A Domain Description Language for Sketch Recognition

Tracy Hammond HAMMOND@AI.MIT.EDU
MIT Artificial Intelligence Laboratory, 200 Technology Square, Cambridge MA, 02139 USA

Abstract

In this paper, we propose a domain description language used to describe domain-specific information to a domain-independent sketch recognition system. Although the language is primarily based on shape, the domain description can include any type of information that would be helpful to the recognition process, such as stroke order or direction. The language consists of pre-defined shapes, constraints, and editing behaviors, as well as a syntax for specifying a domain description.

1. Introduction

Pervasive environments, complete with digital whiteboards and pocket PC's, have increasingly included applications with sketchable interfaces. Sketch recognition applications built for the Oxygen platform include Ligature (Foltz, 2001), Tahuti (Hammond & Davis, 2002), and Assist (Alvarado, 2000) / Assistance (Oltmans, 2000). To date, sketch recognition systems have been domain-specific, with the recognition details of the domain hard-coded into the system. A domain-independent recognition system is advantageous since it may be used for several domains, increasing the flexibility and capabilities of a system. However, the system cannot identify the domain shapes if it doesn't know that they are. In order to properly recognize a sketch of a particular domain, domain-specific information must be supplied to the domain-independent recognition system.

Sketch recognition can be done by measuring many features. Many of these features are not necessarily correlated to the shape of the drawn object and place requirements on the user to draw an object in a single stroke and in a particular direction. By recognizing objects based on shape, we not only ensure correlation between the drawn shape and the recognized shapes, but we also enable designers to draw the shapes as they would naturally.

In this paper, we propose a domain description language used to describe domain-specific information to a domain-independent sketch recognition system. Although the language is primarily based on shape, the domain description can include any type of information that would be helpful to the recognition process, such as stroke order or direction. By describing domain descriptions using the language's syntax, designers can add sketch recognition to their user interfaces.

1.1 Previous Work

Shape description languages have been around for a long time (Stiny & Gips, 1972). These grammars have been studied widely within the field or architecture, and many systems are still built using shape grammars (Gips, 1999). However, they have been developed for design generation rather than recognition, and don't provide for non-graphical information, such as stroke order, that may be helpful in recognition.

Within the field of sketch recognition, there have been other attempts to create shape languages for sketch recognition. Mahoney and Fromherz (2002) use a language to model and recognize stick figures. The language currently is not hierarchical, making large objects cumbersome to describe. Caetano et al. (2002) use fuzzy relational grammars and Bimber et al. (2000) use BNF grammars to describe shape information. Both lack the ability to describe non-shape domain information such as stroke order or direction and editing behavior information.

2. Domain Description Language

The difficulties in determining the language's components and syntax include ensuring that the language allows all common helpful domain information to be specified. The language must also encourage and facilitate the creation of correct programs. For instance, to encourage the reuse of geometric shape definitions, the language distinguishes between geometric shape definitions (shapes usable in many domains) and domain shapes (shapes specific to a domain). The language also provides abstract shape definitions that describe a class of similar shapes to prevent rewriting of identical attributes.

The language consists of pre-defined shapes, constraints, editing behaviors, as well as a syntax for combining them. A domain description is specified by 1) a list of the shapes

and shape compositions in the domain, 2) shape definitions, 3) domain shape definitions, 4) abstract shape definitions, 5) domain shape composition definitions (how shapes interact), 6) constraint definitions, and 7) editing behavior definitions.

2.1 Shape Definitions

A shape definition describes shapes usable in multiple domains. A shape definition is composed of seven components. The *description* (line 1 in Figure 1) is a textual description of the shape. The *is-a* section (line 2) is an indication of any class of abstract shapes that it belongs to. The *components* (line 3) include the geometrical shapes of which this shape is composed. (Shapes are defined hierarchically.) Note that the TriangleArrow is composed of a pre-defined shape Line as well a user-defined shape OpenArrow. The *constraints* (line 4) specify the necessary relationships and can also specify probable, but not required, drawing order. For instance, a probable drawing order may be shaft, head1, l, head2. All constraints in this example are pre-defined. The *derived properties* (line 5) allows us to compute certain properties and name them for use later. The *display* section (line 6) defines what should be displayed on the screen. The default is the original strokes. Generally, the original strokes are shown for all geometrical shapes, and the display only changed for domain shapes. *Editing behaviors* (line 7) can be defined for each shape. The editing behavior below allows the user to move the entire arrow by clicking and dragging the shaft. The user can also click and drag the head or tail of the arrow while the opposite end remains fixed; the shaft stretches and rotates as appropriate.

Figure 1: Shape Definition for a Triangle Arrow.

```
(define sketch-shape TriangleArrow
  (description "An arrow with a triangle head")     %1
  (is-a Arrow)                                      %2
  (components (OpenArrow oa) (Line l))              %3
  (constraints                                      %4
    (meet l.p1 oa.head1.p1) (meet l.p2 oa.head2.p1)
    (angle oa.shaft l 90) (angle l oa.head1 45)
    (angle l oa.head2 45)
    (probable draw-order oa.shaft oa.head1 l oa.head2))
  (derived-properties                               %5
    (Point head oa.shaft.p2) (Point tail oa.shaft.p2)
    (Line shaft oa.shaft) (Line head1 oa.head1)
    (Line head2 oa.head2))
  (display                                          %6
    (cleaned_strokes shaft)(ideal_strokes l head1 head2)
  (editing-behavior                                 %7
    (click_hold_drag head
     (fix tail) (stretch_scale_rotate this) (move head))
    (click_hold_drag tail
     (fix head) (stretch_scale_rotate this) (move head))
    (click_hold_drag shaft (move this))
    (scribble shaft (delete this))))
```

3. Conclusion

3.1 Future Work

In the future, we will test human usability by asking users to develop domain descriptions using the proposed language. We will test that these descriptions agree with the users' intensions by developing a simple domain-independent sketch recognition system.

3.2 Contributions

In this paper we present a language for describing domain-specific information to a domain-independent sketch recognition system. The language is based on shape to ensure correlation between the drawn shape and the recognized shape, and provides for natural sketching interaction. The language is different from other such languages because it can be also be to describe non-shape information, including display information, editing behavior, and drawing order.

Acknowledgements

The language described in this paper is part of a larger project (Hammond et al., 2002) that includes learning shape descriptions, compiling shape descriptions, and recognizing sketches based on these descriptions. This work is being done by members of the Design Rationale Group at MIT, lead by Randall Davis, and include Christine Alvarado, Tracy Hammond, Michael Oltmans, Metin Sezgin, and Olya Veselova. This work is supported in part by the MIT Project Oxygen.

References

Alvarado, C. (2000). A natural sketching environmant: Bringing the computer into early stages of mechanical design. Master's thesis, MIT.

Bimber, O., L.M.Encarnao, & Stork, A. (2000). A multi-layered architecture for sketch-based interaction within virtual environments. *Computer and Graphics*.

Caetano, A., Goulart, N., Fonseca, M., & Jorge, J. (2002). Javasketchit: Issues in sketching the look of user interfaces. *AAAI Spring Symposium on Sketch Understanding*.

Foltz, M. (2001). Ligature, gesture-based configuration of the e21 intelligent environment. *MIT Student Oxygen Workshop*.

Gips, J. (1999). Computer implementation of shape grammars. *NSF/MIT Workshop on Shape Computation*.

Hammond, T., & Davis, R. (2002). Tahuti:a geometrical sketch recognition system for uml class diagrams. *AAAI Spring Symposium on Sketch Understanding*, 59–68.

Hammond, T., Sezgin, M., Veselova, O., Adler, A., Oltmans, M., Alvarado, C., & Hitchcock, R. (2002). Multi-domain sketch recognition. *MIT Student Oxygen Workshop*.

Mahoney, J. V., & Fromherz, M. P. J. (2002). Three main concerns in sketch recognition and an approach to addressing them. *AAAI Spring Symposium on Sketch Understanding*, 105–112.

Oltmans, M. (2000). Understanding naturally conveyed explanations of device behavior. Master's thesis, MIT.

Stiny, G., & Gips, J. (1972). Shape grammars and the generative specification of painting and sculpture. *Information Processing*, 1460–1465.