Chapter 3

Syllogistic Reasoning

This chapter 'opens the box' of propositional logic, and looks further inside the statements that we make when we describe the world. Very often, these statements are about objects and their properties, and we will now show you a first logical system that deals with these. *Syllogistics* has been a standard of logical reasoning since Greek Antiquity. It deals with quantifiers like 'All P are Q' and 'Some P are Q', and it can express much of the common sense reasoning that we do about predicates and their corresponding sets of objects. You will learn a famous graphical method for dealing with this, the so-called 'Venn Diagrams', after the British mathematician John Venn (1834–1923), that can tell valid syllogisms from invalid ones. As usual, the chapter ends with some outlook issues, toward logical systems of inference, and again some phenomena in the real world of linguistics and cognition.

3.1 Reasoning About Predicates and Classes

Aristotle

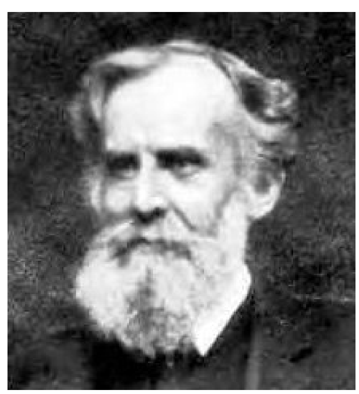

John Venn

The Greek philosopher Aristotle (384 BC – 322 BC) proposed a system of reasoning in his *Prior Analytics* (350 BC) that was so successful that it has remained a paradigm of logical reasoning for more than two thousand years: the *Syllogistic*.

Syllogisms A *syllogism* is a logical argument where a quantified statement of a specific form (the conclusion) is inferred from two other quantified statements (the premises).

The quantified statements are all of the form "Some/all A are B," or "Some/all A are not B," and each syllogism combines three predicates or properties. Notice that "All A are not B" can be expressed equivalently in natural language as "No A are B," and "Some A are not B" as "Not all A are B." We can see these quantified statements as describing relations between predicates, which is well-suited to describing hierarchies of properties. Indeed, Aristotle was also an early biologist, and his classifications of predicates apply very well to reasoning about species of animals or plants.

Your already know the following notion. A syllogism is called *valid* if the conclusion follows logically from the premises in the sense of Chapter 2: whatever we take the real predicates and objects to be: if the premises are true, the conclusion must be true. The syllogism is *invalid* otherwise.

Here is an example of a valid syllogism:

$$\frac{\begin{array}{c}\text{All Greeks are humans}\\ \text{All humans are mortal}\end{array}}{\text{All Greeks are mortal.}} \tag{3.1}$$

We can express the validity of this pattern using the $\models$ sign introduced in Chapter 2:

$$\text{All Greeks are humans, All humans are mortal} \models \text{All Greeks are mortal.} \tag{3.2}$$

This inference is valid, and, indeed, this validity has nothing to do with the particular predicates that are used. If the predicates *human*, *Greek* and *mortal* are replaced by different predicates, the result will still be a valid syllogism. In other words, it is the *form* that makes a valid syllogism valid, not the *content* of the predicates that it uses. Replacing the predicates by symbols makes this clear:

$$\frac{\begin{array}{c}\text{All A are B}\\ \text{All B are C}\end{array}}{\text{All A are C.}} \tag{3.3}$$

The classical quantifiers Syllogistic theory focusses on the quantifiers in the so called *Square of Opposition*, see Figure (3.1). The quantifiers in the square express relations between a first and a second predicate, forming the two arguments of the assertion. We think of these predicates very concretely, as sets of objects taken from some domain of

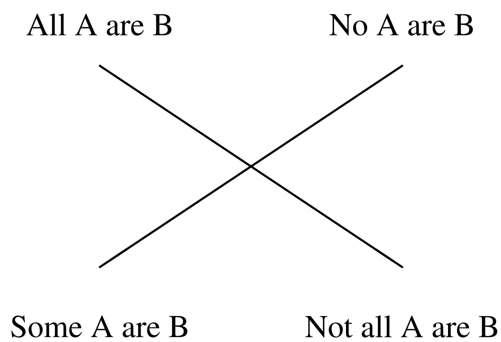

Figure 3.1: The Square of Opposition

discourse that satisfy the predicate. Say, ‘boy’ corresponds with the set of all boys in the relevant situation that we are talking about.

The quantified expressions in the square are related across the diagonals by external (sentential) negation, and across the horizontal edges by internal (or verb phrase) negation. It follows that the relation across the vertical edges of the square is that of internal plus external negation; this is the relation of so-called *quantifier duality*.

Because Aristotle assumes that the left-hand predicate A is non-empty (see below), the two quantified expressions on the top edge of the square cannot both be true; these expressions are called *contraries*. Similarly, the two quantified expressions on the bottom edge cannot both be false: they are so-called *subcontraries*.

Existential import Aristotle interprets his quantifiers with *existential import*: *All A are B* and *No A are B* are taken to imply that there are *A*. Under this assumption, the quantified expressions at the top edge of the square imply those immediately below them. The universal affirmative quantifier *all* implies the individual affirmative *some* and the universal negative *no* implies the individual negative *not all*. Existential import seems close to how we use natural language. We seldom discuss ‘empty predicates’ unless in the realm of phantasy. Still, modern logicians have dropped existential import for reasons of mathematical elegance, and so will we in this course.

The universal and individual affirmative quantifiers are said to be of types **A** and **I** respectively, from Latin **A***ff***I***rmo*, the universal and individual negative quantifiers of type **E** and **O**, from Latin *N***E***g***O**. Aristotle’s theory was extended by logicians in the Middle Ages whose working language was Latin, whence this Latin mnemonics. Along these lines, *Barbara* is the name of the syllogism with two universal affirmative premises and a universal affirmative conclusion. This is the syllogism (3.1) above.

Here is an example of an invalid syllogism:

$$\begin{array}{l} \text{All warlords are rich} \\ \text{No students are warlords} \\ \hline \text{No students are rich} \end{array} \tag{3.4}$$

Why is this invalid? Because one can picture a situation where the premises are true but the conclusion is false. Such a *counter-example* can be very simple: just think of a situation with just one student, who is rich, but who is not a warlord. Then the two premises are true (there being no warlords, all of them are rich – but you can also just add one rich warlord, if you like existential import). This 'picturing' can be made precise, and we will do so in a moment.

3.2 The Language of Syllogistics

Syllogistic statements consist of a quantifier, followed by a common noun followed by a verb: $Q\ N\ V$. This is an extremely general pattern found across human languages. Sentences S consist of a Noun Phrase *NP* and a Verb Phrase *VP*, and the Noun Phrase can be decomposed into a Determiner Q plus a Common Noun *CN*:

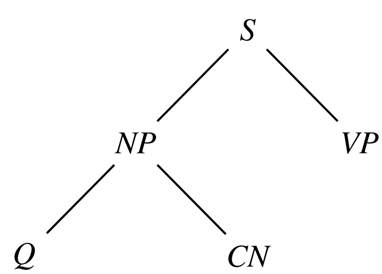

Thus we are really at the heart of how we speak. In these terms, a bit more technically, Aristotle studied the following inferential pattern:

$$\begin{array}{c} \text{Quantifier}_1\ \text{CN}_1\ \text{VP}_1 \\ \text{Quantifier}_2\ \text{CN}_2\ \text{VP}_2 \\ \hline \text{Quantifier}_3\ \text{CN}_3\ \text{VP}_3 \end{array}$$

where the quantifiers are *All*, *Some*, *No* and *Not all*. The common nouns and the verb phrases both express properties, at least in our perspective here ('man' stands for all men, 'walk' for all people who walk, etcetera). To express a property means to refer to a class of things, at least in a first logic course. There is more to predicates than sets of objects when you look more deeply, but this 'intensional' aspect will not occupy us here.

In a syllogistic form, there are two premises and a conclusion. Each statement refers to two classes. Since the conclusion refers to two classes, there is always one class that figures in the premises but not in the conclusion. The *CN* or *VP* that refers to this class is called the *middle term* that links the information in the two premises.

Exercise 3.1 What is the middle term in the syllogistic pattern given in (3.3)?

To put the system of syllogistics in a more systematic setting, we first make a brief excursion to the topic of operations on sets.

3.3 Sets and Operations on Sets

Building sets The binary relation $\in$ is called the element-of relation. If some object a is an element of a set A then we write $a \in A$ and if this is not the case we write $a \notin A$. Note that if $a \in A$, A is certainly a set, but a itself may also be a set. Example: $\{1\} \in \{\{1\}, \{2\}\}$.

If we want to collect all the objects together that have a certain property, then we write:

$$\{x \mid \varphi(x)\} \tag{3.5}$$

for the set of those x that have the property described by φ. Sometimes we restrict this property to a certain *domain of discourse* or *universe* U of individuals. To make this explicit, we write:

$$\{x \in U \mid \varphi(x)\} \tag{3.6}$$

to denote the set of all those x in U for which φ holds. Note that $\{x \in U \mid \varphi(x)\}$ defines a subset of U.

To describe a set of elements sharing multiple properties $\varphi_1, \ldots, \varphi_n$ we write:

$$\{x \mid \varphi_1(x), \ldots, \varphi_n(x)\} \tag{3.7}$$

Instead of a single variable, we may also have a sequence of variables. For example, we may want to describe a set of pairs of objects that stand in a certain relation. Here is an example.

$$A = \{(x, y) \mid x \text{ is in the list of presidents of the US }, y \text{ is married to } x\} \tag{3.8}$$

For example, $(\text{Bill\_Clinton}, \text{Hillary\_Clinton}) \in A$ but, due to how the 2008 presidential election turned out, $(\text{Hillary\_Clinton}, \text{Bill\_Clinton}) \notin A$. Sets of pairs are in fact the standard mathematical representation of binary relations between objects (see Chapter A).

Operations on sets In talking about sets, one often also wants to discuss combinations of properties, and construct new sets from old sets. The most straightforward operation for this is the *intersection* of two sets:

$$A \cap B = \{x \mid x \in A \text{ and } x \in B\} \tag{3.9}$$

If A and B represent two properties then $A \cap B$ is the set of those objects that have both properties. In a picture:

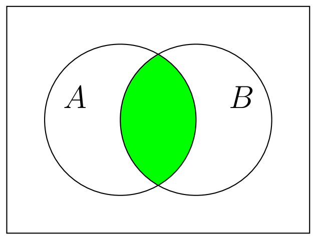

The intersection of the set of 'red things' and the set of 'cars' is the set of 'red cars'.

Another important operation is the *union* that represents the set of objects which have *at least one* of two given properties.

$$A \cup B = \{x \mid x \in A \text{ or } x \in B\} \tag{3.10}$$

The 'or' in this definition should be read in the *inclusive* way. Objects which belong to both sets also belong to the union. Here is a picture:

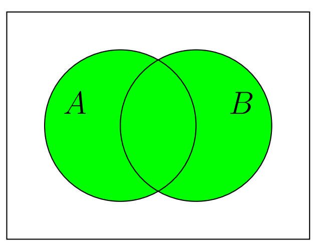

A third operation which is often used is the *difference* of two sets:

$$A \setminus B = \{x \mid x \in A \text{ and } x \notin B\} \tag{3.11}$$

If we think of two properties represented by A and B then $A \setminus B$ represents those things that have the property A but not B. In a picture:

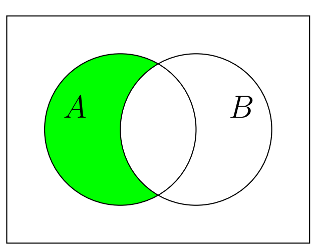

These pictorial representations of the set operations are called *Venn diagrams*, after the British mathematician John Venn (1834 - 1923). In a Venn diagram, sets are represented as circles placed in such a way that each combination of these sets is represented. In the case of two sets this is done by means of two partially overlapping circles. Venn diagrams are easy to understand, and interestingly, they are a method that also exploits our powers of non-linguistic visual reasoning.

Next, there is the *complement* of a set (relative to some given universe U (the domain of discourse):

$$\overline{A} \quad = \quad \{x \in U \mid x \notin A\} \tag{3.12}$$

In a picture:

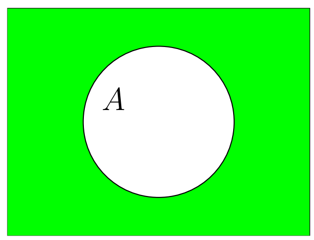

Making use of complements we can describe things that do *not* have a certain property.

The complement operation makes it possible to define set theoretic operations in terms of each other. For example, the difference of two sets A and B is equal to the intersection of A and the complement of B:

$$A \setminus B \quad = \quad A \cap \overline{B} \tag{3.13}$$

Complements of complements give the original set back:

$$\overline{\overline{A}} \quad = \quad A \tag{3.14}$$

Complement also allows us to relate union to intersection, by means of the following so-called *de Morgan equations*:

$$\begin{aligned} \overline{A \cup B} &= \overline{A} \cap \overline{B} \\ \overline{A \cap B} &= \overline{A} \cup \overline{B} \end{aligned} \tag{3.15}$$

From the second de Morgan equation we can derive a definition of the union of two sets in terms of intersection and complement:

$$A \cup B = \overline{\overline{A}} \cup \overline{\overline{B}} = \overline{\overline{A} \cap \overline{B}} \tag{3.16}$$

This construction is illustrated with Venn diagrams in Figure 3.2. Also important are the so-called distributive equations for set operations; they describe how intersection distributes over union and vice versa:

$$\begin{aligned} A \cap (B \cup C) &= (A \cap B) \cup (A \cap C) \\ A \cup (B \cap C) &= (A \cup B) \cap (A \cup C) \end{aligned} \tag{3.17}$$

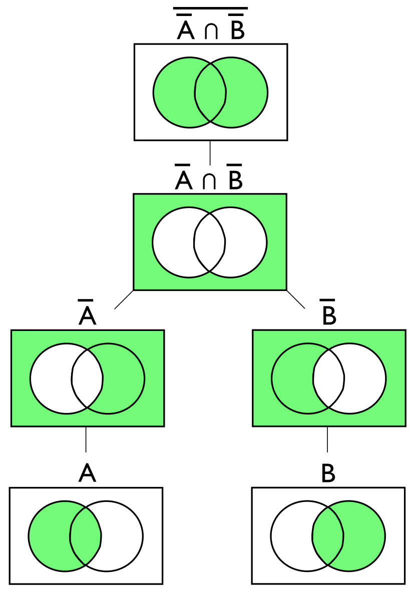

Figure 3.2: Construction of $A \cup B$ using intersection and complement.

Figure 3.3 demonstrates how the validity of the first of these equations can be computed by means of Venn-diagrams. Here we need three circles for the three sets A, B and C, positioned in such a graphical way that every possible combination of these three sets is represented in the diagrams.

The relation between sets and propositions The equalities between sets may look familiar to you. In fact, these principles have the same shape as propositional equivalences that describe the relations between $\neg$, $\wedge$ and $\vee$. In fact, the combinatorics of sets using complement, intersection and union *is* a Boolean algebra, where complement behaves like negation, intersection like conjunction and union like disjunction. The zero element of the algebra is the empty set $\emptyset$.

We can even say a bit more. The Venn-diagram constructions as in Figures 3.2 and 3.3 can be viewed as construction trees for set-theoretic expressions, and they can be reinterpreted as construction trees for formulas of propositional logic. Substitution of proposition letters for the base sets and replacing the set operations by the corresponding connectives gives a parsing tree with the corresponding semantics for each subformula made visible in the tree. A green region corresponds to a valuation which assigns the truth-value 1 to the given formula, and a white region to valuation which assigns this formula the value 0. You can see in the left tree given in Figure 3.3 that the valuations which makes the formula $a \wedge (b \vee c)$ true are abc, $a\bar{b}c$ and $ab\bar{c}$ (see Figure 3.4).

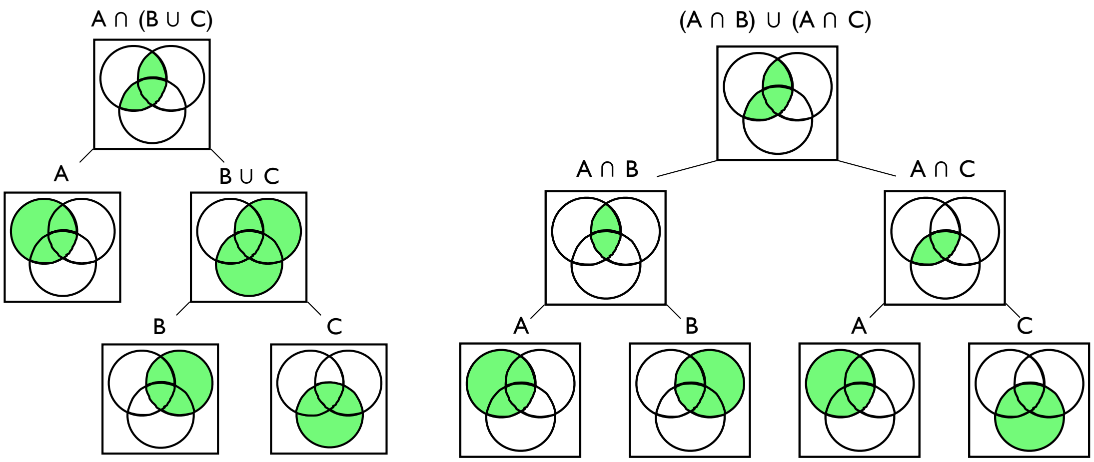

Figure 3.3: One of the distribution laws illustrated by means of Venn diagrams.

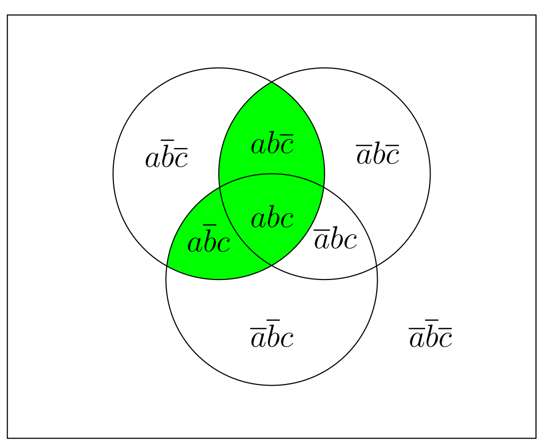

Figure 3.4: The support for $a \wedge (b \vee c)$ in a Venn-diagram.

3.4 Syllogistic Situations

Since all syllogistic forms involve just three predicates A, B and C, we can draw a general picture of a syllogistic situation as the following *Venn Diagram*:

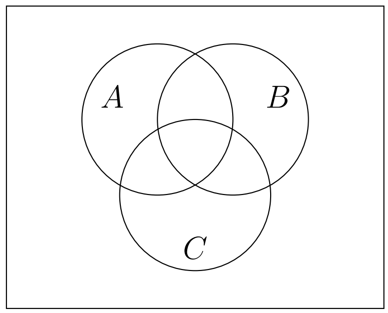

The rectangular box stands for a set of objects that form the domain of discourse, with three possible properties A, B and C. Note that there are 8 regions in all, quite properly, since that is the number of all possible combinations. An individual without any of these properties has to be outside of the three circles, like this:

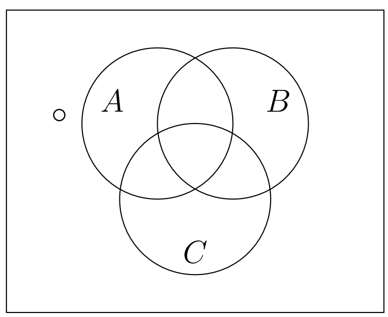

An object with property A but lacking the properties B and C has to be inside the A circle, but outside the B and C circles, like this:

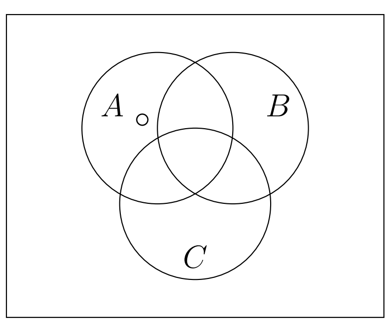

Now let us look in detail at what the Aristotelian quantifiers express. *All A are B* expresses that the part of the A circle outside the B circle has to be empty. We can indicate that in the picture by crossing out the forbidden regions, like this:

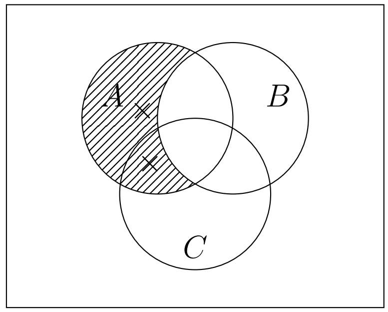

Note that the preceding picture does not take existential import into account. As we already said, we will leave it out in the interest of simplicity. And we lose nothing in this way. If you want to say that a predicate P is non-empty, you can always do so explicitly with a quantifier 'Some'.

No A are B expresses that the part of the A circle that overlaps with the B circle has to be empty. Again, we can indicate this in a picture by crossing out the forbidden areas:

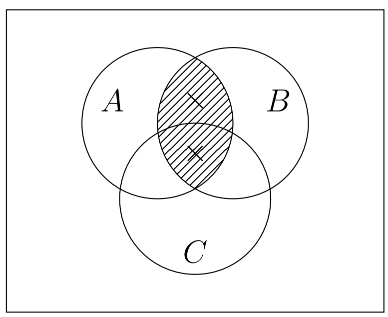

Again, existential import ("there must be A's") is not taken into account by this picture.

Now we move from universal quantifiers to existential ones. *Some A are B* expresses that the part of the picture where the A and the B circles overlap has to be non-empty. We can indicate that in the picture by putting an individual in an appropriate position. Since we do not know if that individual has property C or not, this can be done in two ways:

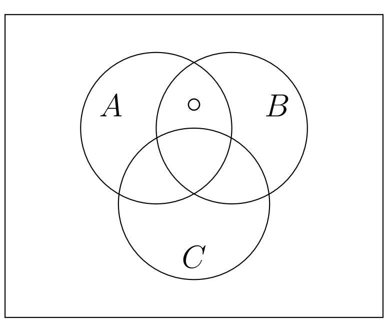

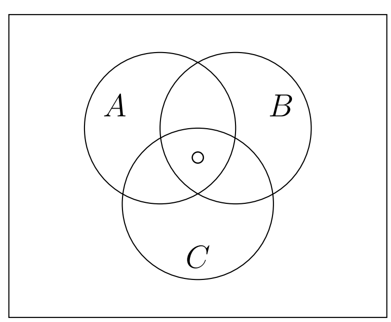

Not all A are B, or equivalently *Some are are not B*, expresses that the part of the A circle that falls outside the B circle has to be non-empty. There has to be at least one individual that is an A but not a B. Since we do not know whether this individual has property C or not, we can again picture this information in two possible ways:

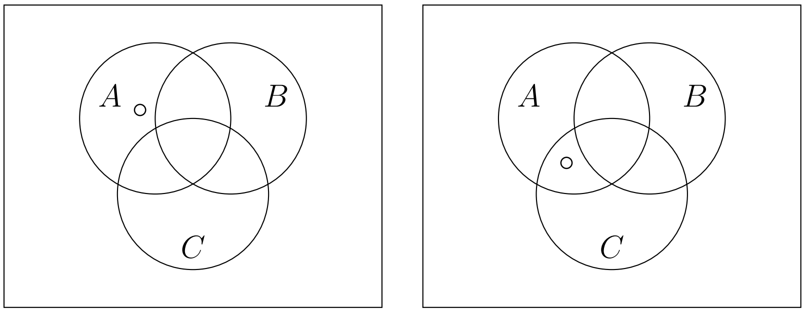

Some authors do not like this duplication of pictures, and prefer putting the small round circle for the individual on the border line of several areas.

You no doubt realize that such a duplication of cases makes the picture method much harder in terms of complexity, and hence, as we shall see, the art in checking validity for syllogisms is avoiding it whenever possible.

3.5 Validity Checking for Syllogistic Forms

The diagrams from the preceding section lead to a check for syllogistic validity:

Working with diagrams We illustrate the method with the following valid syllogism:

$$\frac{\begin{array}{l}\text{All warlords are rich}\\ \text{No student is rich}\end{array}}{\text{No warlord is a student}} \tag{3.18}$$

To carry out the validity check for this inference, we start out with the general picture of a domain of discourse with three properties. Next, we update the picture with the information provided by the premises. Here, the understanding is this:

> Crossing out a region with $\times$ means that this region is empty (there are no individuals in the domain of discourse with this combination of properties), while putting a $\circ$ in a region means that this region is non-empty (there is at least one individual with this combination of properties). Leaving a blank region means that there is no information about this region (there may be individuals with this combination of properties, or there may not).

The method is this: we *update* with the information from the premises, and next *check* in the resulting picture whether the conclusion holds or not. Let A represent the property of being a warlord, B the property of being rich, and C the property of being a student. Then we start with the following general picture:

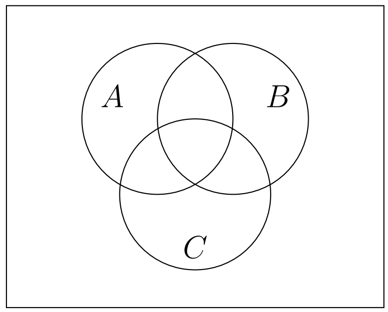

According to the first premise, *All A are B* has to be true, so we get:

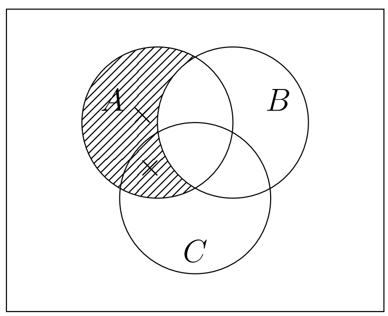

By the second premise, *No C are B* has to be true, so we extend the picture as follows:

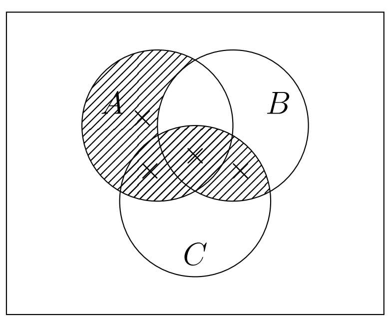

Finally, we have to check the conclusion. The conclusion says that the regions where A and C overlap have to be empty. Well, they are, for both of these regions have been crossed out. So the conclusion has to be true. Therefore, the inference is valid.

The general method The method we have used consists of the following four steps:

Draw the Skeleton Draw an empty picture of a domain of discourse with three properties A, B and C. Make sure that all eight combinations of the three sets are present.

Crossing out – Universal step Take the universal statements from the premises (the statements of the form "All ..." and "No ...", and cross out the forbidden regions in the diagram.

Filling up – Existential step Take the existential statements from the premises (the statements of the form "Some ..." and "Not all ..."), and try to make them true in the diagram by putting a $\circ$ in an appropriate region, while respecting the $\times$ signs. (This step might lead to several possibilities, all of which have to satisfy the check in the next item.)

Check Conclusion If the conclusion is universal it says that certain regions should have been crossed out. Are they? If the conclusion is existential it says that certain regions should have been marked with a $\circ$. Are they? If the answer to this question is affirmative the syllogism is valid; otherwise a counterexample can be constructed, indicating that the syllogism is invalid.

To illustrate the procedure once more, let us now take the invalid syllogism 3.4 that was mentioned before (repeated as 3.19).

$$\begin{array}{l} \text{All warlords are rich} \\ \text{No student is a warlord} \\ \hline \text{No student is rich} \end{array} \tag{3.19}$$

The symbolic form of this syllogism is:

$$\text{All A are B} \quad \text{No C are A} \quad \text{Therefore: No C are B.} \tag{3.20}$$

The premise statements are both universal. Crossing out the appropriate regions for the first premise gives us:

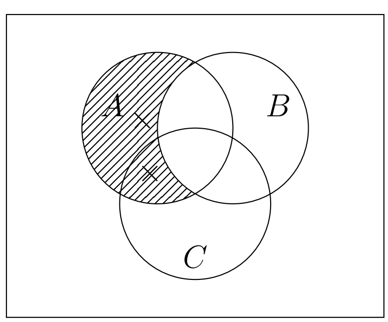

After also crossing out the regions forbidden by the second premise we get:

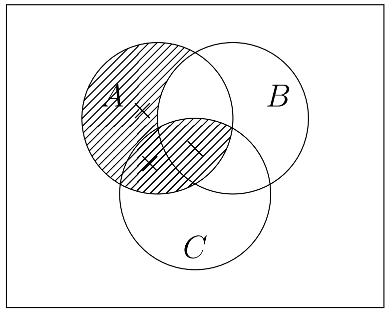

Note that the region for the AC's outside B gets ruled out twice. It looks like the second premise repeats some of the information that was already conveyed by the first premise (unlike the case with the previous example). But though this may say something about presentation of information, it does not affect valid or invalid consequences.

Finally, we check whether the conclusion holds. *No **C** are **B*** means that the regions where C and B overlap are forbidden. Checking this in the diagram we see that the region where A, B and C overlap is indeed crossed out, but the region outside A where B and C overlap is not. Indeed, the diagram does not contain information about this region. This means that we can use the diagram to construct a counterexample to the inference.

> The diagram allows us to posit the existence of an object that satisfies B and C but not A, in the concrete case of our example, a rich student who is not a warlord:

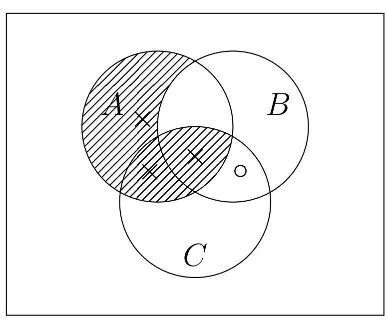

This final diagram gives the shape that all counterexamples to the valididity of 3.19 have in common. All these counterexamples will have no objects in the forbidden regions, and at least one object in the region marked with $\circ$.

Venn diagrams actually have a long history in logic, going back to the 18th century, and they are still an object of study in cognitive science, since they somehow combine visual and symbolic reasoning – a basic human ability that is not yet fully understood..

Exercise 3.2 Check the following syllogistism for validity, using the method just explained.

$$\begin{array}{l} \text{Some philosophers are Greek} \\ \text{No Greeks are barbarians} \\ \hline \text{No philosophers are barbarians.} \end{array} \tag{3.21}$$

Exercise 3.3 Check the following syllogistic pattern for validity.

$$\begin{array}{l} \text{No Greeks are barbarians} \\ \text{No barbarians are philosophers} \\ \hline \text{No Greeks are philosophers.} \end{array} \tag{3.22}$$

Exercise 3.4 Check the following syllogistic pattern for validity.

$$\begin{array}{l} \text{No Greeks are barbarians} \\ \text{Some barbarians are philosophers} \\ \hline \text{Not all philosophers are Greek.} \end{array} \tag{3.23}$$

Exercise 3.5 Can you modify the method so that it checks for syllogistic validity, but now with the quantifiers all read with existential import? How?

More than three predicates What follows is a digression for the interested reader. Venn diagrams were a high point of traditional logic, just before modern logic started. How far does this method take us?

The validity check for syllogistics can be extended to inferences with more than two premises (and more than three predicates). This can still be done graphically (Venn had several beautiful visualizations), but you may also want to think a bit more prosaically in terms of tabulating possibilities. Here is one way (disregarding matters of computational efficiency).

For purposes of exposition, assume that four predicates A, B, C, D occur in the inference. List all possible combinations in a table (compare the tables for the propositonal variables in Chapter 2 – we economized a bit here, writing the property only when it holds):

	A	B	AB
C	AC	BC	ABC
D	AD	BD	ABD
CD	ACD	BCD	ABCD

Take as example the following entailment

$$\text{All A are B}, \text{No C are B}, \text{Some C are D}, \text{Therefore: Not all D are A.} \tag{3.24}$$

Again we can use the update method to check whether this is valid. First update with the information that all A are B. This rules out certain possibilities:

	A ×	B	AB
C	AC ×	BC	ABC
D	AD ×	BD	ABD
CD	ACD ×	BCD	ABCD

All A are B

The information that no C are B also rules out possibilities, as follows:

	A	B	AB
C	AC	BC ×	ABC ×
D	AD	BD	ABD
CD	ACD	BCD ×	ABCD ×

No C are B

Combining these two updates, we get:

	A ×	B	AB
C	AC ×	BC ×	ABC ×
D	AD ×	BD	ABD
CD	ACD ×	BCD ×	ABCD ×

All A are B and No C are B

The third premise, "some C are D," is existential. It states that there has to at least one CD combination in the table. There is only one possibility for this:

	A ×	B	AB
C	AC ×	BC ×	ABC ×
D	AD ×	BD	ABD
CD ○	ACD ×	BCD ×	ABCD ×

Finally, we must check whether "not all D are A" holds in the table that results from updating with the premises. And indeed it does: region CD is non-empty (indicated by the presence of the ○), so it gives us a witness of a D which is not an A. Therefore, the given inference must be valid.

The syllogistic system as such Working through the exercises of this section you may have realized that the diagrammatic validity testing method can be applied to any syllogism, and that, in the terms of Chapter 2:

> The syllogistic is *sound* (only valid syllogism pass the test)
> and *complete* (all valid syllogisms pass the test).

Moreover, the method decides the question of validity in a matter of a few steps. Thus, again in our earlier terms:

> The syllogistic is a *decision method* for validity,

the system of the Syllogistic is 'decidable'. This is like what we saw for propositional logic, and indeed, it can be shown that the two systems are closely related, though we shall not do so here.

Much more can be said about the history of the syllogistic. The website of this course has an improved version of the Venn Diagram method due to Christie Ladd in 1882 which shows how it can be turned into a more efficient 'refutation method' when we picture the premises, but also the *negation* of the conclusion, and then try to spot a contradiction.

As usual, the rest of this chapter explores a few connections with other areas, starting with mathematical systems, then moving to computation, and ending with cognition. These topics are not compulsory in terms of understanding all their ins and outs, but they should help broaden your horizon.

3.6 Outlook — Satisfiability and Complexity

The tabling method for testing the validity of syllogisms suggests that the method behaves like the truth table method for propositional logic: if there are n properties, the method checks 2^n cases. For propositional logic it is an open question whether a non-exponential method exists for checking satisfiability: this is the famous *P* versus *NP* problem. But how about syllogistics? Can we do better than exponential?

Focussing on universal syllogistic forms only, it is easy to see that a set of universal forms is always satisfiable, provided we forget about existential import. The reason for this is that a situation with all classes empty will satisfy any universal form. Therefore:

> A set of syllogistic forms consisting of only universal statements is always satisfiable.

And, as a straightword consequence of this:

> A syllogism with only universal premises and an existential conclusion is always invalid.

The reason for this is that the situation with all classes empty is a counterexample: it will satisfy all the premisses but will falsify the existential conclusion.

If you reflect on this you see that the unsatisfiability of a set of syllogistic forms Σ is always due to the absence of witnesses for some existential forms $\psi_1, ..., \psi_n$ in Σ. Now, since the *number* of witnesses for a particular property does not matter – one witness for some property is as good as many – we can limit attention to situations where there is just a single object in the universe:

> A finite set of syllogistic forms Σ is unsatisfiable if and only if there exists an existential form ψ such that ψ taken together with the universal forms from Σ is unsatifiable.

The interesting thing is that this restricted form of satisfiability can easily be tested with propositional logic, as follows. Remember that we are talking about the properties of a single object x. Let proposition letter a express that object x has property A. Then a universal statement "all A are B" gets translated into $a \rightarrow b$: if x has property A then x also has property B. An existential statement "some A are B" gets translated into $a \wedge b$, expressing that x has both properties A and B. The universal negative statement "no A are B" gets translated into $a \rightarrow \neg b$, and the negative existential statement "some A are not B" gets translated as $a \wedge \neg b$. The nice thing about this translation is that it employs a single proposition letter for each property. No exponential blow-up here.

Note that to test the satisfiability of a set of syllogistic statements containing n existential statements we will need n tests: we have to check for each existential statement whether it is satisfiable when taken together with all universal statements. But this does not cause exponential blow-up if all these tests can be performed efficiently. We will show now that they can.

It may look like nothing is gained by our translation to propositional logic, since all known general methods for testing satisfiability of propositional logical formulas are exponential. But the remarkable thing is that our translation uses a very well-behaved fragment of propositional logic, for which satisfiability testing is easy.

In this outlook, we briefly digress to explain how propositional logic can be written in clausal form, and how satisfiability of clausal forms can be tested efficiently, provided the forms are in a 'nice' shape. Here are some definitions:

literals a literal is a proposition letter or its negation. If l is a literal, we use $\bar{l}$ for its negation: if l has the form p, then $\bar{l}$ equals $\neg p$, if l has the form $\neg p$, then $\bar{l}$ equals p. So if l is a literal, then $\bar{l}$ is also a literal, with opposite sign.

clause a clause is a set of literals.

clause sets a clause set is a set of clauses.

Read a clause as a *disjunction* of its literals, and a clause set as a *conjunction* of its clauses.

Here is an example: the clause form of

$$(p \to q) \wedge (q \to r)$$

is

$$\{\{\neg p, q\}, \{\neg q, r\}\}.$$

And here is an inference rule for clause sets called **Unit Propagation**:

Unit Propagation If one member of a clause set is a singleton $\{l\}$ (a 'unit'), then:

(1) remove every other clause containing l from the clause set (for since l has to be true, we know these other clauses have to be true as well, and no information gets lost by deleting them);

(2) remove $\bar{l}$ from every clause in which it occurs (for since l has to be true, we know that $\bar{l}$ has to be false, so no information gets lost by deleting $\bar{l}$ from any disjunction in which it occurs).

The result of applying this rule is an equivalent clause set. Example: applying unit propagation using unit $\{p\}$ to

$$\{\{p\}, \{\neg p, q\}, \{\neg q, r\}, \{p, s\}\}.$$

yields:

$$\{\{p\}, \{q\}, \{\neg q, r\}\}.$$

Applying unit propagation to this, using unit $\{q\}$ yields

$$\{\{p\}, \{q\}, \{r\}\}.$$

The *Horn fragment* of propositional logic consists of all clause sets where every clause has at most one positive literal. HORNSAT is the problem of checking Horn clause sets for satisfiability. This check can be performed in polynomial time (linear in the size of the formula, in fact).

If unit propagation yields a clause set in which units $\{l\}, \{\bar{l}\}$ occur, the original clause set is unsatisfiable, otherwise the units in the result determine a satisfying valuation. Recipe: for any units $\{l\}$ occurring in the final clause set, map their proposition letter to the truth value that makes l true; map all other proposition letters to false.

> The problem of testing satisfiability of syllogistic forms containing exactly one existential statement can be translated to the Horn fragment of propositional logic.

To see that this is true, check the translations we gave above:

All A are B $\mapsto$ $a \to b$ or equivalently $\{\{\neg a, b\}\}$.

No A are B $\mapsto$ $a \to \neg b$ or equivalently $\{\{\neg a, \neg b\}\}$.

Some A are B $\mapsto$ $a \wedge b$ or equivalently $\{\{a\}, \{b\}\}$.

Not all A are B $\mapsto$ $a \wedge \neg b$ or equivalently $\{\{a\}, \{\neg b\}\}$.

As you can see, these translations are all in the Horn fragment of propositional logic. We conclude that satisfiability of sets of syllogistic forms can be checked in time polynomial in the number of properties mentioned in the forms.

Exercise 3.6 ♠ Consider the following three syllogisms:

$$\begin{array}{c} \text{No } A \text{ are } B \\ \text{Not all } B \text{ are } C \\ \hline \text{Some } C \text{ are } A \end{array} \qquad \begin{array}{c} \text{No } A \text{ are } B \\ \text{Some } B \text{ are } C \\ \hline \text{Not all } C \text{ are } A \end{array} \qquad \begin{array}{c} \text{All } B \text{ are } A \\ \text{Some } C \text{ are } A \\ \hline \text{Some } A \text{ are } B \end{array}$$

(1) One of the three syllogisms is valid. Which one?

(2) Use the diagram method to show the validity of the syllogism you claim is valid.

(3) Use a diagram to show that the other syllogisms are invalid.

(4) Next, show, for these three cases, how the validity of the syllogisms can be checked by translating the premisses and the negation of the conclusion into clausal form, and then using unit propagation to check the resulting clause set for satisfiability. (Note: the clause set is satisfiable *iff* the syllogism is invalid.)

3.7 Outlook — The Syllogistic and Actual Reasoning

Aristotle's system is closely linked to the grammatical structure of natural language, as we have said at the start. Indeed, many people have claimed that it stays so close to our ordinary language that it is part of the *natural logic* that we normally use. Medieval logicians tried to extend this, and found further patterns of reasoning with quantifiers that share these same features of staying close to linguistic syntax, and allowing for very simple inference rules. 'Natural Logic' is a growing topic these days, where one tries to find large simple inferential subsystems of natural language that can be described without too much mathematical system complexity. Even so, we have to say that the real logical hitting power will only come in our next chapter on predicate logic, which consciously deviates from natural language to describe more complex quantifier reasoning of types that Aristotle did not handle.

Syllogistic reasoning has also drawn the attention of cognitive scientists, who try to draw conclusions about what goes on in the human brain when we combine predicates and reason about objects. As with propositional reasoning, one then finds differences in performance that do not always match what our methods say, calling attention to the issue how the brain represents objects and their properties and relations. From another point of view, the *diagrammatic* aspect of our methods has attracted attention from cognitive scientists lately. It is known that the brain routinely combines symbolic language-oriented and visual and diagrammatic representations, and the Venn Diagram method is one of the simplest pilot settings for studying how this combination works.

Summary *In this chapter you have learnt how one simple but very widespread kind of reasoning with predicates and quantifiers works. This places you squarely in a long logical tradition, before we move to the radical revolutions of the 19th century in our next chapter. More concretely, you are now able to*

- *write basic syllogistic forms for quantifiers,*
- *understand set diagram notation for syllogistic forms,*
- *test syllogistic inferences using Venn diagrams,*
- *understand how diagrams allow for update,*
- *understand connections with propositional logic,*
- *understand connections with data representation.*

Further Reading If you wish to be instructed in logic by the teacher of Alexander himself, you should consult the *Prior Analytics* [Ari89] (available online, in a different translation, at `classics.mit.edu/Aristotle/prior.html`). For a full textbook on Aristotelean logic, see [PH91].

Aristotelean logic can be viewed as a logic of concept description. See the first and second chapter [NB02, BN02] of the Description Logic Handbook [BCM$^+$02] for more information about this connection. Connections between Aristotelian logic and predicate logic (see next Chapter of this book) are discussed in [Łuk51]. Extensions of Aristotelian logic in the spirit of syllogistics are given in [PH04] and [Mos08].