

## REASONING OF NON- AND PRE-LINGUISTIC CREATURES: HOW MUCH DO THE EXPERIMENTS TELL US?<sup>1</sup>

**Abstract.** *If a conclusion was reached that creatures without a language capability exhibit some form of a capability for logic, this would shed a new light on the relationship between logic, language, and thought. Recent experimental attempts to test whether some animals, as well as pre-linguistic human infants, are capable of exclusionary reasoning are taken to support exactly that conclusion. The paper discusses the analyses and conclusions of two such studies: Call's (2004) two cups task, and Mody and Carey's (2016) four cups task. My paper exposes hidden assumptions within these analyses, which enable the authors to settle on the explanation which assigns logical capabilities to the participants of the studies, as opposed to the explanations which do not. The paper then demonstrates that the competing explanations of the experimental results are theoretically underdeveloped, rendering them unclear in their predictions concerning the behavior of cognitive subjects, and thus difficult to distinguish by use of experiments. Additionally, it is questioned whether the explanations are rivals at all, i.e. whether they compete to explain the cognitive processes of the same level. The contribution of the paper is conceptual. Its aim is to clear up the concepts involved in these analyses, in order to avoid oversimplified or premature conclusions about the cognitive abilities of pre- and non-linguistic creatures. It is also meant to show that the theoretical space surrounding the issues involved might be much more diverse and unknown than many of these studies imply.*

**Keywords:** *cognitive processes, deduction, probabilistic reasoning, animal cognition, infant cognition*

### Introduction

Theories in cognitive science and philosophy of mind strive to fit the results of experimental psychology. However, the results themselves typically do not hand us conclusive answers, so they have to be analyzed, and the analysis is in turn subject to theoretical assessment. In this paper I examine the assumptions behind the analysis of certain experimental results concerning the possible reasoning mechanisms of non- and pre-linguistic creatures.

---

1 The work on this paper has been supported by the Ministry of Education, Science and Technological Development of the Republic of Serbia through the project Logico-epistemological bases of science and metaphysics (No. 179067). I am also very grateful to Jacob Beck for reading the drafts of the paper and making useful comments about it.

Certain behavior is consistent with a capability for reasoning from an excluded alternative: the recognition that if there are only two possibilities and it is not one of them, it must be the other. Recently, there have been attempts to experimentally test whether some animals, as well as pre-linguistic human infants, are capable of this kind of reasoning. I focus on the analysis of two studies which test the reasoning abilities of these creatures by setting up a task that needs to be solved. Call's (2004) two cups task tested whether great apes are capable of reasoning by exclusion, and since the apes proved to be successful at the task, the other study – Mody and Carey's (2016) four cups task – proposed a way to determine the mechanism responsible for that kind of reasoning.<sup>2</sup> I will demonstrate that there is a mistaken assumption in the analysis of the latter study, and that we cannot decide among the competing explanations based on the strategy proposed by the authors. I will further show that the two main competing explanations cannot be clearly distinguished from each other, because their assumptions, requirements, and implications are not sufficiently defined. As a consequence, the experiments conducted so far, as well as future attempts to decide between them using experimental research, might be misguided.

## Reasoning by exclusion and possible underlying mechanisms

In Call's task, the subjects see the experimenter hiding the reward in one of two cups, not knowing which one. They then receive evidence that one cup is empty. If they reason by exclusion, they should use the information about the empty cup to exclude that location, and instead select the other cup. The subjects proved successful at this task (the rate of correct choices was significantly above chance), and their success can be interpreted and explained by several accounts. The term "reasoning by exclusion" does not make commitments as to the particular reasoning mechanism being used (Mody & Carey, 2016).

The results of this task were taken by many commenters to indicate that the animals manifest a capability for reasoning by the disjunctive syllogism. The subjects supposedly reasoned: The food is either in A or B. It is not in A. Therefore it is in B. This interpretation imposes the highest cognitive requirements: it requires the subjects to be representing the concepts OR and NOT as well as the dependent relationship between A and B (embodied

---

2 The reader will notice that the two studies are done with different subjects (the latter study was done with human children). Moreover, the older children that were successful in the latter experiment clearly do not fall into the group of pre-linguistic creatures. Although this would make a difference for the plausibility of the thesis that is being tested, I, however, do not take this difference to be significant for my analysis, since I am interested in the methodology of these experiments, rather than in the results themselves. I focus on the difficulty of interpreting the behavioral data, regardless of the particular subjects and results of the study.

by the concept OR). The dependent relationship is what leads the subjects to update their representation of B when they see that A is empty, and to conclude that B necessarily contains the reward. Updating of B based on the information about A is referred to in the literature as 'inferential updating (Mody & Carey, 2016).

Another interpretation is named “Maybe A, maybe B”. The subjects initially believe that the reward might be in A, and also that it might be in B. The two possible hiding locations are regarded independently: gaining the information about A does not lead to inferential updating which would form a new appraisal of B. Thus, when they see that A is empty, they do not search in A anymore, and they proceed to B according to their initial premise that it *might* contain the reward (Mody & Carey, 2016).

The results can also be explained with even fewer cognitive demands. According to the “Avoid empty” interpretation, the animals have no particular beliefs about whether B contains the reward. They are merely not searching in A. When they see that A is empty, they avoid searching in it, and instead approach B merely because it is the other salient hiding location available. This does not even require representing the alternatives A and B as potential locations of the reward, so there is also no inferential updating (Mody & Carey, 2016).

Rescorla (2009) proposed another possible mechanism underlying exclusionary reasoning: probabilistic inference over the space of cognitive maps. Rescorla defines cognitive maps as mental representations that represent geometric features of the physical environment. Their most important feature is not having logical form. This allows them to be realized without the subject’s capability for logic. The key part of Rescorla’s analysis is the Bayesian decision theory, which gives calculations for the distribution of probabilities over cognitive maps. Probabilities assigned to cognitive maps can be understood as degrees of belief assigned to different hypotheses. The subjects recognize two possible locations of the hidden reward, represented by two cognitive maps ( $M_1$ ,  $M_2$ ). The maps correspond to two hypotheses concerning which cup has the reward in it. The subjects initially lack evidence regarding where the reward is hidden, so the initial probability distribution treats cups the same:

$$p(M_1) = p(M_2) = 0.5$$

Since they exhaust the space of possibilities, their probabilities sum up to 1.

Also, for each cup there are two possible outcomes:  $y_i$  – the cup has food in it, and  $n_i$  – the cup is empty, and their probabilities also sum up to 1:  $p(y_i) + p(n_i) = 1$ .

Assuming  $M_1$  is true, the probability that the subjects will recognize that the reward is in the cup is taken to be  $p(y_i|M_1) = 0.8$ , since the account allows the small possibility of the subject not seeing the reward even if it is looking in the right cup.

Since  $p(y_i)+p(n_i)=1$ , it follows that the chance of false negatives is  $p(n_i|M_i)=0.2$

There is also a slight chance of false positives, say,  $p(y_i|M_j)=0.1$ , which makes the chance of a correct observation that the cup is empty  $p(n_i|M_j)=0.9$ .

When the subject is presented with the evidence that the first cup is empty, the probabilities are updated over the cognitive maps using Bayes' Law:

$$\begin{aligned} p(M_1|n_1) &= 0.182 \\ p(M_2|n_1) &= 0.818 \end{aligned}$$

The initial probability of 0.5 is lowered for  $M_1$  to 0.182, while  $M_2$  is updated to 0.818, by the process of inferential updating (since the two cups are represented as being in a disjunctive relationship). Thus, the subject reaches a conclusion that it is more probable that the reward is in cup B than that it is in A.

All four interpretations can explain the experimental results, because using any of these approaches would lead subjects to be successful in the task. Subsequent experiments were designed to distinguish between these interpretations. I focus on Mody and Carey's (2016) study, which tested for behavioral evidence of inferential updating. This should then distinguish between the deductive and probabilistic interpretations on one side, and the remaining two interpretations which predict no inferential updating on the other side.

## Distinguishing between the interpretations

The following experiment was conceived as an extended version of Call's task. Two rewards (in this study the rewards were stickers) were hidden in four cups, one reward in each of two pairs. The first pair of cups was covered by a screen so that the subject could not see which of the cups the reward was placed in, and then the same was done with the second pair. The participants were children from 2.5 to 5 years old, divided into four groups by their age. When one of the cups was revealed to be empty, the child was supposed to pick the other cup from the pair, which is certain to contain a reward.

If children were using the deductive syllogism, they would engage in inferential updating, meaning that the information about the empty cup (not-A), in combination with the representation of where the sticker was hidden (A or B), would lead them to conclude that the other cup from the pair necessarily contained the sticker (therefore B). This interpretation predicts children will choose the "target cup" (B).

If they were using the "Maybe A, maybe B" strategy, obtaining the information "not-A" would not lead to updating information about B. The children would still hold on to "Maybe B", and all the remaining cups would seem to be equally good candidates. Thus, the children would choose the target cup at an equal rate as the other two cups.

According to the “Avoid empty” strategy, learning “not-A” would lead to avoiding A, but without representations about other possible locations. Thus, the subjects would also be expected to choose all three remaining cups at an equal rate (Mody & Carey, 2016).

Mody and Carey seem to take the probabilistic account to predict that the children will pick the target cup preferentially to the other cups (due to inferential updating). Still, they seem to take probabilistic reasoning as somewhat “less certain” than deductive reasoning. Thus, the probabilistic interpretation might predict preferential choosing of the target cup, but at a somewhat lower rate than in the deductive scenario. I will address this in more detail in the following section.

Prior to the main task, there was a training phase, which involved only three cups: one pair of cups and one single cup, hidden behind two screens. The procedure of hiding the rewards was the same: the first reward was placed in the single cup, and the second reward was placed in one of the cups from a pair. After removing the screen, the child was asked to choose a cup. This task did not require reasoning by exclusion, but it still required comparing the sure cup to the two uncertain cups.

## Results and analysis

In the training trials chance was established at 33%, since there were three cups to choose among. The children chose the target cup at rates significantly above chance. In the test trials, since the children virtually never chose the empty cup, it was taken that three cups were the remaining options, and chance was also set at 33%. The results were very similar to the training phase. All groups except for the youngest chose the target cup significantly above chance, suggesting that they engaged in inferential updating. The youngest children chose the target cup in only 36% of the cases, not significantly above chance. They behaved in a manner predicted by the “Maybe A, maybe B”, and “Avoid empty” interpretations. The rates of choosing the correct cup in both phases of the experiment are given in the table.

Training trials		Test trials	
Age group	Success rate	Age group	Success rate
2.5	47%	2.5	36%
3	60%	3	58%
4	71%	4	64%
5	72%	5	76%

These results indicate that, except for 2.5-year-olds, the children chose the target cup preferentially, and thus behaved in a manner consistent with inferential updating. This allows accepting only the probabilistic and deductive interpretations as possible, while dismissing the others.

We can sum up the main issue as follows.

In the training trials, an observation that cup A is full leads to representing cup A as sure and cups B and C as unsure, which leads to reaching for cup A.

In the test trials, an observation that cup A is empty leads to representing cup A as sure (empty), cups C and D as unsure, and leads to updating the representation of cup B from unsure to sure (full), and reaching for it.

The two possible interpretations of the results aim to answer the question: What is the cognitive process that leads to updating of cup B from “unsure” to “sure (full)”?

According to the probabilistic interpretation, the subjects engage in Bayesian redistribution of *coupled* probabilities (synchronously lowering the probability of cup A and rising the probability of cup B), while representing cups C and D as *independent* probabilities and as remaining unsure.

According to the deductive interpretation, the subjects engage in a logical inference: “The reward is either in A or B. It is not in A. Therefore it is in B”. Cups C and D are not included in the inference, since they are represented as independent from A and B.

Mody and Carey go further to claim that the deductive interpretation is more plausible. Their strategy was to additionally formulate the difference between these two cognitive processes in terms of *certainty*: deductive reasoning would lead to a choice based on *absolute* certainty that the reward is in cup B, while the probabilistic one would lead to a choice based on *increased* certainty (the subjects are only *more* certain that the reward is in B than that it is in any of the other cups). They then propose a way of distinguishing between these two mechanisms, claiming that one feature of the gathered data indicates absolute certainty behind the children’s choice. Namely, the children chose the target cup just as often in test trials as in training trials – in which they could directly observe (and thus be absolutely certain about) where the sticker was hidden. These results suggest that children were absolutely certain in the test trials, too. In other words, since the rate of choosing the correct cup was the same in the trials which required reasoning as in the trials which did not, their reasoning was interpreted as absolutely certain, and therefore, deductive:

Our design did not allow us to distinguish between a choice based on absolute certainty and one based on increased certainty. The latter would still require that children represented the dependent relationship between the two locations, and that they inferentially updated their assessment (...); however, the inference children made would not be truly deductive. This possibility was put forth by Rescorla (2009), who described it in a Bayesian framework, where the probability associated with one possibility is adjusted up as the probability of another possibility goes down. However, one feature of our data suggests that children were making a deductive inference: 3- to 5-year-old children chose the target cup just as often in test trials as they did in training trials, in which they could directly observe that a sticker was being hidden there (Mody & Carey, 2016, p. 46).

I will proceed to criticize Mody and Carey's characterization of the two reasoning mechanisms by different degrees of certainty. I will first demonstrate that this characterization is false, and then I will show that probabilistic updating of coupled probabilities cannot be distinguished from explicit inferences (for now), and that these two cognitive processes may even be only variants of each other, instead of being independent strategies.

### Analysis of the analysis

Mody and Carey make a mistake of confounding two possible applications of the property of certainty. One is the certainty that defines deduction, and applies to the *transition* from the premises to the conclusion: in a valid deductive argument, the truth of the premises guarantees the truth of the conclusion. Thus, the subject can be certain that if the premises are true, the conclusion must also be true. But this does not tell us how certain the subject was in either of the premises, nor of the conclusion. The latter type of certainty applies to the propositions themselves, and it can have various degrees. The rules of "probability preservation", or "uncertainty propagation" are defined within propositional probability logics. The main idea is that the premises of a valid argument can be uncertain, in which case the conclusion will also be uncertain (Demey et al., 2017). Therefore, the deductive account does not necessarily predict children will be absolutely certain that the reward is in cup B. It needs to additionally postulate that they are absolutely certain of the premises. This type of certainty is what is of interest for the experiment, because the degree of certainty about the final conclusion is what affects the subjects' behavior, and thus the percentages that Mody and Carey appeal to. Let us see how this type of certainty is accommodated within the two accounts.

In the probabilistic account, the distribution of degrees of certainty is determined mathematically, according to the Bayes' Law. We saw in Rescorla's account that it allows the possibility of subjects making both false negative, and false positive judgements, due to their fallibility. This is reflected in not assigning absolute certainty to even the seemingly obvious observations, such as "A is empty." Thus, the "premise" ("Reward is in A") gains a probability slightly higher than 0. This, in turn, renders the probability of the conclusion "Reward is in B" as slightly less than 1. The probabilities assigned to cups C and D should be equal, and the same as in the initial distribution (0.5). Therefore, if the subjects reasoned probabilistically, they would indeed make a choice based on "increased certainty" of B over other options, just as Mody and Carey suggest.

In the deductive account, however, degrees of certainty have not been mentioned. The experimenters expect absolute certainty by default, and the percentages of failure of subjects to perform the task (which were at least 30% of choices, as we saw in the table) are explained by appealing to "noise"

or “performance issues”, such as limitations of attention, working memory, or other factors. (Mody & Carey, 2016, pp. 46–7). So far, it seems that the probabilistic account has a better formal apparatus for dealing with the degrees of certainty. Still, there are probably several ways in which they could be incorporated in the deductive account as well. One way would be to assign probabilities to the propositions, like it is done in the probabilistic semantics. We can have a probability function for the propositional language  $L$ , and the valuations  $v: L \rightarrow \{0,1\}$  of classical propositional logic can be replaced with probability functions  $P: L \rightarrow \mathbb{R}$ , which take values in the real unit interval  $[0,1]$ . The classical truth values of true (1) and false (0) can thus be regarded as the endpoints of the unit interval  $[0,1]$ . This would mean taking classical logic as a special case of probability logic, or equivalently, taking probability logic as an extension of classical logic (Demey et al., 2017). Applying this to the cups task, we can formally express the subject’s deductive reasoning as follows:

$$\begin{aligned} P(A \vee B) &= 1 \\ P(\neg A) &= 0.9 \\ \text{Therefore, } P(B) &= 0.9^3 \end{aligned}$$

In cognitive terms, the probabilities could be defined within a meta-cognitive level, without being explicitly represented by the subject. Thus, even though the subjects reason through a logical inference, each step in the inference (e.g.  $\neg A$ ) may be accompanied by a degree of subject’s certainty about the step. However, it is difficult to specify how the probabilities of two or more premises are to be combined. Some advocates of this kind of extension of classical logic propose a rule that “a p-valid inference cannot take us from low uncertainty in the premises to high uncertainty in the conclusion”. They define the uncertainty of a proposition  $p$  as one minus its probability,  $1 - P(p)$ . Then an inference with two or more premises is p-valid if and only if the uncertainty of its conclusion is not greater than the sum of the uncertainties of its premises for all coherent probability assignments (Evans et al., 2015).

This version of the deductive account needs to be further theoretically developed. Nevertheless, the outline shows a way to implement different degrees of certainty into the deductive account. In application to Mody and Carey’s results, even though it was shown that there are two separate applications of certainty, that still does not prove that they were wrong in assigning absolute certainty to subjects reasoning deductively. Indeed, we do not know how certain the subjects were of any of their propositions.

## Difference between the competing accounts

This brings me to the final point of this paper. What *is* the difference between the deductive and probabilistic accounts? How can we behaviorally

3 The numerical values of probabilities are just an example.

distinguish which of these two reasoning mechanisms the subject is using? I claim that the two accounts are not sufficiently developed, and not clear in their theoretical requirements. This renders them unclear in their predictions concerning the behavior of cognitive subjects, and thus difficult to distinguish by use of experiments. I will show this by presenting several possible candidates by which these accounts might be differentiated.

### *Degrees of certainty*

As I demonstrated, since probabilities (and thus the degrees of certainty) can be accommodated within the deductive account, it is unclear whether the accounts differ in the degrees of certainty assigned to the conclusion. Thus, we do not know whether it is possible to differentiate them empirically – whether they predict different percentages of successful task performances. It is yet to be shown that there would be a difference at all.

### *Format of mental representations*

Another way to distinguish them might be by the format of the mental representations they posit. The accounts were presented as positing different kinds of mental representations: the probabilistic reasoning was presented as defined over *cognitive maps*, while the deductive reasoning is taken to be computed over *proposition-like* mental representations, and made available by language (Bermudez, 2006). Rescorla's probabilistic account was formulated partly as a way to enable computing sophisticated reasoning over non-propositional mental representations. However, neither of these accounts is necessarily tied solely to their respective representational formats. Probabilistic reasoning may also be computed over propositions – the hypotheses which the probabilities are assigned to may as well be in the form of propositions. That is, in fact, exactly how the probability distribution over competing hypotheses is most often presented (Rescorla, 2009).

In addition, even though it is not the most popular opinion among cognitive scientists, there are some authors (e.g. the advocates of diagrammatic reasoning) who claim that logical reasoning can be defined over non-linguistic representations, and that there is no intrinsic difference between symbolic and diagrammatic systems as far as their logical status goes (Shin & Lemon, 2018). This would imply that proposition-like mental representations might not be necessary for logical reasoning. Thus, each of these two accounts could be modelled quite differently from the versions of them proposed so far, and this would certainly reflect on their behavioral implications, significant for the experimental testing.

### *Logical structure*

The most important difference between these accounts is supposed to be whether they commit to logical structure – whether they describe the

subject's reasoning as proceeding by logical rules, or by some other kind of inference. The deductive account clearly appeals to the logical structure of the deductive syllogism. The probabilistic account purportedly does not involve a logical structure, but is instead structured as a distribution of probabilities over a space of hypotheses (which in Rescorla's account have the form of cognitive maps). However, this attempt at differentiating the two accounts also has its difficulties. First, it is not clear how this difference might, if at all, be behaviorally manifested. Thus far, we have no means to experimentally test between these accounts. Second, it is an open question whether Bayesian reasoning is truly an alternative to reasoning by the disjunctive syllogism, or one way of implementing it – which might explain why they are also difficult to differentiate behaviorally. As Mody (2016) observed,

the construction of the hypothesis space [in the probabilistic account] requires that children enumerate the relevant possibilities, and the inference mechanics maintain a fundamentally disjunctive relation between them. Further, the lowering of probability associated with gaining negative information essentially implements negation. Thus, even if reasoning proceeds probabilistically, propositional representations including negation and disjunction might be required to represent the information that the probabilistic mechanism uses as input.

In other words, it is unclear whether probabilistic reasoning is, in fact, dependent on some form of logical reasoning, or on some logical concepts at least. It is at least sophisticated enough to involve the ability to distinguish between coupled and independent probabilities. I agree with Mody that the experimental results presented here do give evidence of representing negation and disjunction in some way. However, they do not necessarily imply full-blown logical inference.

### *Simplicity of explanation*

In referring to the results of similar experiments indicative of the existence of basic logical reasoning in 12- and 19-month-old human infants, Arlotti and colleagues (2018) admit that a Bayesian probabilistic model of reasoning could also explain the results: “Bayesian iterative models (...) could mimic deductive syllogism without assuming a logical inference.” However, they argue that the probabilistic explanation, although compatible with the results, has more requirements than if we only assume the infants are performing a logical inference. It requires that infants represent the space of alternatives (which is equivalent to implementing a disjunctive representation), assign ordered priors to the alternatives, and assess alternative evaluations iteratively (Arlotti 2018).

Arlotti and colleagues thus defend the deductive explanation of their own results, but differently than the previous authors – merely as a more

parsimonious explanation of the results. However, this defence is still not unquestionable. One of the reasons is that it is not clear whether the probabilistic account actually commits to the assumptions brought up by Arlotti, especially if those assumptions imply an explicit representation of assigning probabilities to hypotheses, or of assessing the alternatives. Rescorla (2009) describes this assignment of probabilities as consisting in “suitable function relations between cognitive maps and mental representations denoting numbers”. For a cognitive subject to assign probabilities really means “to enter a mental state bearing appropriate functional relations to other mental states”. Rescorla seems to think that the complex probabilistic computations can be performed at lower-level processes of cognition (and perception), and does not really clarify what exactly the probabilistic account posits as being explicitly represented by the cognitive subjects. Thus, the last possible candidate used for deciding between the deductive and the probabilistic accounts – parsimoniousness of the explanation – is also rendered unusable, due to the lack of knowledge about the exact assumptions of the accounts.

In conclusion, it is difficult to distinguish the reasoning mechanisms by their behavioral signatures when the implications of the theories that posit them have not been made clear. These accounts have not been developed in sufficient detail, and the experimental psychologists should bare this in mind in order to avoid oversimplified or premature conclusions about the cognitive abilities of pre- and non-linguistic creatures. In addition, the theoretical space surrounding these issues might be much more diverse and unknown than these studies imply.

## References

- Bermúdez, José Luis (2006). Animal reasoning and proto-logic. In S. Hurley and M. Nudds (eds.), *Rational Animals?* Oxford: Oxford University Press, 127–138.
- Call, Josep (2004). Inferences about the location of food in the great apes (Pan paniscus, Pan troglodytes, Gorilla gorilla, and Pongo pygmaeus). *Journal of Comparative Psychology*, 117–128.
- Cesana-Arlotti, N., Martín, A., Téglás, E., Vorobyova, L., Cetnarski, R., Bonatti, L. L., (2018) Precursors of logical reasoning in preverbal human infants, *Science* 359, 1263–1266.
- Demey, Lorenz, Kooi, Barteld and Sack, Joshua (2017). Logic and Probability, *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.), Retrieved from <https://plato.stanford.edu/archives/sum2017/entries/logic-probability/>
- Evans, J. St. B. T., Thompson V. A. and Over D. E. (2015). Uncertain deduction and conditional reasoning. *Front. Psychol.* 6, 398.

- Mody, Shilpa & Carey, Susan (2016). The emergence of reasoning by the disjunctive syllogism in early childhood. *Cognition*, 154, 40–48.
- Mody, Shilpa (2016). *The Developmental Origins of Logical Inference: Deduction and Domain-Generality*. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.
- Rescorla, Michael (2009). Chrysippus' dog as a case study in non-linguistic cognition. In Robert W. Lurz (ed.), *The Philosophy of Animal Minds*, Cambridge University Press, 52–71.
- Shin, Sun-Joo, Lemon, Oliver and Mumma, John (2018). Diagrams, *The Stanford Encyclopedia of Philosophy* (Summer 2018 Edition), Edward N. Zalta (ed.), Retrieved from <https://plato.stanford.edu/archives/sum2018/entries/diagrams/>.