

Picturing Problems: Solving Logic Puzzles Diagrammatically

John Howse, Gem Stapleton, Jim Burton, and Andrew Blake

Centre for Secure, Intelligent and Usable Systems, University of Brighton, UK
{John.Howse, g.e.stapleton, j.burton, a.l.blake}@brighton.ac.uk

Abstract. Solving logic puzzles is a popular recreational activity. The solution of a logic puzzle involves understanding and reasoning about the information provided. Diagrammatic logics have been shown to help people to understand and reason about logical information. In this paper we use spider diagrams, a visual logic based on Euler diagrams, to visualize logic problems. Furthermore, we reason with the diagrammatic representation both syntactically and semantically until we reach a solution. We present four example logic puzzles of varying difficulty and produce their solutions in detail using spider diagrams. We suggest that the use of diagrams is helpful in their solution.

1 Introduction

Logic puzzles have a long history. Lewis Carroll was an eminent protagonist, see [4], for example, with logical game-playing also occurring in *Alice in Wonderland* and *Through the Looking Glass*. Logic puzzles are usually expressed in the form of a list of textual statements from which some information can be deduced. The list of statements is frequently called the *premise* (or *premises*) and the information that can be deduced is called the *conclusion*. The solution involves logical reasoning to show that the conclusion follows from the premise. Frequently in logical puzzles, the conclusion is not stated explicitly, but given in the form of a question (or questions). The solver has to answer the question by reasoning about the premise. The simplest of these puzzles are of the form of classical syllogisms. For example, consider the following statements:

1. All Tigers are Cats
2. All Cats are Mammals
3. No Cats are Dogs

Are all Tigers Mammals? Are all Tigers Dogs?

As all Tigers are Cats and all Cats are Mammals, we can deduce that all Tigers are Mammals. Similarly, as all Tigers are Cats and no Cats are Dogs, we can deduce that no Tigers are Dogs.

We can visualise this puzzle using Euler diagrams [5]. In figure 1, the first two statements of the puzzle are presented in an Euler diagram. From this diagram,

we can *observe* that the circle labelled Tiger is within the circle labelled Mammal and hence that all Tigers are Mammals. In figure 2, the first and third statements of the puzzle are presented in an Euler diagram. From this diagram, we can observe that the circle labelled Tiger is disjoint from the circle labelled Dog and hence that no Tigers are Dogs.




Fig. 1. Statements 1, 2




Fig. 2. Statements 1, 3




Fig. 3. Statements 1, 2, 3

In figure 3, all three statements are visualized. From this diagram, we can observe that all Tigers are Mammals and that no Tigers are Dogs. The circles labelled Mammal and Dog intersect because we are not given any information about the relationship between mammals and dogs; some dogs *could* be mammals and some dogs *could not* be mammals. The diagrams we are using do not have existential import; some of the regions in the diagram could represent the empty set. Hence, we *cannot* deduce that all dogs are mammals, even though we *know* from biological taxonomy that dogs are mammals. This is a lesson we can learn when considering logical puzzles and indeed logical statements generally: *we should not make assumptions*. We can only use the information provided – and, of course, anything we can deduce from that information.

There is a widespread belief that the use of diagrammatic logics, such as Euler diagrams, is helpful in allowing people to understand and reason about logical information. One reason often cited for this is that some diagrammatic notations are *well-matched*. A notation is *well-matched to meaning* if its semantic relationships are matched by its syntactic relationships [7]. For example, in figure 1, the semantic relationship ‘All Tigers are Cats’ (statement 1) is matched by the curve labelled Cat containing the curve Tiger in the diagrammatic representation. Similarly, in figure 2, the semantic relationship ‘No Cats are Dogs’ (statement 3) is matched by the curves labelled Cat and Dog being disjoint in the diagrammatic representation. Diagram well-matchedness can help people to *understand* information. Other properties of diagrams can help people to *reason* about information.

Diagrams can also contain *free rides* when compared to the same information represented textually (or symbolically) [11]. Free rides occur when the topological properties of a diagram force some logical consequence to be made explicit, where that consequence would otherwise take some inferential steps. For example, as discussed above, from statement 1, all Tigers are Cats, and statement 2, all Cats are Mammals, we can deduce that all Tigers are Mammals. However, in

figure 1, we can observe directly that all Tigers are Mammals. As such, ‘all Tigers are Mammals’ is an example of a free ride exhibited by the diagram in comparison to the textual representation. A similar free-ride, ‘No Tigers are Dogs’, can be observed in figure 2 compared to the textual representation, where this information has to be deduced. The notion of free rides has been formalized and extended to the theory of *observational advantages* [13]. Thus free rides and observational advantages within diagrams can help people to reason about information. These helpful properties of some diagrams have been known for a long time: the notions of well-matchedness and free-rides are closely related to Peirce’s concept of *iconicity* [10].

Many diagrammatic logics have been developed over the centuries that represent sets and their relationships. Examples include Euler diagrams, *Venn diagrams* [14] and Peirce’s *existential graphs* [10]. Venn diagrams are often confused with Euler diagrams. However, in a Venn diagram all the intersections between sets are represented; shading is used to represent the empty set. *Spider diagrams* [6] are based on Euler diagrams with the addition of syntax to represent individuals. Spider diagrams also use shading, borrowed from Venn diagrams, but to represent an upper bound on the cardinality of sets rather than just emptiness. They have well-defined reasoning rules [8] and are a fully formalized diagrammatic logic [9]. The iconicity and semiotics of spider diagrams are discussed at length in [3]. We will use spider diagrams as the vehicle to visualize and solve logical puzzles.

The mathematician Raymond Smullyan was a very prominent exponent of the art and science of recreational logic puzzle production. He is well-known for puzzles involving knights and knaves; knights always tell the truth and knaves always lie. Some of the puzzles considered in this paper are based on knights and knaves. In particular, they are adapted from Smullyan’s *The Gödelian Puzzle Book* [12], in which he uses puzzles to describe and prove Gödel’s incompleteness theorem. Smullyan died in 2017 at the age of 97. This paper is dedicated to his memory. It begins in the strange land of knights and knaves.

2 Smullyan Island

Smullyan Island is a very peculiar place. All the inhabitants are either Knights or Knaves but not both. Knights always tell the truth but Knaves always lie. The inhabitants of Smullyan Island are a playful lot and like to pose puzzles. However, you can’t always believe what they say – unless they are Knights, of course. Unfortunately, you can’t tell just by looking who is a Knight and who is a Knave. The first puzzle involves inhabitants of Smullyan Island. It is adapted from [12].

2.1 The Prize Puzzle


Three inhabitants of Smullyan Island are named Alice, Bob and Carol. At least one of them is a Knight and at least one is a Knave. Exactly one of them has a prize. They made the following statements:

- **Alice:** Bob doesn't have the prize
- **Bob:** I don't have the prize
- **Carol:** I have the prize

Who has the prize?

Solution


We start the solution by rephrasing the puzzle's context in diagrammatic form:




The diagram, d_1 , states that there are exactly three individuals in the universe that we are considering: Alice, Bob and Carol, which we abbreviate to a, b and c, respectively. The shading tells us that there are no other elements except those represented by the dots. Diagram, d_2 , states that there is at least one Knight and at least one Knave. The shading indicates that all individuals are Knights or Knaves. The diagram, d_3 , states that there is exactly one Prize holder. These statements are invariants: they are assumed to be true.

Having set up the puzzle's context, we can now consider the inhabitants' statements. For each of the inhabitants there are two cases. Either they are a Knight and therefore tell the truth or they are a Knave and therefore lie. We start by considering Alice's statement in the two cases.

Alice is a Knight and therefore tells the truth and so Bob doesn't have the prize:




Alice is a Knave and therefore is lying and so Bob does have the prize:




We now consider Bob's statement in the two cases.

Bob is a Knight and therefore tells the truth and so he doesn't have the prize:




Bob is a Knave and therefore is lying and so he does have the prize:




There are four possible combinations of Alice's and Bob's statements.


Alice and Bob are both Knights so we combine d_4 and d_6 :




Alice and Bob are both Knaves so we combine d_5 and d_7 :



Alice is a Knight and Bob is a Knave so we combine d_4 and d_7 :



Alice is a Knave and Bob is a Knight so we combine d_5 and d_6 :




d_{11}

In the last two cases, d_{10} and d_{11} , we can observe the free ride: **Bob** has the prize and simultaneously doesn't have the prize. Hence, d_{10} and d_{11} are contradictory and we discard them. We are left with two cases: Alice and Bob are both Knights, visualized in d_8 ; and Alice and Bob are both Knaves, visualized in d_9 .


We now consider both cases for Carol's statement.

Carol is a Knight and therefore tells the truth and so does have the prize:



d_{12}


Carol is a Knave and therefore is lying and so doesn't have the prize:



d_{13}


We now combine d_8 and d_9 , derived from Alice's and Bob's statements, with d_{12} and d_{13} , derived from Carol's statement. There are four possibilities.

Alice, Bob and Carol are all Knights so we combine d_8 and d_{12} :




d_{14}

Alice, Bob and Carol are all **Knaves** so we combine d_9 and d_{13} :




d_{15}

Alice and Bob are **Knights**, Carol is a **Knave** so we combine d_8 and d_{13} :



d_{16}


Alice and Bob are **Knaves**, Carol is a **Knight** so we combine d_9 and d_{12} :



d_{17}

In the first case, d_{14} , all three inhabitants are **Knights** and the second case, d_{15} , all three are **Knaves**. However, by d_2 , there must be at least one **Knight** and at least one **Knave**. So these two cases are contradictory and we discard them. We are left with two possibilities: Alice and Bob are **Knights** and Carol is a **Knave**, d_{16} ; or Alice and Bob are **Knaves**, and Carol is a **Knight**, d_{17} .

In d_{17} , both Bob and Carol have the prize, which contradicts d_3 as there must be a unique prize holder. So we can discard this case. Hence, d_{16} is the only valid case. We therefore arrive at the final step in the solution:




d_{18}

As neither Bob nor Carol has the prize, then, by d_1 and d_3 , Alice must have the prize. Furthermore, we have determined that Alice and Bob are **Knights** and Carol is a **Knave**.

The solution was undertaken in a systematic way. We set up the context of the puzzle and then considered each statement in the order it was presented. We considered all of the cases, but ruled out those that were contradictory. In some cases, we could observe the contradictions as free rides; in other cases,

we could derive contradictions from the information observable in the diagrams. The solution was presented diagrammatically in d_{18} with two boxes: one giving information about Knights and Knaves; the other about the prize holder. We can express this information in one box showing the relation between the prize holder and the Knights and Knaves directly:



It is gratifying to know that the prize holder is a Knight. The next puzzle is also adapted from [12] and again features three inhabitants of Smullyan Island. This time one of them is a magician and the statements they make are a little bit more complex, involving a disjunction and a negated conjunction. The diagrams representing these statements are therefore a bit more complicated than those in the Prize Puzzle.

2.2 The Magician Puzzle


Three inhabitants of Smullyan Island are named Ali, Beth and Chris. Exactly one of them is a Magician. They made the following statements:

- **Ali:** Beth is not both a Knave and a Magician
- **Beth:** Either Ali is a Knave or I am not a Magician
- **Chris:** The Magician is a Knave

Who is the Magician?

Solution

Again, we start by setting up the premises diagrammatically:




The diagram d_1 states that there are exactly three individuals in the universe that we are considering: Ali, Beth and Chris, which we again abbreviate to a, b and c, respectively. The diagram d_2 states that every element is either a Knight or a Knave, but not both, and d_3 states that there is exactly one Magician.

The solution to the Prize Puzzle was undertaken in a very systematic way. We will be more strategic in the solution to this puzzle. Chris's statement is the simplest, so we will consider it first. There are two possibilities: Chris is a Knight


or Chris is a Knave. We consider Chris's statement in the two cases.

Chris is a Knight and therefore tells the truth and so the Magician is a Knave:




d_4

Chris is a Knave and therefore is lying and so the Magician is a Knight:



d_5


We observe from d_4 that if Chris is a Knight then he is not the Magician. We also observe from d_5 that if Chris is a Knave then he is not the Magician. So we have deduced that Chris is not the Magician:



d_6

Next, we consider Ali's statement. This time, we consider the cases separately. We consider the case in which Ali is a Knave first, as the negation of Ali's statement is simpler.


Assume that **Ali is a Knave**. He is therefore lying and hence Beth is both a Knave and a Magician:



d_7


We observe from d_7 that Beth is a Knave. We now consider Beth's statement. Beth's statement is true because we are assuming that Ali is a Knave. Hence, Beth is a Knight.

We deduce:




Beth is both a Knave and a Knight which contradicts d_2 . This contradiction arose from assuming that Ali was a Knave. Hence Ali is not a Knave and therefore must be a Knight by d_2 . Now we again consider Ali's statement but this time we know that he is a Knight.

Ali is a Knight and therefore tells the truth. Therefore Beth is not both a Knave and a Magician:




In d_9 , a 3-footed spider is used to convey the information that Beth is therefore a Knave but not a Magician, a Magician but not a Knave, or neither a Knave nor a Magician (that is, Beth is not both a Knave and a Magician).

We now consider Beth's statement. We begin by assuming that Beth is a Knave. The negation of Beth's statement is: "Ali is not a Knave and I am a Magician." From this we deduce that Ali is a Knight and Beth is a Magician:




We visualize the case Ali is a Knight and Beth is a Knave by combining d_9 and d_{10} :




It can be clearly observed that d_{11} is contradictory. Beth is not both a Knave and a Magician and, at the same time, is both a Knave and a Magician. Thus our assumption that Beth is a Knave is false. Therefore Beth a Knight. Now we


consider Beth's statement knowing that she is a Knight. As Ali is not a Knave, we deduce from Beth's statement that Beth is not a Magician:




As we have already deduced that Chris is not a Magician (d_6), the information we have obtained so far is:



Hence, as Beth and Chris are not Magicians, by d_3 , Ali is the Magician:



Finally, Chris stated that the Magician is a Knave. However, Ali is not a Knave. Hence Chris is lying and is therefore a Knave. The complete solution is:



This puzzle was a little bit harder to solve than the first one. Rather than systematically working through the puzzle's statements, we considered the simplest statement first and then strategically determined the order in which the statements were considered.

3 Metapuzzles

A *metapuzzle* is a type of puzzle for which the solution requires knowing that others could or could not solve it. The following example is also adapted from [12] and again considers inhabitants of Smullyan Island. This puzzle requires the involvement of a friendly logician.

3.1 The Medic Puzzle

A logician, *Ray*, meets three inhabitants of Smullyan Island, Al, Bal and Cal. At least two of them are **Knives**. Also, exactly one of them is a **Medic**. Al and Bal made the following statements:

- **Al**: I am the **Medic**
- **Bal**: I am not the **Medic**


Ray then asks Cal, “Is Al really the **Medic**?”

Cal answered and then Ray knew which one was the **Medic**. (We do not know whether Cal answered *yes* or *no*.)

Who is the **Medic**?


Solution

We again visualize the premises:




The diagram, d_1 , states that there are exactly three individuals in the universe that we are considering: Al, Bal and Cal, which we abbreviate to a, b and c, respectively. The diagram, d_2 , states that there are at least two **Knives** and d_3 that there is exactly one **Medic**.

Assume that Al is a **Knight**:




Al is telling the truth and so Al is the **Medic**. As Al is a **Knight**, Bal and Cal must be **Knives** by d_2 . Hence we have:




Bal is lying and so is also a **Medic**. This contradicts d_3 .

Hence, our assumption that Al is a Knight is false, so Al is a Knave:




Al is lying and so is not the Medic. There are two cases to consider: Bal is a Knight or Bal is a Knave. Assume that Bal is a Knight:




Bal is telling the truth and so is not the Medic. We already know that Al is a Knave and not the Medic, so Cal is a Knave and is the Medic.

Now we assume that Bal is a Knave:



Bal is lying and so is the Medic. We cannot determine whether Cal is a Knight or a Knave.

We need to distinguish between the two cases, d_4 and d_5 . If Cal is a Knave, then either of the two cases can occur. So, in order to distinguish between d_4 and d_5 , Ray would have to deduce that Cal is a Knight. In which case, Cal must have answered (truthfully) *no* to Ray's question. Hence, Bal is the Medic. The full solution is:



The next puzzle takes us away from Smullyan Island. The *Hat Puzzle* is another metapuzzle. In its original form, this puzzle dates back to 1940 and probably before. The puzzle presented here is a nice variation on the hat problem in which one of the participants is colourblind. The puzzle is adapted from Alex Bellos's Puzzle column [1] in the UK's Guardian newspaper. The puzzle was originated by the American puzzle enthusiast Jack Lance [2].

3.2 The Hat Puzzle

A box contains two red hats and three green hats. Azalea, Barnaby and Caleb close their eyes, take a hat from the box and put it on. When they open their eyes they can see each other's hats but not their own. They do not know which hats are left in the box.


We can assume that all the protagonists are perfect logicians who tell the truth. They know all the information in the above paragraph. In addition, one of them is colourblind. They all know who the colourblind person is. They made the following statements in order:

- Azalea: I dont know the colour of my hat.
- Barnaby: I dont know the colour of my hat.
- Caleb: I dont know the colour of my hat.
- Azalea: I dont know the colour of my hat.

Who is the colourblind person, and what colour is their hat?


Solution

We start the solution to the puzzle by setting up the premises diagrammatically:




The diagram, d_1 , states that there are exactly three individuals in the universe that we are considering: Azalea, Barnaby and Caleb, which we abbreviate to **a**, **b** and **c**, respectively. The diagram, d_2 , states that every individual is wearing a Red hat or a Green hat, but not both; at most two individuals are wearing Red hats. Diagram d_3 states that there is exactly one ColourBlind individual.

Assume that the ColourBlind individual is wearing a Red hat. There are two cases to consider: another individual is wearing a Red hat or no other individual is wearing a Red hat. The first case is visualized as follows:




In this case, individual z can see that the other two individuals are wearing Red hats. Individual z would therefore know the colour of their own hat which contradicts one of the first three statements, so this case does not occur.


The second case is visualized here:




In this case, assume, without loss of generality, that y 's name occurs before z 's name alphabetically. Then z can deduce that their hat is green as y doesn't know the colour of their own hat. This again contradicts one of the first three statements. Hence this case does not occur. Therefore, the ColourBlind individual is wearing a Green hat:




So, after the first three statements, all three individuals know that the ColourBlind individual is wearing a Green hat. In the fourth statement, Azalea does not know the colour of her hat, hence Azalea is not ColourBlind:



Finally, if Caleb is ColourBlind, then, by the argument above, after the first two statements he would be able to deduce that he is wearing a Green hat. However, this contradicts his statement. So Caleb is not ColourBlind:



Hence, Barnaby is ColourBlind and is wearing a Green hat:



We cannot deduce anything about the colour of the hats worn by Azalea and Caleb. The solution of this puzzle is interesting as it involved the use of variables, x , y , z , to deduce that the ColourBlind individual was wearing a Green hat.

4 Conclusion

We have used spider diagrams to visualize and reason about logical puzzles of varying difficulty. We suggest that the use of diagrams is helpful in their solution. The four examples presented above may offer some evidence for this. Our intention is not to take the fun out of puzzle solving. The diagrams provide a useful tool for visualizing and organizing the solution, but the inherent difficulty in the reasoning involved remains. This paper is based on a talk presented at the British Science Festival in September, 2017. We finish by leaving the reader with another puzzle to solve, again adapted from [12]:

A crime has been committed on Smullyan Island. Three suspects made the following statements:

- Alf: I am guilty.
- Bof: I am the same type [Knight or Knave] as one of the others.
- Chaf: We are all of the same type.

Who is guilty? [Only read the footnote if you want to know who did it¹.]

References

1. <https://www.theguardian.com/science/2018/jan/29/can-you-solve-it-the-puzzle-of-the-red-and-green-hats>, acc. Feb 2018
2. <https://jacoblance.wordpress.com/>, acc. Feb 2018
3. Burton, J., Howse, J.: The semiotics of spider diagrams. *Logica Universalis* 11(2), 177–204 (2017)
4. Carroll, L.: *The Game of Logic*. MacMillan (1886)
5. Euler, L.: *Lettres à une princesse d’Allemagne sur divers sujets de physique et de philosophie*. Letters vol 2, no. 102-108. Birkhauser. (1761)
6. Gil, J., Howse, J., Kent, S.: Formalising spider diagrams. *Proc. Visual Languages (VL99)*. pp. 130–137. IEEE (1999)
7. Gurr, C.: Effective diagrammatic communication: Syntactic, semantic and pragmatic issues. *Journal of Visual Languages and Computing* 10(4), 317–342 (1999)
8. Howse, J., Molina, F., Taylor, J., Kent, S., Gil, J.: Spider diagrams: A diagrammatic reasoning system. *Journal of Visual Languages and Computing* 12(3), 299–324 (2001)
9. Howse, J., Stapleton, G., Taylor, J.: Spider diagrams. *LMS Journal of Computation and Mathematics* 8, 145–194 (2005)
10. Peirce, C.: *Collected Papers*, vol. 4. Harvard University Press (1933)
11. Shimojima, A.: *Semantic Properties of Diagrams and Their Cognitive Potentials*. CSLI Publications, Stanford (2015)
12. Smullyan, R.: *The Gödellian Puzzle Book*. Dover (2014)
13. Stapleton, G., Jamnik, M., Shimojima, A.: What makes an effective representation of information: A formal account of observational advantages. *Journal of Logic, Language and Information* 26(2), 143–177 (2017)
14. Venn, J.: I. on the diagrammatic and mechanical representation of propositions and reasonings. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 10(59), 1–18 (1880)

¹ Alf is guilty