THE FLOW TUTOR: A SCHEMA-BASED TUTORIAL SYSTEM

Donald R. Gentner
Center for Human Information Processing
University of California, San Diego
La Jolla, California 92093

The FLOW tutor system is a simulation of a human tutor which gives advice to a student learning FLOW, a simple computer language. The tutor has a schema-based knowledge structure containing information about the programming language, the student's instruction booklet, and the student's developing knowledge. These schemas form the basis of a distributed intelligence system which uses conceptually guided and data-driven processing to interpret the student's behavior, update the model of the student, and give advice to the student.

The Tutor's Task

The student uses an instruction booklet, containing text, examples and problems, and works at a computer terminal. The automated tutor interprets the student's keypresses and pauses in terms of a student progressing through the FLOW instruction booklet. All schemas which the FLOW tutor creates and works with are part of a multi-level structure, ranging from the instruction booklet and the function of programs to the individual FLOW statements and keypresses.

Conceptually guided prediction provides a context to disambiguate student keypresses. The automated tutor simulates the student's progress through the instruction booklet, allowing pauses for reading, doing the examples, and solving problems, and predicts that the actual student will follow a similar course. The tutor "solves" these simple programming problems by expanding the functional description of the problem into simpler functions, FLOW statements, and finally keypresses. Although this can lead to severe trouble when the predictions are wrong, it normally leads to an easy and efficient interpretation of otherwise ambiguous information.

Data-driven processing, in which the keypresses and simple statements try to assemble themselves into higher level schemas, is used when student behavior does not match the predictions. The identity of two functionally equivalent programs at a higher level of schemas, allows the tutor to interpret alternative correct solutions.

The automated FLOW tutor is intended to simulate an experienced human tutor familiar with FLOW and common student problems, and thus has schemas which represent typical student errors. When unexpected student inputs incorporate themselves into these error schemas, the tutor has effectively recognized a student error. If the student pauses after making an error (the allowed pause length depends on the type of error and the tutor's model of the student), the tutor sends advioe to the student based on its interpretation of the error.

Schemas

An instantiation of a schema corresponding to the FLOW statement 030 DISPLAY "JEAN" is shown below:

```
•DISPLAY-QU0TED-STRING-1932
    schema DISPLAY-QU0TED-STRING
    statement-number 030
    value JEAN
    status SATISFIED
    host *DISPLAY-1911
    element *D-1937
    element *QUOTED-STRING-WI
```

The instance consists of a name followed by a series of slots. The first slot points back to the prototype schema. The next two slots are "argument" slots, used to distinguish individual instances of the schema. The "status" slot indicates that all components of this instance have been found. Schemas and instances are composed of other schemas and instances. The "host" and "element" slots point respectively to the higher and lower level instances. This hierarchical structure of instances plays a major role in the operation of the FLOW tutor. Extension of the hierarchy to higher and lower levels forms the basis for predicting and interpreting student behavior, giving the tutor descriptions of the same information at different conceptual levels.

How Schemas Function,

As a schema-based system, the FLOW tutor has an inherently distributed intelligence. The only central coordination is furnished by an agenda. a simple list of instances waiting to be active.

Each schema has an associated process, called the specialist. When an instance becomes active in the FLOW tutor system, the corresponding specialist is invoked. The specialist may modify an instance, predict new instances of schemas, incorporate the instance into a higher level schema, change the tutor's model of the world, look for input from the student, put instances on the agenda, or send messages to the student. In a typical case, an instance on the agenda might have been predicted by some other instance. When the instance becomes active, its specialist would check to see if all of its elements had been observed. If not, the specialist would predict the next element and put it on the agenda. If all the elements of the instance had been observed, the specialist would place its host on the agenda, or if the instance had no host, the specialist might search for a host and try to incorporate the instance into a suitable higher level instance.

This research was supported by the Advanced Research Projects Agency and the Office of Naval Research of the Department of Defense and was monitored by ONR under Contract No.: N0014-77-C-0628. Donald Norman and Mark Viallen made significant contributions to the research.

For further details see:

Gentner, D. R., & Norman, D. A. The FLOW tutor: Schemas for tutoring (Tech. Rep. 7702). La Jolla, Ca.: University of California, San Diego, Center for Human Information Processing, 1977.