

Do Monkeys Like Elephants or Do Elephants Watch Monkeys? An Empirical Study into the Default Reading of Constraint Diagrams

Andrew Fish
University of Brighton
Andrew.Fish@brighton.ac.uk

Judith Masthoff
University of Aberdeen
jmasthoff@csd.abdn.ac.uk

Abstract

Constraint diagrams complement the Unified Modeling Language that is used in the development of software systems. They generalize Venn and Euler diagrams, and include facilities for quantification and navigation of relations. Difficulties such as ordering quantifiers cause constraint diagrams to have more than one possible meaning. Fish et al. [1, 2] have extended the constraint diagram notation, augmenting it with a reading tree which provides a unique semantics. However, this involves extracting extra information about the reading order from the diagram creator. Progress has been made towards a default reading, which extracts this information implicitly from the diagram [4]. In this paper, we investigate how humans read ambiguous constraint diagrams intuitively, without any instructions on the reading order. We test several principles which could be used in a default reading. We conclude that although some of these principles have significant effects, they also may conflict. Either an explicit reading tree or a taught method of reading a diagram which may conflict with some people's intuition is necessary.

1. Introduction

The Unified Modeling Language (UML) [9] is the Object Management Group's industrial standard for software and system modeling. This has accelerated the uptake in the software industry of diagrammatic notations for designing systems.

Constraint diagrams were introduced in [8] for use in conjunction with UML for object-oriented modelling. Constraint diagrams provide a diagrammatic notation for expressing logical constraints, such as invariants and operation preconditions and postconditions. In UML constraints are expressed using the Object Constraint Language

(OCL) [12], which is essentially a textual, stylised form of first order predicate logic.

Constraint diagrams build on a long history of using diagrams to visualize logical or set-theoretical assertions. They generalize Venn diagrams and Euler circles which are rich research topics, particularly as the basis of visual formalisms and diagrammatic reasoning systems [5, 6, 10].

Constraint diagrams are vastly more expressive than these systems because they can express relations and both universal and existential quantification explicitly, whilst still, according to Fish et al. (2004) retaining the elegance, simplicity and intuitiveness of the underlying diagrammatic systems.

There are, however, some problems with defining the semantics of constraint diagrams: the diagrams can be ambiguous (see Section 2). Fish et al. [1, 2] have extended the constraint diagram notation, augmenting it with a reading tree. However, whilst this results in a unique semantics, it does require some extra notation and effort from the creator. In this paper, we investigate how humans tend to read ambiguous constraint diagrams without such a reading tree.

2. The constraint diagram notation

We informally explain the syntax and semantics of constraint diagrams in an example based manner. The formal notation can be found in [1, 2].

In Figure 1, there are three given contours labeled Teachers, Courses and Students which represent disjoint sets. The dot is an existential spider, which represents existential quantification (there is a student). Its habitat is "inside Students". The asterisk is a universal spider, which represents universal quantification (for all teachers). Its habitat is "inside Teachers". The two arrows labeled "teaches", and "attends" represent relations. The arrow labeled "teaches" is sourced at the universal spider and targets an unlabelled derived contour inside Courses. This

derived contour represents the image of the relation “teaches” (the courses which are taught).

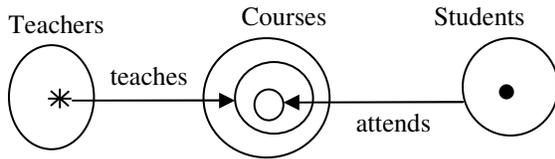


Fig. 1. Example of an ambiguous constraint diagram: the order of the spiders matters

The semantics of this diagram depends upon the order in which it is read. In Figure 1, if we start reading at the universal spider we obtain:

- (1) “Each teacher, t , teaches some courses, and there is a student who attends only courses taught by t .”

Whereas, if we start reading at the existential spider we get:

- (2) “There is a student who attends some courses, and each teacher teaches all of these courses.”

To give another example, the diagram in Figure 2 has two sensible semantics, depending upon whether one starts reading at the spider in Monkey or the spider in Elephant.

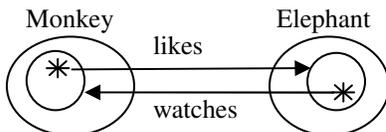


Fig. 2. Another example of an ambiguous constraint diagram: the domain of spiders

If we start reading at the spider in Monkey, we assume the diagram looked like Figure 3 before the second spider was drawn.

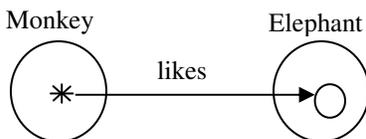


Fig. 3. Starting at the spider in Monkey

Informally, we read:

“Each monkey, m , likes some elephants...”

Then, one reads the other spider:

“...and each of these elephants watches some monkeys including m .”

A spider represents a quantifier, and the domain of a spider represents the set over which quantification takes place. With this reading, the domain of the universal spider in Monkey is the whole of Monkey. The domain of the universal spider in Elephant is the derived contour in Elephant. Notice that in contrast the habitats of both spiders are inside the derived contours, as can be seen in Figure 2.

On the other hand, assume that one starts reading at the spider in Elephant. Then in the same manner one obtains different semantics:

“Each elephant, e , watches some monkeys and each of these monkeys likes some elephants including e .”

Here the domain of the first spider read was the whole of Elephant, whilst the domain of the second spider was inside the derived contour in Monkey.

Fish et al [1,2] have proposed to annotate constraint diagrams with so-called reading trees. Basically this means giving spiders names (like m and e) and explicitly indicating in which order to read them (e.g. m before e).

3. Experiment

3.1. Experimental design

3.1.1. Research questions. The suggested default reading in [4] is an algorithm for determining the order in which the diagram is read using properties of the diagram (which still requires some user input in certain symmetric cases). This algorithm is based on a number of principles. We will test some of these principles, together with some other obvious possibilities, firstly to decide if they have an effect, and secondly to test their prioritization.

We will investigate the effect of:

- *Spider Type.* This principle hypothesizes that subjects have a preference for reading existential spiders before universal ones. Always starting with the existential spider was proposed by Stapleton et al. [11] as a way to reduce ambiguity. In Figure 4 below, this would mean starting at the spider in Monkey.
- *Following Chains of Arrows.* A Chain of Arrows is a sequence of arrows where the source of an arrow is either equal to the target of the preceding arrow or is dependent upon it. For example in Figure 4, the pair (watches, likes) is a chain of arrows since the source of “likes” is equal to the target of “watches”, and the

pair (likes, watches) is a chain of arrows since the source of “watches” is inside Elephant, but outside the derived contour which is the target of “likes”. A chain of arrows gives an ordering of the spiders. For example, in Figure 4, the chain of arrows (watches, likes) orders the spider in Elephant before the one in Monkey, whereas the chain (likes, watches) orders them the other way around. This principle hypothesizes that if there is only one maximal length chain of arrows then this is followed. In Figure 4, no prediction can be made on the basis of Following Chains of Arrows, as there are two conflicting chains. In such cases, the difference between “weakly” and “strongly” bound may be used (see below).

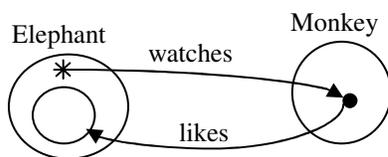


Fig. 4. Example diagram (used in Question 4 of Condition B).

- *Strongly Versus Weakly Bound.* A pair of arrows is *strongly bound* together if the target of the first is the source of the next. A pair is *weakly bound* if the source of the next is not equal to but dependent upon the target of the first. This principle hypothesizes that subjects prefer to follow a strongly bound chain of arrows to a weakly bound chain of arrows. For instance, in Figure 4, the pair of arrows (watches, likes) is strongly bound together but the pair (likes, watches) is weakly bound. So, the prediction is that subjects would start with the spider in Elephant.
- *Domain Equals Habitat.* This principle hypothesizes that subjects would prefer to start with a spider whose domain is equal to its habitat. For instance, in Figure 4, the domain of the existential spider is Monkey, which is equal to its habitat, for either reading order. However, if one starts reading at the universal spider then its domain is the whole of Elephant, which is not equal to its habitat. So, the prediction is that subjects would start with the spider in Monkey.
- *Outside a Derived Contour.* This principle hypothesizes that subjects prefer to start at a spider which is not inside a derived contour.
- *The Positioning in the Plane of the Spiders.* In most western languages, people read from top to bottom

and from left to right. This principle asserts that subjects read spiders in the same manner, starting with the top-left one. For instance, in Figure 4, this would mean starting reading at the spider in Elephant.

3.1.2. Method. A mixture of a within and a between subjects design was used. A within subjects design was used to investigate the principles: Spider Type, Following Chains of Arrows, Strongly Versus Weakly Bound, Domain Equals Habitat, and Outside a Derived Contour. A between subjects design was used to investigate the effect of the positioning in the plane of spiders. Subjects were assigned randomly to one (out of four) experimental condition. The four conditions contained the same diagrams, except for layout changes:

- The diagrams in Condition B are the diagrams of A, horizontally flipped.
- The diagrams in Condition C are the diagrams of B, vertically flipped.
- The diagrams in Condition D are the diagrams of A, vertically flipped.

Subjects were given a questionnaire to fill out, which started with an example-based explanation of the constraint diagram notation (see Appendix A). It used only diagrams with one spider, so that there were no ambiguities or suggestions made here. Subjects were allowed to comment on the notation. Next, a sequence of constraint diagrams was shown. For each diagram, two possible meanings were given, and subjects were asked to select the meaning that they thought the diagram represented most accurately, or to give an alternative meaning.

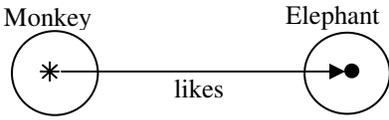
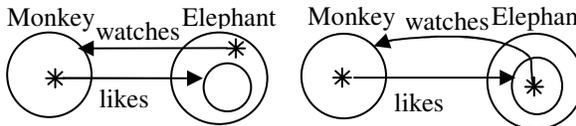
3.1.3. Subjects. Forty subjects participated voluntarily in the experiment. They were undergraduate students of the School of Computing, Mathematical and Information Sciences at the University of Brighton, studying various courses (B.Sc. in Computer Science, Software Engineering, Mathematics, or Mathematics for Finance). The experiment took place in a lecture room. Subjects were assigned randomly to an experimental condition (10 subjects per condition):

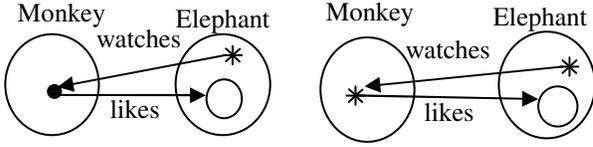
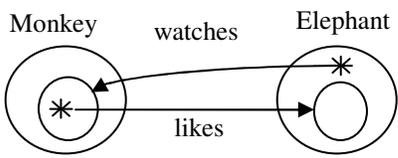
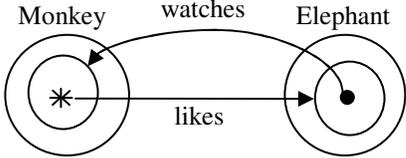
- Condition A: 9 male, 1 female. Average age was 21 with standard deviation 1.1.
- Condition B: 10 male. Average age was 21.8 with standard deviation 2.1.
- Condition C: 9 male, 1 female. Average age was 24.8 with standard deviation 6.
- Condition D: 8 male, 2 female. Average age was 22.1 with standard deviation 3.2.

The spread over courses was similar for all conditions.

3.1.4. Materials. Constraint diagrams which are ambiguous (without a reading tree) have at least two spiders, at least one of which is universal. In order to keep our diagrams as simple as possible, we chose examples which had exactly two spiders, at least one of which was universal. All existential spiders had only one foot, but one universal spider had two feet. All of our diagrams had exactly two given contours, “Monkey” and “Elephant”. These labels were chosen so as not to influence subjects based on the semantics of the labels. For instance, had we chosen “Teacher” and “Course”, then subjects might have started with “Teacher” because they prefer starting with an actor rather than an object. Whether this is true needs investigation, but we did not want this to influence our results. Similarly, we did not want to use abstract labels like “A”, “B”, as subjects might have used alphabetical ordering, and it might have been too difficult for our subjects to understand the semantics. We used “likes” and “watches” as labels for the relations, again taking care that the verbs have no significant semantic relationship (unlike “likes” and “hates”). The diagrams used, and our reasons for using them, are given in Table 1.

Table 1. Diagrams used and reasons

<p>Question 1</p>  <p>Question 1 tests:</p> <ul style="list-style-type: none"> • Following Chains of Arrows: There is only one chain of arrows (likes). So, this principle predicts we start in Monkey. • Spider Type: The two spiders are of different type. This principle predicts we start in Elephant.
<p>Question 2 Question 3</p>  <p>Questions 2 and 3 individually test:</p> <ul style="list-style-type: none"> • Following Chains of Arrows: In both, there is only one maximal length chain of arrows, namely (likes, watches). • Domain Equals Habitat: This chain of arrows also has this property. <p>Both principles predict starting in Monkey.</p> <p>Comparison between the two tests:</p> <ul style="list-style-type: none"> • Outside a Derived Contour: Predicts that more people start in Elephant in Question 2 than in 3.

<p>Question 4 Question 5</p>  <p>Questions 4 and 5 individually test:</p> <ul style="list-style-type: none"> • Domain Equals Habitat • Strongly Versus Weakly Bound <p>In both, (likes,watches) is a weakly bound chain of arrows which has domain equal to habitat, and (watches,likes) is a strongly bound chain of arrows with domain not equal to habitat. So, the two principles give conflicting predictions: Domain Equals Habitat predicts starting in Monkey, while Strongly vs Weakly Bound predicts starting in Elephant.</p> <p>Comparison between the two tests:</p> <ul style="list-style-type: none"> • Spider Type: Predicts that more people start in Monkey in Question 4 than in Question 5.
<p>Question 6</p>  <p>Question 6 individually tests:</p> <ul style="list-style-type: none"> • Outside a Derived Contour: Predicts starting in Elephant. <p>Comparison with Questions 5 and 9 tests:</p> <ul style="list-style-type: none"> • Outside a Derived Contour: The only difference between Questions 5 and 6 is that the spider in Monkey in Question 6 is inside a derived contour. The only difference between Questions 6 and 9 is that the spider in Elephant in Question 6 is outside a derived contour. So, this property predicts that more people start in Monkey in Questions 5 and 9 than in Question 6.
<p>Question 7</p>  <p>Question 7 individually tests:</p> <ul style="list-style-type: none"> • Spider Type: Predicts starting in Elephant. <p>Comparison with Question 9 tests:</p> <ul style="list-style-type: none"> • Spider Type: Predicts that more people start in Monkey in Question 9 than in Question 7.

Question 8

Question 8 individually tests:

- Following Chains of Arrows: There is only one maximal length chain of arrows, namely (watches, likes).
- Domain Equals Habitat: This chain of arrows also has this property.

Both principles predict starting in Elephant.

Comparison with Question 9 tests:

- Domain Equals Habitat: Predicts that more people start in Elephant in Question 8 than in Question 9.

Question 9

Question 9 has no other distinguishing properties, so this question just tests the Positioning in the Plane of Spiders (as will all other diagrams). Comparisons have been discussed above.

As indicated in the reasons given for including the diagrams, we are not only interested in the subjects' answers to the individual questions. We also want to compare their answers between questions. For instance, the only difference between Questions 4 and 5 is that there is an existential spider in Question 4 instead of a universal one. A difference in results between Questions 4 and 5 would therefore indicate an effect of the Spider Type principle.

Table 2 also shows a difficulty in designing this experiment: sometimes it will be hard to distinguish between the effects of principles. For instance, Domain Equals Habitat can not be tested independently from Following Chains of Arrows nor from Strongly vs. Weakly Bound. Unfortunately, this could not have been avoided, given the constraint of diagrams with at most two spiders and arrows which we imposed for simplicity. To test Domain Equals Habitat, we need to have a spider whose domain can be different depending on the order of reading. This requires the spider to be dependent upon a derived contour. This implies we have a weakly bound chain of arrows. If there is no other maximal length chain of arrows, then Following Chains of Arrows is also being tested. If there is a second, strongly bound chain of arrows, then Strongly Bound vs. Weakly Bound is also being tested. Having a second weakly bound chain of arrows would result in no spider with Domain Equals Habitat, so is not an option.

3.1.5. Predictions. Table 2 summarizes the predictions of the principles discussed above with respect to the preferred meanings. "E" and "M" represent starting with the spider in Elephant, and starting with the spider in Monkey, respectively. Shading indicates that the principle is not applicable. The bottom four rows show the predictions of the Positioning in the Plane of Spiders, separated into the predictions of Left to Right and of Top to Bottom, for the various experimental conditions.

Table 2. Predictions of the principles.

	Question								
	1	2	3	4	5	6	7	8	9
Following Chains of Arrows	M	M	M						E
Domain Equals Habitat		M	M	M	M				E
Strongly vs Weakly bound				E	E				
Spider Type	E			M				E	
Outside a Derived contour			M			E			
Left to Right A, D	M	M	M	M	M	M	M	M	M
Left to Right B, C	E	E	E	E	E	E	E	E	E
Top to Bottom A, B		E		E	E	E		E	M
Top to Bottom C, D		M		M	M	M		M	E

3.2. Results

The results of the experiment are given in Figures 5 and 6. We have use a Chi-Square Test for testing the significance of results within a question, and a Wilcoxon Signed Ranks Test to compare between questions.

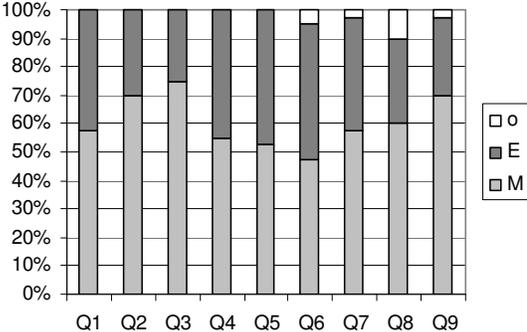


Fig. 5. Results of experiment with all conditions combined. E=starting in Elephant, M=starting in Monkey, o=No answer given.

Figure 5 shows the results of the experiment with all four conditions combined, so that layout effects can be ignored. In the following discussion, we will ignore the results of Question 8, as there is a lot of evidence that subjects had difficulty understanding multi-footed spiders (by comments in the introduction of notation part of the questionnaire, the number of people who did not answer this question, and the number of people who made up their own, incorrect, answers).

3.2.1. Spider Type. Questions 1, 4, and 7 had one existential and one universal spider. However, there is no statistically significant effect for any of these questions. Indeed, in Questions 1 and 7 slightly (but not significantly) more subjects chose the universal spider, but in Question 4 slightly (but not significantly) more subjects chose the existential one. So, there is not even a trend to support this principle. Spider Type is also the only difference between Questions 4 and 5: Question 4 has an existential spider where Question 5 has a universal one. Again, there is no significant difference in the subjects' replies between these two questions. Similarly, there is no statistically significant difference between Questions 7 and 9. Therefore, we conclude that there does not seem to be a preference for Spider Type.

3.2.2. Following Chains of Arrows. This principle applies to Questions 1, 2, and 3. For all three, it predicts that most subjects would start with Monkey. This is indeed the case. However, for Question 1, the difference is not statistically significant. For Questions 2 and 3 it is significant ($p < 0.03$), but this may also be due to the Domain Equals Habitat property. We believe this to be more likely, as otherwise we would have expected a significant difference for Question 1 (after all, to this question only this principle and the Spider Type principle apply, and we had already concluded that an effect of Spider Type is unlikely). Therefore, we do not have conclusive evidence to support the Following Chains of Arrows principle.

3.2.3. Strongly Versus Weakly Bound. Question 4 and 5 test this property, but also the Domain Equals Habitat property. As shown in Table 2, the predictions of these properties conflict. In Questions 4 and 5, there is no significant preference of subjects for Monkey or Elephant. This could mean that both principles do not have an effect, or that they cancel each other out.

3.2.4. Domain Equals Habitat. Questions 2 and 3 tested this property. For both Question 2 and 3, the property predicted that subjects would start with

Monkey, which indeed happened in a statistically significant way ($p < 0.05$). So, there is some evidence that subjects indeed prefer to start with the spider whose domain is equal to its habitat. The results for Questions 4 and 5 may therefore show a canceling preference for Strongly Versus Weakly Bound. However, in contrast with the suggested preferences in the Fish default algorithm [4], if both properties apply it is not the case that the Domain Equals Habitat property wins.

3.2.5. Outside a Derived Contour. This is the only principal which can be applied to Question 6, and there is no significant difference in the preference of subjects for Monkey or Elephant. The only difference between Questions 2 and 3 is the position of a spider: inside a derived contour in Question 3 and outside that derived contour in Question 2. Similar differences exist between Questions 5 and 6, and between Questions 6 and 9. In all three cases, fewer subjects started with this spider when it was inside the derived contour. The difference between Questions 6 and 9 is statistically significant ($p < 0.05$), the difference between the others is not. It should also be noted that in Question 9 significantly more subjects chose to start with Monkey ($p < 0.01$). As the diagram is completely symmetric, the only plausible explanation seems that subjects noticed the fact that both the diagram and the answers were symmetric, and therefore often chose the first answer (which happened to be Monkey for this question)¹. This could then also explain the difference with Question 6. So, there is no conclusive evidence to support that subjects are hesitant to start with a spider inside a derived contour.

3.2.6. Positioning in the plane. The results in Figure 6 lead to the following conclusions on the effect of the Positioning in the Plane of Spiders:

¹ This is a flaw in our design: we should have randomized the order of answers over subjects. Most of the comparisons we have done can, however, not have been affected by this. The exception is the comparison between Questions 4 and 5 which had different answer orders (Elephant was the first answer in Question 4, while it was Monkey in Question 5). This may have prevented us from finding an effect. Given the way subjects answered questions (many provided their own variations of answers), we do not believe strong order effects to have happened in questions other than the symmetrical Question 9 (and perhaps the misunderstood Question 8), but we cannot prove this.

- *Left to right.* There is a statistically significant difference ($p < 0.05$) between the conditions which have Monkey on the left and the conditions which have Monkey on the right. Subjects do indeed seem to prefer to start with the leftmost spider (as we expected on the basis that our western subjects are used to reading from left to right).
- *Top to bottom.* There is no statistically significant difference between conditions with the spider in Monkey on the top and conditions with the spider in Monkey on the bottom. So, there is no evidence to support the hypothesis that the vertical layout of spiders matters.

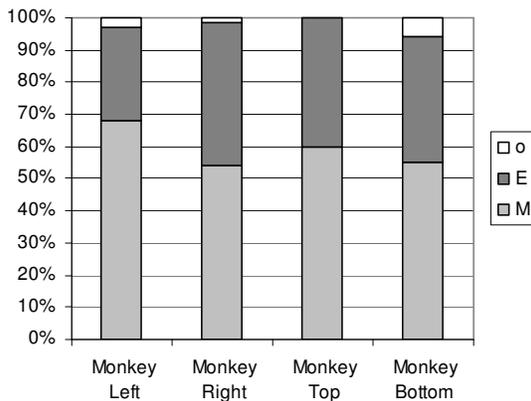


Fig. 6. Results of experiment for all questions combined, dependant on the positioning in the plane of spiders: Monkey Left (Conditions A and D), Monkey Right (Conditions B and C), Monkey Top (Questions 2,4,5,6,8 of A and B; Question 9 of C and D) and Monkey Bottom (Questions 2,4,5,6,8 of C and D; Question 9 of A and B). E=starting in Elephant, M=starting in Monkey, o=No answer given.

It should be noted that having the spider in Monkey on the left always coincided with having the contour Monkey on the left. On the other hand, the two contours were always vertically aligned. It may, therefore, well be the case that the statistically significant effect we have found for horizontal positioning is due to the relative positioning of the contours rather than the spiders. It will need to be investigated also whether an effect of vertical positioning does occur when the contours are vertically positioned.

4. Conclusions

It is vital that diagrams created as part of a system specification or proof are as easy to read as possible. Constraint diagrams, unfortunately, can be ambiguous, so can be read differently to the creator's intention. In a

computer environment, animation can be used to show in which order a diagram has been drawn, so how it should be read. On paper, this is more difficult. It is possible to annotate diagrams, for instance augmenting them with reading trees as proposed by [2]. However, this makes the syntax of a diagram more complicated. Understanding people's intuitive reading of ambiguous diagrams may aid in choosing a default reading (an implicit reading tree) and in displaying diagrams in the most readable manner. The goal of our experiment was to contribute to this understanding.

We have found that (western) people prefer to read diagrams from left to right. If a user has drawn a diagram, an editor could redraw it to better match other readers' intuitions, for instance positioning on the left the spider at which the user should start to read. A reading tree could still be provided, but the user's intuitions are more likely to match this reading tree. The readers' attention could also be drawn to points in which the diagram's meaning deviates from standard intuition. For instance, we found an effect of the Domain Equals Habitat principle, but unfortunately it is not necessarily possible to draw a diagram in such a way as to ensure that the spider to start with has this property. Highlighting unexpected spiders to start with (both in the diagram and in the reading tree) might then help. Similar techniques could be used by an automatic proof writer when generating diagrams and reading trees.

The results of our experiment show that it is very difficult to decide on a default reading. For instance, principles like Domain Equals Habitat and Strongly Versus Weakly Bound tend to cancel each other out. Even for the very simple diagrams used in the experiment, it was impossible for most of them to predict accurately what subjects would tend to do. So, the default algorithm presented in [4] does not correspond to a common intuition. Of course, it may be possible to teach people to apply some of the principles. The Domain Equals Habitat principle might be the best one to start with, as it seems to match intuitions.

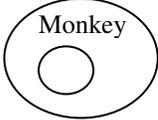
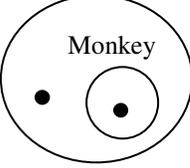
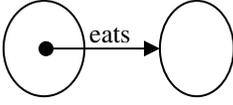
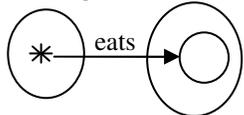
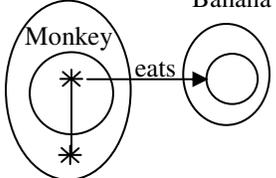
The results of our experiment clearly show the need for an augmentation of constraint diagrams such as the explicit reading trees proposed by [2]. Alternatively, some fixed stated default reading (implicit reading tree) would do, but any choice of this will conflict with some users' intuitive ideas. Maybe users of a diagram editor should be allowed to specify which principles they would like to use for a default reading. The layout of diagrams, highlighting, and reading trees could then be altered accordingly.

The experiment also shows how difficult evaluation can be. We went through six pilot studies (each with two or more subjects) before running this experiment, trying to optimize the experimental material, e.g. the explanation about the notation. Still, we experienced problems with subjects not understanding the two-footed spider. A major problem has also been that the complexity of the diagrams makes it hard to isolate single features. Maybe simpler notations, like Venn/Euler and Spider diagrams (something between Euler and Constraint diagrams, with existential but no universal quantifiers), should be extensively tested first. We are currently undertaking another study, in which we investigate the intuitiveness of the notations used in Spider diagrams, and the ease with which subjects can reason with these diagrams.

5. References

- [1] Fish, A., Flower, J., and Howse, J. (2003). A reading algorithm for constraint diagrams. *IEEE Symposium on Human Centric Computing Languages and Environments*, Auckland, New Zealand, pp161-168.
- [2] Fish, A., and Flower, J., and Howse, J. (in press). *The semantics of augmented constraint diagrams*. To appear in *Journal of Visual Languages and Computing*.
- [3] Fish, A., and Howse, J. (2003). Computing reading trees for constraint diagrams. *Proceedings of AGTIVE 03*, Charlottesville, Virginia, pp260-274.
- [4] Fish, A., and Howse, J. (2004). Towards a default reading for constraint diagrams. *Proceedings of Diagrams 04*, Cambridge, UK, Berlin: Springer Verlag, pp51-65.
- [5] Hammer, H (1995). *Logic and visual information*. CSLI Publications.
- [6] Howse, J. Molina, F., Taylor, J., Kent, S. and Gil, J. (2001). Spider diagrams: A diagrammatic reasoning system. *JVLC*, 12, pp 299-324.
- [7] Howse, J. and Schuman, S. (In press). Precise visual modelling: A case study. To appear in *SoSyM Journal*.
- [8] Kent, S. (1997). Constraint diagrams: Visualising invariants in object oriented models. *Proceedings of OOPSLA 97, ACM SIGPLAN Notices 32 Issue 10*, pp 327-341.
- [9] OMG. *UML Specification, Version 1.4*. Available from www.omg.org.
- [10] Shin, S.-J. (1994). *The logical status of diagrams*. Cambridge University Press.
- [11] Stapleton, G., Howse, J. and Taylor, J. (2003). *A constraint diagram reasoning system*. Proceedings of the 9th International Conference on Distributed Multimedia Systems, International Conference on Visual Languages and Computing, pp263-270.
- [12] Warmer and Kleppe (1998). *The Object Constraint Language*. Addison-Wesley.

Appendix A: Introduction to the notation

	<p>The diagram means: "All monkeys are animals".</p>
	<p>The dot inside "Monkey" means that "There is a (at least one) monkey". The dot outside "Monkey" means that "There is an animal that is not a monkey".</p> <p>So, the diagram means "All monkeys are animals; there is at least one monkey, and one animal that is not a monkey."</p>
	<p>Again, the dot inside "Monkey" means that "there is a monkey". The arrow from that monkey to "Banana" indicates that this monkey eats all bananas (and only bananas).</p> <p>The diagram means: "There is a monkey which eats all bananas (and only bananas)."</p>
	<p>The asterisk inside "Monkey" means "each monkey". The circle inside "Banana" indicates "some (but not necessarily all) bananas". The diagram means: "Each monkey eats only (some) bananas".</p>
	<p>The two asterisks connected by a line mean "each monkey and each animal that is not a monkey"</p> <p>The diagram means: "All monkeys are animals. Each animal (whether a monkey or not) eats only bananas".</p>