Fixed Order Loop	```while (current square is clean)```
	```{ move one square to the left```
	``` move one square up```
	``` move one square to the right```
	``` move one square down```
	``` move one square in a random direction }```

Implied Order Loop	```while (current square is clean)```
	```{ if (certain conditions are true)```
	``` move one square to the left```
	``` if (certain conditions are true)```
	``` move one square up```
	``` if (certain conditions are true)```
	``` move one square to the right```
	``` if (certain conditions are true)```
	``` move one square down```
	``` if (certain conditions are true)```
	``` move one square in a random direction }```

Unordered Loop	```while (current square is clean)```
	```{ Choose only one of the following:```
	``` { if (certain conditions are true)```
	``` move one square to the left```
	``` if (certain conditions are true)```
	``` move one square up```
	``` if (certain conditions are true)```
	``` move one square to the right```
	``` if (certain conditions are true)```
	``` move one square down```
	``` if (certain conditions are true)```
	``` move one square in a random direction } }```

Table 5. Pseudocode for Three Different Types of Looping Behavior Patterns.

ern software-engineering principles to agent-programming environments. Specifically, Herbal leverages the software-engineering principles of high-level languages, maintenance-oriented development environments, and software reuse to simplify agent development, in ways that other systems could do as well.

Early evaluations of the Herbal toolset have been promising. In an early study using six undergraduate computer science students (Cohen, Ritter, and Haynes 2009), participants were asked to create an intelligent agent in Jess that piloted two vacuum cleaner agents through a simulated environment: the first agent was created without the use of PSCM, and the second took advantage of PSCM as a hierarchical behavior organizational tool. Because these agents were not designed to be cognitively plausible, PSCM was used primarily for organizing the rules and for making the problem-solving strategy explicit. Jess's support for the concept of modules and focus (see Friedman-Hill

[2003] for more details) made it relatively easy for the participants to implement PSCM in Jess.

Surveys given to the students at the end of the study showed that they favored the use of PSCM in their programs. Student responses showed that they agreed that the use of PSCM made it possible to break up complicated behavior into smaller, less complicated parts. In addition, the survey showed that the students strongly disagreed with the statement that it would be easier to create complicated agents without the use of PSCM.

A second evaluation, using an early version of Herbal, also showed promise. In a study done by Morgan et al. (2005), a Soar model consisting of 29 productions was created using Herbal. In this study, the authors showed a reduction in the time it took for an undergraduate to create productions as the library of reusable components (for example, conditions and actions) expanded (25 minutes for the first few production pairs decreased to 5 minutes or less beyond the tenth production pair).

This reduction in time is believed to be the benefit of increased reuse. In addition, the overall average time per production was less than that reported in a similar study of graduate students programming in Soar (Yost 1993).

A third, summative evaluation of Herbal was also conducted, and the results were positive (Cohen 2008). Using a cognitive dimensions questionnaire (Blackwell and Green 2000), 24 undergraduate students majoring in computer science, computer information science, management information science, and psychology were asked to create an agent using Herbal. The task was broken into three subtasks: creating a reusable library, creating an agent using the library, and finding and fixing a bug in the resulting agent. Data were collected using participant observation and a user reaction survey based on cognitive dimensions (Blackwell and Green 2000). Herbal scored high in five of nine dimensions: the ability easily to make changes to a model (viscosity); the conciseness of the language (diffuseness); the ability to evaluate and obtain feedback from an incomplete agent (progressive evaluation); the closeness of the language to the way the agent behavior is described naturally (closeness of mapping); and the lack of hard mental processing required at the notational level (hard-mental operations). In addition, there was no statistical evidence that there is a correlation between the number and type of observed events during task completion and the participants' major. This is especially interesting because it suggests that Herbal may help make cognitive modeling more accessible to students majoring in areas other than computer science.

In addition to the positive preliminary empirical results, the implementation of the Herbal toolset also provided some important lessons. For example, the trade-off between the power of programming close to the architecture and the simplicity of programming at a higher-level was continually reinforced. On the one hand, basing the Herbal high-level language on PSCM provided some much needed structure and organization to a traditionally rule-based programming environment. However, the absence of the underlying architectural support for the PSCM in Jess created a need to limit or simulate portions of PSCM.

Interestingly, our high-level language also generated opportunities for improving PSCM. As illustrated in tables 2, 3, and 4, the addition of conditions and actions as first-class objects of Herbal added another level of granularity, which allowed for better reuse. While PSCM is originally agnostic about the implementation of operators, this level is real for programmers and the code. In other words, operators that utilize similar conditions and actions no longer need to duplicate the whole operator (previously the smallest unit in PSCM). This

feature was not easy to implement because of the dependencies between actions and conditions (some actions are designed to work with specific conditions). These dependencies were reduced by providing language support for wiring conditions to actions at the time they are used. This suggests that the preconditions and actions of operators, previously on the level below PSCM, can be promoted from the symbol level to the PSCM level at least in system design. Without the higher-level constructs provided by Herbal, this would not be possible in Soar models.

Another interesting lesson was the need to support structured programming techniques in what is traditionally an unstructured rule-based environment. For example, many of the agents created by our students in various class projects implemented different types of looping constructs. However, creating these constructs in a rule-based language was quite challenging. By including high-level support for structured programming paradigms, such as looping constructs, the agent-programming task was significantly simplified, and yet the rule-based paradigm was extended, not discarded.

Some unexpected instructional benefits of the Herbal toolset were also discovered during this project. Initially designed to reduce the need to program at a low level, the Herbal high-level language and GUI editor also appear to be valuable for teaching low-level rule-based programming. Working with the Herbal GUI editor and the Herbal high-level language editor side-by-side, programmers can learn the Herbal language by making changes graphically and then viewing the generated XML code. In addition, by editing the Herbal XML code directly and then viewing the generated low-level productions, programmers can learn native Jess and Soar programming.

Herbal has also been used to create several models examining learning. Cohen, Ritter, and Haynes (2007) created a competitive, reflective learning model as well as several opponent models in an afternoon. Friedrich (2008; Friedrich and Ritter 2009) created a problem-solving model that learned the Diag task using five different problem-solving strategies. The first model took six months to write; the revised models and strategies took about a month. Haynes et al. (2008) used Herbal to create a model of an antiterrorism force protection planner as part of the Rampart project. This model assists users in selecting between various resource allocation options. This model has over 100 operators; maintaining an equivalent model in Soar would be difficult. Finally, we have recently created an Herbal to ACT-R compiler, and it has helped us create a 500-rule model with 10 different processing times and learning rates (different mixes of procedural and declarative knowledge).

Developing intelligent agents is a complex software-engineering activity. Creating complex software is not a new problem, and the software-engineering community has developed strong theories and principles about how to solve complex problems with software solutions. The Herbal toolset is one example of how these lessons from software engineering can be applied to help simplify the task of agent development. Some of the benefits of applying software-engineering principles are being realized with the creation of the Herbal toolset. The complete Herbal toolset is currently available and can be downloaded from the Herbal website.[9]

Acknowledgements

Notes

1. See www.w3.org/XML.
2. See sitemaker.umich.edu/soar.
3. See herzberg.ca.sandia.gov/jess.
4. See www.w3.org/Style/XSL.
5. See www.w3.org/Graphics/SVG.
6. See www.w3.org/XML/Schema.
7. See acs.ist.psu.edu/herbal.
8. See eclipse.org.
9. See acs.ist.psu.edu/herbal.

References

Beck, L. L., and Perkins, T. E. 1983. A Survey of Software Engineering Practice: Tools, Method, and Results. *IEEE Transactions on Software Engineering* 9(5): 541–561.

Blackwell, A. F., and Green, T. R. G. 2000. A Cognitive Dimensions Questionnaire Optimised for Users. Paper presented at the 12th Annual Meeting of the Psychology of Programming Interest Group, Cosenza, Italy, 10–13 April.

Boehm, B. W. 1987. Improving Software Productivity. *IEEE Computer* 20(9): 43–57.

Boehm, B. W. 1988. Understanding and Controlling Software Costs. *IEEE Transactions on Software Engineering* 15(10): 1462–1477.

Brooks, F. P. 1987. No Silver Bullet: Essence and Accidents of Software Engineering. *IEEE Computer* 20(4): 10–19.

Brooks, F. P. 1995. *The Mythical Man-Month: Essays on Software Engineering*. Reading, MA: Addison-Wesley.

Clancey, W. J. 1981. The Epistemology of a Rule-Based Expert System: A Framework for Explanation. Technical Report STAN-CS-91-896, Stanford University, Stanford, CA.

Cohen, M. A. 2008. A Theory-Based Environment for Creating Reusable Cognitive Models. Ph.D. diss., College of Information Sciences and Technology, The Pennsylvania State University, University Park, PA.

Cohen, M. A. 2005. Teaching Agent Programming Using Custom Environments and Jess. *The Newsletter of the Society for the Study of Artificial Intelligence and the Simulation of Behavior,* 120 (Spring): 4.

Cohen, M. A.; Ritter F. E.; and Haynes S. R. 2009. Evaluating Design: A Formative Evaluation of Agent Development Environments Used for Teaching Rule-Based Programming. In *Proceedings of the Information Systems Education Conference 2009,* 1542–7382, Washington, DC. Chicago, IL: Education Special Interest Group of the Association of Information Technology Professionals.

Cohen, M. A.; Ritter, F. E.; and Haynes, S. R. 2007. Using Reflective Learning to Master Opponent Strategy in a Competitive Environment. In *Proceedings of the 8th International Conference on Cognitive Modeling,* 157–162. Oxford, UK: Taylor and Francis.

Cooper, R. P., and Fox, J. 1998. Cogent: A Visual Design Environment for Cognitive Modeling. *Behavior Research Methods, Instruments, and Computers* 30(4): 553–564.

Daly, E. B. 1977. Management of Software Development. *IEEE Transactions on Software Engineering* 3(3): 229–242.

Friedman-Hill, E. 2003. *Jess in Action: Rule-Based Systems in Java.* Greenwich, CT: Manning Publications Company.

Friedrich, M.; Cohen, M. A.; and Ritter, F. E. 2007. *A Gentle Introduction to XML within Herbal*. University Park, PA: ACS Lab, The Pennsylvania State University.

Friedrich, M. B. 2008. Implementierung von schematischen Denkstrategien in einer höheren Programmiersprache: Erweitern und Testen der vorhandenen Resultate durch Erfassen von zusätzlichen Daten und das Erstellen von weiteren Strategien (Implementing Diagrammatic Reasoning Strategies in a High-Level Language: Extending and Testing the Existing Model Results by Gathering Additional Data and Creating Additional Strategies). Faculty of Information Systems and Applied Computer Science, University of Bamberg, Germany.

Friedrich, M. B., and Ritter, F. E. 2009. Reimplementing a Diagrammatic Reasoning Model in Herbal. Paper presented at the Ninth International Conference on Cognitive Modeling, Manchester, England, 24–26 July.

Gamma, E.; Helm, R.; Johnson, R.; and Vlissides, J. 1995. *Design Patterns: Elements of Reusable Object-Oriented Software*. Reading, MA: Addison-Wesley.

Haynes, S. R. 2005. Three Studies of Design Rationale as Explanation. In *Rationale Management in Software Engineering,* ed. R. Dutoit, I. Mistrik, and B. Paech, 53–71. Berlin: Springer.

Haynes, S. R.; Cohen, M. A.; Ritter, F. E. 2009. Design Patterns for Explaining Intelligent Systems. *International Journal of Human-Computer Studies* 67(1). 99–110.

Haynes, S. R.; Kannampallil, T. G.; Cohen, M. A.; Soares, A.; and Ritter, F. E. 2008. Rampart: A Service and Agent-Based Architecture for Anti-Terrorism Planning and Resource Allocation. In *Proceedings of the First European Conference on Intelligence and Security Informatics,* Euroisi 2008 (Esbjerg, Denmark 2008), 260–270. Berlin: Springer.

Hordijk, W., and Wieringa, R. 2005. Surveying the Factors That Influence Maintainability. In *Proceedings of the European Software Engineering Conference (ESEC) and the ACM SIGSOFT Symposium on the Foundations of Software*

Engineering (FSE), 385–388. New York: Association for Computing Machinery.

Ivory, M. Y., and Hearst, M. A. 2001. The State of the Art in Automating Usability Evaluation of User Interfaces. *Computing Surveys* 3(4): 470–516.

John, B. E., and Kieras, D. E. 1996. The GOMS Family of User Interface Analysis Techniques: Comparison and Contrast. *ACM Transactions on Computer-Human Interaction* 3(4): 320–351.

John, B. E.; Prevas, K.; Salvucci, D. D.; and Koedinger, K. 2004. Predictive Human Performance Modeling Made Easy. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems 2004,* 455–462. New York: Association for Computing Machinery.

Jones, R. M.; Crossman, J. A. L.; Lebiere, C.; and Best, B. J. 2006. An Abstract Language for Cognitive Modeling. In *Proceedings of the International Conference on Cognitive Modeling,* 160–165. Mahwah, NJ: Lawrence Erlbaum.

Jones, R. M.; Laird, J. E.; Nielsen, P. E.; Coulter, K. J.; Kenny, P.; and Koss, F. V. 1999. Automated Intelligent Pilots for Combat Flight Simulation. *AI Magazine* 20(1): 27–41.

Ko, A. J.; Aung, H. H.; and Myers, B. A. 2005. Eliciting Design Requirements for Maintenance-Oriented IDEs: A Detailed Study of Corrective and Perfective Maintenance Tasks. In *Proceedings of the International Conference on Software Engineering,* 126–135. New York: Association for Computing Machinery.

Krueger, C. W. 1992. Software Reuse. *ACM Computer Surveys* 24(2): 131–183.

Laird, J. E. 2001. It Knows What You're Going to Do: Adding Anticipation to a Quakebot. In *Proceedings of the Fifth International Conference on Autonomous Agents,* 385–392. New York: Association for Computing Machinery.

LaToza, T. D.; Venolia, G.; and DeLine, R. 2006. Maintaining Mental Models: A Study of Developer Work Habits. In *Proceedings of the 28th International Conference on Software Engineering,* Shanghai, China, 492–501. New York: Association for Computing Machinery.

Lehman, J. F.; Laird, J. E.; and Rosenbloom, P. S. 1996. A Gentle Introduction to Soar, an Architecture for Human Cognition. In *Invitation to Cognitive Science,* ed. S. Sternberg and D. Scarborough, Volume 4, 211–253. Cambridge, MA: The MIT Press.

Maxwell, K. D.; Wassenhove, L. V.; and Dutta, S. 1996. Software Development Productivity of European Space, Military, and Industrial Applications. *IEEE Transactions on Software Engineering* 22(10): 706–718.

McIlroy, M. D. 1968. Mass Produced Software Components. In *Proceedings of Software Engineering—Report on a Conference by the NATO Science Committee*, 138–150. Brussels, Belgium: NATO Scientific and Environmental Affairs Division, North Atlantic Treaty Organization.

Morgan, G. P.; Cohen, A. M.; Haynes, S. R.; and Ritter, F. E. 2005. Increasing Efficiency of the Development of User Models. In *Proceedings of the IEEE System Information and Engineering Design Symposium,* Charlottesville, VA. Piscataway, NJ: Institute of Electrical and Electronics Engineers, Inc.

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

Newell, A.; Yost, G. R.; Laird, J. E.; Rosenbloom, P.;, and Altmann, E. 1991. Formulating the Problem Space Computational Model. In *Carnegie Mellon Computer Science: A 25-Year Commemorative,* ed. R. F. Rashid, 255–293. Reading, MA: Addison-Wesley.

Pew, R. W., and Mavor, A. S., eds. 1998. *Modeling Human and Organizational Behavior: Application to Military Simulations.* Washington, DC: National Academy Press.

Ritter, F. E.; Haynes, S. R.; Cohen, M. A.; Howes, A.; John, B. E.; Best, B.; Lebiere, C.; Jones, R. M.; Lewis, R. L.; St Amant, R.; McBride, S. P.; Urbas, L.; Leuchter, S.; and Vera, A. 2006. High-Level Behavior Representation Languages Revisited. Paper presented at the Seventh International Conference on Cognitive Modeling. Trieste, Italy, 5–8 April.

Ritter, F. E.; Kase, S. E.; Bhandarkar, D.; Lewis, B.; and Cohen, A. M. 2007. Dtank Updated: Exploring Moderator-Influenced Behavior in a Light-Weight Synthetic Environment. In *Proceedings of the 16th Conference on Behavior Representation in Modeling and Simulation,* 51–60. Orlando, FL. Red Hook, NY: Curran Associates Inc.

Ritter, F. E.; Shadbolt, N. R.; Elliman, D.; Young, R.; Gobet, F.; and Baxter, G. D. 2003. *Techniques for Modeling Human and Organizational Behavior in Synthetic Environments: A Supplementary Review.* Wright-Patterson Air Force Base, OH: Human Systems Information Analysis Center.

Salvucci, D. D., and Lee, F. J. 2003. Simple Cognitive Modeling in a Complex Cognitive Architecture. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems,* 265–272. New York: Association for Computing Machinery.

St. Amant, R.; Freed, A. R.; and Ritter, F. E. 2005. Specifying ACT-R Models of User Interaction with a GOMS Language. *Cognitive Systems Research* 6(1): 71–88.

Swartout, W. R. 1983. Xplain: A System for Creating and Explaining Expert Consulting Programs. *Artificial Intelligence* 21(3) 285–325.

Tambe, M.; Johnson, W. L.; Jones, R. M.; Koss, F. V.; Rosenbloom, P.; and Schwamb, K. 1995. Intelligent Agents for Interactive Simulation Environments. *AI Magazine* 16(1): 15–39.

Tassey, G. 2002. *The Economic Impacts of Inadequate Infrastructure for Software Testing* (No. 7007.011). Washington, DC: National Institute of Standards and Technology.

Yost, G. R. 1993. Acquiring Knowledge in Soar. *IEEE Expert: Intelligent Systems and their Applications* 8(3): 26–34.

Mark Cohen is an associate professor in the Business Administration, Computer Science, and Information Technology Department at Lock Haven University. His current research efforts include developing software that simplifies the creation and maintenance of cognitive models. His email address is mcohen@lhup.edu.

Frank Ritter is one of the founding faculty of the College of Information Sciences and Technology, an interdisciplinary academic unit at Pennsylvania State University to study how people process information using technology. He works on the development, application, and methodology of cognitive models, particularly as applied to interfaces and emotions. His email address is frank.ritter@psu.edu.

Steven Haynes is a professor of practice in the College of Information Sciences and Technolgy at Pennsylvania State University. He researches system design, modeling, and development; human-computer interaction; design rationale; system explanation; and the philosophy of technology. Prior to entering academe he worked at Apple Computer, Adobe Systems, and several smaller technology companies in the United States and in Europe. His email address is shaynes@ist.psu.edu.