Building a Phonological Inventory

Veer, B.M. van 't

Author: Veer, B.M. van 't
Title: Building a phonological inventory : feature co-occurrence constraints in acquisition
Issue Date: 2015-01-07

CHAPTER 1

Introduction

Every language has an inventory of phonological segments. Every child acquires a language. Hence, every child must acquire the inventory of phonological segments of her language. This, in a nutshell, is the topic of the current thesis: to arrive at a theory that describes the phonological segment inventory, and describes how it is acquired.[1] This theory is dubbed Feature Co-occurrence Constraint Theory. We shall develop a theory with minimal ingredients: distinctive features, a generator mechanism which (limitlessly) combines these features, and two types of constraints on the co-occurrence of features. One of the main innovations of the theory is that it views acquisition as the tandem development of representational and computational aspects of phonology, rather than focussing on one aspect in particular, as previous approaches tended to do. Each of the ingredients is in and of itself uncontroversial; combining them is uncontroversial, as well. This thesis attempts to be as explicit about assumptions and implications as is possible.

The Feature Co-occurrence Constraint Theory does not stand in isolation; it has a number of predecessors (see section 1.2). One of these, Hayes (1999), provides a phonetically motivated alternative to the innate constraint hypothesis (see also Hayes & Wilson, 2008; Adriaans & Kager, 2010). Furthermore, Fikkert and Levelt (2008) propose a lexicon-driven constraint generation mechanism (see section 4.6). In our theory, constraint templates are innately present, and these templates are populated (with features) and activated as soon as the conditions for population (i.e., the acquisition of two features) are present. In this sense, the current proposal is firmly couched within the tradition of gen-

[1] In this thesis, we consider only spoken language.

erative approaches to phonology.

The natural habitat for generators and constraints in contemporary phonology is Optimality Theory (Prince & Smolensky, 1993/2004, henceforth: OT). Most certainly, there is no objection to implement the theory of Feature Co-occurrence Constraints in that framework; in fact, we borrow many formalisms and insights from the rich Optimality Theory literature. First, there is a generator that supplies the full permutation of feature combinations, where feature combinations (in the simultaneous sense) stand for possible segments. This generator, to which we shall not devote a great deal of attention, has an obvious analogue in Optimality Theory: GEN. Next, Feature Co-occurrence Constraints are invoked and revoked; the revocation of constraints can be modeled in OT by constraint demotion. The invocation of constraints is somewhat less straightforward; in 'classical' OT, all constraints are thought to be universal and innate (Prince & Smolensky, 1993/2004), and while it has been proposed that in the initial state, markedness constraints collectively outrank faithfulness constraints (Gnanadesikan, 1995; Boersma, 1998; Tesar & Smolensky, 2000), some have argued for the opposite (Hale & Reiss, 1998). In their computational model, Boersma and Hayes (2001) show that (under some assumptions at least) an initial state in which constraints are unranked (or: ranked in the same stratum) can lead to a successful final state grammar. With respect to our current proposal, where constraints are invoked, we can assume that it makes no great difference how constraints are ranked in the initial state. We assume that constraints can only be invoked at one point in the development of the grammar: namely, the point at which the latter of the two features the constraint refers to is acquired (we will argue for innate features in chapter 2). Even if the constraints are innate and every possible Feature Co-occurrence Constraint is present in the initial state, we will only see its effect if its structural description (that is, both features) is present. Hence, while we assume that constraints are invoked (see section 4.6), even if they are in fact innate, their ranking prior to the invocation point is irrelevant.

Whereas OT provides a ranking scale for constraints (be it absolute or gradual (Boersma, 1998)), in other frameworks constraints are either active and not violated, or inactive. A theory of constraints that aspires to be compatible with both Optimality Theory and non-constraint-ranking theories must forego the possibilities offered by a full-fledged infinite constraint hierarchy, and be limited to constraints divided over no more than two strata: one for violated constraints, one for satisfied constraints. Constraints in the latter strata are those that are active in non-OT frameworks. We will see, in chapter 3, that a two-strata OT grammar results from generating the Feature Co-occcurrence Constraints necessary for any inventory, if no faithfulness constraints are considered. Hence, the theory is transportable between frameworks. In chapter 5 we will return to this issue.

1.1 Feature Co-occurrence Constraints: the proposal

As mentioned above, the Feature Co-occurrence Constraint Theory rests on features and constraints. The central idea is quite simple: the acquisition of the segment inventory can best be described in terms of the tandem development of both the representational and the output components of the phonological system (see also Levelt & van Oostendorp, 2007). The representational component is taken to be the lexicon, which is populated by lexical items consisting of structured bundles of monovalent features. Monovalency is not a choice on which the theory depends; it is, however, well motivated (section 3.4.1) and has repercussions for the system of constraints that we propose (chapter 3).

The output system is represented by a system of highly simple Feature Co-occurrence Constraints. Only two types of constraints are necessary; these are represented in the form of constraint templates in (1) below.

(1) a. *FG
assign a violation mark for every segment Σ iff [F] is in Σ and [G] is in Σ (*c-constraint*)

b. F→G
assign a violation mark for every segment Σ iff [F] is in Σ and [G] is not in Σ (*i-constraint*)

We call these constraints 'c-constraints' (where 'c' stands for 'co-occurrence'), and i-constraints (where 'i' stands for 'implicational'). With these two constraint templates, we will show that the inventory of adult language can be described (chapter 3), and likewise the acquisition of that inventory. The c-constraints are straightforward: they are violated by the combination in one segment of the features that are listed in the constraint's structural description. The addition of i-constraints is closely tied to the adoption of monovalent, rather than binary features (section 3.4.1). In monovalent feature theory, the complement of the presence of a feature is its absence, whereas in binary feature theory the complement of [+F] can be expressed by [-F]. Hence, if a feature [F] cannot co-occur without feature [G], this can be formulated in binary feature theory by *[+F, -G], which fits the ordinary c-constraint template. The expression [-G] is unavailable in monovalent feature theory, but the logical formula $\neg(F \wedge \neg G)$, which is the logical representation of the c-constraint mentioned above, is equivalent to the statement $F \rightarrow G$, which is the logical representation of i-constraints (this argument will be explored further in chapter 3).

A crucial assumption about Feature Co-occurrence Constraints is that they can be invoked, but also revoked (or demoted). Many individual constraints remain active throughout the final state (adult grammar), but some are transient and only active for some period during acquisition. These constraints are invoked when a feature has been acquired, but is still limited in its combinational options. Revocation implies that the feature becomes more free in its association with other features.

Features are assumed to be strictly monovalent, and innate. In section 2.3 we will review a large body of literature that in some way pertains to the acquisition of features, and we will see that the evidence in favor of innateness is fairly strong. Monovalency will be defended more thoroughly in chapter 3. Furthermore, some properties are not represented in featural terms. There is no feature [coronal], and neither is there a feature [stop]. This implies that the empty segment (which cannot be excluded by the types of FCCs developed here) is interpreted by the phonetics as /t/.

Feature Co-occurrence Constraints act as a filter to separate legal from illegal segments, or grammatical from ungrammatical segments. Even illegal segments in one language, however, are still possible segments, and may be grammatical in another language. We define a 'possible segment' as any combination of distinctive features. This implies that there is no principled difference between universal gaps and language specific gaps: both arise through acquisition, and hence the motivation for universal gaps must be sought elsewhere.

Feature geometry and major class features are difficult to combine with the Feature Co-occurrence Constraint Theory as it is proposed in the current thesis, but at the same time, it remains important to be able to group features into classes, and refer to these classes. This issue is resolved by adopting a version of Feature Class theory (Padgett (2002), see also chapter 3).

Finally, the Feature Co-occurrence Constraint Theory is a theory not so much about the inventory, but about the gaps in the inventory that can occur. Every constraint describes a gap; a fully symmetrical inventory can be described with no FCCs at all. A drive towards as little active (or high ranked) FCCs as possible is also the drive towards more symmetry. A fully symmetrical system is also fully maximally economic, and every gap represents a decrease in economy.

We will see that the number of Feature Co-occurrence Constraints that are activated is relatively small. Only 19 are needed to accurately describe the inventory of Dutch, and the number for each stage in acquisition is of a similar magnitude. Compared to the impressively large inventory of possible constraints, this is number is perhaps in itself an argument against the innateness of substantive Feature Co-occurrence Constraints, instead of the innateness of FCC templates (as we assume here). Before turning to the main matter of the dissertation, we shall briefly consider the earlier uses of Feature Co-occurrence Constraints in adult phonology, and previous approaches to the acquisition of the segment inventory.

1.2 Earlier uses of Feature Co-occurrence Constraints in adult phonology

An early predecessor of Feature Co-occurrence Constraints is the system of Marking Conventions of Chomsky and Halle (1968, henceforth: SPE)'s chapter nine. These state universal constraints on feature co-occurrences, in both the paradigmatic and syntagmatic sense. Interestingly, the Marking Conventions are both violable (they indicate markedness preferences rather then hard restrictions), and universal. Many other conceptions of constraints in the phonological literature take them to be inviolable, language-specific generalisations over the (underlying) structure of morphemes (Stanley, 1967) or surface phonotactics (Shibatani, 1973). It is not until Singh (1987) and Paradis (1988) that the we once more encounter the idea of the universal phonological constraint.

1.2.1 Feature combinations in rule-based phonology

Admittedly not technically constraints by format, the Marking Conventions serve a similar role as our Feature Co-occurrence Constraints: to limit the generative power of a grammar. This is not *entirely* correct, however, because in SPE the marking conventions are "...not part of the grammar but rather conventions for the interpretation of a grammar"(Chomsky & Halle, 1968, p. 403) whereas in the current proposal, the Feature Co-occurrence Constraints *are* the grammar. In SPE, it is a system of highly formal ("overly formal" in the words of Chomsky & Halle, 1968) rules, constituting the phonological grammar, that are capable of generating much more than is attested; in our case, there is a generator function, much like GEN in Optimality Theory, that is equally blind to the "intrinsic content" of features.

Although the Marking Conventions in SPE all take the shape of implicational statements (like our i-constraints), the reliance on binary feature values presents us with an interesting reversal of the situation outlined in section 1.1 above. Whereas we adopt i-constraints to accommodate co-occurrence restrictions that could be expressed as c-constraint only if using binary features, some of the Marking Conventions in SPE utilise the implicational format (or rewrite format) to express what we would express in c-constraints. Take, for example, the "universal constraint on feature combination" that states that it is impossible to have segments that are both [+low] and [+high] (Chomsky & Halle, 1968, p. 404). Both features being positively valued, the constraint can be rephrased in FCC terms as *[low, high]. Chomsky and Halle (1968, p. 405), however, formalise the constraint by two Marking Conventions:[2]

(2) Marking Conventions

a. Marking Convention 6VII
[+low] → [-high]

[2] See also the paragraph on Kiparsky and Pajusalu (2003) below.

b. Marking Convention 6IX
[+high] → [-low]

In the theory proposed in the current thesis, the segment [low, high] is a possible segment and will be generated; it will also be ungrammatical by virtue of the c-constraint *[low, high], which, we predict, will be activated and remain active in the grammar of learners of every language where no mid vowels exist.[3]

Close in spirit and format to the universal Marking Conventions of SPE are the redundancy and predictability rules in the various flavours of Underspecification Theory (Archangeli, 1984, 1988, etc.). These too serve to limit the generative power of the derivational mechanism, and they do so in a similar way: by limiting the amount of information specified in the input of phonological rules, the range of possible output structures is likewise limited. Among the differences between these types of rules and the SPE Marking Conventions is that, typically, the redundancy and predictability rules of Underspecification Theory are part of the set of operations performed by the grammar. This includes insertion *during* (or as part of) derivation.

As in SPE, redundancy and predictability rules are not constraints but rather implicational statements. In (3) we see an example of fill-in rules in Radical Underspecification (example from Archangeli, 1988). A five-member vowel inventory is specified as in (3a), and the missing values are inserted by the set of rules in (3b).

(3) Fill-in rules in Radical Underspecification

a. Inventory and specifications

	i	e	ɑ	o	u
[high]		-		-	
[low]			+		
[back]				+	+
[voice]					

b. Fill-in rules

i. [+low] → [-high]
ii. [+low] → [+back]
iii. [ ] → [-low]
iv. [ ] → [+high]
v. [ ] → [-back]
vi. [ ] → [+voice]

The main difference between the rules in (3b) and the Feature Co-occurrence Constraints proposed here, apart from the fact that the former are not technically constraints, is that FCCs act as a filter; they come in *after* the fact

[3] It should be noted that mid vowels are represented as the presence of both height features in the current privative approach, contrary to the standard practice in binary feature theory, where mid vowels are [-high, -low].

whereas fill-in rules are applied in the course of derivation. Fill-in rules do not ban what is already generated; they postpone certain types of generation to the point that it no longer applies. We will see that in the current proposal, some traits are unspecified (as in Underspecification Theory), but unspecification is permanent, and not undone by the FCCs. Feature Co-occurrence Constraints do not fill in, they filter out. Another way of stating the difference is to say that FCCs do not limit the generative power of the grammar *per se*, but rather undo some of it.

1.2.2 Constraints on feature combinations in harmony

Proper constraints on the co-occurrence of distinctive features in both OT and non-Optimality Theory phonology we encounter in the treatment of (vowel) harmony. In a government-based account of vowel harmony, Van der Hulst and Smith (1986) and Van der Hulst (1988) employ feature co-occurrence constraints in combination with strict locality to model opacity (see also Piggott & van der Hulst, 1997). In this analysis, the incompatibility of, for example, [low] and [ATR] in the surface vowel inventory blocks [ATR] spreading when a [low] vowel is encountered; strict locality prohibits it from looking further. It should not go unmentioned that these authors employ co-occurrence constraints in a way that is remarkably similar to the current proposal. Van der Hulst employs elements rather than features, and discusses vowel inventories rather than consonant inventories (as the subject of his study is vowel harmony), and does not explicitly limit the types of constraint as we shall do in chapter 3, but otherwise his system is very close indeed to the current proposal.[4] The spirit of Van der Hulst (2012) is very similar to the theory explored in Van der Hulst and Smith (1986) and Van der Hulst (1988), namely that the behaviour of non-alternating vowels in vowel harmony languages is predictable: inert vowels that are compatible with the harmonising feature are transparent, whereas those that are incompatible with the harmonising feature are opaque. In Van der Hulst (2012), incompatibility is explicitly modeled in terms of FCCs.

Let us move on to an Optimality Theoretic account of vowel harmony in which constraints on feature co-occurrences play an important role. Kiparsky and Pajusalu (2003) aim to derive a typology of neutral vowels in [back] harmony. They do so from the Factorial Typology of a harmony driving constraint (Agree[F]), positional faithfulness constraints and constraints on feature co-occurrence. The latter are anti-harmonic: they act against the surfacing of marked segments:

(4) Featural markedness constraints in Kiparsky and Pajusalu (2003)

a. [–low, –round] → [–back]
If a vowel is non-low and ungrounded, it must be front

[4]It should be noted that Van der Hulst (2012) also uses syntagmatic FCCs, see the discussion on Hayes (1999).

b. [–back] → [–low, –round]
If a vowel is front, it must be non-low and unrounded

Although similar to our Feature Co-occurrence Constraints, the featural markedness constraints in Kiparsky and Pajusalu (2003) are different in several respects: first, they refer to three features rather than to two (see section 3.3.1 for a discussion). Second, they only act against marked segments, whereas our Feature Co-occurrence Constraints are blind to markedness: every combination of two features is a possible candidate for an FCC.

Feature Co-occurrence Constraints are employed in the studies of Turkish vowel harmony by Padgett (2002) and Linke and van Oostendorp (in preparation). The facts of Turkish vowel harmony, very briefly, are that all suffix vowels agree with the stem vowel in terms of backness, and high vowels additionally agree in roundness (remember that Turkish has a perfectly symmetrical eight-vowel inventory in stems). Building on substantial evidence, Padgett (2002) argues that both forms of harmony should be analysed as one and the same process, dubbed 'color' harmony, where [color] stands for the natural class formed by [back] and [round] vowels.[5] The fact that low vowels do not harmonise in roundness is captured by the Feature Co-occurrence Constraint *[low, round]. Padgett's proposal is couched in Optimality Theory, and the co-occurrence constraint outranks the constraint enforcing [color] harmony. Low round vowels do occur, however, albeit that their distribution is limited to stems. Hence, a set of faithfulness constraints for [low] and [round] outranks both the harmony and the co-occurrence constraint.

Linke and van Oostendorp (in preparation) employ a set of constraints that is minimally different from the set in (1). Using the same format of c-constraints, the definition is slightly different (the formulation in (5) is my own paraphrase of Linke and van Oostendorp (in preparation)'s constraint definition):

(5) a. *FG
assign a violation mark for every feature [F] that co-occurs with feature [G] (*c-constraint*)

Although the different formulation is not important to the proposal made in the current theory, it is crucial to Linke and van Oostendorp (in preparation) that violations are assigned on a per-feature base. This is because features that violate a c-constraint in the stem, but do this to avoid violating a higher ranked faithfulness constraint, do not incur additional violations when spreading to suffix vowels. Under the segment-based definition in (1), the c-constraint in a harmonic form would be violated twice: once by the stem vowel, and once by the suffix vowel. Under the alternative definition, conversely, the fact that the

[5] A more extensive discussion of Padgett's proposal will be provided in chapter 3. The reader is reminded that in Element Theory (see, e.g., Harris, 1994; Harris & Lindsey, 1995; Backley, 2011) both backness and roundness are represented by one and the same element, |U|.

feature in the stem vowel is linked to two segmental slots is of no effect to the amount of violations: there is only one offending feature, and the number of association lines is irrelevant for the c-constraint.

1.2.3 Feature Co-occurrence Constraints and specificational variation

Laterals are notoriously variable in their typological behaviour. In her overview, Yip (2011) proposes to capture this behaviour by means of two universal ranking schemes involving simple feature co-occurrence constraints referring to the feature [lateral]; one for place (6a), one for major class (6b):

(6) Universal ranking schemas for [lateral]

a. Place of Articulation
[Lateral, Labial]≫[Lateral, Dorsal]≫*[Lateral, Coronal]

b. Major classes
[Lateral, Obstruent]≫[Lateral, Sonorant]

By interspersing these universal rankings of c-constraints with faithfulness constraints, different types of inventories are derived. For example, if faithfulness is ranked below the three constraints in (6a), the resulting inventory will have either no laterals at all (as in Maori), or placeless laterals only (Cambodian). Ranking faithfulness between *[Lateral, Dorsal] and *[Lateral, Coronal] rules out labial and dorsal laterals, but allows for laterals with a coronal place specification (Yip gives English as an example). Interaction with Share-F constraints results in a typology of lateral spreading behaviour.

Although Yip (2011)'s constraints are formally very similar (if not equal) to our c-constraints, it must be noted that the way in which they are employed is conceptually somewhat different. The main difference is that the constraints are ranked with respect to each other (or, more correctly, with respect to faithfulness or spreading constraints) in a universal hierarchy. Nevertheless, Yip (2011) shows that c-constraints can be used to derive inventories.

So far we have mostly seen variants of c-constraints, or i-constraints to harness the effects of c-constraints (as in Kiparsky & Pajusalu, 2003). An early OT-based use of i-constraints is Itô, Mester, and Padgett (1995), who focus on the effects of the constraint [sonorant]→[voice]. Apart from stating that the presence of [sonorant] implies the presence of [voice], Itô et al. (1995) note that the constraint states that voicing is redundant, or non-contrastive, for sonorants: sonorants are always voiced. From this, they deduce that voicing is not *licensed* in sonorants; and in interaction with the constraint License[voice], which requires that the feature [voice] always be licensed, the implicational constraint drives surface underspecification of sonorants for [voice]. At the same time, there is no incompatibility between [voice] and [sonorant] (which would follow from a c-constraint such as *[sonorant, voice]), which would equally derive phonologically voiceless sonorants at the surface.

The discussion in Itô et al. (1995) takes place in the context of Yamato Japanese, where the initial obstruents of the second member of compounds becomes voiced (a process known as Rendaku) *unless* there is a voiced consonant elsewhere in the same morpheme (a generalisation known as Lyman's Law). An example of Rendaku is given in (7a), whereas (7b) shows how Lyman's Law interferes.

(7) Compounds in Yamato Japanese
  a. yama tera → yama+dera 'mountain temple'
  b. ore kugi → ore+kugi 'broken nail' (*ore+gugi)
  c. ori kami → ori+gami 'paper folding'

Crucially, example (7c) shows that nasals do not block Rendaku; hence, they must be underspecified for [voice].

The problem is that Yamato Japanese has a different property involving nasals and [voice], which leads to what Itô et al. (1995) call an 'underspecification paradox'. This property is post-nasal voicing, as in (8).

(8) Post-nasal voicing in Japanese
  a. tombo 'dragonfly' (*tompo)
  b. /yom+te/ → yonde 'reading'

Example (8a) shows that post-nasal voicing holds true of monomorphemic forms, whereas (8b) shows that it is active in derived forms, as well. The paradox is that whereas nasals act as voiceless for Rendaku and Lyman's Law, they seem to act as phonologically voiced in post-nasal voicing.

Without doing justice to the full analysis in Itô et al. (1995), the solution lies in the interpretation of the implicational constraint [sonorant]→[voice] (which, for reasons we will not go into here, only pertains to nasals in their analysis). The implication entails a 'licensing cancellation': as indicated above, sonorants do not independently license [voice], but they are not incompatible with [voice] either. This leads to the non-specification of sonorants for [voice] in the general case. Obstruents, however, are able to license [voice]. In NC clusters, therefore, the nasal acts as if it is voiced, because it is linked to a [voice] feature. That same feature is also linked (by spreading) to the obstruent, which licenses it (remember that high-ranked LICENCE[VOICE] prohibits that the optimal candidate contains unlicensed instances of [voice]).

Although Itô et al. (1995) employ an implicational constraint that is formally similar to the i-constraint in the current thesis, there are important differences in interpretation. We will see, in chapter 3, that the description of the Dutch consonant inventory involves one i-constraint: [approximant] → [continuant]. We shall follow Itô et al. (1995) in observing that this constraint (when unviolated) states that [continuant] is redundant on [approximant], but we do not adhere to their notion of 'licensing cancellation', nor will we explicitly state any licensing requirements for the realisation of features in the surface form. Hence, in Feature Co-occurrence Theory, [apprx]→[cont] does not entail that approximants are un(der)specified for [continuant]. Incidentally, an interesting learnability issue arises under the Itô et al. (1995, et passim) interpretation of implicational constraints: if sonorants are surface-underspecified for [voice] except in NC clusters, then how will the learner know that [sonorant]→[voice]?

1.2.4 Feature Co-occurrence Constraints and the inventory

We have now seen a number of specific applications of feature co-occurrence constraints in Optimality Theory, from vowel harmony to typological variation and questions regarding (under)specification. Before we turn to the an overview of earlier approaches to the acquisition of the inventory, it must be mentioned that from the early days of OT, constraints on feature combinations were understood to govern the structure of the adult inventory (see, for example, Kager, 1999 who uses feature co-occurrence constraints in various forms). At the same time, employing constraints on feature co-occurrences to model the consonant inventory is by no means a practice exclusive to Optimality Theory; the reader is reminded of the discussion on Van der Hulst (2012) in section 1.2.2.

In section 4.6, we will discuss the theory of Inductive Grounding (Hayes, 1999) in some detail. What is important to our present discussion is that Hayes (1999) uses both paradigmatic and syntagmatic feature co-occurrence constraints, where I take 'paradigmatic' to mean infra-segmental, and 'syntagmatic' to mean intersegmental.[6] Four types of constraints are proposed to be emergent:[7]

(9) a. *[F][G]
assign a violation mark for every output where [F] precedes [G]

b. *[F]
assign a violation mark for every [F]

c. *[F, G]
assign a violation mark for every co-occurrence of [F] and [G]

d. *[F] unless [G]
assign a violation mark for every [F] that occurs without [G]

The latter two constraints are directly related to our c-constraints and i-constraints, respectively; the second will be shown in chapter 3 to be a special version of a c-constraint, and the first is a syntagmatic constraint, pertaining

[6]Kiparsky and Pajusalu (2003) follow the same nomenclature.

[7]Hayes does not provide formal definitions, and to avoid going beyond a general impression into a distracting discussion about technicalities the definitions given here are admittedly somewhat sloppy.

to phonotactic restrictions (which we are not primarily concerned with in the current thesis).

Hayes' constraints live in an OT ecosystem, and function much to the same end as the FCCs proposed in the current thesis, namely to limit the range of grammatical combinations of features. An important difference is that Hayes' proposal builds on a notion of phonetic difficulty, represented in a 'difficulty map'. Of all possible feature co-occurrence constraints, only those are selected to participate in the actual grammar that are more 'effective' than constraints that are minimally different, where effectiveness is measured by how well the constraint predicts the relative difficulty of two points on the difficulty map (the reader is referred to section 4.6 for a more elaborate discussion of Inductive Grounding).

1.3 Two perspectives on the acquisition of the segment inventory

The acquisition of language is of interest to science in at least two ways: first, the process it worth studying in its own right. Acquiring a language implies interacting with the world, uncovering knowledge and codifying it, *et cetera*. At the same time, any theory of the final state, that is, the adult grammar, must show that it is learnable for the theory to be a candidate model of human linguistic competence. These two perspectives do not always converge, and consequently Ingram (1989) distinguishes between the 'developmental problem of language acquisition', which concerns the study of language acquisition in the first sense, and the 'logical problem of language acquisition', concerning learnability issues.

In this section, we will review some of the past literature on the acquisition of the segment inventory, classifying studies in one of Ingram's two categories, or, when appropriate, as aiming for a synthesis. It is perhaps to the latter category that the current study belongs.

We have already mentioned Jakobson (1941/1968), the study that directly inspired much of modern day language acquisition research. As mentioned by Ingram (1989, section 6.5.2) in his thorough review of Jakobson's proposal, the Jakobsonian theory of language acquisition is incomplete and imprecise at some points. In fact, Ingram proposes to read Jakobson (1941/1968) as the *outlines* of a theory rather than a theory in its own right – a proposal which I think is correct. Nevertheless, *Kindersprache* set the agenda for language acquisition research for decades after its publication, or after the publication of the English translation in 1968. One of the most spectacular claims made by Jakobson is that the acquisition of *oppositions* ('contrasts') is fixed and universal in terms of precedence. That is to say, the Jakobsonian hypothesis with respect to the acquisition of oppositions is that the order of acquisition proceeds such that the greatest opposition is acquired first, and subsequent oppositions are acquired in decreasing orders of magnitude. The point of departure for the child is the CV syllable, meaning that the opposition between consonants and vowels is prime. For consonants, the first two oppositions are nasal vs. oral and labial vs. dental. The first vocalic opposition is high versus low, or /i/ versus /a/. The following vowel stage introduces either a front-back opposition (resulting in /i, u, a/) or a further refinement of the height dimension (such that the child has /i, e, a/).

This claim is a rather strong one at first glance, and it has triggered a host of reactions. However, as pointed out by Ingram (1989), it is not a straightforward hypothesis to test, for being somewhat underdeveloped and not very specific in some aspects. The Jakobsonian hypothesis holds that the order of acquisition of phonemic contrasts is related to typological markedness statements, or, in the words of Jakobson (1941/1968), 'laws of irreversible solidarity'. A result of this claim is that a great deal of importance is ascribed to the order of acquisition of the segment inventory, and even though the idea that the same principles govern both the order of acquisition and typological distribution patterns was not always acknowledged explicitly in subsequent research, the order of acquisition question remained high on the agenda.

In this thesis, we will not have much to say about the order of acquisition. We will make little or no claims about the relation of order of acquisition and the inherent content of features, or their (relative) markedness. Rather than focusing on the order of acquisition of the segment inventory, we will extend the Jakobsonian hypothesis to the domain of the *phonological mechanisms underlying acquisition.* In short, one of the central claims of the current thesis is that the acquisition of the segment inventory can be described in the same terms as the structure of the adult inventory, where these terms are formalised as the Feature Co-occurrence Constraints in example (1) (chapter 1).

Having said that, it must be acknowledged that the (universality of) order of acquisition of the inventory remains a research topic in its own right. Perhaps the main reason for that is that the variation that has been shown to exist between children (both in the same and in different language communities) indicates that the order of acquisition of the inventory is not the most promising avenue for investigating the Jakobsonian claim; in addition, the order of acquisition question needs to be tackled employing a uniform methodology for analysing acquisition and adult phonology; a goal of the current thesis is to contribute to the development of such a methodology.[8]

1.3.1 The developmental problem of language acquisition

A highly influential study of the development of the segment inventory is Ferguson and Farwell (1975), who introduce the notion of Phone Classes, which

[8] With respect to the order of acquisition of the inventory, this is perhaps the moment to note that Levelt and van Oostendorp (2007) did not find that any of the three measures of distributional frequency is a good predictor, either. Hence, the factors determining the order of acquisition remain to be determined.

can be ordered into Phone Trees. Phone classes are groupings of variants of sounds (the authors focus on word-initial consonants, as do we), such that all variants occurring in the same word are grouped together in a phone class, together with all variants of words with the same initial consonant. Example (10) gives a phone class for one of Ferguson and Farwell (1975)'s subjects, T.:

(10)	*Phone class*	*Lexical items*
	[b ∼ β ∼ bw ∼ p^h ∼ ɸ ∼ ∅]	baby, ball, blanket, book, bounce, bye-bye, paper

As we can see, the adult target forms are all labial-initial words, and so are all the variants in the phone class (with the obvious exception of the null-realisation). Phone classes are construed for the entire lexicon, and by listing the phone classes on a horizontal axis, and subsequent recording sessions on a vertical axis, the longitudinal development of a child's inventory can be plotted. Phone classes in successive recording sessions are connected when they contain the same words, the result of which is a so-called Phone Tree (the graphic representation of Phone Trees is rather complex, or, in the words of Ingram (1989), they "look like the wiring diagram of a television set" (p. 201). For this reason, the reader is referred to either Ferguson & Farwell, 1975 or Ingram, 1989, p. 202 for examples).

Ingram (1989) notes a number of problems with phone trees, the most important of which is their sensitivity to variation of one single lexical item. Going back to example (10), there is no way of telling whether variation in the first column corresponds to different tokens of each type in the second column, or whether tokens belonging to one type exhibit more variation than those belonging to other types. Furthermore, with its focus on surface variability, it is difficult to see what phone trees are meant to represent other than exactly that, surface variability. With respect to the latter criticism, it should be noted that Ferguson and Farwell (1975) explicitly reject the continuity hypothesis and propose that the early lexicon (at least up to about containing 50 items) is organised no further than individual word forms (i.e., without any type of sub-segmental representational units or rules). Phonological organisation and representation is constructed by the individual during the course of linguistic development. If this is true, however, the question arises whether it is warranted to group into phone classes surface variants of adult targets if these occur in different lexical items.

Nevertheless, one of the successes of Ferguson and Farwell (1975) and successive applications of the phone tree methodology was that it showed systematically that individual differences in the order of acquisition exist within learners of the same language, and hence, that the strong Jakobsonian hypothesis must be amended.

A different method for analysing the acquisition of the segment inventory is developed by Ingram (1981, 1988, 1989). The aim of Ingram's program is virtually opposite to that of Ferguson and Farwell (1975): whereas the latter incorporate as much surface variability as possible, and reject the hypothesis that this reflects some underlying unity (although their method suggests otherwise), Ingram explicitly aims to filter out variability in order to obtain insights in the child's phonological competence.

Ingram (1989) lists a number of criteria for selecting types on which to base the assessment of the inventory. These are reproduced below:

1. If a phonetic type occurs in a majority of the phonetic tokens, select it.
2. If there are three or more phonetic types, select the one that shares the most segments with the others
3. If there are two phonetic types, select the one that is not pronounced correctly
4. If none of the above work, select the first type listed

Using these criteria as a heuristic, a representative sample of words can be selected, without incorporating incidental performance-induced variants. This is something that Ferguson and Farwell (1975) do not attempt, or consider to be relevant. The method proposed in Ingram (1981, 1988, 1989) is similar in spirit to the one we employ in the current thesis (see section 4.2). The developmental problem of language acquisition is also tackled frequently from the perception point-of-view. We will discuss a large body of literature below (see chapter 2.3)

1.3.2 The logical problem of language acquisition

As pointed out by Levelt (1994, p. 3), the logical problem of language acquisition can in principle be resolved without reference to actual child language data whatsoever. As noted before, in order to obtain cognitive credibility, any theory of the final state (adult grammar) must provide an account of learnability. Learning algorithms for Optimality Theory, for example, are proposals to resolve the logical problem of language acquisition (Tesar & Smolensky, 2000; Boersma, 1998; Boersma & Hayes, 2001), even if they are shown to be adequate predictors of actual language acquisition (Boersma & Levelt, 1999).

In section 2.2 we will see the learnability argument of Boersma and Hamann (2008), which pertains to the acquisition of the segment inventory. Perhaps the most important and influential solutions to the logical problem of language acquisition is the Modified Contrastive Hierarchy (Dresher, 2009), which employs the Successive Division Algorithm to arrive at the underlying structure of the segment inventory. We will devote considerable attention to both theories throughout the thesis, but a brief introduction is in order at this point.

The Modified Contrastive Hierarchy holds that only contrastive features are specified, in other words, that rules or constraints can only refer to contrastive features. Hall (2007, p. 20) goes even further, and formalises the 'contrastivist hypothesis':

(11) Contrastivist Hypothesis
The phonological component of a language L operates only on those features which are necessary to distinguish the phonemes of L from one another.

Superficially, the Modified Contrastive Hierarchy is similar to various forms of Underspecification Theory, and to be sure, it does incorporate underspecification. It goes beyond Underspecification Theory (Archangeli, 1984, 1988, among others), however, in that features may not be filled-in at some stage in the derivation to allow for certain rules to operate on them. The entirety of representational information is codified in the lexicon, and no more is codified than what is necessary to uniquely specify each member of the inventory.

The logical problem of language acquisition in this respect can be rephrased as the problem of algorithmically arriving at the correct feature specifications for the forms in the lexicon. After showing that the so-called 'pairwise algorithm' (the practice of determining contrast based on the comparison of minimal pairs) fails to deliver, Dresher (2009, footnote 3) proposes that contrast and feature specification is determined via the Successive Division Algorithm:

(12) Successive Division Algorithm

a. In the initial state, all tokens in inventory I are assumed to be variants of a single member. Set I = S, the set of all members.
b. i. If S is found to have more than one member, proceed to (c).
   ii. Otherwise, stop. If a member, M, has not been designated contrastive with respect to a feature G, then G is *redundant* for M
c. Select a new n-ary feature F, from the set of distinctive features. F splits members of the input set, S, into n sets, $F_1 - F_n$, depending on what value of F is true of each member of S.
d. i. If all but one of $F_1 - F_n$ is empty, then loop back to (c). (That is, if all members of S have the same value of F, then F is not contrastive for this set.)
   ii. Otherwise, F is *contrastive* for all members of S
e. For each set F_i, loop back to (b), replacing S by F_i.

In prose, the Successive Division Algorithm assumes an initial state in which all members of the inventory I are present, but allophonic to each other. Next, a feature is selected (e.g., [±voice]), and the set of inventory members is exhaustively assigned to either [+voice] or [-voice]. If the inventory contains no more than two members, the resulting subsets (designated by [+voice] and [-voice] respectively) each contain one uniquely specified member. However, if I contains more than two members, the resulting subsets must be further subdivided.

Take, for example, the subset inventory of labial stops in French (we adopt this example from Dresher, 2009, who adapts it from Jakobson & Lotz, 1949), which consists of three members: /p, b, m/. This subset can be specified by using two features: [±voice] and [±nasal], and hence, two different feature orders are possible [±voice]≫[±nasal], or [±nasal]≫[±voice]. Example (13) gives a graphical representation of the first option, together with the resulting feature specifications, and example (14) does the same for the second possibility (where branching to the left implies the negative value of the feature, and branching to the right implies the positive value).

(13) [±voice]≫[±nasal]

a. Feature hierarchy

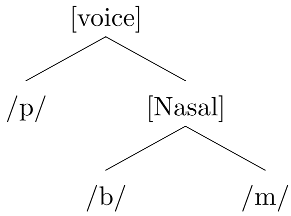

b. Contrastive specification

	p	b	m
[voice]	-	+	+
[nasal]		-	+

(14) [±nasal]≫[±voice]

a. Feature hierarchy

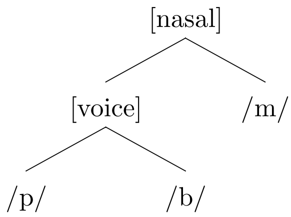

b. Contrastive specification

	p	b	m
[voice]	-	+	
[nasal]	-	-	+

For every set of n features, there are $n!$ possible orderings. The (hypothesised) learner and the linguist are thus faced with the question which feature to select first, which second, *et cetera*. Multiple strategies are possible (trial-and-error, parallel evaluation, to name but a few), but the crucial point is that the source of the evidence is always the same: contrast and phonological activity. In terms of examples (13) and (14), let us assume that the learner has found evidence that the three members of the subset inventory are contrastive. If there is evidence that /p/ functions with /b/ to the exclusion of /m/, the order in [±nasal]≫[±voice] is the correct one, as it assigns [-nasal] to both obstruents, whereas under the ordering [±voice]≫[±nasal] /p/ remains unspecified for [nasal] and hence does not form a natural class with /b/.

1.3.3 Syntheses

The deductive methods employed to tackle the developmental problem of language acquisition and the theory-driven approaches to the logical problem are both valid and valuable scientific enterprises in their own right, but ideally, the two should converge. We briefly mentioned Boersma and Levelt (1999) who show a convergence of the acquisition of syllable structure in Dutch and the predictions made by the Gradual Learning Algorithm (Boersma, 1998; Boersma & Hayes, 2001). Dresher (2009), too, mentions that the Successive Division Algorithm predicts the same feature hierarchy when applied to the inventory of Dutch, as the order of acquisition of distinctive features reported in Fikkert (1994), but it should be noted that we have no *a priori* reason to assume that the hierarchies derived by the Successive Division Algorithm represent the developmental order as uncovered in longitudinal child language studies. We will have more to say about this in chapter 5.

Having said that, a number of studies do tackle the developmental problem of language acquisition informed by hypotheses derived from the logical problem of language acquisition (Smith, 1973; Fikkert, 1994; Pater, 1997; Rose, 2000, among many others). Two of these specifically concern issues pertaining to the acquisition of the segment inventory and hence merit some discussion. Incidentally, both also concern the acquisition of Dutch.

Focusing on place of articulation (PoA) features exclusively, Levelt (1994) develops an analysis of the developing phonology of Dutch by means of a system of output conditions. In her system, features are specified only for words in their entirety, whereafter the left word-edge may become specified independently, and finally, the nucleus and right edge may receive independent specifications for PoA. Under this analysis, combined with a feature system in which vowels and consonants are specified by the same feature set (Levelt employs the model of Lahiri & Evers, 1991) Levelt (1994) shows that consonant harmony, a process previously analysed as long-distance assimilation (Vihman, 1978; Stemberger & Stoel-Gammon, 1989, 1991), is in fact subject to strict locality (i.e., adjacency) restrictions: the harmonising features are always shared between both the harmonising consonants and the intervening vowel. At the same time, specific output constraints limit the distribution of features over the various positions that become available to the child stepwise: initial dorsals are forbidden, and [labial] must be aligned with the left edge.

As we can see, Levelt (1994) is not concerned primarily with the acquisition of the inventory, but with her focus on the developing system of place of articulation features, she does provide a contribution. Furthermore, the investigations in Levelt (1994) are theoretically informed and data oriented, and as such, the author tackles both the developmental and the logical problem of language acquisition.

In a cross-sectional study, Beers (1995) analyses the acquisition of the Dutch inventory in both normally developing and language-impaired children. Like Levelt (1994), Beers assumes a model of feature geometry, and predicts that the acquisition of a system of contrasts that characterises the inventory proceeds in terms of a top-down fashion along the lines of that geometry. Higher-level features indicate more coarse contrasts (e.g., [labial] ∼ [dorsal] contrast is represented higher in the geometry then the [nasal] ∼ [lateral] contrast, and hence, it is predicted to be acquired earlier). Acquisition patterns are grouped in three categories: 'expected', 'unlikely', and 'abnormal', in descending degrees of convergence with the feature geometry. The multitude of children in the study display patterns that fall under the 'expected' ledger.

A major difference with the current study is that Beers (1995) focusses on the system of contrasts, rather than on the inventory proper. What this means is that when evidence of a contrast (say, [labial] ∼ [coronal]) at any manner of articulation reaches the pre-set criterion, this is sufficient for that contrast to be considered acquired. In other words, Beers focusses less on gaps, and more on the underlying representation. The difference with the current study is mainly that we propose, closer in spirit to Levelt (1994), that the acquisition of the segment inventory is best described in terms of a tandem development of underlying representations and output constraints, rather than focusing on the representational aspect of acquisition exclusively.[9]

As we discussed above, Levelt and van Oostendorp (2007) first proposed the idea that the acquisition of the inventory can be described in terms of the acquisition of features and constraints on feature combinations. We aim to contribute to resolving both the developmental and the logical problem of language acquisition – limited to the consonant inventory, that is.

1.3.4 Levelt and van Oostendorp (2007)

The current study is based on a pilot study published as Levelt and van Oostendorp (2007). Before exploring the formal properties and practical application to child language data in following chapters, we must briefly reflect on the similarities and differences between the Levelt and van Oostendorp study and the theory proposed here.

Apart from the theoretical coverage, which is both deeper and broader in the current study, and some minor differences such as the different feature sets, there are two more fundamental differences.

First, the difference in the features employed is more than minor in at least one respect, namely that Levelt and van Oostendorp adhere to a full three-way place specification, whereas we have adopted coronal underspecification. As we will see, this choice leads us to predict the early acquisition of /t/. It should be noted that Levelt and van Oostendorp assume underspecification of non-continuancy, as we do in the present thesis (that is to say, there is no feature

[9] Another difference is that we will have little to say about the order of acquisition.

[stop]). The issue of underspecification is not fully explored in the pilot study, presumably due to reasons of space.

Another difference relates to assumptions about the role and scope of FCCs. In the current thesis, we have adopted the minimalist view that there is no difference between language-specific gaps and universal restrictions on feature co-occurrence. Both are governed by FCCs. Levelt and van Oostendorp, however, note that there need be no FCCs referring to both [coronal] and [labial], as "... place features [...] exclude each other" (Levelt & van Oostendorp, 2007, p. 168).

In addition to the differences mentioned here, there are also some similarities. Both studies, for example, assume that features are privative. Both studies, too, adopt a definition of the Feature Co-occurrence Constraints that is essentially segment-driven. Consider the definitions given in Levelt and van Oostendorp (2007, ex. 5):

(15) *Assumptions about Feature Co-occurrence Constraints (FCCs)*
Where F, G denote features; FCCs in the constraint set Con are of the following type (only):

a. *[F,G]: No segment has both F and G

b. [F]⊃[G]: If a segment has F, it should also have G (no segment has [F] without having [G])

If we compare these definitions to the definitions in 1, we see that both assume the segment as the domain of the constraints. Levelt and van Oostendorp are less explicit about this than chapter 3, but in their case, too, we can read 'segment' as 'dominating the root node'. That is to say, the 'segment' in this thesis is nothing more than the simultaneous expression of phonological features, namely those that are dominated by the same root node.

The alternative would be to define the constraints with the feature as 'subject'. The definition for c-constraints would then read along these lines: "assign a violation mark to every feature [F] that is dominated by the same root node as feature [G]." The difference may seem one of little importance, but while the current, segment-driven definitions do not allow for asymmetrical c-constraints, the alternative would. One area where this becomes visible is feature spreading. Say a given form has two segments, A{[F,G]} and B{[F]}, and a constraint that enforces spreading of [G], such that [G] is dominated by both A and B, the constraint *[F, G] is violated by both A and B under the current definitions. A hypothetical *[G, F] behaves exactly the same and hence is redundant. Under the feature-driven view, however, even though there are two root nodes dominating both [F] and [G], the constraint *[F, G] is violated by both A and B, but its mirror image, *[G, F] is only violated once (remember there is only one instance of [G]). A more thorough exploration of the feature-driven view can be found in Linke and van Oostendorp (in preparation).

1.4 A brief history of thought on the segment inventory

The subject matter of the current thesis is the acquisition of the segment inventory, as we have seen. In the previous section, we discussed some earlier approaches to this problem, but before turning to the main matter of the thesis, it is important to reflect for a moment on the history of the study of the segment inventory *per se*.

The systemic, synchronic study of language begins, in modern times, with de Saussure. Distinguishing *langue* and *parole*, the *phonème* (speech sound) and sound images, de Saussure opened the door to an abstract approach to phonology. For our present purposes, which concern the phonological inventory, the most important contribution of de Saussure is his emphasis of language as a system: "*dans la langue, il n'y a que des differences [...] sans termes positifs*". In other words, the segment inventory is defined in terms of the contrasts it expresses, although Anderson (1985) emphasises that for de Saussure, the study of the speech sounds themselves is an important prerequisite to linguistic analysis, for it is only by knowing the positive definitions of speech sounds that we can study their relations. Thus contrast, or *oppositions*, became an important focus of phonology.

The inventory was of central concern for (American) structuralism. For American structuralists, in fact, the study of language was to construct inventories, not merely of phonemes, but also of morphemes, constructions, etc. (Anderson, 2000). The construction of these inventories was regulated by strict principles, of most concern to us at this point, the segregation of levels, and the bi-uniqueness criterion. According to the latter, for a speech sound to be considered a phoneme, it was required to stand in a special relation to its phonetic counterpart: For any phonetic event X, corresponding to phoneme A, every such event X corresponds to phoneme A. In other words, there is a unique relation between phonetic event and phoneme. Conversely, for any phoneme B, corresponding to phonetic event Y, every instance of B corresponds to event Y. Again, the relation is unique, hence the term bi-uniqueness. The bi-uniqueness criterion is a strict requirement, and various effects follow from it. It must be understood, however, that it does not entail that every phonetic detail or contrast is represented phonemically.

The importance of Jakobson's *Kindersprache, Aphasie und Allgemeine Lautgesetze* (Jakobson, 1941/1968) cannot be underestimated. Several important traits of modern linguistics come together in the work of Jakobson: first, it signals the beginning of feature theory (see also Jakobson, Fant, and Halle (1952); Jakobson and Halle (1956)). The introduction of feature theory not only allowed phonologists to study the internal structure of the phoneme; it also provided an apparatus to formalise natural classes, where a natural class is defined as the set of phonemes having a feature in common. Secondly, in addition to the Saussurian paradigm of linguistics as the synchronic study of language as a system, Jakobson proposed that the systems of individual languages were in fact not unrelated, but governed by universal principles, such as the set of laws of irreversible solidarity. These laws are implicational universal statements over the structure of the inventory, and can be re-interpreted as relative markedness generalisations ("if a language has X, it must also have Y"). Thirdly, Jakobson proposed to integrate the study of child language (and aphasic language) in the domain of general linguistics. Specifically, the laws of irreversible solidarity were intended to govern not only the sound systems of the world's languages (typology), but also (order of) acquisition, and loss of contrast in aphasia. Thus, Jakobson proposed an integrated theory of the segment inventory as a structured set of phonemes defined by distinctive features, and in addition, that language (or, at least, phonology) is driven by universal principles that make themselves known in typology, language acquisition, and language attrition.[10]

A true shift away from the primacy of the inventory was introduced with the advent of Generative Phonology (Chomsky, 1951; Halle, 1957, 1959; Chomsky & Halle, 1968). In (early) generative phonology, a view was propagated that the optimal way to explain linguistic structures was fundamentally derivational, or, in the words of Goldsmith and Laks (to appear),

> Of the four major tenets [...] of the SPE model, the most important was the view that the best explanation was algorithmic explanation. An algorithmic explanation is one which provides an account of the data which satisfies the conditions for being an algorithm: it is a fully explicit description of a process that can be carried out in a finite amount of time on a computational device such as a Turin machine or its equivalent.

A similar account is given by Anderson (2000), in his evaluation of the impact of Halle's The Phonetic Rules of Russian (Halle, 1957), later published as The Sound Structure of Russian (Halle, 1959). With the predominance of generativism in the ensuing years, phonology became derivationally oriented. In fact, SPE placed so much importance on rules and derivations, that positing new underlying segments was not an issue, as long as the system of rules would generate the attested surface form, and no more. The importance of the internal structure of the inventory was pushed to the background.

[10]With respect to aphasia, it should be noted that it is an abrupt condition, resulting from a stroke or lesion in the brain. It is not a progressive condition, and the mirror of child language is thus not proposed to be an individual's development (as it is for child language) but rather the states of individual patients ranked with respect to severity of the condition. In that respect, it could prove interesting to apply the Jakobsonian hypothesis to Altzheimer's disease, which is a progressive degenerative disease of the nervous systems, of which the progression of cognitive and functional symptoms is often described as mirroring the cognitive and functional development of early childhood, and which effects linguistic abilities as well as other cognitive and physical capacities (Blair, Marczinski, Davis-Faroque, & Kertesz, 2007).

This is not to say that the inventory no longer was considered; it was merely less prominent as an object of phonological study. In addition, it appears important to note that the introduction of feature theory, in combination with the developments in derivational analyses and the abandonment of the segregation of levels that was adhered to by structuralism, enabled phonologists to define segments and their contrastive and allophonic relations no longer with reference to phonetic correlates and systemic properties (contrasts), but also with respect to their phonological behaviour (morphophonology). To this day, phonologists do not agree on how to evaluate the different metrics to arrive at feature specification. Dresher (2009)'s Contrastive Hierarchy, for example, relies much on contrast, where contrast is defined in a specific way. For Morén (2003)'s Parallel Structures model, behaviour and economy are paramount, and the phonetic correlates are close to irrelevant. Conversely, much work in Element Theory (see for example Backley (2011)) holds that the behaviour and phonetic (acoustic) signature of a segment determine its subsegmental makeup, and contrast is almost epiphenomenal.

1.5 Overview of the thesis

Chapter 2 deals with the inventory, its description and its parts (features). We will begin by examining studies of the acquisition of the inventory and of features. A great deal of work has been done in this field, often focussing on the order of acquisition. This latter issue is of relatively minor concern to the current thesis; rather than testing a hypothesis about the order of acquisition, we provide a theory about the (former) mechanisms of acquisition. On the other hand, while the focus is often different, previous studies also had to provide some assumed or hypothesised formalism.

After discussing the inventory from an acquisitionist perspective, the chapter continues by exploring some aspects of the inventory *per se*. We will make a distinction between the *shape* of the inventory, and its *structure*, where the latter refers to the phonological organisation of the inventory, and the former to its phonetic implementation. Key aspects relating to the structure of the inventory are discussed: these include contrast, (under-)specification, symmetry and economy. Next, several theories about the shape and the structure of the inventory are reviewed: Dispersion Theory (Flemming, 2004), Parallel Bidirectional Phonetics and Phonology (Boersma & Hamann, 2008), the Modified Contrastive Hierarchy (Dresher, 2009; Hall, 2007) and expanding on that the theory of Dispersedness proposed in Hall (2011). These theories are elected because a) they explicitly concern themselves with the inventory and b) they are representative for the breadth of the phonological spectrum: from the fully functionalist to the formalist, from emphasising the shape of the inventory to focusing on its structure, and, via Hall (2011), back again. Furthermore, they each differ from the Feature Co-occurrence Constraint Theory in that the latter is anti-holistic: the inventory is fully epiphenomenal, it has no ontological status and no reference can be made to it by phonological rules or constraints. It is simply emergent from the system of freely combining distinctive features and filtering constraints. At the same time, we explore areas of overlap and conclude that some of the proposals reviewed are indeed compatible with Feature Co-occurrence Constraint Theory.

Next, we turn to the question of the developmental origins of features. The innateness debate is neither new nor settled, but after we acknowledge that the question is difficult to resolve empirically, we review a large body of literature. Rather than assuming that emergence is the default hypothesis and hence should be adopted in the absence of conclusive evidence indicating otherwise, we adopt a working definition of innateness, which states that a property is innate if the child exhibits knowledge of that property at or before the first encounter with it. This definition allows us to systematically examine the question of the innateness of features from three different perspectives, based on the different functional roles of distinctive features: categorisation/distinction, lexical organisation and phonological patterning. We find enough evidence to assume innate features.

Chapter 3 introduces the theory of Feature Co-occurrence Constraints, by outlining the Final State of the acquisition process of Dutch: the adult inventory. Only a small set of constraints is shown to be necessary, and no possible segments are either over- or underpreditcted. With these results, we proceed to an exploration of the formal properties of Feature Co-occurrence Constraints. It is demonstrated that the two constraint types are sufficient; and that it is neither necessary nor desirable for constraints to refer to more than two features. Single feature constraints can be modeled by assuming constraints such as *[F, F] and [F]→[F]. The latter is obviously always vacuously satisfied, but the former has a its own sets of effects, which we explore.

Next, we turn to the assumptions and implications regarding feature theory. Monovalency is defended, and so is universality. The constraint set introduced above is unable to rule out the empty segment, and hence the empty segment must be interpreted phonetically. There is a large body of literature on the non-specification of coronality and stopness; we assume indeed that both traits are not represented by distinctive features. The same goes for major class features. Rather than adopt Feature Geometry, we opt for Feature Classes. First, the latter is designed with Feature Co-occurrence Constraints in mind; secondly, and more importantly, it allows us to do away with major class features while retaining the insight that there are different *types* of features. Whereas Feature Co-occurrence Constraints do not reference these types but rather individual features directly (as per Feature Classes), they must be visible to the grammar or at the phonology-phonetics interface in order to avoid the situation where a segment containing, for example, [continuant] is also interpreted as a stop. In other words, feature classes are necessary to block the default interpretation when it is unnecessary. Also, it is shown how the Feature Co-occurrence Constraint Theory provides a way of implementing both inventory symmetry and

Feature Economy.

Acquisition of the inventory is the theme of **Chapter 4**, using data from the CLPF corpus of Dutch (Fikkert, 1994; Levelt, 1994). The main difference between the adult grammar and acquisition is that the latter entails a chronological component, whereas the former is assumed to be stable. Chronology is segmented into *stages*, which are defined by a change in the inventory. Constraints are invoked and can be revoked, but the time at which they are invoked is severely limited: no later than the time at which both features that the constraint refers to are acquired. This is borne out in the data. Next, the methods by which the data for this chapter are obtained are explained. As an example, we go through every stage of Noortje's actual productions. Her development is shown to be representative for the other children, and, other than in the case of Adult Dutch, some overpredictions occur. Considering all the children in the study, only six different segments (feature combinations) are ever overpredicted. This is not borne out for every child. Importantly, however, most instances of overpredictions concern cases where the segment is in fact in the child's inventory, but not represented in the sample because it is not robust enough to meet the inclusion criteria, which are subsequently discussed. Only four different segments are ever actually overpredicted. One of these is /t/, which, due to being represented as the empty segment, is predicted to be present in every inventory from the first stages. Three definitions are given for the contexts in which segments can be overpredicted, and the featural contents of overpredicted segments are discussed.

We conclude the chapter by considering the innateness of constraints, a difficult issue. For one thing, there is no way of knowing whether a constraint is present in the constraint set if its structural description has not been acquired yet. For example, the constraint *[F, G] has no effect on the grammar until both [F] and [G] are acquired. We conclude that while it is not unreasonable to assume that constraint templates are innate, there is no evidence in favour of innate feature co-occurrence constraints.

The main body of the thesis is concluded by **Chapter 5**. Some issues remain, which are dealt with there. We will explore to what degree the theory developed in this thesis is compatible with existing theories, such as Parallel Bidirectional Phonetics and Phonology (Boersma & Hamann, 2008), Inductive Grounding (Hayes, 1999) and the Modified Contrastivist Hierarchy/Successive Division Algorithm (Dresher, 2009; Hall, 2011). The theory of Feature Co-occurrence Constraints that is proposed owes much to these, and we see that especially in the first two cases, a high degree of compatibility exists. Finally, we will discuss the relation between emerging constraints and negative evidence, the different frameworks that the theory can be implemented in (mainly: constraint-ranking and non-constraint-ranking frameworks), and the role of perception.

Finally, **Appendices A, B** and **C** list for each child the succession of stages (A), the different FCCs that are activated during acquisition (B), and the inventories, features, constraints and overpredictions per child, per stage and per level of description (C).