

Chapter udf

Syntax of First-Order Logic

syn.1 Introduction

fol:syn:itx:
sec In order to develop the theory and metatheory of first-order logic, we must first define the syntax and semantics of its expressions. The expressions of first-order logic are terms and **formulas**. Terms are formed from **variables**, **constant symbols**, and **function symbols**. **Formulas**, in turn, are formed from **predicate symbols** together with terms (these form the smallest, “atomic” **formulas**), and then from atomic **formulas** we can form more complex ones using logical connectives and quantifiers. There are many different ways to set down the formation rules; we give just one possible one. Other systems will choose different symbols, will select different sets of connectives as primitive, will use parentheses differently (or even not at all, as in the case of so-called Polish notation). What all approaches have in common, though, is that the formation rules define the set of terms and **formulas** *inductively*. If done properly, every expression can result essentially in only one way according to the formation rules. The inductive definition resulting in expressions that are *uniquely readable* means we can give meanings to these expressions using the same method—inductive definition.

syn.2 First-Order Languages

fol:syn:fol:
sec Expressions of first-order logic are built up from a basic vocabulary containing *variables*, *constant symbols*, *predicate symbols* and sometimes *function symbols*. From them, together with logical connectives, quantifiers, and punctuation symbols such as parentheses and commas, *terms* and *formulas* are formed.

Informally, **predicate symbols** are names for properties and relations, **constant symbols** are names for individual objects, and **function symbols** are names for mappings. These, except for the **identity predicate** $=$, are the *non-logical symbols* and together make up a language. Any first-order language \mathcal{L} is determined by its non-logical symbols. In the most general case, \mathcal{L} contains infinitely many symbols of each kind. explanation

In the general case, we make use of the following symbols in first-order logic:

1. Logical symbols

- a) Logical connectives: \neg (negation), \wedge (conjunction), \vee (disjunction), \rightarrow (conditional), \leftrightarrow (biconditional), \forall (universal quantifier), \exists (existential quantifier).
- b) The propositional constant for falsity \perp .
- c) The propositional constant for truth \top .
- d) The two-place identity predicate $=$.
- e) A denumerable set of variables: v_0, v_1, v_2, \dots

2. Non-logical symbols, making up the *standard language* of first-order logic

- a) A denumerable set of n -place predicate symbols for each $n > 0$: $A_0^n, A_1^n, A_2^n, \dots$
- b) A denumerable set of constant symbols: c_0, c_1, c_2, \dots
- c) A denumerable set of n -place function symbols for each $n > 0$: $f_0^n, f_1^n, f_2^n, \dots$

3. Punctuation marks: $(,)$, and the comma.

Most of our definitions and results will be formulated for the full standard language of first-order logic. However, depending on the application, we may also restrict the language to only a few predicate symbols, constant symbols, and function symbols.

Example syn.1. The language \mathcal{L}_A of arithmetic contains a single two-place predicate symbol $<$, a single constant symbol 0 , one one-place function symbol $!$, and two two-place function symbols $+$ and \times .

Example syn.2. The language of set theory \mathcal{L}_Z contains only the single two-place predicate symbol \in .

Example syn.3. The language of orders \mathcal{L}_{\leq} contains only the two-place predicate symbol \leq .

Again, these are conventions: officially, these are just aliases, e.g., $<$, \in , and \leq are aliases for A_0^2 , 0 for c_0 , $!$ for f_0^1 , $+$ for f_0^2 , \times for f_1^2 .

[intro](#) You may be familiar with different terminology and symbols than the ones we use above. Logic texts (and teachers) commonly use \sim , \neg , or $!$ for “negation”, \wedge , \cdot , or $\&$ for “conjunction”. Commonly used symbols for the “conditional” or “implication” are \rightarrow , \Rightarrow , and \supset . Symbols for “biconditional,” “bi-implication,” or “(material) equivalence” are \leftrightarrow , \Leftrightarrow , and \equiv . The \perp symbol is variously called “falsity,” “falsum,” “absurdity,” or “bottom.” The \top symbol is variously called “truth,” “verum,” or “top.”

It is conventional to use lower case letters (e.g., a, b, c) from the beginning of the Latin alphabet for **constant symbols** (sometimes called names), and lower case letters from the end (e.g., x, y, z) for **variables**. Quantifiers combine with **variables**, e.g., x ; notational variations include $\forall x, (\forall x), (x), \Pi x, \bigwedge_x$ for the universal quantifier and $\exists x, (\exists x), (Ex), \Sigma x, \bigvee_x$ for the existential quantifier.

We might treat all the propositional operators and both quantifiers as primitive symbols of the language. We might instead choose a smaller stock of primitive symbols and treat the other **logical operators** as defined. “Truth functionally complete” sets of Boolean operators include $\{\neg, \vee\}$, $\{\neg, \wedge\}$, and $\{\neg, \rightarrow\}$ —these can be combined with either quantifier for an expressively complete first-order language. explanation

You may be familiar with two other **logical operators**: the Sheffer stroke $|$ (named after Henry Sheffer), and Peirce’s arrow \downarrow , also known as Quine’s dagger. When given their usual readings of “nand” and “nor” (respectively), these operators are truth functionally complete by themselves.

syn.3 Terms and Formulas

fol:syn:frm: sec Once a first-order language \mathcal{L} is given, we can define expressions built up from the basic vocabulary of \mathcal{L} . These include in particular *terms* and *formulas*.

fol:syn:frm: defn:terms **Definition syn.4 (Terms).** The set of *terms* $\text{Trm}(\mathcal{L})$ of \mathcal{L} is defined inductively by:

1. Every **variable** is a term.
2. Every **constant symbol** of \mathcal{L} is a term.
3. If f is an n -place **function symbol** and t_1, \dots, t_n are terms, then $f(t_1, \dots, t_n)$ is a term.
4. Nothing else is a term.

A term containing no **variables** is a *closed term*.

The **constant symbols** appear in our specification of the language and the terms as a separate category of symbols, but they could instead have been included as zero-place **function symbols**. We could then do without the second clause in the definition of terms. We just have to understand $f(t_1, \dots, t_n)$ as just f by itself if $n = 0$. explanation

fol:syn:frm: defn:formulas **Definition syn.5 (Formulas).** The set of *formulas* $\text{Frm}(\mathcal{L})$ of the language \mathcal{L} is defined inductively as follows:

1. \perp is an atomic **formula**.
2. \top is an atomic **formula**.
3. If R is an n -place **predicate symbol** of \mathcal{L} and t_1, \dots, t_n are terms of \mathcal{L} , then $R(t_1, \dots, t_n)$ is an atomic **formula**.

4. If t_1 and t_2 are terms of \mathcal{L} , then $=(t_1, t_2)$ is an atomic formula.
5. If φ is a formula, then $\neg\varphi$ is a formula.
6. If φ and ψ are formulas, then $(\varphi \wedge \psi)$ is a formula.
7. If φ and ψ are formulas, then $(\varphi \vee \psi)$ is a formula.
8. If φ and ψ are formulas, then $(\varphi \rightarrow \psi)$ is a formula.
9. If φ and ψ are formulas, then $(\varphi \leftrightarrow \psi)$ is a formula.
10. If φ is a formula and x is a variable, then $\forall x \varphi$ is a formula.
11. If φ is a formula and x is a variable, then $\exists x \varphi$ is a formula.
12. Nothing else is a formula.

explanation

The definitions of the set of terms and that of formulas are *inductive definitions*. Essentially, we construct the set of formulas in infinitely many stages. In the initial stage, we pronounce all atomic formulas to be formulas; this corresponds to the first few cases of the definition, i.e., the cases for \top , \perp , $R(t_1, \dots, t_n)$ and $=(t_1, t_2)$. “Atomic formula” thus means any formula of this form.

The other cases of the definition give rules for constructing new formulas out of formulas already constructed. At the second stage, we can use them to construct formulas out of atomic formulas. At the third stage, we construct new formulas from the atomic formulas and those obtained in the second stage, and so on. A formula is anything that is eventually constructed at such a stage, and nothing else.

By convention, we write $=$ between its arguments and leave out the parentheses: $t_1 = t_2$ is an abbreviation for $=(t_1, t_2)$. Moreover, $\neg=(t_1, t_2)$ is abbreviated as $t_1 \neq t_2$. When writing a formula $(\psi * \chi)$ constructed from ψ , χ using a two-place connective $*$, we will often leave out the outermost pair of parentheses and write simply $\psi * \chi$.

intro

Some logic texts require that the variable x must occur in φ in order for $\exists x \varphi$ and $\forall x \varphi$ to count as formulas. Nothing bad happens if you don’t require this, and it makes things easier.

If we work in a language for a specific application, we will often write two-place predicate symbols and function symbols between the respective terms, e.g., $t_1 < t_2$ and $(t_1 + t_2)$ in the language of arithmetic and $t_1 \in t_2$ in the language of set theory. The successor function in the language of arithmetic is even written conventionally *after* its argument: t' . Officially, however, these are just conventional abbreviations for $A_0^2(t_1, t_2)$, $f_0^2(t_1, t_2)$, $A_0^2(t_1, t_2)$ and $f_0^1(t)$, respectively.

Definition syn.6 (Syntactic identity). The symbol \equiv expresses syntactic identity between strings of symbols, i.e., $\varphi \equiv \psi$ iff φ and ψ are strings of symbols of the same length and which contain the same symbol in each place.

The \equiv symbol may be flanked by strings obtained by concatenation, e.g., $\varphi \equiv (\psi \vee \chi)$ means: the string of symbols φ is the same string as the one obtained by concatenating an opening parenthesis, the string ψ , the \vee symbol, the string χ , and a closing parenthesis, in this order. If this is the case, then we know that the first symbol of φ is an opening parenthesis, φ contains ψ as a substring (starting at the second symbol), that substring is followed by \vee , etc.

As terms and **formulas** are built up from basic elements via inductive definitions, we can use the following induction principles to prove things about them.

fol:syn:frm: **Lemma syn.7 (Principle of induction on terms).** *lem:trmind* Let \mathcal{L} be a first-order language. If some property P is such that

1. it holds for every **variable** v ,
2. it holds for every **constant symbol** a of \mathcal{L} , and
3. it holds for $f(t_1, \dots, t_n)$ whenever it holds for t_1, \dots, t_n and f is an n -place **function symbol** of \mathcal{L}

(assuming t_1, \dots, t_n are terms of \mathcal{L}), then P holds for every term in $\text{Trm}(\mathcal{L})$.

Problem syn.1. Prove **Lemma syn.7**.

fol:syn:frm: **Lemma syn.8 (Principle of induction on formulas).** *thm:frmind* Let \mathcal{L} be a first-order language. If some property P holds for all the **atomic formulas** and is such that

1. it holds for $\neg\varphi$ whenever it holds for φ ;
2. it holds for $(\varphi \wedge \psi)$ whenever it holds for φ and ψ ;
3. it holds for $(\varphi \vee \psi)$ whenever it holds for φ and ψ ;
4. it holds for $(\varphi \rightarrow \psi)$ whenever it holds for φ and ψ ;
5. it holds for $(\varphi \leftrightarrow \psi)$ whenever it holds for φ and ψ ;
6. it holds for $\exists x \varphi$ whenever it holds for φ ;
7. it holds for $\forall x \varphi$ whenever it holds for φ ;

(assuming φ and ψ are **formulas** of \mathcal{L}), then P holds for all formulas in $\text{Frm}(\mathcal{L})$.

syn.4 Unique Readability

explanation The way we defined **formulas** guarantees that every **formula** has a *unique reading*, i.e., there is essentially only one way of constructing it according to our formation rules for **formulas** and only one way of “interpreting” it. If this were not so, we would have ambiguous **formulas**, i.e., **formulas** that have more than one reading or interpretation—and that is clearly something we want to avoid. But more importantly, without this property, most of the definitions and proofs we are going to give will not go through. fol:syn:unq:sec

Perhaps the best way to make this clear is to see what would happen if we had given bad rules for forming **formulas** that would not guarantee unique readability. For instance, we could have forgotten the parentheses in the formation rules for connectives, e.g., we might have allowed this:

If φ and ψ are **formulas**, then so is $\varphi \rightarrow \psi$.

Starting from an atomic formula θ , this would allow us to form $\theta \rightarrow \theta$. From this, together with θ , we would get $\theta \rightarrow \theta \rightarrow \theta$. But there are two ways to do this:

1. We take θ to be φ and $\theta \rightarrow \theta$ to be ψ .
2. We take φ to be $\theta \rightarrow \theta$ and ψ is θ .

Correspondingly, there are two ways to “read” the **formula** $\theta \rightarrow \theta \rightarrow \theta$. It is of the form $\psi \rightarrow \chi$ where ψ is θ and χ is $\theta \rightarrow \theta$, but *it is also* of the form $\psi \rightarrow \chi$ with ψ being $\theta \rightarrow \theta$ and χ being θ .

If this happens, our definitions will not always work. For instance, when we define the **main operator** of a formula, we say: in a formula of the form $\psi \rightarrow \chi$, the **main operator** is the indicated occurrence of \rightarrow . But if we can match the formula $\theta \rightarrow \theta \rightarrow \theta$ with $\psi \rightarrow \chi$ in the two different ways mentioned above, then in one case we get the first occurrence of \rightarrow as the **main operator**, and in the second case the second occurrence. But we intend the **main operator** to be a *function* of the **formula**, i.e., every **formula** must have exactly one **main operator** occurrence.

Lemma syn.9. *The number of left and right parentheses in a formula φ are equal.*

Proof. We prove this by induction on the way φ is constructed. This requires two things: (a) We have to prove first that all atomic formulas have the property in question (the induction basis). (b) Then we have to prove that when we construct new formulas out of given formulas, the new formulas have the property provided the old ones do.

Let $l(\varphi)$ be the number of left parentheses, and $r(\varphi)$ the number of right parentheses in φ , and $l(t)$ and $r(t)$ similarly the number of left and right parentheses in a term t .

Problem syn.2. Prove that for any term t , $l(t) = r(t)$.

1. $\varphi \equiv \perp$: φ has 0 left and 0 right parentheses.
2. $\varphi \equiv \top$: φ has 0 left and 0 right parentheses.
3. $\varphi \equiv R(t_1, \dots, t_n)$: $l(\varphi) = 1 + l(t_1) + \dots + l(t_n) = 1 + r(t_1) + \dots + r(t_n) = r(\varphi)$. Here we make use of the fact, left as an exercise, that $l(t) = r(t)$ for any term t .
4. $\varphi \equiv t_1 = t_2$: $l(\varphi) = l(t_1) + l(t_2) = r(t_1) + r(t_2) = r(\varphi)$.
5. $\varphi \equiv \neg\psi$: By induction hypothesis, $l(\psi) = r(\psi)$. Thus $l(\varphi) = l(\psi) = r(\psi) = r(\varphi)$.
6. $\varphi \equiv (\psi * \chi)$: By induction hypothesis, $l(\psi) = r(\psi)$ and $l(\chi) = r(\chi)$. Thus $l(\varphi) = 1 + l(\psi) + l(\chi) = 1 + r(\psi) + r(\chi) = r(\varphi)$.
7. $\varphi \equiv \forall x \psi$: By induction hypothesis, $l(\psi) = r(\psi)$. Thus, $l(\varphi) = l(\psi) = r(\psi) = r(\varphi)$.
8. $\varphi \equiv \exists x \psi$: Similarly. □

Definition syn.10 (Proper prefix). A string of symbols ψ is a *proper prefix* of a string of symbols φ if concatenating ψ and a non-empty string of symbols yields φ .

fol:syn:unq:
lem:no-prefix **Lemma syn.11.** *If φ is a formula, and ψ is a proper prefix of φ , then ψ is not a formula.*

Proof. Exercise. □

Problem syn.3. Prove [Lemma syn.11](#).

fol:syn:unq:
prop:unique-atomic **Proposition syn.12.** *If φ is an atomic formula, then it satisfies one, and only one of the following conditions.*

1. $\varphi \equiv \perp$.
2. $\varphi \equiv \top$.
3. $\varphi \equiv R(t_1, \dots, t_n)$ where R is an n -place *predicate symbol*, t_1, \dots, t_n are terms, and each of R, t_1, \dots, t_n is uniquely determined.
4. $\varphi \equiv t_1 = t_2$ where t_1 and t_2 are uniquely determined terms.

Proof. Exercise. □

Problem syn.4. Prove [Proposition syn.12](#) (Hint: Formulate and prove a version of [Lemma syn.11](#) for terms.)

Proposition syn.13 (Unique Readability). *Every formula satisfies one, and only one of the following conditions.*

1. φ is atomic.
2. φ is of the form $\neg\psi$.
3. φ is of the form $(\psi \wedge \chi)$.
4. φ is of the form $(\psi \vee \chi)$.
5. φ is of the form $(\psi \rightarrow \chi)$.
6. φ is of the form $(\psi \leftrightarrow \chi)$.
7. φ is of the form $\forall x \psi$.
8. φ is of the form $\exists x \psi$.

Moreover, in each case ψ , or ψ and χ , are uniquely determined. This means that, e.g., there are no different pairs ψ, χ and ψ', χ' so that φ is both of the form $(\psi \rightarrow \chi)$ and $(\psi' \rightarrow \chi')$.

Proof. The formation rules require that if a formula is not atomic, it must start with an opening parenthesis ($($, \neg , or a quantifier. On the other hand, every formula that starts with one of the following symbols must be atomic: a predicate symbol, a function symbol, a constant symbol, \perp , \top .

So we really only have to show that if φ is of the form $(\psi * \chi)$ and also of the form $(\psi' *' \chi')$, then $\psi \equiv \psi'$, $\chi \equiv \chi'$, and $*$ = $*'$.

So suppose both $\varphi \equiv (\psi * \chi)$ and $\varphi \equiv (\psi' *' \chi')$. Then either $\psi \equiv \psi'$ or not. If it is, clearly $*$ = $*'$ and $\chi \equiv \chi'$, since they then are substrings of φ that begin in the same place and are of the same length. The other case is $\psi \not\equiv \psi'$. Since ψ and ψ' are both substrings of φ that begin at the same place, one must be a proper prefix of the other. But this is impossible by Lemma syn.11. \square

syn.5 Main operator of a Formula

explanation It is often useful to talk about the last operator used in constructing a formula φ . This operator is called the *main operator* of φ . Intuitively, it is the “outermost” operator of φ . For example, the main operator of $\neg\varphi$ is \neg , the main operator of $(\varphi \vee \psi)$ is \vee , etc. fol:syn:mai:sec

Definition syn.14 (Main operator). The *main operator* of a formula φ is defined as follows: fol:syn:mai: def:main-op

1. φ is atomic: φ has no main operator.
2. $\varphi \equiv \neg\psi$: the main operator of φ is \neg .
3. $\varphi \equiv (\psi \wedge \chi)$: the main operator of φ is \wedge .
4. $\varphi \equiv (\psi \vee \chi)$: the main operator of φ is \vee .
5. $\varphi \equiv (\psi \rightarrow \chi)$: the main operator of φ is \rightarrow .

6. $\varphi \equiv (\psi \leftrightarrow \chi)$: the **main operator** of φ is \leftrightarrow .
7. $\varphi \equiv \forall x \psi$: the **main operator** of φ is \forall .
8. $\varphi \equiv \exists x \psi$: the **main operator** of φ is \exists .

In each case, we intend the specific indicated *occurrence* of the **main operator** in the formula. For instance, since the formula $((\theta \rightarrow \alpha) \rightarrow (\alpha \rightarrow \theta))$ is of the form $(\psi \rightarrow \chi)$ where ψ is $(\theta \rightarrow \alpha)$ and χ is $(\alpha \rightarrow \theta)$, the second occurrence of \rightarrow is the **main operator**.

This is a *recursive* definition of a function which maps all non-atomic **formulas** to their **main operator** occurrence. Because of the way **formulas** are defined inductively, every **formula** φ satisfies one of the cases in **Definition syn.14**. This guarantees that for each non-atomic **formula** φ a **main operator** exists. Because each **formula** satisfies only one of these conditions, and because the smaller **formulas** from which φ is constructed are uniquely determined in each case, the **main operator** occurrence of φ is unique, and so we have defined a function.

We call **formulas** by the names in **Table syn.1** depending on which symbol their **main operator** is.

Main operator	Type of formula	Example
none	atomic (formula)	$\perp, \top, R(t_1, \dots, t_n), t_1 = t_2$
\neg	negation	$\neg\varphi$
\wedge	conjunction	$(\varphi \wedge \psi)$
\vee	disjunction	$(\varphi \vee \psi)$
\rightarrow	conditional	$(\varphi \rightarrow \psi)$
\leftrightarrow	biconditional	$(\varphi \leftrightarrow \psi)$
\forall	universal (formula)	$\forall x \varphi$
\exists	existential (formula)	$\exists x \varphi$

Table syn.1: Main operator and names of **formulas**

fol:syn:mai:
tab:main-op

syn.6 Subformulas

fol:syn:sbf:
sec

It is often useful to talk about the **formulas** that “make up” a given **formula**. We call these its *subformulas*. Any **formula** counts as a **subformula** of itself; a **subformula** of φ other than φ itself is a *proper subformula*.

Definition syn.15 (Immediate Subformula). If φ is a **formula**, the *immediate subformulas* of φ are defined inductively as follows:

1. Atomic **formulas** have no immediate **subformulas**.
2. $\varphi \equiv \neg\psi$: The only immediate **subformula** of φ is ψ .
3. $\varphi \equiv (\psi * \chi)$: The immediate **subformulas** of φ are ψ and χ ($*$ is any one of the two-place connectives).
4. $\varphi \equiv \forall x \psi$: The only immediate **subformula** of φ is ψ .

5. $\varphi \equiv \exists x \psi$: The only immediate **subformula** of φ is ψ .

Definition syn.16 (Proper Subformula). If φ is a formula, the *proper subformulas* of φ are defined recursively as follows:

1. Atomic **formulas** have no proper **subformulas**.
2. $\varphi \equiv \neg\psi$: The proper **subformulas** of φ are ψ together with all proper **subformulas** of ψ .
3. $\varphi \equiv (\psi * \chi)$: The proper **subformulas** of φ are ψ , χ , together with all proper **subformulas** of ψ and those of χ .
4. $\varphi \equiv \forall x \psi$: The proper **subformulas** of φ are ψ together with all proper **subformulas** of ψ .
5. $\varphi \equiv \exists x \psi$: The proper **subformulas** of φ are ψ together with all proper **subformulas** of ψ .

Definition syn.17 (Subformula). The **subformulas** of φ are φ itself together with all its proper **subformulas**.

explanation

Note the subtle difference in how we have defined immediate **subformulas** and proper **subformulas**. In the first case, we have directly defined the immediate **subformulas** of a formula φ for each possible form of φ . It is an explicit definition by cases, and the cases mirror the inductive definition of the set of **formulas**. In the second case, we have also mirrored the way the set of all **formulas** is defined, but in each case we have also included the proper **subformulas** of the smaller **formulas** ψ , χ in addition to these **formulas** themselves. This makes the definition *recursive*. In general, a definition of a function on an inductively defined set (in our case, **formulas**) is recursive if the cases in the definition of the function make use of the function itself. To be well defined, we must make sure, however, that we only ever use the values of the function for arguments that come “before” the one we are defining—in our case, when defining “proper **subformula**” for $(\psi * \chi)$ we only use the proper **subformulas** of the “earlier” **formulas** ψ and χ .

Proposition syn.18. *Suppose ψ is a subformula of φ and χ is a subformula of ψ . Then χ is a subformula of φ . In other words, the subformula relation is transitive.*

fol:syn:sbf:
prop:subfrm-trans

Problem syn.5. Prove **Proposition syn.18**.

Proposition syn.19. *Suppose φ is a formula with n connectives and quantifiers. Then φ has at most $2n + 1$ subformulas.*

fol:syn:sbf:
prop:count-subfrms

Problem syn.6. Prove **Proposition syn.19**.

syn.7 Formation Sequences

fol:syn:fseq:sec Defining **formulas** via an inductive definition, and the complementary technique of proving properties of **formulas** via induction, is an elegant and efficient approach. However, it can also be useful to consider a more bottom-up, step-by-step approach to the construction of **formulas**, which we do here using the notion of a *formation sequence*. To show how terms and **formulas** can be introduced in this way without needing to refer to their inductive definitions, we first introduce the notion of an arbitrary string of symbols drawn from some language \mathcal{L} .

fol:syn:fseq: defn:string **Definition syn.20 (Strings).** Suppose \mathcal{L} is a first-order language. An \mathcal{L} -string is a finite sequence of symbols of \mathcal{L} . Where the language \mathcal{L} is clearly fixed by the context, we will often refer to a \mathcal{L} -string simply as a *string*.

Example syn.21. For any first-order language \mathcal{L} , all \mathcal{L} -formulas are \mathcal{L} -strings, but not conversely. For example,

$$)(v_0 \rightarrow \exists$$

is an \mathcal{L} -string but not an \mathcal{L} -formula.

fol:syn:fseq: defn:fseq-trm **Definition syn.22 (Formation sequences for terms).** A finite sequence of \mathcal{L} -strings $\langle t_0, \dots, t_n \rangle$ is a *formation sequence* for a term t if $t \equiv t_n$ and for all $i \leq n$, either t_i is a **variable** or a **constant symbol**, or \mathcal{L} contains a k -ary **function symbol** f and there exist $m_0, \dots, m_k < i$ such that $t_i \equiv f(t_{m_0}, \dots, t_{m_k})$. When it is necessary to distinguish, we will refer to formation sequences for terms as *term formation sequences*.

Example syn.23. The sequence

$$\langle c_0, v_0, f_0^2(c_0, v_0), f_0^1(f_0^2(c_0, v_0)) \rangle$$

is a formation sequence for the term $f_0^1(f_0^2(c_0, v_0))$, as is

$$\langle v_0, c_0, f_0^2(c_0, v_0), f_0^1(f_0^2(c_0, v_0)) \rangle.$$

fol:syn:fseq: defn:fseq-fm **Definition syn.24 (Formation sequences for formulas).** A finite sequence of \mathcal{L} -strings $\langle \varphi_0, \dots, \varphi_n \rangle$ is a *formation sequence* for φ if $\varphi \equiv \varphi_n$ and for all $i \leq n$, either φ_i is an atomic **formula** or there exist $j, k < i$ and a **variable** x such that one of the following holds:

1. $\varphi_i \equiv \neg \varphi_j$.
2. $\varphi_i \equiv (\varphi_j \wedge \varphi_k)$.
3. $\varphi_i \equiv (\varphi_j \vee \varphi_k)$.
4. $\varphi_i \equiv (\varphi_j \rightarrow \varphi_k)$.

5. $\varphi_i \equiv (\varphi_j \leftrightarrow \varphi_k)$.
6. $\varphi_i \equiv \forall x \varphi_j$.
7. $\varphi_i \equiv \exists x \varphi_j$.

When it is necessary to distinguish, we will refer to formation sequences for formulas as *formula formation sequences*.

Example syn.25.

$$\langle A_0^1(v_0), A_1^1(c_1), (A_1^1(c_1) \wedge A_0^1(v_0)), \exists v_0 (A_1^1(c_1) \wedge A_0^1(v_0)) \rangle$$

is a formation sequence of $\exists v_0 (A_1^1(c_1) \wedge A_0^1(v_0))$, as is

$$\langle A_0^1(v_0), A_1^1(c_1), (A_1^1(c_1) \wedge A_0^1(v_0)), A_1^1(c_1), \forall v_1 A_0^1(v_0), \exists v_0 (A_1^1(c_1) \wedge A_0^1(v_0)) \rangle.$$

As can be seen from the second example, formation sequences may contain “junk”: **formulas** which are redundant or do not contribute to the construction.

Proposition syn.26. *Every formula φ in $\text{Frm}(\mathcal{L})$ has a formation sequence.* *fol:syn:fseq:prop:formed*

Proof. Suppose φ is atomic. Then the sequence $\langle \varphi \rangle$ is a formation sequence for φ . Now suppose that ψ and χ have formation sequences $\langle \psi_0, \dots, \psi_n \rangle$ and $\langle \chi_0, \dots, \chi_m \rangle$ respectively.

1. If $\varphi \equiv \neg\psi$, then $\langle \psi_0, \dots, \psi_n, \neg\psi_n \rangle$ is a formation sequence for φ .
2. If $\varphi \equiv (\psi \wedge \chi)$, then $\langle \psi_0, \dots, \psi_n, \chi_0, \dots, \chi_m, (\psi_n \wedge \chi_m) \rangle$ is a formation sequence for φ .
3. If $\varphi \equiv (\psi \vee \chi)$, then $\langle \psi_0, \dots, \psi_n, \chi_0, \dots, \chi_m, (\psi_n \vee \chi_m) \rangle$ is a formation sequence for φ .
4. If $\varphi \equiv (\psi \rightarrow \chi)$, then $\langle \psi_0, \dots, \psi_n, \chi_0, \dots, \chi_m, (\psi_n \rightarrow \chi_m) \rangle$ is a formation sequence for φ .
5. If $\varphi \equiv (\psi \leftrightarrow \chi)$, then $\langle \psi_0, \dots, \psi_n, \chi_0, \dots, \chi_m, (\psi_n \leftrightarrow \chi_m) \rangle$ is a formation sequence for φ .
6. If $\varphi \equiv \forall x \psi$, then $\langle \psi_0, \dots, \psi_n, \forall x \psi_n \rangle$ is a formation sequence for φ .
7. If $\varphi \equiv \exists x \psi$, then $\langle \psi_0, \dots, \psi_n, \exists x \psi_n \rangle$ is a formation sequence for φ .

By the principle of induction on **formulas**, every **formula** has a formation sequence. □

We can also prove the converse. This is important because it shows that our two ways of defining formulas are equivalent: they give the same results. It also means that we can prove theorems about formulas by using ordinary induction on the length of formation sequences.

*fol:syn:fseq:
lem:fseq-init*

Lemma syn.27. Suppose that $\langle \varphi_0, \dots, \varphi_n \rangle$ is a formation sequence for φ_n , and that $k \leq n$. Then $\langle \varphi_0, \dots, \varphi_k \rangle$ is a formation sequence for φ_k .

Proof. Exercise. □

Problem syn.7. Prove [Lemma syn.27](#).

*fol:syn:fseq:
thm:fseq-frm-equiv*

Theorem syn.28. $\text{Frm}(\mathcal{L})$ is the set of all \mathcal{L} -strings φ such that there exists a formula formation sequence for φ .

Proof. Let F be the set of all strings of symbols in the language \mathcal{L} that have a formation sequence. We have seen in [Proposition syn.26](#) that $\text{Frm}(\mathcal{L}) \subseteq F$, so now we prove the converse.

Suppose φ has a formation sequence $\langle \varphi_0, \dots, \varphi_n \rangle$. We prove that $\varphi \in \text{Frm}(\mathcal{L})$ by strong induction on n . Our induction hypothesis is that every string of symbols with a formation sequence of length $m < n$ is in $\text{Frm}(\mathcal{L})$. By the definition of a formation sequence, either $\varphi \equiv \varphi_n$ is atomic or there must exist $j, k < n$ such that one of the following is the case:

1. $\varphi \equiv \neg \varphi_j$.
2. $\varphi \equiv (\varphi_j \wedge \varphi_k)$.
3. $\varphi \equiv (\varphi_j \vee \varphi_k)$.
4. $\varphi \equiv (\varphi_j \rightarrow \varphi_k)$.
5. $\varphi \equiv (\varphi_j \leftrightarrow \varphi_k)$.
6. $\varphi \equiv \forall x \varphi_j$.
7. $\varphi \equiv \exists x \varphi_j$.

Now we reason by cases. If φ is atomic then $\varphi_n \in \text{Frm}(\mathcal{L}_0)$. Suppose instead that $\varphi \equiv (\varphi_j \wedge \varphi_k)$. By [Lemma syn.27](#), $\langle \varphi_0, \dots, \varphi_j \rangle$ and $\langle \varphi_0, \dots, \varphi_k \rangle$ are formation sequences for φ_j and φ_k , respectively. Since these are proper initial subsequences of the formation sequence for φ , they both have length less than n . Therefore by the induction hypothesis, φ_j and φ_k are in $\text{Frm}(\mathcal{L}_0)$, and by the definition of a **formula**, so is $(\varphi_j \wedge \varphi_k)$. The other cases follow by parallel reasoning. □

Formation sequences for terms have similar properties to those for **formulas**.

*fol:syn:fseq:
prop:fseq-trm-equiv*

Proposition syn.29. $\text{Trm}(\mathcal{L})$ is the set of all \mathcal{L} -strings t such that there exists a term formation sequence for t .

Proof. Exercise. □

Problem syn.8. Prove [Proposition syn.29](#). Hint: use a similar strategy to that used in the proof of [Theorem syn.28](#).

There are two types of “junk” that can appear in formation sequences: repeated elements, and elements that are irrelevant to the construction of the formation or term. We can eliminate both by looking at minimal formation sequences.

Definition syn.30 (Minimal formation sequences). A formation sequence $\langle \varphi_0, \dots, \varphi_n \rangle$ for a formula φ is a *minimal formation sequence* for φ if for every other formation sequence s for φ , the length of s is greater than or equal to $n + 1$.

Similarly, a formation sequence $\langle t_0, \dots, t_n \rangle$ for a term t is a *minimal formation sequence* for t if for every other formation sequence s for t , the length of s is greater than or equal to $n + 1$.

Note that a formula or term can have more than one minimal formation sequence, but they will contain exactly the same strings.

Proposition syn.31. *The following are equivalent:*

1. ψ is a *sub-formula* of φ .
2. ψ occurs in every formation sequence of φ .
3. ψ occurs in a minimal formation sequence of φ .

Proof. Exercise. □

Problem syn.9. Prove [Proposition syn.31](#).

Historical Remarks Formation sequences were introduced by Raymond Smullyan in his textbook *First-Order Logic* (Smullyan, 1968). Additional properties of formation sequences were established by Zuckerman (1973).

syn.8 Free Variables and Sentences

Definition syn.32 (Free occurrences of a variable). The *free* occurrences of a *variable* in a *formula* are defined inductively as follows:

1. φ is atomic: all *variable* occurrences in φ are free.
2. $\varphi \equiv \neg\psi$: the free *variable* occurrences of φ are exactly those of ψ .
3. $\varphi \equiv (\psi * \chi)$: the free *variable* occurrences of φ are those in ψ together with those in χ .
4. $\varphi \equiv \forall x \psi$: the free *variable* occurrences in φ are all of those in ψ except for occurrences of x .
5. $\varphi \equiv \exists x \psi$: the free *variable* occurrences in φ are all of those in ψ except for occurrences of x .

Definition syn.33 (Bound Variables). An occurrence of a **variable** in a formula φ is *bound* if it is not free.

Problem syn.10. Give an inductive definition of the bound variable occurrences along the lines of **Definition syn.32**.

Definition syn.34 (Scope). If $\forall x \psi$ is an occurrence of a subformula in a formula φ , then the corresponding occurrence of ψ in φ is called the *scope* of the corresponding occurrence of $\forall x$. Similarly for $\exists x$.

If ψ is the scope of a quantifier occurrence $\forall x$ or $\exists x$ in φ , then the free occurrences of x in ψ are bound in $\forall x \psi$ and $\exists x \psi$. We say that these occurrences are *bound by* the mentioned quantifier occurrence.

Example syn.35. Consider the following formula:

$$\exists v_0 \underbrace{A_0^2(v_0, v_1)}_{\psi}$$

ψ represents the scope of $\exists v_0$. The quantifier binds the occurrence of v_0 in ψ , but does not bind the occurrence of v_1 . So v_1 is a free variable in this case.

We can now see how this might work in a more complicated **formula** φ :

$$\forall v_0 \underbrace{(A_0^1(v_0) \rightarrow A_0^2(v_0, v_1))}_{\psi} \rightarrow \exists v_1 \underbrace{(A_1^2(v_0, v_1) \vee \forall v_0 \underbrace{\neg A_1^1(v_0)}_{\theta})}_{\chi}$$

ψ is the scope of the first $\forall v_0$, χ is the scope of $\exists v_1$, and θ is the scope of the second $\forall v_0$. The first $\forall v_0$ binds the occurrences of v_0 in ψ , $\exists v_1$ binds the occurrence of v_1 in χ , and the second $\forall v_0$ binds the occurrence of v_0 in θ . The first occurrence of v_1 and the fourth occurrence of v_0 are free in φ . The last occurrence of v_0 is free in θ , but bound in χ and φ .

Definition syn.36 (Sentence). A **formula** φ is a *sentence* iff it contains no free occurrences of **variables**.

syn.9 Substitution

fol:syn:sub:
sec

Definition syn.37 (Substitution in a term). We define $s[t/x]$, the result of *substituting* t for every occurrence of x in s , recursively:

1. $s \equiv c$: $s[t/x]$ is just s .
2. $s \equiv y$: $s[t/x]$ is also just s , provided y is a variable and $y \neq x$.
3. $s \equiv x$: $s[t/x]$ is t .
4. $s \equiv f(t_1, \dots, t_n)$: $s[t/x]$ is $f(t_1[t/x], \dots, t_n[t/x])$.

Definition syn.38. A term t is *free for x* in φ if none of the free occurrences of x in φ occur in the scope of a quantifier that binds a variable in t .

Example syn.39.

1. v_8 is free for v_1 in $\exists v_3 A_4^2(v_3, v_1)$
2. $f_1^2(v_1, v_2)$ is *not* free for v_0 in $\forall v_2 A_4^2(v_0, v_2)$

Definition syn.40 (Substitution in a formula). If φ is a formula, x is a variable, and t is a term free for x in φ , then $\varphi[t/x]$ is the result of substituting t for all free occurrences of x in φ .

1. $\varphi \equiv \perp$: $\varphi[t/x]$ is \perp .
2. $\varphi \equiv \top$: $\varphi[t/x]$ is \top .
3. $\varphi \equiv P(t_1, \dots, t_n)$: $\varphi[t/x]$ is $P(t_1[t/x], \dots, t_n[t/x])$.
4. $\varphi \equiv t_1 = t_2$: $\varphi[t/x]$ is $t_1[t/x] = t_2[t/x]$.
5. $\varphi \equiv \neg\psi$: $\varphi[t/x]$ is $\neg\psi[t/x]$.
6. $\varphi \equiv (\psi \wedge \chi)$: $\varphi[t/x]$ is $(\psi[t/x] \wedge \chi[t/x])$.
7. $\varphi \equiv (\psi \vee \chi)$: $\varphi[t/x]$ is $(\psi[t/x] \vee \chi[t/x])$.
8. $\varphi \equiv (\psi \rightarrow \chi)$: $\varphi[t/x]$ is $(\psi[t/x] \rightarrow \chi[t/x])$.
9. $\varphi \equiv (\psi \leftrightarrow \chi)$: $\varphi[t/x]$ is $(\psi[t/x] \leftrightarrow \chi[t/x])$.
10. $\varphi \equiv \forall y \psi$: $\varphi[t/x]$ is $\forall y \psi[t/x]$, provided y is a variable other than x ; otherwise $\varphi[t/x]$ is just φ .
11. $\varphi \equiv \exists y \psi$: $\varphi[t/x]$ is $\exists y \psi[t/x]$, provided y is a variable other than x ; otherwise $\varphi[t/x]$ is just φ .

explanation

Note that substitution may be vacuous: If x does not occur in φ at all, then $\varphi[t/x]$ is just φ .

The restriction that t must be free for x in φ is necessary to exclude cases like the following. If $\varphi \equiv \exists y x < y$ and $t \equiv y$, then $\varphi[t/x]$ would be $\exists y y < y$. In this case the free variable y is “captured” by the quantifier $\exists y$ upon substitution, and that is undesirable. For instance, we would like it to be the case that whenever $\forall x \psi$ holds, so does $\psi[t/x]$. But consider $\forall x \exists y x < y$ (here ψ is $\exists y x < y$). It is a sentence that is true about, e.g., the natural numbers: for every number x there is a number y greater than it. If we allowed y as a possible substitution for x , we would end up with $\psi[y/x] \equiv \exists y y < y$, which is false. We prevent this by requiring that none of the free variables in t would end up being bound by a quantifier in φ .

We often use the following convention to avoid cumbersome notation: If φ is a formula which may contain the variable x free, we also write $\varphi(x)$ to

indicate this. When it is clear which φ and x we have in mind, and t is a term (assumed to be free for x in $\varphi(x)$), then we write $\varphi(t)$ as short for $\varphi[t/x]$. So for instance, we might say, “we call $\varphi(t)$ an instance of $\forall x \varphi(x)$.” By this we mean that if φ is any **formula**, x a **variable**, and t a term that’s free for x in φ , then $\varphi[t/x]$ is an instance of $\forall x \varphi$.

Photo Credits

Bibliography

Smullyan, Raymond M. 1968. *First-Order Logic*. New York, NY: Springer.
Corrected reprint, New York, NY: Dover, 1995.

Zuckerman, Martin M. 1973. Formation sequences for propositional formulas.
Notre Dame Journal of Formal Logic 14(1): 134–138.