Teleological Representations for Multimodal Design Explanations

Jon Wetzel and Ken Forbus
Qualitative Reasoning Group, Northwestern University
2133 Sheridan Road, Evanston, IL, 60201, USA
jw@northwestern.edu forbus@northwestern.edu

Abstract

An important problem in teaching engineering design is helping students learn to communicate via sketching. A good explanation of a design should include how the structure and behavior of the design enable it to achieve its intended purpose. Thus the ability to represent and reason about teleology becomes important for creating software coaches for engineering design. This paper describes a simple representation for teleology that enables a sketch-based coaching system to provide better feedback on student designs. We show the sufficiency of this representation by evaluating it over a corpus of student-generated sketches.

Introduction

One important skill for many engineers is the use of sketches to communicate a design to others. At Northwestern University, freshmen in engineering majors are taught this skill in their introductory course, EDC (Engineering Design and Communication). However, instructors in EDC found that students have trouble learning to communicate with sketches. Our goal is to help these students by creating an intelligent coach that can give them feedback on their engineering design explanations.

These engineering design explanations are multimodal, consisting of both sketches and language. In human-to-human sketching, language is used to disambiguate and complete the information depicted in the sketch. This is particularly true of *teleological* knowledge, which explains the function of the design. The general problem of representing and reasoning about teleology is extremely difficult, in part because of the extreme breadth of purposes that people design things for. Fortunately, for any particular design exercise, there is a reasonably small set of teleological concepts that is needed. By providing a general framework that can be filled in incrementally as needed, we can develop and extend teleological representations as needed to cover new classes of designs.

This paper presents simple representations and reasoning that are sufficient to do the teleological evaluation of multimodal design explanations for a particular class of designs. We begin by briefly reviewing CogSketch, qualitative mechanics, our previous work in engineering design understanding, and design rationale. Then we describe our approach to teleological representations and using them in critiquing. We demonstrate its performance on a corpus of sketches drawn by engineering students. Finally we conclude with related and future work.

Background

Design Coach[1] is a sketch-based educational software system for coaching students in explaining engineering designs. We summarize briefly the ideas underlying it, and its prior capabilities, to set the stage for explaining the new capabilities.

Design Coach is built on CogSketch [Forbus et al 2011], an open domain sketch understanding system. As students sketch, CogSketch creates a qualitative representation of their sketch. CogSketch also incorporates contents from the OpenCyc[2] knowledge base, augmented with domain-specific knowledge for qualitative reasoning and design. This includes a model of qualitative mechanics (QM) [Nielsen 1988][Kim 1993]. Given a qualitative representation of a mechanical system, QM can predict how objects will move. In [Wetzel & Forbus 2009], we adapted this model to work with CogSketch sketches as input. Qualitative representations are a natural fit for our work both for their explanatory power and because the designs EDC students are sketching are in the conceptual stages, where the quantitative information about parameters and shapes that would be required for numerical simulation simply is not available.

Critiquing Engineering Design Explanations

Students describe their design by drawing one or more interlinked *subsketches*, creating a *comic graph* (Figure 1). They also use the *Tell window* to create sentences from

[1] This system was previously called Design Buddy.
[2] www.opencyc.org

Wetzel, J. and Forbus, K. (2012). Teleological Representations for Multi-Modal Design Explanations. Proceedings of the 26th International Workshop on Qualitative Reasoning. Los Angeles, California.

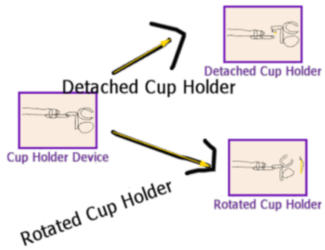

Figure 1: An example of a design, with three subsketches

templates (Figure 2) that are directly converted into predicate calculus representations that are understandable by the system. Students can request feedback at any time. Feedback is generated by two techniques [Wetzel & Forbus 2010]: (1) *State Transition Verification* (STV) which uses qualitative mechanics to tell if the motions depicted in the series of sketches would actually occur. (2) *Sequential Explanation Analysis* (SEA) analyzes the predicate calculus generated for their sentences to determine if each statement is true (i.e. provable using qualitative mechanics, the sketch, and KB contents).

While these methods gave feedback on the motions of parts in a design, students could not describe the intended purpose of their designs, nor could the coach give them feedback based on their intent. The extensions we describe next provide this ability.

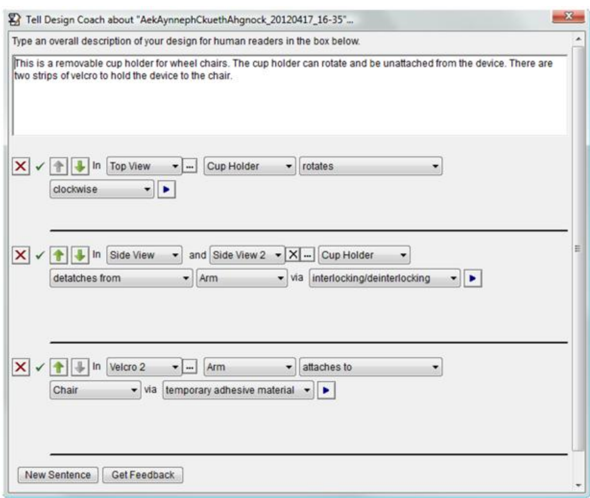

Figure 2 : The Tell window enables students enter to sentences about their design via a dynamic form

Design Rationale and Teleological Representations

A number of ontologies have been proposed for teleology and design rationale. After considering several alternatives, we selected a part of the ontology of device function developed by Kitamura et al [2006] to use as a framework for our representations. They describe artifacts using *function decomposition trees*, in which the *goal function* of a system is recursively decomposed into a tree

- **Easy to attach using Velcro**

- **Cup holder rotates giving user flexibility in placement**

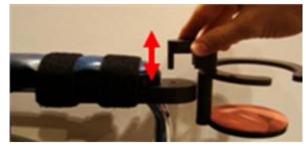

- **Cup holder slides off for quick attachment/removal**

Figure 3: Removable cup holder device, which students explained to Design Coach for the Fall and Winter homework assignments.

of *method functions* which, when performed, achieve their parent goal. In these trees, parents are related to their child functions via a type of relationship called a *way-of-function-achievement*, or a way for short. In their ontology, each way has a list of generic functions which accomplish its goal. For example, the "heat water" function can be achieved by a "heat transfer" way which specifies two generic child functions: "transform electricity to heat" and "give heat to water". They showed that engineers at a semiconductor manufacturing plant could use the ontology to build their own catalogue of functional knowledge covering several kinds of manufacturing tasks.

In Design Coach we employ this function/way model for our internal teleological ontology, but unlike with the aforementioned engineers, the students are never directly exposed to it. Also, the design models created by the engineers were much deeper than our first version, but our current application, described next, limits the domain Design Coach must cover for now.

Approach

Adding a new teleological ontology for Design Coach included building the ontology itself, and extending our structured language input system to include those terms.

Design Coach in the Classroom

We performed an in-class experiment in Fall 2011 to gather information about what teleological language students were likely to use on a typical design problem. Working with EDC instructors, the following assignment was created:

1. Create a sketch of the given design (a removable cup holder for a wheelchair, see Figure 3) and explain it to the Design Coach.

2. Add a specific use-case scenario to help illustrate benefits or limitations of this design. For example, it might be used in a hospital setting for a range of patients, or it could be used outdoors to carry a cold drink on a hot, humid day.
3. Consider possible refinements to the design. Use Design Coach to draw and describe one variation, and get feedback on it.

For each of the three sketches, the students were also asked to write a 3-5 sentence description in the free text area in English for their instructor, as well as write about the Design Coach's feedback. Two sections of 16 students each were given the assignment, and of these, 24 students submitted sketches.

Teleology for a design

From the English descriptions written in the free text area of the Tell Window, we collected a list of functions students mentioned in both the initial design and their refinements. For example, one student wrote:

> "The Under Armrest Rotating Cup Holder design allows it to be easily attached to an armrest using Velcro. The cup holder can be rotated for flexibility in placement. The cup holder also detaches from the cup holder arm."

From such descriptions we derived the following list:

- Attaching/detaching parts
- Rotating cup into desired position for user
- Accommodating multiple cup sizes
- Changing the angle of attachment
- Spill prevention
- Securing the cup holder so it doesn't fall out during chair movement
- Adding additional functionality (e.g. another compartment for holding snack foods)
- Changing the shape or configuration of objects to make the user more comfortable in some way

Using this list as a source, we focused on the tasks which involved mechanics or space, since we have representations for those domains already. We created a set of more general functions and ways of achieving them (Table 1).

Table 1: Teleological Ontology

Function	Ways of Achieving Function
Attachment	Interlocking Parts Adhesive Material Temporary Adhesive Material
Detachment	De-interlocking Parts Temporary Adhesive Material
Adapting in Size	Adjustable Parts
Containment	Enclosure Prohibiting Downward Motion
Comfort	Change in Shape Moving within reach

Formally represented versions of these function and way concepts were integrated into the Cyc ontology. A new relationship was added to link them up to a design, named functionAchievedVia:

```
(functionAchievedVia
    <function> <way> <parts> <contexts>)
```

- *<function>* is the teleological function
- *<way>* is the way *<function>* is achieved
- *<parts>* is the list of parts involved
- *<contexts>* is the list of logical contexts involved, i.e. which subsketch(es) depict the achievement of *<function>*

For example, the student who drew the sketch in Figure 1 constructed a sentence which said the cup holder detaches from its base:

```
(functionAchievedVia Detachment-TF
  InterlockingPartsWay
  (TheList Object-96 Object-130)
  (TheList BCase-3530397068
           BCase-3530397204))))
```

Detachment-TF is the detachment teleological function, InterlockingPartsWay is the way of achieving it, Object-96 and Object-130 are the cup holder and the base, and the two BCase-s are the names of the logical contexts representing the subsketches.

To complete the teleological knowledge for the Design Coach, we created rules to handle queries involving functionAchievedBy. These rules check for behaviors which fulfill the function in question, and/or for required sub-functions. For example, for detachment to be achieved a state of attachment between the parts must have been achieved and another state must exist demonstrating their separation.

Structured Language Input

We incorporated the above ontology into our structured language input system by extending its grammar and vocabulary. First we added a new verb phrase for each of the functions. For example, the attachment function was added as "attaches to" and the comfort function as "is made more comfortable". For the functions that require a subject and object (e.g. attachment), the object field is displayed and populated with eligible items from the sketch. When a function verb is selected, a new form for selecting the way of achieving that function appears.

The sample student English description can be expressed in our structured language using the following three statements[3]:

- In state 1, Velcro[4] is velcro which causes Cupholder Arm attaches to Arm of Chair via adhesive material.

[3] We note that the sentences are sometimes clumsy grammatically; this has not been a problem, based on our observations of students using it.

- In state 2, Cup Holder rotates clockwise which causes Cup Holder is made more comfortable via moving it.
- In state 1 and state 3, Cupholder detaches from Cupholder Arm via interlocking/deinterlocking.

In the first sentence, the first clause tells Design Coach that the object in the sketch named “Velcro” is of type VelcroTheFastener. That, combined with the spatial information about the three objects, explains to Design Coach how the cup holder arm is attached to the arm of the chair. In the second sentence, we express the causal link between the rotating motion and its purpose.

Some functions, like detachment in the third sentence, occur over multiple states in the comic graph. For these we have changed the tell window to allow users to select multiple states for a given sentence instead of just one.

Feedback Generation

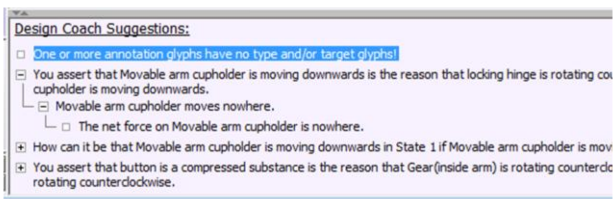

Figure 4: Drill-down on explanations

Teleological statements are entered through the structured language input and feedback on them is generated through Sequential Explanation Analysis, which reads through the sentences one by one and checks for contradictions. For the teleological clauses, Design Coach queries to see if the function is achieved in the stated way.

The feedback is converted to English using templates. In the winter quarter, we also added a structured explanation system, similar to that [Forbus et al 1999], where students can click on feedback in an outline format to learn more about it. In doing so, they are traversing justification structures created as Design Coach inferred the facts which caused it to give that specific piece of feedback. However, it is a guided tour: rules tell the system what justifications to leave out and what additional information to add to make the explanation concise but complete for a student’s purposes.

Evaluation and Discussion

This section explains our evaluation method and presents the overall results, using examples to demonstrate the range of new feedback given by the Design Coach.

Method

After implementing the ontology and adding it to the structured language, we entered the language from the student’s free text as structured sentences in their existing sketches, to see if the system would successfully give the right kind of feedback. We did this for parts 1 and 3 of their design, as most students just used part 2 to set up a design problem that they then solved in part 3. This gave us an initial corpus of 40 multimodal explanations. In winter quarter 2012, we continued our classroom activities with the new ontology in place. At the instructors’ recommendation, we also modified the assignment by eliminating the second step. We received from 25 students a total of 48 sketch files, which brought us to a corpus of 88 explanations.

We evaluated the new teleological extensions by running the feedback system on the corpus of 88 student multimodal explanations. Each explanation is in a single sketch file, and contains a comic graph (like in Figure 1) of one of more subsketches, and a set of sentences entered through the sketch’s Tell window.

Accuracy was measured by checking each sentence with SEA and against the kind of feedback given (negative or positive) and whether it was true or not (valid or invalid). Thus, for any sentence there are for possible types of feedback: a valid positive, a valid negative, an invalid positive, and an invalid negative. The same was done for feedback about transitions between states using STV.

Overall Results

Across the 88 sketches, 52 contained no invalid feedback and at least one positive feedback, 32 had at least one piece of invalid feedback, and five had zero of both (in four, the student didn’t include text/sentences and in the fifth an error halted the feedback mechanism). There were a total of 374 choices of feedback (positive or negative) made by the Design Coach. The breakdown is shown in Table 2.

Table 2: Overall Feedback Validity and Error Rates

	Total	Valid	Invalid	Error %
Both	**347**	**291**	**56**	**16.1%**
STV	68	52	16	23.5%
SEA	279	239	40	14.3%
SEA positive	163	158	5	3.1%
SEA negative	116	81	35	30.1%

Table 3: Invalid Feedback by Source

	Count	% of Errors
Teleological representations	16	28.6%
STV	16	28.6%
Various bugs	14	25%
Quantities not implemented	6	10.7%
Visual processing	4	7.1%

[4] This clause states that the object named “Velcro” in the sketch is a member of the collection “VelcroTheFastener”

The 56 invalid pieces of feedback came from four categories of issues, enumerated in Table 3. The first source of errors was limitations in our teleological representations. Some of these were due to students using the teleological vocabulary in ways we did not expect, either in their sketches or in sentences. For the teleology rules to work, students' sentences had to mention the correct number of subsketches in the correct order, and/or their subsketches would have to use the same object in both subsketches without redrawing it from scratch. However, we did not provide enough information in the feedback for students to realize these requirements. In some situations, these explanations did include enough information, so we counted them as false negatives by the Design Coach.

In other cases, students' sentences expressed causal relationships where functions affected behavior (e.g. saying attachment causes the cup holder to not move). But because functions are inferred from behaviors and not the other way around, the justification structure never reflects this causal pattern. We plan to increase flexibility by adding more rules for teleological reasoning.

The second source of invalid responses was produced by a non-teleological issue within State Transition Verification, which was counting rotations about points not at the center of the object as translations.

The third category, various bugs, included: a UI bug where rotational force arrows were reified with invalid values, a reasoning bug where stale facts were being believed, and a representational bug where some subsketches believed facts they should not.

The fourth source was the fact that our vocabulary included quantities (e.g. the amount of force) but these representations are not currently handled in our coach (though we plan to in the future).

The last type of invalid feedback, visual processing, was cases where CogSketch could not properly determine the location or direction of surface contacts. However, we have new techniques for segmenting glyphs into edges which should allow us to handle these cases in the future.

Some Examples

The sketch in Figure 5 demonstrates the use of the attachment and detachment. The student explained the removable cup holder using six subsketches and three sentences. Two of the three sentences involved the new teleological representations:

- In Side View and Side View 2 Cup Holder detaches from Arm via deinterlocking
- In Velcro 2 Arm attaches to Chair via temporary adhesive material

To verify the first sentence, the detachment, the Design Coach checks the first state mentioned to see if the cup holder and arm are attached via interlocking. Due to the limitations of the visual spatial representations, touching or overlapping of objects is sufficient to satisfy attachment via interlocking, so this check succeeds. Next, it checks to see if the objects are no longer attached in the second state. The system views the first sentence as successful.

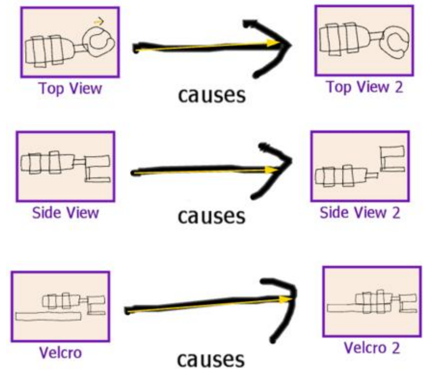

Figure 5: Attachment and detachment

In the second sentence, the student indicates that the arm attaches to the chair via a temporary adhesive material. Our rule for this case is: if the two objects mentioned are touching the same adhesive material, then they are attached. The velcro strips (which the student resized in the second state, "Velcro 2") were labeled with the collection "VelcroTheFastener" by the student, so Design Coach can infer they are attached.

Figure 6: Adjusting to different sized cups

Figure 6 is a student's refinement of the cup holder design in which they made the jaws of the cup holder able to adjust to hold different size cups. In this example, Design Coach affirmed the student's sentence:

- In State 1 and State 2 cup holder attacher adapts to change in size of cupholder base via being adjustable

Adjusting to the size via being adjustable is seen as true when the object adapting is moveable. It must also attach to the target object somehow (as evidence of its interaction with the object).

Related Work

Yaner and Goel [2006] proposed a five level Drawing-Shape-Structure-Behavior-Function model for linking drawings in a diagram to teleological function. In Design Coach, CogSketch's representations are like their drawing, shape, and structure levels, the comic graph and qualitative mechanics are similar to their behavior level, and the structured language input is similar to their function level.

Wang and Kim [2007] created an ontology for supporting form-function reasoning. Their representation conflated shape and generic functions of objects. For example, gravity is assumed so one function of containers is to limit downward motion of the continued substance. Given that CogSketch produces spatial representations, we can opt to focus on the spatial functions of mechanisms, like as containment. Using the function/way model allows us to also tackle non-spatial functions, like improving comfort, by finding ways of achieving that function that fit within CogSketch's reasoning abilities.

Conclusions and Future Work

There are a number of extensions that we plan to make. First, the rules which check teleological information are sometimes generous (e.g. using "touches" as criterion for interlocking) and sometimes overly strict (e.g. requiring explicit mention of states in order in the sentence) in what they allow or do not allow. We can use more of the spatial representations in CogSketch and additional reasoning to make the existing teleology better. Second, this particular design exercise only requires a small vocabulary of functions and ways. We plan to expand it by working with instructors and by collecting more student sketches involving novel designs, thereby ensuring that coverage matches what is relevant for our target population. Designs in EDC include devices for improving the lives of people with disabilities, such as stroke victims who have lost the use of a limb. Many of these are mechanisms and their functions will involve the manipulation of objects and comfort. We expect that, over time, the number of new functions and ways needed for each new design exercise will drop, but at the moment, we are in the steep part of the knowledge capture curve. Additionally, to accommodate more complex designs our teleology will need to include a functional hierarchy, where the functions have sub-functions. In function-way ontology, ways can be accomplished via functions, so our teleology can be extended hierarchically. Also, using our metalayer, we will be able to allow students to depict sub-functions in subsketches and use arrows to relate them. Third, we also plan to incorporate Qualitative Process Theory [Forbus 1984] which will allow us to support more teleological functions involving quantities (e.g. energy, work, speed). Finally, instructors find that students often jump right into describing the behaviors of the design without telling them what the overall goal is. In the future, Design Coach could check for the absence of teleological sentences, and for behaviors or states which aren't linked to any sort purpose, and ask the student to provide one.

Acknowledgements

This research was supported by the Spatial Intelligence and Learning Center, NSF SLC Grant SBE-0541957. We thank Bruce Ankenman, John Anderson, Stacy Benjamin, Dan Brown, and David Sabo for their feedback and help, and for allowing us to work with their EDC students.

References

[Forbus 1984] Forbus, K. (1984). Qualitative process theory. *Artificial Intelligence*, 24, 85-168.

[Forbus et al 1999] Forbus, K.D., Whalley, P., Everett, J., Ureel, L., Brokowski, M., Baher, J. and Kuehne, S. (1999). CyclePad: An articulate virtual laboratory for engineering thermodynamics. *Artificial Intelligence*, 114, 297-347.

[Forbus et al 2011] Forbus, K., Usher, J., Lovett, A., Lockwood, K., & Wetzel, J. (2011). CogSketch: Sketch understanding for Cognitive Science Research and for Education. *Topics in Cognitive Science*. 3(4), pp 648-666.

[Kim 1993] Kim, H. (1993). Qualitative reasoning about fluids and mechanics. Ph.D. dissertation and ILS Technical Report, Northwestern University. Evanston, IL.

[Kitamura et al 2006] Kitamura, Y., Koji, Y., & Mizoguchi, R. (2006). An ontological model of device function: industrial deployment and lessons learned. *Applied Ontology* , vol. 1, no. 3-4, pp. 237-262

[Nielsen 1988] Nielsen, P.E. (1988). A qualitative approach to rigid body mechanics. (Tech. Rep. No. UIUCDCS-R-88-1469; UILU-ENG-88-1775). Urbana, Illinois: University of Illinois at Urbana-Champaign, Department of Computer Science.

[Wang & Kim 2007] Wang, E., Kim, Y., (2007). Form-Function Reasoning for Product Shape Ontology. *Proceedings of the 6th International Semantic Web Conference*, Busan, South Korea

[Wetzel & Forbus 2009] Wetzel, J. and Forbus, K. (2009). Automated Critique of Sketched Mechanisms. Proceedings of the 21st Innovative Applications of Artificial Intelligence Conference. Pasadena, CA.

[Wetzel & Forbus 2010] Wetzel, J. and Forbus, K. (2010). Design Buddy: Providing Feedback for Sketched Multi-Modal Causal Explanations. In *Proceedings of the 24th International Workshop on Qualitative Reasoning*. Portland, Oregon.

[Yaner & Goel 2006] Yaner, P., Goel, A., (2006). From Diagrams to Models by Analogical Transfer, *Proceedings of the 4th International Conference on Diagrammatic Representation and Inference*, pp. 55-69