MEDIEVAL SYLLOGISMS

LUKE COLLINS

CONTENTS

1. INTRODUCTION

A *syllogism* is a simple logical argument with two premisses and a conclusion, whose origins trace back to antiquity, in the writings of Aristotle. A famous example is the following.

$$\frac{\begin{array}{c}\text{All men are mortal.}\\ \text{All Greeks are men.}\end{array}}{\therefore \quad \text{All Greeks are mortal.}}$$

Despite its simplicity, we have not yet developed enough logical theory to be able to describe this in symbols. So far, we have seen the logical junctors $\neg, \wedge, \vee, \rightarrow, \leftrightarrow$; but to formalise syllogisms, we need to upgrade our logical framework by introducing two new symbols called *quantifiers*. The logical theory we have developed so far is usually referred to as propositional or *zeroth order* logic. Equipped with the two symbols we will describe in these notes, our theory graduates to the so-called *predicate* or *first order* logic.

1.1. **Predicates.** A *predicate* captures the idea of a proposition with "holes". For instance, consider the propositions

(i) *4 is an even number*,
(ii) *8 is an even number*,
(iii) *5 is an even number*.

Date: 15th June, 2020.

We can identify these as different instances of the predicate

$$\mathfrak{p}(x) = \text{``}x \text{ is an even number''},$$

where x is a "hole" we plug things into; we call x a *variable*. Indeed, we would write (i) as $\mathfrak{p}(4)$, (ii) as $\mathfrak{p}(8)$ and (iii) as $\mathfrak{p}(5)$. Notice that for different values of x, the predicate is sometimes true, and sometimes false. The distinction between a proposition and a predicate is that a proposition has no variables; for instance, "Today is a Monday" is a proposition, whereas "x is a Monday" is a predicate.

1.2. **Quantifiers.** Even though, in the general case, a predicate's truth-value (i.e., whether it is true or false) depends on what we substitute for its variable, sometimes we want to make claims such as "there is at least one x such that $\mathfrak{p}(x)$ is true" or "for every x, $\mathfrak{p}(x)$ is false". This leads us to our two new symbols. The first is the so-called *universal quantifier*, $\bigwedge_x$, which we write as a prefix to a predicate $\mathfrak{p}(x)$ as

$$\bigwedge\nolimits_x \mathfrak{p}(x),$$

and read as "for all x, $\mathfrak{p}(x)$". For example, if we are talking about flowers in a garden, and the predicate $\mathfrak{p}(x)$ is "x has red petals", then $\bigwedge_x \mathfrak{p}(x)$ is "for all x, x has red petals", or stated more simply, "all flowers have red petals".

Next we have the *existential quantifier*, $\bigvee_x$, which we write as a prefix to a predicate $\mathfrak{p}(x)$ as

$$\bigvee\nolimits_x \mathfrak{p}(x),$$

and read as "for some x, $\mathfrak{p}(x)$" or (equivalently) as "there exists at least one x such that $\mathfrak{p}(x)$". For instance, say we are talking about whole numbers, and that $\mathfrak{p}(x)$ is "x is prime and even", then $\bigvee_x \mathfrak{p}(x)$ is "for some x, x is prime and even", or simply, "some number is prime and even".

Let us make a few observations:

(i) Notice that when we say "for all x" or "for some x", we are being ambiguous unless we specify what values x is allowed to take on. Notice in the examples we gave, we always provided context about what x can be; e.g., whether x is confined to flowers in a particular garden, or socks in a drawer, or the whole numbers.

There is much more nuanced theory which goes into this (namely the theory of sets), but we will not go into it here. For our purposes, we will always assume that what x is allowed to be can be inferred from the context.

(ii) The symbols $\bigwedge_x$ and $\bigvee_x$ resemble $\wedge$ and $\vee$: this is not unintentional. When x ranges over a finite number of possibilities, then we can interpret these quantifiers as repeated iterations of $\wedge$ and $\vee$ respectively. For instance, suppose we are talking about the chairs in a classroom, and suppose $\mathfrak{a}(x)$ is "x is made of wood" and $\mathfrak{b}(x)$ is "x is new". Then $\bigwedge_x \mathfrak{a}(x)$, i.e., all the chairs in this room are made of wood, is equivalent to saying "this chair is made of wood $\wedge$ that chair is made of wood $\wedge \cdots$". Similarly, $\bigvee_x \mathfrak{b}(x)$, i.e., some chair in this room is new, is equivalent to saying "this chair is new $\vee$ that chair is new $\vee \cdots$."

(iii) When a predicate is quantified, i.e., when one of $\bigwedge_x$ or $\bigvee_x$ is placed in front of a predicate $\mathfrak{p}(x)$, the result is no longer a predicate, but becomes a proposition. This is because there is no longer a variable which can be substituted for. For example, if $\mathfrak{a}(x)$ is "I rolled an x on the dice", then we

can substitute a value for x in $\mathfrak{a}(x)$ (e.g., $\mathfrak{a}(4)$ is "I rolled a 4 on the dice"), but

$$\bigvee_x \mathfrak{a}(x) \text{ is } \mathfrak{a}(1) \vee \mathfrak{a}(2) \vee \mathfrak{a}(3) \vee \mathfrak{a}(4) \vee \mathfrak{a}(5) \vee \mathfrak{a}(6),$$

i.e., I rolled a 1 or I rolled a 2 or, ..., or I rolled a 6,

which contains no variable. In this context, x is called a *dummy variable*, because it appears in the notation $\bigvee_x \mathfrak{a}(x)$ but there isn't actually a variable there.

2. Beth Tableaux

The Beth tableaux rules for quantifiers are quite straightforward. We outline the reasoning behind them in the following points.

- If $\bigwedge_x \mathfrak{p}(x)$ is true, then $\mathfrak{p}(x)$ is true for all x, so we can plug in anything we want for x in $\mathfrak{p}(x)$ and it will still be true. This leads to the rule

$$\frac{\Sigma_1(\bigwedge_x \mathfrak{p}(x)) \,\Big\|\qquad \Sigma_2}{\mathfrak{p}(n) \qquad\Big\|}$$

  where there is no restriction on n.

- If $\bigvee_x \mathfrak{p}(x)$ is true, then there is some value $x = n$ such that $\mathfrak{p}(n)$ is true. We don't know what n is, so we can't assume it is some value of our choosing. In particular, we insist that n is not a letter which has already appeared in the Beth tableau so far (this does not exclude the possibility that it equals some of the other letters, but it doesn't imply it either). Thus we have the rule

$$\frac{\Sigma_1(\bigvee_x \mathfrak{p}(x)) \,\Big\|\qquad \Sigma_2}{\mathfrak{p}(n) \qquad\Big\|} \qquad (\text{where } n \text{ is new})$$

- If $\bigwedge_x \mathfrak{p}(x)$ is false, it means that it is not the case that $\mathfrak{p}(x)$ is true *for each* x, so in particular, there must be at least one x, say $x = n$, for which $\mathfrak{p}(n)$ is false. Therefore, just as in the previous rule, we insist that n is a new variable name so that it is not identified with any which have appeared so far in the Beth tableau, obtaining the rule

$$\frac{\Sigma_1 \qquad\Big\|\; \Sigma_2(\bigwedge_x \mathfrak{p}(x))}{\Big\|\quad \mathfrak{p}(n)} \qquad (\text{where } n \text{ is new})$$

- Finally, if $\bigvee_x \mathfrak{p}(x)$ is false, then there does not exist an x such that $\mathfrak{p}(x)$ is true, in other words, $\mathfrak{p}(x)$ is false for any x we choose. Thus we get the rule

$$\frac{\Sigma_1 \qquad\Big\|\; \Sigma_2(\bigvee_x \mathfrak{p}(x))}{\Big\|\quad \mathfrak{p}(n)}$$

  where there is no restriction on n.

Example. If $\mathfrak{a}(x) \vee \mathfrak{b}(x)$ is true for some x, then at least one of $\mathfrak{a}(x)$ is true, or one of $\mathfrak{b}(x)$ is true. In symbols,

$$\bigvee_x(\mathfrak{a}(x) \vee \mathfrak{b}(x)) < \left(\bigvee_x \mathfrak{a}(x)\right) \vee \left(\bigvee_x \mathfrak{b}(x)\right).$$

As usual, we start by drawing two columns, placing the premiss on the left and the conclusion on the right.

$$\frac{\bigvee_x(\mathfrak{a}(x) \vee \mathfrak{b}(x)) \,\Big\|\, \big(\bigvee_x \mathfrak{a}(x)\big) \vee \big(\bigvee_x \mathfrak{b}(x)\big)}{}$$

Now we seem to have a lot of options as to the rules we can apply here. As usual, we want to avoid branching early on, but an additional good rule of thumb is to use rules which require "new" variables before we use the rules with no restrictions. Let's start with $\Sigma_2(\,\cdot\,\vee\,\cdot\,)$:

$$\frac{\bigvee_x(\mathfrak{a}(x) \vee \mathfrak{b}(x)) \,\Big\|\, \big(\bigvee_x \mathfrak{a}(x)\big) \vee \big(\bigvee_x \mathfrak{b}(x)\big)}{\begin{array}{c} \bigvee_x \mathfrak{a}(x) \\ \bigvee_x \mathfrak{b}(x) \end{array}}$$

Now at this stage, we can invoke our law $\Sigma_2(\bigvee_x\,\cdot\,)$ on either of the two new lines, but first notice that we can apply $\Sigma_1(\bigvee_x\,\cdot\,)$, and that the latter insists that we have a "new" variable n, so let's use that first, in accordance with our rule of thumb:

$$\frac{\bigvee_x(\mathfrak{a}(x) \vee \mathfrak{b}(x)) \,\Big\|\, \big(\bigvee_x \mathfrak{a}(x)\big) \vee \big(\bigvee_x \mathfrak{b}(x)\big)}{\begin{array}{c} \\ \mathfrak{a}(n) \vee \mathfrak{b}(n) \end{array}}$$

Now we invoke $\Sigma_2(\bigvee_x\,\cdot\,)$, and since that law poses no restrictions on the variable we can use, let's use the same letter n:

Finally, branching the $\vee$ (i.e., using $\Sigma_1(\,\cdot\,\vee\,\cdot\,)$), closes the tableau:

which completes the proof. $\qquad\square$

Example (Negation of Quantified Statements). If $\mathfrak{a}(x)$ is not true for all x, then there is some x such that $\mathfrak{a}(x)$ is false, i.e., such that $\neg\mathfrak{a}(x)$ is true. In symbols,

$$\neg\big(\textstyle\bigwedge_x \mathfrak{a}(x)\big) \asymp \textstyle\bigvee_x \neg\mathfrak{a}(x)$$

Similarly,

$$\neg\big(\textstyle\bigvee_x \mathfrak{a}(x)\big) \asymp \textstyle\bigwedge_x \neg\mathfrak{a}(x).$$

We provide Beth tableaux proving the first equivalence.

Which proves the $<$ direction, next we have

for the $>$ direction. Since both implications are valid, we conclude that the equivalence is valid. □

Exercises. (i) Prove the reverse implication of the first example:

$$\big(\textstyle\bigvee_x \mathfrak{a}(x)\big) \vee \big(\textstyle\bigvee_x \mathfrak{b}(x)\big) < \textstyle\bigvee_x(\mathfrak{a}(x) \vee \mathfrak{b}(x)).$$

(ii) Prove the following analogous equivalence for the universal quantifier:

$$\textstyle\bigwedge_x(\mathfrak{a}(x) \wedge \mathfrak{b}(x)) \asymp \big(\textstyle\bigwedge_x \mathfrak{a}(x)\big) \wedge \big(\textstyle\bigwedge_x \mathfrak{b}(x)\big).$$

(iii) Prove the analogous equivalence for negating existentially quantified statements, i.e.,

$$\neg\big(\textstyle\bigvee_x \mathfrak{a}(x)\big) \asymp \textstyle\bigwedge_x \neg\mathfrak{a}(x).$$

(iv) Show that

$$\textstyle\bigwedge_x(\mathfrak{p}(x) \to \neg\mathfrak{m}(x)) \,,\, \textstyle\bigvee_x(\mathfrak{s}(x) \wedge \mathfrak{m}(x)) < \textstyle\bigvee_x(\mathfrak{s}(x) \wedge \neg\mathfrak{p}(x))$$

is valid.

3. Aristotelian Propositions

In this section, we treat four specific instances of quantified statements which are of historical significance. These are known as the *Aristotelian propositions* (or sometimes, the *categorical* propositions). These are:

- a. All S are P. (Universal affirmative)
- e. No S is P. (Universal negative)
- i. Some S is P. (Particular affirmative)

o. Some S is not P.   (Particular negative)

Usually S is called the *subject* and P is called the *predicate*,[1] and both are called the *terms* of the proposition. The use of the vowels A, E, I, O as labels for these four propositions comes from the corresponding Latin words for "I affirm" (corresponding to the first and third) and "I deny" (corresponding to the second and fourth); we have **a**ffirmo, n**e**go, aff**i**rmo, and neg**o**.

Some more terminology: the adjectives universal/particular, when applied to an Aristotelian proposition, are referred to as the *quantity* of that proposition; whereas the adjectives affirmative/negative are referred to as its *quality*.

In modern symbols, if $\mathfrak{s}(x)$ and $\mathfrak{p}(x)$ are the subject and predicate respectively, the four propositions are

a. $\bigwedge_x(\mathfrak{s}(x) \to \mathfrak{p}(x))$.   (SaP)
e. $\bigwedge_x(\mathfrak{s}(x) \to \neg\mathfrak{p}(x))$.   (SeP)
i. $\bigvee_x(\mathfrak{s}(x) \wedge \mathfrak{p}(x))$.   (SiP)
o. $\bigvee_x(\mathfrak{s}(x) \wedge \neg\mathfrak{p}(x))$.   (SoP)

We use the abbreviations SaP, SeP, SiP and SoP as shorthand names for these.

So for instance, if $\mathfrak{s}(x)$ is "x is a (hu)man" and $\mathfrak{p}(x)$ is "x is white", then the four corresponding Aristotelian propositions are:

a. All men are white.   (SaP)
e. All men are not white.   (SeP)
i. Some men are white.   (SiP)
o. Some men are not white.   (SoP)

3.1. **Vacuous Truths.** If $\mathfrak{t}(x)$ is a term in an Aristotelian proposition, we say that $\mathfrak{t}(x)$ is *occupied* or *non-empty* if $\bigvee_x \mathfrak{t}(x)$ is true, and *empty* otherwise (i.e., if $\bigvee_x \mathfrak{t}(x)$ is false). Although, intuitively, a "for all" statement is generally a stronger assertion than a "for some" one, the former does not necessarily imply the latter. For example, the statement

"All dinosaurs alive today can fly"

(i.e., DaF) is true, since there are no dinosaurs alive today. In particular, it is *less* strong than the statement

"Some dinosaur alive today can fly"

(i.e., DiF), which implies that some dinosaur alive today exists.

It might seem strange that the first statement is true, but perhaps this special case is better understood when we represent the statements symbolically. If we let $\mathfrak{d}(x)$ be "x is a dinosaur alive today" and $\mathfrak{f}(x)$ be "x can fly", then the two statements we have are

$$\bigwedge_x \mathfrak{d}(x) \to \mathfrak{f}(x) \qquad \text{and} \qquad \bigvee_x \mathfrak{d}(x) \wedge \mathfrak{f}(x).$$

Here it is more clear, we see that since $\mathfrak{d}(x)$ is never true for any x, the "if...then" in the first statement is automatically true for any x, which is why the for all statement is true, without ever needing $\mathfrak{d}(x)$ to be true itself. On the other hand, the second statement implies that $\mathfrak{d}(x)$ is true for at least one x.

---

[1]This terminology is a bit confusing, but it is what was used historically. Nowadays, in most contexts of logic, a predicate is any proposition $\mathfrak{a}(x)$ in which a variable occurs, but in the context of Aristotelian propositions, it refers specifically to the second term.

Conversion Type				Conditions
simple	SeP	$\rightarrow$	PeS	
	SiP	$\rightarrow$	PiS	
accidental	SaP	$\rightarrow$	PiS	S is occupied
	SeP	$\rightarrow$	PoS	P is occupied

TABLE 1. Conversions between Aristotelian Propositions

We can also look at it this way: for any statement $\mathfrak{p}$, either $\mathfrak{p}$ is true, or its negation $\neg\mathfrak{p}$ is. In the case of our "for all" statement, we saw in a previous exercise that

$$\neg\left(\bigwedge_x \mathfrak{d}(x) \rightarrow \mathfrak{f}(x)\right) \asymp \bigvee_x \neg(\mathfrak{d}(x) \rightarrow \mathfrak{f}(x)),$$

and it can easily be seen (by constructing a Beth tableau, say) that $\neg(\mathfrak{a} \rightarrow \mathfrak{b})$ is equivalent to $\mathfrak{a} \wedge \neg\mathfrak{b}$, so we therefore have that

$$\neg\left(\bigwedge_x \mathfrak{d}(x) \rightarrow \mathfrak{f}(x)\right) \asymp \bigvee_x \mathfrak{d}(x) \wedge \neg\mathfrak{f}(x)$$

(i.e., $\neg\mathsf{DaF} \asymp \mathsf{DoF}$), which in words, translates to

"There is some x such that x is a dinosaur alive today, and x cannot fly",

or more simply, "Some dinosaur alive today cannot fly". This is clearly false, since it implies the existence of a dinosaur alive today, so we are forced to say that its negation "All dinosaurs alive today can fly" must be true.

Such "for all" statements are said to be *vacuously true*, since they are true solely because the subject is empty.

3.2. **Conversion.** An Aristotelian proposition tells us something about its subject. For instance, "All men are white" tells us something about men (i.e., about those x which satisfy $\mathfrak{s}(x)$ where $\mathfrak{s}(x)$ is "x is a man").

If we are given a true Aristotelian proposition about a term $\mathfrak{s}(x)$, can we obtain an Aristotelian proposition about its predicate $\mathfrak{p}(x)$? In other words, can we obtain another Aristotelian proposition where the roles of the terms are interchanged? It turns out the answer is almost always yes—we call this process *conversion*. Table 1 summarises the possible conversions for Aristotelian propositions. Simple conversion simply interchanges the subject and predicate, whereas accidental conversion applies only to universal ($\bigwedge_x$) formulæ and changes them to particular ($\bigvee_x$), without changing their quality. Accidental conversion is subject to conditions on S and P, as summarised in the table.

Exercises. (i) Use Beth Tableaux to prove the validity of the simple conversion rules $\mathsf{SeP} \rightarrow \mathsf{PeS}$ and $\mathsf{SiP} \rightarrow \mathsf{PiS}$, i.e., show that

$$\bigwedge_x(\mathfrak{s}(x) \rightarrow \neg\mathfrak{p}(x)) < \bigwedge_x(\mathfrak{p}(x) \rightarrow \neg\mathfrak{s}(x))$$

and

$$\bigvee_x(\mathfrak{s}(x) \wedge \mathfrak{p}(x)) < \bigvee_x(\mathfrak{p}(x) \wedge \mathfrak{s}(x))$$

are valid.

(ii) Write out the two accidental conversion rules as implications. (Be careful, remember there are conditions.) Hence, prove that they are both valid using Beth tableaux.

3.3. **The Square of Opposition.** We have the following terminology to describe relationships among propositions.

   (i) $\mathfrak{a}$ is *superaltern* to $\mathfrak{b}$ if $\mathfrak{a} < \mathfrak{b}$ is valid.
  (ii) $\mathfrak{a}$ is *subaltern* to $\mathfrak{b}$ if $\mathfrak{b} < \mathfrak{a}$ is valid.
 (iii) $\mathfrak{a}$ is *contrary* to $\mathfrak{b}$ means that $\mathfrak{a} < \neg\mathfrak{b}$ is valid.
  (iv) $\mathfrak{a}$ is *subcontrary* to $\mathfrak{b}$ means that $\neg\mathfrak{a} < \mathfrak{b}$ is valid.
   (v) $\mathfrak{a}$ is *contradictory* to $\mathfrak{b}$ means that $\mathfrak{a}$ is both contrary and subcontrary to $\mathfrak{b}$.

These terms are enough to characterise the relationships between each of the four Aristotelian propositions, giving us the so-called *square of opposition*, illustrated in figure 1. The grey arrows require S to be occupied in order to travel along. The


FIGURE 1. The Square of Opposition

black arrows are unconditionally true.

In summary, the square tells us that:

   (i) SaP and SoP are contradictories, and so are SiP and SeP, i.e., if one is true, the other is false, and vice-versa,

and, on the condition that S is occupied,

  (ii) SaP and SeP are contraries, i.e., if one is true, the other is false,
 (iii) SiP and SoP are subcontraries, i.e., if one is false, the other is true,
  (iv) SiP is subaltern to SaP, and SoP is subaltern to SeP.

Exercises. Use Beth tableaux to prove the claims of the square of opposition.

4. Syllogisms

A syllogism is a special type of implication, made up of two Aristotelian propositions as premises, and one Aristotelian proposition as a conclusion. The two premisses share one term, called the *middle term*, the other two terms in them are respectively the subject and predicate of the conclusion. Here is an example:

$$\frac{\begin{array}{l} \text{All men are mortal.} \\ \text{All Greeks are men.} \end{array}}{\therefore \quad \text{All Greeks are mortal.}}$$

In the above instance, we have the subject $\mathfrak{s}(x) =$ "x is Greek", the middle term $\mathfrak{m}(x) =$ "x is a man" and the predicate $\mathfrak{p}(x) =$ "x is mortal". It is conventional to order the premisses so that the premiss made up of the subject of the conclusion and the middle term comes first, and the one made up of the middle term, and the predicate of the conclusion is second; as we have done above.

With this convention in mind, there are two things which determine a syllogism:

(i) *The figure.* The middle term can be either the subject or the predicate of either of the premisses, which gives us four possibilities. The corresponding possibility is called the figure of the syllogism, and we associate each with a number, as indicated in figure 2. Focusing on the position of the middle

|										
M	P		**M**	P		P	**M**		P	**M**
M	S		S	**M**		**M**	S		S	**M**
∴ S	P		∴ S	P		∴ S	P		∴ S	P
Figure 3			Figure 1			Figure 4			Figure 2	

FIGURE 2. Different figures of a syllogism

terms, they form the shape of the letter M, which gives us the mnemonic in figure 3 for the different possible figures.


FIGURE 3. Mnemonic for the different figures of a Syllogism

(ii) *The mood,* i.e., the type of formula (a, e, i, o) of the first premiss, the second premiss and of the conclusion. There are thus $4^3 = 64$ possible moods for each figure.

It turns out that for each figure, precisely six of all the possible moods are valid, however some of them require that a term be occupied. Medieval logicians assigned mnemonics to each valid syllogism: the mood is given by extracting the first three vowels in the name in order, corresponding to the first premiss, second premiss and the conclusion.

We list the syllogisms corresponding to each figure below, denoting any terms which are required to be occupied in brackets after the name where necessary.

Figure 1. *Barbara, Barbari* (S), *Celarent, Celaront* (S), *Darii, Ferio.*
Figure 2. *Cesare, Cesaro* (S), *Camestres, Camestros* (S), *Festino, Baroco.*
Figure 3. *Darapti* (M), *Felapton* (M), *Disamis, Datisi, Bocardo, Ferison.*
Figure 4. *Bamalipton* (P), *Camentes, Camentos* (S), *Dimatis, Fesapo* (M), *Fresison.*

So for instance, we recognise that our example with men, Greeks and mortals is an instance of the Barbara syllogism. Another example: Camestres is the following syllogism.

$$\begin{array}{c} \mathsf{PaM} \\ \dfrac{\mathsf{SeM}}{\therefore\quad \mathsf{SeP}} \end{array} \qquad \text{which is, in words,} \qquad \begin{array}{c} \text{All } \mathsf{P} \text{ are } \mathsf{M} \\ \dfrac{\text{No } \mathsf{S} \text{ is } \mathsf{M}}{\therefore\quad \text{No } \mathsf{S} \text{ is } \mathsf{P}} \end{array}$$

e.g.,

$$\begin{array}{c} \text{All men are mortal} \\ \dfrac{\text{No god is mortal}}{\therefore\quad \text{No god is a man}} \end{array}$$

Written out formally, Camestres is the implication

$$\bigwedge\nolimits_x \mathfrak{p}(x) \to \mathfrak{m}(x) \,,\, \bigwedge\nolimits_x \mathfrak{s}(x) \to \neg\mathfrak{m}(x) < \bigwedge\nolimits_x \mathfrak{s}(x) \to \neg\mathfrak{p}(x),$$

which we can prove by Beth tableau.

References

[1] Riolo, V. *Introduction to Logic*, second ed. Malta University Press, 2001.
[2] Smullyan, R. M. *First-Order Logic*, first ed. Springer–Verlag, 1968.

D2.08, Mathematical Institute, University of Warwick, Coventry, CV4 7AL, UK
Email address: luke.collins@warwick.ac.uk
URL: https://drmenguin.com