Time and Phonology: Precedence-Based Representations

Rim Dabbous

A Thesis in the
Individualized Program
(INDI)

Presented in Partial Fulfillment of the Requirements
For the Degree of Master of Arts (Individualized Program) at
Concordia University
Montreal, Quebec, Canada

March 2023

CONCORDIA UNIVERSITY
School of Graduate Studies

This is to certify that the thesis prepared

By: Rim Dabbous

Entitled: Time and Phonology: Precedence-Based Representations

and submitted in partial fulfillment of the requirements for the degree of

Master of Arts (Individualized Progam)

complies with the regulations of the University and meets the accepted standards with respect to originality and quality.

Signed by the final examining committee:

______________________________ Examiner
Alan Bale

______________________________ Examiner
Dana Isac

______________________________ Thesis Supervisor(s)
Charles Reiss

______________________________ Thesis Supervisor(s)

Approved by ______________________________
Felice Yuen, Chair of Department or Graduate Program Director

____
Effrosyni Diamantoudi, Dean of Graduate Studies

Abstract

Time and Phonology: Precedence-Based Representations

Rim Dabbous

A major factor hindering the establishment of a successful neuroscience of phonology centers around the biological viability of a given phonological framework. The ultimate aim of this project is to find potential alignments between linguistics and neuroscience. In this vein, the main topic of the thesis rests upon establishing the minimal complexity requirements for a phonological representation that is biologically plausible, cognitively sound, and empirically motivated. Heeding Minimalist proposals (Chomsky, 1995) that encourage efficiency in computation and economy in representation, I embark on an in-depth exploration of the parameters of cognition that are necessary and sufficient in a phonological representation while discounting the processes and parameters that can be said to be "domain-general". To that end, I take seriously Ernst Pöppel's (2004) exhortation to consider the role of temporal events like *linear order* and *precedence* in the study of cognitive systems like phonology by surveying the literature on time perception. The conclusions support a separation of *order* from phonological representations, extending the scope of *substance-freeness* (Hale and Reiss, 2000) by characterizing *order* as substance. Such an approach can contribute to thoroughly defining the object of study and offer insight that narrows the search space for potential bridges.

Acknowledgements

I would like to express my deepest gratitude to my principal supervisor, Professor Charles Reiss, for his support, guidance, and encouragement throughout the course of this project. His exceptional insights and expertise as well as the opportunities I was afforded working alongside him have been invaluable in shaping the work presented in this thesis.

I would also like to thank Professors Alan Bale and Dana Isac for their ongoing support and generosity in imparting knowledge and insights that have been instrumental for my academic development and the completion of this thesis. Additionally, I am grateful to the staff at Concordia's Indi department for their support in managing logistics and administrative tasks related to the completion of my program.

I also want to express my gratitude to my family, friends and fellow students for the moral support and engaging conversations. I am especially indebted to Sarah, my unofficial therapist, for always lending an ear to my musings and rants; to Sam for his ongoing feedback and advice, and to Jenna for her much-appreciated levity when it was needed. Finally, my deepest gratitude goes to my husband and my children whose encouragement, support and patience sustained me throughout my academic career.

This research was supported by the Master's Canada Graduate Scholarship from the Social Sciences and Humanities Research Council of Canada(SSHRC), and the Concordia University entrance scholarship. To SSHRC and Concordia University, I submit my genuine gratitude.

Contents

List of Figures

Chapter 1

Introduction

> Space and Time, as conditions of
> the possibility of the presentation of
> objects, are valid no further than
> for objects of sense, consequently,
> only for experience. Beyond these
> limits, they represent to us nothing,
> for they belong only to sense, and
> have no reality apart from it.
>
> Immanuel Kant

Cognitive neuroscience is the scientific field that attempts to unify cognitive science and neuroscience by uncovering the biological mechanisms underpinning cognitive functions and representations (Posner and Digirolamo, 2000). Although existing linguistics models demonstrate a high degree of sophistication and predictive power, a cognitive neuroscience of language that links phonological computations to neurobiological mechanisms is still lacking. The problems facing cognitive neuroscientists in elucidating how the brain represents and carries out cognitive tasks reflect the divergent origins of the two fields, tracing back to the mind-body problem that sought to explain the nature of knowledge and its sources.

Cognitive science and linguistics historically originate from the “mind” answer to the mind-body problem and mirror the philosophical and methodological commitments of rationalists, namely, that reason is the source for knowledge and experience of reality is mediated by logical rules inherent in the mind, an immaterial entity that can be studied separately from its relation to the world. Neuroscience on the other hand, the study of the nervous system and its organization, historically stems from the “body” side of the mind-body debate and from an empiricist perspective favoring explanations based on the senses and basic laws of association as the source for knowledge (Fodor, 1981).

While the sophistication of linguistic models and neuroscientific methods continues to be refined, an explicatory mechanism linking linguistic computations to neurobiological processes has yet to be clearly and fully defined. The highly detailed maps of cortical activation that technical advancements in neuroimaging permitted disproportionately fixated interest on localization of function and spatial characterization, while overlooking descriptive mechanisms to ground cognition in the brain beyond a correlating capacity (Poeppel, 2012). As such, linguistic processes of the mind have yet to be adequately reduced to neurobiological processes in the brain. This modern manifestation of the mind-body problem has been recently articulated by Poeppel and Embick (2005) as two problems preventing a successful founding of a cognitive neuroscience of language: *The Granularity Mismatch Problem*, a conceptual mismatch between the levels of description employed in fine-grained linguistic approaches versus coarse-grained neuroscientific approaches, and the *Ontological Incommensurability Problem*, an ontological misalignment of the primitives of analysis available in neuroscience (i.e. dendrites, neuron, long term potentiation etc.) with those available in linguistics (i.e. features, syllables, noun phrase etc.) preventing matching/unification from one to the other. (Poeppel and Embick, 2005; Embick and Poeppel, 2014).

A major factor hindering a successful establishment of a neuroscience of phonology centers around the biological viability of a given phonological framework. The ultimate aim of this project is to find potential alignments between linguistics and neuroscience. In this vein, the main topic of the thesis rests upon establishing the minimal complexity requirements in a phonological representation that is biologically plausible, cognitively sound, and empirically motivated. Heeding Minimalist proposals (Chomsky, 1995) that encourage optimization and computational efficiency, I embark on an in-depth exploration of the parameters of cognition that are necessary and sufficient for phonology to delineate the required criteria in a phonological representation while discounting processes and parameters that can be said to be “domain-general”, such as time.

The experimental literature on time makes it apparent that at the lowest levels, perception of time is not veridical or in accordance with common sense. However, discussions of higher order processes like phonological rules and representations have ignored these insights, and discourse persists in terms of simplistic views of time perception and cognition.

To remedy this, I take seriously Ernst Pöppel’s exhortation to consider the place of temporal events in the study of cognitive systems by acknowledging the necessity of separating temporal experiences like *order* from the study of cognitive systems like *phonology*. Only when such a distinction is made can a cognitive process be satisfactorily defined, affirmed by Pöppel (2004):

> It has become clear, that perceptual or cognitive processes can only be understood if the dimension of time is taken more seriously. The reduction of complexity in neuronal systems is for instance, achieved by temporal integration mechanisms which are independent of the content of a percept or a cognitive act but are presemantical operations. (p. 295)

The goal of this thesis centers around using the insights offered in the experimental literature on temporal experiences like *sequence computation, linear order* and *precedence* to 'rebuild' phonology and prevent simplifications and errors concerning the lower levels of perception from being passed up to analyses of higher levels.

This new approach will provide evidence that certain parameters that have been assumed to be fundamental to many models of phonology like *linear order* are, in fact, epiphenomenal. Much like Hale and Reiss (2000) consider that features are *substance-free* and devoid of phonetic properties, I argue that a phonological representation is similarly devoid of temporal properties which are, in fact, substance. If successful, the current work will offer simpler and more economical theoretical alternatives that are better suited to address the granularity mismatch and ontological incommensurately problems. Such an approach can contribute in thoroughly defining the object of study and, in principle, narrow the search space for potential bridges that can unify neuroscience and linguistics.

This thesis will unfold as follows. Chapter 2 provides a brief overview of the field of phonology, as well as an introduction to a particular phonological model which includes fundamental operations that can account for segmental alternations, both locally and at a distance.

Chapter 3 discusses the nature of a *mental representation* in general, building the definition from the ground up to clearly establish an explicit characterization of this loaded term and provide the relevant terminology for exposition in future chapters where the representations of concern are phonological structures, e.g. strings of segments or ultimately syllabified strings of feature bundles.

Chapter 4 outlines the experimental landscape surrounding the perception and cognition of time to motivate the claim that temporal order is divorced from phonology. Readers interested only in the phonological model, and not its grounding in cognitive science, can skip chapters 3 and 4.

Chapter 5 presents a phonological model featuring temporality-free non-linear phonological representations based on (logical) precedence, along with an adaptation of the SEARCH and CHANGE functions which lessens the computational burden of SEARCH by virtue of a new representational apparatus, as well as by relying on domain general temporal mechanisms. Finally, chapter 6 contains discussions and conclusions.

Chapter 2

Phonology

Consider the Turkish words *ev* ‘house’ and *dev* ‘giant’ whose plurals are *evler* and *devler*, respectively. A simple cross-comparison suggests that the plural marker is likely *ler*. Now consider the word *at* ‘horse’ that has the plural *atlar*. This new data throws a wrench in our hypothesis, providing two possibilities for a plural marker to consider, *-ler* and *-lar*. Upon examination of more singular/plural data like *cep/cepler, kek/kekler, can/canlar, kap/kaplar*, this alternation between *ler* and *lar* is determined to be the result of the nature of the vowel in the root, namely ‘e’ or ‘a’. This is known as ‘vowel harmony’, a pattern endemic to multiple natural languages such as the above examples from Turkish.

The study of these patterns of sound and their generation is called phonology, which incidentally, is also the name of the mental organ which computes these patterns. A *phonology* is a computational mental system that transforms an underlying phonological representation into a surface representation (and vice-versa) to be translated into motor commands during production or to link between sound and a discrete feature-based representation during speech perception.

A variety of phonological frameworks exist, generally sorted into two main categories: rule-based derivational and constraint-based (like Optimality Theory). While differing over the fundamental assumptions they make to account for the sound distributions observed across languages, they no less agree that the patterns are an outcome of transformations from underlying representations to surface representations. The approach my model departs from —the SEARCH & CHANGE—builds upon the derivational framework and substance-free assumptions of Hale and Reiss (2008), as well as the set-theoretic approaches to phonology outlined in Bale and Reiss (2018) who provide a mathematically motivated and computationally explicit framework to model phonological patterns, two necessary requirements outlined by Poeppel (2012) in the quest to bridge linguistics and neuroscience.

In this chapter, I will start with a brief description of what constitutes phonological transformations and representations. Next I present the SEARCH and CHANGE model which forms the basis for my own model. Finally, I provide two implementations for the SEARCH, an iterative one and a recursive one to showcase the computational requirements for each, and appeal to enriching the representation to offset the trade-off in complexity.

2.1 Transformation and Representation

The two components of a *phonology* are the transformations—the rules that map input to output; and the input—the phonological representations over which the rules apply.

2.1.1 Representation

Let's consider representations first. A closer examination of the Turkish example hints at more complexity underpinning the patterns. Recall that the plural is formed by appending the plural suffix with an alternating vowel depending on the vowel in the root: the plural of the word *ev* 'house' is *evler*, but the plural of *at* 'horse' is *atlar*. One might hypothesize that the vowel in the suffix is copied from the root. However, this pattern is not simply a matter of strict identity since all roots with vowels articulated with the tongue placed forward, like [i,y,e,ø] take *-ler* as plural (*ipler* 'ropes') while all roots with vowels articulated with the tongue placed towards the back, like [ɨ,u,a,o] take *-lar* as plural (*pullar* 'stamps'). This suggests that the patterns in Turkish plurals are described in terms of more basic units than whole segments.

Now let's consider an example from English. Plural formation in English consists of appending the plural suffix /-z/ onto a root in either its voiced [z] or voiceless [s] manifestation depending on the voicing status of the final sound of the root.[1] The words *cat* and *dock*, ending in voiceless segments ([t] and [k]), license the plurals *cats* and *docks*, pronounced [kæts] and [dɔks], while *pad* and *dog*, ending in voiced segments ([d] and [g]), form the plurals *pads* and *dogs* pronounced [pædz] and [dɔgz].

These examples show that the basic symbols of phonological computation are not the speech segments themselves, but rather the distinctive features that make up these segments (voiced/unvoiced, front/back etc.). Features are therefore the atomic primitives in terms of which the the harmonic agreement in Turkish and voicing agreement in English are computed.

---

[1] This specific sound pattern similarly holds for the possessive -z and the 3rd person present singular agreement -z.

It is critical to note that -Z devoicing does not happen in Turkish and that vowel harmony does not happen in English. These sound patterns are unique to these languages, hence the rules cannot be derived from phonetic mechanisms. The features, while descriptively correlating with the articulatory properties that define them, are in fact symbolic and content-less. Hale and Reiss (2008) describe phonology as "substance-free" and phonological segments as divorced from their phonetic "contents". Thus a (voiced) segment which is specified +VOICE does not encode voicing by the vibrations of the vocal tract, rather it is a contentless symbol which gets translated into a command to vibrate the vocal tract by transducers.

A phonological segment is thus a bundle of symbolic features bound simultaneously in time to form a distinct structure. However, since the agreement is calculated with respect to features of other segments, a phonological representation requires at least another segment that determines the proper environment. The English plural /z/, for instance, requires information from the segment which precedes it which in turn, informs how it surfaces. Similarly, the vowel segment in the Turkish plural marker requires information from the vowel that precedes it in the root.

Characterisations of phonological representations are therefore two-fold; on the one hand are the simultaneous or *static structures*, the segments comprised of features bound together in time, and on the other, the *dynamic structures*, a collection of at least two segments that are sequentially ordered. The topic of this thesis are the dynamic structures that present over time, and I will thus abstain from outlining the particulars of features to simplify my exposition of dynamic representations by only referring to segments when possible while acknowledging that the account I provide remains simplified due to the omission of features. [2]

---

[2] For a thorough mathematical account of features and the operations that alter them, refer to Bale and Reiss (2018).

2.1.2 Transformation

Let us now consider transformations. A transformation is a mapping from an underlying phonological form to a surface form. Generativity experiments in which participants produce outputs to novel stimuli that are consistent with phonological rules show that representations are stored in their most abstracted version, i.e. without any details that do not contribute to a contrast in meaning. In the English plural example above, the plural morpheme surfaces as either [s] or [z] depending on some property of the final root segment. Both forms need not be stored in memory however; the existence of computations to generate the surface forms accounts for the observed outputs without needlessly storing both forms. Thus only /-z/ is stored and the [s] computed in the phonology.

Chomsky and Halle's (1968) "The sound patterns of English" (henceforth SPE) provides a rule template which describes the structural changes that transform an underlying phonemic representation into a phonetic surface form in a manner that accounts for the patterns observed across many languages. A rule, also known as a *map* or *transformation*, takes as input a string of symbols, the underlying form, and outputs another string, the surface form.[3] For example, in English plural formation, such a rule can be listed as the following:

- $\text{z} \mapsto \text{s} / \text{[-V\textsc{oice}]} \_\_$

This rules states that the segment /z/ becomes an [s] when it is directly preceded by a voiceless segment.

However, as the Turkish example shows, some segments targeted by rules occur at

---

[3] It is important to note that the description I present is simplified and tailored for the informational requirements of this thesis; Rule application is more complex than a simple rule changing an UR to SR since some rules apply to the output of other rules, thus intermediary representational levels are available.

a "distance" from the conditioning environment. In the plural *atlar* of *at*, the vowel of the suffix surfaces as [a] as a result of the nature of the vowel of the root, which does not directly precede it. Phonological rules thus require sensitivity not only to local environments, but also "long-distance" ones. In SPE, the use of the Kleene star achieves this by describing the segments to be skipped. For the Turkish plural, a simplified version of the rule would be stated as such:

- a ↦ e / e (C)* __ Where C stands for consonant.

This rule states that an /a/ surfaces as [e] when it is preceded by /e/ with the possibility of any number of consonants—zero or more—intervening.

While phonology seems to require two segments, the segment undergoing the change (target/effect) and the environment conditioning the change (trigger/cause), most phonological frameworks operate over the assumption that a phonological representation (input to the rule) is a string of segments. This type of data structure seems suitable for so called *canonical* phonological processes like *assimilation* where some feature in a segment assimilates to the feature of the segment adjoining it. *Locality* in such processes seems over-represented when contrasted to the non-local processes. Adjacent phenomena have therefore been considered to be the baseline in phonological computation. But as the Turkish example can attest, phonology needs to be equipped with the means to compute over "long-distances".

To account for non-canonical phonological patterns, either the computations or the representations need to be enriched. An abundance of models have been proposed that enrich the computations. One such account is the Search and Change (Dabbous et al., 2021) which develops the 'Search and Copy' model outlined in Mailhot and Reiss (2007) and Nevins (2010) to cover more empirical ground.

2.2 Search and Change

Search and Change comprises two processes that take as input a string and output another string. The first sub-part is the Search wherein an initiator segment (INR)–coinciding with the target of traditional rules–launches a search in the string for its respective terminator segment (TRM)–coinciding with the trigger in traditional rules–in a specified direction. If successful, the INR will proceed to the Change which alters the INR using set theoretic operations depending on additional conditions to be satisfied. Accordingly, a phonology is a set of INR, TRM pairs specifying which segments can initiate Searches, and which segments satisfy the criteria for terminating a Search; as well as a Change which specifies the changes to the INR and outputs the new one.

Two main assumptions distinguish the Search and Change from other models. The first one is the 'non-locality assumption' which considers that phonological rules are actually built upon unbounded search procedures, and adjacency is a 'special case'. The second assumption is the *substance-Freeness of Structural Changes* (SFSC), which breaks a phonological rule's environment down into two distinct sub-parts: the WHAT and the WHERE, stated as the following:

> "Principle of Substance-Freeness of Structural Changes (SFSC): The features added to segments by the application of a rule need not be found in the rule environment"
>
> (Dabbous et al., 2021)

This breakdown offers an economical yet computationally explicit model that accounts for a wide range of phenomena by the combination of the following basic parameters: the specifications of the INR, the specifications of the TRM, the direction of the Search (DIR), the specifications of the Change, and conditions on the Change.

2.2.1 Toy Examples

Consider a language with six symbols, each defined by a shape and a color. This language has an alphabet Σ = { ■, ●, ▲, □, ○, △} and the following rule:

(1) Rule 1: Simple example

- Turn a ■ into a ○ if there is a ▲ anywhere to its right.
- Rule 1 Parameters

(where?) Search: INR: ■, TRM: ▲, DIR: R

(what?) Change: INPUT: INR, OUTPUT: ○

In rule (1), the specification of INR is ■, the specification for DIR is R (which means that the INR will search rightward for a TRM), the specification of TRM is ▲ and the specification of the Change is change the INR into a ○.

If we apply rule (1) to the string in (2), then the black squares in positions 2 and 5 fit the specification for INR and will launch a Search rightward, 'see' the same TRM, the black triangle in position 7, as the arrows in (2) show.

(2)

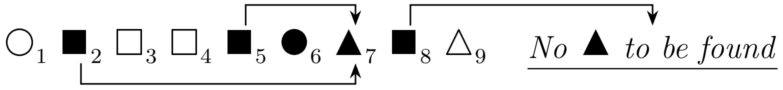

Crucially, while this process is generally described as "regressive harmony" in the literature where a 'donor' segment parses a string right to left to find a recipient, in this system the INR (recipient) scans the string rightwards (left to right) for a TRM (donor).

Rule (1) thus performs a mapping like (3) with this input string.

(3) ○$_1$ ■$_2$ □$_3$ □$_4$ ■$_5$ ●$_6$ ▲$_7$ ■$_8$ △$_9$ $\rightsquigarrow$ ○$_1$ ○$_2$ □$_3$ □$_4$ ○$_5$ ●$_6$ ▲$_7$ ■$_8$ △$_9$

Using the above parameters as well as partial descriptions, an abundance of phonological patterns like transparency, opaqueness, icy targets and adjacency can be accounted for. I refer the reader to Dabbous et al. (2021) for a thorough account of the combinatorial possibilities, however, I will provide an example of how adjacency, under this account, is a 'special case' that results from the combination of providing a partial specification for a TRM as well as conditions on the Change.

Adjacency is a special case of long-distance Suppose we want a rule similar to rule (1), with the difference of changing a ■ to a ○ *only if* there is a ▲ to its immediate right. In this system, adjacency requires more specification than long-distance environments. However, the Search model need not be more complex—adjacency conditions can be easily captured with the available machinery simply by partially describing the TRM. If a TRM in a rule is specified as SHAPE, the Search will terminate on the first segment that satisfies this condition, thus any symbol it encounters, regardless of its shape or color. Rule (1) can be altered to the following new rule which provides the necessary specifications for the required mapping:

(4) Rule 2

- From a ■ Search for a SHAPE to the right and turn the ■ into a ○ if TRM is ▲
- Rule 2 Parameters

  Search: INR: ■, TRM: SHAPE, DIR: R

  Change: INPUT: INR, OUTPUT: ○, CONDITION: TRM: ▲

In rule 2, the specifications for the TRM is the maximally general category SHAPE. A black square launches a Search, and stops at the first segment specified as SHAPE, which happens to be any segment, thus terminating at the first segment it meet.

However, since the ‘black triangle’ specification is now a condition on the Change, the INR will only become a white circle if the first segment to its right is a black triangle.

Consider rule (2) applied to the following string in (5):

(5) Mapping from (4)

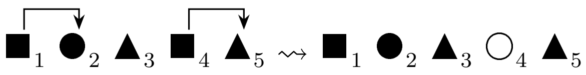

The black square in position 4 satisfies the INR specification. This INR finds the TRM in position 5, which is a black triangle, and position 4 is rewritten as a white circle. However, the black square INR in position 1 initiates a Search that terminates on the black circle in position 2. Since this TRM does not meet the condition of the change (it is not a black triangle), the INR in position 1 fails to become a white circle.

Within this system, a TRM located at an unspecified and unbounded distance from an INR can be found either by iterating through the string, which requires keeping track of indices, or alternatively, via recursively structuring the string such that INR,TRM pairs become structurally adjacent. In this latter case, a SEARCH can be construed as a representation instead of a computation.

2.2.2 Iterative SEARCH

Under an iterative account, the input to the SEARCH is a string, i.e. an ordered set of segments, each bearing an index which specifies its position. The Search and Change consists in transforming a string $< x_1, x_2, x_3, ..., x_n >$ of length n to the string $< y_1, y_2, y_3, ..., y_n >$.

Suppose a given segment x_i is an INR and the segment x_{i+k} the terminator. [4] The SEARCH will operate on the string starting from position i, and incrementally stopping at $i+1$, $(i+1)+1$ until it reaches $i+k$ such that $k = 1+1+1...+1$.

The following flowchart demonstrates the logic:

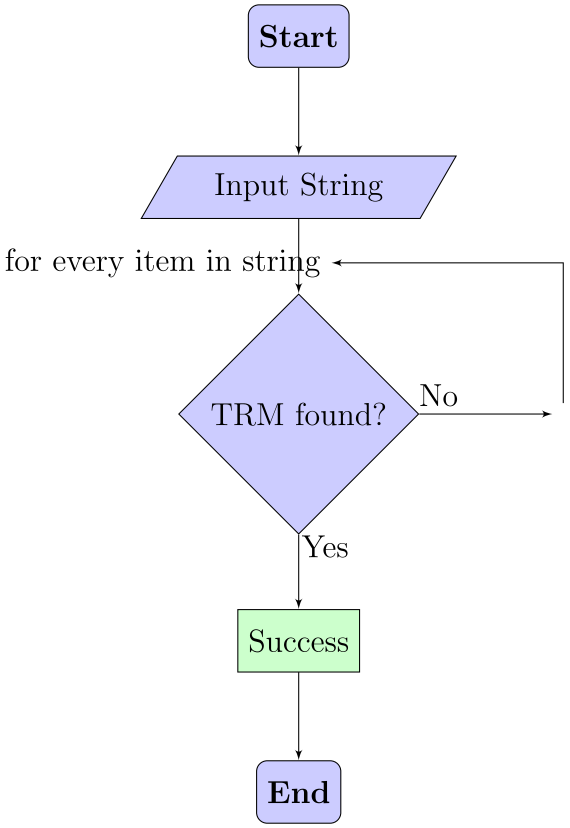

Consider the string *badac*. Suppose there is a rule which changes a b into c when there is a c anywhere to its right ($b \mapsto c/\_\_(\text{SEGMENT})^*c$).

The parameter specifications for this rule are the following:

---

[4] the direction of this SEARCH is specified by the addition in the index as a right-sided SEARCH. In a left-sided SEARCH, the TRM will bear an index that is less than the index of the INR, namely the TRM will be x_{i-k}.

Rule 3:

- Turn a b into a c if there's a c anywhere to its right.
- Rule 3 Parameters

  SEARCH: INR: b, TRM: c, DIR: R

  CHANGE: INPUT: INR, OUTPUT: c, $\boxed{\text{CONDITION: no condition}}$

According to this account, each segment of the string will carry an index specifying its location. Therefore $x_1 = b$, $x_2 = a$, $x_3 = d$, $x_4 = a$, and $x_5 = c$. Since b is an initiator, the segment x_1 launches a search in the string. In this particular example, it so happens that the INR is at index 1, which means that the SEARCH will start from x_1:

$x_1 = b$: Set $x_1 = INR$

For each i in the range of $(2,\ n)$ such that $x_1 = INR$ and $n =$ length of input, if $x_i = c$, set $y_1 = c$, else $y_1 = x_1$.

1. $x_{1+1} = x_2 = a$. Since $x_2 \neq c$, move onto x_{2+1}
2. $x_2 + 1 = x_3 = d$. since $x_3 \neq c$, move onto x_{3+1}
3. $x_3 + 1 = x_4 = a$. since $x_4 \neq c$, move onto x_{4+1}
4. $x_4 + 1 = x_5 = c$. since $x_5 = c$, set INR $y_1 = c$

The output of this change will be the string *cadac*.

2.2.3 Recursive Search

A recursive SEARCH can be modeled by a context free grammar which takes as input a string and outputs a hierarchical representation in which the INR and TRM are structurally adjacent.

Suppose a language has the same rule (3) which changes a b into c when there is a c anywhere to its right ($b \mapsto c/\_\_(SEGMENT)^{*}c$). According to this rule, $< b, c >$ form an initiator/terminator pair.

Let be $G =< V, \Sigma, R, S >$ be a context free grammar. Let S be the start symbol, Σ be the set of terminal symbols $\{a, b, c, d\}$, and V be the set of non-terminal symbols $= \{X\}$, and R the following production rules:[5]

$$\begin{array}{lll} S & \mapsto & INR \quad TRM \\ INR & \mapsto & b \\ TRM & \mapsto & X \quad TRM \\ TRM & \mapsto & c \\ X & \mapsto & a, d \end{array}$$

Consider the same string *badac*. It will be transformed by the SEARCH into the following structure:

---

[5] Note that this is a simplified example which parses an input only from the first INR encountered and ending at the first TRM that satisfied the criteria, while disregarding material before the INR and after the TRM.

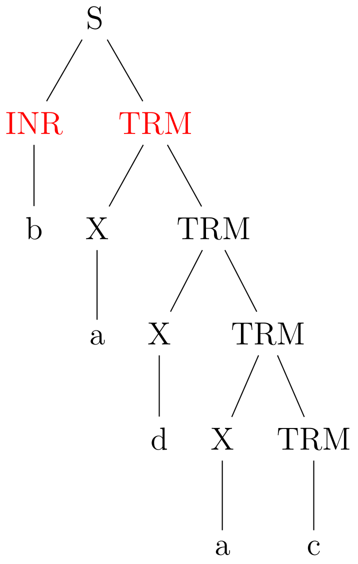

Crucially, the semantic rules of interpretation will consist of a mechanism which percolates the value of the TRM until the semantic value of the red TRM node is equal to that of the lowest. The SEARCH in this case consists in casting INR TRM pairs adjacently to be fed to the CHANGE module, which in turn acts on this simplified input, and outputs a new value for INR if the CHANGE conditions are met.

2.3 Tradeoff between representation and computation

Although this model provides a simplified account that covers a wide range of phonological processes, it is no less computationally demanding. Andersson et al. (2020) provide an analysis which describes the 'Search and Copy' system upon which SEARCH and CHANGE is based as computationally inefficient. However, as mentioned previously, to account for *non-canonical* phonology, either the computations—or alternatively, the representations—can be enriched. Considering the trade-off in complexity between a data-structure and the computations over it, exploring different representations can potentially offer frameworks where phonological complexity can be reduced by finding the optimal trade-off between representation and computation. However, while any model can offer parsimonious coverage of the empirical landscape, biological plausibility remains the end-goal for cognitive-neuroscientists interested in finding potential links between cognitive models and how the brain implements them. As such, a thorough and in-depth investigation of the parameters of *a representation* and the relevant cognitive properties to be included or discounted in a given representation can offer insights into biological plausibility of distinct models and shed light on the required complexity in a given system to distinguish its requirements from those found across cognition as a whole.

To that end, in what follows I will endeavor to describe a *mental representation* from the ground-up in an attempt to determine the minimal requirements for a phonological representation.

Chapter 3

Representation and Reality

> To think is to forget differences, to generalize, to abstract. In the teeming world of Funes, there were only details, almost immediate in their presence.
>
> *Funes The Memorious*, Jorge Luis Borges

Jorge Luis Borges' Ireneo Funes is a fictional teenager suffering an injury that gave him the capacity to eidetically remember every experience. He spends his days laying in his cot in a dark room, sometimes reconstructing entire past days with a level of detail that renders the construction analogous to reliving the past; Other times, he attempts to create languages in which each instance of an object has its own name. The narrator recalls that "not only was it difficult for him to see that the generic symbol 'dog' took in all the dissimilar individuals of all shapes and sizes,it irritated him that the 'dog' of three-fourteen in the afternoon, seen in profile, should be indicated by the same noun as the dog of three-fifteen, seen frontally". Funes claims that each leaf in each tree in each forest is different at different times, licensing each instance across time to be uniquely identified. He discards this endeavor as impossible, however, due to the fact that immortality would be a requirement.

Borges's reflections on the nature of thought and contents of knowledge are not novel. Philosophers and cognitive scientists alike dating from the pre-Socratics onward have held similar epistemological concerns regarding knowledge of the "essence" of things. While fictional, the narrator articulates a related concern, a concept fundamental for understanding the mind and its operations: Mental representation. In its technical definition (within cognitive/computational frameworks), a representation is a knowledge or information structure in the mind. Representation, along with the closely intertwined concept of computation– the procedures operating on representations–form a fundamental foundation for cognitive science, the study of the mind and its operations (Thagard, 2023).

The term *representation* however remains loaded. In what follows, I attempt to define the term from the ground up to determine the criteria that most adequately characterize it. I will start by outlining its necessity as a mediator with reality. Next I sketch out some of the prominent views of what constitutes a representation, namely, whether a representation is analog or symbolic, and provide arguments that representations comprise discrete symbols. Finally, I attempt to relate the symbolic aspects of a representation with the concept of *modularity* to determine what properties a particular representation should encode.

3.1 Representation as Mediator between Mind and World

To unambiguously “know” something, Funes eschews invariance and believes one must specify each instance of the thing, its spatial configuration, and its temporal tag. Funes’ requirements for a relation to the external world is thus an analogical one wherein each state of affairs in the world requires a one-to-one mapping with some physiological reaction and mental state without loss of information. While impossible due to the finiteness of the human brain, this quest for the conciliation of subjective experience with objective reality no less raises a very important and contentious question that philosophers and scientists have long grappled with: How is knowledge of reality achieved.

Funes’s proposed relation to reality is a “direct” one, supposing that humans’ experience is unmediated by any mental processes that might structure the stimuli engendered by the external world in any way. While such a radical lack of invariance would make it impossible to generalize and thus “know” anything, similar proposals by-passing the necessity of mental representations have been made. Proponents of embodied cognition continuing in the phenomenological traditions of Husserl, Heidegger and Merleau-Ponty, for example, took inspiration from ecological psychology (e.g., Gibson, 1979; Palatinus and Michaels, 2014) to conclude that properties of the environment are directly experienced/perceived by the organism as a whole. While a full review of the literature on the existence or lack-thereof of mental representations is beyond the scope of this thesis, I will no less provide a simplified account of why such views are untenable.

3.1.1 The Problem of Reality

A "direct experience" of external reality, one unmediated by any mental constructs, presupposes the existence of an external reality, and furthermore that this reality is faithful to our experience of it. Two arguments demonstrate the intractability of this latter claim: 1- the contents of experience (phenomenal knowledge) do not match what is scientifically known about reality (relational knowledge) and 2- any attempts to define reality outside our perceptions will inevitably lead to circularity.

Mismatch between the relational and the phenomenal

Frameworks operating over the assumption that experience of reality is isomorphic to external reality fall short when one considers how lacking our phenomenal experience of reality really is. But as the linguist and anthropologist Edward Sapir points out, "no entity in human experience can be adequately defined as the mechanical sum or product of its physical properties…it is notorious how many of these physical properties are, or may be, overlooked as irrelevant" Sapir (1949). Jackendoff (1991) lists three cases that aptly demonstrate how disparate external reality is from our experience of it: under-representations, illusions, and constructions.

Under-representations: Under-representations occur when parts of what is known scientifically about external reality fail to figure in our perceptual experience. For example, while the existence of ultraviolet or infrared light is scientifically established as a fact, humans lacking visual receptors capable of picking these wavelengths fail to experience them.

Illusions: The second set of examples Jackendoff mentions are illusions. As anyone familiar with illusions can attest, the fascination that these engender results from the dissonance between our expectation of reality and our subjective experience. Illusions reveal how sensory information received from the external world is organized and interpreted. Famous examples include the Müller-Lyer illusion (figure 3.1) and the Ames room (figure 3.2). In the Müller-Lyer illusion the direction of the arrowheads influences the perception of the length of the lines, one being perceived as longer than the other, in stark contrast to their length outside our direct experience which is the exact same. The Ames room (figure 3.2) is designed such that while one corner is farther than another, equidistant perception no less occurs from a specific angle. The room is not a square, but is perceived as one, causing the scale mismatch that results from the varying distances. These illusions demonstrate how the rules of vision can be manipulated to construct our spatial experience.

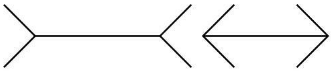

Figure 3.1: Müller-Lyer Illusion

Figure 3.2: Ames Room

Constructions: Finally, the existence of a mediating interface between reality and its experience is solidified when one considers cases where the physical stimulus doesn't correspond to our phenomenal experience. Examples demonstrating the breadth of the construction of our experience of reality includes amodal completion in vision and audition wherein an object is perceived (seen or heard) despite parts of it being occluded (figure 3.3) or constructed from a stimulus lacking it (figure

3.4); so called phonemic restoration effects in which an obscured sound segment is filled in by the mind regardless of its absence in the signal (Warrren, 1970); or even the bi-stability of a Necker cube (figure 3.5) where the same stimulus engenders an ambiguous alternating perception (the red dot can be either on the front face or the back face).

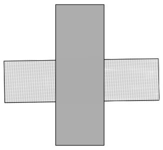

Figure 3.3: Amodal Completion

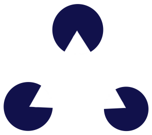

Figure 3.4: Illusory contours

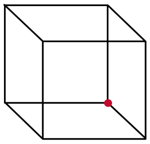

Figure 3.5: Necker cube

The mismatch between reality and the experience of reality is best exemplified by the hilariously tragic fate of beetles facing extinction due to the trickery their minds play in interpreting visual stimuli. In his *Interface Theory of Perception*, Hoffman (2016) explains that the beetles' mating call in the form of a dimpled body shares a property with beer bottle stubs, leading the male beetles to misconstrue beer bottles as female beetles ready to mate. The beetles are only sensitive to this particular property to guide the mating behavior and consequently constructed a reality completely different than that of humans or other species, further proving that the parts of reality accessible to an organism are relative to evolutionary pressures guiding a species' survival.

3.1.2 Circular definition: *Reality* is whatever is experienced as reality

From a more formal standpoint, attempts to demonstrate a one-to-one correspondence between an external reality and our experience of it will inevitably end circularly. The relation between reality and an agent inhabiting it as conceived by ecological psychologists consists of an exchange of information between "reality" and the system/agent whereby "visual invariant ecological properties" are picked-up and transduced by the system (Gibson, 1979). Aside from the problems arising from the assumption that environmental stimuli contain invariant properties, [1] the notion that invariant properties can be "picked-up" by a system would necessarily demand that said system be equipped with the means to accept/parse this class of properties. As Fodor and Pylyshyn (1981) point out, "there is a circle of inter-defined notions here, a directly detected property being one to which a transducer responds, and a transducer being a mechanism that responds directly to the properties it detects" (1989).

This line of arguments extends to any information processing system (IPS). An IPS can only accept inputs in terms of the representational apparatus inherent to it, and thus reality can only be defined in terms of the representational language that the system is equipped with (Hammarberg, 1981). "Raw" data are inaccessible to an IPS, thus "any IPS, including one that is a sentient being—is a prisoner of its own representational processes: we can never escape a point of view" (Hammarberg, 1981). Reality can only be defined by the system that apprehends it, rendering any claims of a direct correspondence between an external reality and experience impossible as reality "as" is cannot be known outside the limits of the perceptual/cognitive apparatus we

---

[1] For a more thorough discussion about the problems arising from assumptions of invariance in speech stimuli, refer to Appelbaum's 'The problem of Lack of invariance' (1996).

are endowed with.

In fact, the mismatch between reality and our perception of it is even more pronounced than previously outlined when one considers modern scientific advancements that posit a very counter-intuitive model of time and space (more in next chapter), leading to the need for more fine-grained distinctions in order to appropriately delineate our objects of study. The inaccessibility of external reality substantiates the necessity of a Kantian approach in the quest to uncover the mind. Emphasizing the existence of mental categories that interpret an unchanging and inaccessible external world, a Kantian approach places the contents of experience at the forefront of inquiry. When confronted by the prescient knowledge that “Space and time are the framework within which the mind is constrained to construct its experience of reality” (Kant), a statement that is vindicated in light of the scientific advancements in theoretical physics which characterize time and space as solely dependent on a frame of reference and matter on a conscious observer (Hammarberg, 1981), the only conclusions we can draw about our interaction with reality are limited to the manner by which our mind constructs it and the categories it imposes to interpret it, an approach that Jackendoff calls the “psychological approach” (1991).

3.2 Analog versus Symbolic Representations

As the above examples demonstrate, the mind is necessarily equipped with categories or “representations” by which the world is parsed. This raises a new set of questions regarding the nature and contents of these representations, the types of operations they license, and the interplay between both.

Descriptions of representations generally characterize them as “stand-ins” for states of affair in the world, some additionally positing that the structure of the distal system is maintained in the representation. In a famous series of experiments probing visual imagery by Kosslyn et al. (1978), participants were asked to memorize a map marked with different locations and asked to mentally scan from one location to the other. Elapsed time to the scanning was recorded and was found to systematically vary in proportion to the spatial distance in the stimulus. Locations which were farther apart on the map took longer to scan. This led researchers to conclude that the representation of space contains spatial properties (Kosslyn, 1980).

While similar to the direct view advocated by ecological psychologists and proponents of embodied cognition, this "pictorial" or "analog" conception of representation differs in that mental representations are still posited. Aside from the fact that claims about a resemblance between a representation and properties of the world cannot be made due to inaccessibility to the external world (which I revisit in the following section), the conclusions that representations "match" the distal systems they enter into a relationship with is contested by some as unjustified. Pylyshyn (1981, 2007) points out that the results of the above experiments can be explained in terms of artefacts of the experimental design such as the prompt requiring participants to "scan" the mental image which by its definition entails a temporal dimension to allow a point-by-point investigation of a path. When similar experiments were duplicated with a different prompt bypassing the "scan" requirement, elapsed time no longer correlated with distance in the stimulus. In the new experiments, participants were asked to imagine a map with multiple locations, but instead of scanning from location A to location B, they were asked to imagine a light turning on at location B while the light at location A goes off and were asked to press a button when the "distance" was mentally traveled. This was repeated for different locations that varied in distance.

Results revealed that the distance between locations did not correlate with elapsed time, which remained constant, demonstrating that the representation of space does not in fact contain spatial properties, comparably to how a line can be defined as an equation and a point as a Cartesian pair on a Cartesian plane (Pylyshyn, 1981, 2007).

Pylyshyn's results support the claims that representations are discursive or discrete. Contrary to pictorial/analog representations which suppose that representations are continuous (Goodman, 1976), a discrete representation encodes specific properties by the presence or absence of meaningless symbols and their combinations. As Jackendoff puts it, "it is not the symbols themselves that are significant, but rather the range of distinctions possible in the system of symbols we adopt" (Jackendoff, 1991). Accordingly, a representation can be defined a symbolic expression comprising discrete symbols that encode particular phenomenal properties.

3.3 Representation as stand-in for Phenomenal States

The presence or absence of directly meaningful information about a particular property is what I will refer to as a representation's *decoupling* or *decoupleability*, [2] the separation of a representation from its stimulus. A symbolic representation of distance is thus completely decoupled from distance since it lacks spatial extension. This definition however raises a critical concern when one considers that space itself, as a Kantian approach implies, is constructed by the mind. Thus the same issues that riddle non-representational frameworks similarly plague either of the above characterizations of a representation, namely, the inaccessibility to the structural properties of external reality which precludes comparisons to be made with either phenomenal experience

---

[2] This term is borrowed from Chemero (2009), barring the causal links to external reality.

or the representations that the mind operates over. This raises the question of what "distance" refers to when describing the stimulus in the experiment above, or the nature of the "distance" that enters into conflict with the perception of "distance" in the illusions if "distance" in external reality is inaccessible.

3.3.1 Problem of the definition of reality

Those intentionally characterizing representations suppose that a representation is a stand-in for states of affairs in the world, which I have argued is an impossible claim to make, due to our inaccessibility to the latter. However, the experiments and illusions described above do seem to reference genuine mismatches between phenomenal experience and some objective[3] measure. This raises the question of what our phenomenal experience conflicts with if not external reality. Jackendoff (1991)'s arguments delineating the problems of reality thus need a re-classification which includes this observation. While under-representations do demonstrate the inaccessibility of some properties of external reality, illusions reference not the mismatch between external reality and phenomenal experience, but between two distinct phenomenal states: egocentric and allocentric. Egocentric phenomenal states concern the experiences in which the contents depend on the subject's specific point of view, while allocentric experiences are independent of one's perspective and thus can be said to be objective. The description of the Ames room illusion, for example, references the egocentric experience of distance as perceived from a specific angle as being different from a unit of distance as measured objectively by a standardized metric. Critically, the allocentric experience, while transcending a person's specific spatio-temporal circumstances, is still a phenomenal state and subject to the filters imposed by the

---

[3] While I employ the term "objective" due to its denotation, it is in fact an "inter-subjective consensus" amongst humans with similar mental rules as true objectivity would have to include the experiences of all species on earth.

mind.

3.3.2 Four definitions constituting reality

There are thus four distinct objects that enter into relationships in the construction of experience:

- Relational reality: (meta)physical properties of the external world (inaccessible and unknowable) that might trigger certain phenomenal properties through an opaque interaction. In the case of distance for example, external distance is unknowable and may or may not correspond to our experience of it. What is known scientifically as a ramification of Einstein's Special Relativity (the phenomenon of *length contraction* for example) is that "distance" is not absolute and varies depending on frames of reference.
- Phenomenal egocentric reality: the subjective and point-of-view dependent phenomenal experiences that the properties of relational reality engender (the meaningful contents of our experience). In the case of "distance", it is distance as experienced from one's specific perspective. This corresponds to experiencing a line in the Müller-Lyer illuson as shorter than another.
- Phenomenal allocentric reality: objective phenomenal knowledge, outside one's perspective but still mediated by the mind. This is a standardized measure achieved via humans' capacity for generating meta-mental states which abstract overarching properties from groups of phenomenal experiences. In the example of "distance", this corresponds to a mental construct which standardizes the experience of distance by achieving an inter-subjective consensus amongst humans on an absolute unit of distance. This "distance" is what relates the lines as

equal in Muller-Lyer illusion, in contrast to the "distance" of our phenomenal egocentric experience.

- Symbolic representations: mediating between the above to facilitate the construction of experience. In the case of distance, it is a data-structure comprising meaningless and spatially-empty symbols.

Note that while some aspects of relational reality such as the existence of infra-red or ultraviolet light are "known", these can never be experienced and are purely "known" due to our capacity for science. The distinction between what is known about relational reality on the one hand, and allocentric reality on the other, is that knowledge of relational reality can be said to be "the properties of reality that we know exist because of our capacity for science/reason" while allocentric reality is "the capacity for science/reason which permits us to reason and generalize about a reality we have no access to".

Critically, the representations can only be said to be "stand-ins" at most for our phenomenal experience, and not, as opposed to intentional characterizations of representations, to states of affairs in the world. For example, distance, in the above experiments, is a phenomenal property, either egocentric or allocentric. When two locations are said to be at a certain distance from each other, it is in fact the experience of distance that is referred to and not distance as a property in the relational world. A representation of distance is therefore not a stand-in for (meta)physical distance, which may or may not correspond to how the mind conceives distance, rather, it's a stand-in for our experience of distance as constructed by the mind. The conflict of illusions is between egocentric and allocentric experiences of distance, both phenomenal.

> Any mentions of decoupleability can therefore only be made in terms of how separate a representation is from its phenomenal properties, and not from a (meta)physical stimulus. This conflation of allocentric phenomenal experience with the external relational world is the cause of Jackendoff (1991)'s hesitation when using the label "representation" as these do not "represent" anything from the external world, rather, they are stand-ins for phenomenal experiences, both egocenric and allocenric.

3.4 Representation and Modularity

If one considers that relational reality is the totality of all possible properties, then phenomenal (allocentric and egocentric) *reality* are the subset of those properties specific to each species depending on their endowment; *representations* are thus the partitions of this subset (of phenomenal properties) which contain the properties only relevant for a single modality. A representation is "symbolic" in the sense that it encodes a particular aspect of phenomenal experience by inclusion or exclusion of a property, only meaningful by its presence or absence, and not its content. Accordingly, *decoupleability* refers to the encapsulation of a representation and its independence from other properties deemed irrelevant for a particular computation. This new definition of *decoupleabilit*y is intricately related to the concept of "substance freeness" from the previous chapter which describes phonological segments as devoid of "phonetic properties". I extend the definition of *substance* to include any property deemed irrelevant for a given computation. Thus the claim that "phonetic substance" is irrelevant for phonological computation is equivalent to asserting the decoupleability of phonological representations from phonetic substance.

Decoupleability and *substance* therefore relate to a concept essential for partitioning the study of specific cognitive processes: *modularity.*

Fodorian Modularity and Substance A module is a mental system responsible for the computations of specific cognitive tasks Fodor (1983). For example, *language* is processed by a specific module that is independent from other modalities like *audition*, or *vision.* In *modularity of the mind* (1983), Fodor lists the criteria that a module possesses, which include, amongst other criteria, encapsulation and domain specificity. These criteria require that representations be specified only for particular computations and be informationally encapsulated; representations should symbolically encode only the properties relevant for a particular computation and nothing else. Fodor, however, makes distinctions between "vertical processes" and "horizontal" ones, the former presented as the domain-specific computations particular to a module, for example, the computations that are specific to phonology like the mapping from underlying to surface form, and the latter explained as the domain-general processes that all modules draw from, like memory or attention, which are necessary resources required by all modules in their functioning.

How decoupled are Space and Time?

Two interconnected implications arise from modularity. The first is the identity condition necessary for the generalization that engenders a property to be "captured" by a representation. This refers to the shared properties which allows a module of the mind to consider all instantiations captured by a representation as "same" for the purpose of some mental computation, making these properties an 'equivalence class' (Pylyshin, 1984, Isac & Reiss, 2013). A representation, therefore, is a symbolic expression individuated in terms of the cognitive processes that operate over it. In other words, a representation is a data-structure that can act as input to a specific module. For the computation of distance, for example, representations are individuated in terms of the properties that are relevant for distance and nothing else.

The second implication which results from this characterization is that some information is considered irrelevant for determining membership in an equivalence class (color information, for example is discounted in the computation of distance). In Funes' example, a dog seen from the front is different from a dog seen from the side, and a dog seen in the morning is different from the same dog seen in the afternoon; however for most humans, these temporal and spatial dimensions are "forgotten differences" for the purpose of identifying a "dog" and are discounted from the representation which collapses all experiences into a single category. Thus, the spatial and temporal properties that are necessary for the construction of our experience of 'dog' do not seem relevant for the representation of 'dog'—in other words, are substance/ are decoupled from the representation of 'dog'; this raises the question of how intertwined "space" and "time" are with mental representations in general. Put simply, how modular/decoupled are mental representations from "space" and "time"? Are "space" and "time" domain-general horizontal processes or are they included in each representation?

In order to answer the above questions, an exploration of time, the different characteristics comprising our experience of time, as well as the associated computations follow in the next chapter.

Chapter 4

Time and Cognition

> If future and past events exist, I want to know where they are. If I have not the strength to discover the answer, at least I know that wherever they are, they are not there as future or past, but as present. For if there also they are future, they will not yet be there. If there also they are past, they are no longer there. Therefore, wherever they are, whatever they are, they do not exist except in the present
>
> ---
>
> *The Confessions*, Saint Augustine

Time has been a philosophical obsession since ancient philosophers began ruminating about the unchanging essence of things when change due to the time passage was inevitable. From Heraclitus' epistemological concern about knowledge when change was the only certainty, to Zeno's preoccupation with the apparent impossibility of motion, unease with respect to the problems engendered by time were evident even then.

More recent philosophical discussions about time have been generally described in terms of "reductionism with respect to time" or "substantivism with respect to time", the former ascribed to philosophers like Aristotle and Leibniz who argue that time is not independent from events and thus can be described in terms of relationship between events, and the latter adopted by Plato and Newton who subscribe to the view that time is independent of events and can be likened to "an empty container into which things and events may be placed" (Emery et al., 2020).

Modern understanding of time reveals more intricacy in light of the scientific advancements made with respect to its true nature. In fact, time perception is so multifaceted that most introductory perception and cognition textbooks completely omit chapters on the topic! In what follows, I will attempt to describe time from the three standpoints of reality outlined in the previous chapter (Relational, phenomenal allocentric and phenomenal egocentric) and provide evidence that certain mental representations and computations constituting time, in particular *order* is substantivist and thus modality independent.

4.1 Definitions

4.1.1 Allocentric Time: clocks

The most common-sensical and easiest to understand definition of time is the allocentric one, that is, time as a ticking clock, understood as the passage of absolute units of measurements such as seconds, minutes, hours, months, years etc. Those measurements are standardized and do not vary according to a subject's specific circumstances, rather, they designate measurements on which an inter-subjective consensus amongst the human species has been achieved and which can be measured by various instruments that are subject to Newtonian laws. Crucially, this conception of time is still phenomenal in that the concepts mediating its understanding are still mind-dependent. Experiments that measure reaction time, for example, refer to this definition of time.

4.1.2 Relational time: Eternalism and the block universe

The setting for many sci fi movies and novels, the understanding of time resulting from the ramifications of Einstein's special relativity and its reformulation as a fourth dimension to the three spatial dimensions marked a shift in attitude about the true nature of time. Einstein noticed that for special cases where two frames of reference are inertial (without acceleration), maintaining Maxwell's conclusions about a constant speed of light for both frames of reference while simultaneously preserving Newtonian laws of motion led to paradoxical conclusions that could only be resolved when positing that time and space are intricately linked and, in fact, form a single space-time. While counter-intuitive to say the least, some consequences of special relativity include time dilation and length contraction to accommodate a constant speed of light, as well as relative simultaneity whereby the order of events varies from observer to observer depending on the frame of reference.

Accordingly, space-time is conceived as an unchanging four-dimensional "block" universe where all events in time, past, present and future, exist at once. This "eternalist" view of time is what allegedly led Einstein to famously (and tactlessly) write that "The distinction between past, present, and future is only an illusion, however persistent" to reassure his friend upon the death of his sister that she is simultaneously alive somewhere in this "block" universe (Emery et al., 2020).

4.1.3 Egocentric phenomenal time: our object of study

Einstein's quote about time being a persistent illusion further expresses the mismatch between "real" reality and phenomenal reality outlined in the previous chapter. While the Tralfamagorian alien species from Kurt Vonnegut's *Slaughterhouse Five*, capable of seeing all time at once, are privy to the secrets of space-time, us humans on earth are sadly epistemically constrained by our mental constructs to be blind to these truths. Instead, what we experience is a 3-dimensional space modulated by the passage of time. The fabric of space-time described above is not apparent to humans who evolved under conditions which did not necessitate such an understanding for their survival, leading to an "illusory" or "distorted" experience of reality. Thus time cannot be perceived by humans who lack "time" receptors (Pöppel, 1997), only change is perceived, which we conflate with time. As French philosopher Jean-Marie Guyau writes in *la Genese de l'idee du temps*, "It is movement in space which creates time in human consciousness. Without movement there is no time" (Guyau cited from Wearden (2016)).

Jackendoff's argument applied to time Jackendoff's arguments from the previous chapter (under-representation, illusions and constructions) can similarly illustrate the constructed nature of temporal experience. For instance, contrasting what we know about the a-temporal and unchanging universe with our own tensed experiences exposes the partial/underrepresented interpretation of external reality that our existence engenders, hinting at the properties of reality which are missing from our experience, for example, future events.

Another set of evidence comes in the form of illusions. For instance, in the *Cutaneous Rabbit* illusion (Geldard and Sherrick, 1972), two spots on the wrist and elbow are tapped at a certain speed, resulting in an illusory perception of a sequence of touches on the forearm between these two spots, likened to the hops of a rabbit. Similar effects were replicated in the auditory (Shore et al., 1998) and visual (McFarland, 1970; Geldard, 1976) modalities where similar saltation effects were constructed from a stimulus lacking it, demonstrating that temporal allocentric stimuli can be manipulated akin to distance in the Ames Room or Müller-Lyer illusion to yield specific perceptions of order and succession, and thus confirming that perception of sequential order, much like perception of distance or space, is mind-dependent.

4.2 Taxonomy of Temporal Events

To better understand the mental processes underlying the construction of our experience of time, psychologist and neuroscientist Ernst Pöppel describes a psychological approach to the study of time which emphasizes our experience.

Mirroring Augustine's concern that "events are perceivable, but time is not", Pöppel (1988) substitutes the study of "time perception" by "event perception" and proceeds to characterize temporal events by establishing a taxonomy of elementary time experiences. These include (subjective) present, temporal continuity and duration, (non-)simultaneity, successiveness, temporal order, and change. The concern of this thesis, however, is on the representation of order and its relation to phenomena that appear to be comprised of successive changes over intervals such as the perception of motion, melodies, or words and I will thus focus on that particular elementary time experience.

4.2.1 The specious present, Order and Modularity

Consider the sequence of letters "blicket". More than likely, this sequence was experienced at once as a single event, the sub-components temporally integrated in a 'present' state in which the entity as a whole was assessed. The entity however comprises multiple sub-events (letters) and is characterized by the changes between each, so how can it be apprehended in a 'present' if the assessment of change requires contrast between consecutive events? To understand what order is, or rather, is not, the question of what the 'present' is necessary.

Guyau conceived of *the present* as an indivisible instant capturing the transition from past to future (Wearden, 2016). That characterization, as demonstrated above, leads to paradox as perception of change or motion happens over an interval, requiring *the present* to be a duration-less interval. This paradox of the nature of the present is best expressed by Augustin in the following passage:

> But to what period do we relate time when we measure it as it is passing? To the future, from which it comes? No: because we cannot measure what does not exist. To the present, through which it is passing? No: because we cannot measure what has no duration. To the past, then, towards which it is going? No again: because we cannot measure what no longer exists. (*Confessions*, Book XI, p. 27)

Those particular concerns about the nature of the present led to descriptions of "the specious present" by William James in his foundational volume *The Principles of Psychology* as "the short duration of which we are immediately and incessantly sensible" (James, 1890) , reformulated more recently as a "psychological moment" or "subjective present". Attempts to correlate this interval with allocentric quantitative measurements link it to attentional mechanisms and characterize it as the minimal time required for two events to be considered separate, with experiments showing that this interval of co-temporality occurred below100 ms (Brescher, 1932). More recently, the interval in which events were considered simultaneous and thus as belonging to the same event-structure[1] have been found to correlate with even smaller allocentric durations between 53-55 ms, coinciding with the neural mechanisms implementing the construction of precepts, namely, "the minimum number of oscillations required for two or more neurons to form an assembly that allows for the coding of perceptual structure" (Elliott and Giersch, 2016).

The contents of the specious present are characterized as simultaneous or "atemporal" (Elliott and Giersch, 2016), which means that they lack order. Order however is no less psychologically real, its perception necessary for the detection of patterns associated with particular sequences, indicating that the computation of order is separate from the contents of the present and leading to the conclusion that the events comprising the taxonomy above are independent computations.

4.2.2 Two sides of modularity

The absence of *order* in the perception of "present" raises questions about the independence of the computations underpinning these temporal events from each other on the one hand, and their independence from the events over which they operate on the other (i.e., the contents of experience over which temporal properties are imposed). In the following section, I will outline each facet of modularity and provides arguments supporting both instances.

---

[1] I employ this term in the psychological sense as designating the perceptual structure underpinning events and not in its mathematical and linguistic sense.

Modularity of elementary time experiences The modularity of the elementary time experiences outlined in Pöppel's taxonomy is empirically motivated. Pöppel (1997) argues that while these elementary temporal experiences are hierarchically connected, their neuro-cognitive implementations are independent. In other words, the sub-components of the mental system associated with time perception require distinct representations and operations. He lists patients with certain brain disorders such as Korsakoff syndrome where temporal ordinality—the perception of order—is affected, but the perception of successiveness is not, demonstrating that *order* and *succession* are independently computed (Pöppel, 1997). Additionally, experiments on temporal order judgement (TOJ) by Hirsh and Sherrick (1961) where participants were presented with two audio pulses of high and low pitch separated by varying temporal intervals show that whereas participants were capable of discriminating both sounds as distinct (i.e. non-simultaneous), they were unable to determine which sound, the high or low pitch, occurred first. This lead the researchers to conclude that non-simultaneity and order are serviced by separate mechanisms (Hirsh and Sherrick, 1961). These results are robustly supported in subsequent experiments by Basharat et al. (2018) who not only conclude that simultaneity judgements (SJ) and TOJs are served by distinct neural mechanisms, but that SJs are preserved with age while TOJs decay. And finally, evidence from double dissociations from patients with damage to specific brain areas— the gold standard in establishing that distinct neural mechanisms underpin different computations—has demonstrated that patients with damage to the left superior pre-frontal cortex suffer from “accelerated time phenomenon” where the duration of events seems compacted, but without additional temporal defects, confirming that the computation of duration can be impaired while other temporal experiences remain unaffected (Binkofski and Block, 1996). Thus *order* forms its own module consisting of specific representations and computations that underpin our constructed phenomenal experience.

Order versus Precedence Some experimental work on order link it to causality. For example, Bechlivanidis and Lagnado (2013, 2016) determined that when participants were presented with a simple three-item sequence of causally related visual stimuli A, B and C such that B causes C, they report perceiving events that are presented in a sequence ACB as ABC, conforming the order to causal patterns. For example, when shown a three object pseudo-collision where the third object starts moving before the second object moves and hits it (the effect happens before the cause), participants no less report that the second object started moving before the third. The researchers conclude that causality influences the perception of temporal order, and not the other way around as has been previously assumed.

However, causality has been generally attributed to precedence (Mellor, 1985). The perception of order thus needs to be distinguished from precedence between two events to unambiguously draw conclusions. Whereas temporal sequences determine order, no causal claim can be made between the events (Karimi, 2010). A consensus over a definition for causality is multi-faceted; however, it has generally been described by Hume by the following three intuitive conditions: A causes B if 1) A precedes B in time, 2) A and B are contiguous in space and time, and 3) A and B always co-occur or neither of them occurs (Karimi, 2010). The crucial takeaway is the observation that events for the purpose of precedence as it relates to causality have generally been described in terms of two events: *cause* and *effect*, necessitating the positing of no more than two events that enter this relation. Precedence, therefore, determines the relation between two events in an event-structure, while order determines the relation of one event relative to all other events comprising the same event-structure. This distinction between order and precedence adds a new element to the taxonomy.

Modularity and Decoupleability: how much *order* is in an event? A temporal event can thus be described in terms of more primitive components such as *duration, succession, order*, and *precedence* among others. Some events are grouped together in time as a single perceptual structure according to certain mechanisms which integrate distinct features into a coherent unit perceived simultaneously as an a-temporal entity. For example, *illusory conjunction* experiments in which two different shapes with different colors flashed on a screen led to the mismatched perception of an illusory combination of one shape with the other's color, supporting the *feature integration* hypothesis proposed in *binding theory* (Treisman and Gelade, 1980). Similar illusions have been elicited when the features were temporally spread out, no less resulting in a cohesive simultaneous perception of a single object. *Illusiory conjunctions* thus demonstrate how the brain assembles distinct features into a single structure. I will refer to the event-structures which are grouped together in time as "static representations". Other events are perceived temporally and are characterized by ordered changes over an interval. I will label these event-structures which comprise non-simultaneous events "dynamic representations". However as argued above, the property characteristic of dynamic representations, namely *order*, does not "exist" in external reality and is thus a product of the mind. In an event comprising change such as a melody, motion, or a word, the order in which the changes are perceived are therefore imposed by the mind which provides the filters according to which each sub-event is sorted. Ordering thus consists in an assessment of each sub-event event at a certain level of the cognitive architecture in which it is given ordinality. One representation utilizes time tags wherein an ordinal number defines the position of an event within a string of events (Pöppel, 1997).

Considering our approach to representations as abstract and symbolic, divorced from "meaning" and the contents of phenomenal perception, the next consideration concerns the encapsulation of this *order* representation and how interleaved it is with the representations of the events over which it imposes an ordering filter. Simply put, are the order indices included in dynamic representations (mirroring the reductionism with respect to time view) or is order independently represented to be incorporated at a different interfacing level in the cognitive architecture (similar to the substantivism with relation to time view)? For example, in the perception of a melody, is the order in which the sub-events (notes) arranged included in the representational input to the music module, or is there a "substance-free" representation of order which abstracts from the modality-specific tokens and operates over symbolic indices? In other words, how decoupled is the phenomenal perception of order in a specific modality from the representation of order?

4.3 Evidence for Substantivism

Rat timing experiments have shown that the computation of *duration* is modality independent. When presented with visual and auditory stimuli, rats were capable of estimating the full duration by adding the duration of each stimuli, showcasing that the representation of duration is modality-independent (Roberts, 1998). Similar conclusions about modality have been made by Pöppel (2004). In addition to the conclusion that neuro-cognitive implementations of order are independent from other temporal events, Pöppel believed that temporal mechanisms are "contentless" or "pre-semantic". That is, the computations implementing temporal experiences do not operate over representations that contain modality specific information, rather, they are substance-free, i.e. decoupled from the contents of a particular modality.

Discrepancies between egocentric order and allocentric order have been noted since the advent of experimental psychology in the 19th century after the astronomer Bessel noted that "objectively simultaneous stimuli fails to be subjectively simultaneous" (Sternberg and Knoll, 1973). Accordingly, a number of experiments on temporal order judgements (TOJ) have been undertaken as a means to understand the origins of this discrepancy. While some have attributed variability in order perception to what is known as "arrival latency" (i.e. the difference in time it takes a signal to arrive to brain centers that judge order), experiments have shown that TOJ is primarily controlled by a central mechanism independent of modality. For example, in a series of TOJ experiments, Hirsh and Sherrick (1961) presented participants with ordered stimuli from different modalities (vision, audition, touch) and found that although temporal resolution of succession—a prerequisite for order resolution—varied across modalities, judgement of temporal order did not, concluding that "whereas the time between successive stimuli that is necessary for the stimuli to be perceived as successive rather than simultaneous may depend upon the particular sense modality employed, the temporal separation that is required for the judgement of perceived temporal order is much longer and is independent of the sense modality employed" (1961). Additional experiments by Eijkman and Vendrik (1965) and Cheatham and White (1952) support the modality independence of temporal perceptions. The experimental results were taken as evidence for the existence of an "independent-channels" central timing mechanism that is responsible for order (Sternberg and Knoll, 1973).

In addition to these experiments, evidence from patients with brain damage in specific areas also seems to support the modularity of temporal order perception. For example, patients with unilateral brain damage in the left hemisphere showed impaired perception of temporal order in both auditory and visual stimuli (Carmon and Nachshon, 1971). This was corroborated by another set of experiments by Szelag et al. (1997) in which participants with Broca's aphasia were shown to report deficiencies in temporal sequential processing and employed different temporal integration strategies. In the tasks requiring sequences of tones to be integrated into order-sensitive higher order perceptual structures, aphasic patients with lesions in Broca's area (and not other areas such as Wernicke's) were found to rely on mental counting as opposed to automatic temporal integration in both language and non-language functions, hinting at the substance-free nature of ordered representations.

More recently, reasearch has linked the hippocampus and the entorhinal cortex with sequential activation of neuronal assemblies. In an experiment probing sequential recall, patients with lesions in the hippocampus were incapable of describing a sequence of events in the correct order, although the descriptions of the contents of each event did not vary from the control group (Dede et al., 2016). The hippocampus was described as a *sequence generator* in which sequences are stored seperately from content in the neocortex Buzsaki (2019), evoking the "what" and "where" streams which independently service distinct functions, one being content and the other sequence/order generation. Buzsaki (2019) gives the example of the generation of a sequence of words via appeal to both mechanisms, claiming that "Instead of storing every possible neuronal sentence, the cortex can separately store all the words, while the hippocampus concatenates their sequential order" (p.274).

Linear order is therefore arguably not part of the representational apparatus required by the distinct modalities that the mind utilizes in its functioning. In a science attempting to understand the mind and bridge the divide between the brain and the mind, a thorough and complete understanding of cognitive architecture requires a more fine-grained conception of the distinct cognitive representations and inputs to modules that takes order and other temporal features into account, reiterated from

Pöppel (2004):

> It has become clear, that perceptual or cognitive processes can only be understood if the dimension of time is taken more seriously. The reduction of complexity in neuronal systems is for instance, achieved by temporal integration mechanisms which are independent of the content of a percept or a cognitive act but are presemantical operations (p.295).

Now that we have established that order is decoupled from representations which seem to phenomenally contain it, we are in a position to start revising conceptions of representations which incorporated order. As this is a thesis on phonology, my concern is the input to the phonology module, a word.

Chapter 5

SEARCH Revisited: Non-linear Phonological Representations

As outlined in chapter 1, SEARCH & CHANGE is a formalism in which phonological dependencies between target and trigger are construed as unbounded. Locality under that account is epiphenomenal, with the preponderance of apparent local phenomena explained as resulting from non-linguistic pressures favoring the acquisition of certain patterns over others, what Hale (2003) calls the "diachronic filter"(p.363).

Crucially, the input to the phonological module, in this case the SEARCH, is a linear string such that each segment has its own index that locates it within the string. Implementing a SEARCH on a string can thus be done iteratively or recursively, as illustrated in chapter 2. In what follows, I will outline a simplified account of what a phonological formalism might look like if we consider the evidence outlined in the previous chapter that order can be decoupled from phonological representations.

Not only would such a reassessment offer more essential primitives to link phonological representations to potential neuro-physiological candidates, but it would also provide a stronger commitment to Minimalism wherein economy and simplicity are key, where representations that are not conceptually necessary would be avoided and symbols simplified (Lasnik, 2002), and where redundancy would be removed by importing certain overarching operations from other cognitive domains when possible (Hornstein, 2008), the residue of remaining computation and representational apparatus to be considered as purely linguistic, and thus clearly distinct from the content of other cognitive domains.

I will start by explaining why "order" is not necessary for phonology by proposing an architecture that places computation of temporal events outside the language module and distinguishes precedence from linear order. Next, I propose a simplified ordered pair representation similar to proposals made in Chandlee et al. (2019), Raimy (2000a) and Papillon (2020) and adapt the SEARCH formalism accordingly. I will follow that up with some simplified examples from real language (bypassing features for ease of exposition). And finally, I will explain how an ordered pair representation relates to the discussion in the literature on about the non-importance of the distinction between iterative versus simultaneous rule application given non-string representations.

5.1 Temporal Order is a *Substance*

Substance-free phonology, as seen in chapter 1, argues that the static phonological representations, the segments, are composed of symbolic features, and are thus *decoupled* from the corresponding physiological and acoustic mechanisms. But what about the dynamic phonological representation, the underlying representation that makes up the input to the phonological module? The conflation of temporal order in the representation with distance is rampant in phonology literature. For example, Nevins' 2010 *relativized minimality* [1] characterizes the relationship between target and trigger by analogies using distance on a map. The terminology employed in general, e.g. *locality*, *long distance*, all hint not only at the integration of "substance" into phonological representations but also at the false conceptual equivalence between temporal succession and distance in a representation.

Pushing the commitment to substance-freeness as well as the Search to its logical conclusion, we end up in a position in tandem with the conclusions from the previous chapter evidencing the existence of an independent sequential system, namely, that commitment to ordered segments/string representations as inputs to phonology is unwarranted.

Similar proposal of a separation of linear order from syntactic rules of combination were originally found in earlier versions of generative grammar (Chomsky, 1965) where the categorial component of the base specified the hierarchical structure as well as linearization. [2] While Kayne's (1994) *Linear Correspondence Axiom* had attempted to account for linear order representationally, the necessity of reconsidering the place of linear order in syntax has been reiterated in Berwick and Chomsky (2008) who assert that linearization is part of externalization, and in Idsardi and Raimy (2013) who claim that "The removal of linear order in narrow syntax provides the basis for a deeper investigation on how language is represented in the brain". As evidence from the previous chapter shows, similar claims can be made for the role of linear order in phonological representations.

---

[1] This term comes from syntactic literature, particularly Rizzi (1990).

[2] A separation of linearization from syntax could offer potential insights into non-compositional meaning for strings that are ungrammatical yet meaningful, see Lau (2023) on meaning and hippocampal memory

5.1.1 Order and Meaning

A formal language is the set of strings produced by a given formal grammar. A grammar consists of an alphabet and a set of production rules. In natural languages, the alphabet of morpho-syntax which forms words and sentences[3] is the lexicon, the "storage system" housing the set of morphemes of a given language. For phonology, the alphabet is the segment inventory of a given language.

A morpheme is generally described as the smallest unit of sound associated with a particular meaning. For example, the word *dogs* is made up of two morphemes, the root morpheme *dog*, and the suffix *-z* which encodes 'plural'.

Bale and Reiss (2018) describe a morpheme as a tripartite structure comprising a syntactic representation, a semantic/conceptual representation, and a phonological representation. Since *dog* and *god* are two distinct morphemes, the sound/meaning correspondence crucially necessitates that the order in which the segments occur be encoded in the morpheme to relay the proper interpretation.

5.1.2 Dual-system Representations

Generally, the order of segments is believed to be encoded in the phonological form. SPE, for example, assumes that morphemes are composed of ordered segments. OT describes correspondences and alignment constraints that refer to string position. A phonological representation of *dog* under those assumptions is a representation in which the segment d is given index 1, the segment o given index 2, and the segment g given index 3. Thus a phonological representation of *dog* could look like this $\{\langle d, 1\rangle, \langle \text{o}, 2\rangle, \langle g, 3\rangle\}$.

---

[3] for the sake of simplicity, I will follow Distributed Morphology and Nanosyntactic frameworks and assume that words and sentences are generated by the same module.

Chandlee et al. (2019) propose a more elaborate model-theoretic representation comprising an alphabet Σ (of segments), a domain of ordinal numbers and a relation R over the domain. They define a model signature as a *tuple* S such that:

$S = \langle D; R_1, R_2, ...R_m \rangle$ where the domain D is a finite set, and each R_i is a $n_i - ary$ relation over the domain (Chandlee et al., 2019).

A precedence model under that framework is defined as $\langle D; <, [R_\sigma]_{\sigma \in \Sigma} \rangle$ where $<$ is defined as $\{(i, j) \in D \times D \mid i < j\}$

Chandlee et al. (2019) provide the example of a string *abba*, given $\Sigma = \{a, b, c\}$, as the following:

$M^<(abba) = \langle D = \{1, 2, 3, 4\}; <= \{(1, 2), (1, 3), (1, 4), (2, 3), (2, 4), (3, 4)\}, R_{\mathrm{a}} = \{1, 4\}, R_{\mathrm{b}} = \{2, 3\}, R_{\mathrm{c}} = \emptyset \rangle$. In addition to the precedence, this specifies that a occurs in positions 1 and 4, b in positions 2 and 3, and c does not occur in this string.

Under such an account, a simplified representation of the string *dog* would be the model $M^<(dog) = \langle D = \{1, 2, 3\}; <= \{(1, 2), (1, 3), (2, 3)\}, R_{\mathrm{d}} = \{1\}, R_{\mathrm{o}} = \{2\}, R_{\mathrm{g}} = 3 \rangle$. [4]

Crucially, the precedence pairs produced by $<$ in this model are “substance-free” in that the specifications of their content is independently defined in R. The relevant information for meaning is the position of a segment in a string, thus this representation already offers two distinct components that can be broken down, one which links to “meaning”, and the other to specify the phonological environments.

As the above model shows, ordinality and precedence can be divorced. The indexed position of segments that links to meaning does not necessarily need to be included in

---

[4] Note that this is a simplified model that omits the position of the rest of the segment inventory Σ of English.

the input to the phonology module, which only requires information about precedence to determine the proper environments for rule application. Building upon *autosegmental phonology* (Goldsmith, 1979), similar proposals utilizing precedence-based morphophonological representations that capture a wide range of morphological and non-linear phonological phenomena have been made in Raimy (2000b), Idsardi and Raimy (2013) and Papillon (2020). Idsardi and Raimy (2013) describe a phonological representation as containing an X-tier (timing tier) with X slots from which precedence is derived, as well as the feature bundles (segments) that each X-slot contains.

Under such accounts, a representation of *dog*= $\langle P, A\rangle$ where P is the set of precedence statements and A is a set of autosegmental associations, which look similar to the model proposed in Chandlee et. al as the following:

- $\langle\{\langle \text{START}, d\rangle, \langle d, \text{o}\rangle, \langle \text{o}, g\rangle, \langle g, \text{END}\rangle, \langle \text{START}, \times_1\rangle, \langle\times_1, \times_2\rangle, \langle\times_2, \times_3\rangle, \langle\times_3, \text{END}\rangle\}, \{\langle d, \times_1\rangle, \langle \text{o}, \times_2\rangle, \langle d, \times_3\rangle\}\rangle$

This precedence model similarly provides a representation wherein the "timing" in which segments occur is divorced from the contents of the segments, in other words, supporting the substantivist views outlined in the previous chapter. The autosegmental graph representations however differ from the precedence models outlined in Chandlee et. al in that only immediate precedence is encoded in the representation.

The burden of generating the linearized output proposed in Idsardi and Raimy (2013) falls within the purview of the language module, divided into three distinct procedures detailing the transitions of the representations from one linguistic sub-module to the next to generate a linear phonological string at the end, copied in figure below :

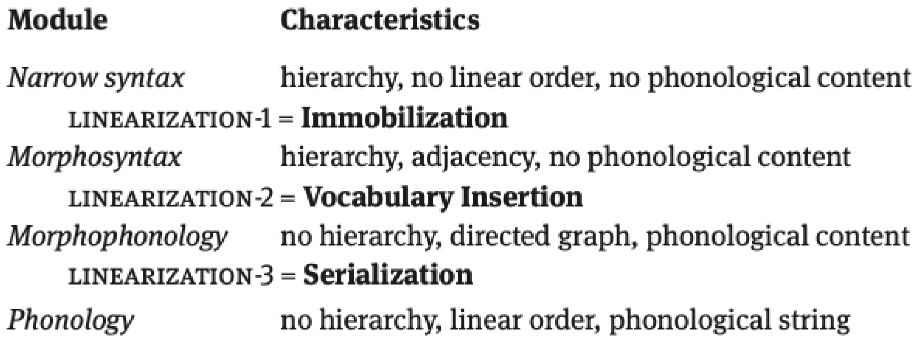

Module	**Characteristics**
Narrow syntax	hierarchy, no linear order, no phonological content
LINEARIZATION-1 = **Immobilization**	
Morphosyntax	hierarchy, adjacency, no phonological content
LINEARIZATION-2 = **Vocabulary Insertion**	
Morphophonology	no hierarchy, directed graph, phonological content
LINEARIZATION-3 = **Serialization**	
Phonology	no hierarchy, linear order, phonological string

Figure 5.1: Path from narrow syntax to the phonological representation (Idsardi & Raimy, 2013)

Under the above account, the final linearization step from a morpho-phonological immediate-precedence representation to a phonological linearly ordered representation is labeled “serialization”, producing a linear string of phonological segments. While Papillon (2020)’s model utilizes graph representations within the phonology to account for non-canonical phonological phenomena like tone and vowel harmony, Idsardi and Raimy (2013) maintain that strict linear order is present at the end of a derivation, specifying that by the end of phonology, a representation must be linear.

Whereas the model I present utilizes the distinction between ordinality (linking to meaning) and precedence (acting as input to phonology), my proposal diverges from both the above in two ways: 1- Both input and output to phonology are a full precedence representation mediated by the temporal center which transitivizes the immediate-precedence relation -2: The existence of a modality independent linearization/serialization mechanism that takes the ordered pair output of phonology and generates a linear string. In other words, it’s “precedence all the way down”, with the burden of serialization and transitivity generation falling outside the language module.

As such, the model I describe attempts to provide a higher-level description of linguistic objects and processes that incorporates the temporal taxonomy outlined in the previous chapter. Thus the input to phonology I describe via precedence differs in that the domain of the relation specifying the contents of each precedence token is a subset of the segment inventory Σ, namely, the set of segments that comprise the "word" being inputted to phonology which delineates it as a distinct psychological event-structure. The partitioning of the segment inventory for a particular phonological representation provides an account for the much-needed interaction between the language systems and the attentional/temporal mechanisms underpinning mental operations, which requires a more abstracted description of mental objects such as a *word* independent of their contents that takes into account the empirical evidence supporting the ubiquity of mental temporality.

5.2 What's in a *Word*?

The term "word" is often avoided in linguistic literature, replaced by more clearly defined and technically explicit terms like "morpheme" or "output of morphology". A thorough and explicit linguistic characterization of what constitutes a *word* remains vague. Clear and unambiguous definitions of *word* are even more complicated when polysynthetic languages are considered. In such languages, "words" are composed of many morphemes following combination rules isomorphic to sentence construction. Yet the distinction between "word" and "sentence" is no less valid for speakers making that distinction, signifying that a "word" is a psychologically real construct, its stand-alone experience suggesting boundaries that delineate it as a distinct perceptual object.

5.2.1 A word is more than just a linguistic object

A "word" for the purposes of phonology is circularly defined as the "minimal domain of phonological computation". This input is considered to be the output of *Spell-out* (Chomsky, 1998), the concatenation of multiple morphemes in (morpho)syntax which cyclically produces outputs that determine the domain of the phonological representation.

This characterization however includes multiple caveats and challenges necessitating the positing of extra machinery within phonology. While a morpheme is the smallest unit of meaning, a "word" is not necessarily a single morpheme, suggesting that the input to phonology is not the "phonological representation" associated with a morpheme, rather, it's the concatenation of multiple morphemes, which requires that either indexation be updated if committing to string representations or the precedence relations between all morphemes calculated to generate the necessary phonological environments. For example, in the formation of the word *dogs*, the phonological representations underpinning the morphemes "dog" and "-z" must be combined, and either the indexation of each segment updated (if we maintain that input to phonology is a string) or precedence relation updated to include the pairs specifying the relation of /-z/ with respect to the segments of *dog* (if we consider that only precedence matters) prior to being inputted into the phonology module.

Presuming that the new index allocation (for strings) or precedence relations (for precedence representations) be updated within the phonology conflicts with Minimalist assumptions and adds more machinery to phonology. Committing to a precedence model as outlined by Chandlee et al. (2019) for example necessitates a *transitivization function* to generate the new precedence pairs of the combined graphs. Under the architecture of Idsardi and Raimy (2013), the concatenation and update of the immediate-precedence graph falls within morpho-phonology (via anchor points). The precedence representation is then serialized to create a linear string within phonology. While both the above are possible, this nonetheless seems to render the phonology module more computationally complex than it needs to be. In light of empirical evidence hinting that the mind is likely supplied with a modality-independent temporal center endowed with the requisite functions, positing the existence of supplementary serialization or transitivization mechanisms within linguistic centers as well leads to undue redundancy.

Additionally, the definition of "word" as input to phonology from the output of Spell-out is one-directional and only descriptive in the process of production. When language perception is considered, a "word" is characterized as an event-structure constructed using temporal integration mechanisms. As the evidence from experiments described in the previous chapter demonstrates, the distinct properties activated in the brain during the perception of an event need to be sorted according to whether they belong to the static or dynamic representation. During word perception, the properties activated are the features[5] which will be grouped either simultaneously in a single bundle to form a distinct "segment" X, or will be deemed to be simultaneous with another group of features, thus forming another segment Y which, while distinct from X, no less belong to the same event (in working memory). While experimental evidence suggests that the perception of simultaneity is modality specific (which suggests that the binding of distinct features into a static segment is likely done within the language module), other temporal properties necessary for the construction of the dynamic event-structure (like order) are modality independent. A *word* therefore necessarily interfaces with temporal centers as well as linguistic centers. A *word* is thus a temporal object with linguistic content, a dynamic event-structure comprising

---

[5] This describes the perception at a more advanced level in which auditory signals have undergone some processing

segments as static sub-events. Mediation between language and time is necessarily present in the direction of perception, bolstering the claim that linguistic processes, whether in production or perception, are mediated by temporal mechanisms.

5.2.2 Temporal Interfaces

A *word* can thus be characterized as an event-structure that is psychologically experienced as a single entity. By the definitions of the previous chapter, a *word* can be construed as the contents of a "psychological moment/specious present", a system-state dictated by the oscillations that correlate with neural assembly, a product of temporal/event integration, and not solely a linguistic object.

As argued previously, the contents of a moment, order, and precedence are served by distinct neuro-cognitive systems, therefore an event-structure (in this case a word) comprises at least three distinct representations: the domain of the event-structure making up the components of the dynamical representation (a set E of sub-events);[6] an ordering relation on E which assigns indices to each sub-event; and a precedence relation on E which defines the ordered pairs that satisfy precedence. The crucial difference between linear order and precedence is that linear order determines the position of a sub-event (segment) relative to the entire set comprising the string while an immediate precedence relation determines the relation of a segment relative to only one other segment. For cognitive systems that are sensitive to causality, the linear representation is not relevant and only the precedence representation is. The association between a "word" and the domain of input to phonology (the set E) arises from extra-linguistic factors that restrict the maximal domain (Σ) via attentional mechanisms and memory resources required for the task of percept building (the

---

[6] In the case of phonology, the sub-events are segments.

specious present). This new dimension to the input of phonology could potentially provide insights into phonological processes which seem to cross the boundaries of what is described solely as the phased output of *Spell-out*.

The input to phonology is thus mediated by the temporal centers, the mechanisms of which defines the domain E. In the definitions of reality (relational, phenomenal and representational), one could say that relational reality is the set of all possible properties (the universal set U), phenomenal reality is the subset of these properties that humans are capable of apprehending (the set of alphabets Σ, primitives for a given modality), and the domain of a given representation a subset E of Σ defined by the restriction on Σ that the specious present provides. The specious present is thus a domain restrictor that limits the maximal domain Σ comprising phenomenal reality to a particular sub-set correlating with a slice of time, thus modulating its experience.

5.3 Precedence and Causality in SEARCH

Both of the outlined characterizations of SEARCH in chapter 1 (iterative vs recursive) are descriptive at an algorithmic level. Let's consider recursive SEARCH at a more implementational level. It describes a system wherein two segments are fetched from memory in a specified order – not necessarily equivalent to the perceived linear order– to be contrasted to the phonology's inventory of INR (initiator),TRM (terminator) pairs. Crucially, the "intervening material" is irrelevant for the purposes of this fetching mechanism much like *MERGE* disregards whatever material we consider to be intervening in the linear realization of a sentence.

5.3.1 Characterizing intervening material

One of the ways in which the Search differs from other formalisms dealing with long-distance dependencies is that instead of attempting to characterize the intervening material in order to determine what segments should be skipped, only the potential TRM is characterized, resulting is the desired outcome. This is made possible as a result of the inversion of the relation underpinning target and trigger such that the burden of setting-off the transformation falls on the target (the INR) as opposed to the trigger (the TRM), allowing each target to be mapped to at most one trigger, without the issue of other potential targets figuring in the intervening material when rules that include parenthesis star notation are used, a problem outlined in Jensen (1974). A Search can thus be said to encode causality between two events, with one causing another to change.

5.3.2 Search requires only two events

Consider an example from chapter 1 of a rule that changes a b to a c if it occurs (anywhere) before a c. Under a Search & Change account, this rule is described as the following:

(6) Rule 1

- Turn a b into a c if there is a c anywhere to its right.
- Rule Parameters
  - Search: INR: b, TRM: c, DIR: R
  - Change: INPUT: INR, OUTPUT: c

Iteratively, this rule states that when /badac/ is inputted, b will initiate a search rightwards, skip a, d and a because they do not comply with the specifications of a

TRM, and stop at c which becomes TRM. Finding a TRM satisfies the condition to move on to the CHANGE. In the CHANGE, the INR b becomes a c, and the outputted string is thus /cadac/. Alternatively the SEARCH can act as a procedure that hierarchically structures the string such that the b and c become sisters, with the CHANGE applying to the structurally adjacent INR TRM while a, d and a are omitted.

In either case, the relevant information for the rule in implementation is the ordered pair $\langle b, c\rangle$ which indicates that the INR b is followed by a TRM, c. This is the conclusion that arises when the assumption underpinning SEARCH & CHANGE which characterizes non-local phenomena as primary and adjacency as epi-phenomenal is pushed to its limit. Thus the only relevant segments for any SEARCH portion of a rule are the ⟨ INR, TRM ⟩ pairs specified by the language, namely the segment that changes (effect), and the environment required for the change (cause).[7]

In short then, a SEARCH only cares about precedence among two segments (the segment that changes and the environment) and not their ordinality within the word. To this end, the only parameters that a phonology model operating under SEARCH and CHANGE require is which segment precedes/follows which specific other. The precedence model of a word under that account contains the segment pairings to be compared against a language's inventory of $\langle INR, TRM\rangle$ pairs [8] corresponding with the parts of a rule that designate the target of change and the environment (the effect

---

[7] While the examples I have employed thus far are assimilatory where the cause and effect seem to be, on the surface, related by the same feature, the principle of *substance-freeness of structural change* mentioned in chapter 2 still applies; it is not the case that the contents of the cause (the trigger/where) are necessarily related to the changes to the effect (target/what). Cases where they are related (assimilation, for example a nasal trigger causing a target to became nasal) are a subset of all cases and are only over-represented in the inventory of rules due to the *diachronic filter* which favors the acquisition of certain patterns over others.

[8] Note that INR and TRM are not phonological categories, but descriptive tools for ease of exposition, correlating with the sub-event that changes and the cause of the change.

and the cause).

5.4 Architecture

An architecture that captures temporally experienced events is multifaceted and needs to account for the empirical observations outlined in the previous chapter by temporally sorting properties depending on modality.

Consider the universal set $U = \{a, b, c, d, e..., z\}$ of all properties. The genetic endowment of a species restricts the universal set U such that only a subset X_i of it can be apprehended. Thus $U = \bigcup\{X_i : i \in I\}$ where $I = \{1...n\}$ and X_i is the set of properties accessible to a particular species.

Suppose that $X_3 = \{a, b, c, x, y, z\}$— one of the subsets of U— is the set of phenomenal properties accessible to humans. The partitions of X_3 designate the domains (of properties) to a given modality. For example, one partition of X_3 is the set $\Sigma_1 = \{a, b, c\}$, the subset of X_3 containing phonological features. Another partition $\Sigma_2 = \{x, y, z\}$, is the subset of X_3 containing visual properties. Crucially, the Σ sets do not overlap since each contains the primitives of a given modality. Finally, the set E which designates the domain of a single event-structure restricts Σ_1, limiting access to a specified portion of all possible events, thus giving rise to illusions of temporality that result from the contrast that the absence of parts of Σ_1 engenders.

Thus $E \subset \Sigma_j \subset X_i \subset U$.

In summary then, the genetic endowment of a species restricts the universal set U to a subset X_i which is the set of properties that a given species is sensitive to; a mechanism that "sorts" properties depending on the relevant modality is available, one which partitions the set X_i of phenomenal properties into distinct sets Σ_j, each constituting the primitives for a given modality; and finally the content of a given event E experienced in the present results from a domain restriction to Σ_j. The mechanisms of the restriction to Σ_j go far beyond the scope of this thesis into the realm of quantum physics and the mechanisms that determine wave-function collapse. A brief overview of the distinct systems underpinning the interactions between time and phonology is outlined below.

5.4.1 The computational systems

A simplified architectural configuration mirroring the empirical evidence that supports the separation of temporal information from other representations is one in which the domain of the event-structure E is sent to a *precedence generator* in the temporal center to be transformed into a set of ordered pairs encoding precedence that can act as input to phonology and a *sequence generator* that determines the order of segments that links to meaning. The computational systems required to implement such a division of labor are the following:

- The ***present*** **generator**: A system that is equipped with the means to restrict the maximal domain of each modality which results in the subset E of a given Σ set. The upper bound on set cardinality (length) is contingent upon attentional resources. This is where distinct properties are "bound" together to form coherent perceptual objects, whether static or dynamic.[9] Note that this system is not an independent module, but a complex resource that relays with distinct modalities in the process of construction.

- **Sequence/order generator (SG)**: Allocates indices to sub-events in E to

---

[9] For phonology, the components of static representations are features, and the components of dynamic representations are segments.

individuate them relative to all the other sub-events in E. The output of this system correlates with our "experience" of linear order. Neuropsychological evidence tasks the hippocampus with these computations (see previous chapter). This system interfaces with the conceptual system to determine meaning,[10] and the sensorimotor system to organize gestural sequences. In the case of words, this corresponds to sequences of articulatory gestures. For example, if the domain of an event-structure is the set $E = \{a, b, c, d\}$, the SG assigns each element of that set an index that defines its location relative to the entire event-structure. One possible output is the ordered set $\langle b_1, c_2, a_3, b_4 \rangle$ which assigns b indices 1 and 4, c index 2, and a index 3.

- **Precedence generator (PG)**: Sorts sub-events in E into a binary precedence relation, and contains a transitivization mechanism which generates precedence relations among concatenated events. The output of this system is a set of binary ordered pairs defined over empty "slots" that link to the relevant static sub-events depending on modality. Sub-events are thus not locally copied, but referred to globally, i.e. via their address in memory. Since the relation is defined over substance-free "empty slots" and not over the sub-events in E themselves, sub-events which seemingly occur more than once in a precedence relation are no less "distinct" for precedence (since they are not defined solely by their content, but by also what precedence slot they occupy). For ease and clarity of exposition, I will use subscripts to describe equivalent segments which are linked to distinct precedence slots.

This system relays with modules that are sensitive to order between two events

---

[10] complex meaning is computed compositionally from the syntactic representation, however, some meaning can be generated from linear order as well as evidenced by the existence of meaningful ungrammatical sentences.

(e.g. phonology). For example, if $E = \{a, b, c, d\}$, the string b_1cab_2 is stored as the immediate precedence pairs $\{\langle b_1, c\rangle, \langle c, a\rangle, \langle a, b_2\rangle\}$, with the rest of the precedence pairs generated by transitivization mechanism as the following:

$<= \{\langle b_1, c\rangle, \langle b_1, a\rangle, \langle b_1, b_2\rangle, \langle c, a\rangle, \langle c, b_2\rangle, \langle a, b_2\rangle\}$[11]

- **The phonology module**: Takes as input the phased output of PG (coinciding with the output at each transitivization step) and applies set-theoretic operations to determine whether conditions on transformation are met (more later).
- **The Sensorimotor system**: takes as input the output of SG and generates sensori-motor commands.
- **Conceptual system**: relays with SG to determines the proper conceptual address.
- **Memory**: Stores phonological representations as immediate precedence pairs.

5.4.2 The transitivization mechanism of PG

A precedence relation is a binary relation on a set, generally defined as $R = \langle D, G\rangle$ where D is the domain (in our case, the domain of an event-structure E) and G is the graph of the binary relation, the set of ordered pairs related to each other by R. A (full) precedence relation $<$ is transitive. This means that if $\langle x, y\rangle \in G$ and $\langle y, z\rangle \in G$, then $\langle x, z\rangle \in G$.

As mentioned above, the representation of an event-structure consists of a domain, a linear order, and a set of precedence pairs. The precedence relation however need

---

[11] Note that while I use the segments directly in the ordered pairs, this precedence relation is not defined over the segments themselves, rather over the precedence slots and would look something more similar to the description provided in Chandlee et al. (2019). I use the segments directly however for ease of exposition.

not be stored fully in memory for a given sub-event, which would require unnecessary storage. The potential existence of a mechanism that transitivises a relation by generating the missing pairs from a representation that lists immediate precedence pairs lessens the burden on memory while preserving the relevant environments between events. For example, while a full precedence representation of *dog* would include the following pairs $\{\langle \mathrm{d,\ o}\rangle, \langle \mathrm{o,\ g}\ \rangle, \langle\ \mathrm{d,g}\ \rangle\}$, a stored representation need not include the pair $\langle$ d,g $\rangle$ which can be computed.

Idsardi and Raimy (2013) describe an architecture where the graphs of distinct morphemes combine at the level of morpho-phonology to form a new connected graph using anchor points. The combined graph representation then gets serialized for phonology to be shipped to the motor system for production. Considering that the precedence relation within their system is not defined over segments but over X tiers, placing the burden of generating the connected graph of two or more morphemes on a domain-specific mechanism conflicts with Minimalist assumptions, particularly when one considers that the operation is defined over the "temporal X slots", and not over the phonological information itself. Under their account, word boundaries # and % are precedence slots that contain no phonological information (an empty set), and morpheme types are defined by those word boundaries: root moprhemes are connected between # and %, but affixes are not; they only contain one boundary that acts as an "anchor point" which determines where it connects to a root. The implication arising from such a system for representing morphological information is that while "anchor points" are available in the morpho-phonological representations to inform word formation, the (combined) connected graphs are not necessarily generated within the language module, but possibly within a general-purpose mechanism which uses the anchoring information to determine where the graphs connect.

This general-purpose mechanism, in addition to producing the connected graph, transitivizes any input requiring decisions about causality between multiple events. To achieve this, the PG needs to be equipped with the tools to generate a combined graph (the mechanisms of which I omit in this thesis for the purpose of focusing on the phonological aspects), as well as a mechanism to compute all required pairs to determine the proper environments among events.

Note that for ease of exposition, I omit word boundaries and assume that anchor point information (in the form of empty content) is used to determine how to connect the graphs. Moreover, I employ the precedence relation directly over the domain of the event structure purely for exposition; in actuality, the relation and ensuing representation are more complex as the precedence is defined over a presemantic structure that links to the phonological content via pointers. Thus when I represent a string *abc* as the precedence model $\{\langle a, b\rangle, \langle b, c\rangle\}$, the actual model is more complex, akin to the models outlined in Chandlee et al. (2019) and Idsardi and Raimy (2013) where precedence and content are independently defined.

Consider the strings *abc* and *de* forming *abcde*, defined as the concatenated relation $R_{abcde} = \langle D, G\rangle$, where $D = \{a, b, c, d, e\}$, and $G = \{\langle a, b\rangle, \langle b, c\rangle, \langle c, d\rangle, \langle d, e\rangle\}$. [12]

- **Transitivization of G:**

  The transitivization of a graph of a binary relation is defined as:

  $Trn(R) = \langle D, \bigcup G_n\rangle$ where $G_1 = G$ and $G_{n+1} = G_n \cup G_nG_n$ (GG is the composition of G with itself).

---

[12] I am omitting here the anchor-point mechanisms which connect the graphs of the relations R_{abc} and R_{de} associated with the strings *abc* and *de* to directly tackle the transitivization mechanism that generates the relevent phonological environments for a given rule.

This transitivization proceeds in multiple passes, [13] generating the composition of G with itself incrementally until the relation is fully transitive. Each binary relation on a domain with $n+2$ members requires n iterations of the procedure to render it transitive. Each iteration of the procedure is shipped off to the relevant module, thus creating a tiered domain of precedence upon which distinct classes of rules can apply. The final application of the procedure leads to G_n, the set containing all precedence pairs. For the relation R, the passes proceed as follows:

1. $G_1 = G = \{\langle a,b\rangle, \langle b,c\rangle, \langle c,d\rangle, \langle d,e\rangle\}$

   and $GG = \{\langle a,c\rangle, \langle b,d\rangle, \langle c,e\rangle\}$

2. $G_2 = G \cup GG = \{\langle a,b\rangle, \langle a,c\rangle, \langle b,c\rangle, \langle b,d\rangle, \langle c,d\rangle, \langle c,e\rangle, \langle d,e\rangle\}$

   and $G_2G_2 = \{\langle a,d\rangle, \langle a,e\rangle, \langle b,e\rangle\}$

3. $G_3 = G_2 \cup G_2G_2 = \{\langle a,b\rangle, \langle a,c\rangle, \langle a,d\rangle, \langle a,e\rangle, \langle b,c\rangle, \langle b,d\rangle, \langle b,e\rangle, \langle c,d\rangle, \langle c,e\rangle, \langle d,e\rangle\}$

   The procedure ends here for this relation since G_3 is fully transitivized. For relations with larger domains, the transitivization requires more steps and thus more phases to output. What is crucial in the above procedure is that each G_i that results from the union of the graph with the composition of the previously generated graph with itself gets passed on to the respective module as a domain for the relevant rules to apply. As we will see later, this separation provides the proper environments for classes of rule without the need to posit interactions from other cognitive modules such as morphological boundaries or syllable/foot structure as rule environments.

The phased outputs from step 3 are the representations that will interface with

---

[13] While transitive closure can apply as a single operation, transitivization can also be computed incrementally.

relevant modules that require sensitivity between specific of events. In the case of phonology, the output of each procedure will be parcelled to the phonology individually corresponding with distinct domains for rules within a word.

5.5 SEARCH formalism

As a ramification of the non-locality assumption underpinning SEARCH and CHANGE, and in compliance with evidence of a domain-general *sequence generator* responsible for order, input to phonology, as argued above, need not be a string. A phonological representation under those conditions is a set of ordered pairs determined by the precedence relation. For example, what we perceive as the string $< b, a, d, a, c >$ would be generated by PG as the family of sets $\{G_1...G_n\}$ such that $G_1 = \{\langle b, a\rangle, \langle a, d\rangle, \langle d, a\rangle, \langle a, c\rangle\}$, and $\bigcup_1^n G = \{\langle b, a\rangle, \langle b, d\rangle, \langle b, c\rangle, \langle a, d\rangle, \langle a, a\rangle, \langle a, c\rangle, \langle d, a\rangle, \langle d, c\rangle\}$.

A language's *phonology* P is thus defined as $P = \langle \Sigma, \zeta, \textsc{Change}\rangle$ where Σ is the segment inventory, the set $\zeta = \{\zeta_1...\zeta_n\}$ where each set ζ_i lists the $\langle$ INR,TRM $\rangle$ pairs[14] for a specific tier G_i of the output of PG that correlate with the "target" and "trigger" of traditional rules,[15] a SEARCH defined as the intersection of a "word" with the relevant ζ tier: SEARCH= $\zeta_i \cap G_i$, and a CHANGE function which takes as input the $\langle$ INR,TRM $\rangle$ pairs and outputs another pair in which the INR is rewritten. Since the precedence relation is defined over X-slots which link to phonological content rather than the content itself, the changes to the contents of an X-slot result in the automatic update of the contents associated with each occurrence of that slot within the graph. Thus when a b associated with slot X_3 gets updated to a c, every occurrence of b associated with X_3 in other pairs within the graph similarly gets updated to a c.

---

[14] Note that INR TRM categories are not phonological categories, they are purely descriptive tools to facilitate exposition.

[15] For the purpose of introducing the model, I assume very simple rule types.

Note that the direction of the Search (left versus right) is easily encoded in the ordered pair: Rightward searches correspond with the INR being the first member of the pair, while leftward searches correspond with the INR occupying the position of second member of the pair.

5.5.1 Search as set intersection

Now that the terminology and parameters for this new system have been outlined, we are in a position to focus on the Search equivalent for ordered-pair representations: set-intersection on transitivization passes. Suppose a language has an alphabet $\Sigma = \{a, b, c, d, e\}$. The rules in this language are:

- c becomes d directly before a d
- b becomes d directly before an d
- a becomes e before an e

This language's phonology is defined as $P = \langle \Sigma, \zeta, \textsc{Change} \rangle$ where

$\zeta = \{\zeta_1, \zeta_2, \zeta_3\}$ such that $\zeta_1 = \{\langle b, d\rangle, \langle c, d\rangle\}$, $\zeta_2 = \emptyset$ and $\zeta_3 = \langle a, e\rangle$.

and Change is the following mappings:

$$\begin{aligned} \textsc{Change:}\quad & \langle c, d\rangle \mapsto \langle d, d\rangle \\ & \langle b, d\rangle \mapsto \langle d, d\rangle \\ & \langle a, e\rangle \mapsto \langle e, e\rangle \end{aligned}$$

This means that any members $\langle b, d\rangle$ and $\langle c, d\rangle$ of G_1 or member $\langle a, e\rangle$ of G_3 will be outputted by the Change as $\langle d, d\rangle$, $\langle d, d\rangle$ and $\langle e, e\rangle$ respectively.

The Search consists solely of the set-theoretic operation of *intersection* which finds all the $\langle INR, TRM\rangle$ pairs that a phonological representation possesses in a given tier in bulk.

The string *abcde* above will be generated by the PG as the sets:

- $G_1 = \{\langle a, b\rangle, \langle b, c\rangle, \langle c, d\rangle, \langle d, e\rangle\}$
- $G_2 = \{\langle a, b\rangle, \langle a, c\rangle, \langle b, c\rangle, \langle b, d\rangle, \langle c, d\rangle, \langle c, e\rangle, \langle d, e\rangle\}$
- $G_3 = \{\langle a, b\rangle, \langle a, c\rangle, \langle a, d\rangle, \langle a, e\rangle, \langle b, c\rangle, \langle b, d\rangle, \langle b, e\rangle, \langle c, d\rangle, \langle c, e\rangle, \langle d, e\rangle\}$

The Search applied to *abcde* will be:

1. $\text{Search}_1 = \zeta_1 \cap G_1 = \{\langle b, d\rangle, \langle c, d\rangle\} \cap \{\langle a, b\rangle, \langle b, c\rangle, \langle c, d\rangle, \langle d, e\rangle\} = \{\langle c, d\rangle\}$

   Search_1 will be the input to the Change which will change $\langle c, d\rangle$ to $\langle d, d\rangle$. The updated G_1 will be outputted to the linealization module to reconstruct[16] the string using this new ordered pair into *abdde*.

2. $\text{Search}_2 = \zeta_2 \cap G_2 = \emptyset \cap \{\langle a, b\rangle, \langle a, c\rangle, \langle b, c\rangle, \langle b, d\rangle, \langle c, d\rangle, \langle c, e\rangle, \langle d, e\rangle\} = \emptyset$

   No input to Change and G_2 is outputted as is.

3. $\text{Search}_3 = \zeta_3 \cap G_3 = \{\langle a, e\rangle\} \cap \{\langle a, b\rangle, \langle a, c\rangle, \langle a, d\rangle, \langle a, e\rangle, \langle b, c\rangle, \langle b, d\rangle, \langle b, e\rangle, \langle c, d\rangle, \langle c, e\rangle, \langle d, e\rangle\}$ $= \{\langle a, e\rangle\}$

   The output of Search_3 is inputted to the Change which rewrites it as $\langle e, e\rangle$. The changed G_3 is shipped to the linearlization module which uses the new output to update the string, which comes out as *ebdde*.

---

[16] Note that the reconstruction of the string from a full graph will require that every pair that contains an INR that changes to similarly be updated, which can, on the surface, render the linearlization process complex. However as the Change operation merely reroutes the pointer of an X-slot to another content address, the contents of that same X-slot within other pairs of the graph will be automatically updated. For example when $\langle c, d\rangle$ is rewitten as $\langle d, d\rangle$, the contents of the X-slot associated with the *c*— call it X_3— will change to a *d*, thus all ordered pairs in G which contain X_3 will have their contents automatically updated/rerouted as a result.

Dialectical Variation Suppose that P_2 is a dialect of P where ζ_1 and ζ_2 are flipped such that $\zeta_1 = \emptyset$ and $\zeta_2 = \{\langle b, d\rangle, \langle c, d\rangle\}$.

The Search portion would vary in that Search_1 would result in no change to G_1 and Search_2 would result in the set containing both $\langle b, d\rangle$ and $\langle c, d\rangle$. The Change would thus rewrite both into $\langle d, d\rangle$ in G_2, leading the string to be reconstructed as *addde* after the second phase and as *eddde* after the third phase.

Under that account, a transformation is neither a segment to segment nor a string to string mapping, but a cyclical ordered pair to ordered pair mapping.

5.5.2 Toy Examples

Consider a language $L = \langle \Sigma, \zeta, \text{Change}\rangle$ with the following rules:

Rule 1:

- Turn a □ to a ○ only if the shape directly to its right is a ▲
- Rule 1 Parameters

  Search: INR: □, TRM: SHAPE, DIR: R

  Change: INPUT: INR, OUTPUT: ○, CONDITION: TRM: ▲

and rule 2:

- Turn a ■ into a ○ if there is a ▲ anywhere to its right.
- Rule 2 Parameters
  - Search: INR: ■, TRM: ▲, DIR: R
  - Change: INPUT: INR, OUTPUT: ○

This translates to the following parameters of the phonology:

- $\Sigma = \{ \bigcirc, \square, \triangle, \bullet, \blacksquare, \blacktriangle \}$

- $\zeta_1 = \{\langle \square, \blacktriangle \rangle\}$
- $\zeta_2 = \emptyset$
- $\zeta_3 = \{\langle \blacksquare, \blacktriangle \rangle\}$

The transitivization tier allocation corresponds to the scope of rule application when Search is described over a string. Rule 2 specifies a TRM without any conditions on change, which results in an unbounded Search that can skip any intervening segment and seemingly operates over long-distance. This rule applies on the final tier ζ_3 which encodes precedence among all static-events comprising the representation. Rule 1 however seemingly targets adjacent segments, which coincides with tier 1 that encodes immediate precedence.

- Change:

(Rule 1) $\langle \square, \blacktriangle \rangle \mapsto \langle \bigcirc, \blacktriangle \rangle$

(Rule 2) $\langle \blacksquare, \blacktriangle \rangle \mapsto \langle \bigcirc, \blacktriangle \rangle$

Consider the following string:

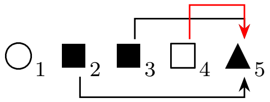

Under the string-based Search & Change, applying rules 1 (red arrows) and 2 (black arrows) results in the string $\bigcirc_1 \bigcirc_2 \bigcirc_3 \bigcirc_4 \blacktriangle_5$

The black squares in positions 2 and 3 launch searches (rule 2) and find a black triangle in position 5, changing into white circles. The white square in position 4 launches a search (rule 1 in red) for any shape and terminates on the first shape in position 5. Since TRM is a black triangle, the condition on change stipulating that the shape must be a black triangle is met, and square 4 turns into a white circle. Equivalent outcomes are achieved via the transitivization tiers in the precedence relation to result in the same output.

Under a phased precedence model, the above string is represented as [17]:

$G_1 = \{\langle \bigcirc, \blacksquare_1\rangle, \langle \blacksquare_1, \blacksquare_2\rangle, \langle \blacksquare_2, \square\rangle, \langle \square, \blacktriangle\rangle\}$

$G_2 = \{\langle \bigcirc, \blacksquare_1\rangle, \langle \bigcirc, \blacksquare_2\rangle, \langle \blacksquare_1, \blacksquare_2\rangle, \langle \blacksquare_1, \square\rangle, \langle \blacksquare_2, \square\rangle, \langle \blacksquare_2, \blacktriangle\rangle, \langle \square, \blacktriangle\rangle\}$

$G_3 = \{\langle \bigcirc, \blacksquare_1\rangle, \langle \bigcirc, \blacksquare_2\rangle, \langle \bigcirc, \square\rangle, \langle \bigcirc, \blacktriangle_1\rangle, \langle \blacksquare_1, \blacksquare_2\rangle, \langle \blacksquare_1, \square\rangle, \langle \blacksquare_1, \blacktriangle_2\rangle, \langle \blacksquare_2, \square\rangle, \langle \blacksquare_2, \blacktriangle\rangle, \langle \square, \blacktriangle\rangle\}$

The search component consists of the intersection of the set of ordered pairs designated by the language ζ with the set of ordered pairs constituting the the word. Thus for rule 1, the following intersection takes place:

- $\textsc{Search}_1 = \zeta_1 \cap G_1$

$= \{\langle \square, \blacktriangle\rangle\} \cap \{\langle \bigcirc, \blacksquare\rangle, \langle \blacksquare, \blacksquare\rangle, \langle \blacksquare, \square\rangle, \langle \square, \blacktriangle\rangle, \langle \blacktriangle, \blacksquare\rangle\}$

$= \{\langle \square, \blacktriangle\rangle\}.$

The output of this Search is fed to the Change, which effects the following change:

- Change: $\langle \square, \blacktriangle\rangle \mapsto \langle \bigcirc, \blacktriangle\rangle$

---

[17] Note that while I label black squares to demonstrate their distinction, this is purely for descriptive purposes. They are considered distinct on account of the distinctness of the X-slots into which the contents are sorted.

The second rules consists of intersection of ζ_3 with G_3: • $\textsc{Search}_2 = \zeta_3 \cap G_3$

$= \{\langle \blacksquare, \blacktriangle\rangle\} \cap \{\langle \bigcirc, \blacksquare_1\rangle, \langle \bigcirc, \blacksquare_2\rangle, \langle \bigcirc, \square\rangle, \langle \bigcirc, \blacktriangle_1\rangle, \langle \blacksquare_1, \blacksquare_2\rangle, \langle \blacksquare_1, \square\rangle, \langle \blacksquare_1, \blacktriangle_2\rangle, \langle \blacksquare_2, \square\rangle, \langle \blacksquare_2, \blacktriangle\rangle, \langle \square, \blacktriangle\rangle\}$

$= \{\langle \blacksquare_1, \blacktriangle\rangle, \langle \blacksquare_2, \blacktriangle\rangle\}.$

- Change: $\langle \blacksquare_{1,2}, \blacktriangle\rangle \mapsto \langle \bigcirc, \blacktriangle\rangle$

The new graph with the the ordered pairs replaced is then fed to the SG which reconstructs a linear order to be send to the sensori-motor system.

Each single operation generates a new ordered pair that replaces the old one. One distinction between the linear Search and this approach is that the output of the change is not a single segment (the INR), but rather the set of ordered pairs consisting of both INR and TRM. This allows the Change to use information from the TRM when extra conditions for change are stipulated to determine whether the change takes place. The Change can remain encapsulated, requiring no additional information aside from the information inside the bundle of ⟨INR,TRM ⟩ pairs. [18] This output in turn serves in the reconstruction of the string at the interface (SG) that lineralizes the representation to be transduced into sensory-motor commands.

---

[18] A reminder that these examples are simplified with precedence defined directly over the segments; In actuality precedence is defined over X-slots as opposed to the phonological content itself, thus when an INR associated with a given X-slot changes, the content is automatically updated in all the pairs containing that same X-slot.

5.5.3 Real Language Examples

Crimean Vowel Harmony

Consider the following data from the Southern and Central dialects of Tatar from McCollum and Kavitskaya (2018).

3SG.POSS

	UR	SR C	SR S	gloss
a.	tuz-lIɣ-I	tuz-luɣ-ɯ	tuz-luɣ-u	‘salt’
b.	kyz-lIg-I	kyz-lyg-i	kyz-lyg-y	‘autumn’
c.	toz-lIɣ-I	toz-luɣ-ɯ	toz-luɣ-u	‘dust’
d.	køz-lIg-I	køz-lyg-i	køz-lyg-y	’eye’
e.	baʃ-I	baʃ-ɯ	baʃ-ɯ	‘color’
f.	tiʃ-I	tiʃ-i	tiʃ-i	‘tooth’

The Southern (S) and Central (C) dialects of Crimean Tatar both show so-called ‘iterative’ Back harmony. In contrast, Round harmony is ‘iterative’ in S but not C. In this framework iterativity is merely apparent—rules apply simultaneously to all parts of an input string.

Assume that the vowel specifications are the same in both dialects, as the following:

- 3 underspecified vowels:

  I = {+Hi }

  Y= { +Hi, +Rnd }

  U= { +Hi, -Rnd }

Using the Search and Unify approach, this microvariation in Round harmony is captured by the following rules:

- Central Crimean Tatar:

  INR $\supseteq$ { +Syl }

  DIR: LEFT

  TRM $\supseteq$ { +Syl }

  INR $\sqcup$ { +Rnd } if TRM $\supseteq$ { +Rnd}

  Default rule: INR $\sqcup$ { -Rnd }

- Southern Crimean Tatar:

  INR $\supseteq$ { +Syl }

  DIR: LEFT

  TRM $\supseteq$ { +Syl, αRnd }

  INR $\sqcup$ { αRnd }

As the mappings in [tuz-luɣ-u] and [tuz-luɣ-ɯ] show, in the central dialect, a suffix will agree with respect to round only if it immediately follows the root vowel. The second suffix vowel does not appear to copy the +Rnd value of the root vowel, instead surfacing as the -Rnd [ɯ]. INRs launch a leftwards search, terminating on any vowel (+ Syl) in the central dialect, but on a vowel specified for a Rnd value in the southern dialect. The non-iterativity in C is accounted for by stipulating a condition which allows unification with +Rnd only if the segment that terminates the search is also +Rnd, and will otherwise resort to a default feature filling rule that unifies it with -Rnd. As such, the under-specified first-suffix vowel will terminate the search launched by the second-suffix vowel, but will prevent a +Rnd vowel from surfacing due to the condition on +Rnd not being met. In S, in contrast, no such stipulations to unification exist. The iterativity is explained by a condition on the segment that can terminate the search, namely, that it must be specified for a Rnd value, which unifies with the initiating segment.

Under a transitivization approach Under a transitivization approach, this variation is captured by the tier upon which the search applies. The parameters of both are specified as the same with the exception of ζ[19] which varies.

---

[19] the direction of the search is leftward, coinciding with a $\langle$ TRM, INR$\rangle$ format for the pairs

Thus $P = \{\Sigma, \zeta_{S,C}, \textsc{Change}\}$ such that $\zeta_C = \{\zeta_{C1} ... \zeta_{C4}\}$ and $\zeta_S = \{\zeta_{S1} ... \zeta_{S4}\}$

Where the Central dialect has the followings ζ_C specifications :

$\zeta_{C3} = \{\langle [\alpha\textsc{Rnd}] \text{ , I } \rangle\}$

$\zeta_{C4} = \{\langle [\alpha\textsc{Back}] \text{ , I } \rangle\}$

And the Southern dialect has the following ζ_S specifications:

$\zeta_{S4} = \{\langle [\alpha\textsc{Back}] \text{ , I } \rangle, \langle [\alpha\textsc{Rnd}] \text{ , I } \rangle\}$.

Where $[\alpha\textsc{Back}]$ and $[\alpha\textsc{Rnd}]$ designate any segment that is specified for a Back and Rnd value respectively.

The Change is specified as the same for both dialects as the following:

- Change:

(Back Harmony) $\langle [\alpha\textsc{Back}] \text{ , I } \rangle \mapsto \langle [\alpha\textsc{Back}] \text{ , I} \sqcup \alpha\textsc{Back} \rangle$

(Rnd Harmony) $\langle [\alpha\textsc{Rnd}] \text{ , I } \rangle \mapsto \langle [\alpha\textsc{Rnd}] \text{ , I} \sqcup \alpha\textsc{Rnd} \rangle$

Consider the underlying form specified in (b) /kyz-lIg-I/ surfacing as *kyz-lyg-i* in the Central dialect *kyz-lyg-y* in the southern dialect. A phased precedence model for the underlying representation is the following:

$G_1 = \{\langle k, y\rangle, \langle y, z\rangle, \langle z, l\rangle, \langle l, I_1\rangle, \langle I_1, g\rangle, \langle g, I_2\rangle\}$

$G_2 = \{\langle k, y\rangle, \langle k, z\rangle, \langle y, z\rangle, \langle y, l\rangle, \langle z, l\rangle, \langle z, I_1\rangle, \langle l, I_1\rangle, \langle l, g\rangle, \langle I_1, g\rangle, \langle I_1, I_2\rangle, \langle g, I_2\rangle\}$

$G_3 = \{\langle k, y\rangle, \langle k, z\rangle, \langle k, l\rangle, \langle y, z\rangle, \langle y, l\rangle, \langle y, I_1\rangle, \langle y, g\rangle, \langle z, l\rangle, \langle z, I_1\rangle, \langle z, g\rangle,$

$\langle l, I_1\rangle, \langle l, g\rangle, \langle l, I_2\rangle, \langle I_1, g\rangle, \langle I_1, I_2\rangle, \langle g, I_2\rangle\}$

$G_4 = \{\langle k, y\rangle, \langle k, z\rangle, \langle k, l\rangle, \langle k, I_1\rangle, \langle k, g\rangle, \langle k, I_2\rangle, \langle y, z\rangle, \langle y, l\rangle, \langle y, I_1\rangle, \langle y, g\rangle, \langle y, I_2\rangle, \langle z, l\rangle, \langle z, I_1\rangle$

$\langle z, I_2\rangle, \langle z, g\rangle, \langle l, I_1\rangle, \langle l, g\rangle, \langle l, I_2\rangle, \langle I_1, g\rangle, \langle I_1, I_2\rangle, \langle g, I_2\rangle\}$

Central Dialect

In the Central dialect, the Search proceeds as the following Search_1 for the Rnd harmony rule and Search_2 for the Back harmony rule:

- $\text{Search}_1 = \zeta_{C3} \cup G_3$ (Rnd Harmony)

  $= \{\langle[\alpha\text{Rnd}]\text{ , I }\rangle\} \cup \{\langle k, y\rangle, \langle k, z\rangle, \langle k, l\rangle, \langle y, z\rangle, \langle y, l\rangle, \langle y, I_1\rangle, \langle y, g\rangle, \langle z, l\rangle, \langle z, I_1\rangle, \langle z, g\rangle\}$

  $=\{\langle y, I_1\rangle\}$

- $\text{Search}_2 = \zeta_{C4} \cup G_4$ (Back Harmony)

  $= \{\langle[\alpha\text{Back}]\text{ , I }\rangle\} \cup \{\langle k, y\rangle, \langle k, z\rangle, \langle k, l\rangle, \langle k, I_1\rangle, \langle k, g\rangle, \langle k, I_2\rangle, \langle y, z\rangle, \langle y, l\rangle, \langle y, I_1\rangle, \langle y, g\rangle,$

  $\langle y, I_2\rangle, \langle z, l\rangle, \langle z, I_1\rangle\}$

  $= \{\langle y, I_1\rangle, \langle y, I_2\rangle\}$

The outputs of Search_1 and Search_2 are fed to the change, which output $\langle y, I_1 \sqcup \{+\text{Rnd},+\text{Back}\}\rangle = \langle y, y\rangle$ and $\langle y, I_2 \sqcup \{+\text{Back}\}\rangle = \langle y, i\rangle$ respectively.

Southern Dialect In the southern dialect, there is only one search on the last tier, proceeding as follows:

- $\text{Search} = \zeta_{S4} \cup G_4$ (Back and Rnd Harmony)

  $= \{\langle[\alpha\text{Back}]\text{ , I }\rangle\} \cup \{\langle k, y\rangle, \langle k, z\rangle, \langle k, l\rangle, \langle k, I_1\rangle, \langle k, g\rangle, \langle k, I_2\rangle, \langle y, z\rangle, \langle y, l\rangle, \langle y, I_1\rangle, \langle y, g\rangle,$

  $\langle y, I_2\rangle, \langle z, l\rangle, \langle z, I_1\rangle\}$

  $= \{\langle y, I_1\rangle, \langle y, I_2\rangle\}$

The output of the SEARCH is fed to the change, which outputs $\langle y, I_1 \sqcup \{+\textsc{Rnd}, +\textsc{Back}\}\rangle = \langle y, y\rangle$ and $\langle y, I_2 \sqcup \{+\textsc{Rnd}, +\textsc{Back}\}\rangle = \langle y, y\rangle$.

5.6 Global versus Local Application

One ramification of non-string representations is how it bears on the discussion on simultaneous/global versus iterative rule application. In global rule application, the entire string is first scanned for segments that satisfy the environmental constraints of the rule. After all such segments have been identified in the string, the changes required by the rule are applied simultaneously' (Chomsky and Halle, 1968). In local/iterative application, the change takes place as the string is scanned, possibly creating new environments for the next pass.

Consider the string *baaa* such that $b_1a_2a_3a_4$ and the rule $a \mapsto b \;/\; b \;\_\_$ where an a becomes b if there's a b directly preceding it.

Application of the rule can hypothetically result in two possible outputs that are the outcome of three potential distinct application types:

- Simultaneous application: The string is scanned globally for any a preceded by b, and once the scan is complete, each a found is changed into b simultaneously, resulting in the string *bbaa*.
- Iterative left to right application: The string is scanned from the left to the right, and each segment satisfying the rule is rewritten as the scan proceeds. a_2 is preceded by b so it gets rewritten as a b, a_3 is preceded by b which changed from a_2, so it also gets rewritten as b, and a_4 is preceded by the b that changed from a_3 and is also rewritten as b. The resulting string is *bbbb*.

- Iterative right to left application: The string is scanned left to write and segments updated in a similar on the go manner as 2 above, however the scan proceeds from the right to the left. Thus the scan starts with a_4 which is preceded by a_3, so it doesn't change. Next is a_3 which is preceded by a_2, it also doesn't change. And finally a_2 is preceded by b_1, so it changes to b. The resulting string is *bbaa* .

The distinct rule application methods are particular to string representations. In a precedence-based model, only specific ordered pairs undergo change, and so called 'iterative' outcomes the result of higher tier generating more environments for rule application. Rule application therefore proceeds 'globally' on the output of each SEARCH.

This model provides an account that parsimoniously explains multiple observations in phonology while simultaneously taking into account the evidence for the independence of temporal order from phonological representations.

Chapter 6

Discussion

As the “granularity mismatch” and “ontological incommensurately” problems outlined by Poeppel and Embick demonstrate, biological plausibility of a given phonological framework and the corresponding representational assumptions are a main concern for successfully bridging the gap between computation and implementation. Outlining a template for a phonological representation which takes into account the experimental literature on extra-linguistic processes like time perception narrows the hypothesis space for the necessary requirements in a phonological representation by abstracting away from the irrelevant information for phonology. Experimental evidence robustly supports the separation of order from a phonological representation. A precedence based representation along with domain general linearization and transitivization mechanisms seems more in line with experimental evidence while still maintaining a degree of predictive power.

6.1 Future Directions

An interesting corollary that the transitivization mechanism, specifically the transitivization tiers provide is their capacity to account for the scope of rule application with minimal machinery within the phonology. Rules that are seemingly sensitive to additional structures like syllable structure or morphological boundaries could possibly correlate with a given tier. While the examples provided in this thesis are simple for illustrative purposes, applying the transitivation approach to more language data could possibly provide insight into the types of environments that can be modeled using transitivization tiers and cover more empirical ground with the same machinery.

Considering the memory requirements for storing a representation, an interesting direction arises if the precedence relation were to be defined without the requirement of specifying and repeating the distinct precedence tokens. For example, representing the string *bababa* as $\{\langle b, a\rangle, \langle a, b\rangle\}$ without additionally specifying each subsequent occurrence, which maintains the proper phonological environments with minimal redundancy. This provides the phonology with the means to capture more generalizations while limiting memory and computational requirements. However, the serialization process can become overly complex, requiring additional machinery to determine proper order based on the precedence pairs. A means to overcome the added complexity could be to pass to the phonology at each step the new pairs generated by the composition of the graph of the previous pass with itself (each G_iG_i tier) and not the full graph, and stipulating whether a phonological rule applies to a single tier or multiple ones, with the rules targeting the 'longest distances' applying to the set resulting from the union of all tiers. This however seems to add more complexity to the phonology module, as well as forcing a certain pattern of rule ordering based on tier. Furthermore, the transitivization applied to such relations overgeneralizes.

For example, the string *badac* stored as $\{\langle b,a\rangle, \langle a,d\rangle, \langle a,d\rangle, \langle a,c\rangle\}$ produces the pair $\langle d,d\rangle$ in the second transitivization pass, which is an environment not available in the string representation.

Additionally, biological plausibility of this model can be experimentally tested by creating toy language stimuli to probe if precedence is sufficient for phonology. 'Repetition blindness'— the phenomenon whereby a repeated stimulus fails to be registered as having occurred—outlined in Kanwisher (1987) as the result of "types/token" distinction, can potentially be used as an experimental paradigm. Linear order requires 'tokens': each segment should necessarily be instantiated in order to be distinguished (for the purposes determining meaning for example); precedence, on the other hand, might operate over types. Determining whether repetition blindness occurs in different environments (phonological and non-phonological) may shed light on the true nature of the representation.

Finally, as the conclusions on time perception transcend phonology, similar approaches across cognitive domains, one where time is taken seriously, can provide further insight and shed light on overarching themes in cognition to further help bridge the gap between mind and brain. As the ultimate goal of cognitive neuroscience is to ground cognition in the physical brain, Poeppel's (2012) encouragement on cross-domain collaboration is key for such endeavors.

6.2 Conclusion

In this thesis, I have outlined the experimental evidence on time perception that supports the separation of temporal order from representations that seem to phenomenally contain it. Phonological representations have generally been described as linearly ordered segments, and the computations defined in terms of the status of the representation, with a trade-off existing between both. Accounting for long-distance rules under such models requires either enriching the representation or alternatively, the computations. The adoption of a precedence-based representation is sufficient for accounting for distinct classes of phonological patterns parsimoniously, while simultaneously alleviating the computational burden due to the existence of domain-general central timing mechanisms.

Chapter 2 introduced phonology and the Search and Change model that accounts for long-distance phonological processes by enriching the computation. Chapter 3 provided a thorough description of the requirements in a representation and the properties that a particular representation should capture. Chapter 4 offered a background on time perception as well as experimental evidence demonstrating the modality independent nature of temporal order. Chapter 5 outlined a simplified architecture that takes into account the decoupleability of temporal order from phonological representations by adopting a precedence-based model and a central transitivization mechanism which translates a stored immediate-precedence graph into a fully transitive precedence relation in phases. Finally, chapter 6 contains a discussion section which offers potential directions to explore such as the empirical coverage that the transitivization mechanisms provides as well as possible experimental ideas to support testing whether phonology operates over types or tokens.

To that end, the neural mechanisms underpinning phonological computations, should be a bit easier to identify due to a clearer delineation, taking into account the temporal integration mechanisms that should be discounted.

Bibliography

Andersson, Samuel, Hossep Dolatian, and Yiding Hao. 2020. Computing vowel harmony: The generative capacity of search & copy. *Proceedings of the Annual Meetings on Phonology* 8.

Appelbaum, Irene. 1996. The lack of invariance problem and the goal of speech perception. *Proceeding of Fourth International Conference on Spoken Language Processing. ICSLP '96* 3:1541–1544 vol.3.

Bale, Alan, and Charles Reiss. 2018. *Phonology: A formal introduction*. Cambridge, MA: MIT Press.

Basharat, Aysha, Meaghan Adams, William Staines, and Michael Barnett-Cowan. 2018. Simultaneity and temporal order judgments are coded differently and change with age: An event-related potential study. *Frontiers in Integrative Neuroscience* 12.

Bechlivanidis, Christos, and David Lagnado. 2013. Does the "why" tell us the "when"? *Psychological Science* 24:1563–1572.

Bechlivanidis, Christos, and David A. Lagnado. 2016. Time reordered: Causal perception guides the interpretation of temporal order. *Cognition* 146:58–66.

Berwick, Robert, and Noam Chomsky. 2008. The biolinguistic program: The current state of its evolution and development. *In Di Sciullo Boeckx* .

Binkofski, Ferdinand, and Richard A. Block. 1996. Accelerated time experience after left frontal cortex lesion. *Neurocase* 2:485–493.

Buzsaki, G. 2019. *The brain from inside out*. OXFORD University Press.

Carmon, Amiram, and Israel Nachshon. 1971. Effect of unilateral brain damage on perception of temporal order. *Cortex; a journal devoted to the study of the nervous system and behavior* 7 4:411–8.

Chandlee, Jane, Rémi Eyraud, Jeffrey Heinz, Adam Jardine, and Jonathan Rawski. 2019. Learning with partially ordered representations. 91–101.

Cheatham, Paul G., and C. T. White. 1952. Temporal numerosity: I. perceived number as a function of flash number and rate. *Journal of Experimental Psychology* 44:447.

Chemero, Anthony. 2009. *Radical embodied cognitive science*. Bradford.

Chomsky, Noam. 1965. *Aspects of the theory of syntax*. Cambridge, MA: MIT Press.

Chomsky, Noam. 1995. *The minimalist program*. Cambridge, Massachusetts: MIT Press.

Chomsky, Noam. 1998. Some observations on economy in generative grammar. In *Is the best good enough?*, ed. Pilar Barbosa, Danny Fox, Paul Hagstrom, Martha McGinnis, and David Pesetsky 115–128. Cambridge, Massachusetts: MIT Press.

Chomsky, Noam, and Morris Halle. 1968. *The sound pattern of English*. New York: Harper & Row.

Dabbous, Rim, Marjorie Leduc, Fatemeh Mousavi, Charles Reiss, Ta-Chun David, and Shen. 2021. Satisfying long-distance relationships (without tiers) a strictly anti-local approach to phonology *.

Dede, Adam, Jennifer Frascino, John Wixted, and Larry Squire. 2016. Learning and remembering real-world events after medial temporal lobe damage. *Proceedings of the National Academy of Sciences of the United States of America* 113.

Eijkman, E., and A. Vendrik. 1965. Can a sensory system be specified by its internal noise? *The Journal of the Acoustical Society of America* 37:1102–9.

Elliott, Mark, and Anne Giersch. 2016. What happens in a moment. *Frontiers in Psychology* 6.

Embick, David, and David Poeppel. 2014. Towards a computational(ist) neurobiology of language: Correlational, integrated, and explanatory neurolinguistics. *Language, Cognition and Neuroscience* 30:1–10.

Emery, Nina, Ned Markosian, and Meghan Sullivan. 2020. Time. In *The Stanford encyclopedia of philosophy*, ed. Edward N. Zalta. Metaphysics Research Lab, Stanford University Winter 2020 edition.

Fodor, Jerry A. 1981. The mind-body problem. *Scientific American* 244 1:114–20, 122–3.

Fodor, Jerry A. 1983. *The modularity of mind: an essay on faculty psychology*. Cambridge, MA: MIT Press.

Fodor, Jerry A., and Zenon W. Pylyshyn. 1981. How direct is visual perception?: Some reflections on gibson's "ecological approach". *Cognition* 9:139–196.

Geldard, F. 1976. The saltatory effect in vision. *Sensory processes* 1:77–86.

Geldard, Frank A, and Carl E Sherrick. 1972. The cutaneous" rabbit": a perceptual illusion. *Science* 178:178–179.

Gibson, James J. 1979. *The ecological approach to visual perception.* Boston: Houghton Mifflin.

Goldsmith, John A. 1979. *Autosegmental phonology.* Garland Publishers.

Hale, Mark. 2003. Neogrammarian sound change. In *Handbook of historical linguistics*, ed. Richard D. Janda Brian D. Joseph 343–368. Oxford, UK: Blackwell Publishing Ltd.

Hale, Mark, and Charles Reiss. 2000. "Substance abuse" and "dysfunctionalism": Current trends in phonology. *Linguistic Inquiry* 31:157–169.

Hale, Mark, and Charles Reiss. 2008. *The phonological enterprise.* Oxford University Press UK.

Hammarberg, Robert. 1981. The cooked and the raw. *Journal of Information Science* 3:261–267.

Hirsh, Ira, and C.E. Sherrick. 1961. Perceived order in different sense modalities. *Journal of experimental psychology* 62:423–32.

Hoffman, Donald D. 2016. The interface theory of perception. *Current Directions in Psychological Science* 25:157–161.

Hornstein, Norbert. 2008. A theory of syntax: Minimal operations and universal grammar. *A Theory of Syntax: Minimal Operations and Universal Grammar* 1–194.

Idsardi, William, and Eric Raimy. 2013. Three types of linearization and the temporal aspects of speech. In *Challenges to linearization*, ed. Theresa Biberauer and Ian Roberts 31–56. Berlin, Boston: De Gruyter Mouton.

Jackendoff, Ray. 1991. The problem of reality. *Noûs* 25:411–33.

James, William. 1890. *The principles of psychology* volume 1. New York: Henry Holt.

Jensen, John. 1974. A constraint on variables in phonology. *Language* 50:675.

Kanwisher, Nancy G. 1987. Repetition blindness: Type recognition without token individuation. *Cognition* 27:117–143.

Karimi, Kamran. 2010. A brief introduction to temporality and causality. *Computing Research Repository - CORR* arXiv:1007.2449.

Kayne, Richard S. 1994. *The antisymmetry of syntax*. Cambridge, Massachusetts: MIT Press.

Kosslyn, Stephen, Thomas Ball, and Brian Reiser. 1978. Visual images preserve metric spatial information: Evidence from studies of image scanning. *Journal of experimental psychology. Human perception and performance* 4:47–60.

Kosslyn, Stephen Michael. 1980. *Image and mind*. Harvard University Press.

Lasnik, Howard. 2002. The minimalist program in syntax. *Trends in cognitive sciences* 6:432–437.

Mailhot, Frederic, and Charles Reiss. 2007. Computing long-distance dependencies in vowel harmony. *Biolinguistics* 1.1:28–48.

McCollum, Adam G., and Darya Kavitskaya. 2018. Non-iterative vowel harmony in Crimean Tatar.

McFarland, Joseph. 1970. A visual illusion of succession. *Perception & Psychophysics* 7:43–46.

Nevins, Andrew. 2010. *Locality in vowel harmony*. Cambridge, MA: MIT Press.

Palatinus, Zsolt, and Claire Michaels. 2014. *A ten commandments for ecological psychology*.

Papillon, Maxime. 2020. Precedence and the lack thereof: Precedence-relation-oriented phonology. Doctoral Dissertation University of Maryland.

Poeppel, David. 2012. The maps problem and the mapping problem: Two challenges for a cognitive neuroscience of speech and language. *Cognitive Neuropsychology* 29:34–55.

Poeppel, David, and David Embick. 2005. Defining the relation between linguistics and neuroscience. *Twenty-first century psycholinguistics: Four cornerstones* 103–118.

Posner, Michael, and Gregory Digirolamo. 2000. Cognitive neuroscience: Origins and promise. *Psychological bulletin* 126:873–89.

Pylyshyn, Zenon. 1981. The imagery debate: Analogue media versus tacit knowledge. *Psychological Review* 88:16–45.

Pylyshyn, Zenon W. 2007. *Things and places: how the mind connects with the world*. The Jean Nicod lectures. Cambridge, Mass.: MIT Press. URL `http://www.loc.gov/catdir/toc/ecip073/2006035507.html`.

Pöppel, Ernst. 1988. Time perception. 124–125. Springer.

Pöppel, Ernst. 1997. A hierarchical model of temporal perception. *Trends in cognitive sciences* 1:56–61.

Pöppel, Ernst. 2004. Lost in time: A historical frame, elementary processing units and the 3-second window. *Acta neurobiologiae experimentalis* 64:295–301.

Raimy, Eric. 2000a. *The phonology and morphology of reduplication*.

Raimy, Eric. 2000b. *The phonology and morphology of reduplication*. Berlin: M. de Gruyter.

Rizzi, Luigi. 1990. *Relativized minimality*. Cambridge, Massachusetts: MIT Press.

Roberts, Seth. 1998. The Mental Representation of Time. In *An Invitation to Cognitive Science: Methods, Models, and Conceptual Issues*. The MIT Press.

Sapir, Edward. 1949. The psychological reality of phonemes. In *Selected writings of Edward Sapir*, ed. D. Mandelbaum 46–60. Berkeley and Los Angeles: University of California Press.

Shore, David, Susan Hall, and Raymond Klein. 1998. Auditory saltation: A new measure for an old illusion. *The Journal of the Acoustical Society of America* 103:3730–3.

Sternberg, Saul, and Ronald Knoll. 1973. The perception of temporal order: Fundamental issues and a general model. *Attention and Performance IV* 4.

Szelag, Elzbieta, Nicole von Steinbüchel, and Ernst Pöppel. 1997. Temporal processing disorders in patients with broca's aphasia. *Neuroscience Letters* 235:33–36.

Thagard, Paul. 2023. Cognitive Science. In *The Stanford encyclopedia of philosophy*, ed. Edward N. Zalta and Uri Nodelman. Metaphysics Research Lab, Stanford University Spring 2023 edition.

Treisman, Anne, and Garry Gelade. 1980. A feature-integration theory of attention. *Cognitive psychology* 12:97–136.

Warrren, Richard M. 1970. Perceptual restoration of missing speech sounds. *Science* 167:392–393.

Wearden, John. 2016. *The psychology of time perception*. Springer.