Vrije Universiteit Brussel

Faculteit Wetenschappen en Bio-ingenieurswetenschappen
Vakgroep Computerwetenschappen

Simulating the Shift Towards Semantic Gender in Dutch

A Multi-Agent Language Game Approach

Proefschrift ingediend met het oog op het behalen van de graad van
Master of Science in Applied Sciences and Engineering: Computer Science

Roxana-Teodora Rădulescu

Promotor: Prof. Dr. Katrien Beuls
Begeleider: Dr. Remi van Trijp

Augustus 2015

Vrije Universiteit Brussel

Faculty of Science and Bio-Engineering Sciences
Department of Computer Science

Simulating the Shift Towards Semantic Gender in Dutch

A Multi-Agent Language Game Approach

Thesis submitted in partial fulfilment of the requirements for the degree of
Master of Science in Applied Sciences and Engineering: Computer Science

Roxana-Teodora Rădulescu

Promotor: Prof. Dr. Katrien Beuls
Advisor: Dr. Remi van Trijp

August 2015

Abstract

Studying language origin and evolution seems to be one of the hardest endeavours that human kind has faced so far. On the one hand, language has an ephemeral character, having left no traces behind that pre-date written language. On the other hand, studying such a complex system engages the need of an interdisciplinary framework.

In the context of cultural evolution, evolutionary linguistics proposes such a framework, in the form of a multi-agent model, that allows, through pair-wise interactions among individuals in a population, a self-organized linguistic system to emerge. If the created artificial languages exhibit features that resemble those of natural ones, then one can state to have built a model of how such a feature can culturally develop.

Our attention was grasped by the particularities of the Flemish dialects. At a closer glance, one can discover a language system in mid-transition process in terms of gender agreement systems, under the pressure of standardization, marker erosion and cross-dialect influences.

We build here a multi-agent language game, in order to simulate the transition from a syntactic towards a semantic gender agreement system. We evaluate and demonstrate what are the minimal abilities agents require, such that a gender mapping over the semantic space can emerge at population level.

Declaration of Originality

I hereby declare that this thesis was entirely my own work and that any additional sources of information have been duly cited. I certify that, to the best of my knowledge, my thesis does not infringe upon anyone's copyright nor violate any proprietary rights and that any ideas, techniques, quotations, or any other material from the work of other people included in my thesis, published or otherwise, are fully acknowledged in accordance with the standard referencing practices. Furthermore, to the extent that I have included copyrighted material, I certify that I have obtained a written permission from the copyright owner(s) to include such material(s) in my thesis and have included copies of such copyright clearances to my appendix.

I declare that this thesis has not been submitted for a higher degree to any other University or Institution.

Acknowledgements

First of all, I would like to express my profound gratitude towards my advisor, Katrien Beuls. Thank you for giving me the opportunity of engaging into language games, and also for guiding me through the maze of the Babel framework. Thank you Remi and Katrien for the endless patience of explaining Fluid Construction Grammar to me, it might have taken a while, but I did get there in the end.

A big thank you should go to all my Dutch speaking friends, as not only once have I challenged their language proficiency in terms of agreement systems.

Lastly, I have to express my infinite gratitude towards my parents, without whom my Belgian adventure would have never been possible in the first place.

Contents

List of Figures

List of Tables

1
Introduction

There are no verbs preserved in amber, no ossified nouns, and no prehistorical shrieks forever spread-eagled in the lava that took them by surprise.

– Christine Kenneally, The First Word:
The Search for the Origins of Language (2007)

It is odd to think that we are able to send robots to explore our neighbour planets and the vastness of space, while we still cannot explain how and why we came about to using such a complex communication system as language. To be fair, however, any task in this setting would be equivalent to simply stumbling in the dark, as there is little to no trace related to language development and evolution to follow up on, as Christine Kenneally marvellously stated. It should not come as a surprise that even the history of language study is an intricate maze with many dead ends and diverging paths. After Darwin, in his *Origin of Species*, 1859, first associated the idea of evolution with languages, a chaotic agglomeration of theories were launched. The situation seemed so tensed but hopeless at the same time that, in 1866, the high language authority at that time, the Linguistic Society of Paris, banned any discussion or speculation on the origin of language, as it was not considered to be worth as a scientific endeavour, given the lack of evidence

in this regard. For nearly a century, peace settled over this realm, before further advancements and ideas from parallel fields allowed for any step forward to be taken in language study.

Even so, Christiansen & Kirby (2003b) have launched the question *is language evolution the hardest problem in science?* The encumbrance of language study comes from the nature of the task involved – a multi-disciplinary effort in understanding the human brain, its evolution and the mechanisms behind language, language itself, its purpose, structure and shaping forces (biological, cultural, social).

Hard as it may be, unravelling the mystery of language, its origin and evolution process, is a task that humanity will never abandon, as it also has its subtle enchantment in making us feel unique among all the other creatures in this world.

Thus, how did language emerge, what shapes and constrains it? Van Trijp (2008, p.30-33) distinguishes between three theories in order to provide a possible answer: language as a result of adaptation (Pinker & Bloom, 1990), language as a secondary product of another function's evolution (Gould & Vrba, 1982) and lastly, the recruitment theory of language (Steels, 2007). Dwelling a bit more on the last concept, as it is also the point of view adopted for this work, the recruitment theory proposes languages as a 'dynamic configuration of brain mechanisms', forever shifting, changing and evolving such that communication can be carried out at optimal levels (power of expression and successful interactions, while maintaining a low cognitive effort for all interlocutors).

1.1 Artificial Intelligence and Language Study

Steels' approach on language study imposes not only a reversed view on causality and effect (as we do not search for the genetic basis of language, but rather for the minimal cognitive mechanisms that can give rise to certain linguistic properties), but also shifts this study into the field of artificial intelligence.

Artificial languages, emerged from computational linguistic models under the form of multi-agent decentralized systems, can be studied in order to evaluate if and how cultural evolution can be held responsible for features we notice in natural languages, if they are exhibited by the artificially emerged system. Tasks that need to be performed in order to build such a system include: defining an agent's cognitive abilities, providing an interaction script for the model, defining constraints and learning mechanisms.

For this work, we set out to build a multi-agent system that attempts to simulate a linguistic phenomenon currently observed in present day Dutch, more specifically in some Flemish dialects.

1.2 Case Study – Dutch

Even though the Dutch speaking population and its territory are not very spread, the variation and diversity of the dialects that compose it make it an interesting object of study. Dutch has a rich, tumultuous history, that shaped it in an unprecedented manner: with the southern part under foreign occupation, separated from the linguistically thriving north, lagging behind even nowadays in terms of linguistic processes and standardization. This region thus offers us today a direct window towards language change and evolution and the opportunity to observe system shifts in mid-transition.

We focus here on the southern Dutch, more specifically on the Dutch speaking region of Belgium, Flanders. We are interested in the mechanisms of gender agreement in Flemish, namely how the gender of a noun is reflected in a pronominal referent. A small example is presented in sentence 1. One can immediately notice the difference between the Dutch pronominal gender system (which in this case causes a use of the masculine pronoun 'hij' for the word 'koffie') and that of English (for which the correct form should be the neutral 'it').

(1) A: Is de **koffie** klaar?
'A: Is the **coffee** ready?'

B: Ja, **hij** staat op de tafel.
'B: Yes, **he** is standing on the table.'

However, because Flemish dialects are also starting to suffer from a lack of gender indicators (event that already took place in northern Dutch), the language users are faced with the problem of an increasingly opaque system, making it more and more difficult to make the correct agreement choices. This task is even harder for non-native Dutch speakers, for which even the two-gender system (common 'de' and neuter 'het') of Dutch is hard to acquire.

Studies on this issue have revealed that a possible solution is represented by a shift towards the semantic properties of the nouns, as indicators for gender mapping, an idea that we examine in our work, under the form of a computational linguistic model.

1.3 Research Question and Thesis Structure

We have tried to create a fine outline of the context that our research is taking place in, crossing the borders of evolutionary linguistics, multi-agent models and gender agreement system in linguistics.

What we attempt here is to build a multi-agent model, in the context of evolutionary language games, that can simulate a simple form of pronominal reference. We wish to engage a competition between syntactic and semantic gender agreement systems, while seeking out language strategies that can perform the transition from the syntactic to the semantic mapping of pronominal gender. We perform this study starting from a syntactically marked gender vocabulary of Dutch nouns, selected to be distributed over a chosen semantic space.

As for the structure of the thesis, Chapter 2 grounds this work in one of the approaches for language study, evolutionary linguistics. We present a deeper view of language and language study movements, define the fundamental concepts of evolutionary linguistics and also the tools we use to build our model.

Chapter 3 emerges the reader into a study of gender system in Dutch (northern and southern), its evolution, current state and arising issues. We explain how we inspire our model from these problems and what are the mechanisms we implement to simulate it.

Chapter 4 becomes more practical oriented, as it presents a full description of the multi-agent system we built for our case study, all possible configurations, strategies and mechanisms. Additionally we describe our language game script, the vocabulary of the agents and the tools to evaluate our system.

Chapter 5 presents all the performed experiments, under a methodical approach, such that we study the impact of all our parameters and strategies, plus a discussion for the observed results.

Finally, Chapter 6 summarizes our accomplishments, presenting additional directions for enhancing our linguistic model.

2

Agent-Based Modelling in Evolutionary Linguistics

We start by exploring language, as a complex natural system, and approaches to study its continuous evolutionary process. We define and present, in the context of evolutionary linguistics, available methods, frameworks and explain all the necessary concepts used to develop the current work.

Language is a human communication system which has grown to be one of the most elaborate living networks ever to have arisen in nature. Language has a dual nature, on the one hand there is the external manifestation that comes across in the form of speech or gestures to build a speaker's utterance and on the other hand it has an internal representation personalized in each individual (Kirby, 1999). Language is a *complex adaptive system*, as, firstly, it aggregates at a microscopic level a population of individual language users, who *adapt* according to previous experienced interactions amongst each other as well as to environmental pressures. Secondly, the *complexity* emerges from the impossibility of trivially predicting the outcome and evolution trajectory of such a system, by only looking at the individual level, as there is an ongoing process of change and self-organization at both micro- and macroscopic scale (Smith, 2014).

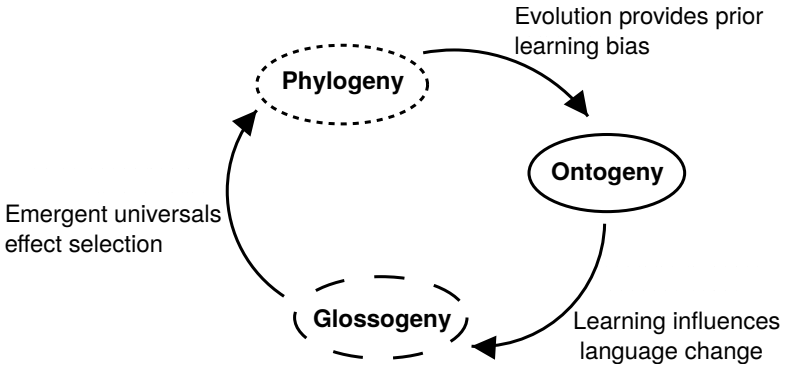

Figure 2.1: An overview of the three complex adaptive systems (and their interactions) responsible for the perpetual language evolution and change (Kirby & Hurford, 2002). Phylogeny, the biological evolution, offers the physical adaptation for language use. Ontogeny is the learning done at individual level, from which consensus and new strategies can arise, elements that can then be spread and persist through cultural evolution, glossogeny.

2.1 Studying Linguistic Evolution

In order to fully comprehend language evolution, Kirby & Hurford (2002) has placed language at the crossing of three complex adaptive systems (Figure 2.1) that are in a continuous spiralling motion, pushing it upon perpetual innovation, development and change, each at a different time-scale:

- *Phylogeny* or the biological evolution provides individuals with the mechanisms necessary for language acquisition and use, adapting continuously, over large time periods, in order to cope with environmental challenges and communication needs.

- *Ontogeny* is represented by the learning, development and adaptation of linguistic skills at individual level, in response to ones experiences and encountered situations.

- *Glossogeny* is the historical or cultural change of languages, ranging from minor elements, like spelling or pronunciation, to the shifting or expansion of meanings, word acquisitions, loss or change of language systems (e.g., gender, case systems).

There is a strong connection between these systems, as they are interacting and pushing each other towards increasing complexity. As it has been pointed out

(Christiansen & Kirby, 2003a; Hauser et al., 2002), studying linguistic evolution becomes an inter-disciplinary task, as it crosses the boundaries of fields like evolutionary biology, neuroscience, anthropology, linguistics (Figure 2.2). Building realistic models brings forth a difficult task, as one needs to consider constraints and apply concepts from all areas, by understanding how the human brain might have adapted for language formulation and comprehension, what are the aspects that differ between different stages of a language throughout history, how language can develop and change inside a population and so on.

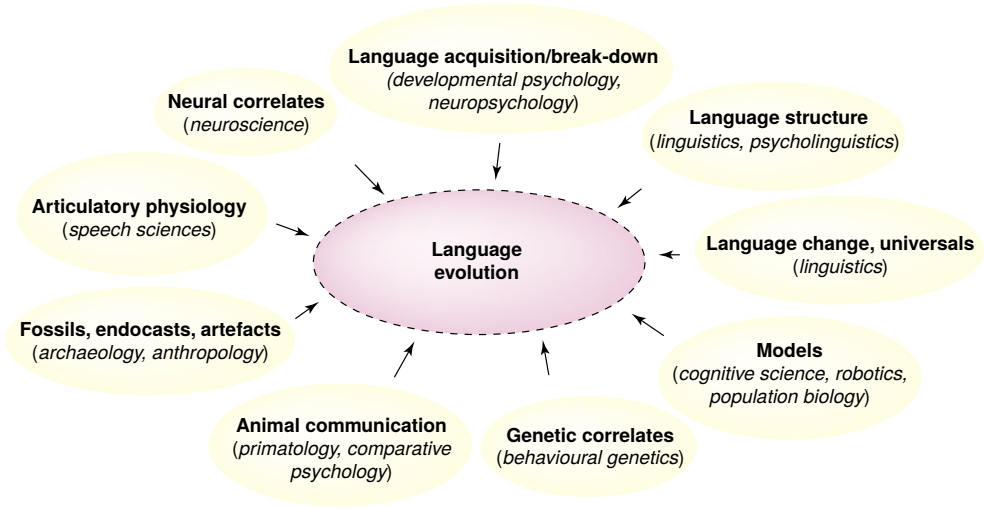

Figure 2.2: Language evolution seen as an inter-disciplinary effort (Christiansen & Kirby, 2003a). The figure presents concepts and their originating fields necessary for studying the evolution of language.

We can distinguish different movements for studying the emergence and evolution of language, depending on which language scale is emphasised.

2.1.1 Biological Evolution

The biological evolution direction is centred around the concept of linguistic nativism. It is considered that the ability to acquire language and complex grammar despite the lack of examples and clear instructions, impediment defined as the 'poverty of stimulus' (Chomsky, 1980), is endowed by a biological adaptation of the human brain, providing humans with an innate universal structure underlying all languages, the 'Universal Grammar' (Chomsky, 1965).

Additional contribution for this direction was made by Pinker & Bloom (1990), who explain, in the context of evolutionary theory, how adaptations for language

learning can become innate through the Baldwin effect. This phenomenon claims that the ability of learning a behaviour can become genetically encoded, if that provides a species with a higher fitness to survive. Their theory thus suggests that individual language acquisition and expertise from a young age is possible due to language specific cognitive abilities that have been developed and tuned in the biological driven process of natural selection.

2.1.2 Cultural Evolution

Rather than considering that the human brain has adapted to learn language, this direction conveys the idea that languages are the product of a continuous evolutionary process in a cultural realm. Sound, form and structure are being put under a perpetual selection pressure in terms of ease of acquisition, communication or transmission, in the context of human cognition and vocal tract capabilities.

A first important idea behind this direction comes from the fact it is impossible to observe the immensity of form and meaning combinations of a language. This causes a *transmission bottleneck* to occur (Kirby, 2000), phenomenon that then becomes a cultural force that drives the need for the ability to generalize beyond the perceived data, thus for a systematic mapping, a structured language.

Another element that defines a language that emerges from cultural evolution comes from the fact that speakers do not have access to each others internal representation of meanings. Each individual has to infer an approximate meaning given a common context within an interaction, thus each language user ends up having a personalized internal mapping (Smith, 2014). However, again during an interaction, for a successful communication between speakers, the exhibited external behaviour of this representation has to yield a meaning reproducible from the observed common context, grounding thus language in the outer world.

2.2 Evolutionary Linguistics

> *The goal of evolutionary linguistics is not to build systems that can parse or produce English, but to understand the generic forces and mechanisms that give rise to these capabilities.*
>
> – Steels (2003)

Evolutionary Linguistics is a selectionist theory of the evolution of language which aims to unravel the origins of language and explain what are the minimum cognitive abilities that allow this system to evolve and exhibit certain features

(e.g., agreement, case system, constituent structure). As tools, it proposes a mapping of evolutionary biology concepts to the realm of linguistics, in the cultural evolution context.

Languages should be unrestricted as speakers need the ability to convey any new encountered reality from the existing forms. This communication system is not equipped with any central control mechanism, thus is subjected to a lack of uniformity within a population, as it is impossible for the members to be aware of all deprecations or inventions taking place. These attributes have inspired the idea of selectionism as a viable approach for language design (Steels, 2012). Building a system in which each individual is equipped with human level cognitive abilities and allowing them to interact within a specific goal realization, can give rise to surprising results in terms of the emerged self-organized communication system. We discuss next a few additional concepts that are fundamental in designing and building such a system.

2.2.1 Language Systems and Strategies

A *language system* extends across morphological, syntactic and phonological levels, to comprise all the rules that create a systematic module of a language in terms of both meaning and form. For example, in order to express participant roles for an action, Romanian has a five case system, inherited from Latin (nominative, accusative, genitive, dative and vocative). In the sentence "I-a înapoiat cartea colegului său." ("He returned the book to his colleague") we can notice that 'cartea' is the direct object in singural, accusative and 'colegului' is the indirect object, again singular, marked with the specific ending of dative '-lui'. For the current work, we focus on the Dutch gender system, which we will thoroughly describe in the following chapter.

On the other hand, a *language strategy* is the underlying approach towards realizing the language system. For example in our case, we will notice two language strategies (syntactic and semantic) competing for the pronominal gender system in Dutch.

Figure 2.3 presents an overview of how language change takes place at the level of language systems and strategies. In the context of interactions within a population, the language system is directly used in utterance production. In case of successful interactions, the used norms will be adopted and transmitted. On the long run, the strategy underlying the most successful language system will spread and be adopted inside the population.

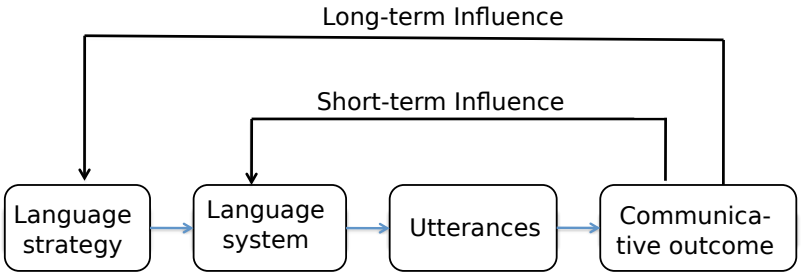

Figure 2.3: The reinforcement loops created between communicative outcomes and language strategies and systems (Steels, 2012). Language systems that produce a successful communication are rewarded by being spread in the population, along with the underlying strategy that generates it.

2.2.2 Evolutionary Linguistics Paradigms

Whether an interaction between language users is successful or not can be measured in terms of goal accomplishment (performing a certain action, or receiving an answer to a question), cognitive effort (the mental endeavour that user has to perform in order to convey his message within one or multiple language systems), learnability (ease of extending ones vocabulary and acquiring new forms). All these factors act as *selection criteria* at population level, leading to a fit linguistic system under some imposed constraints or objective.

Lastly, we mention the role of *cognition* for evolutionary linguistics. To begin with, we consider that all language users share a common set of cognitive abilities that form the basis for a social interaction. These abilities are not necessarily adapted exclusively for linguistic capacity. Furthermore, these capabilities can act as monitoring and intervention mechanisms when producing or parsing speech. This can be represented in an evolutionary linguistic system as a two-layer architecture, where the first level is used in *routine formulation and comprehension* of utterances, while a learning meta-level is used for *diagnostics* and *repairs* when necessary (e.g., encountering an unknown word, adding agreement markers for ease of disambiguation).

2.3 Multi-agent Models

Agent-based models for studying linguistic evolution have proven to be successful in demonstrating how language emergence and acquisition is possible using general learning strategies, removing the need for any specific language abilities (Smith, 2014).

The general setting for these types of computational models consists usually of

a population of agents, playing the role of language users, equipped with chosen cognitive abilities such that a certain goal can be achieved. They share the same world and each interaction is grounded within a context (a subset of the world). These models make use of social transmission mechanisms such as teaching (vertical transmission) or imitation (horizontal transmission) (Mesoudi, 2011). Depending on the focused mechanism, we distinguish here between two main types of models: iterated learning (vertical transmission) and language games (horizontal transmission).

The iterated learning framework (Kirby, 2002b,a) focuses mainly on *vertical transmission*, considering a split in the population between adults and learners. It was initially used to demonstrate how the transmission bottleneck (introduced by the E-Language, Figure 2.4) gives rise to structured languages that exhibit two main characteristics: compositionality and recursion.

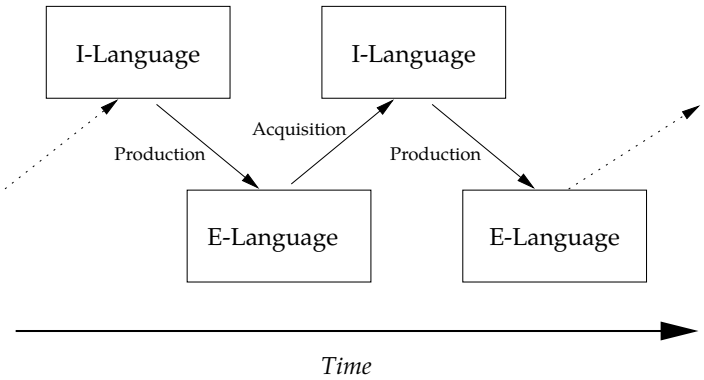

Figure 2.4: Iterated learning transmission model in time (Kirby, 2002a). It applies the idea of language existing in two forms, internal representation I-Language and external manifestation, E-Language. The acquisition is done through observational learning of the external behaviour.

2.3.1 Language Games

In order to study language, we desire to have a fully situated environment, meaning interacting partners (speaker and hearer), a shared world and context, a communication goal and a population of language users. Thus we make use of a concept introduced by (Wittgenstein, 1958), a *language game*, to denote a grounded interaction between a pair of language users from within the population. Our evolutionary model is composed by a sequence of language games, between randomly chosen interlocutors, during which we monitor the system and then report the observed outcome.

In contrast with iterated learning, language games focus more on *horizontal transmission*, therefore on how the system can evolve and change, achieving a social coordination at macroscopic level, during the interaction between individual members of the same generation. Additional mechanisms can be introduced to have, for example, a mortality rate, in order to study the effect of occasional vertical transmission inside this system, but in general this is not done using any fitness measures for the agents, but just in the form of a flow of agents from and into the population.

The goal of these models is to determine what initial cognitive abilities the individual agents need, such that, under a certain goal or interaction script, the population can converge towards a stable communication system. In case this system exhibits aspects comparable or mappable to human language, we can infer that we obtained a possible model of how that feature can culturally develop.

2.3.2 The Semiotic Cycle

As mentioned before, a language game consists in an interaction between two randomly selected agents, marked as speaker and hearer (listener), and it follows a script composed of several well-defined steps. The *semiotic cycle* is represented by this sequence of steps, as it is depicted in Figure 2.5. The agents build an internal model of the world, which is then used to either conceptualize and produce an utterance (in the speaker's case) or to parse and interpret the utterance (in the hearer's case).

2.3.3 Learning and Alignment

During or after an interaction agents may need to recourse to their additional cognition layer (implemented in a language game model as a double-layered architecture presented in Figure 2.6) in order to extend their vocabularies and express additional meanings or to add new encountered words or structures, process denoted at learning. Additionally, at the end of each encounter, each agent makes use of an alignment mechanism, depending on the outcome of the interaction, to try to bring their internal representations closer to each other. In a following chapter, we will present in more details the implemented learning and alignment procedures used for the current work.

2.4 Construction Grammar

Construction grammar starts from the Saussurean notion of linguistic sign (De Saussure, 1916), unit that comprises an arbitrary association between form (conven-

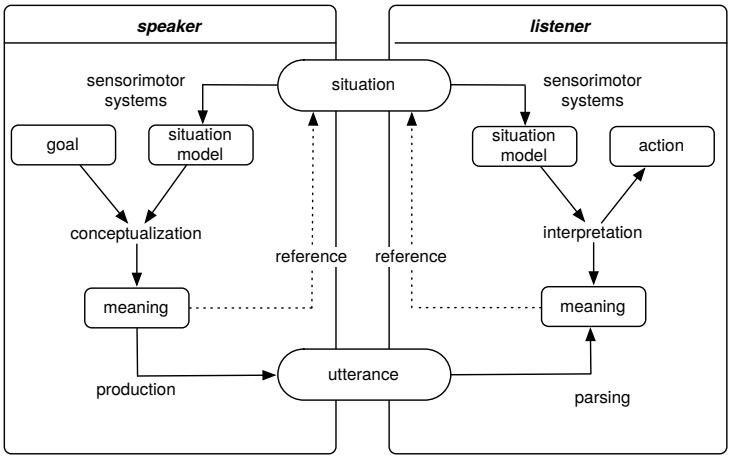

Figure 2.5: The semiotic cycle (Beuls et al., 2012) portrays the connection between a speaker, a hearer and the world, in the context of an interaction. This is a generalized skeleton of a language game, upon which one can build as needed in the case of specific experimental models.

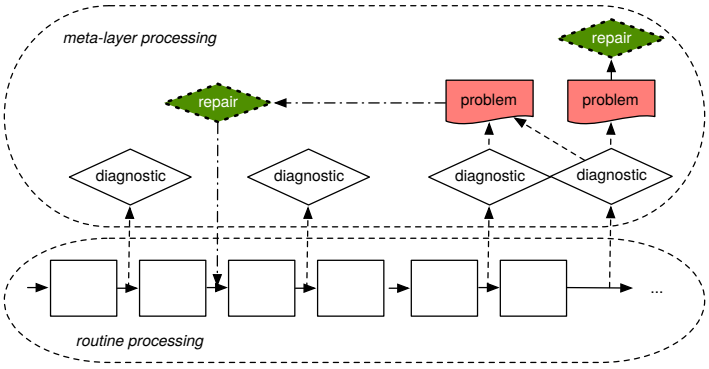

Figure 2.6: The double-layered architecture in a language game model (Beuls et al., 2012) allows for problem diagnosis and repair during or at the end of an interaction.

tional symbol) and meaning (internal representation). As a simple example to demonstrate this concept, lets consider the words 'tree' in English, 'copac' in Romanian, 'boom' in Dutch (and the list can go on), forms that all have the same associated meaning of a *tree*. Extending the notion of linguistic sign over all grammatical levels, one can define a *construction* (that can represent thus anything from a simple morpheme, to words, or even entire phrasal structures).

Consequently, construction grammar proposes a different view over the former lexicon – syntax separation, modelling grammar as a mental network of con-

structions (Hoffmann & Trousdale, 2013), defined as the‘constructicon’. A construction thus represents a rich structure containing information about and also establishing a mapping between the meaning and the form of a word (or phrase, or any other abstraction one needs in order to deal with functional, argument or information structure), spanning across multiple linguistic layers (morphology, phonology, syntax, semantics).

2.4.1 Fluid Construction Grammar

Fluid Construction Grammar (FCG) (Steels, 2011) is a computational grammar formalism developed as a tool for implementing grammars in the context of linguistic cultural evolution models (Van Trijp et al., 2012). Taking a cognitive-functional approach towards language (Wellens et al., 2013), FCG provides bidirectional processing functionalities, being capable of both parsing (form $\rightarrow$ meaning) and producing (meaning $\rightarrow$ form) utterances, while utilising a single construction inventory.

The vocabulary of the agents is build using FCG. In order to produce and parse utterances the agents use the mechanisms encapsulated by FCG (centred around the unification operation, implemented in two steps, matching and merging). A full description, supported by examples, of how the inventory of an agent is build and handled will follow in the next chapters.

All in all, for the current work, we build an evolutionary language model, focusing on the cultural evolution aspect of language, without putting entirely aside the other perspectives. We introduce constraints in order to avoid building an unrealistic model, such as the lack of a centralized control mechanism (non-existent in a real social environment) or we avoid equipping our abstract language users with abilities beyond the current human brain limitations (e.g., individuals do not have access to each others internal representation of the world) (Steels, 2012).

3

Grammatical Gender in Dutch

(2) **De** **hond** achtervolgde **de** **kat**, maar **hij**
The-SG.C dog-SG.C chased the-SG.C cat-SG.C, but he-SG.M
kon **haar** niet vangen.
could her-SG.F not catch.

'The dog chased the cat, but he could not catch her.'

Sentence 2 presents a simple case of agreement between pronouns and their preceding nouns, in Dutch. However, one can immediately notice that the gender for 'hond' and 'kat' is *common*, clearly reflected in the accompanying determiners 'de', while the chosen pronouns are later 'hij' ('he'), thus *masculine*, for 'hond' and 'haar' ('her'), thus *feminine*, for 'kat'. The language speaker must dig further into his knowledge in order to acquire this information and make the correct choices.

We introduce here an interesting case study in Dutch, determined, as exemplified above, by the decision the speaker has to make in order to select the correct gender for the pronominal reference, given the sparseness or total lack of gender markers in the current state of the system.

3.1 Grammatical Gender

The nominal layer of a language comprises all the elements of the world, the concrete objects, the abstract entities, the animated beings, or concepts like feelings, beliefs, states. Gender comes as a reflection of how the speakers of a language visualize and characterize this layer, as it acts as a catalyst towards building a categorization of all the elements, based on the speakers' most salient dimensions (animacy, human, non-human, tangibility, etc.) (Janse et al., 2011).

In linguistic terms, *gender* is a grammatical category that divides the nouns in a language into mutually exclusive classes, based on different criteria. It varies greatly from one language to another, ranging from a classical two (masculine - feminine) or three-way (masculine - feminine - neuter) system, to twenty noun classes (in the Kordofanian language family), or in some cases even lacking completely (Corbett, 2006b).

We can now start asking what are the criteria employed in order to assign the gender label for the nouns of a language. The allocation rules are gathered under the term of gender assignment systems. We explain next the approach of Corbett (2006a) for describing and categorizing these systems.

Gender assignment systems are models that reflect the mechanisms underlying the gender distribution inside a language. We can distinguish two main categories, one derived from the meaning of the nouns, the other from the form.

1. Semantic Gender Systems

   The gender assignment for these types of systems is done based on the semantic features exhibited by the nouns. Allocation principles can vary, but some of the recurring dimensions are biological gender (female - male), animacy (animate - inanimate), tangibility (abstract - concrete). Depending on whether there are violations of the assignment rules, these types of models can further be split in strict semantic systems (the principles cover the entire vocabulary) or predominantly semantic systems (containing exceptions from the allocation rules).

   Due to the fact that humans can infer without difficulties the semantic dimensions of words, semantic systems should generally be stable and easily extendible for new word acquisitions inside a language, as long as the assignment rules are still known to the speakers.

2. Formal Gender Systems

   These categories of assignment systems are not standalone, they are complementary in case the semantic assignment strategy does not cover the entire vocabulary and additional rules are necessary. They appear at two

levels, phonological (usually based on the final sound unit of the word) or morphological (for example based on the inflectional classes of nouns, like in Russian).

Currently, a bit over half of the world's languages have a mixture of semantic and formal assignment systems, while the rest have just a semantic assignment system (Corbett, 2006a, p. 130).

3.2 The Origins of Gender

Regarding the arguments for the appearance of gender as a nominal category, we mainly note two theories, as presented by Kraaikamp (2012).

An initial motivation was thought to be represented by the reflection of biological sex in the language (Grimm, 1831). This initial classification was then extended to the rest of the inanimate objects, depending on how they were seen by the speakers in the context of their evolving society and culture.

On the other hand, Brugmann (1897) was arguing that this natural gender categorization came only after an initial phonetic-semantic mapping. For example certain endings could show that the denoted element was 'collective' and 'abstract'. Only after this part of the semantic space was attributed to femininity were the female entities attracted to this category. When additional phonetic evidence came into reach, it could be shown that this theory stands on firmer ground. Lehmann (1958) consolidated the idea, showing that Proto-Indo-European nouns followed a pattern in terms of phonetic ending and meaning correlation. He argues that there were three main suffixes, each corresponding to a certain category (individual, collective and mass), later mapped to the classes we know today (masculine, feminine, neuter). Figure 3.1 depicts the resulting meaning-gender correspondence that later became the basis for many gender systems in Indo-European languages.

3.3 Gender Evolution in Dutch

Generally, it is thought, looking at the current gender distribution in Dutch nouns, that the assignment system is an arbitrary one. However, Kraaikamp (2012) argues that, like many Indo-European languages, Dutch gender also draws its roots from a semantically based assignation, but was then subjected to disruptions that produce deviations from the original categorization. A main cause could be attributed to gender moving from a label denoting the placement of the word in the semantic space (as explained in the previous section), to a fixed feature of the noun (process called 'lexicalization'), as meanings are known to unavoidably evolve (widen or

The semantic basis of Indo-European gender

count/individual	collective/abstract	mass
male individual	female individual	sexless
↓	↓	↓
'masculine gender'	'feminine gender'	'neuter gender'

Figure 3.1: Mapping between semantic categories and gender in Indo-European languages (Kraaikamp, 2012), showing the original meaning for each of these nominal categories.

narrow) or shift, leading to discrepancies with the original semantic based mapping.

Currently, Dutch has a two-way nominal gender system, 'common' (for feminine and masculine) and 'neuter', denoted by the respective definite articles 'de' and 'het'. However, this was not always the case, as Middle Dutch was equipped with a three-way system: masculine, feminine and neuter, marked on the noun itself, and also on any accompanying determiner or adjective. Erosion of the gender markers have determined the feminine and masculine categories to collapse under one banner, as there were no longer any indications present for distinguishing them (De Vos, 2013).

Despite the reorganization of nominal gender in the standard Dutch, one can notice a difference between northern and southern Dutch language varieties. While northern Dutch complied to the changes, in southern Dutch regions (Belgium, but also North Brabant and Limburg from Netherlands), marking of the nominal syntactic gender is still present to some degree in determiners and adjectives, offering the speakers the ability to still distinguish the three-way system. On the other hand, pronominal gender continues to function under the three initial gender categories for both language regions.

3.3.1 Pronominal Reference in Dutch

Agreement can be defined as a correspondence between words in terms of their syntactic features (number, person, gender, case, etc.). A noun usually takes the role of the controller, imposing an agreement necessity with surrounding syntactic categories, such as articles, adjectives, or pronouns. For example in the sentence 'The boy kicks the ball.' the verb takes the suffix '-s' in order to agree with the singular 'boy'.

For the current work we are mainly interested in gender agreement and what

strategies are employed in order to achieve accordance between a noun and a pronoun used to refer to it subsequently, thus the pronominal gender.

In the case of an anaphoric reference, when a pronoun is used to refer to a previously mentioned noun, there are two possibilities for the gender agreement: syntactic (or lexical) agreement, where the grammatical gender of the noun is used, or semantic agreement, when the gender is chosen based on the properties of the referent. Usually, in the case of animate elements, the agreement is done semantically, using the biological gender of the referent, while for inanimate, elements syntactic agreement applies (Siemund, 2008). For Dutch, the default strategy is the syntactic agreement, leaving the semantic one just for cases in which the biological gender of the referent is known (Figure 3.2). For example, when talking about a specific professor, the pronoun used would be 'hij' ('he') or 'zij' ('she'), depending on the natural gender of the person, while the grammatical gender of the noun 'professor' is masculine.

Lexical gender agreement (inanimate referents):
common gender noun → masculine personal pronoun
neuter gender noun → neuter personal pronoun

Semantic gender agreement (animate referents):
male → masculine personal pronoun
female → feminine personal pronoun

Figure 3.2: Lexical (syntactic) and semantic agreement in Spoken Dutch (Kraaikamp, 2012). As one can notice, the employment of the feminine pronoun is only done in case of semantic agreement for animate referents.

3.3.2 Resemanticization

Because there is no longer a clear marking of the nominal gender in northern Dutch, leading to an overall reduced lexical gender system (Kraaikamp, 2012), anaphoric pronominal reference can raise discrepancies between the gender of the chosen pronoun and the one of the noun, creating a competition between syntactic and semantic agreement strategies. Dutch corpus studies revealed that the language is undergoing the process of *resemanticization*, through which the opaque syntactic agreement system is being replace with a semantic-based one (Audring, 2006; De Vos, 2013). Audring (2009) thoroughly presents this phenomenon and introduces an instance of the *individuation hierarchy* (Audring, 2006) to which the new system seems to be mapping (Figure 3.3). Again, for animate referents (considered to have the highest individuation), the biological gender comes into action. Moving towards the middle of the hierarchy, we go into masculine gender

territory, reserved for bounded, tangible objects, while the end of the scale (the unspecific masses, abstracts) triggers the neuter gender.

Human	>	Other animate	>	Bounded object/ Abstract	>	Specific mass	>	Unspecific mass, Unbounded abstract
		feminine/ masculine		*masculine*				*neuter*

Figure 3.3: Individuation Hierarchy presented by Audring (2006) in order to show how the new pronominal gender allocation is done in Dutch, after resemanticization, replacing the previous syntactic agreement system between nouns and their pronominal referents: biological gender (masculine or feminine) for animate referents, masculine for bounded, tangible objects and neuter for abstract and masses.

3.3.3 Gender in Southern Dutch

After presenting the current situation of nominal and pronominal genders in northern Dutch, we turn our attention to southern Dutch, especially towards Flemish dialects. As we mentioned before, in these regions a more rich morphological marking of the gender is available, making syntactic agreement still a strong competitor. Examples for each gender case are presented in Figure 3.4.

masculine:	Waar	is	de	stoel?	**Hij**	staat	buiten.
	where	is.3sg	the	chair	he	stand.3sg	outside
feminine:	Waar	is	de	tafel?	**Ze**	staat	buiten.
	where	is.3sg	the	table	she	stand.3sg	outside
neuter:	Waar	is	het	rek?	**Het**	staat	buiten.
	where	is.3sg	the	rack	it	stand.3sg	outside

'Where is the chair/table/rack? It is standing outside.'

Figure 3.4: A sample of pronominal reference in Belgian Dutch (De Vos, 2013), showing how syntactic agreement occurs between the preceding noun and the referencing pronoun.

Figure 3.5 presents a general map of the main Flemish dialects, i.e., West Flemish (region A), East Flemish (region B), Brabantic (region C), and Limburgian (region D). We note that the Brabantic dialect is considered to be the most stable but also the most influential among them all. In between we can notice the transition zones (I, II and III). During a preliminary study of the West and East

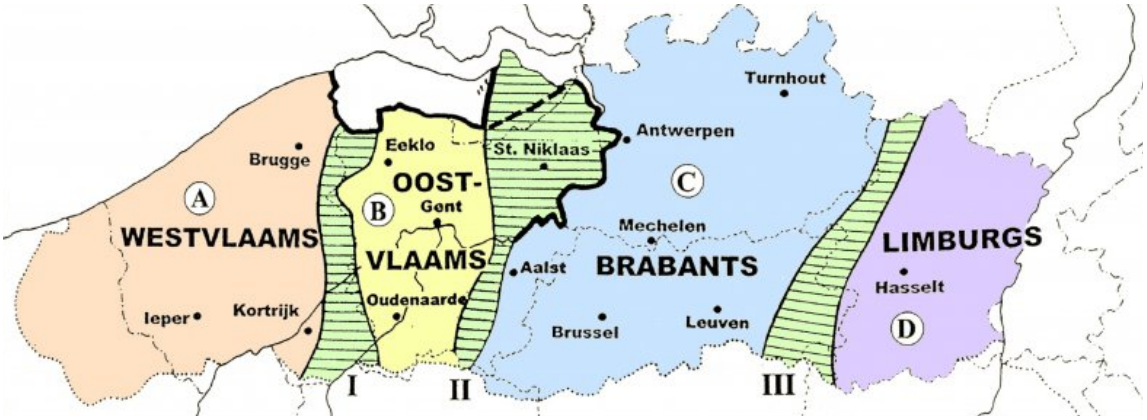

Figure 3.5: Map presenting the main Flemish Dialects (Taeldeman, 2001)

Flemish dialects, De Vogelaer (2009) has noted that indeed the three-way gender system is still preserved in the corresponding regions (A and B), however there seems to be a lot of variation in term of individual words' gender. As an example, we extracted the gender map for the word 'sneeuw' ('snow'), presented in Figure 3.6. In Standard Dutch the noun is feminine, as well as in the majority of West and East Flemish dialects (marked by the prominent grey square). In a small part of East Flemish dialect region the noun is masculine (the area is marked by a grey circle and it is delimited from the previous one by the black dotted line). We notice on both sides deviation from the standard gender, especially from the feminine one (in West Flemish), replaced by masculine and neuter referents. We can now look further for possible motives that cause these deviations.

A deeper study of Belgian Dutch dialects reveals that changes in the pronominal gender, discussed previously, are also taking place here, but with additional triggers and motivation. De Vogelaer (2009) describes three main factors that are speculated to cause the commotion within the Flemish language region: firstly, one must consider the impact of Standard Dutch (that is taught in all the academic institutions), secondly, the inter-dialectal influences, especially the susceptibility to converge towards the central Brabantic dialect and lastly, there seems to also be a parallel tendency to shift towards the semantic system (resemanticization). The first two elements are attributed to diffusion (i.e., the transfer of features between adjacent language related communities, belonging to related branches of a language family tree), while the last one is considered a case of imperfect transmission (i.e., incomplete or faulty language acquisition between different generations of native speakers)(Labov, 2007).

In this new light we can now return to our small case study for the word 'sneeuw', as presented by De Vogelaer (2009). In the context of standardization,

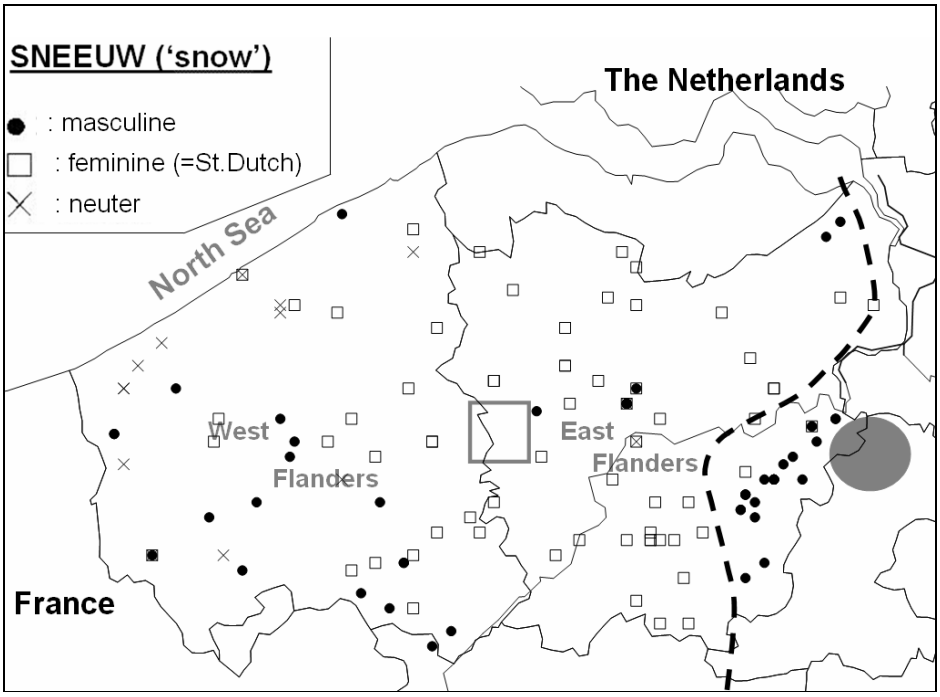

Figure 3.6: Geographical gender distribution for a pronoun referring to the word 'snow' in West and East Flemish dialect regions (De Vogelaer, 2009). We can notice that there is a lack of uniformity for the gender of a pronominal referent for 'sneeuw', even though in standard Dutch this noun is feminine.

we would expect an increase towards the standard gender of the noun, feminine, which is not the case. Furthermore, the general tendency of southern Dutch is to diminish the use of feminine when the referent is not animate, lacking the biological gender. The Brabantic influence, meaning a deviation towards masculine, is noticed in East Flemish, at the border between the regions. However it would be hard to attribute this cause to the perturbation we notice in the isolated far west corner of the map, for West Flemish, as the border with Brabant lies further away, with a block of feminine gender usage lying in between. Lastly, a plausible reason for the deviation towards neuter could be an incipient resemanticization, as 'snow' is a mass denoting noun. Additionally, as a parallel process, a weakening of the lexical system, resulting in a tendency of merging the masculine and feminine could explain the takeover of masculine in this case for the West Flemish region (De Vos & De Vogelaer, 2011).

The current work draws inspiration from the third factor presented before, resemanticization in Southern Dutch, as this is an ongoing process, clearly lagging behind Northern Dutch, thus making it a very interesting study case in the context

of evolutionary linguistics. As already mentioned, a main cause for this phenomenon is considered the imperfect transmission of the syntactic gender knowledge between consecutive generations of native speakers. We try to build a simulation based on these principles and discover possible mechanisms in order to repair the mismatches and create a stable system.

4
Dutch Gender Language Game

Language games, along with Fluid Construction Grammar (FCG), have become the main tools for the multi-agent based approach for studying historical and evolutionary linguistics (Steels & Vogt, 1997; Steels, 2001; Loreto & Steels, 2007). This work adheres to the increasing interest in grammatical agreement studies (Beuls & Höfer, 2011; Beuls, 2011; Beuls & Steels, 2013). As described in the previous chapter, we focus here on the pronominal agreement in Southern Dutch, namely on the shift from syntactic agreement to semantic agreement.

We present in the following sections our language game, an *anaphoric reference game*, that uses pronouns (marked by one of the three genders, masculine, feminine or neuter) to refer back to a heard noun, thus simulating a simplified case of pronominal agreement. Our entire model is implemented in Common Lisp, within the Babel framework (Steels & Loetzsch, 2010; Loetzsch et al., 2008), that can be retrieved and installed from `https://ai.vub.ac.be/trac/babel2`.

4.1 Description and Purpose

We devise a new language game, in order to allow the agents to compare their knowledge about the gender of the nouns in their vocabularies, in the case of Dutch. We simulate an anaphoric reference, such that, after hearing an utterance that consists of a noun, an agent has to choose one of the personal pronouns in the

oblique case, presented in Table 4.1: 'hem' ('him') for masculine, 'haar' ('her') for feminine and 'het' ('it') for neuter.

Personal pronoun	**Masculine**	**Feminine**	**Neuter**
Nominative	hij	zij	het
Oblique	hem	haar	het

Table 4.1: The personal pronoun in Dutch, for nominative and oblique cases.

Figure 4.1 presents an overview of an anaphoric reference language game. Each agent has an internal representation of the world. At the start of each interaction two agents are chosen randomly from the population. We consider the speaker to be the guide of the conversation, holding the role of a tutor we could say, such that, in case of disagreement over the gender of the noun, we take his option to be the correct one. Thus, the speaker chooses a topic, conceptualizes the meaning and the gender of the chosen noun, then produces an utterance. The hearer parses the utterance then tries to interpret the gender of the noun, based again on its own model of the world.

In a daily social interaction, a successful interaction of this game can be seen as an abbreviation for the following dialogue:

(3) A: Waar is de **tafel**?
'A: Where is the **table**?'

B: Ik bracht **haar** naar buiten.
'B: I took **her** outside.'

In our case, we strip down speaker's A utterance to simply the noun, additionally removing the definite article (to remove any auxiliary gender indication) and also speaker's B response consists of just the oblique pronoun:

(4) A: **Tafel**?
'A: Table?'

B: **Haar**.
'B: Her.'

Sentence 5 presents a possible example for a failed interaction.

(5) A: Waar is de **tafel**?
'A: Where is the *table*?'

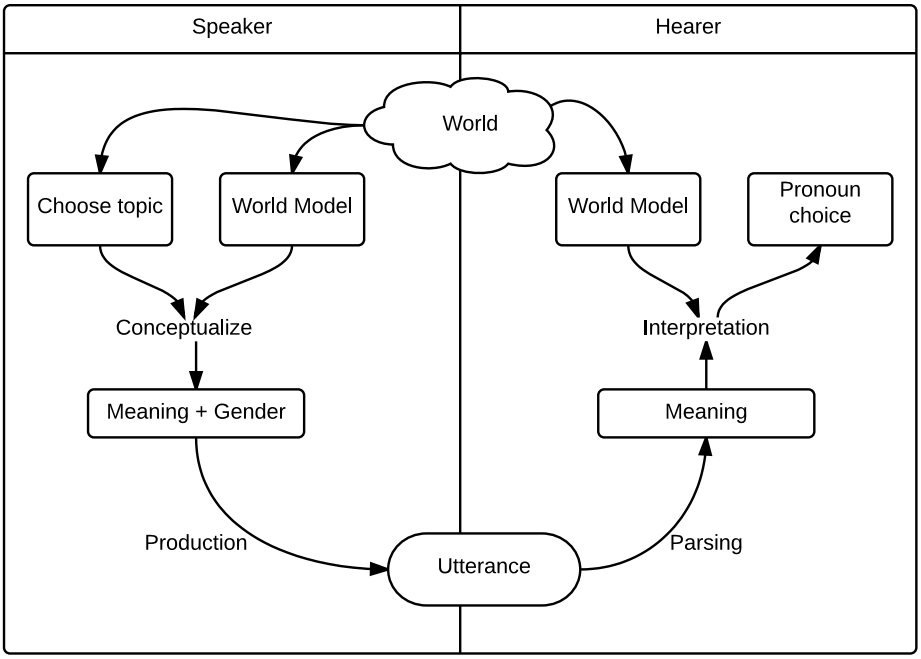

Figure 4.1: The anaphoric reference language game constructed in order to simulate the shift from an increasingly opaque syntactic gender system towards a semantic one, in Dutch.

B: Ik bracht **het** naar buiten.
'B: I took **it** outside.'

A: **Haar**, bedoel je.
'A: **Her**, you mean.'

Notice thus how person A (speaker) plays the role of the tutor and corrects person B (hearer). Thus after each interaction, feedback is received from the speaker and the hearer can then align his world model accordingly.

The agents start with a full knowledge over the syntactic gender of the nouns in their vocabulary. Through population turnover, we simulate *imperfect transmission*, such that after, a certain number of interaction, the agents, having lost the gender knowledge about the noun, are forced to apply different strategies in order to produce the utterances. Our goal is thus to test various mapping mechanisms and observe the outcome of a sequence of language games, in terms of distribution over a defined semantic space, simulating a shift from syntactic gender agreement towards a semantic system.

We present next every aspect of our evolutionary language model, introduce strategies and explain our system parameters.

4.2 Vocabulary

4.2.1 Dutch Nouns

In order to build the vocabulary of our agents we used constructions under the FCG framework. Before looking at our lexical constructions, we first present the choice of nouns, retrieved from the work of De Vos (2008). The nouns are grouped according to multiple semantic dimensions, that were chosen to compose our semantic space: tangibility (abstract – concrete), countability (countable – uncountable), animacy (animate – inanimate) and lastly ontological category (not specified – animal – collective).

	Concrete - Countable - Inanimate				
Masculine	*Dutch*	sleutel	brief	bloesem	olm
	English	key	letter	blossom	elm
Feminine	*Dutch*	bank	fles	klier	vod
	English	bench	bottle	gland	rag
Neuter	*Dutch*	dak	schip	kadaver	sas
	English	roof	ship	corpse	flood gate

Table 4.2: Concrete, countable, inanimate nouns

	Concrete - Uncountable - Inanimate				
Masculine	*Dutch*	steen	rijst	kaviaar	blubber
	English	stone	rice	caviar	blubber
Feminine	*Dutch*	sla	melk	zeep	henna
	English	lettuce	milk	soap	henna
Neuter	*Dutch*	hout	vlees	koper	beleg
	English	wood	meat	copper	filling

Table 4.3: Concrete, uncountable, inanimate nouns

There are in total 72 nouns, thus at the start, each agent has a vocabulary size of 72, composed by the above noun along with their syntactic gender and semantic properties. We remark here that we have an equal distribution of the three genders (masculine, feminine and neuter) over our semantic space. During our experiments we track the evolution of individual nouns, in terms of gender, but also the gender distribution over the semantic dimensions along with the gender distribution over the whole population.

	Abstract - Countable - Inanimate				
Masculine	*Dutch*	vorm	dag	bluts	cultus
	English	form	day	dent	cult
Feminine	*Dutch*	kleur	ruzie	vondst	gave
	English	color	quarrel	discovery	gift
Neuter	*Dutch*	cijfer	weekend	atoom	taboe
	English	digit	weekend	atom	taboo

Table 4.4: Abstract, countable, inanimate nouns

	Abstract - Uncountable - Inanimate				
Masculine	*Dutch*	adem	hulp	hinder	advent
	English	breath	help	nuisance	advent
Feminine	*Dutch*	rust	kou	griep	do
	English	rest	cold	flu	do
Neuter	*Dutch*	weer	geluk	proza	applaus
	English	weather	happiness	prose	applause

Table 4.5: Abstract, uncountable, inanimate nouns

	Concrete - Countable - Animate - Animal				
Masculine	*Dutch*	hond	vis	panter	gorilla
	English	dog	fish	panther	gorilla
Feminine	*Dutch*	kat	koe	zwaluw	krokodil
	English	cat	cow	swallow	crocodile
Neuter	*Dutch*	paard	schaap	hert	kuiken
	English	horse	sheep	deer	chicken

Table 4.6: Concrete, countable, animate, animal nouns

	Abstract - Countable - Inanimate - Collective				
Masculine	*Dutch*	groep	raad	senaat	vijfling
	English	group	board	senate	quintuplets
Feminine	*Dutch*	brandweer	politie	kliek	kudde
	English	firedepartment	police	gang	herd
Neuter	*Dutch*	leger	gezin	korps	comite
	English	army	family	corps	committee

Table 4.7: Abstract, countable, inanimate, collective nouns

4.2.2 FCG Constructions

As mentioned before, we implement the vocabulary as a construction inventory within the FCG framework, for each of our agents.

To exemplify our method, we present next the code used to generate the construction for the word 'hond' ('dog'):

```
(def-lex-cxn dog-lex
  (def-my-lex-skeleton dog-lex
      :meaning (== (dog ?ref))
      :args (?ref)
      :string "hond"
      :sem-cat ((sem-class identifier)
                (obj-type concrete)
                (ontological-cat animal)
                (animacy +)
                (countability +))
      :syn-cat ((lex-class noun)
                (gender m))))
```

Thus the construction string is 'hond', the lexical class is noun, gender masculine, while the semantic categories are: object type – concrete, countability – countable, animacy – animated and ontological category – animal. A better visualization for the construction can be extracted from the available web interface and it is presented in Figure 4.2.

When the agent has in his world model multiple competitors for the gender of a noun, this is reflected in his inventory by the presence of a construction for each of the considered genders, each having the same string and meaning for the noun. We have to also mention that each construction has a score that reflects the agents confidence for that word-gender association. In case the agent does not know the gender of a construction, this is shown in his inventory by a variable for the gender feature: `(gender ?gender)`.

4.3 Population Turnover

The first mechanism we present is the population turnover, which introduces a vertical transmission in our system, in order to simulate the imperfect transmission phenomenon discussed earlier in Section 3.3.3. There are three parameters that regulate this mechanism:

Turnover fraction determines what percentage of the population will be replaced.

We note here that our agents do have an age and, during each population

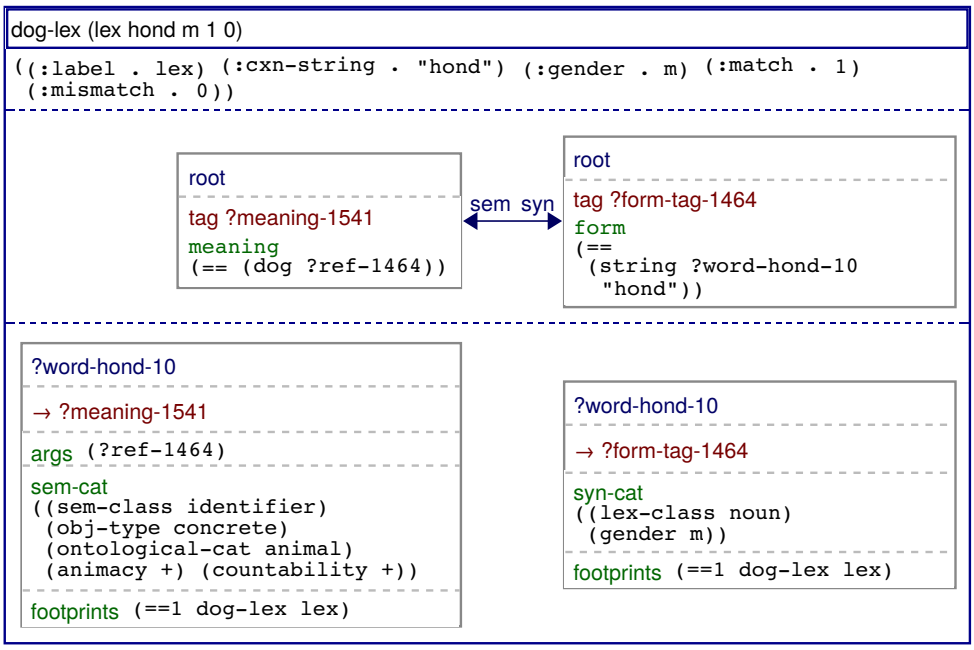

Figure 4.2: The lexical construction for the word 'hond' ('dog'). We can notice the semantic properties of the noun on the left side of the construction, and the gender on the right side.

turnover, the oldest agents are always replaced first. In the place of each deprecated agent, we create an offspring which inherits his vocabulary from his parent.

Transmission rate sets the fraction of the parent-agent vocabulary that is transmitted to a child. For example, if the transmission rate is 0.5, 50% of the vocabulary is stripped of gender knowledge, while for the rest the new agent inherits, in case of multiple competitors, the highest scored construction, or simply the only construction available for the word. Additionally, the new agent enters the system having all the construction scores initialized (with a value associated to the used scoring strategy, discussed later in Section 4.6).

Turnover step defines once in how many interactions the population turnover should take place.

4.4 Learning Operators

After a population turnover, it is then possible that a speaker encounters a word he wants to utter, about which he has no gender knowledge. In this case, the

meta-level learning comes into action, diagnoses the unknown gender problem and, using the selected repair strategy, maps a new gender for the word.

For our language model, we created three possible learning strategies for selecting a gender for a word: random, history and best representative.

The random strategy is straightforward: when a word with unknown gender is encountered, it simply chooses randomly with equal probabilities among the three possible variants: masculine, feminine or neuter.

4.4.1 Construction Similarity

Before getting into more details regarding the other two learning strategies, we need to explain how we measure the similarity between two constructions. Because we are interested in a mapping over the semantic space, we define for each construction a numeric feature vector with the following format: [object type, countability, animacy, ontological category].

In order to measure the similarity between two feature vectors, X and Y, we use the cosine similarity, obtaining a score between -1 and 1:

$$similarity_score = \frac{X \cdot Y}{||X|| \cdot ||Y||} = \frac{\sum_{i=1}^{n} X_i \times Y_i}{\sqrt{\sum_{i=1}^{n} (X_i)^2} \times \sqrt{\sum_{i=1}^{n} (Y_i)^2}} \tag{4.1}$$

We map each semantic feature to a numeric value as follows:

		Value
Object type	abstract	0
	concrete	1
Countability	uncountable	0
	countable	1
Animacy	inanimate	0
	animate	1

Table 4.8: Feature vector mapping of the semantic categories

As for the ontological category, we have three values: not specified, collective and animal. Although we could simply associate the values 0, 1 and 2, we need to take into consideration the following issue. Lets consider two nouns with the common semantic properties: concrete, countable, inanimate, while the ontological category for one is collective and for the other one is animal. Their respective feature vectors are then $[1, 1, 0, 1]$ and $[1, 1, 0, 2]$. When we compare these nouns

against a concrete, uncountable, inanimate, with a non specified ontological class, thus $[1, 0, 0, 0]$ we obtain the following scores: 0.42 and 0.59. Looking at the vectors, we notice that they both differ in terms of two values and as we consider all semantic features to have the same weight, this result is not in accordance with our expectation. In order to solve this issue we extend the numerical mapping of the ontological category to a three dimensional vector as follows:

		Value
Ontological category	Not specified	1 0 0
	Collective	0 1 0
	Animal	0 0 1

Table 4.9: Feature vector mapping of the ontological category

4.4.2 History Strategy

For the history gender mapping strategy, we offer each agent the ability to memorize a number of previous encountered constructions. This memory size is configurable through a parameter.

Whenever the agent encounters a word for which he does not know its gender, the history strategy computes the pairwise similarity between the respective construction and each of the constructions stored in the memory. The next step is to aggregate the scores, as more than one construction in the memory can have the same gender, by simply averaging the scores belonging to the same gender. The final step is to pick the gender which corresponded to the highest similarity score. In case of score equality, the gender is randomly picked among the tying candidates.

4.4.3 Best Representative Strategy

The best representative strategy requires the agent to memorize, for each gender, the construction with the highest score encountered so far. Thus each agent, keeps in his internal representation, a candidate which best models the semantic space – gender correspondence, for each case.

When faced with the unknown gender problem, the agent can then compute the similarity between the encountered word and each of the gender representatives from his memory. He can then select the gender corresponding to the highest score.

4.5 Alignment

So far we have described all the operations and mechanisms available for the speaker to be able to arrive at the production step of the language game. Next, it is the hearer's turn to parse the utterance and try to refer to it using a gender-marked pronoun. In case the selected pronoun, thus gender, does not correspond to the speaker's, the hearer will adopt the speakers solution, if he did not have it already in its vocabulary. This gives rise to competition among genders for a word.

4.5.1 Scoring Strategies

Additionally, the final step of the alignment, at the end of an interaction, also involves updating the scores of the constructions. We present next three types of strategies for score alignment:

Frequency – introduced by Steels & McIntyre (1998), this approach implies updating a counter, for the hearer only, each time a construction is heard (the construction equivalent to the speaker's solution). The scores, in this case, are integer values between 1 and ∞.

Success rate – implies keeping a counter for the matches and mismatches, for each construction. During alignment (done only for the hearer), the speaker's solution gets an increase for the match count, while all the other competitors have their mismatch counter increased. The success rate is then defined as:

$$success\_rate = \frac{match}{match + mismatch} \tag{4.2}$$

The scores will thus have values between 0 and 1.

Lateral Inhibition – was introduced by Steels & Kaplan (1999) and it is already a spread strategy in the field of language games. It implies the idea of inhibiting (decreasing the scores for) the competitors, in the case of a successful interaction, while entrenching (increasing the score for) the correct solution, for a hearer. For a failed interaction, the wrongly chosen solution gets its score decreased. For our model we introduce two learning rates: δ_+ (used in successful interaction updates) and δ_- (used in failed interactions updates) and the scores are altered using the interpolated lateral inhibition method (Wellens, 2010), keeping them between 0 and 1, as follows:

- Successful interaction - entrenching solution

$$new\_score = \delta_+ + old\_score(1 - \delta_+) \tag{4.3}$$

- Successful interaction - inhibiting competitors

$$new\_score = old\_score(1 - \delta_+) \tag{4.4}$$

- Failed interaction - inhibiting the wrong choice

$$new\_score = old\_score(1 - \delta_-) \tag{4.5}$$

As De Vylder (2007)'s study of conventions in multi-agent systems revealed that updating speakers, thus reinforcing their own preferences, leads to less adaptable agents, in our model all the scoring strategies are only applied to the hearer.

4.6 Solving Competition

The agents, both the speaker and hearer, can face the issue of making a choice between multiple gender competitors for a word. In order to deal with this situation, we have introduced three possible strategies: random, probability, and highest score.

Again the random strategy is straightforward, by choosing randomly with equal probabilities between the competitors.

Probability – this strategy implies making a probabilistic choice among competitors, based on their scores. We normalize the scores using the softmax function:

$$P(C_k) = \frac{exp(\frac{score(C_k)}{\tau})}{\sum_{i=1}^{n} exp(\frac{score(C_i)}{\tau})} \tag{4.6}$$

The parameter τ is the temperature and it controls the uniformity of the mapping between 0 and 1 of the scores.

Highest score – this strategy simply chooses the competitor with the highest score. This approach is the one usually adopted in most language game models that imply competition (Steels, 2001; Lenaerts et al., 2005; Steels & Loetzsch, 2012; Beuls & Steels, 2013).

4.7 Measurements

Next we introduce a few measures that will be used to evaluate all our strategies, as well as the final outcome of our language model.

4.7.1 Communicative Success

Communicative success is used in order to determine if an interaction was successful or not, meaning if the agents could communicate properly. This can happen in two cases: either the speaker and the hearer have the same preferred gender for a word, or when the hearer has the speaker's solution in its inventory, but he did not choose it as its preference. After signalling the interaction's outcome, the alignment procedure can start for the involved agents.

4.7.2 Alignment Success

Alignment success measures the degree to which the agents have the same preferences over the genders for the elements of the world. This means that alignment success is reached at the end of an interaction only if the agents have the same solution.

4.7.3 Gender Distribution

Gender distribution is measured at two levels, always over the entire population. Firstly, we want to see the gender distribution over all the elements of the world. Secondly, we are also interested in the gender distribution over the semantic space. Again we remark here that the initial state of the system has equal gender distribution in both cases.

4.7.4 Gender Competition per Word

Lastly, we monitor the gender evolution for each of the elements in the world. We want to observe the gender competition taking place at word level, what was the initial syntactic gender and which is the final gender upon which the population reached an agreement.

5

Anaphoric Reference Language Game Experiments

This chapter presents a step by step analysis of our evolutionary language game model of the shift towards semantic agreement for pronominal in Dutch. We choose to build the study of the impact of our strategies and parameters sequentially, as it would be difficult to evaluate the dynamics of the system as a whole directly. We present our experimental settings and results for each study case, trying to observe and discuss impact and influences of the strategies along the way.

Beforehand, we present the initial state of the system, with a population having full knowledge of the syntactic nominal gender. The communication and alignment success are always true, while the gender distribution over the agents' vocabularies, as well as over the semantic space are shown in Figures 5.1 and 5.2.

We notice that we start from a balanced system, in which all genders are equally represented over the population's vocabulary and the available semantic categories. Figure 5.2 presents the six categories of nouns we use: C-C-I (concrete countable inanimate), C-U-I (concrete uncountable inanimate), A-C-I (abstract countable inanimate), A-U-I (abstract uncountable inanimate), C-C-A-A (concrete countable animate animals) and A-C-I-C (abstract countable inanimate collective), categories which will always be compared across all experiments.

Additionally, we will present resulting gender distributions over an aggregated

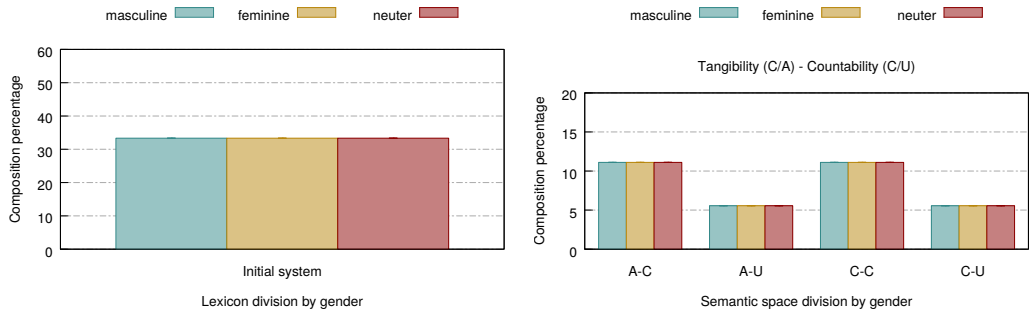

(a) Gender distribution over the population vocabulary

(b) Gender distribution over the tangibility–countability axes

Figure 5.1: Initial gender distribution

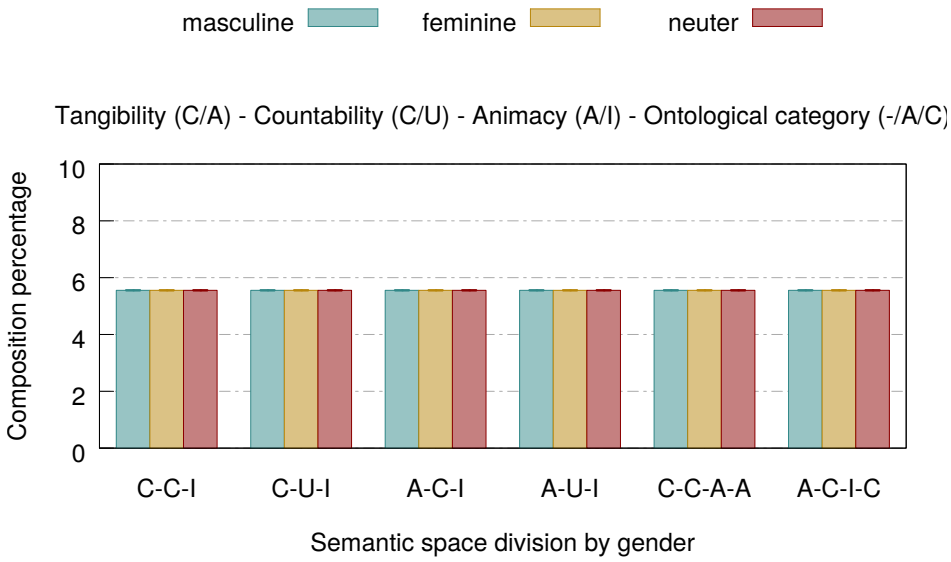

Figure 5.2: The initial gender distribution over the semantic space.

semantic space, for the tangibility – countability axes. All the plots for communicative and alignment success are averaged over the series and are accompanied by the standard deviation in the form of a semi-transparent area surrounding each line, thus indicating the amount of variation over the trials. As for the bar plots, the results are averaged over the last 1000 interactions for the last series of the language game, over which the system had already converged towards a stable gender distribution.

Lastly, we also list here the parameters that remain constant throughout our entire set of experiments: population size – 10, number of interactions – 2 $\times$

10^4, turnover step – 10^4, number of series – 5. Thus, there will be 2 population turnovers during each run, the first turnover takes place after interaction 100 (as we also want to represent in our plots the initial state of the system), the second one at interaction 10^4.

For the first experiments we are mainly interested in observing the impact of the tested strategies on the convergence of the population towards full communicative and alignment success, leaving the gender distribution and gender competition per word for last.

5.1 Experiment 1 – Scoring Strategies

We analyse here what is the impact of the various score alignment methods implemented for the system: frequency, success rate and lateral inhibition. As a starting point, we expect lateral inhibition to perform the best, as it has been shown in many previous language game models (Wellens, 2010).

5.1.1 Experimental Setting

We present below the list of parameters used for this experiment. We mention that for population turnover and transmission rate we simply picked middle values.

- Scoring strategy – frequency / success rate / lateral inhibition
- Competitor selection strategy – highest score
- Gender mapping strategy – best representative
- Population turnover fraction – 0.5
- Transmission rate – 0.5
- Lateral inhibition – δ_+ 0.2 and δ_- 0.1

5.1.2 Results and Discussion

We can observe from Figure 5.3 that all the strategies reach full communicative success, under the chosen learning mechanism (best representative).

A more interesting perspective is shown in Figure 5.4 for the alignment success, thus the degree to which the agents share the same preferences for genders.

Both lateral inhibition and success rate are performing well, with the interpolated lateral inhibition appearing to be faster in terms of reaching convergence. The success rate is a simplified, weaker version of lateral inhibition, if we consider

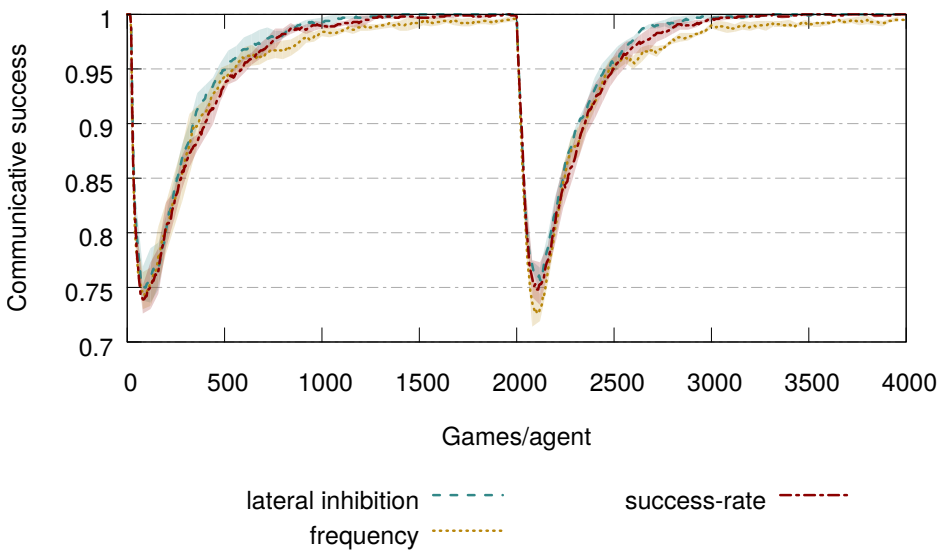

Figure 5.3: Communicative success of the scoring strategies

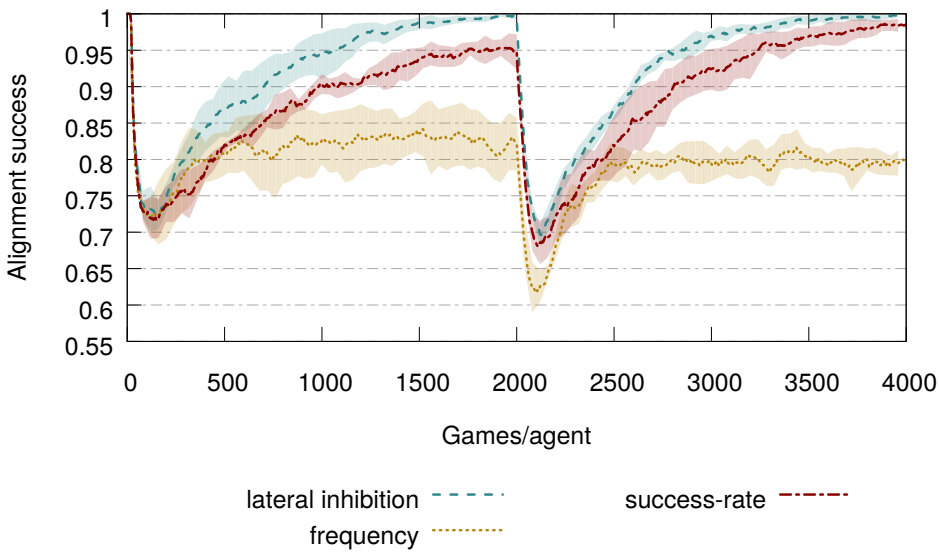

Figure 5.4: Alignment success of the scoring strategies

the match count to play the role of entrenchment and the mismatch the inhibitor for a construction. However, we see that it is slower in terms of building the same gender predilections for the agents, as probably the upgrades are more drastic at the start (for example 1 match - 0 mismatch yields a success rate of 1, while 1 match - 1 mismatch reduces the success rate all the way to 0.5) and slow down

considerably with the number of interactions (for example 58 match - 41 mismatch gives a success rate of 0.586, 58 match - 42 mismatch yields 0.58, thus just a small difference of 0.006). Thus, the success rate strategy loses flexibility, making it harder for the agents to switch preferences, with the increasing number of interactions.

Lastly, we observe that the frequency scoring approach does not manage to allow the agents to reach the same preferences in terms of gender for the nouns. Wellens (2010, p. 96) also encountered this issue and attributes it to a lack of flexibility, as the frequencies increase high enough, the agents become too robust to acquire emerging trends.

5.2 Experiment 2 – Solving Competition

In the previous chapter we explain that competition can arise for a word, between multiple gender possibilities. In order to deal with this problem, we implemented three strategies: random, probability and highest score. As mentioned before, the usual approach is the last one, while the probability strategy is our own contribution for this model.

Before starting the comparison of the three approaches, we need to also include a short study on the impact of the temperature parameter, for the probability strategy. We recall that the temperature was used in the softmax function in order to normalize the scores, before doing a probabilistic selection among them. In general, the lower the temperature is, the higher the difference between probabilities after the mapping, and vice versa. Thus we expect that for a high temperature this strategy becomes comparable to the random one, while for low temperatures it should tend towards the highest score approach.

The idea of introducing stochasticity in a language game is not new Steels & Kaplan (1998), as in real-life interactions, the utterance is not always transmitted without noise, or even the context can be perceived with faults (thus like most natural systems, we deal with inherent stochasticity). In general, in computational models, thus in the case of language games as well, stochasticity can be used to achieve innovation, variation or exploration.

5.2.1 Experimental Setting

- Competitor selection strategy – random / probability / highest score
- Temperature τ – 0.1 / 0.3 / 0.5 / 1
- Scoring strategy – lateral inhibition (δ_+ 0.2 and δ_- 0.1)

- Gender mapping strategy – best representative
- Population turnover fraction – 0.5
- Transmission rate – 0.5

5.2.2 Results and Discussion

Again we can see from Figure 5.5 that all the strategies reach full communicative success.

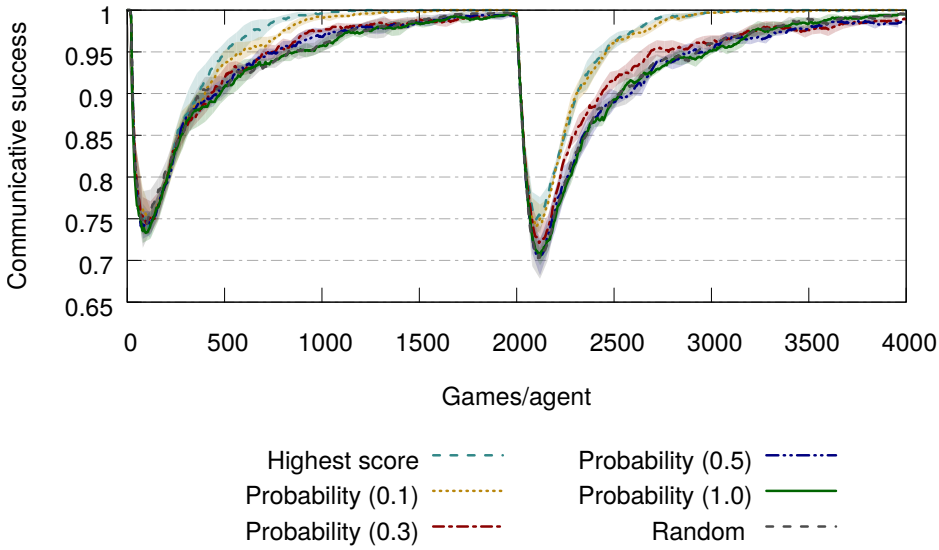

Figure 5.5: Communicative success of the competitor selection strategies

Turning now our attention to Figure 5.6, we see that the temperature imposes the expected behaviour on the probability strategy. Highest score converges the fastest, followed closely by the probability with τ 0.1. The rest of the strategies do not manage to reach full alignment success. Furthermore, we notice that the probability with τ 1 is already equivalent with the random strategy. We remark here that the temperature study was made in the context of lateral inhibition as a scoring strategy, thus all scores were always between 0 and 1. It might be different for other scoring strategies, thus additional tests should be made when choosing new strategy combinations.

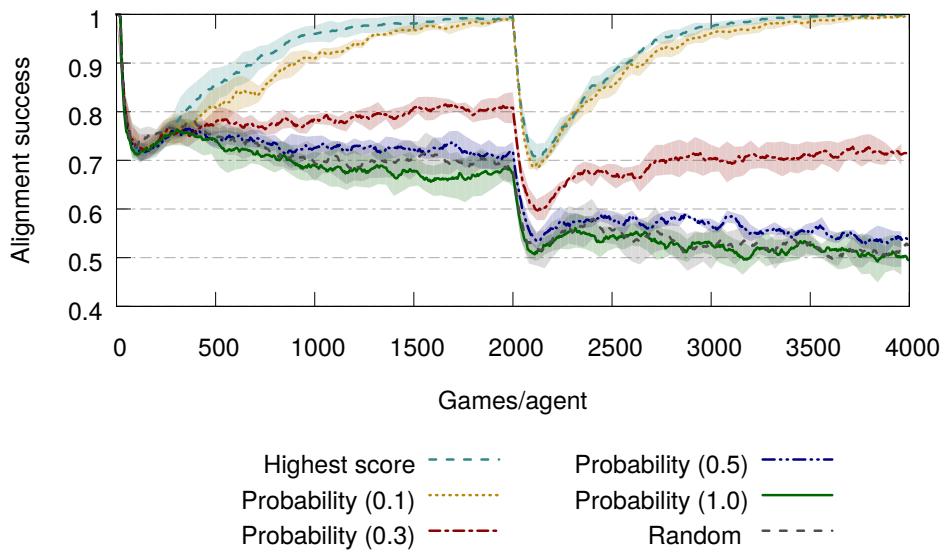

Figure 5.6: Alignment success of the competitor selection strategies

5.3 Experiment 3 – Population Turnover

For this experiment we investigate the impact of the population turnover fractions: 0.2, 0.5, 0.8 and 1.0. As for the transmission rate, we decided to take the value 0.0, in order to see how much drop in communicative and alignment success we obtain and, if and how fast the population can reach again full agreement. We look with great interest at the combination 1.0/0.0, as this implies that all the agents are replaced and no gender knowledge is transmitted.

5.3.1 Experimental Setting

- Population turnover fraction – 0.2 / 0.5 / 0.8 / 1.0
- Transmission rate – 0.5
- Gender mapping strategy – best representative
- Competitor selection strategy – highest score
- Scoring strategy – lateral inhibition (δ_+ 0.2 and δ_- 0.1)

5.3.2 Results and Discussion

Figures 5.7 and 5.8 present the communicative and alignment success. As we expected, the higher the turnover fraction is, the lower the success drops. Addition-

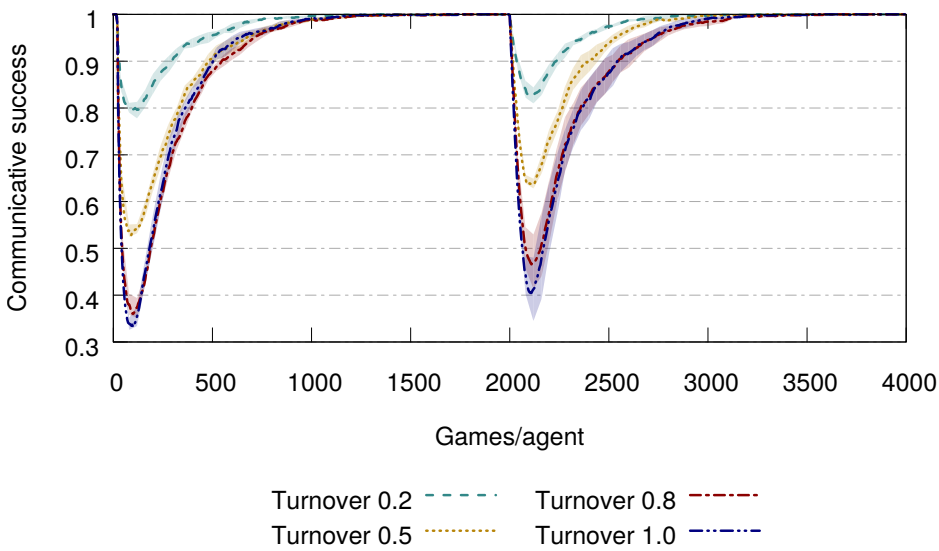

Figure 5.7: Communicative success for various population turnover fraction values, with transmission rate 0.5

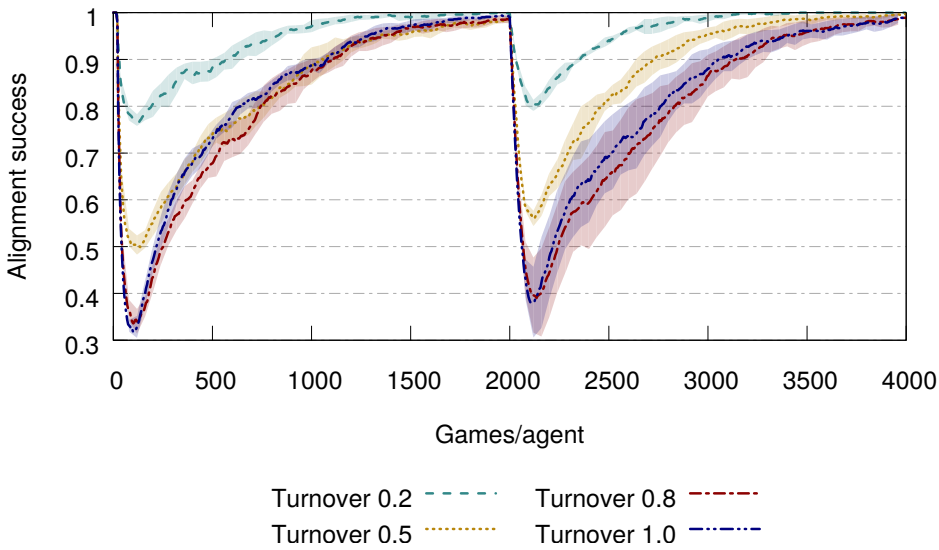

Figure 5.8: Alignment success for various population turnover fraction values, with transmission rate 0.5

ally, the agents managed to reach communicative success, even for the extreme case, with all agents being replaced and no knowledge transmitted.

5.4 Experiment 4 – Transmission Rate

For this small experiment, we analyse the impact of the transmission rate on the system's convergence. We take the values: 0.0, 0.2, 0.5 and 0.8. For the population turnover we choose the extreme value 1.0.

5.4.1 Experimental Setting

- Transmission rate – 0.0 / 0.2 / 0.5 / 0.8
- Population turnover fraction – 1.0
- Gender mapping strategy – best representative
- Competitor selection strategy – highest score
- Scoring strategy – lateral inhibition (δ_+ 0.2 and δ_- 0.1)

5.4.2 Results and Discussion

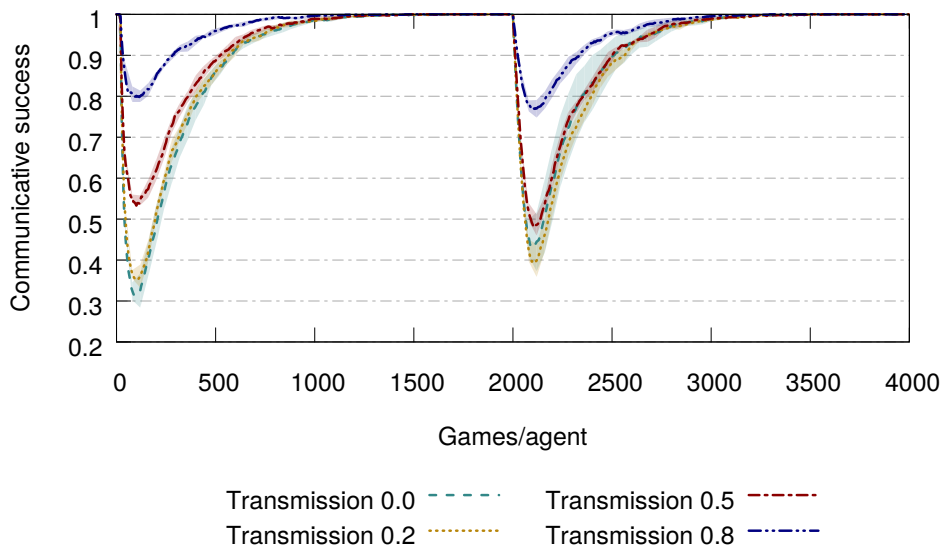

Figure 5.9: Communicative success for various transmission rates, with a population turnover fraction of 1.0

Figures 5.9 and 5.10 present the communicative and alignment success for the transmission rate variation case. We notice that the higher the transmission rate, the less decrease in success is registered. Also, after the second turnover, all the

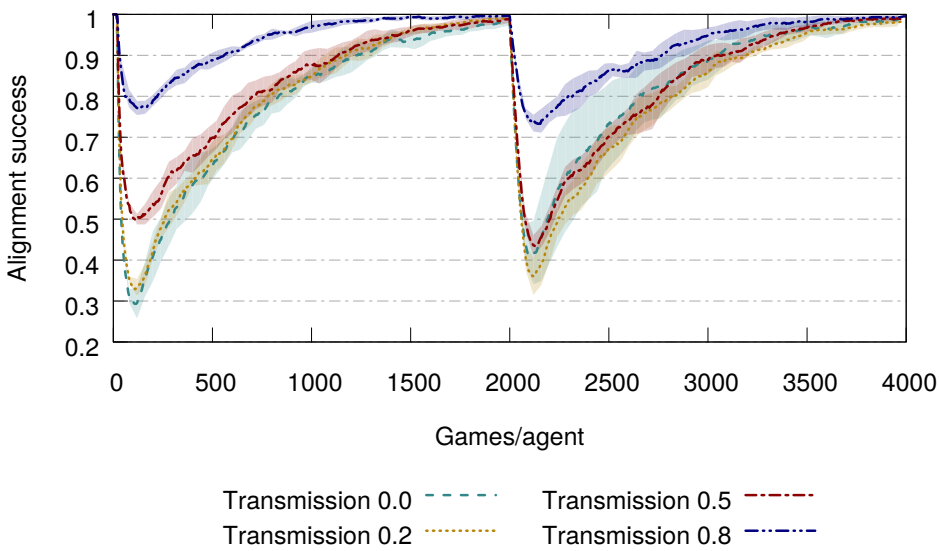

Figure 5.10: Alignment success for various transmission rates, with a population turnover fraction of 1.0

values below 0.8 seem to behave in the same manner for both communicative and alignment success.

5.5 Experiment 5 – Mapping Gender onto Semantic Space

The analysis of the learning strategies for gender mapping will also include a study for the impact of the memory size of an agent for the history strategy. The agent's memory, in this case, plays the role of a window that offers a limited view over the semantic space – gender correspondence from the agent's internal world representation.

5.5.1 Experimental Setting

- Gender mapping strategy – best representative / history / random
- History memory size – 5 / 50 / 100
- Competitor selection strategy – highest score
- Scoring strategy – lateral inhibition (δ_+ 0.2 and δ_- 0.1)

- Population turnover fraction – 1.0
- Transmission rate – 0.0

We decided to take the extreme values for the population turnover and transmission rate, in order to eliminate any inherited mapping between gender and the semantic space from the initial syntactic system, such that we observe how the gender mapping is done from scratch by the learning strategies strategy.

5.5.2 Results and Discussion

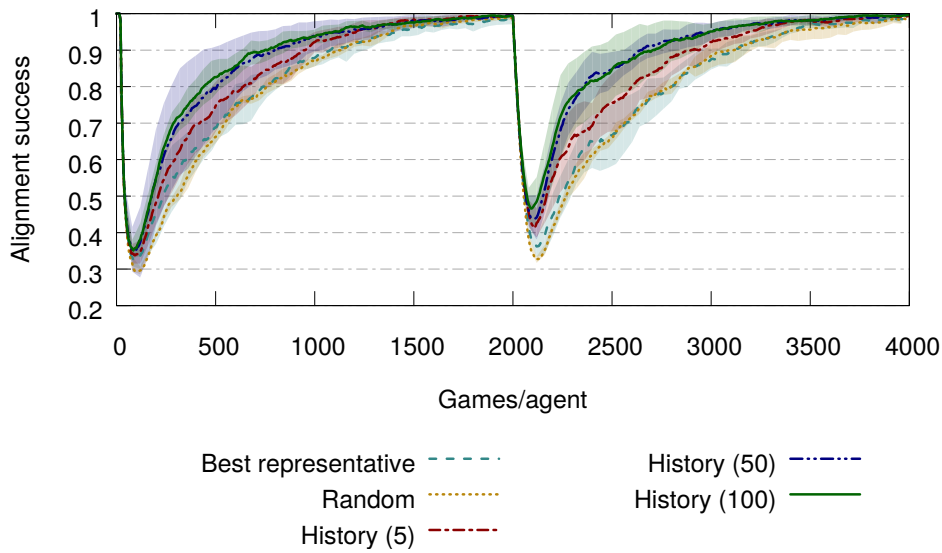

Figure 5.11: Alignment success of the history gender mapping strategy, for memory sizes 5, 50 and 100

Figure 5.11 shows that all tested strategies lead to full alignment success (and thus communicative success, chosen not to show here as it does not provide any additional information). One might wonder how the random strategy achieves this, but we must keep in mind that the learning strategies are simply in charge of gender selection for unknown cases, leaving the transmission and alignment to the scoring and competitor selection mechanisms. Thus a more interesting perspective in this case is offered by the gender distribution over the population vocabulary, as well as the mapping over the semantic space.

Figure 5.12 presents the final outcome of the gender distribution over the agents' inventories. Here, we can notice the difference between the random approach and the other mapping mechanisms. The random strategy will redistribute

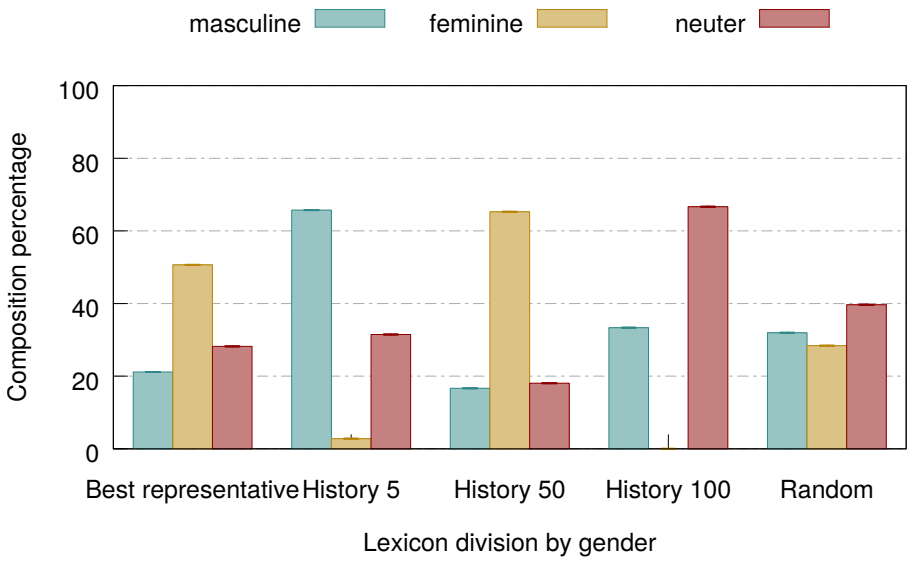

Figure 5.12: Gender distribution over the entire population for different gender mapping strategies

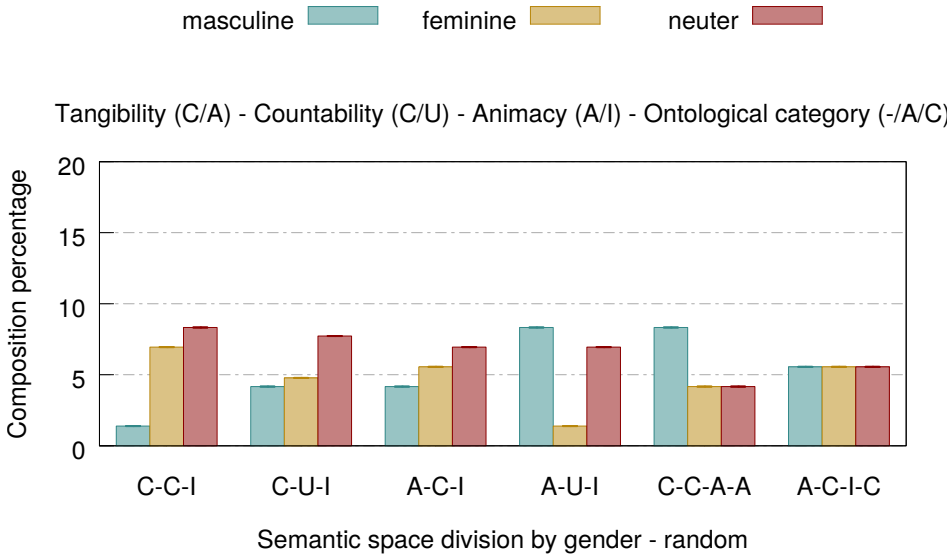

Figure 5.13: Gender distribution over the semantic space for the random gender mapping strategy

the nouns almost equally over the entire vocabulary and also over the semantic space, as presented in Figure 5.13, without making any connection between the semantic knowledge of a word and its gender. Looking back at the initial state of

the system, we see that the random strategy outcome closely resembles it.

Regarding the best representative, we also remark a tendency towards a more equal redistribution of gender over the vocabularies, Figure 5.12, as agents always have an example for each gender in their memory. Additionally, we remark that we have equally distributed nouns over the semantic space (regardless of the initial gender).

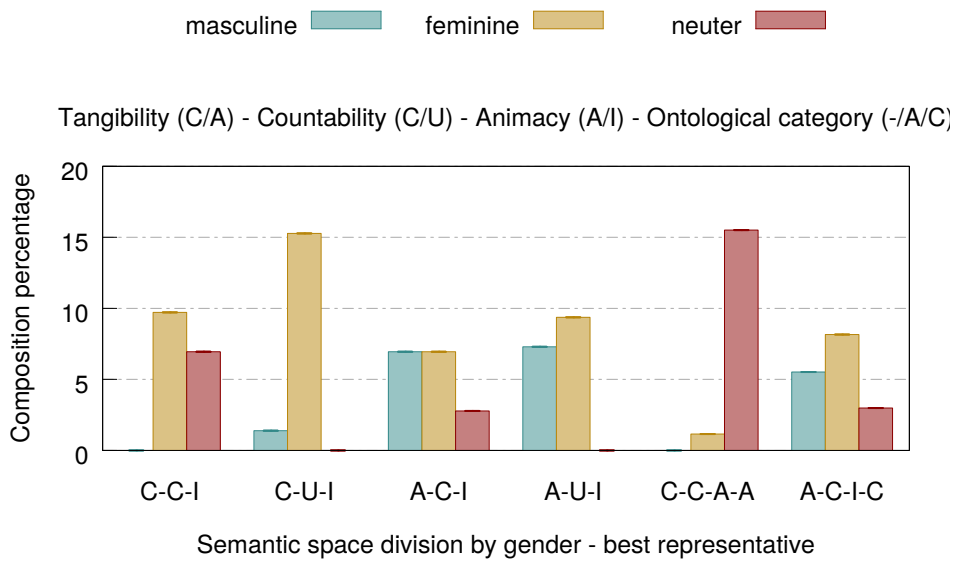

Figure 5.14: Gender distribution over the semantic space for the best representative gender mapping strategy

However, looking at Figure 5.14, we notice how the best representative approach attempts to cluster the genders over the semantic dimensions (for example, concrete - uncountable - inanimate is mapped to feminine, while the animals to neuter).

As for the history mapping mechanism, we can see a lot of variation in the results from Figure 5.12. History with memory size 5 has as dominant genders neuter and masculine, leaving feminine close to zero. History with memory 50, on the other hand, has feminine as dominant, while history 100 lost feminine as a gender altogether.

Looking at the distribution over the semantic space for each variant of the history strategy, we notice that history-5, Figure 5.15, chose masculine for all the cases without a specified ontological category and neuter for the rest (animal and collective), with small exceptions in each case. History-50 on the other hand, Figure 5.16, went with feminine for the first four groups, masculine for animals and neuter for collectives. Lastly, the two-gender system resulted in the case of

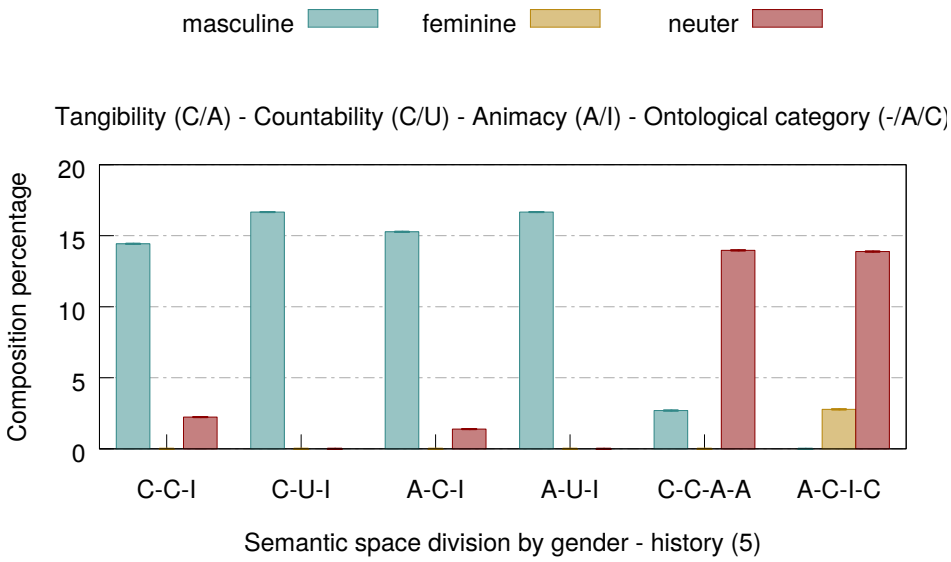

Figure 5.15: Gender distribution over the entire semantic space for the history gender mapping strategy, memory size 5

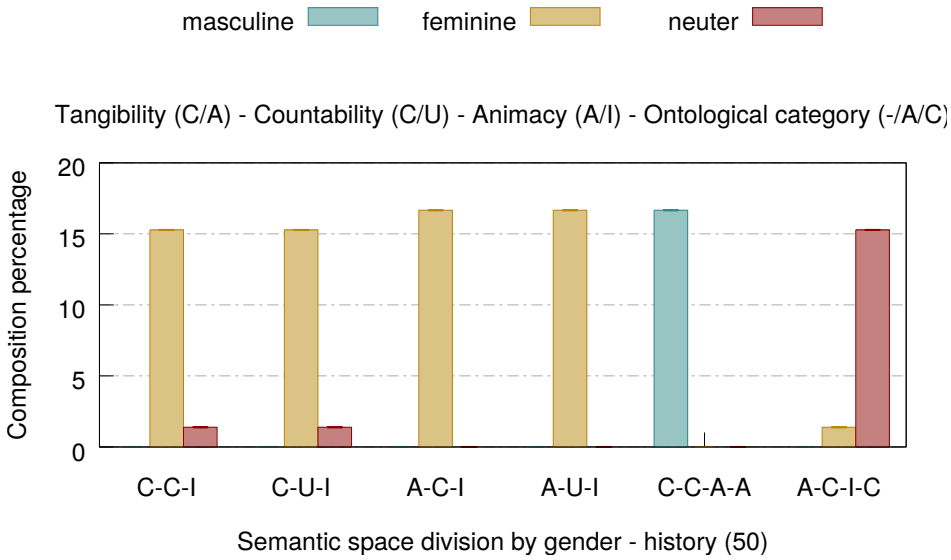

Figure 5.16: Gender distribution over the entire semantic space for the history gender mapping strategy, memory size 50

history-100, Figure 5.17, masculine was chosen for animals and collectives and neuter for the rest.

The memory size does not have a linear impact on the gender mapping. If it is

too small, there might be the case that a gender is not represented, thus no mapping is formed for it. A too large memory again tends to neutralize and eliminate the gender with the least representatives (our purpose for the highest memory size was to simulate an ubiquitous agent over his entire internal representation). The middle size memory seems to perform the best, in the context of our mechanisms, but further tuning is needed in order to see the optimal value.

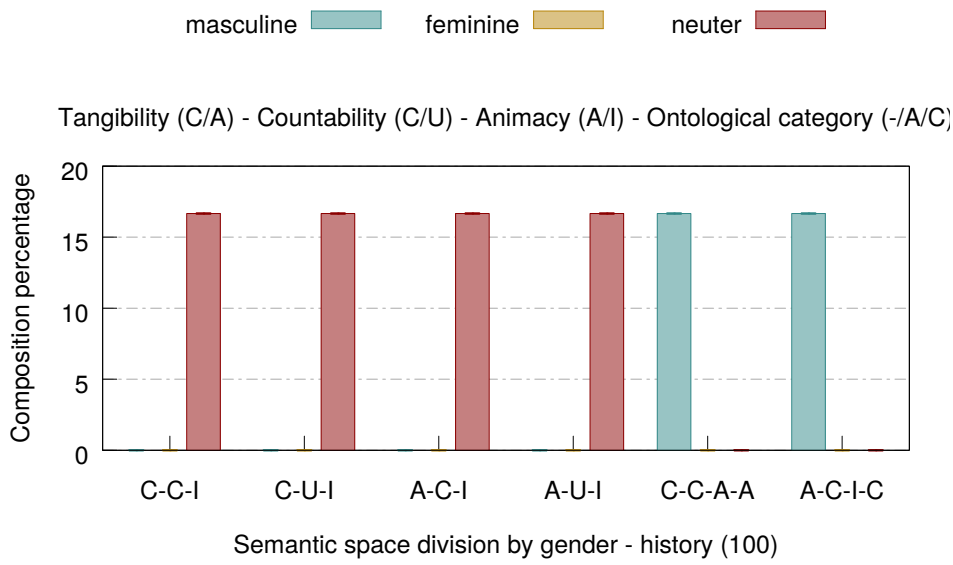

Figure 5.17: Gender distribution over the entire semantic space for the history gender mapping strategy, memory size 100

After seeing the results, we notice that our history strategies manage to get a clean correspondence between gender and semantic space, while starting from scratch after each population turnover. However, we consider this strategy to lack flexibility and also to be radical, in terms of its capacity for gender elimination from the system. Further study is needed in order to establish an optimal memory size for the agents.

As for the best representative, the result are more loose, meaning that the mapping is not completely exclusive for the gender – semantic space, thus maybe bares a closer resemblance to a real-world system. In order to explain this, we can think about the dynamics of the language game. During each interaction, two agents are randomly chosen from the population. Thus, after a generation change, the gender mappings will start emerging from multiple groups inside the population, from the agents that were chosen as speakers and then passed on to their interlocutors. The end result comprises a merged version (agreed upon by the population) of all these incipient mappings, as each agent is forming probably

a different model for its initial gender representatives.

In order to obtain a mapping comparable to the individuation hierarchy (Audring, 2006), agents should be aware of this hierarchical structure and have salient dimensions and, perhaps, genders. Further experiments in this direction might output interesting results.

5.6 Experiment 6 – Stochasticity and Word Level Gender Competition

For the last experiment, we want to investigate further the idea of introducing stochasticity in the system, through the probabilistic competitor selection strategy. We compare again the highest score and the probability (τ 0.1) strategies, but this time looking into final gender distribution over vocabulary and semantic space and also at the gender competition per word.

5.6.1 Experimental Setting

- Competitor selection strategy – highest score / probability (τ 0.1)
- Gender mapping strategy – best representative
- Scoring strategy – lateral inhibition (δ_+ 0.2 and δ_- 0.1)
- Population turnover fraction – 1.0
- Transmission rate – 0.0

Again we decided to take extreme values for the turnover and transmission, such that the strategies can start the gender mapping from scratch, as we only incorporate two turnover events during our run.

5.6.2 Results and Discussion

Gender Distribution

Figure 5.18 provides us with an overview of how the strategies distributed to gender over the entire population's vocabulary. We can notice that for the highest score approach, feminine was the least considered gender, while for the probability strategy we have a more equal distribution.

Looking at Figures 5.19 and 5.20, we can already notice a more consistent gender–semantic space correspondence that resulted from the probability approach,

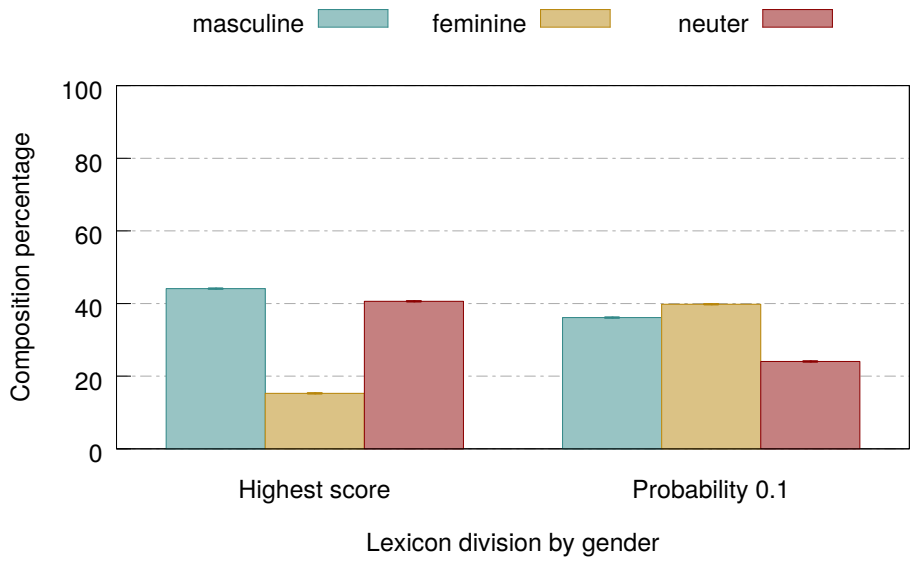

Figure 5.18: Gender distribution over the entire population for the highest score and probability (τ 0.1) selection strategies

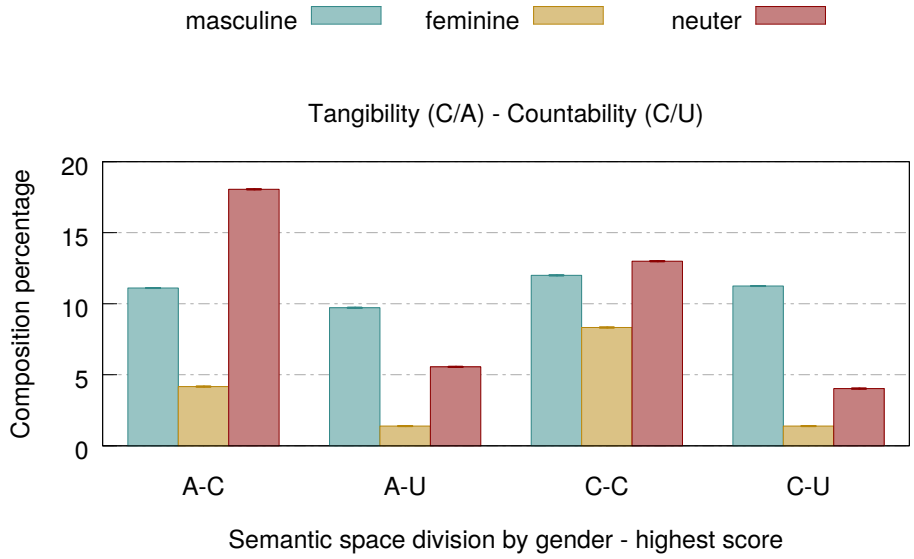

Figure 5.19: Gender distribution over the tangibility–countability dimension for the highest score selection strategy

compared to the highest score one. The neuter and feminine were mostly reserved for the abstract - countable axes and masculine and feminine for the concrete - countable. In order to further explore the results, we look into the distribution

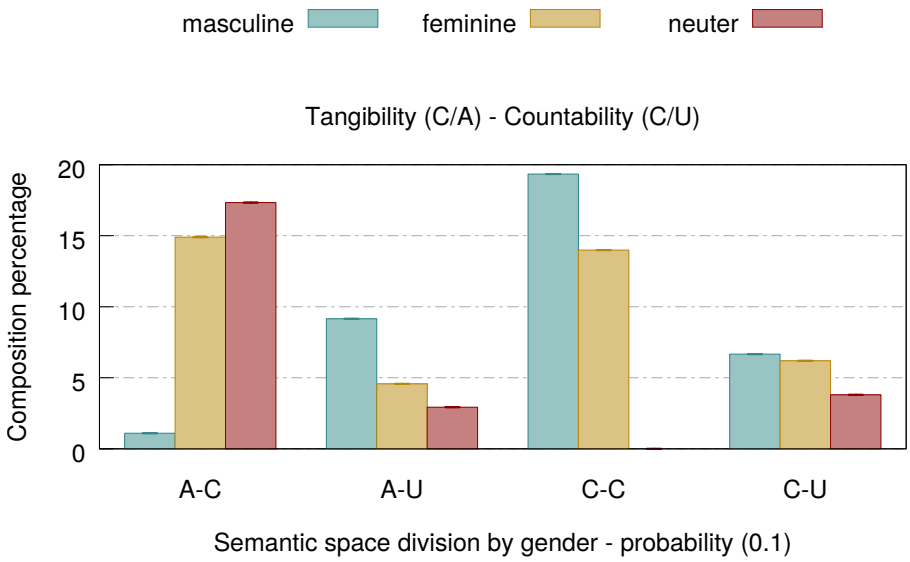

Figure 5.20: Gender distribution over the tangibility–countability dimension for the probability (τ 0.1) selection strategy

over our full semantic space.

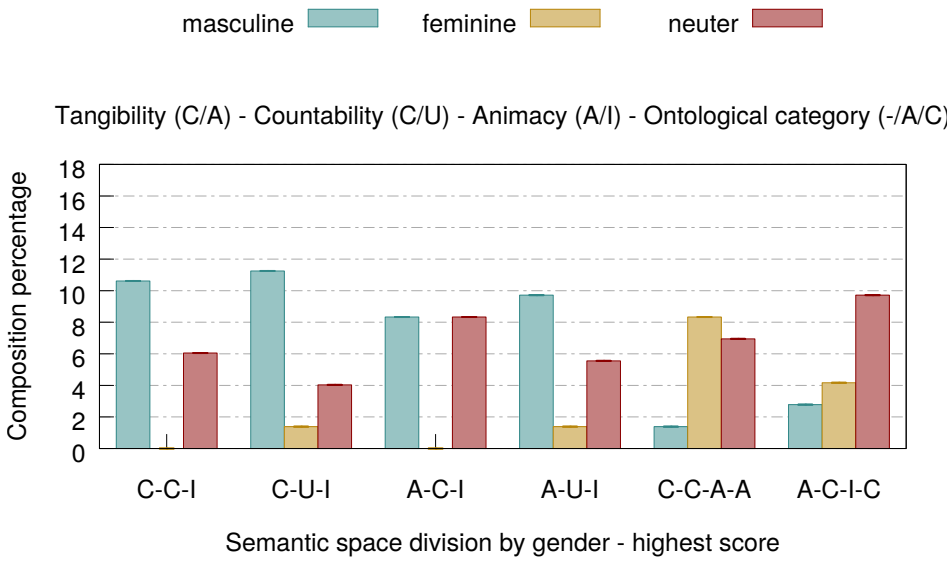

Figure 5.21: Gender distribution over the semantic space for the highest score selection strategy

Figures 5.21 and 5.22 present the gender mapping over the full semantic space

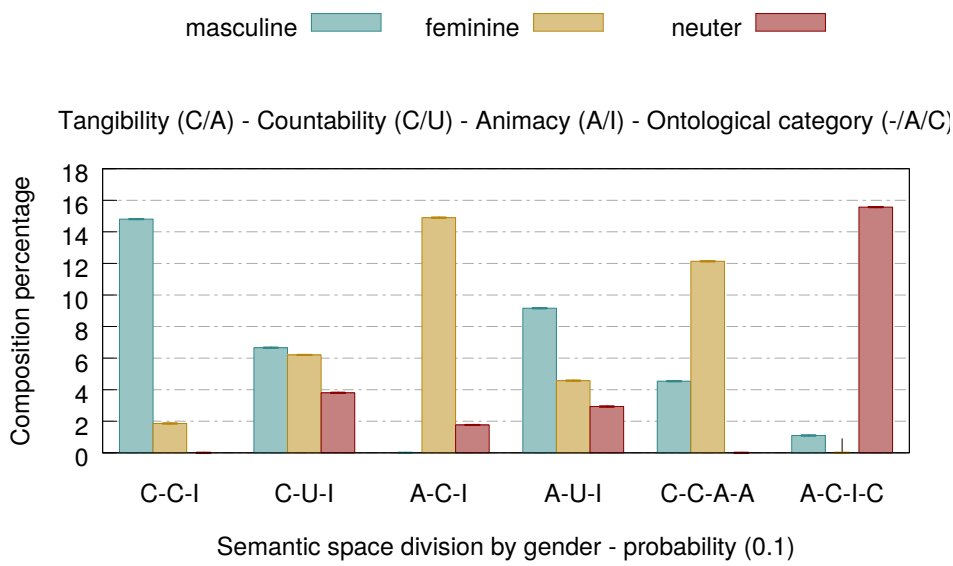

Figure 5.22: Gender distribution over the semantic space for the probability (τ 0.1) selection strategy

we defined. We can notice that the highest score mapped masculine for the groups without a specified ontological category (the first four groups of Figure 5.21), leaving feminine and neutral for the animals and mostly neuter for the collective group. Additionally, neuter was mapped fairly equally over the entire space.

Looking at Figure 5.22 at the resulting mapping of the probability approach, we notice a clear preference of the masculine for the concrete-countable-inanimate, feminine for abstract-countable-inanimate and also animals, while neuter was mapped mostly for the collective nouns.

In order to better understand the effect of stochasticity inside our system, we will look at word level gender competition, expecting there to see how probability selection can contribute to turning over a competition situation, or even maintaining it.

Word Level Gender Competition

	Tangibility - Countability - Animacy - Ontological category					
	C-C-I	**C-U-I**	**A-C-I**	**A-U-I**	**C-C-A-A**	**A-C-I-C**
Dutch	brief	henna	weekend	griep	hert	gezin
English	letter	henna	weekend	flu	deer	family

Table 5.1: Words selected for gender competition analysis

Table 5.1 presents the words selected for the gender competition study. Next, we display and analyse the figures obtained for each of the words, under each considered strategy, in terms of gender preference (the gender belonging to the competitor with the highest score, for each agent) over the population. Each figure has four components, the three genders, along with the status of 'unknown', present in the agents inventory after a population turnover. Thus, in each plot, the two population turnovers are marked by the rise of the unknown gender preference.

Looking at Figure 5.23a, we notice an interesting dynamic after the second turnover, around the 50% level, in the form of a tight competition between masculine and feminine. As for Figure 5.23b, we notice that the probability selection strategy can maintain an equilibrium between gender competitors when the scores reach equality. After the second turnover, we even notice all three genders competing, however, masculine ends up as the winner.

Figure 5.24a also displays an interesting competition between masculine and neuter, showing that the highest score approach can also allow for equilibrium between genders in the population, as long as the usage variation in one direction is not too high (i.e., agents from both sides are chosen equally as speakers) and the preference percentage is almost equal between competitors.

Figure 5.24b presents another interesting dynamic that arises under the probability selection strategy: an equilibrium between competitors, even if their preference percentage is not close to equality.

For the word 'weekend', in Figure 5.25a, we notice probably the most common and expected behaviour: rise of competition after the turnover, with a gender gaining fast control over the system.

Figure 5.26 shows us once again the different behaviours that emerge under the two selection strategies. In the case of the highest score approach, once two preferences diverge far enough, the system seems to head for convergence, while the probability strategy offers more equilibrium opportunities for the competitors.

Figure 5.27a presents however a case of equilibrium between two genders under the highest score strategy as well. We need to also understand that important factors are also the agent and topic selection dynamics, which introduce another form of stochasticity.

Lastly, Figure 5.28b shows again the interesting dynamics that can take place under the probabilistic competitor selection approach, with competition turnarounds and equilibrium between multiple genders.

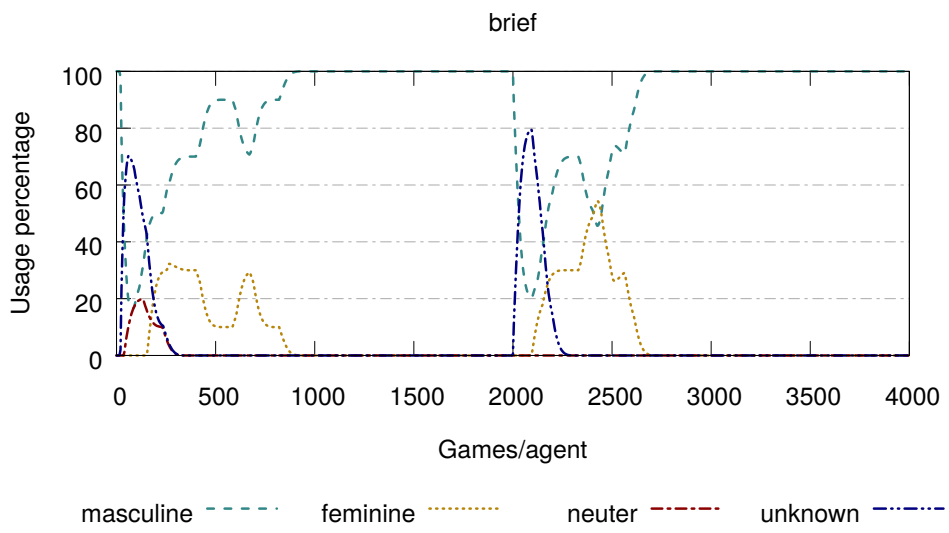

(a) Highest score

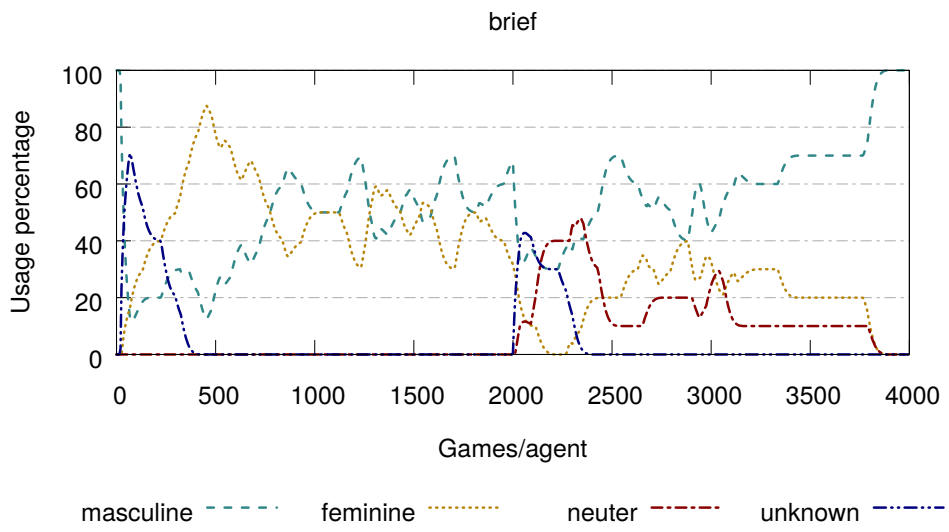

(b) Probability (τ 0.1)

Figure 5.23: Gender competition for the word ‘brief’ (letter)

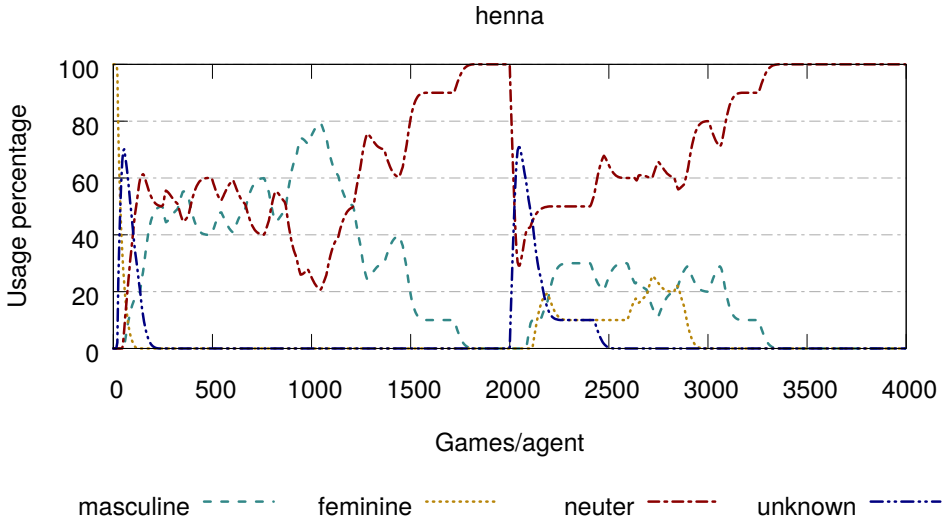

(a) Highest score

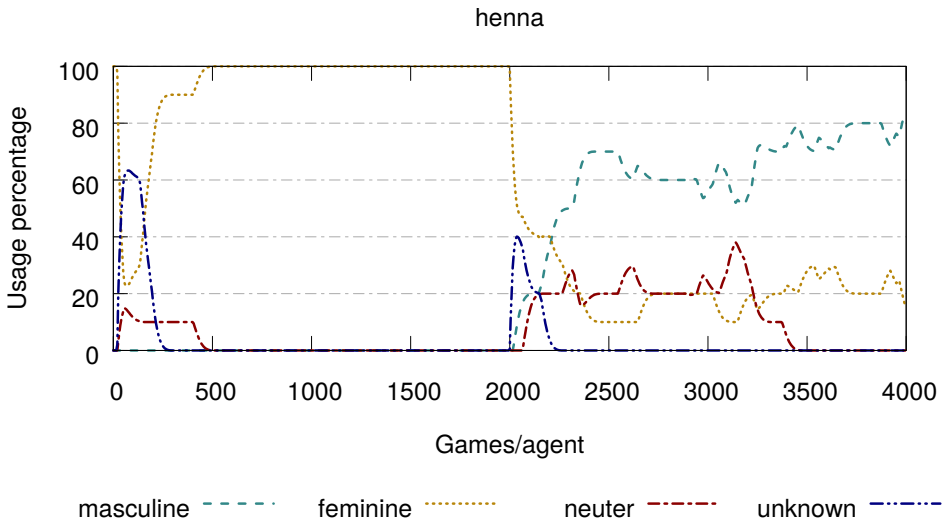

(b) Probability (τ 0.1)

Figure 5.24: Gender competition for the word ‘henna’ (henna)

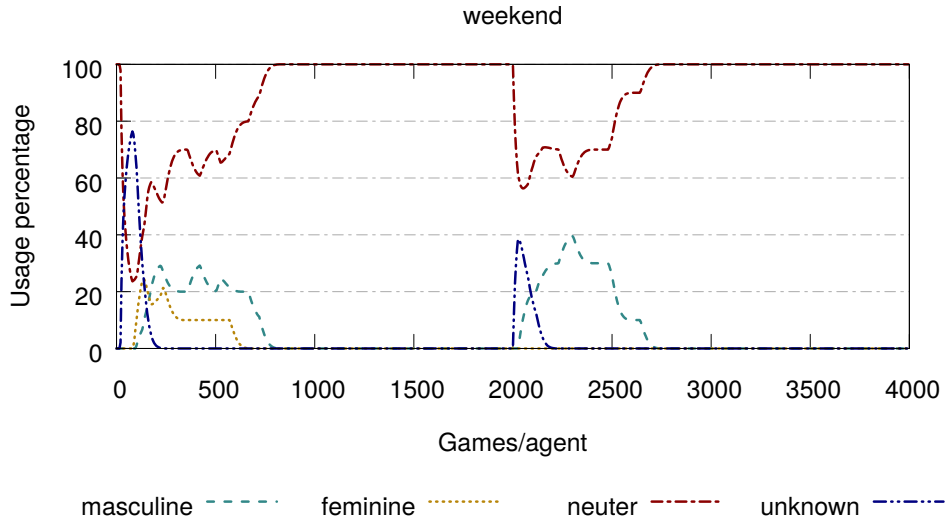

(a) Highest score

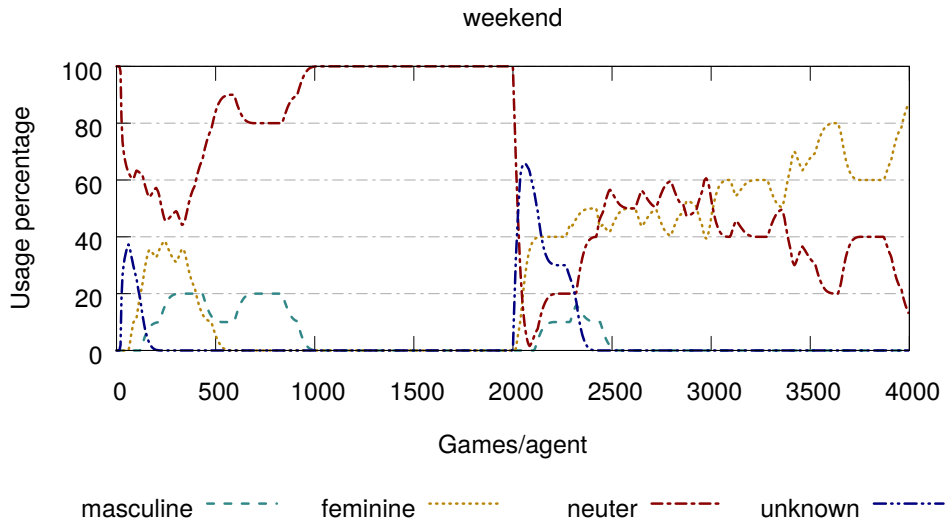

(b) Probability (τ 0.1)

Figure 5.25: Gender competition for the word ‘weekend’ (weekend)

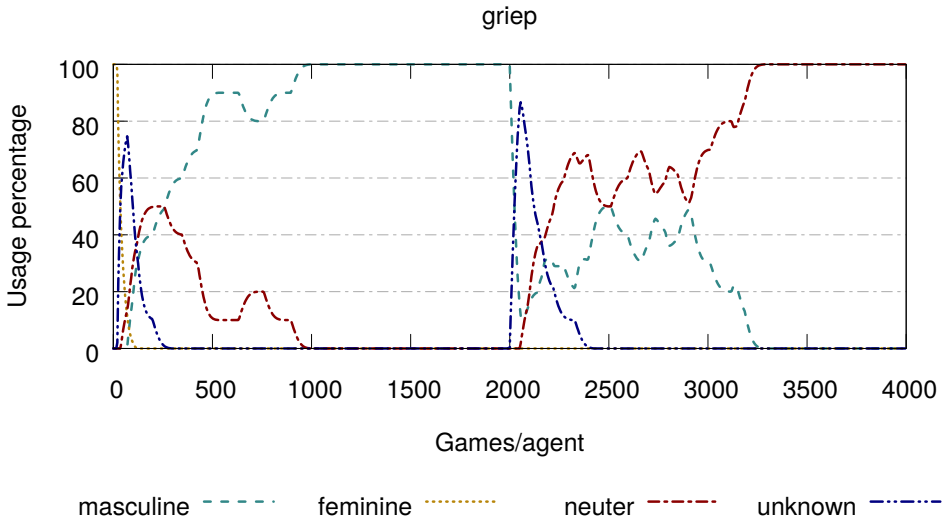

(a) Highest score

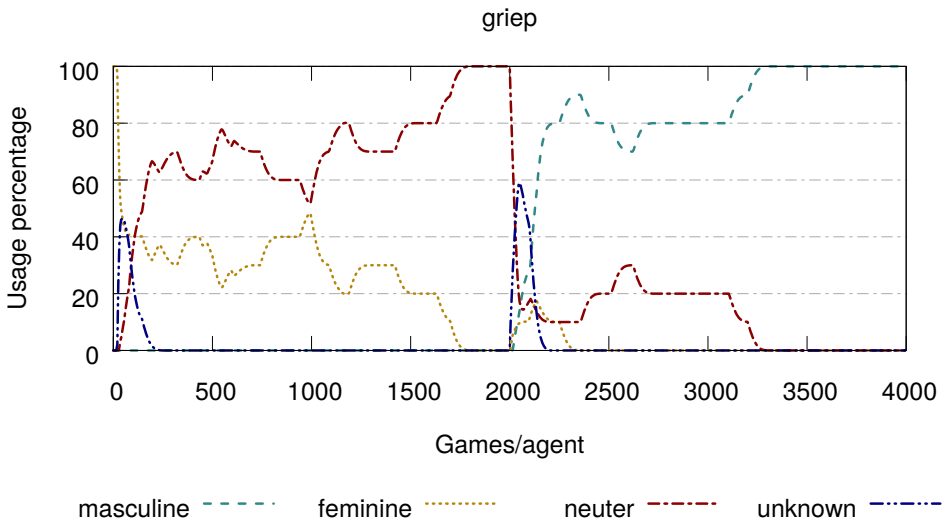

(b) Probability (τ 0.1)

Figure 5.26: Gender competition for the word ‘griep’ (flu)

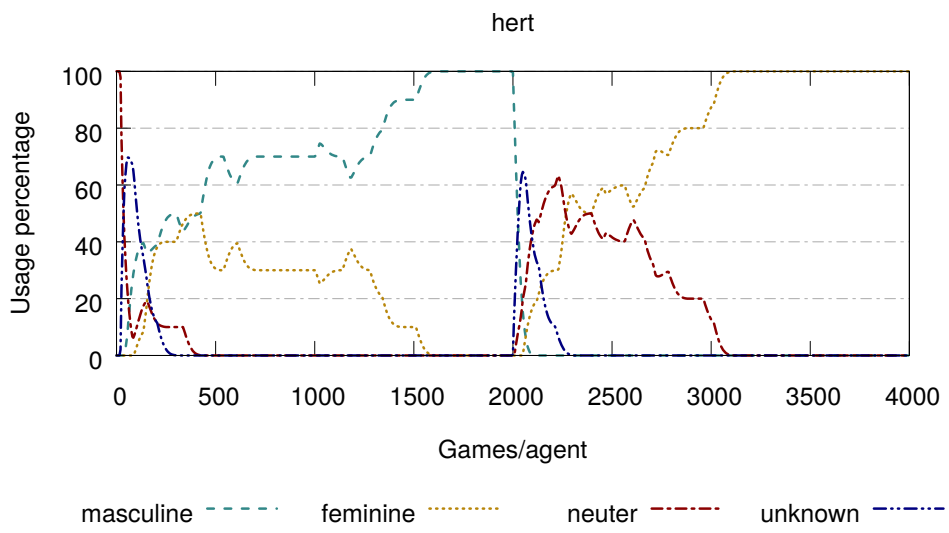

(a) Highest score

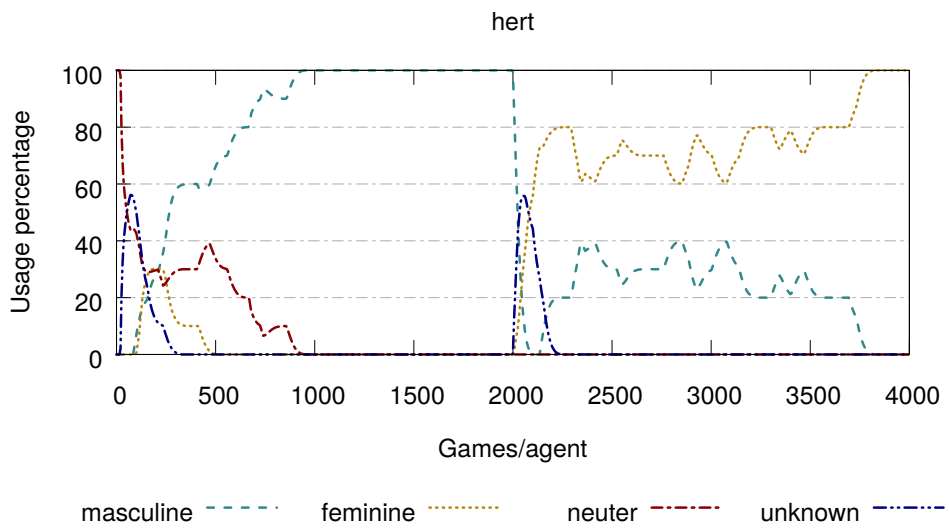

(b) Probability (τ 0.1)

Figure 5.27: Gender competition for the word ‘hert’ (deer)

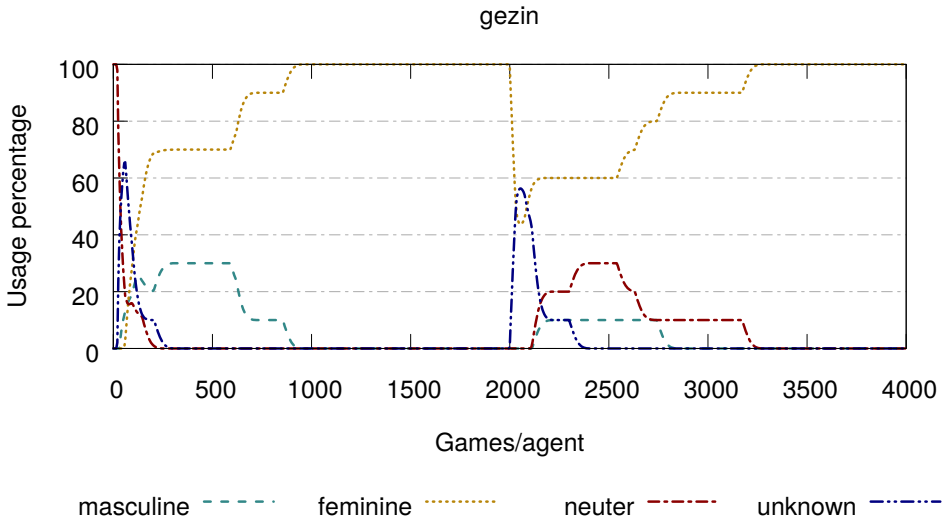

(a) Highest score

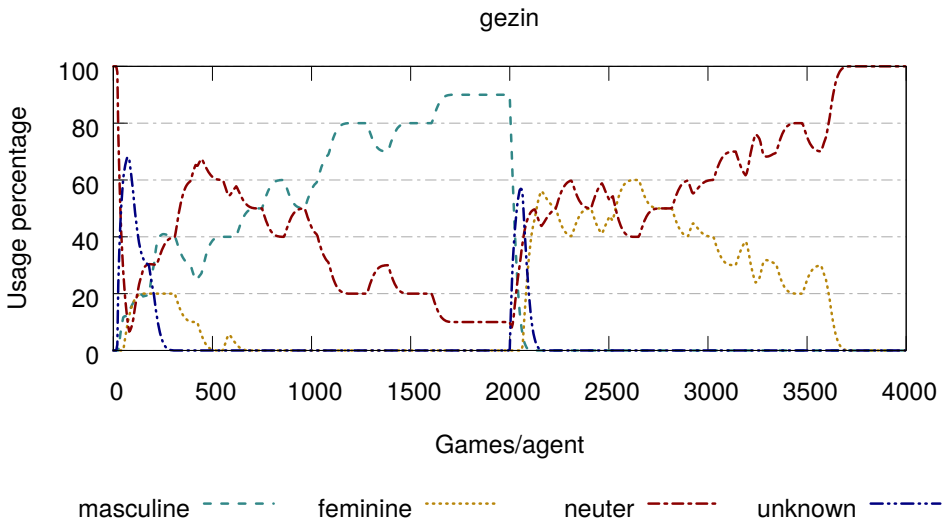

(b) Probability (τ 0.1)

Figure 5.28: Gender competition for the word ‘gezin’ (family)

5.7 Concluding Remarks

We have presented here a thorough step-by-step analysis of our language game, considering all the parameters and strategies that were implemented.

We have managed to show that the system can reach full communicative and alignment success under the provided strategies, even for the extreme case of entire population replacement with no gender knowledge transmission.

The implemented gender mapping strategies are still in a simplistic phase, but they do manage to obtain a semantically driven mapping of gender. We consider that the agents can additionally be equipped with salient semantic dimensions and also repair mechanisms that can enforce a better clustering and balance between the gender models represented in their memory.

Additionally, we have shown that stochasticity can also introduce new dynamics at both population and agent level, taking the system a step further towards a real language in terms of eternally present exceptions from the underlying rules.

All in all, we have presented the ability of mapping gender onto a semantic space and reaching agreement in the system, while making use of a limited memory capacity and the ability of comparing element kinship based on semantic properties.

6
Conclusion

We began our quest for developing an evolutionary language game that can simulate the shift from a syntactic gender agreement system to a semantic one by presenting the research context that this work adheres to. Evolutionary linguistics proposes a research framework in which the systematic nature of language does no longer require a genetic basis, but rather we start asking what are the general cognitive mechanisms necessary for a language system that exhibits certain features to emerge. The driving force behind this approach is cultural evolution, modelled in our case using a multi-agent language game, that allows, through pair-wise interactions among agents in a population, a self-organized system to emerge. Even though the nature of the arisen artificial languages is more simplistic than the natural ones, it is enough to observe features that are shared between the two in order to be able to state that we have obtained a model for the feature in question to culturally develop. An important tool, that this approach makes use of, is Fluid Construction Grammar (FCG), formalism that has at its core the concept of construction (as the basic linguistic unit that provides a mapping between meaning and form, spanning across any linguistic layer as needed – semantic, syntactic, phonetic, morphological). FCG thus allows the implementation of complex linguistic inventories for which it also provides production and parsing mechanisms.

The next step was to define and understand the grammatical gender category and search for its origins. We then shifted our attention to Dutch and how its gender system emerged and evolved. Special attention was given to pronominal

reference and the process of resemanticization that has been observed in northern Dutch. We then turn to southern Dutch, a system in mid-transition regarding pronominal gender among its dialects, from which we draw our inspiration for constructing an evolutionary language game that simulates resemanticization (a shift from an old syntactic system to a new semantic one), caused by imperfect transmission between generations.

After establishing the theoretical grounds of our work, we moved on to present in detail our devised language game: the script, the used Dutch nouns and the semantic space that was defined based on the tangibility - countability - animacy - ontological category dimensions. We showed how the agents' vocabularies were created and manipulated in terms of gender in FCG. We have introduced, in the context of our anaphoric reference language game, the concept of construction similarity, which on a cognitive level can be interpreted as element kinship inference, based on semantic knowledge. Other novel strategies are the probability-based selection between competitors, instead of the usual highest score approach and the success rate as a scoring strategy.

Our two learning strategies, history mapping and best representative, are based on the previously mentioned similarity computation, and function as methods of incipient clustering the nouns with related semantic features to the same gender.

The final step was a methodical analysis of our system. We ran multiple experiments, in order to visualise the impact of our parameters and strategies on the communicative and alignment success and also on the final gender distribution of the gender over the semantic space.

In our case, the frequency scoring mechanism does not converge, in terms of alignment success (thus the agents do not seem to reach the same gender preferences), as they lack flexibility of acquiring new trends emerging in the population, once the frequencies increase high enough.

The competition solving experiments showed that highest score and only the probability approach having a small enough temperature can reach full alignment.

All combinations of transmission and population turnover values manage to reach full convergence, but at different rates, depending on the selected amplitudes.

It is hard to draw solid conclusions about the learning strategies, after just a single run. Unfortunately, we are unable to aggregate over multiple trials, as we desire to see the state of the gender distribution after the convergence of each trial, and not variations in the gender mapping results. But we can say that the learning strategies seem to be able to create a fairly cohesive gender – semantic space mapping, considering the fact that only local clustering methods are applied by the speaker agents.

Lastly, looking into the gender competition at word level can cast an insight

of the current state over the Flemish dialect, as they are at a point where nouns have more than one gender, depending on the considered dialect (as for example with the word 'sneeuw'). Stochasticity can introduce interesting dynamics in the population, which we consider to be worth looking into in more detail.

All in all, we managed to build an initial multi-agent language model that is capable of dealing with the loss of gender marking for the nouns by mapping a new gender according to the semantic dimensions of the represented element.

6.1 Future Work

There are many extensions that can be applied to our incipient linguistic model. We present here a few ideas or directions one can take in order to study further the shift from syntactic to semantic agreement.

To begin with, an aspect that lacks in our system is the notion of biological gender. Thus one can extend the world, such that it also incorporates new objects that make use of this dimension.

An extension to our gender mapping strategies could consist in additional constraints in order to make sure there is enough distance between the gender models formed by each agent. One could introduce clustering quality measurements in order to implement these constraints.

We have talked in Chapter 3 about the individuation hierarchy and how the gender seems to be mapped on those scales. Agents could be made aware of such an hierarchical structure and also have salient dimensions or genders. Applying then a new mapping strategy that makes use of this knowledge could result in a more consistent overall mapping in the population.

Lastly, the agents, instead of memorizing a new gender for each word, could learn a mapping function for gender – semantic space, and only keep in memory exception cases. This might be the source of a cognitive effort measurement, however it implies different mechanisms and an entire restructuring of the model.

Bibliography

Audring, J. (2006). Pronominal gender in spoken Dutch. *Journal of Germanic Linguistics*, *18*(02), 85–116.

Audring, J. (2009). *Reinventing pronoun gender*. Utrecht: LOT Dissertation Series.

Beuls, K. (2011). Disambiguating stickers grammatical agreement as a design pattern. In P. De Causmaecker, J. Maervoet, T. Messelis, K. Verbeeck, & T. Vermeulen (Eds.), *Proceedings of the twenty-third benelux conference on artificial intelligence.*

Beuls, K., & Höfer, S. (2011). Simulating the emergence of grammatical agreement in multi-agent language games. In T. Walsh, NICTA, & U. of NSW (Eds.), *Ijcai proceedings-international joint conference on artificial intelligence* (Vol. 22, p. 61).

Beuls, K., & Steels, L. (2013). Agent-based models of strategies for the emergence and evolution of grammatical agreement. *PloS one*, *8*(3), e58960.

Beuls, K., Van Trijp, R., & Wellens, P. (2012). Diagnostics and repairs in Fluid Construction Grammar. In L. Steels (Ed.), *Language grounding in robots* (pp. 215–234). Berlin: Springer.

Brugmann, K. (1897). *The nature and origin of the noun genders in the Indo-European languages. A lecture delivered on the occasion of the sesquicentennial celebration of Princeton University*. New York: Charles Scribner's Sons.

Chomsky, N. (1965). Aspects of the theory of syntax cambridge. *Multilingual Matters: MIT Press*.

Chomsky, N. (1980). Rules and representations. *Behavioral and brain sciences*, *3*(01), 1–15.

Christiansen, M. H., & Kirby, S. (2003a). Language evolution: Consensus and controversies. *Trends in cognitive sciences*, *7*(7), 300–307.

Christiansen, M. H., & Kirby, S. (2003b). Language evolution: The hardest problem in science? *Studies in the Evolution of Language*, *3*, 1–15.

Corbett, G. G. (2006a). *Agreement*. Cambridge: Cambridge University Press.

Corbett, G. G. (2006b). Gender, grammatical. *The Encyclopedia of Language and Linguistics. 2nd Edition. Oxford: Elsevier*, 749–756.

De Saussure, F. (1916). Nature of the linguistic sign. In D. H. Richter (Ed.), *The critical tradition: Classic texts and contemporary trends* (pp. 832–835). 1998, Boston: Bedford/St. Martin's Press.

De Vogelaer, G. (2009). Changing pronominal gender in Dutch. In S. Tsiplakou, M. Karyolemou, & P. Pavlou (Eds.), *Language variation–European perspectives ii: Selected papers from the 4th International Conference on Language Variation in Europe (iclave 4), Nicosia, june 2007* (Vol. 2).

De Vos, L. (2008). *Pronominale hersemantisering in het Zuid-Nederlandse genussysteem* (Unpublished master's thesis). Universiteit Gent – Faculteit Letteren en Wijsbegeerte.

De Vos, L. (2013). On variation in gender agreement: The neutralization of pronominal gender in Dutch. *Synchrony and Diachrony: A dynamic interface*, *133*, 237.

De Vos, L., & De Vogelaer, G. (2011). Dutch gender and the locus of morphological regularization. *Folia Linguistica*, *45*(2), 245–281.

De Vylder, B. (2007). *The evolution of conventions in multi-agent systems* (Doctoral dissertation). Vrije Universiteit Brussel – Faculteit Wetenschappen.

Gould, S. J., & Vrba, E. S. (1982). Exaptation-a missing term in the science of form. *Paleobiology*, 4–15.

Grimm, J. (1831). *Deutsche Grammatik, Dritter Theil.* Gottingen: Dieterich.

Hauser, M. D., Chomsky, N., & Fitch, W. T. (2002). The faculty of language: What is it, who has it, and how did it evolve? *Science*, *298*(5598), 1569–1579.

Hoffmann, T., & Trousdale, G. (2013). Construction Grammar: Introduction. In T. Hoffmann & G. Trousdale (Eds.), *The oxford handbook of construction grammar.* Oxford University Press.

Janse, M., Joseph, B. D., & De Vogelaer, G. (2011). Changing gender systems: A multidisciplinary approach. *Folia Linguistica*, *45*(2), 237–244.

Kirby, S. (1999). *Function selection and innateness: the emergence of language universals*. Oxford University Press.

Kirby, S. (2000). Syntax without natural selection: How compositionality emerges from vocabulary in a population of learners. In C. Knight (Ed.), *The evolutionary emergence of language: Social function and the origins of linguistic form* (pp. 303–323). Cambridge University Press.

Kirby, S. (2002a). Learning, bottlenecks and the evolution of recursive syntax. In T. Briscoe (Ed.), *Linguistic evolution through language acquisition: Formal and computational models* (pp. 173–204). Cambridge University Press.

Kirby, S. (2002b). Natural language from artificial life. *Artificial life*, *8*(2), 185–215.

Kirby, S., & Hurford, J. (2002). The emergence of linguistic structure: An overview of the iterated learning model. In A. Cangelosi & D. Parisi (Eds.), *Simulating the evolution of language* (pp. 121–147). Springer London.

Kraaikamp, M. (2012). The semantics of the Dutch gender system. *Journal of Germanic linguistics*, *24*(03), 193–232.

Labov, W. (2007). Transmission and diffusion. *Language*, *83*, 344–387.

Lehmann, W. P. (1958). On earlier stages of the Indo-European nominal inflection. *Language*, 179–202.

Lenaerts, T., Jansen, B., Tuyls, K., & De Vylder, B. (2005). The evolutionary language game: An orthogonal approach. *Journal of Theoretical Biology*, *235*(4), 566–582.

Loetzsch, M., Wellens, P., De Beule, J., Bleys, J., & Van Trijp, R. (2008). *The babel2 manual* (Tech. Rep.). Brussels: Tech. Rep. AI-Memo 01-08, AI-Lab VUB.

Loreto, V., & Steels, L. (2007). Social dynamics: Emergence of language. *Nature Physics*, *3*(11), 758–760.

Mesoudi, A. (2011). *Cultural evolution: how Darwinian theory can explain human culture and synthesize the social sciences*. University of Chicago Press.

Pinker, S., & Bloom, P. (1990). Natural language and natural selection. *Behavioral and brain sciences*, *13*(04), 707–727.

Siemund, P. (2008). *Pronominal gender in English: A study of English varieties from a cross-linguistic perspective*. Routledge.

Smith, A. D. (2014). Models of language evolution and change. *Wiley Interdisciplinary Reviews: Cognitive Science*, *5*(3), 281–293.

Steels, L. (2001). Language games for autonomous robots. *IEEE Intelligent systems*, *16*, 16–22.

Steels, L. (2003). Evolving grounded communication for robots. *Trends in cognitive sciences*, *7*(7), 308–312.

Steels, L. (2007). The recruitment theory of language origins. In C. Lyon, C. L. Nehaniv, & A. Cangelosi (Eds.), *Emergence of communication and language* (pp. 129–150). Springer.

Steels, L. (2011). Introducing Fluid Construction Grammar. In L. Steels (Ed.), *Design patterns in fluid construction grammar* (pp. 3–30). Amsterdam: John Benjamins.

Steels, L. (2012). Self-organization and selection in cultural language evolution. In L. Steels (Ed.), *Experiments in cultural language evolution* (pp. 1–37). Amsterdam: John Benjamins.

Steels, L., & Kaplan, F. (1998). Stochasticity as a source of innovation in language games. In *Proceedings of artificial life vi* (pp. 368–376).

Steels, L., & Kaplan, F. (1999). Collective learning and semiotic dynamics. In D. Floreano, J.-D. Nicoud, & F. Mondada (Eds.), *Advances in artificial life* (pp. 679–688). Springer Berlin Heidelberg.

Steels, L., & Loetzsch, M. (2010). Babel: A tool for running experiments on the evolution of language. In S. Nolfi & M. Mirolli (Eds.), *Evolution of communication and language in embodied agents* (pp. 307–313). Berlin: Springer.

Steels, L., & Loetzsch, M. (2012). The grounded naming game. In L. Steels (Ed.), *Experiments in cultural language evolution* (pp. 41–59). Amsterdam: John Benjamins.

Steels, L., & McIntyre, A. (1998). Spatially distributed naming games. *Advances in complex systems*, *1*(04), 301–323.

Steels, L., & Vogt, P. (1997). Grounding adaptive language games in robotic agents. In *Proceedings of the fourth european conference on artificial life* (pp. 474–482). MIT Press.

Taeldeman, J. (2001). De regenboog van de vlaamse dialecten. In J. De Caluwe, M. Devos, & J. Taeldeman (Eds.), *Het taallandschap in vlaanderen* (pp. 49–58). Academia Press.

Van Trijp, R. (2008). *Analogy and multi-level selection in the formation of a case grammar - a case study in Fluid Construction Grammar* (Doctoral dissertation). Universiteit Antwerpen – Faculteit Letteren en Wijsbegeerte.

Van Trijp, R., Steels, L., Beuls, K., & Wellens, P. (2012). Fluid construction grammar: The new kid on the block. In W. Daelemans, M. Lapata, & L. Màrquez (Eds.), *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics* (pp. 63–68).

Wellens, P. (2010). *Adaptive strategies in the emergence of lexical systems* (Doctoral dissertation). Vrije Universiteit Brussel – Faculteit Wetenschappen.

Wellens, P., Van Trijp, R., Beuls, K., & Steels, L. (2013). Fluid Construction Grammar for historical and evolutionary linguistics. In *Acl (conference system demonstrations)* (pp. 127–132).

Wittgenstein, L. (1958). *Philosophical investigations* (G. E. M. Anscombe, Trans.). New York: Prentice Hall.