Understanding Universal Disjunction via Normal Forms using Venn Diagrams

Garry Goodwin

January 28, 2026

Abstract

The difference in meaning between the two versions of universal disjunction is a fundamental feature of quantificational logic, yet it is rarely subjected to sustained semantic analysis. Working within a limited fragment of quantificational logic, a Venn diagram semantics is used to identify twelve distinct meanings. This yields normal forms that distinguish disjunctions of universally quantified sentences from cases in which a universal quantifier is distributed over a disjunctive matrix. The difference between the two is reflected in the differing lengths of their respective disjunctive normal forms.

In quantificational logic (QL), *universal disjunction* may take the form of a disjunction between universally quantified sentences, for example,

$$\forall x(Bx) \vee \forall x(Rx).$$

Alternatively, the universal quantifier may be placed in front of a disjunction, thereby distributing its scope across a disjunctive sentence:

$$\forall x(Bx \vee Rx).$$

These two forms of disjunction are not equivalent; the direction of entailment may be expressed in QL as Equations (1) and (2).

$$\forall x(Bx) \vee \forall x(Rx) \;\vdash\; \forall x(Bx \vee Rx), \tag{1}$$

$$\forall x(Bx \vee Rx) \;\nvdash\; \forall x(Bx) \vee \forall x(Rx). \tag{2}$$

Lemmon illustrates the failure of the reverse entailment with the familiar example of a domain consisting of the positive integers.[1] From the fact that every number is either even or odd, it does *not* follow that all numbers are even, nor that all numbers are odd. While textbook examples like this bring the difference between the two forms of universal disjunction into view, sustained attempts to explain *why* their meanings are not equivalent are comparatively rare. Much of the underlying semantic structure is usually left underexplored.

To make this structure explicit, it is helpful to introduce the concept of a *semantic cell*. This can be thought of as doing the work of a valuation cell: a single unit in a finite semantic partition, familiar from truth tables. Consider the limited set of symbols containing just two propositional variables together with negation, conjunction, and disjunction:

$$p, q, \neg, \wedge, \vee.$$

The fragment restricted to combinations of the two atomic sentences is capable of generating infinitely many syntactically distinct yet semantically redundant sentences, while expressing only four distinct meanings:

$$p \wedge q, \quad p \wedge \neg q, \quad \neg p \wedge q, \quad \neg p \wedge \neg q.$$

The distinguishing feature of a valuation cell is that its surrounding truth-table semantics, beginning with atomic sentences, is combinatorial from the ground up. Yet, within the confines of a limited fragment, each such cell is entailed only by sentences that are logically equivalent to it, aside from the trivial case of contradiction. If lone atomic sentences are excluded, the meanings represented by these four conjunctions are not further divisible. As a result, every sentence expressible in the fragment is equivalent to some combination of these four partitions.

However, there is always a background assumption that each fragment is, in principle, constructible from atomic sentences. More generally, what counts as a cell depends on the logical fragment in which it is defined. For instance, if we add one more propositional variable, the set of partitions expands to eight. A valuation cell is therefore relative to a specific setting within an open-ended system.

The term *semantic cell*, as used here, is a purely semantic notion: it concerns sets of partitions while ignoring the role of atomic sentences. Whereas valuation cells are conveniently scalable, a system based on semantic cells may lack the *algorithmic expandability* characteristic of a truth table. The notion of a semantic cell is therefore intended to be more general, encompassing cells defined within systems built upon atomic sentences, while also including those defined within systems that are not.

Deflating the role of atomic sentences and focusing solely on what counts as a partition makes a semantic cell a useful notion for studying QL. If we know a fragment's set of cells, it becomes possible to delineate QL sentences by examining when their cells coincide and when they do not.

The pursuit of a finite semantics also tightens the scope of the investigation, preventing the analysis from straying into cases where a finite stock of semantic cells is unavailable. In the present case, we are specifically focusing on the fragment able to express the two forms of universal disjunction for two predicates. Accordingly, the following stock of symbols confines our investigation:

$$\forall,\ \exists,\ B,\ R,\ x,\ a,\ \neg,\ \wedge,\ \vee,\ \rightarrow,\ (,).$$

Here B and R are unary predicate letters, a is a constant symbol, and x is the sole variable. These are accompanied by the standard propositional connectives and parentheses. Formulae in which x occurs free are not counted as sentences in this fragment.

As was the case with the simpler example of the propositional fragment, standard syntax permits infinitely many well-formed sentences. Once again, only a finite number of logically distinct meanings are expressible. Each semantic cell for this restricted QL fragment may therefore be identified with the set of sentences that are logically equivalent to one another. In richer languages, a determinate classification need not be available; here the semantic cells fall into finitely many equivalence classes.

To fully decompose the possible meanings of a QL sentence, the fragment's complete set of cells must first be identified. Once this set is made explicit, any sentence within a finite fragment can be expressed as a disjunction of cells, thereby yielding a disjunctive normal form. In finite systems in which every sentence is closed, each sentence is determinately either true or false. Establishing a sentence's disjunctive normal form is therefore sufficient for determining differences in meaning between particular QL sentences.

A system of Venn diagrams is used to make this determination. What is modelled is the relational structure presented by a diagram. It is this structure, not the diagram's individual parts, that gives a cell its meaning. A diagram is then evaluated as either true or false. Boolean computation therefore does not occur below the level of the complete diagram.

A familiar limitation now forces our hand. Two intersecting circles are able to represent only three cells in total. Allowing a diagram to also denote an element of the domain is a simple yet conservative extension, sufficient to exhaust every meaning expressible within this fragment. With this addition, the full set of cells expands to thirty-two. To organise this number in a way that makes patterns of entailment more transparent we we arrgange them into grid patterns. Given two predicates, the relatively straightforward disjunction of universally quantified sentences naturally fits into a simple 4×4 array.

As we proceed, a graphical presentation makes the differences between the two forms of disjunction easier to follow. Each Boolean value is shown as a "tile": a white (ivory) tile represents Boolean 1 (true), while a black tile represents Boolean 0 (false). By the usual truth-functional expansion 2^n, sixteen tiles admit 2^{16} distinct Boolean patterns. It is therefore natural to organise these patterns by reference to the four arrays corresponding to the sentences Ba, Ra, $\forall x(Bx)$, and $\forall x(Rx)$, illustrated in Figure 1.

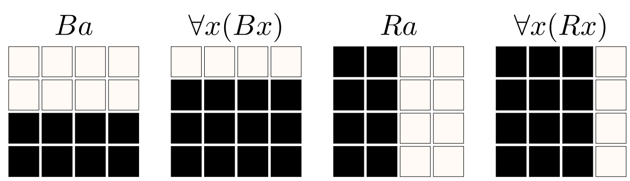

Figure 1

Distributed universal disjunction requires all thirty-two cells. These may be arranged as in Figure 2.

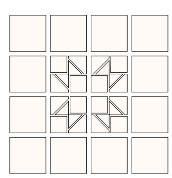

Figure 2

The four central tiles are now each replaced by a system of five tiles; there are only five because the double triangle (hourglass) represents a single proposition and is counted as a single tile. Although its chosen shape is arbitrary, as will be seen later with P11 and P12, each of these propositions represents two mutually exclusive possibilities. The triangular halves of the hourglass reflect this exclusive disjunction. The triangular cells also reveal which predicate is positive and which is negative.

From the patterns they produce, it seems natural to refer to expanded arrays as *mosaics*. A mosaic's thirty-two tiles give rise to 2^{32} distinct Boolean patterns, and hence to a number approaching 4.3 billion possible combinations of meanings expressible within this fragment. This now includes both forms of universal disjunction. Depending on which disjuncts are negated, the mosaic patterns are organised so that they *point* to the appropriate corner of the array, as shown in Figure 3.

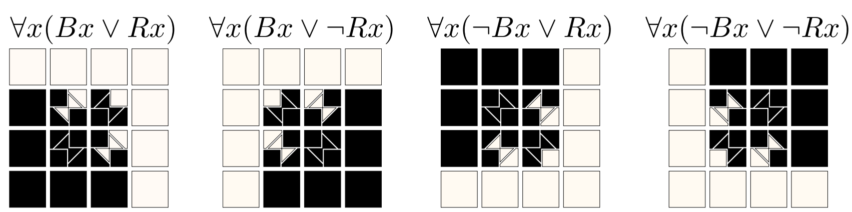

Figure 3

Mosaics arise from a meta-analysis of a system of Venn diagrams and function as a bookkeeping device for tracking entailment relations that are difficult to follow from the diagrams alone. Simple cases can be handled by pattern recognition.

Although their meanings are not further decomposable into a sub system of cells, the following Venn diagrams may still be expressed as a conjunction of true statements. It is possible to recover a natural-language counterpart of the concise QL formula with attentive analysis of the diagram alone. Examples are provided in Appendix A. This information has already been built into the accompanying mosaic, and it is often easier to retrieve the concise QL formula directly from there. However, because a Venn diagram places the relevant logical relationships directly on view, the assurance that all logical possibilities have been exhausted is more convincing than that afforded by its mosaic counterpart. The diagram and its meta-mosaic should therefore be understood as complementary.

To keep matters simple, we narrow our attention to the positive predicates represented by the mosaics shown in Figure 4.

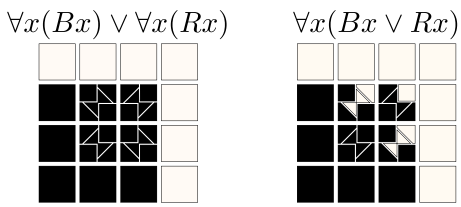

Figure 4

The domain we shall consider consists of flower stalks, each bearing blue (B) or red (R) petals. Working through the full set of possibilities, each diagram represents the relevant truth conditions. The interpretation of the first three diagrams aligns with a simple Venn diagram reading: each models a semantic cell, and their meanings are primitive. Superimposing these diagrams therefore yields a construction of semantic cells, rather than a single cell. The inclusion of an element extends the expressive power of the diagrams, revealing universal disjunction as defined over twelve distinct meanings.

P1. $\forall x(Bx) \wedge \forall x(Rx)$

Every stalk has a blue petal and a red petal.

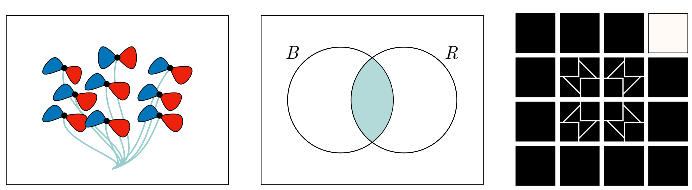

P2. $\forall x(Bx) \wedge \forall x(\neg Rx)$

Every stalk has a blue petal, none has a red petal.

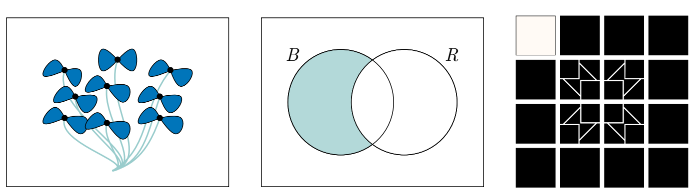

P3. $\forall x(Rx) \wedge \forall x(\neg Bx)$

Every stalk has a red petal, none has a blue petal.

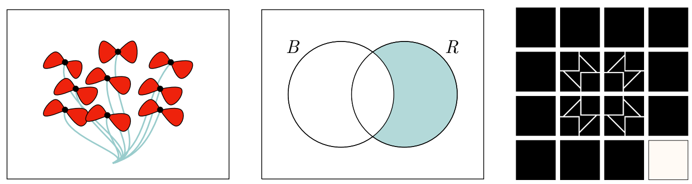

P4. $\forall x(Bx) \wedge Ra \wedge \exists x(\neg Rx)$

Every stalk has a blue petal. While stalk a has a red petal, some stalks do not.

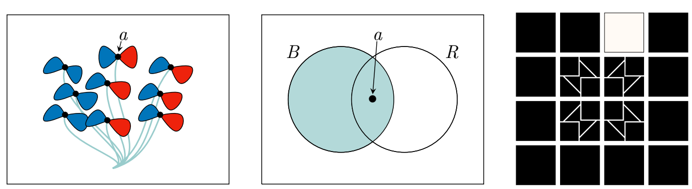

P5. $\forall x(Bx) \wedge \neg Ra \wedge \exists x(Rx)$

Every stalk has a blue petal. While some stalks have a red petal, stalk a does not.

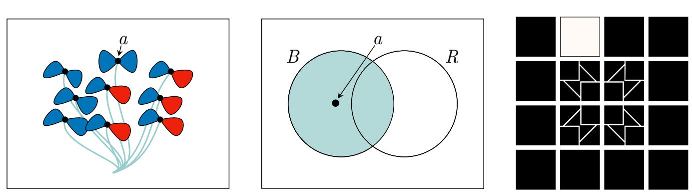

P6. $\forall x(Rx) \wedge Ba \wedge \exists x(\neg Bx)$

Every stalk has a red petal. While stalk a has a blue petal, other stalks do not.

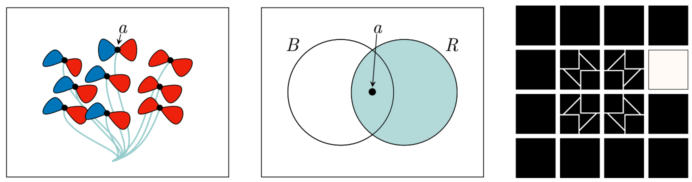

P7. $\forall x(Rx) \wedge \neg Ba \wedge \exists x(Bx)$

Every stalk has a red petal. While some stalks have a blue petal, stalk a does not.

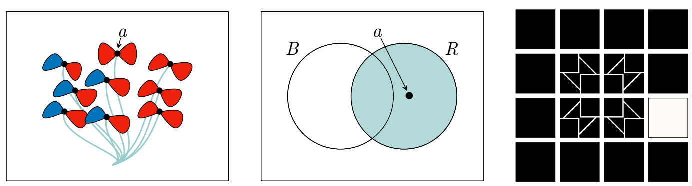

P8. $\forall x(Bx \vee Rx) \wedge Ba \wedge Ra \wedge \exists x(\neg Bx) \wedge \exists x(\neg Rx)$

Every stalk has a blue or red petal. While stalk a has both, some stalks do not have a blue petal and some do not have a red petal.

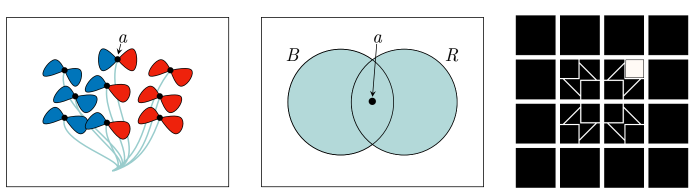

P9. $\forall x(Bx \lor Rx) \land \neg Ra \land \exists x(Bx \land Rx) \land \exists x(\neg Bx)$

Every stalk has a blue or red petal and some stalks have both a blue petal and a red petal. Stalk a does not have a red petal, while other stalks do not have a blue petal.

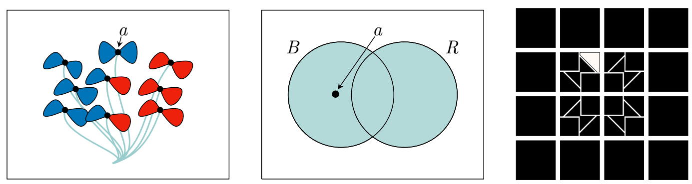

P10. $\forall x(Bx \lor Rx) \land \exists x(Bx \land Rx) \land \neg Ba \land \exists x(\neg Rx)$

Every stalk has a blue or red petal and some stalks have both a blue petal and a red petal. Stalk a does not have a blue petal, while other stalks do not have a red petal.

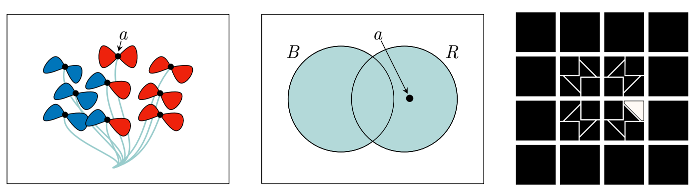

P11. $\forall x(Bx \lor Rx) \land \forall x(Bx \to \neg Rx) \land Ba \land \exists x(Rx)$

Stalks either have a blue or a red petal, and if a stalk has a blue petal then it does not have a red petal. While stalk a has a blue petal, other stalks have a red petal.

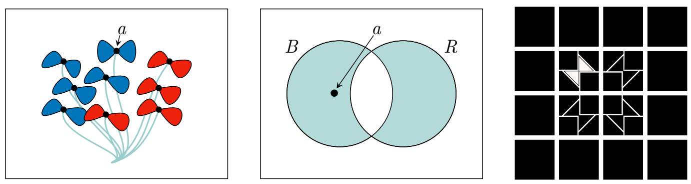

P12. $\forall x(Rx \lor Bx) \land \forall x(Bx \to \neg Rx) \land Ra \land \exists x(Bx)$.

Stalks either have a blue or a red petal, and if a stalk has a blue petal then it does not have a red petal. While stalk a has a red petal, other stalks have a blue petal.

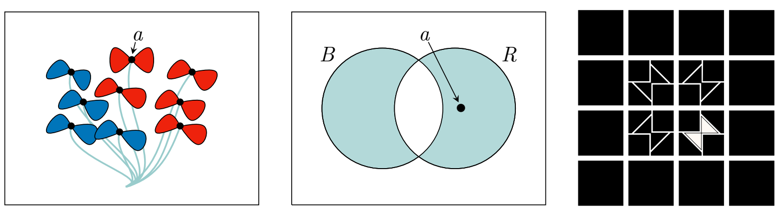

This set of twelve diagrams accounts for all the possibilities of positive universal disjunction. While a purely syntactic approach may arrive at the same determination, it does so less systematically. Without semantic support, any decomposition of universal disjunction must be carried out manually, by working through the relevant permutations of well-formed formulae. Once completed, this process yields four valid arguments that are already implicit in the cellular framework. The routine formal proofs are omitted here due to their length.

The first of the four establishes that any two propositions of the set P1–P12 cannot both be true simultaneously.

$$\vdash \neg(Pn \land Pm),$$

where Pn and Pm are any two QL sentences taken from P1 to P12.

This metatheorem is needed if the twelve propositions are semantic cells. Although the proof is not reproduced here, careful attention to the meaning of the sentences themselves will show that this is a set of contrary statements.

The next argument, valid in QL, confirms the disjunction of two universally quantified sentences is equal to the disjunction P1 to P7.

$$\forall x(Bx) \vee \forall x(Rx) \dashv\vdash \bigvee (P1 - P7)\,.$$

Given that each of P1 through P7 is a minimal statement (in the sense that it defines a semantic cell implied only by a logically equivalent sentence or by a contradiction), and that any two of these propositions form a contrary pair, their full disjunction provides a semantic decomposition of $\forall x(Bx) \vee \forall x(Rx)$ in disjunctive normal form.

The third of the four valid arguments confirms that distributed universal quantification is equal to the disjunction P1 to P12.

$$\forall x(Bx \vee Rx) \dashv\vdash \bigvee (P1 - P12)\,.$$

Since $\bigvee (P1 - P7)$ obviously entails $\bigvee (P1 - P12)$, and both $\forall x(Bx) \vee \forall x(Rx)$ and $\forall x(Bx \vee Rx)$ decompose into these respective normal forms, the last two entailments, valid in QL, confirm our semantic analysis while affording a clear view of the difference between the two forms of universal disjunction.[1]

Finally, there are the remaining diagrams and mosaics P13–P32 that complete the full set of possibilities. These are found in Appendix B. This final set of thirty-two cells provides a semantic foundation for the following QL theorem:

$$\vdash \bigvee (P1 - P32)\,.$$

A disjunction of thirty-two propositions would clearly constitute a cumbersome derivation. However, the eight propositions that form a mosaic's top-right quadrant, shown in Figure 5, are jointly equivalent to $Ba \wedge Ra$.

[1]Standard normal forms in first-order logic, such as prenex and Skolem normal form, are syntactic transformations designed to facilitate proof-theoretic or computational tasks. While various semantic techniques—such as semantic tableaux, Hintikka-style normal forms, and locality results for restricted fragments—provide ways of analysing model classes, they are not typically formulated so as to provide a finite, cell-based decomposition of the meanings of universally quantified disjunctions of the kind developed here, nor do they distinguish the two forms of universal disjunction by means of differing semantic normal forms.

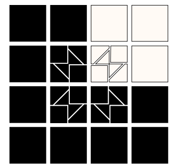

Figure 5

As a matter of symmetry, the $Ba \wedge Ra$ cells stand as a template for the remaining three quadrants. The eight propositions forming the $Ba \wedge Ra$ quadrant are:

$$
\begin{aligned}
P1 &: \forall x(Bx) \wedge \forall x(Rx),\\
P4 &: Ra \wedge \forall x(Bx) \wedge \exists x(\neg Rx),\\
P6 &: Ba \wedge \forall x(Rx) \wedge \exists x(\neg Bx),\\
P8 &: Ba \wedge Ra \wedge \forall x(Bx \vee Rx) \wedge \exists x(\neg Bx) \wedge \exists x(\neg Rx),\\
P14 &: Ra \wedge \forall x(Rx \to Bx) \wedge \forall x(\neg Rx \to \neg Bx) \wedge \exists x(\neg Bx),\\
P20 &: Ra \wedge \forall x(Rx \to Bx) \wedge \exists x(\neg Bx) \wedge \exists x(Bx \wedge \neg Rx),\\
P23 &: Ba \wedge \forall x(Bx \to Rx) \wedge \exists x(\neg Rx) \wedge \exists x(Rx \wedge \neg Bx),\\
P26 &: Ba \wedge Ra \wedge \exists x(Bx \wedge \neg Rx) \wedge \exists x(\neg Bx \wedge Rx) \wedge \exists x(\neg Bx \wedge \neg Rx).
\end{aligned}
$$

As before, each of these propositions is a semantic cell and any two form a contrary pair. Thus,

$$Ba \wedge Ra \;\dashv\vdash\; \bigvee(P1, P4, P6, P8, P14, P20, P23, P26).$$

This result reminds us that conjunctions of closed formulae that are syntactically quantifier-free may encode genuinely non-redundant quantificational content at the level of meaning, despite the absence of any overt quantificational syntax.

Conversely, this semantic framework makes it evident that twenty-eight of the thirty-two cells rely on the closed atomic sentences Ba and Ra, and therefore require a referent. This is especially evident when distinguishing between two mutually exclusive subsets within a domain, as illustrated by the example of even and odd positive integers with which we commence this investigation. To articulate the difference, at least one element must be fixed as the referent; without it, despite being syntactically well formed, the meanings of the relevant QL sentences do not settle into a semantics able to track the valid argument structure of the present QL fragment. Once the semantics is sufficiently fine-grained, it becomes clear that universal disjunction relies on a constant to resolve its meaning.

The role of a referent as a precondition for the evaluation of truth-conditions is scaffolding that passes largely unnoticed without this semantic emphasis. Once pointed out, the different commitments of syntax and semantics begin to surface. But perhaps the practical advantage of taking the semantic route makes itself felt if a visualisation like that in Figure 4 is found useful for explaining why the entailment shown in Equation (1) holds in only one direction.

All theorems and valid forms of argument stated in this paper have been verified using the automated theorem prover Prover9.

Appendix A: Redundancy Analysis

To the right of each Venn diagram is a list of true statements. The headline QL formula provides a minimal yet precise description of the content of the diagram using the simplest available syntax. Statements shown in red are redundant, while those shown in blue are interchangeable and may be substituted for one another without loss of precision. The analysis of the first few diagrams is straightforward.

P1: $\forall x(Bx) \wedge \forall x(Rx)$

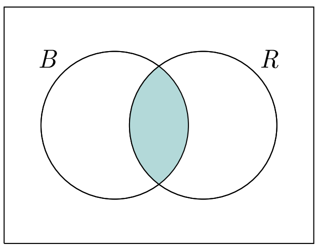

- Everything is blue / nothing is not blue.

  Something is blue.

  Element a is blue.

- Everything is red / nothing is not red.

  Element a is red.

  Something is red.

P2: $\forall x(Bx) \wedge \forall x(\neg Rx)$

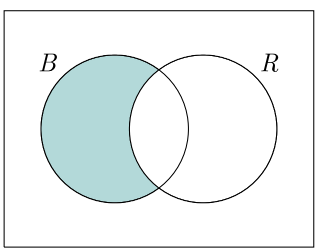

- Everything is blue / nothing is not blue.

  Something is blue.

  Element a is blue.

- Everything is not red / nothing is red.

  Element a is not red.

  Something is not red.

P3: $\forall x(Rx) \wedge \forall x(\neg Bx)$

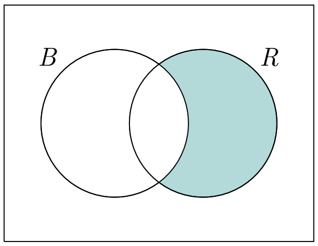

- Everything is red / nothing is not red.

  Something is red.

  Element a is red.

- Everything is not blue / nothing is blue.

  Element a is not blue.

  Something is not blue.

  Something is red and not blue.

For those diagrams that explicitly reference an element, this indicates that one or both predicates do not fall within the scope of a universal quantifier in the region in which the element appears.

P4: $\forall x(Bx) \wedge Ra \wedge \exists x(\neg Rx)$

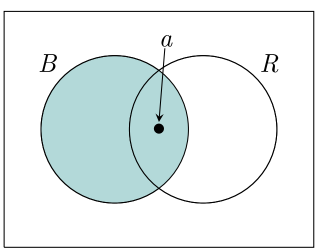

- Everything is blue / nothing is not blue.

  Something is blue.

  Element a is blue.

- Element a is red.

  Something is red.

  Something is red and blue.

  If it is red then it is blue.

- Something is not red.

  Something is blue and not red.

P5: $\forall x(Bx) \land \neg Ra \land \exists x(Rx)$

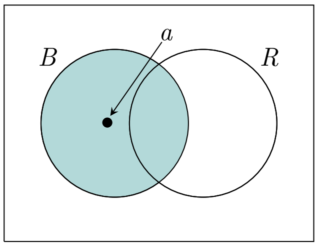

- Everything is blue / nothing is not blue.

  Something is blue.

  Element a is blue.

- Element a is not red.

  Something is not red.

  Something is both blue and not red.

- Something is red.

  Something is red and blue.

  If it is red then it is blue.

P6: $\forall x(Rx) \land Ba \land \exists x(\neg Bx)$

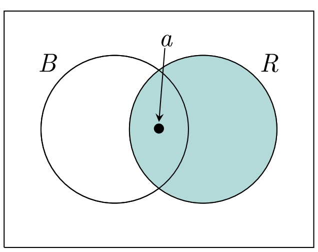

- Everything is red.

  Something is red.

  Element a is red.

- Element a is blue.

  Something is blue.

  Something is both blue and red.

  If it is blue then it is red.

- Something is not blue.

  Something is red and not blue.

P7: $\forall x(Rx) \land \neg Ba \land \exists x(Bx)$

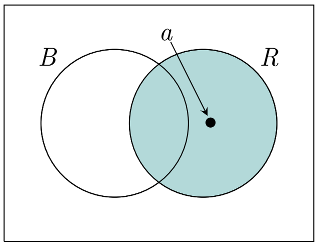

- Everything is red / nothing is not red.

  Something is red.

  Element a is red.

- Element a is not blue.

  Something is not blue.

  Something is red and not blue.

- Something is blue.

  Something is both blue and red.

  If it is blue then it is red.

P8: $\forall x(Bx \lor Rx) \land Ba \land Ra \land \exists x(\neg Bx) \land \exists x(\neg Rx)$

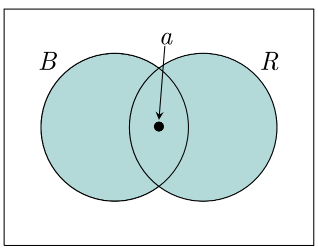

- Everything is blue or red.

- Element a is blue.

  Something is blue.

- Element a is red.

  Something is red.

  Something is both blue and red.

- Something is not blue.

  Something is red and not blue.

- Something is not red.

  Something is blue and not red.

P9: $\forall x(Bx \vee Rx) \wedge \neg Ra \wedge \exists x(Bx \wedge Rx) \wedge \exists x(\neg Bx)$

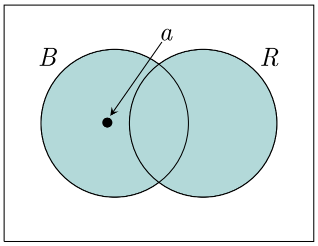

- Everything is blue or red.
- Element a is not red.

  Something is not red.

  Element a is blue.

  Something is blue.
- Something is both blue and red.
- Something is not blue.

  Something is red and not blue.

To describe this diagram it may not be immediately obvious that $\neg Ra$ is a necessary assumption. For example, the following two sentences

$$\forall x(Bx \vee Rx) \wedge Ba \wedge \exists x(Bx \wedge Rx) \wedge \exists x(Rx),$$

$$\forall x(Bx \vee Rx) \wedge Ba \wedge \exists x(Bx \wedge Rx) \wedge \exists x(\neg Bx),$$

both leave open the possibility that everything might be red. To rule this out we assume $\neg Ra$. The conjunction $\forall x(Bx \vee Rx) \wedge \neg Ra$ functions as a form of disjunctive syllogism and determines that a is blue. Similarly, $\forall x(Bx \vee Rx) \wedge \exists x(\neg Bx)$ entails $\exists x(\neg Bx \wedge Rx)$. Thus, both conclusions are redundant descriptions of the diagram.

However, we do not have to rely on formal derivation rules to determine redundancy. Rather, the semantic analysis is complete once each region of the diagram has been delineated.

To do this, we remove sentences individually from the set of facts and check to see whether the remaining list continues to describe the diagram accurately. Of the four non-redundant sentences, the first three identify three of the diagrams four regions. We then need only identify the fourth region to complete the analysis. The diagram clearly shows that the statement something is not blue must mean something is red and not blue. At this point, each region of the diagram has been described. Careful study shows that whenever the diagram is described in full, these four sentences are always part of that description.

P10: $\forall x(Bx \lor Rx) \land \exists x(Bx \land Rx) \land \neg Ba \land \exists x(\neg Rx)$

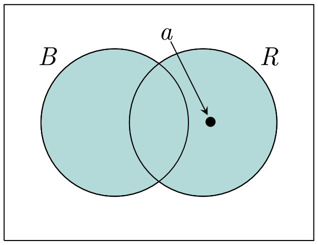

- Everything is blue or red.
- Element a is not blue.

  Something is not blue.

  Element a is red.

  Something is red.
- Something is both blue and red.
- Something is not red.

  Something is not red and blue.

P11: $\forall x(Bx \lor Rx) \land \forall x(Bx \to \neg Rx) \land Ba \land \exists x(Rx)$

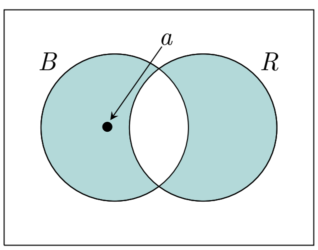

- Everything is blue or red.
- Everything blue is not red.
- Everything red is not blue.

  Something is not both red and blue.
- Element a is blue.

  Element a is not red.

  Something is not red.
- Something is red.
- Something is not blue.

In this example the sentences "something is red" and "something is not blue" are interchangeable relative to the background assumptions P. Let

$$P := \forall x(Bx \lor Rx) \land \forall x(Bx \to \neg Rx).$$

These diagrams and their mosaic, model the following pattern of valid QL entailments:

$$P \wedge Ba \wedge \exists x(Rx) \dashv\vdash P \wedge \neg Ra \wedge \exists x(Rx),$$
$$P \wedge Ba \wedge \exists x(Rx) \dashv\vdash P \wedge Ba \wedge \exists x(\neg Bx),$$
$$P \wedge Ba \wedge \exists x(Rx) \dashv\vdash P \wedge \neg Ra \wedge \exists x(\neg Bx),$$
$$P \wedge Ba \wedge \exists x(Rx) \quad \vdash P \wedge \neg Ra \wedge \exists x(\neg Rx).$$

Again correctly, there is no model supporting entailment in the following direction, which is invalid in QL:

$$P \wedge \neg Ra \wedge \exists x(\neg Rx) \nvdash P \wedge Ba \wedge \exists x(Rx).$$

P12: $\forall x(Bx \vee Rx) \wedge \forall x(Bx \rightarrow \neg Rx) \wedge Ra \wedge \exists x(Bx)$.

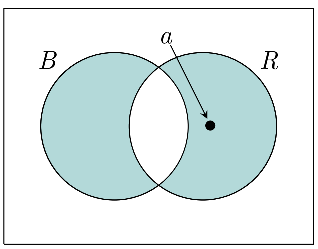

- Everything is blue or red.
- Everything blue is not red.
- Everything red is not blue.

  Something is not both red and blue.
- Element a is blue.

  Element a is not red.

  Something is not red.
- Something is red.
- Something is not blue.

This completes the analysis for universal disjunction when R and B are both positive.

Appendix B: Remaining Models

P13. $\forall x(\neg Bx) \land \forall x \neg(Rx)$.

No stalk has a blue petal and no stalk has a red petal.

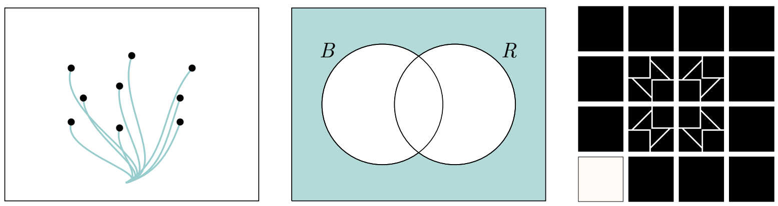

P14. $\forall x(Rx \to Bx) \land \forall x(\neg Rx \to \neg Bx) \land Ra \land \exists x(\neg Bx)$.

If stalks have a red petal then they have a blue petal, and
if they do not have a red petal they do not have a blue petal.
Stalk a has a red petal, while other stalks do not have a blue petal.

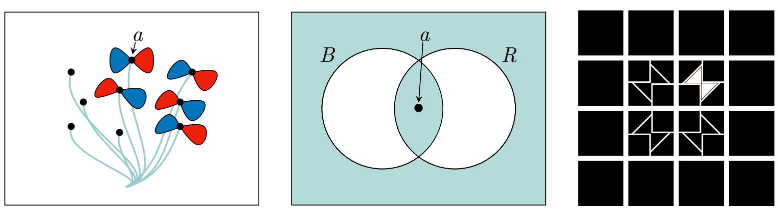

P15. $\forall x(Bx \to Rx) \land \forall x(\neg Bx \to \neg Rx) \land \neg Ba \land \exists x(Bx)$.

If stalks have a blue petal then they have a red petal, and
if they do not have a blue petal they do not have a red petal.
Stalk a does not have a red petal,
while some other stalks do not have a blue petal.

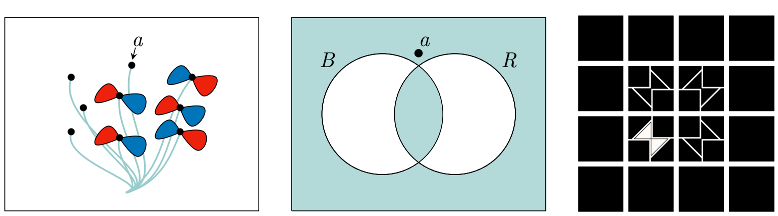

P16. $\forall x(\neg Rx) \wedge Ba \wedge \exists x(\neg Bx)$.

No stalk has a red petal. Stalk a has a blue petal
but some stalks do not have a blue petal.

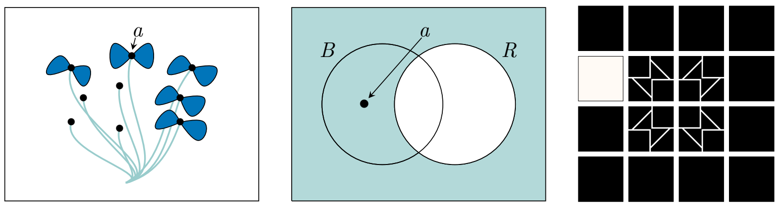

P17. $\forall x(\neg Rx) \wedge \neg Ba \; \wedge \exists x(Bx)$.

No stalk has a red petal.
Stalk a does not have a blue petal but some stalks have a blue petal.

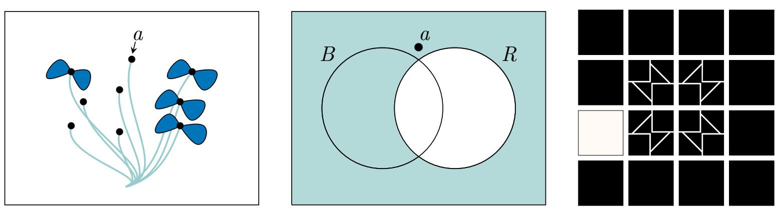

P18. $\forall x(\neg Bx) \wedge Ra \; \wedge \exists x(\neg Rx)$.

No stalk has a blue petal. Stalk a has a red petal
but some stalks do not have a red petal.

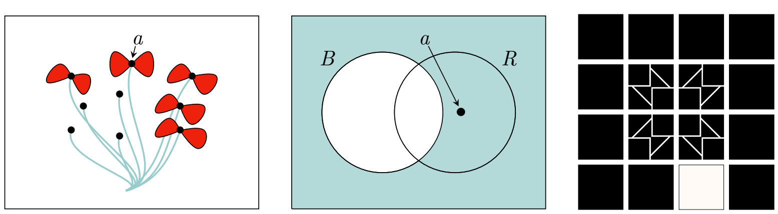

P19. $\forall x(\neg Bx) \land \neg Ra \land \exists x(Rx)$.

No stalk has a blue. petal Stalk *a* does not have a red petal.
but some stalks have red petals.

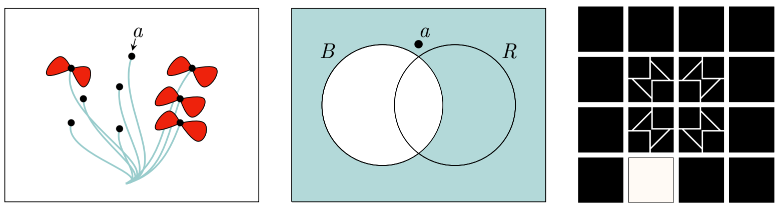

P20. $\forall x(Rx \to Bx) \land \exists x(\neg Bx) \land Ra \land \exists x(Bx \land \neg Rx)$.

If a stalk has a red petal then it has a blue petal.
Although some stalks do not have a blue petal, stalk *a* has a red,
while some stalks have a blue but not a red petal.

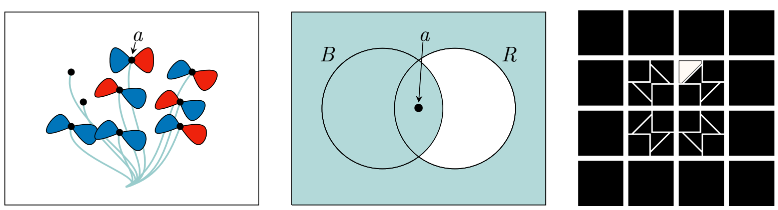

P21. $\forall x(Rx \to Bx) \land Ba \land \neg Ra \land \exists x(\neg Bx) \land \exists x(Rx)$.

If a stalk has a red petal then it has a blue petal.
Stalk *a* has a blue but not a red petal.
Some stalks do not have a blue petal,
while some stalks have a red petal.

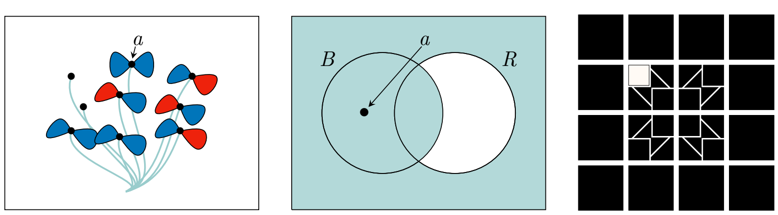

P22. $\forall x(Rx \rightarrow Bx) \wedge \neg Ba \wedge \exists x(Rx) \wedge \exists x(Bx \wedge \neg Rx)$.

If a stalk has a red petal then it has a blue petal.
Although stalk a does not have a blue petal, some have red,
and some stalks have a blue but not a red petal.

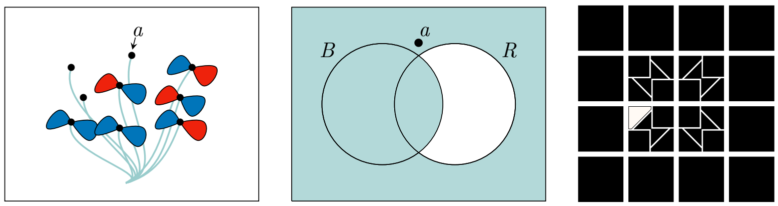

P23. $\forall x(Bx \rightarrow Rx) \wedge Ba \wedge \exists x(\neg Rx) \wedge \exists x(Rx \wedge \neg Bx)$.

If a stalk has a blue petal then it has a red petal.
While stalk a has a blue petal, some do not have a red petal,
and some stalks have a red but not a blue petal.

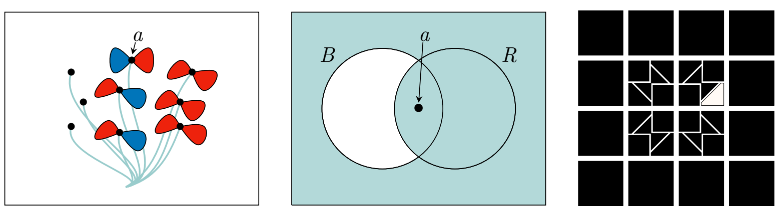

P24. $\forall x(Bx \rightarrow Rx) \wedge \neg Ba \wedge Ra \wedge \exists x(Bx) \wedge \exists x(\neg Rx)$.

If a stalk has a blue petal then it has a red petal.
Although stalk a has a red petal, but not blue,
some have a red and others do not have a blue.

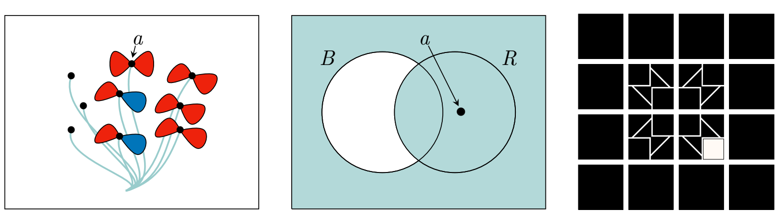

P25. $\forall x(Bx \rightarrow Rx) \wedge \neg Ra \wedge \exists x(Bx) \wedge \exists x(Rx \wedge \neg Bx)$.

If a stalk has a blue petal then it has a red petal. Although stalk a does not have a red petal, some have blue, while some have a red but not a blue petal.

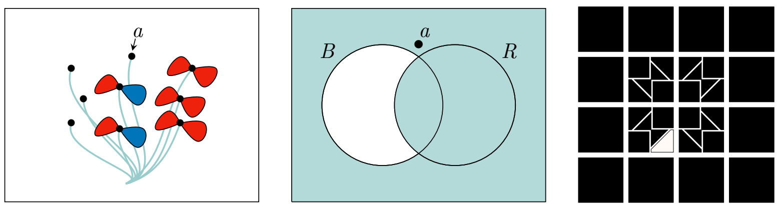

P26. $Ba \wedge Ra \wedge \exists x(Bx \wedge \neg Rx) \wedge \exists x(\neg Bx \wedge Rx) \wedge \exists x(\neg Bx \wedge \neg Rx)$.

Stalk a has both a blue and a red petal, some stalks have blue but not red petals, some have red but not a blue, others have neither a blue nor a red petal.

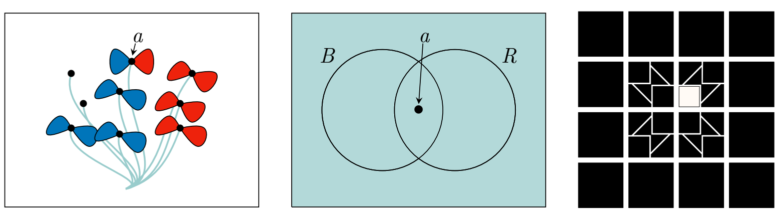

P27. $Ba \wedge \neg Ra \wedge \exists x(Bx \wedge \ Rx) \wedge \exists x(Rx \wedge \neg Bx) \wedge \exists x(\neg Bx \wedge \neg Rx)$.

Stalk a has a blue but not a red petal, some have both a blue and a red petal, some have red but not blue, others have neither a blue nor a red petal.

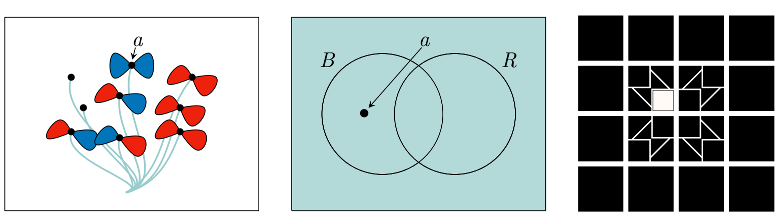

P28. $\neg Ba \wedge Ra \wedge \exists x(Bx \wedge Rx) \wedge \exists x(Bx \wedge \neg Rx) \wedge \exists x(\neg Bx \wedge \neg Rx)$.

> Stalk a has a red but not a blue petal, some have both a blue and a red petal, some have blue but not red, others have neither a blue nor a red petal.

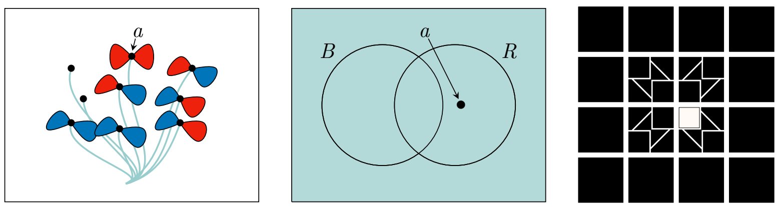

P29. $\neg Ba \wedge \neg Ra \wedge \exists x(Bx \wedge Rx) \wedge \exists x(Bx \wedge \neg Rx) \wedge \exists x(Rx \wedge \neg Bx)$.

> Stalk a has neither a blue nor a red petal, some have both a blue and a red petal, some have blue but not red, others have a red petal but not a blue.

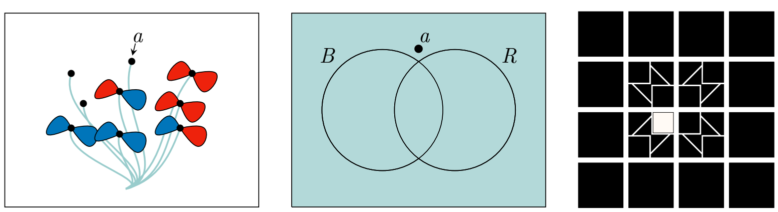

P30. $\forall x(Rx \rightarrow \neg Bx) \wedge Ba \wedge \exists x(Rx) \wedge \exists x(\neg Bx \wedge \neg Rx)$

> If a stalk has a red petal it does not have a blue petal. Stalk a has a blue petal, while some have a red. Some stalks have neither a blue nor a red petal.

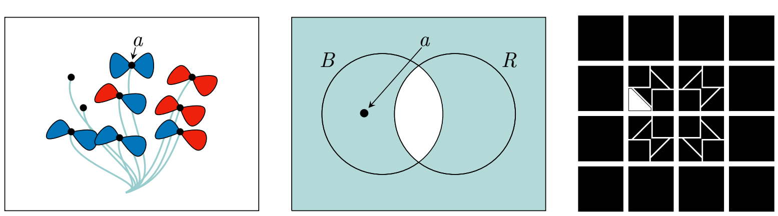

P31. $\forall x(Rx \rightarrow \neg Bx) \wedge Ra \wedge \exists x(Bx) \wedge \exists x(\neg Bx \wedge \neg Rx)$.

If a stalk has a red petal it does not have a blue petal.
Stalk a has a red petal, while some have a blue.
Some stalks have neither a blue nor a red petal.

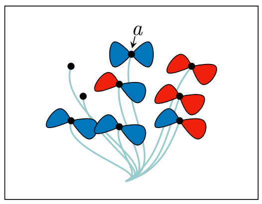

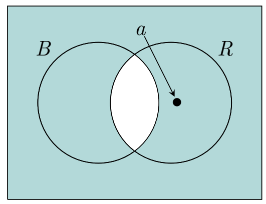

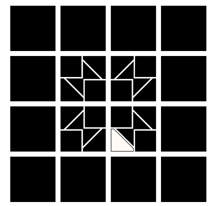

P32. $\forall x(Rx \rightarrow \neg Bx) \wedge \neg Ba \wedge \neg Ra \wedge \exists x(Bx) \wedge \exists x(Rx)$.

If a stalk has a red petal it does not have a blue petal.
Some stalks have a blue petal, while some have a red.
Stalk a has neither a blue nor a red petal.

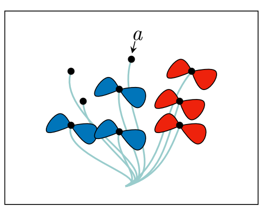

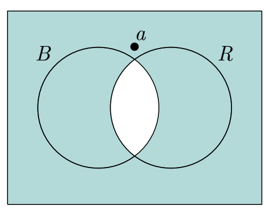

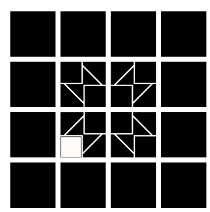

References

[1] Edward John Lemmon. *Beginning logic*. CRC Press, 1971.