

An Interactive Activation Model of Affix Stripping

Brechtsje Kingma

August 2013

Master Thesis
Human-Machine Communication
Department of Artificial Intelligence
University of Groningen, The Netherlands

Internal Supervisor:

Dr. Hedderik van Rijn (University of Groningen, The Netherlands)

External Supervisor:

Prof. dr. Jonathan Grainger (CNRS, Aix-Marseille Université, France)



**university of
 groningen**

**faculty of mathematics
 and natural sciences**

Abstract

Morphology influences the structuring of form and meaning representations in the brain. These conclusions are based on a specific type of behavioural experiment, the masked morphological priming (MMP) lexical decision task. In this task, three types of priming are compared: transparent priming, with real suffixed prime-target word pairs (*farmer-farm*), opaque priming, with pseudo-suffixed pairs (*corner-corn*), and orthographic priming, a control condition without legitimate suffixes (*scandal-scan*). A robust finding in these tasks is that priming is stronger in the affixed conditions, whether they are real affixed or pseudo-affixed, than in the non-affixed condition. A second, less robust, finding is that priming is stronger in real affixed than in pseudo-affixed word pairs.

To explain these morphological influences in the larger framework of word recognition, several interactive activation models have been proposed in previous literature. However, none of these models have been implemented, so that the models' assumptions, implications and predictions can't be tested.

This thesis introduces the interactive activation affix stripping (IAAS) model, which is a computational interactive activation model that has been adapted to enable processing of stems and affixes. To this end, a layer with affix nodes is included. These nodes, when activated, strip the affix from the incoming stimulus by inhibiting sublexical nodes. In order to simulate orthographic as well as semantic priming effects, affix nodes can be activated in two ways. They can be activated by the affix's orthographic presence through sublexical nodes, and by the affix's semantic presence through morpho-semantic nodes. A comparison between different types of sublexical representations revealed that precise positional information at the sublexical level is necessary for correct orthographic affix detection and inhibition of sublexical nodes by the affix nodes.

The IAAS model has successfully simulated stronger transparent and opaque priming than orthographic priming in the MMP task. Therefore, orthographic affix stripping can be used as a mechanism to explain differences in processing of pseudo-affixed and non-affixed words. The model has not successfully simulated stronger transparent priming than opaque priming. Therefore, the presented simulations with the IAAS model are inconclusive on whether processing differences between pseudo-affixed and real affixed words can be explained with affix stripping. Finally, a confounding variable, the length of the affix, was detected.

Future research should investigate whether a mechanism like orthographic affix stripping is used in the brain. In addition, research should investigate whether improvements in the model can be made to simulate morpho-semantic affix stripping, or whether alternative accounts are more favourable explanations. Finally, further research should investigate whether confounding variables, such as affix length, also influence human behaviour.

Acknowledgements

First of all, I would like to thank Professor Jonathan Grainger, my supervisor in Marseille, for his endless guidance and support. During my entire stay, he made sure that I was fully included in the research group, that I became acquainted with all aspects of being a researcher, and that my contribution was worthwhile. In addition, I'm thanking my second mentor, Dr. Thomas Hannagan, for his excellent introduction into new research topics and programming languages, and for his patient answers to my many questions.

Then, I wish to thank Dr. Hedderik van Rijn, my supervisor in Groningen. Not only did he keep an eye on the project and was he always available for advice, he also inspired me to switch from Psychology to Artificial Intelligence three years ago, which I didn't regret for a moment.

My thanks go out to my colleagues in Marseille, Dr. Maria Ktori, Marianna Boros, Dr. Lisi Beyersmann, and so many other researchers in the Laboratoire Psychologie Cognitive. Studying in Marseille wouldn't have been as much fun without your advice, interesting collaborations and nice talks during lunch and the department drinks.

I would like to thank my new friends, Sonja, Dmitrij, Max, Luuk, and all other neighbours from around the world who made the year pass so quickly. I thank my friends and family back home, who stood by me during my whole study period, and didn't forget me while I was away. I also thank Trent Reznor, for keeping me concentrated for months on end. Finally, a huge thank you to my soon to be husband, Gert, for encouraging me to go abroad, even though he wouldn't like it, for teaching me so many things, and for helping me in so many ways.

Contents

1	Introduction	9
1.1	Background of the Problem	9
1.2	Modelling Human Behaviour	10
1.3	Research Questions	12
1.4	Thesis Structure	12
2	Literature Review	13
2.1	Background	13
2.1.1	Morphology	13
2.1.2	Masked Morphological Priming	14
2.1.3	Neural Network Models	15
2.1.4	The Interactive Activation Model	17
2.2	Representations of Words and Morphemes	18
2.3	Connections between Word and Morpheme Representations	19
2.3.1	IA Proposals	19
2.3.2	Hybrid Accounts	21
2.4	Orthographic Representations	23
2.4.1	Position-Based Coding	24
2.4.2	Local-Context Coding	25
2.4.3	Dual-Route Approach	25
2.5	Summary	26
3	IAAS	27
3.1	Requirements	28
3.2	General Outline	28
3.3	Sublexical Representations	30
3.3.1	Both-End Position Coding	30
3.3.2	Open-Bigram Coding	32
3.4	Masked Field Weighting	32
3.5	Affix Stripping	34
3.6	Requirement Evaluation with Both-End Coding	38
3.6.1	Suffixes are stripped correctly	38
3.6.2	Words are recognized correctly	40
3.6.3	The model can be primed	41
3.7	Requirement Evaluation with Open-Bigram Coding	41
3.8	Simulation Studies	43
3.8.1	Procedure	44
3.8.2	Results Simulation 1	44

3.8.3	Results Simulation 2	45
3.8.4	Discussion	46
3.9	Summary	47
4	Exploration of the Model	49
4.1	Simulation without Morpho-Semantic Affix Stripping	49
4.1.1	Results Simulations	49
4.1.2	Discussion	51
4.2	Simulation without Affix Stripping	51
4.2.1	Results Simulations	51
4.2.2	Discussion	53
4.3	Summary	53
5	General Discussion	55
5.1	Evaluation of Affix Stripping	55
5.2	Confounding Variables	56
5.3	Stimulus Selection	58
5.4	Interactive Activation	59
5.5	Parameter Setting	59
5.6	Behavioural Plausibility	60
5.7	Conclusions	62

Chapter 1

Introduction

Reading is a complex task, that requires identification of letters, recognition of words and understanding of sentences within seconds. With a brain that seems perfectly suited for this task and years of practice, people are able to process written text in more efficient ways than just scanning a line of letters from left to right. Efficiency is raised by the ways letters are discriminated, words and their meanings are stored, and grammatical structures are processed. Comparing word recognition times of different word types can reveal processing mechanisms and brain structures involved in reading.

According to Diependaele et al. (2009), morphology, the structuring of linguistic units in a language, imposes constraints on the structuring of both orthographic and semantic representations in the brain. This conclusion was based on experimental findings of lexical decision tasks with masked priming, so called masked morphological priming (MMP) tasks (e.g. Rastle et al., 2000, in English; Longtin et al., 2003, in French; Diependaele et al., 2005, in Dutch and French; Kazanina et al., 2008, in Russian). These experiments suggest that humans apply morphological decomposition if a possible affix is detected within a word, even if this word has no semantic relationship with its affixless counterpart.

The aim of this thesis is to explain the influence of morphology on orthographic and semantic representations in the larger framework of word recognition. To accomplish this, an interactive activation (IA) model is proposed to enable processing of stems and affixes. Simulations that resemble MMP experiments will be carried out with this model and the results will be compared to data from human subjects.

1.1 Background of the Problem

The results of MMP experiments can be divided in two main findings. First, a target is generally recognized more quickly if the preceding prime is built up of the target plus a legitimate affix, such as *farmer-farm* or *corner-corn*, than if the prime consists of the target plus some additional letters that do not form a legitimate affix, such as *scandal-scan*. Note that this is an orthographic effect, since the ending *-er* in *corner* isn't an instance of the suffix *-er*: *corner* isn't semantically derived from *corn*. This is in contrast to *-er* in *farmer*, since *farmer* is derived from *farm*. A letter combination that could form an affix, but doesn't

have the function of an affix in that particular word, is called a pseudo-affix. This distinction between pseudo-affixes and real affixes gives rise to the second finding, which is that semantic overlap between the prime and target further increases the speed of target recognition. Reaction times on word pairs that are semantically related, containing real affixes, such as on *farmer-farm*, are shorter than on semantically unrelated word pairs, containing pseudo-affixes, such as on *corner-corn*.

A great number of the MMP studies are accompanied by the proposal of a scientific theory, often in the form of a model (e.g. Grainger et al., 1991; Taft, 1994; Diependaele et al., 2009). These are abstract models that describe the relations between cognitive representations. They are the result of theoretical conclusions, based on experimental data. They are designed in a general framework of assumptions, and are often adaptations of existing models.

However, none of the IA models that are proposed to explain morpho-orthographic and morpho-semantic effects in MMP experiments have been implemented and tested. To design an abstract model is one thing, to implement this model and to simulate behaviour with a real set of stimuli is another. The implementation forces the designer to think about every detail and its exact implications. In addition, simulations from different models can be compared, which makes it easier to reveal the consequences of certain choices in the proposed architecture.

This thesis aims to fill this gap in knowledge by implementing an IA model to simulate MMP experiments and in that way give explanations for the effects of morphology on both orthographic and semantic processing.

1.2 Modelling Human Behaviour

Conclusions from modelling might shed light on the broader debate of how letters and words are processed. Investigating whether morphological decomposition exists and how it works has the general purpose of increasing our knowledge about the neuronal background of reading. Furthermore, as Rastle and Davis (2008) point out, modelling morphological decomposition has the specific purpose of shedding light on the processing advantages that decomposition has and on how children develop the mechanism, which might change reading education methods.

Simulating the MMP task is data-driven, because the data to be simulated already exists. As Figure 1.1 illustrates, behavioural data is acquired by feeding input, which is the task and its stimuli, into a system, which is the group of participants in the study, and measure the outcomes, which are the reaction times (Cohen, 1995). The same input will be fed to the model. The model's output, which is measured in discrete time steps, reflecting processing cycles, is compared to the behavioural output, which is measured in continuous time. If the comparison is unsatisfying, the model might be changed, after which new output is obtained, which is again compared with the behavioural data. With data-driven modelling it is common to perform many iterations of design, testing, implementation and evaluation. A danger of performing a lot of iterations is that the model is too much tuned to the specific data set which is used to test the model with, so that conclusions cannot be generalized to any other data set. Therefore, it is important to use only parts of the data during most of the

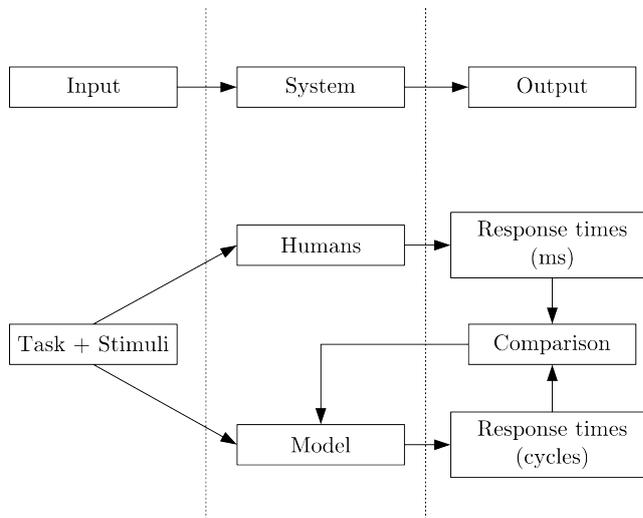


Figure 1.1: A systematic overview of the modelling process in simulating human behaviour

iterations, thereby checking the variation in outcome between different sets of stimuli.

The behavioural data is not the only constraint the model should take into account. For the model to make sense theoretically, the model must fit in a theoretically plausible framework. When an implemented framework exist, the new model can be nested into this existing model (Grainger and Jacobs, 1996). If the new behaviour can be successfully modelled, and the behaviour from the old model remains intact, conclusions can be linked to these other behaviours that previously have been explained in this framework. If modelling isn't successful, either the new addition to the model should be reconsidered, or the framework's assumptions are falsified. In this way, it not only helps understand the current behaviour, but also provides insight in the concepts with which we explain other behaviours. It is therefore useful to take an existing framework, a model that has shown to simulate other behavioural aspects in the same domain, and adapt this to answer the current research questions.

The framework this thesis focuses on is the IA model (McClelland and Rumelhart, 1981), a model widely applied in the domain of visual word recognition. The model that will be proposed in this thesis is the first IA model of morphological processing that is actually implemented. It might incorrectly suggest that no models of morphological processing have been implemented at all. However, several models outside the IA framework have been implemented, such as Koskenniemi (1984) and Baayen et al. (2011). The former of these models is symbolic. Symbolic models aren't very concerned with neurological plausibility, which makes it more difficult to relate the models' outcomes to brain processes. The latter of the models is a neural network model, which is slightly more concerned with neurological constraints than symbolic models. Neither of these models is an IA model, which itself is a neural network model. The IA model imposes additional restrictions on the model, based on experimental findings, which can more easily provide insight into neurological processes that

are involved. In addition, the IA framework has the advantage that certain experimental tasks, such as a lexical decision task with priming, can be simulated more easily than the models mentioned above, so that results of the model can be compared directly to behavioural results.

1.3 Research Questions

The main research question that will be answered in this thesis is:

Can morpho-orthographic and morpho-semantic effects in human processing of (pseudo-)affixed words, as displayed in masked morphological priming experiments, be simulated by incorporating an affix stripping mechanism in an interactive activation model?

To answer this question, the basic lay-out of the original IA model must be extended to enable processing of stems and affixes. This extension will be an instance of affix stripping, which is the inhibition of the sublexical nodes that represent the affix. The extension should be in line with the theoretical constraints the model imposes and should reproduce the results that were obtained in the behavioural MMP research. The first criterion can be met by careful implementation and consideration of the theoretical framework. The second criterion can be met by statistical comparison of the behavioural data and the results from the model.

An important decision in the model's design is what kind of representations will be used in the sublexical level. The choice between alternatives reflects an ongoing debate about how words are represented sublexically. An important difference between proposals is how precisely positional information is coded in these sublexical representations. In relation to this debate, an additional research question will be answered:

Is precise positional information necessary for affix detection?

In order to answer this question, two sublexical coding schemes will be implemented and their ability to recognize affixes correctly will be compared. The first coding scheme, both-end position coding, uses precise positional information, whereas the second scheme, open-bigram coding, only uses a coarse form of location information.

1.4 Thesis Structure

The outline of the thesis is as follows: Chapter 2 gives an overview of the main debates regarding morphological processing and of the interactive activation (IA) models that have been proposed to explain it. Chapter 3 introduces the interactive activation affix stripping (IAAS) model and presents simulations of two masked morphological (MMP) priming experiments. Chapter 4 explores the model by removing affix stripping mechanisms, to determine their contribution to the simulatory results. Finally, Chapter 5 evaluates the model and discusses topics for further research.

Chapter 2

Literature Review

Over the past four decades, many possible explanations for morphological processing have been given, each with its own pros and cons. The explanations mainly differ in the way in which morphemes and words are stored in the brain and the way these representations are connected. Related to this is the debate about how orthographic units are represented in the brain. These issues will be discussed in Section 2.2, Section 2.3 and Section 2.4, respectively. First, Section 2.1 defines the fundamental terminology.

2.1 Background

This section starts with a short description of morphology. Then, it explains the MMP task, a common task to study morphological effects in visual word recognition, and the basic results that have been obtained with it. The section ends with a detailed description of Neural Network models in general and the IA model, the specific framework that will be used to implement the IAAS model.

2.1.1 Morphology

Words are built from morphemes, the smallest semantic units in a language. Words that are morphologically complex, contain two or more morphemes. These morphemes can be free or bound. Free morphemes could exist independently as a word. For example, the word *keyboard* is made up of *key* and *board*, which are words in their own right.

In contrast to free morphemes, bound morphemes, also called affixes, are only used by attaching them to another morpheme. In the word *co-worker*, *work* is the stem. *Work* with the suffix *-er* attached to it means *someone who works* and the prefix *co-* changes the meaning of *worker* to *someone who works with you*. The affixes *-er* and *co-* are not words in their own right and are therefore bound morphemes. Not every bound morpheme can be attached to every other morpheme. The suffix *-er* means *someone who or something that...*, after which a verb is inserted. This suffix can be added to a verb by rule, but not to other word categories.

Words can be categorized consciously according to grammatical rules as provided above. However, grammatical rules are just one side of the morphological

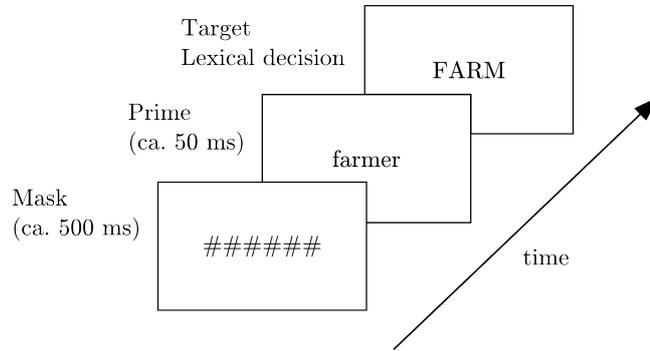


Figure 2.1: A trial in a masked morphological priming-lexical decision task.

story, as human processing doesn't need to reflect the grammatical categorization. To acquire knowledge about how morphological information is stored and to what extent people use this information during reading, behaviour should be studied. A popular device for this is the MMP task.

2.1.2 Masked Morphological Priming

Priming is the modulation of the processing of a given stimulus (the target) by another stimulus (the prime). In the masked priming paradigm (Forster and Davis, 1984), prime presentation is so short that the participant doesn't report having seen the stimulus. The prime is typically presented right before the target. The participant has to perform a task with this target stimulus. The participant is mostly unaware of the prime, because its conscious processing is masked by the presentation of the target. However, if priming influences the execution of the task, the stimulus must be subconsciously processed to a certain degree.

The lexical decision task is a paradigm in which the participant has to decide whether the target is an existing word or a non-word. Figure 2.1 gives a schematic overview of stimuli through time in a trial. Reaction times and error rates of the participant are measured. These are typically compared between conditions that differ in the relation between the prime and target. This relation can cause facilitation or interference in processing. To obtain the absolute effect of priming, each priming condition is subtracted from a condition in which the prime and target are totally unrelated, both semantically and orthographically. A difference between the experimental and control conditions is an indication for overlapping representations of prime and target in the brain.

In masked morphological priming (Forster et al., 1987; Grainger et al., 1991), the different conditions are defined by the way the prime is morphologically related to the target. For example, the word *farmer* is morphologically close to its stem, *farm*, since a *farmer* is someone who works on a *farm*. *Farmer-farm* is a prime-target word pair in the transparent priming condition. In contrast, a *corner*, which also contains the legitimate suffix *-er*, isn't someone who *corns*. It even has nothing to do with *corn*. *Corner-corn* is a word pair in the opaque priming condition. Many researchers carried out lexical decision experiments with MMP (e.g. Rastle et al. (2000); Longtin et al. (2003); Diependaele et al.

(2005, 2009) in English, French and Dutch, respectively). In these experiments, the two conditions mentioned above are compared with each other, as well as with a third condition; *orthographic* or *form* priming. In this last condition, the prime is also created by adding letters to the target, but these letters aren't a legitimate affix. For example, *scandal* could be seen as composed of the target word *scan*, and an illegitimate suffix *-dal*. This condition is included to measure the priming that is purely orthographic. The difference between this condition and the opaque and transparent conditions reflects the additional priming, on top of orthographic priming.

The majority of the studies focused on suffixes. To study the generalization to prefixed words, Diependaele et al. (2009) used prefixed primes. Their results reflect priming patterns from studies with suffixed primes. Therefore, they suggest that similar underlying processes lead to the priming patterns in both the prefixed and the suffixed word studies.

Rastle and Davis (2008) have collected the results from 19 MMP studies. The average priming effects per condition are 2 ms in the orthographic condition, 23 ms in the opaque condition and 30 ms in the transparent condition. In each study, opaque priming was included, as well as orthographic priming or transparent priming, or both. From the 14 studies that included the orthographic condition, orthographic priming was smaller than the opaque and transparent priming in all but one studies. This is therefore a strong indication that the orthographic presence of an affix, regardless of whether it is a real affix or a pseudo-affix, influences visual word recognition. Furthermore, the 16 studies that included the transparent condition, transparent priming was stronger than opaque and orthographic priming in 11 studies. This indicates that the semantic presence of a real affix also influences word processing, although this effect isn't as robust as the effect from orthographic presence of the affix.

Many of the articles that present behavioural findings on morphological processing include a theoretical proposal of how this processing might occur in the brain. These proposals vary from rough indications about what factors might influence processing to detailed descriptions with instructions for implementation into a computational model. A fair amount of these models, including the model that will be implemented in this thesis, are designed within the framework of the IA model.

2.1.3 Neural Network Models

The IA model (McClelland and Rumelhart, 1981) is a connectionist parallel distributed processing model, also known as a neural network model. Neural network models are composed of nodes, that spread activation to connected nodes. Each cycle, each node i calculates its difference in activation Δa_i as follows;

$$\Delta a_i = \begin{cases} (\max_i - a_i)\text{input}_i - \text{decay}_i(a_i - \text{rest}_i), & \text{if } a > 0, \\ (a_i - \min_i)\text{input}_i - \text{decay}_i(a_i - \text{rest}_i), & \text{otherwise.} \end{cases} \quad (2.1)$$

where \min_i and \max_i are the node's minimum and maximum allowed activation values and a_i its current activation. decay_i is the node's strength to return to its resting value rest_i . input_i is constructed as follows;

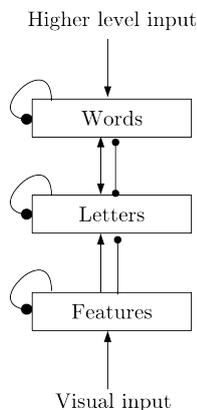


Figure 2.2: The interactive activation model

$$\text{input}_i = \sum_j w_{ij} \text{output}_j + \text{extinput}_i + \epsilon \sim \mathcal{N}(0, \sigma^2) \quad (2.2)$$

where for each other node j the weight of the connection between the two nodes w_{ij} is multiplied by the other node's activation output_j . A positive weight resembles an excitatory connection and a negative weight an inhibitory connection. The sum of these input values is added to the combined input from external sources extinput_i and noise ϵ from a Gaussian distribution.

These dynamics result in decreasing change in activation as a node's activation moves closer to its minimum or maximum value. If the input is stable, the change will ultimately be negligible. There is an equilibrium if all nodes are stabilized.

“Neural network model” implies that it resembles neural function. This is true to some extent. Like a computational neural network, brain tissue consists of cells that become active and spread their activation by exciting or inhibiting other cells (Ellis and Humphreys, 1999). However, one node in a neural network model doesn't need to reflect a single brain cell, but can also represent a group of brain cells. Although modelling at a more abstract level might reflect the human brain less accurately, it is easier to implement, so that the modeller can focus on the architecture of the model instead of the particular biological aspects of the human brain. In addition, the exact behaviour of brain cells and their connections are not yet fully understood. Modelling might therefore be a good way to hypothesize certain brain lay-outs.

Another similarity with the brain is that nodes in neural network models are grouped together in layers that are connected with neighbouring layers. In brain tissue, activation also flows through different groups of cells with distinctive functions in processing. In general, neural network modellers aspire to design a simplified version of certain brain segments and the connections between them. The aim is not to build a realistic model of the brain, but to focus on a specific aspect of behaviour. However, constraints imposed by brain research are incorporated to limit the number of design options and to arrive at a simulation that fits in the general line of brain research.

2.1.4 The Interactive Activation Model

The original IA model was introduced within the domain of visual word recognition. It can account for a number of behavioural effects in reading, that has not been explained within a single model before.

In the IA model, nodes, as described above, are grouped together in layers (see Figure 2.2). The lowest level contains nodes that represent letter features. The features in this layer represent basic geometrical shapes, such as lines in certain directions. Evidence for specialized cells that are activated upon presentation of such geometrical shapes is obtained in single cell recording studies (e.g., Hubel and Wiesel, 1959, 1968, studied, respectively, in cats and monkeys). When a letter is seen, the features that are present in the input become active. They spread their activation further to the letter level. This level has a node for each letter at each position. For every position, the letter that contains most active features, gets activated most. For instance, presenting the stimulus *book* results in the strongest activation of B1, O2, O3 and K4.

Connections between letter nodes are mutually inhibitory. This results in an effect called competition, or lateral inhibition. The node with the strongest inhibitory power decreases activations in neighbouring letter nodes, and in that way decreases the inhibitory power from these nodes. This amplifies the difference in activations between the nodes. In such a winner-take-all system, the most activated node eventually inhibits all other nodes. Competition is also present in the next level: the word level. The letters with the highest activations activate the words that contain these letters. Upon presenting *book*, B1, O2, O3 and K4 activate BOOK. However, O2, O3 and K4 also activate COOK and NOOK, and B1, O2 and O3 also activate BOOM, although activation of these nodes is less than of BOOK. The words compete with each other, so that one word will be activated most.

Between the feature and letter levels, connections are exclusively bottom-up; the features activate or inhibit the letters. However, between the letter and word levels, activation also runs top-down from the word to the letter nodes. This means that when a word becomes activated, it increases activation of the letter nodes that make up the word. The mutually excitatory connections between the levels increase each other's activation, an effect called resonance.

An important finding resonance can account for is the word superiority effect. This implies that letter recognition is generally faster if the letter is presented in a word than if just a single letter is presented or the letter is part of a non-word (Reicher, 1969; Wheeler, 1970). The IA model explains this by resonance between letter and word level (see Figure 2.3). The resonance increases the top-down activation, which is added to the bottom-up activation from the feature nodes and facilitates letter recognition. If only a single letter is presented, bottom-up activation is the only source of activation, so that no resonance takes place. This causes the activation to increase more slowly and therefore causes recognition to be slower.

With this background information in mind, we can look at the main topics of debate on morphological processing. The next section describes some views on the way in which complex words are stored in the brain.

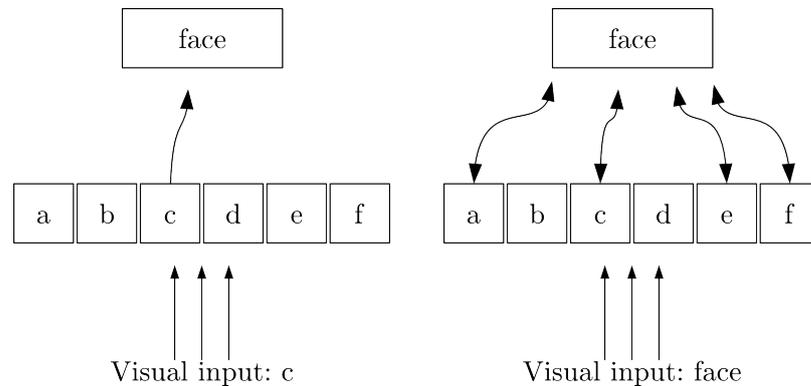


Figure 2.3: The word superiority effect; letter nodes get only bottom-up activation if a letter is presented, but get bottom-up and top-down activation if a word is presented and a word node is sufficiently activated

2.2 Representations of Words and Morphemes

Like most languages, English language allows us to combine and recombine morphemes to form new words. Seeing a new combination of morphological units for the first time, for example the recently introduced word *unfriend*, doesn't require one to look it up in a dictionary. Instead, the meaning can be derived from the knowledge one has about the components and the context it appears in. The question that arises is how these words are stored in memory. Are morphologically complex words generally represented as whole words, as combinations of their morphological parts, or both?

There is a trade-off between lexicon storage efficiency and language processing efficiency. At one end of the spectrum exclusively whole-word representations are stored. This doesn't require difficult processing steps, because each known word with each meaning is available. The right entry only has to be retrieved from memory. However, if morphologically complex words are stored as whole word forms, every stem-affix combination must be remembered. This requires a lot of memory space, since there exist far more morpheme combinations than single morphemes. In addition, it increases the effort of learning new morpheme combinations, since a complete word must be stored, instead of the combination between the morphemes.

At the other end of the spectrum exclusively morphemes are stored. If complex words are stored as combinations of morphemes, only the separate morphemes need to be stored. Although this is memory efficient, additional information should be stored as well, because the relation between two morphemes isn't always that clear (Sandra, 1994). For example, one needs to know whether or not a morpheme combination is legal. *Bitter-ness* is allowed, whereas *winter-ness* isn't. Furthermore, the meaning of a morpheme combination might mean something else than the sum of its parts. For example, the meaning of *re-send*, *re-boot* or *re-fill* can be easily derived with the knowledge that the prefix *re-* means "again". However, this knowledge can't be used in *request*, *reveal* or *reduce*, in which *re-* is a pseudo-affix. In other cases, the original meaning of *re-still* applies, but the stem isn't used any more with the original meaning, such

as in *repeat*, from Latin *petere* (“attack, beseech”).

Most contemporary researchers do not regard storage as the limiting factor in word representations (Nooteboom et al., 2002). Models are proposed with morpheme representations as well as whole-word representations, and even separate idiom and high frequent phrase representations (Sprenger and van Rijn, 2013). Words with pseudo-affixes, as well as highly frequent real-affixed words, have whole-word representations. The meaning of other words can be derived from their parts. However, opinions differ on how these representations are connected.

2.3 Connections between Word and Morpheme Representations

The first proposed model to explain morphological processing was presented almost four decades ago by Taft and Forster (1975). In a lexical decision experiment without priming, participants took longer to classify non-words of which a prefixed form exists (e.g., *scure* from *obscure*, ob- is a prefix) than non-words of which no prefixed word exists (e.g., *zette* from *gazette*, ga- isn’t a prefix) (Taft and Forster, 1975; Taft, 1994). From this and other results, the authors concluded that stems and prefixes of morphologically complex words are stored separately in memory.

They proposed that, upon seeing a complex word, the prefix is stripped from its stem. Then, the stem is retrieved from a lexicon, after which is tested whether the prefix and the stem can co-exist in one word. This is an example of a serial search model, in which a word is matched to the lexical entries in a left-to-right manner. Prefixed words form large categories of entries that start with equal letter combinations. Stripping of the prefix before lexical matching reduces the number of entries the word needs to be compared to. This increases processing speed for prefixed words.

The model describes seven different processing steps and makes quite a few assumptions. Implementing and testing such a model might give a critical view on the model’s details. However, the model is theoretical, it has not been implemented.

In contrast to this serial search model, parallel access models of word recognition have been proposed. The original IA model (McClelland and Rumelhart, 1981) that has been explained in Section 2.1.4 is a parallel access model. In the IA model, all letters in four letter words become available at the same moment. The word-superiority-effect that can be explained by the parallel processing in their model is therefore evidence for parallel access in reading.

A range of IA models have been proposed to explain morphological processing. The next section will explain a number of these models in more detail.

2.3.1 IA Proposals

IA models of morphological processing generally agree on the fundamental principles of the original IA model. Like the original model, they place word representations in a layer above one or more sublexical levels, such as the feature and letter levels in the original model (see Section 2.1.4). The first IA model contained higher level top-down input to the word layer. Most subsequent morpho-

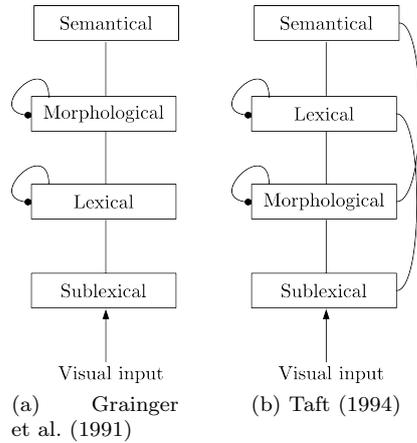


Figure 2.4: Simplified schematic IA proposals of morphological processing

logical models have replaced this input with a layer of semantic representations on top of the word level.

Grainger et al. (1991) were among the first to study morphological representations with the masked priming paradigm (see Section 2.1.2). Grainger et al. (1991) describe two experimental effects that influence morphological processing. The first effect was a faster target recognition if prime and target were morphologically related (transparent condition) than if they were only orthographically related (orthographic condition) or not related at all. This suggests that words within the same morphological family are interconnected and activate each other. The second observed effect was slower recognition of targets with orthographically related primes than with unrelated primes. This means that words that are orthographically overlapping inhibit each other.

To incorporate both types of connections in a single model, the authors propose an IA model with a morphological level above the lexical level (see Figure 2.4a). This level contains morphological family nodes, with excitatory connections from and to the whole-word representations that belong to these families. This results in facilitatory effects between morphologically related words and therefore explains the first effect. The second effect can be explained by the inhibitory between-word connections, that were already present in the original IA model. If words are orthographically overlapping, but morphologically unrelated, they inhibit each other so that recognition is slower.

To implement the model, one could adapt the original IA model. “[O]ne could simply add a morphological level of representation situated directly above the word level” (p. 381). Nevertheless, like the serial access model, this IA model hasn’t been implemented.

The placement of the morpheme level above the word level is in contrast to the serial access model, where prefixes are stripped from the word before lexical access. Based on the experimental data, the authors can’t differentiate between the two proposals. However, as Sandra (1994) points out, affixes can’t be stripped off prelexically just on an orthographic basis, because pseudo-affixes would also be stripped off. Taft (1994) adds to this that people would need a low level prefix detection mechanism, which seems unlikely since people had

great difficulty in identifying letter strings as prefixes.

Therefore, Taft also proposes an IA model. In line with his earlier proposal (Taft and Forster, 1975), the morphological level is placed sublexically (see Figure 2.4b). This differs from sublexical prefix stripping, because prefixes are not treated as a separate category that are removed from the rest of the word. Instead, bound morpheme nodes are activated by grapheme nodes in the bottom level. In their turn, these morpheme nodes activate nodes in the lexical level. Free morphemes bypass the morpheme level. Prefixes on themselves can't be words, but they bear semantic information, so the bound morpheme nodes have direct connections with the concept level as well. Theoretically, the bound morpheme nodes could be established by Hebbian learning; connections between cells that fire simultaneously are strengthened (Hebb, 1949). The nodes act as a connection between the affixes' letter combinations and their semantic representations. These connections are strengthened when the semantic representation is active simultaneously with the letter combinations. Because affixes are highly frequent letter combinations, the connections are quite strongly established.

A nice feature of this model is that it not only explains how prefixed and non-prefixed words are processed, but also suggests how pseudo-prefixed words are represented. According to Taft, pseudo-prefixed words are in the lexical level built from the morphological units they consist of (e.g. *relate* as RE and LATE), similar to the way prefixed words are connected. However, in contrast to prefixed words, pseudo-prefixed words and their affixless counterparts both have their own unit in the semantical level. Prefixed words are connected to the same semantical unit.

Although the model is described in quite some detail, the model hasn't been implemented and tested. Therefore, it remains to be seen whether the model can make correct predictions in a real language, with its linguistic categories and exceptions.

2.3.2 Hybrid Accounts

As described in Section 2.1.2, the general finding of MMP studies is that a transparent morphological relation between prime and target facilitates target recognition. A pure orthographic relation, without legitimate affix, doesn't facilitate and might even inhibit target recognition, compared to unrelated prime-target combinations. The reaction time in the opaque condition, with a legitimate affix but without a morphological prime-target relation, is generally somewhere between the other two conditions.

In the research that accompanied both IA proposals described above (Graininger et al., 1991; Taft, 1994), an opaque condition wasn't incorporated, although the sublexical account described positive recognition of pseudo-affixed words. With both models, predictions on facilitation and inhibition can be made, although these predictions remain vague. Placing the morphological level below the lexical level would result in decomposition on an orthographical basis, thus one would expect the opaque condition to be relatively similar to the transparent condition. In contrast, placing morphology above the lexical level would result in decomposition on a lexical basis, thus when a real affix is present. This would therefore predict processing in the opaque condition to resemble the orthographic condition.

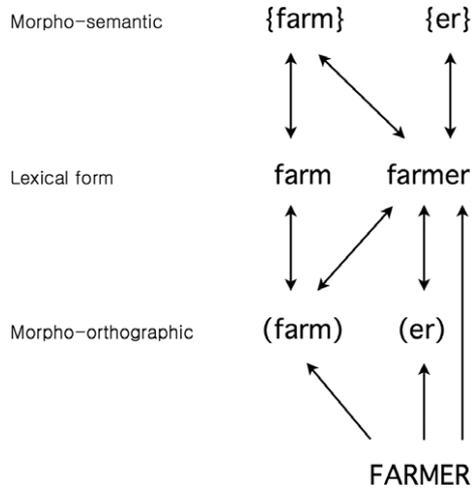


Figure 2.5: The hybrid model of morphological processing (from Diependaele et al., 2009)

The findings support neither of the models exclusively, since the opaque reaction times are somewhere between the predictions of both models. This led to the proposal of a hybrid account of morphological decomposition (Diependaele et al., 2005; Rastle and Davis, 2008). This account was drawn up in what is called the hybrid model of morphological processing (Diependaele et al., 2009) (see Figure 2.5).

The hybrid IA model incorporates both sublexical and supralexical morphological decomposition. The morpho-orthographic decomposition resembles the sublexical decomposition from Taft (1994). It reacts to the presence of the letter combination that forms the affix, but is insensitive to its semantic function. It explains why target recognition in the opaque condition is faster than in the orthographic condition as follows: The affix is stripped from a prime in the opaque and transparent conditions. This decreases activation of the lexical prime node. During prime presentation and once the target is presented, the remaining activity in the prime node inhibits the lexical target node. The lower the activation in the prime node, the lower its inhibition of the target node and the faster the target activation can reach the recognition threshold. If no (pseudo-)affix is stripped off, as in the orthographic condition, activation of the lexical prime node is higher, which results in stronger inhibition of the target node and therefore slower target recognition.

Morpho-semantic relations operate on a different level. They react to the presence of the semantics of the affix, as exists in the transparent prime. If prime and target are morphologically related, they share a semantic representation. In addition, the prime is also connected with a morpho-semantic affix unit. Upon prime presentation, positive feedback connections from activated morpho-semantic units to all morphologically related lexical forms result in excitation of the lexical target node through the shared morpho-semantic unit. This accelerates target recognition. In the transparent and opaque condition, no morpho-semantic units are shared, and no facilitatory effect from the semantic level occurs.

Like the previously proposed architectures, the hybrid model is hypothetical. It has not been implemented nor presented with a set of real stimuli. However, a hierarchical account of morphological decomposition with two mechanisms has been supported by examination of brain potentials (Lavric et al., 2012). Participants performed a lexical decision experiment with transparent, opaque en orthographic conditions, but without priming. The presented stimuli were comparable to the primes in the MMP studies (e.g. *darkness*, *corner* and *brothel*). The ERP study revealed a divergence in potentials between the orthographic condition on the one hand and both the opaque and transparent conditions on the other hand after ~ 190 ms. 60-70 ms later, the opaque and transparent conditions diverged. This supports a low-level early morpho-orthographic decomposition and a high-level later morpho-semantic decomposition.

This notion, however, is in contrast with findings in Diependaele et al. (2005). In the experiment, an instance of the MMP lexical decision experiment, different prime durations, 40 ms and 67 ms, were tested. In the orthographic condition, facilitatory effects were absent with both prime durations. In the transparent condition, facilitatory effects were present with both prime durations, although the effect was stronger in the 67 ms condition. The opaque condition also showed a facilitatory effect, but only with a 67 ms duration. The effect was absent with 40 ms prime duration. This suggests that morpho-semantic effects, perceptible only in the transparent condition, can be more quickly established and possibly kick in earlier than morpho-orthographic effects. However, one should be cautious with comparing ERP and behavioural results, especially if the design of the study differs. Nevertheless, it presents an extra factor that influences morphological processing in priming studies; something to keep in mind by performing and comparing these studies with one another.

Before moving on to the description of the IAAS model, one last matter of debate should be considered: the representations of sublexical orthographic nodes. This is an important topic, because these representations influence the way morphological nodes can be activated.

2.4 Orthographic Representations

Humans can read in various fonts, sizes and orientations on different locations on the retina. This information must be mapped from the retina on to neurons. This retinotopic mapping should normalize the input for all kinds of variations, without loss of essential information.

As mentioned in Section 2.1.4, brains possess cells that are specialized in certain simple geometric shapes (Hubel and Wiesel, 1959, 1968). These shapes, such as lines in different orientations, can be combined to reconstruct letters. The IA model incorporates such a transition from the bottom feature level to the letter level. The letters that are activated are identity and location specific. From this information, letters can be combined, to reconstruct words.

Generally, IA modellers agree that lexical nodes have positive connections to multiple sublexical nodes. The pattern of activity from these nodes results in more activation of specific word nodes than others. However, there is debate about what exactly the nodes in the sublexical level represent. A distinction can be made between position-based and context-based representations.

2.4.1 Position-Based Coding

It's important that letters are represented on the right location in the word. Otherwise, one couldn't see the difference between anagrams, such as *stop* and *spot*. However, the system turns out to be quite flexible. For example, one is able to read a sentence in which the letters of the words are mixed up.

The original IA-model uses position-based coding (McClelland and Rumelhart, 1981), also referred to as slot-based coding. A node is present for each letter on each position (see Section 2.1.4). The model only handles four-letter words, so that the total amount of letter nodes is 4×26 . The principle has been extended to words of different lengths (Coltheart et al., 1993). However, some problems arise with this.

The word *tomato* contains a typographical error. Nevertheless, most people understand quickly which word is meant. However, the word wouldn't be recognized with the orthographic representations in the IA model. The activated nodes would be T1, M2, A3, T4 and O5, whereas the lexical node *tomato* would have excitatory connections to T1, O2, M3, A4, T5 and O6. Only one node activates the correct word node.

This problem also applies to affixed words, especially prefixed words. Studies of MMP (Diependaele et al., 2009) (see Section 2.1.2) suggest that whole-word representations of (pseudo)-prefixed words and their stems share connections to the same sublexical nodes. However, using position-based coding, a word as *relate* wouldn't share any connections with *late*.

To overcome this problem, various alternatives to the left-to-right position-based coding scheme have been proposed. Examples are slot coding with the centre as reference point (Caramazza and Hillis, 1990), the closest edge as reference point (Jacobs et al., 1998) or both-end representations; two representations of each letter, once with each word edge as reference point (Glasspool and Houghton, 2005; Fischer-Baum et al., 2010, 2011). These proposals have been compared by evaluating how well each model could simulate errors of patients with dysgraphia (Fischer-Baum et al., 2010), assuming that the patterns of dysgraphia represented deficits in sublexical representations. Both-end representations could simulate the dysgraphic errors most successfully. In addition to writing, both-end coding simulated errors in a reading task with healthy participants better than the other coding schemes (Fischer-Baum et al., 2011). Therefore, this seems to be the most plausible position-based coding mechanism.

Prefixed and suffixed words and their stems always share connections when both-end coding is used. If the affix is a suffix, the left-to-right coding supplies the shared connections. If the affix is a prefix, the right-to-left coding provides them. However, simulation gets stuck with circumfix processing. A circumfix is an affix that attaches letters in front, as well as behind the word. The English language doesn't know any true circumfixes, although other Germanic languages do, such as Dutch and German. A Dutch example is the regular past-participle, which is formed with *ge+stem+d/t*. For instance, the past-participle of *gooi-en* (to throw) is *ge-gooi-d*. Another Dutch circumfix transforms a noun's meaning into a range of ...s and is formed with *ge+noun+te*, although its frequency is low. For instance, a range of mountains is *ge-berg-te* and a range of bones is *ge-been-te*.

Priming effects have been found in circumfixed words (Heide et al., 2010). These effects can't be simulated with both-end coding. A whole other ap-

proach to position-based coding is taken with local-context coding. This approach isn't influenced by positional references.

2.4.2 Local-Context Coding

Not the absolute position in the word, but information about nearby letters are used as references to the letter's location in local-context coding. The nodes in the sublexical level don't represent single letters and their positions, but combinations of letters. Various models have been proposed in this domain. For instance, a letter's location can be represented with respect to its adjacent preceding letter (Whitney, 2001; Grainger and Van Heuven, 2003). Presenting the word *blue* results in activation of nodes BL, LU and UE. These nodes are called bigrams. Another option is to represent a letter's location with respect to both its adjacent preceding and succeeding letters (Seidenberg and McClelland, 1989). Presenting the same word to the second instance results in activation of trigrams; BLU and LUE.

A third alternative is called open-bigram coding (Grainger and Van Heuven, 2003). Upon presenting a word in a model with open-bigrams, adjacent as well as non-adjacent letter combinations are activated. The nodes are called open-bigrams, because two letters can be intervened by other letters. Beforehand, it should be decided what is the maximum allowed distance between two letters to be included in the representation. This distance is called the gap. $Gap = 0$ results in inclusion of only adjacent letters. This option is similar to the closed bigram model. $Gap = 1$ adds BU and LE to the closed bigram combinations, when *blue* is presented. $Gap = 2$ includes also BE.

Several behavioural effects are successfully modelled with open-bigrams, such as transposition priming effects and relative-position priming effects (Grainger and Van Heuven, 2003). These effects seem more dependent on the direct context of a letter than on its absolute position in the word (Lété and Fayol, 2013).

However, the local-context coding approach isn't suitable to simulate every behavioural effect. The model comparison of reading error simulations that was mentioned in the previous section (Fischer-Baum et al., 2011) also included local-context models. These models were outperformed by the both-end coding model. Because both approaches outperform each other in different types of simulations, it has been suggested that both types of orthographic representations might be present.

2.4.3 Dual-Route Approach

The dual-route approach of orthographic processing (Grainger and Ziegler, 2011), distinguishes a coarse-grained and a fine-grained route from the visual stimuli to the lexical level. Both of these paths operate in parallel. The coarse-grained route is assumed to use local-context representations, such as open-bigrams, which allows quick recognition of easily distinguishable words. This pathway induces the flexibility in word recognition of misspelled words. In contrast, the fine-grained path uses position-based coding, such as both-end coding. This path is necessary in order to process words when precise information of letter order and location is needed. This might take more time, but is more precise.

The authors suggest that this fine-grained path is used in the processing of morphologically complex words. They argue that information of letter posi-

tion is needed for morpho-orthographic segmentation. Nevertheless, there is the theoretical possibility that a location independent approach could also account for the morphological priming effects. This issue reflects the second research question of this thesis. To rule the possibility without precise location information out, simulations with a position-based and a context-based coding scheme can be compared. Successful simulations with position based coding, and not with context based coding, would indicate that precise positional information is needed for affix recognition.

2.5 Summary

Interactive activation models, which are neural network models, have been proposed to explain results from MMP experiments, that are used to study how morphology is represented in the human brain.

In the context of word representations, a computational trade-off exists between storage efficiency and processing efficiency. Researchers generally agree on the presence of both morpheme representations and whole-word representations in memory. Morphologically simple words, including pseudo-affixed words, as well as highly frequent morphologically complex words have whole-word representations, whereas the meaning of other morphologically complex words can be derived from their morphemes.

Researchers disagree on how lexical and morphological representations are connected with each other and with orthographic and semantic representations. In the IA framework, morphological nodes are proposed to be placed between the orthographical and lexical layer, or between the lexical and semantical layer. Experimental results suggest that morphological representations are present close to orthographic as well as semantic representations. The hybrid model of morphological processing blends morphological information in both the orthographic and the semantic layer. Position-based coding is expected to be necessary in the sublexical level of this model to establish accurate affix representations, because of affixes are location specific.

None of the proposed models have been implemented, which makes it difficult to draw conclusions about the contents of representations and the connectivity between them. To remedy this, the next chapter introduces the IAAS model. A detailed description of the implementation will be provided, as well as simulations of MMP experiments.

Chapter 3

IAAS

The model we present in this chapter, the interactive activation affix stripping (IAAS) model, is designed to simulate human processing of morphologically complex words, as is displayed in masked morphological priming (MMP) tasks. Grainger and Ziegler (2011) state that affix detection is at the heart of morpho-orthographic decomposition, which implies that different mechanisms might be used for processing affixed words than for other morphologically complex words. Because the MMP task focuses on processing of affixed words, we could focus on a model specialized in handling affixed words.

The basis for IAAS model was the hybrid model of morphological processing from Diependaele et al. (2009) (see Figure 2.5), because this model integrated morphological information in both the sublexical and the semantical level, which corresponds to the orthographical and semantical effects in the MMP task. However, when preparing implementation, the model turned out to have a major drawback. In the hybrid model, morpho-orthographic units receive bottom-up activation from lower processing levels and spread activation to lexical nodes. Because there are no connections with semantical information, morpho-orthographical nodes should be activated on an orthographic basis. However, morphology can't be defined by simple orthographic rules. For instance, the model should be able to discriminate between *singer* and *ginger*. Although these words differ only by one letter, the first word consists of two morphemes, that should both become active, whereas the last word consists of only one morpheme. These differences in semantics between words that are orthographically quite similar make it hard to differentiate between these words on an orthographic basis. Wiring which combinations are allowed and what levels of inhibition should be used between specific morphological nodes could be implemented by training the model. That, however, is beyond the scope and purpose of this study, mainly because we attempt to explain how and why affixed words are processed in a certain way, and not just implement a model that simulates behaviour without providing information on how and why this is done.

This chapter first presents a number of requirements the model should fulfil. Then it gives an overview of the IAAS model, which is followed by a detailed description of all the model's components. The mechanisms of the model are illustrated by some examples of how the model reacts to the different priming conditions of the experiment. Next, the requirements will be tested for two

options of model design. By comparing these options, the research question about whether positional information is needed in order to process affixed words can be answered. The chapter ends with the presentation of two simulations of MMP tasks, with which we try to answer the main research question of this thesis, whether morpho-orthographic and morpho-semantic effects in human processing of (pseudo-)affixed words can be simulated by implementation of affix stripping in the IA model.

3.1 Requirements

Before implementing the model, a list was set up of the main requirements the model should fulfil. Without fulfilment of these requirements, no proper conclusions can be drawn from the simulations.

- The first requirement is that, upon presentation of (pseudo-)suffixed words, suffixes are correctly stripped. The first stage is the detection of the affix. This should be correct for both pseudo-affixed words, when the affix is recognized orthographically, and real affixed words, when the affix is recognized both orthographically and semantically. Furthermore, not too many false affixes should be recognized. The second stage is the inhibition of corresponding sublexical nodes, which starts when an affix node rises in activation above a threshold. The differences in the first stage between real and pseudo-affixed words should be reflected in the second stage by a higher number of inhibitory cycles when real-affixed words than when pseudo-affixed words are presented.
- The second requirement is that the words in the lexicon, including the stimuli that will be presented to the model in the experiment, are recognized correctly. These stimuli include morphologically complex words, from which a suffix might be stripped off.
- The last requirement is that the model can be primed. Upon presenting a prime for a small numbers of cycles, the prime node must gain sufficient activation, in order to influence the target node. However, the prime's activation shouldn't be too high, because the prime needs to be inhibited by the target during target presentation, or the prime will be recognized instead of the target.

3.2 General Outline

The IAAS model consists of IA nodes as described in Section 2.1.4. Like in the original IA model, the nodes are grouped together in layers. The configuration of the layers remains close to the original IA model, by leaving the sublexical→lexical→(morpho-)semantical path intact. Figure 3.1 depicts how these layers are connected to an affix layer. Upon presentation of a stimulus, nodes in the sublexical layer are activated (connection 1 in Figure 3.1). These nodes are comparable to letter nodes in the original IA model. Section 3.3 describes what these sublexical nodes in the IAAS model represent. The sublexical layer has excitatory connections to two other layers: to the lexical layer

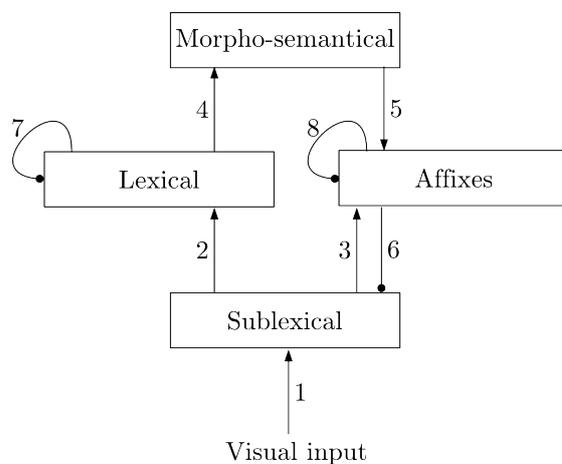


Figure 3.1: Schematic representation of the IAAS model. Arrows represent excitatory connections and lines with circle ends represent inhibitory connections. Numbers 7 and 8 represent lateral inhibition within the group of nodes.

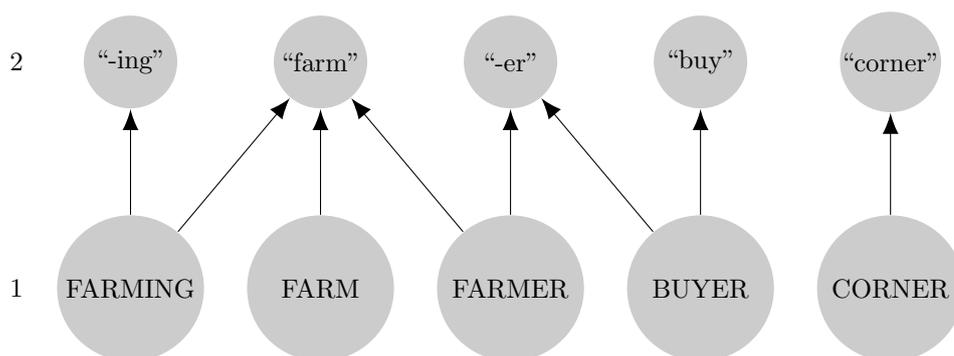


Figure 3.2: Schematic illustration of how lexical nodes (1) are connected to morpho-semantic nodes (2)

(2), which contains a lexicon of whole word forms, and to the affix layer (3), which contains a lexicon of affixes. Each node in the lexical and affix layers represents one entry in the lexicon it belongs to. Both layers are connected to the morpho-semantic layer. This morpho-semantic layer contains semantic representations of morphological families. Each node represents the meaning of a certain stem or affix. The lexical node of each word that contains a certain morpheme is connected to its morpho-semantic node. Figure 3.2 demonstrates some of these connections. Affixed word nodes, like FARMING and BUYER, are connected to morpho-semantic stem and affix nodes. Note that pseudo-affixed word nodes, like CORNER, are not connected to affix nodes, since the affix is only orthographically present, not semantically. Activation in the lexical layer spreads to the morpho-semantic layer through bottom-up connections (step 4 in Figure 3.1). Morpho-semantic affix nodes spread their activation to their corresponding nodes in the affix layer by excitatory top-down connections (5).

Affix stripping is defined as the inhibition of sublexical nodes that represent the affix. This inhibition is established by the inhibitory top-down connections between the affix and sublexical nodes (6). An affix node starts inhibiting its corresponding sublexical nodes when its activation value rises above a threshold. An affix node can be activated by two paths: a morpho-orthographic path, which runs directly from the sublexical to the affix nodes, and a morpho-semantic path, which runs to the affix nodes via the lexical and morpho-semantic level. Activation from both paths, individually or combined, might result in affix stripping: morpho-orthographic and morpho-semantic affix stripping, respectively. Section 3.5 explains in detail how affix stripping is implemented.

Like in the affix layer, nodes in the lexical layer also have a threshold. When the activation of a lexical node rises above the threshold, the model is assumed to have recognized the word this node represents. In the lexical layer, as well as in the affix layer, the threshold should preferably be reached by just one node, to avoid ambiguity in word or affix detection. To increase the differences in activation levels between nodes, the lexical and affix layer both have lateral inhibitory connections between their nodes (7,8). The inhibitory connections in the lexical layer also play a major role in the priming effect, as will be illustrated in Section 3.5. Because the model has to handle words of different lengths, lateral inhibition was altered from the way it was implemented in the original IA model. Section 3.4 explains this adaptation.

3.3 Sublexical Representations

As explained in Section 2.1.4, upon stimulus presentation in the original IA model (McClelland and Rumelhart, 1981), a set of nodes in the feature level was activated, which spread their activation onward to the letter level. For sake of simplicity, this feature level is omitted in the IAAS model, as in Grainger and Van Heuven (2003). Instead, each stimulus is transformed to a format that can be directly matched against the sublexical nodes. The nodes that are matched positively to the input receive direct activation from a virtual external source. The external input is constant for as long as the stimulus is present.

As mentioned in Section 2.4.3, Grainger and Ziegler (2011) argue that precise information about the position of letters at the sublexical level is necessary for morphological processing. However, there is a theoretical possibility that characteristics of the stimuli, such as frequencies of letters or letter combinations, already give rise to differences between the experimental conditions, without precise information of letter positions.

In order to test this, two variations of the IAAS model are implemented, with differing modes of representations in the sublexical level. The first option is both-end position (BEP) coding, in which positional information determines how activation spreads through the model (see Section 2.4.1). The second option is open-bigram (OB) coding, in which frequencies of adjacent and non-adjacent letter combinations determine how activation will spread (see Section 2.4.2).

3.3.1 Both-End Position Coding

As stated in Section 2.4.1, position-based coding as used in the original IA model is not sufficient in the length-dependent IAAS model. The letter level

Table 3.1: An example of both-end coding of affixed words

		S	E	A	R	C	H			
Left-end coding		1	2	3	4	5	6			
Right-end coding		-6	-5	-4	-3	-2	-1			
		S	E	A	R	C	H	I	N	G
Left-end coding								7	8	9
Right-end coding		-9	-8	-7	-6	-5	-4	-3	-2	-1
	R	E	S	E	A	R	C	H		
Left-end coding	1	2	3	4	5	6	7	8		
Right-end coding	-8	-7								

in the original model contains a node for each letter at each possible position. Since only four-letter words were used, the letter level consisted of 4×26 nodes. The task to simulate in the current study requires the model to handle words of different lengths. For this, the number of position nodes can be increased. The first two lines in Table 3.1 illustrate coding of the word *search* with additional left-end position units; the nodes S1, E2, A3, R4, C5 and H6 would be activated.

In order to establish morpho-orthographic affix stripping, specific nodes in the sublexical level need to spread their activation to an affix node if an affix is present. If the word *searching* is presented to a model with left-end coding, the affix is represented with the nodes I7, N8 and G9. These nodes could be connected with the affix node -ING. However, stems can differ in length. If the affix *-ing* is added to a stem of four letters, the nodes I5, N6 and G7 would be activated. This leads to an inconsistent representation of the same affix, which makes consistent activation of affix nodes impossible.

To overcome this problem, a second coding scheme is added in the IAAS model; right-end coding. Combined left- and right-end coding is called both-end coding Fischer-Baum et al. (2011). As Table 3.1 shows, the affix *-ing* is represented as I-3 (I minus 3), N-2 and G-1 with right-end coding. These nodes activate the affix node -ING. This representation is the same for every word length. In addition, the affix node -ING is only activated if the letter combination *ing* is present at the end of the word. For instance, the affix node -ING isn't activated if the word *bingo* is presented, because this words activates I-4, N-3 and G-2, which are not connected to the affix node -ING. Although only suffixes are implemented, prefixes could be added as well. If we assume that the stripping mechanism of prefixes is comparable to that of suffixes, prefixes could be represented with left-end coding, such as R1 and E2, representing the prefix *re-* in the word *research*.

The coding scheme by Fischer-Baum et al. (2011) is graded. Grading is the distribution of activation across nodes of the same letter but different positions. For example, upon presenting the word *ace* to a discrete coding scheme, not only the C at positions 2 and -2 are activated, but also C at positions 1 and 3, and -1 and -3, although less strongly. Although graded coding captures certain behavioural effects better than discrete coding (Fischer-Baum et al., 2010), grading is left out in order to reduce the noise in affix detection.

3.3.2 Open-Bigram Coding

OB coding is implemented according to the model described in Grainger and Van Heuven (2003). If a word is presented to the model, the input is converted into a list of present bigrams, of which the corresponding nodes in the sublexical layer are activated. For example, upon presenting the word *scan*, the bigrams SC, SA, SN, CA, CN and AN receive activation. The list of bigrams present in a word depends on the maximum allowed distance between two letters to form a bigram; the gap (see Section 2.4.2). The gap parameter can be set to a specific number of letters or to infinity, for which all non-adjacent letter combinations are included.

The activation of word nodes is not fundamentally different for both types of sublexical representations. For this reason, OB coding and BEP coding are interchangeable and can even be switched on simultaneously. The amount of activation that is spread from each representation mode can be altered by changing the proportion of external input to all nodes in each representation.

Executing simulations with these coding schemes had implications for the activation of lexical nodes and lateral inhibition between them. The next section describes the problems that arose and the mechanism that is implemented to solve these problems.

3.4 Masked Field Weighting

With the introduction of the IA model (McClelland and Rumelhart, 1981), only four-letter words were presented. Equation 2.2 reflected the input a certain node receives in a cycle. In all following simulations with the IAAS model, the noise is left out, for sake of simplicity. If a node in the IAAS model i is a word node, which doesn't get any external input, the input looks like this:

$$\text{input}_i = \sum_j w_{ij} \text{output}_j \quad (3.1)$$

If a word node receives its maximum amount of input, it gets activation from all nodes with which it has excitatory connections, which are the corresponding letter nodes, and no activation from inhibiting nodes. In simple word presentation, the weights of excitatory connections between a word node and its corresponding letter nodes, as well as the output from each letter node, are equal across connections. Therefore, they can be represented as one variable; the excitation from a corresponding letter node j to word node i , excitation_{ij} . Note that the excitation is still dependent on the output values of the letter nodes, which changes between cycles, but is equal across nodes within each cycle. Consequently, the maximum amount of bottom-up excitation a word node can receive is dependent of word length, length_i :

$$\text{input}_{i,\max} = \text{length}_i \text{excitation}_{ij} \quad (3.2)$$

If only four-letter words are presented, the maximum input is equal for each letter node. However, if words of different lengths are presented, the maximum excitation is higher for longer words than for shorter words. Moreover, words in which a shorter word is embedded receive the same amount of bottom-up excitation upon presenting the shorter word. For example, *cow* is embedded

in *coward*. Upon presenting *cow* to the original IA model, both lexical nodes COWARD and COW receive activation from three letter nodes: C1, O2 and W3. In this instance, COW should become higher in activation, in order to reach the threshold before COWARD. To establish this, the letter to word excitation can be normalized by dividing the summed activation a word node receives by the number of excitatory connections between the sublexical level and this word node, which is equal to the word length, length_i , in the original IA model:

$$\text{input}_{i,\text{max}} = \text{length}_i \text{excitation}_{ij} / \text{length}_i = \text{excitation}_{ij} \quad (3.3)$$

The word node COW would have 3 letter connections and COWARD 6. In this way, COW will get $3/3 = 1 * \text{excitation}_{ij}$ and COWARD only $3/6 = 0.5 * \text{excitation}_{ij}$. This has a drawback, because, upon presenting the longer word, *coward*, the word node COWARD receives $6/6 = 1 * \text{excitation}_{ij}$, but also COW gets $3/3 = 1 * \text{excitation}_{ij}$.

One way to control for this, is by adding bottom-up inhibitory connections between letter nodes and their negative matching word nodes. In this way, COW would get inhibitory activations from A4, R5 and D6. These inhibitory connections were already present in the original IA model. A problem with this, however, is that these connections can't be normalized very easily. On one hand, the inhibition should be strong enough to correct for word length, but on the other hand, it shouldn't be too strong in order to establish a priming effect, for which the target, the shorter word, should be activated as well. Especially when bigram coding is used, small differences between prime and target can change the proportion of shared bigrams a lot. For example, presentation of *corner* leads to 6 excitatory and 9 inhibitory connections between the bigrams and the word node CORN, whereas *pigment* leads to 3 excitatory and 18 inhibitory connections between bigrams and PIG. Therefore, no constant normalization value, nor a value that is dependent on the number of bottom-up connections can be applied.

Another solution is the implementation of masked field weighting (MFW) (Davis, 2010). This solution isn't dependent on the connections with the sublexical layer and therefore isn't dependent on the coding scheme, so that both-end coding and bigram coding can be used interchangeably. Instead, MFW is dependent on word length. The masked field weight adjusts lateral inhibitory values between the word nodes corresponding to their word length. In this way, upon receiving the same amount of bottom-up excitation, inhibition from longer to shorter words is stronger than in the opposite direction. The inhibitory weight from word node i to word node j is multiplied by the masked field weight, m_i , which is calculated as follows:

$$m_i = 1 + (\text{len}_i - \text{len}_{m=1})w_m \quad (3.4)$$

in which $\text{len}_{m=1}$ is the length of the word for which m_i is 1; the length that keeps its original weight. Davis (2010) did set this default length to 4 and w_m to 0.35, so that a seven-letter word inhibits approximately twice as strongly as a four-letter word.

In the IAAS model, no inhibitory bottom-up connections are implemented, because of the difficulty of normalization. Instead, masked field weighting is implemented. $\text{len}_{m=1}$ was set to 4, like in Davis (2010). However, w_m was set to 0.1, because higher values, like the original value of 0.35, were too strong for

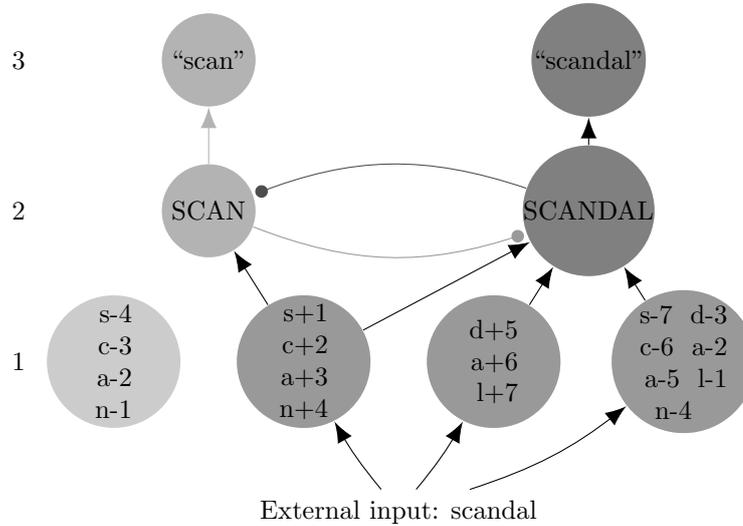


Figure 3.3: Schematic illustration of activation flow upon presentation of a non-affixed word. 1: BEP nodes, 2: lexical nodes, 3: morpho-semantic nodes. The darker, the stronger the activation.

prime-target word pairs with a large difference in length, like *pig* and *pigment* (opaque) or *work* and *workable* (transparent). In these pairs, the shorter target node was too much inhibited by the long prime node during and after priming, which impaired target recognition.

Figure 3.3 illustrates how the lexical nodes SCAN and SCANDAL are activated when *scandal*, a word without affix that is used as prime in the orthographic condition in the MMP task, is presented. The external input activates the BEP nodes (1) that correspond to the input. These nodes activate lexical nodes (2). The maximum amount of activation spreads to SCANDAL, and also some activation spreads to SCAN. Lateral inhibition between SCAN and SCANDAL, reinforced by MFW, increases the difference between both nodes, so that SCANDAL will be the node that eventually will cross the recognition threshold. *Scan* and *scandal* are semantically unrelated, so they both have their own morpho-semantic node (3).

When presenting orthographic prime-target pairs, without (pseudo-)affixes, affix nodes aren't involved. The next section describes how affixes are detected and stripped, and how activation flows through the model in the opaque and transparent conditions.

3.5 Affix Stripping

As mentioned above, affix stripping in the IAAS model is the inhibition of sublexical nodes by an activated affix node. The affix nodes represent chunks, which are higher-level representations of letter combinations. Theoretically, the frequent co-occurrence of these letter combinations and their meanings would lead to the establishment of specific nodes connecting the letter combinations directly to their meaning in the morpho-semantic level. In the IAAS model,

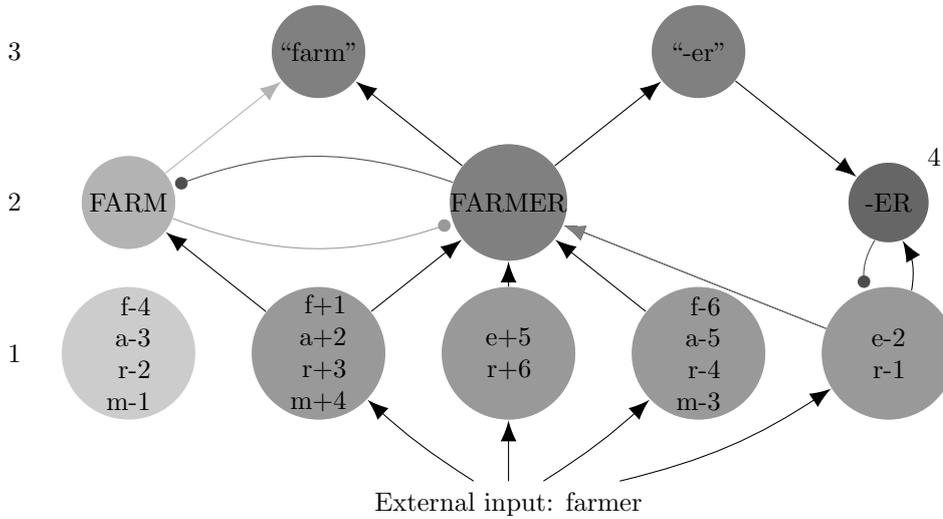


Figure 3.4: Schematic illustration of activation flow upon presentation of a real-affixed word. 1: BEP nodes, 2: lexical nodes, 3: morpho-semantic nodes, 4: affix node. The darker, the stronger the activation.

the lexicon of affix nodes is fixed. Morpho-orthographic and morpho-semantic affix stripping use the same inhibitory connections. Therefore, both types of affix stripping are interconnected. The difference between them is the way in which affix nodes are activated. Morpho-orthographic affix stripping reacts to the orthographic presence of the affix. The affix node is activated directly by its corresponding sublexical nodes. In real- as well as pseudo-affixed word recognition, the affix node is activated through this path. Morpho-semantic affix stripping reacts to the semantic presence of the affix. Activation spreads from the lexical level, through the morpho-semantic affix node to the affix node. The morpho-semantic affix node is only activated when real affixed words are presented.

Figure 3.4 gives a simplified illustration of how nodes are activated when a real-affixed word, in this case *farmer*, is presented. This is an example of a prime in the transparent condition in the MMP task. Similar to a word without an affix, sublexical nodes (1) activate whole-word representations (2). *Farmer* is derived from *farm*, so both word nodes have connections to the same morpho-semantic affix node (3). In addition FARMER, is connected to the morpho-semantic node “-er”. This morpho-semantic node and the sublexical nodes E-2 and R-1 are all connected to the affix node -ER. When activated, they spread their activation to this affix node.

Figure 3.5 shows the activation levels of all nodes when the prime-target pair *farmer-farm* is presented. It takes a couple of cycles to bring the activation of the affix node above the threshold. When more affix nodes get activated, they compete with each other through lateral inhibition, just like lexical nodes, including the same mechanisms of normalization and MFW (see Section 3.4). For affix nodes, $\text{len}_{m=1}$ was set to 2, and w_m to 0.35. When the threshold is reached, -ER starts inhibiting E-2 and R-1. This decreases the bottom-up activation from the two sublexical nodes to FARMER. Note that the nodes E5

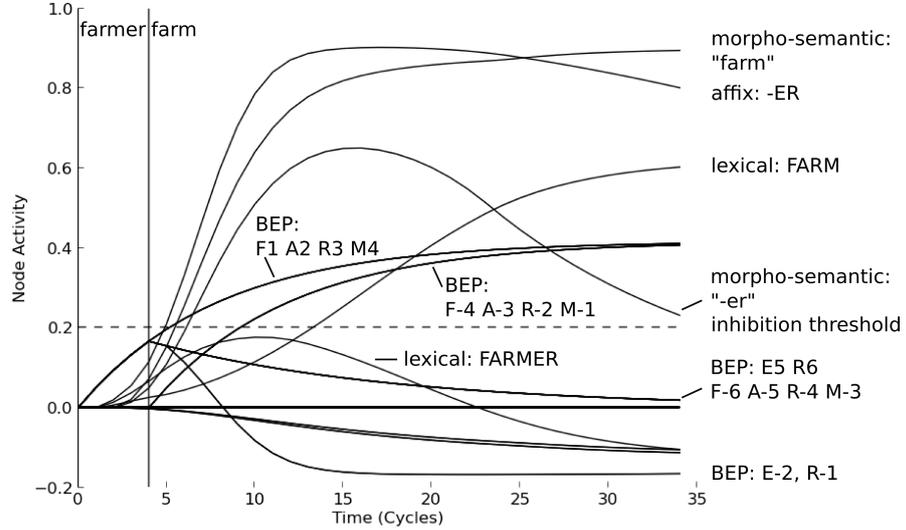


Figure 3.5: Node activity during presentation of the prime-target pair *farmer-farm*, for respectively 4 and 30 cycles.

and R6 are not inhibited and still activate FARMER with other uninhibited nodes. Therefore, as long as *farmer* is presented, the amount of activation that FARMER receives will always be higher than FARM, even if E-2 and R-1 are fully inhibited.

FARM and FARMER inhibit each other. Because the activation from FARMER decreases, the inhibition from FARMER to FARM decreases as well. This leads to an increased activation of FARM. In a simulation of the MMP task, this change in lateral inhibition between the word nodes results in the priming effect. If the prime *farmer* is presented, the activation of FARMER is lower compared to words of which no affix is stripped off. Because FARM is less strongly inhibited by FARMER, its activation is higher. This decreased lateral inhibition and increased activation of FARM continues if the target *farm* is presented.

The increased early activation of FARM leads to faster word recognition, because its activation reaches the recognition threshold earlier. In addition, FARM has more power to inhibit FARMER. This decreases the inhibition of the remaining activation in FARMER, and therefore increases the activation of FARM even more. In this way continues, even after the prime is presented, the remaining activation in the prime node to influence target recognition.

During target presentation, prime activation drops, because of its inherent decay, its decreased bottom-up activation and its increased lateral inhibition from the target word node. Therefore, priming is most strongly affected by affix stripping in the beginning of the trial. To increase the size of the priming effect, the affix node should start to inhibit early in the trial. This means that the activation threshold of the affix nodes should be reached quickly. However, affix nodes need a couple of cycles to compete through lateral inhibition, so that the threshold is reached only by the correct node. This displays a trade-off between

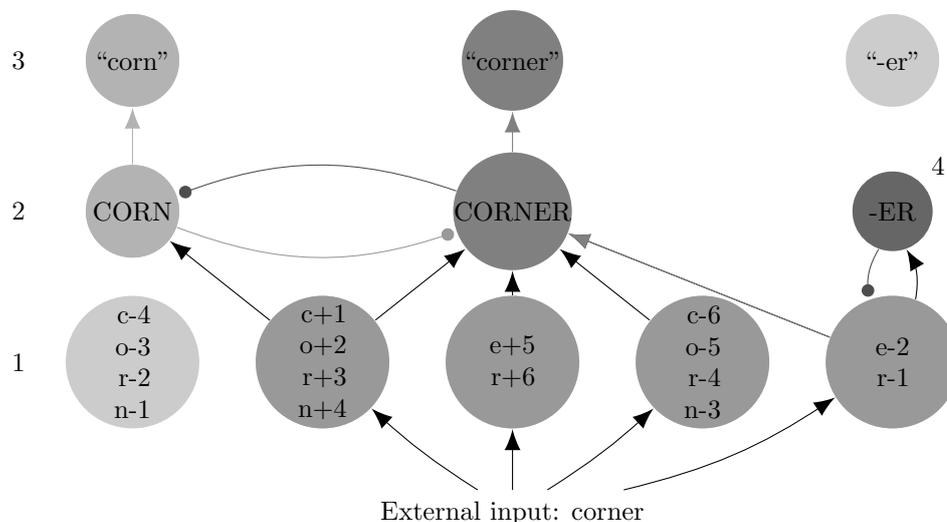


Figure 3.6: Schematic illustration of activation flow upon presentation of a pseudo-affixed word. 1: BEP nodes, 2: lexical nodes, 3: morpho-semantic nodes, 4: affix node. The darker, the stronger the activation.

speed and accuracy of affix identification: the faster an affix node can reach its inhibition threshold, the larger the probability that an incorrect affix (also) reaches the threshold. The optimal balance between speed and accuracy can be established by systematic parameter testing.

Morpho-semantic affix stripping continues as long as the morpho-semantic affix node sufficiently activates the affix node. This generally means that the affix is inhibited from the moment the affix node becomes active until the word is recognized. In pseudo-affixed words, affix stripping isn't that long-lasting.

Figure 3.6 gives a simplified illustration of how nodes are activated when a pseudo-affixed word from the opaque condition, in this case *corner*, is presented. Similar to a word without an affix, sublexical nodes (1) activate whole-word representations (2). *Corner* isn't derived from *corn*, so both word nodes have connections to different morpho-semantic affix nodes (3). *Corner* ends with *-er*, which isn't a real affix. Therefore, CORNER isn't connected to the morpho-semantic node “-er”. However, the sublexical nodes E-2 and R-1 are connected to the affix node -ER. They activate the affix node on an orthographic basis.

There are two reasons why affix stripping isn't as long-lasting in pseudo-affixed words, compared to real-affixed words. First, a few more cycles might be needed for the affix node to reach its inhibition threshold, because there is only bottom-up activation of the affix node, and no top-down activation. Second, the inhibition might stop after a couple of cycles, because the affix node inhibits its own and only source of activation. Due to the decay of activation in a node, the activation level in the affix node will drop below its inhibition threshold. This also drops the inhibition of sublexical nodes, after which these nodes can again gain activation and activate the affix nodes. Figure 3.7 illustrates how these periods with and without sublexical inhibition alternate. However, before the ending of the first inhibition period, a word is generally already recognized.

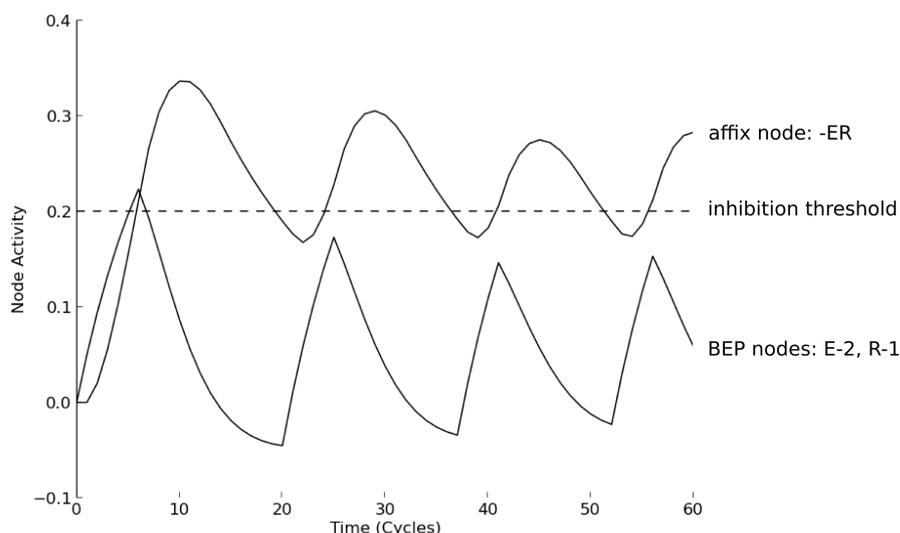


Figure 3.7: Node activity during presentation of the pseudo-affixed word *corner*. As long as the activation of affix node -ER rises above the inhibition threshold, it inhibits its corresponding BEP nodes; E-2 and R-1.

3.6 Requirement Evaluation with Both-End Coding

The requirements that are stated in the beginning of this chapter are evaluated many times during the process of implementation. The evaluations that are described below are executed with the final parameter settings of the IAAS model. Table 3.2 gives an overview of these values.

3.6.1 Suffixes are stripped correctly

To test whether suffixes are stripped correctly, real affixed and pseudo-affixed words, as well as non-affixed words, were presented to the model. This was simple word presentation without priming, lasting 30 cycles. For each cycle was stored which affixes' activations have risen above the inhibition threshold.

A total of 222 words were presented, 74 of each prime type; orthographic, opaque and transparent. For orthographic, non-affixed primes, 43 of 74 words were correctly detected as non-affixed, but in 31 trials, an affix node did rise above the inhibition threshold. In a number of these trials, an affix node could reach the threshold because it partially overlapped with the input (e.g. -ING in *china*, -LY in *single*). However, a number of recognitions were in fact correct recognitions, even though this was the condition with non-affixed words. These words were non-affixed if the target was considered to be the prime minus the removed letters. For example, one prime-target pair is *textile-text*. *textile* minus *text* is *-ile*, which isn't an affix. However, an orthographic affix is present; *-le*. The model doesn't take into consideration which word might be presented as

Table 3.2: Parameter settings in the IAAS model.

Parameter name	Equation	Value
Node parameters		
\min_i	2.1	-0.20
\max_i	2.1	1.00
rest_i	2.1	0.00
decay_i	2.1	-0.07
σ Noise	2.2	0.00
Lexical threshold		0.40
Affix threshold		0.20
Connection parameters		
extinput_i	2.2	0.05
w_{ij} Sublex. - Lexical	2.2	0.28/(No. Exc. Conn.*)
w_{ij} Sublex. - Affix	2.2	0.40/(No. Exc. Conn.*)
w_{ij} Lexical - Morpho-sem.	2.2	1.00
w_{ij} Morpho-sem. - Affix	2.2	1.00
w_{ij} Affix - Sublex.	2.2	-0.40
w_m Lexical - Lexical	3.4	-0.21
m_i Lexical	3.4	0.1
$\text{len}_{m=1}$ Lexical	3.4	4
w_m Affix - Affix	3.4	-0.60
m_i Affix	3.4	0.35
$\text{len}_{m=1}$ Affix	3.4	2

*Normalization for number of excitatory bottom-up connections. This is the number of sublexical units to represent the word or affix.

target, so the affix node -LE is activated like a pseudo-affix. Other examples of trials like these were -Y in *galaxy* (target is *gala*), -ER in *beer* (target is *bee*).

For the 148 real and pseudo-affixed words, all affixes were detected correctly. However, in 58 trials, more than one affix node reached the threshold (e.g. -Y and -LY in *mainly*). This is a consequence of the speed-accuracy trade-off in affix detection, described in Section 3.5. The inhibition threshold is put quite low, in order to detect an affix quickly. If the threshold was higher, the affix nodes would have had more time to inhibit each other, so that only the longest present affix would remain. However, affix detection would then be too slow to establish the priming effect. In some cases, this sloppy detection lead in fact to activation of the correct affix. These were words from which the last letter of the stem could be the first letter of the affix (e.g. *armoury* activated -Y correctly, but also -RY). The longest activated affix in this case isn't the real affix. An extreme case of multiple affix detection is the pseudo-affixed word *factory*, which in the experiment is presented with *fact* as target. This word activated the affix nodes -ORY, -RY and -Y. These could all be correct pseudo-affixes, since *fact*, but also *facto* (in *de facto*) and *factor*, could be presented as target words.

Finally, some affix nodes that had no orthographic overlap with the presented stimuli were activated. This was a false activation through the lexical and morpho-semantic path. For example, upon presenting the word *hungry*, the affix nodes -Y, -RY and -ER reached the threshold. The sublexical units activated HUNGRY, but also, in smaller amount, HUNTER. This activated the morpho-semantic affix node “-er”. Another example was *poster*, which also activated -AL, since POSTAL, which is connected to the morpho-semantic node “-al”, was also activated. For most of these false affixes, their sublexical inhibition was short lived, since the correct lexical node inhibits the incorrect nodes, and therefore the activation in the morpho-semantic affix node decays after a couple of cycles.

The average number of cycles with sublexical inhibition per trial was 8.12 for the words that will be presented orthographic primes, 23.30 for opaque primes and 25.05 for transparent primes. This corresponds to the requirement for the second stage of affix stripping that orthographic primes lead to the lowest number of cycles with inhibition and transparent primes to the highest.

Although a number of false or double affixes were detected, these were considered to be natural consequences of the words in the lexicon and of the choice to give priority to fast affix detection at the cost of accurate detection. Because of this, in combination with the correct pattern in number of cycles with inhibition, the requirement is regarded sufficiently fulfilled.

3.6.2 Words are recognized correctly

In order to test this requirement, each word in the model's lexicon was presented to the model. This lexicon includes the stimuli that are presented during the simulations, primes and targets, as well as the 1000 most frequent words in the online CELEX database (Max Planck Institute for Psycholinguistics, 2011). There was no priming included for testing this requirement. A recognition was considered correct when the activation from the correct word node would reach the recognition threshold first. If another node would reach the threshold first, the response was incorrect. In addition, if no word node was recognized within 30 cycles, recognition was also negative.

A total of 1666 words were presented. Most words were recognized correctly within the presentation time ($M=16.92$, $SD=1.84$). One word was recognized incorrectly; *industrialization*. This is also the longest word that was presented. The given response was *in*. Two words were not recognized within the allowed response time of 30 cycles; *arbitrary* and *injection*. These three words won't be presented to the model as stimuli during following simulations. Because of the low number of incorrect recognitions, and because these words won't be presented, the requirement is regarded sufficiently fulfilled.

3.6.3 The model can be primed

This requirement is tested by presenting prime-target combinations to the model. Of each prime target pair was measured whether the prime node received sufficient activation in order to influence the target node. By changing the prime duration systematically, the value for which the prime has maximum influence over the target without preventing the target from being recognized was investigated.

A total of 222 prim-target word pairs were presented. The longest prime duration for which each target was recognized within 30 cycles from target presentation was 5 cycles. Figure 3.8 shows an example of activations over time from twelve prime nodes (ending below zero) and their corresponding target nodes (ending above the recognition threshold of 0.4). The prime activity over all words was a bit lower than in this figure, due to the lack of inhibition from a large lexicon in the example, but still high enough to influence target nodes (maximum prime activity: $M=0.179$, $SD=0.035$).

With these results, also the third requirement is fulfilled. This makes BEP coding a scheme that can be used to simulate MMP tasks. To check whether a less precise position coding scheme would also be feasible, the requirements were also evaluated for OB coding.

3.7 Requirement Evaluation with Open-Bigram Coding

One of the research questions of this thesis is whether precise positional specific information is needed for affix stripping. This question was answered in an early stage of implementing, when an attempt was made to implement affix stripping with open-bigram coding. The dynamics of open-bigram affix stripping are largely the same as orthographic affix stripping with both-end coding, as described in Section 3.5. However, affixes are represented with one or more bigrams instead of right-end positioned letters. Only bigram nodes of which both letters were present in the bigram, were connected to affix nodes (e.g. the affix node -ORY was connected to the bigrams OR, OY and RY). If a word was presented, activated bigrams spread their activation to the affixes that contained these bigrams. A large sample of affixed as well as non-affixed words were presented to the model over many implementation iterations. These were the targets as well as the primes of the behavioural experiment simulated in this study. Single words were presented, as well as prime-target combinations, to focus on the priming effect. The main findings during this process, encountered by evaluation of the requirements summed up in Section 3.1, are as follows:

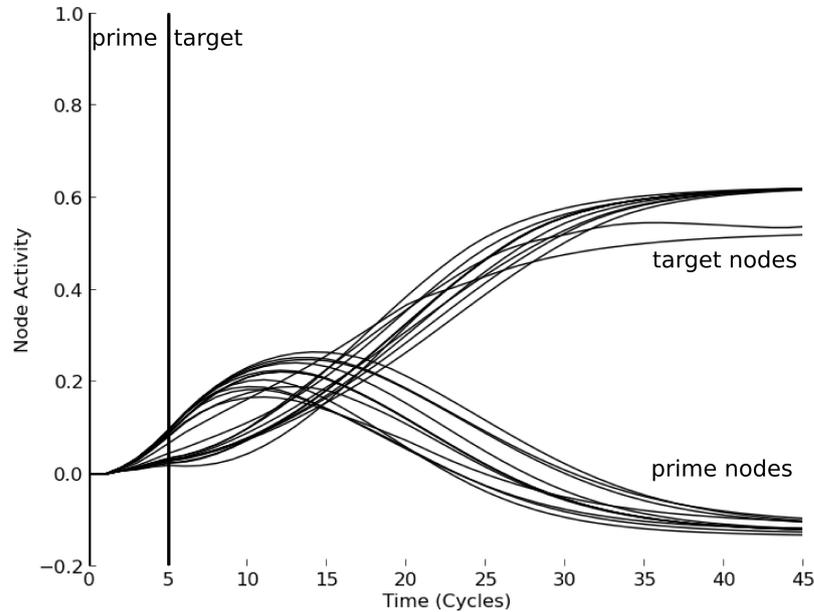


Figure 3.8: Prime and target node activity during presentation of twelve word-pairs.

The requirement that the presented words are recognized correctly, the second requirement, was satisfied. Upon presenting a morphological complex word, such as *farmer*, this word is recognized and not its affixless counterpart, *farm*. Although its affix is inhibited, only those bigrams are inhibited in which both letters are present in the affix, in this case ER. The bigrams that consist of one letter in the stem and one in the affix, such as RE, ME and MR, are not inhibited. This increases the activation of the word node FARMER above FARM.

The requirement that the model can be primed, the third requirement, was also satisfied. The target node was activated during prime presentation, because of the shared bigrams between prime and target. In addition, there was lateral inhibition between words, which caused orthographically related primes to inhibit the targets stronger than unrelated primes.

However, the requirement that suffixes are correctly stripped, the first requirement, raised a number of problems:

- Present affixes are recognized correctly. However, a lot of false affixes are recognized as well. These are typically letter combinations that are, adjacent or non-adjacent, present somewhere else in the word (e.g. *-er* in *bear* and *-ed* in *edit*).
- When the gap limit is set to infinity, the number of open bigrams increases with $n - 1$ additional bigrams for every additional letter in a word or affix. A two-letter affix contains one bigram, whereas a four-letter affix contains already six bigrams. In this way, the proportion of bigrams that

is inhibited differs a lot between trials. For instance, the word *shoulder*, a prime from the opaque condition, contains 28 bigrams. One of these bigrams is connected to the affix node -ER, and will be inhibited. The word *finish*, also an opaque prime, contains 15 bigrams. The affix node -ISH inhibits 3 bigrams, so that 1 in 5 active bigrams is inhibited. Parameter values couldn't be balanced in such a way that for every prime-target combination a priming effect could be established and the target could be recognized correctly as well.

- There is an affix that consists of only one letter and therefore cannot be represented by a bigram: *-y* (e.g. *creamy*, *lucky* and *snowy*).

To overcome the last problem, an extra character was added after each word, which was also included in the bigrams to mark the end of the word. The affix *-y* was now represented with the bigram *-Y#*. However, this changed the proportions of inhibited bigrams, and therefore the size of priming effects.

A solution to both the first two problems could be to decrease the gap limit. This decreases the chance that an affix is constructed from letters that are spread out in the word. In addition, the number of additional OBs for an increased word length decreases, which flattens the differences in proportion of inhibited bigrams. However, even if only adjacent bigrams were allowed, still a lot of false affixes were detected (e.g. 'al' in *stalk*, 'ic' in *nick* and 'et' in *ether*). Furthermore, problems to establish correct recognition and a priming effect for each prime-target pair remained to some extent.

The conclusion of the simulations is that affix stripping couldn't be implemented with open bigram coding, because too many letter combinations were falsely recognized as affixes. This suggests that a location specific coding scheme is necessary for affix stripping. Therefore, further adaptations to the model were performed with BEP coding. Open bigram coding is still an optional function in the model, although OB nodes are not connected any more to affix nodes. These nodes are switched off during the following the simulations. We can now look into the simulation studies that were performed with this model.

3.8 Simulation Studies

This section describes the final two simulations that were executed with the full model, as described earlier in this chapter. Different sets of stimuli were presented in both simulations, in which all other variables were kept constant (see Table 3.2 for the parameter values). For both simulations, the model's lexicon included all words that are presented in the testing phase, as well as an additional lexicon of the thousand most frequent words in the online CELEX database (Max Planck Institute for Psycholinguistics, 2011) to obtain more realistic lateral inhibition at the lexical level. Both stimuli sets reflect the stimuli that were presented in two independently executed behavioural studies. The stimuli that were used for Simulation 1 were obtained from Beyersmann et al. (2012a) and the stimuli that are used for Simulation 2 from Beyersmann et al. (2012b). Both these experiments did find a difference between the orthographic condition and the other conditions, reflecting that the orthographic presence of the affix influences priming. However, a noticeable difference between the studies is that the first study did find a difference between the opaque and

transparent condition, reflecting that the semantic presence of the affix influences priming, whereas the second study didn't find a significant difference. Trials with non-word stimuli were left out in both simulations.

3.8.1 Procedure

As long as the visual input is presented, the sublexical nodes that correspond to the visual input receive external activation. The prime was presented for 5 cycles, which corresponds to the priming time from the behavioural experiment of 50 ms. Immediately after prime presentation, without resetting any value in the model, the target was presented for 40 cycles. The only thing that changed in this transition, was that different nodes, corresponding to the target stimulus, received activation.

To simulate lexical decision, a threshold was set for the activation in lexical nodes at 0.4. If a node crosses this threshold, the word this node represents is the given response and the number of cycles until the threshold is crossed is the reaction time. This roughly corresponds to the reaction time in the MMP task, although additional processes, such as stimulus processing before and in the primary visual cortex and the initiation and execution of the motor response, are not reflected in the model's reaction time. The original IA model contains a feature level below the letter level. Because this level is left out in the IAAS model, to reduce complexity of the model, the reaction time decreases even more. If within 40 cycles of target presentation no word node activation had crossed the threshold, this was considered a negative lexical decision or a non-response.

To measure the priming effect, a control condition is also simulated, in which the prime and target are semantically and orthographically unrelated. The priming effect is the difference in priming between related and unrelated primes. An interaction of item type (orthographic vs. opaque vs. transparent) and prime type (related vs. unrelated) means a difference of priming effects between conditions. In addition, the priming effect can also be obtained per target, because the same target stimuli were presented in the related priming and unrelated priming condition. To calculate this priming effect, the reaction time of the trial with the related prime (e.g. *farmer-farm*) is subtracted from the corresponding trial with the unrelated prime (*stormy-farm*).

3.8.2 Results Simulation 1

All stimuli were recognized correctly by the model. Figure 3.9 shows the means and standard deviations of the experimental data and the data from the model.

The main analysis in the experimental data was a generalized linear mixed-effects model, with random intercepts for subjects and items, on the latency data (Beyersmann et al., 2012a). This analysis showed significant interactions between prime type and item type for all combinations of the prime type; orthographic-opaque, orthographic-transparent and opaque-transparent. The experimental study also reports two ANOVAs; one by subject (collapsed across item data) and one by item (collapsed across subject data). The latter corresponds to the simulation, with the average per item across participants. Although the by subject ANOVA did find significant effects of item type, the by item ANOVA didn't show a main effect of item type. However, both ANOVAs

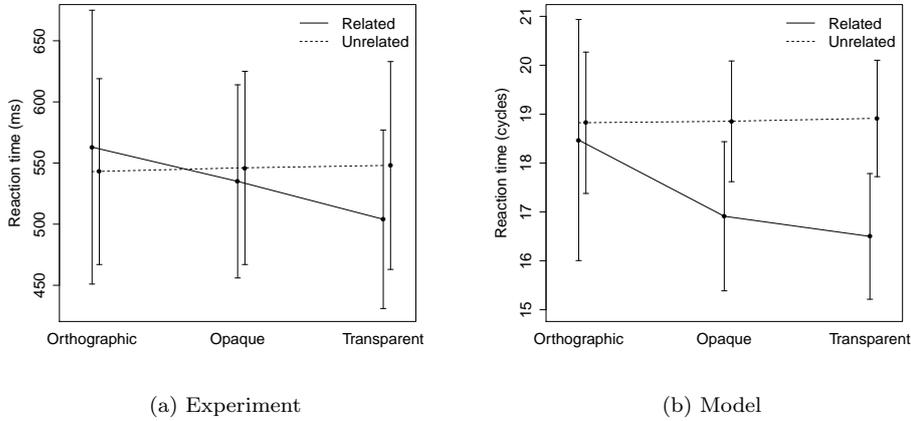


Figure 3.9: Means and standard deviations of reaction times from Beyersmann et al. (2012a) and the IAAS model to the same set of stimuli.

did show significant interaction effects between prime type and item type and significant main effects of prime type. No post-hoc analyses are reported.

There are no subjects included in the simulation. Instead, the response per item in the simulation is assumed to resemble the average across subjects. This makes the linear mixed-effects analysis unsuited for the data from the model, as well as the ANOVA by subject. Therefore, a type III item type by prime type ANOVA is performed to resemble the by item ANOVA from the experimental data. The ANOVA showed a significant interaction effect ($F(2, 198) = 34.92, p < 0.001$). In addition, the main effect for item type was significant ($F(2, 198) = 13.32, p < 0.001$), as well as the main effect for prime type ($F(1, 198) = 67.25, p < 0.001$).

In addition, a one-way item type ANOVA of the pre-calculated priming effect was also significant ($F(2, 99) = 14.42, p < 0.001$). With this ANOVA, post-hoc comparisons can be performed on the priming effects. A Tukey's HSD test revealed that the priming effect in the orthographic condition was significantly smaller than in the opaque condition ($p < 0.001$) and in the transparent condition ($p < 0.001$). However, the difference between the opaque and transparent condition wasn't significant ($p = 0.392$).

3.8.3 Results Simulation 2

All stimuli were recognized correctly by the model. Figure 3.10 shows the means and standard deviations of the experimental data and the data from the model.

For the second data set, the only reported analysis in the article was a linear mixed-effects model (Beyersmann et al., 2012b), comparable to the analysis of the first data set. This model revealed significant differences between the orthographic and opaque condition and the orthographic and transparent condition. However, the difference between the opaque and transparent condition wasn't significant.

A one-way ANOVA on the priming effects in the simulation data showed a significant effect for item type ($F(2, 117) = 4.95, p = 0.0097$). A Tukey's HSD

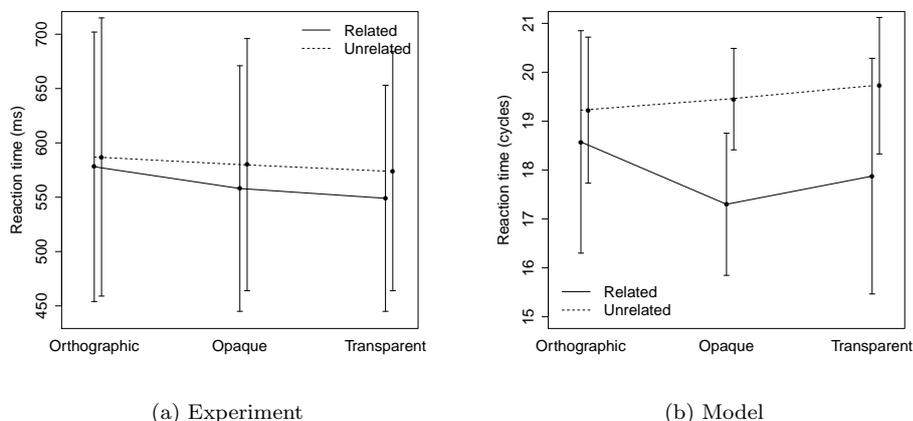


Figure 3.10: Means and standard deviations of reaction times from Beyersmann et al. (2012b) and the IAAS model to the same set of stimuli.

test showed that the priming effect in the orthographic condition is significantly smaller than in the opaque condition ($p = 0.010$) and in the transparent condition ($p = 0.049$). Like in the first simulation, the difference between the opaque and transparent condition wasn't significant ($p = 0.823$).

3.8.4 Discussion

The significant interaction effect between prime type and item type in the ANOVA by item, performed in Beyersmann et al. (2012a), is successfully simulated with the IAAS model. The same interaction effect is also simulated with the set of stimuli from Beyersmann et al. (2012b), although no ANOVA was reported for the behavioural data in that study.

A linear mixed-model analysis did show a significant difference in priming effects between the opaque and the transparent condition in the behavioural study of Beyersmann et al. (2012a). No such difference was found in Beyersmann et al. (2012b). This is in line with the conclusions from other MMP experiments that the difference in priming between affixed and pseudo-affixed words isn't as robust as the difference in priming between non-affixed and real or pseudo-affixed primes (Rastle and Davis, 2008).

The simulations suggest that morpho-orthographic affix stripping establishes a larger priming effect in the opaque and transparent conditions than in the orthographic condition. In addition, morpho-semantic affix stripping didn't capture the significant difference between the opaque and the transparent condition in any of the simulations, although this effect was reported in the behavioural study corresponding to the first set of stimuli. To determine what effects the affix stripping mechanisms exactly have on the results, additional simulations should be performed with alternative versions of the IAAS model.

3.9 Summary

This chapter introduced the IAAS model. The presentation of an affixed word to this model leads to the activation of an affix node, which inhibits its corresponding sublexical nodes. This decreases the activation of the word node at the lexical level and in that way increases the priming effect in the MMP task.

To establish a morpho-orthographic affix stripping mechanism, affix nodes are activated directly by sublexical nodes if the affix is present orthographically, independent of its semantic function. Positional information at the sublexical level is necessary in order to correctly activate affix nodes along the orthographical path. Therefore, the IAAS model is implemented with both-end position coding at its sublexical level, instead of open bigram coding. Upon presentation of real affixed words to the IAAS model, affix nodes are, in addition to activation by sublexical nodes, also activated by morpho-semantic nodes, which establishes morpho-semantic affix stripping.

The IAAS model has demonstrated to successfully simulate a larger priming effect in real and pseudo-affixed words than in non-affixed words, which corresponds to the effects reported in behavioural studies that were performed with the same sets of stimuli. One of the behavioural studies also reported a larger priming effect with real affixed than with pseudo-affixed primes. The other study didn't find a significant difference. However, this effect wasn't successfully simulated.

To investigate what effect morpho-semantic, as well as morpho-orthographic affix stripping exactly has on the priming effects, additional simulations have been executed, which are described in the following chapter.

Chapter 4

Exploration of the Model

Performing the simulations in the previous chapter raised a number of questions. The model is designed around the idea that two mechanisms, morpho-orthographic and morpho-semantic affix activation, cooperate in affix stripping. Morpho-orthographic affix stripping was reflected well in these simulations. However, morpho-semantic affix stripping didn't seem to establish the larger priming effect in the transparent condition than in the opaque condition that was reported for the first set of stimuli (Beyersmann et al., 2012a). This effect wasn't reported in the article from which the second set of stimuli was simulated (Beyersmann et al., 2012b). In order to look at what morpho-semantic affix stripping contributes to the simulation results, a simulation without morpho-semantic affix stripping was executed.

4.1 Simulation without Morpho-Semantic Affix Stripping

In order to remove morpho-semantic affix stripping from the IAAS, the excitation parameter from morpho-semantic nodes to affix nodes was set to zero (see Figure 4.1). In this way, morpho-semantic affix activation was turned off. All other parameters were kept equal to the values in Table 3.2.

4.1.1 Results Simulations

In the simulations with both sets of stimuli, all stimuli were recognized correctly by the model. Figure 4.2 shows the means and standard deviations of both simulations.

A one-factor ANOVA of the priming effect (unrelated minus related) did reveal a significant effect for item type (orthographic, opaque or transparent) in the first set of stimuli ($F(2, 99) = 12.94, p < 0.001$). Post-hoc comparisons with the Tukey HSD test showed a significant smaller priming effect in the orthographic condition than in the opaque condition ($p = 0.006$) and the transparent condition ($p < 0.001$). The comparison between the opaque and transparent condition wasn't significant ($p = 0.176$).

A direct comparison between this model and the full IAAS model (see Section 3.8.2) for the first set of stimuli with a two-way priming type by model

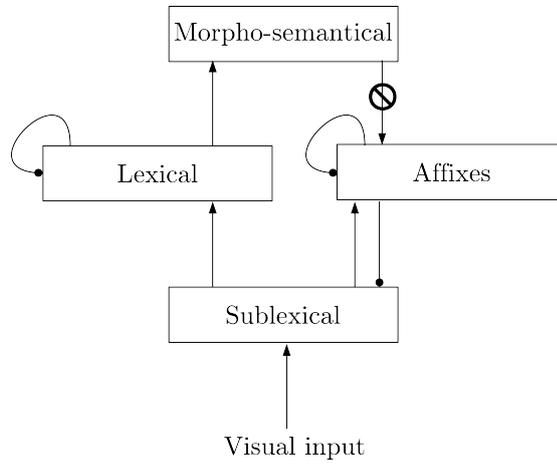
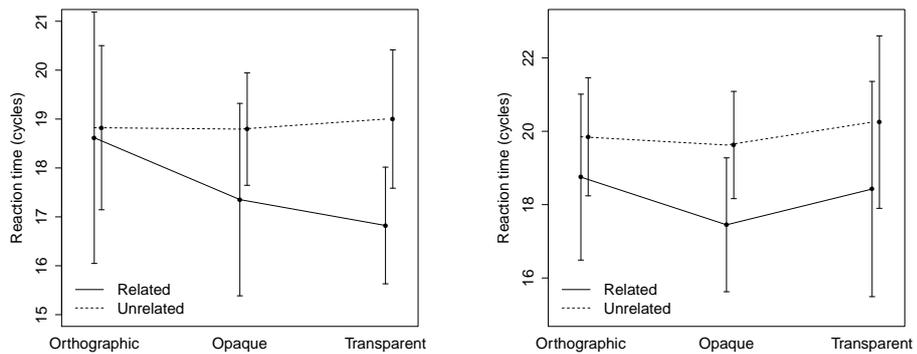


Figure 4.1: The IAAS model without morpho-semantic affix stripping. The weights of the connections from morpho-semantic nodes to affix nodes are set to zero.



(a) Stimuli from Beyersmann et al. (2012a)

(b) Stimuli from Beyersmann et al. (2012b)

Figure 4.2: Means and standard deviations of reaction times from the IAAS model without morpho-semantic affix stripping

type ANOVA revealed no significant interaction effects between both models ($F(2, 198) = 0.10, p = 0.905$). The main effect of model type was not significant ($F(2, 198) = 8.34, p = 0.196$).

The same ANOVA with the priming data from the second set of stimuli didn't show a significant effect of item type ($F(2, 117) = 2.89, p = 0.060$). Post-hoc comparisons revealed only a significant larger priming effect in the opaque condition than in the orthographic condition ($p = 0.049$), and no difference between the transparent condition and the orthographic ($p = 0.322$) and opaque ($p = 0.614$) conditions.

A direct comparison for this set of stimuli with the results from the full IAAS revealed no significant interaction effect ($F(2, 234) = 0.36, p = 0.696$), but a significant main effect for model type ($F(1, 234) = 4.16, p = 0.043$). Post-hoc comparisons showed that the overall priming effect was significantly larger in the full model than in the model without morpho-semantic affix stripping.

4.1.2 Discussion

The tests from the first simulation without morpho-semantic affix stripping show the same effects as the tests from the simulation with the full model. Although the ANOVA from the second simulation without morpho-semantic affix stripping was not significant, whereas this effect was significant with the full model, a direct comparison of these models only showed an overall difference in priming, no interaction across conditions.

From these results, we can conclude that removing morpho-semantic affix stripping doesn't change differences between priming conditions. To find out whether orthographic affix stripping does establish the differences between the orthographic and the other two conditions, another simulation will be performed in which morpho-orthographic affix stripping will be switched off.

4.2 Simulation without Affix Stripping

Removing orthographic affix stripping in addition to morpho-semantic affix stripping was done by changing the weights of the connections from sublexical nodes to affix nodes to zero (see Figure 4.3).

4.2.1 Results Simulations

In the simulations with both sets of stimuli, all stimuli were recognized correctly by the model. Figure 4.4 shows the means and standard deviations of both simulations.

For the first, as well as for the second set of stimuli, a one-way ANOVA didn't reveal a significant effect of item type (respectively $F(2, 99) = 2.46, p = 0.090$, and $F(2, 117) = 1.145, p = 0.322$).

A two-way prime type by model type ANOVA between the model without affix stripping and the model without morpho-semantic affix stripping for the first set of stimuli revealed a significant interaction-effect ($F(2, 198) = 3.092, p = 0.048$), as well as significant main effects for prime type ($F(2, 198) = 13.35, p < 0.001$) and model type ($F(1, 198) = 23.62, p < 0.001$). For the second set of

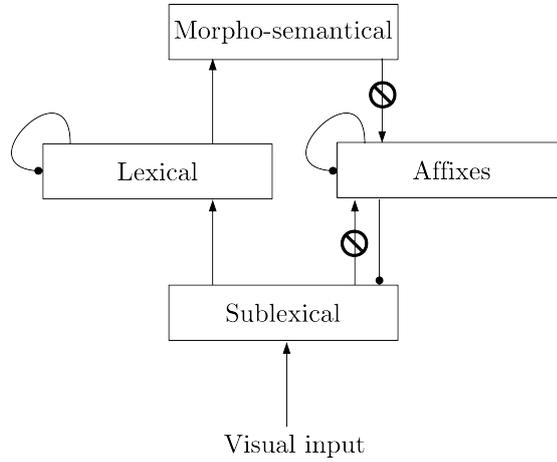
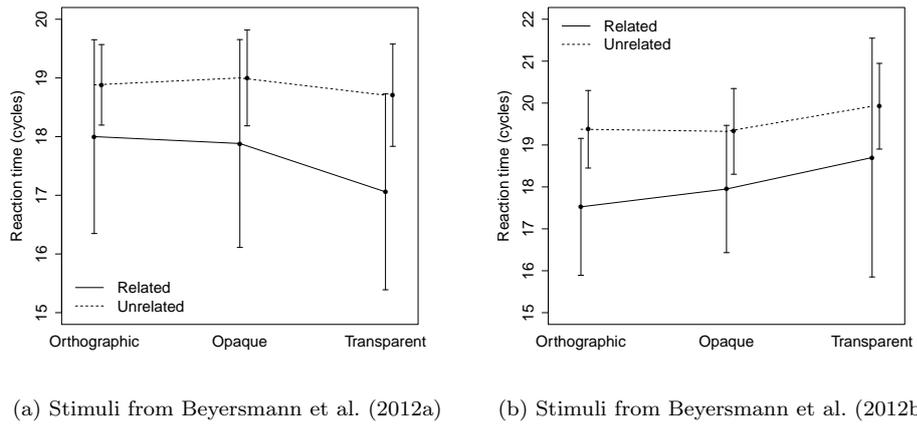


Figure 4.3: The IAAS model without morpho-semantic and orthographic affix stripping. The weights of the connections from morpho-semantic nodes to affix nodes and from sublexical nodes to affix nodes are set to zero.



(a) Stimuli from Beyersmann et al. (2012a)

(b) Stimuli from Beyersmann et al. (2012b)

Figure 4.4: Means and standard deviations of reaction times from the IAAS model without morpho-semantic and morpho-orthographic affix stripping

stimuli, the interaction-effect was also significant ($F(2, 234) = 3.74, p < 0.025$). The main effect for model type was not significant ($F(1, 234) = 3.66, p < 0.057$).

4.2.2 Discussion

No significant differences in priming effects have been found between conditions. Since removing morpho-orthographic affix stripping lead to this disappearance of these effects, we can conclude that morpho-orthographic affix stripping works as a mechanism to simulate orthographic priming effects from the MMP task.

Although no significant differences between priming results are found, there are differences noticeable between the means of these particular data sets. Looking at the first simulation, shows that the mean priming effect for this particular set of stimuli is larger in the transparent condition than the opaque ($M = 0.50$ cycles) and orthographic ($M = 0.74$ cycles) condition. Means in the second simulation show a larger priming effect in the orthographic condition than the opaque ($M = 0.48$ cycles) and transparent condition ($M = 0.53$ cycles). Since no mechanism can cause these differences, all differences must be caused by orthographic differences between conditions. Although these differences didn't lead to significant biases, the vulnerability of morpho-semantic effects in experimental research is a reason to look more closely at what confounding variables influence priming effects. If these orthographic effects also influence human behaviour, controlling for them might reduce variance or remove false effects.

4.3 Summary

This chapter explored the IAAS model, as described in Chapter 3, by removing affix stripping mechanisms in order to determine their exact contributions to the effects described in the previous chapter.

Removing morpho-semantic affix stripping from the IAAS model didn't change differences in priming between conditions. Therefore, morpho-semantic affix stripping as implemented in the IAAS model doesn't establish a larger priming effect for real affixed primes than for pseudo-affixed primes in the MMP task.

With the removal of morpho-orthographic affix stripping from the model without morpho-semantic affix stripping, the difference in priming between the orthographic condition and the opaque and transparent conditions disappeared. Therefore, morpho-orthographic affix stripping does establish a larger priming effect for real and pseudo-affixed primes than for non-affixed primes.

Although differences in priming effects weren't significant when affix stripping was turned off, these simulations did show that orthographic characteristics of the stimuli did influence means of the priming effects. If these characteristics can be identified, the influence of such confounding variables on human behaviour can be investigated. If a similar influence is found, future studies can control for this, so that variance might be reduced or false effects removed.

Chapter 5

General Discussion

This chapter provides an overview of the main discussion points related to the IAAS model. It evaluates the model, its implementation process and its theoretical plausibility, as well as some aspects of stimulus selection that might bias priming effects. The chapter ends with the conclusions from this thesis.

5.1 Evaluation of Affix Stripping

Morphological and pseudo-morphological priming effects of the orthographic presence of the affix were successfully modelled by the IAAS model. Therefore, morpho-orthographic affix stripping offers an explanation for the differences in priming effects between non-affixed and (pseudo-)affixed primes in the MMP task.

The simulations with the IAAS model and the further explorations of the model revealed that morpho-semantic affix stripping didn't establish semantic priming effects. The reason for this is probably that the time window in which the priming effect can be established is too small to differentiate between the opaque and transparent condition. In order to establish a priming effect between affixed and non-affixed conditions, affix stripping must start very early in a trial, so that enough time remains for the lexical prime nodes to influence the target nodes sufficiently. Semantic affix stripping cannot be established quickly enough to overpower this effect and differentiate between real words and pseudo-affixed words.

The short time in which an affix needs to be detected in the prime also decreases the accuracy with which affixes can be detected orthographically. The main mechanism of this trade-off between speed and accuracy is the competition that takes place between affix nodes. The inhibition threshold should be reached as quickly as possible, in order to establish the priming effect. Therefore, this threshold is put at a low level of activation. However, it takes some time for the correct affix node to inhibit all other affix nodes sufficiently. Because the threshold is low, incorrect affix nodes can reach the threshold before they are inhibited. Putting the threshold at a higher level would increase the room for competition and therefore make the accuracy higher, but detection slower, in which case the priming effect would decrease. Therefore, the inhibition threshold is put at a low level, and affix detection is sloppy.

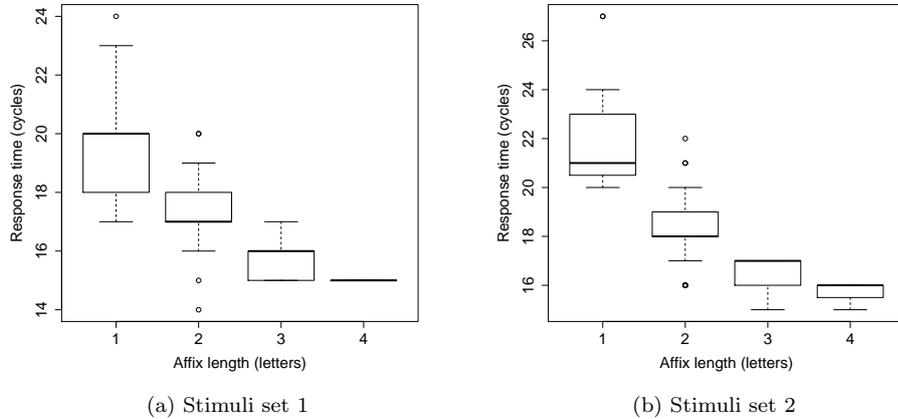


Figure 5.1: Response time per affix length

Another source of incorrect affix detection was that affix nodes that partially overlapped with the input could reach the threshold (e.g. -ING upon presentation of *china*). The increase in activation of affix nodes over time is amplified by the cumulative increase in activation of sublexical nodes. In an attempt to decrease these incorrect affix detections, in an early stage of implementation, toggle units were temporarily implemented at the sublexical level. These are units that can be either 1, when active, or 0, when inactive. This stabilizes the output from sublexical units. Affix and word units remained nodes with continuous values, in which activation accumulates according to neural network principles (see Section 2.1.3). This made affix detection more accurate. However, it also changed activation patterns and therefore distorted priming effects. In addition, it made the model less comparable to the original IA model, and therefore further removed from the models that were proposed in the literature. Therefore, toggle units weren't a satisfying solution, and sublexical nodes were restored as nodes with continuous values, at the cost of a higher number of incorrect affix detections.

5.2 Confounding Variables

In the simulations without morpho-orthographic and morpho-semantic affix stripping, no mechanism could cause a difference in priming effects between conditions. However, differences in means were visible (see Section 4.2). Therefore, the difference between conditions must be the consequence of orthographic characteristics of the stimuli. Theoretical and statistical consideration has led to a characteristic that seems to influence the priming effect: the difference in length between the prime and the target. In the opaque and transparent condition, this is the length of the (pseudo-)affix. Therefore, we call this variable the affix length.

The theoretical argument is as follows: during target presentation, the lexical prime node also gets some bottom-up activation from sublexical nodes, which is normalized for its own word length. If the difference in word length between

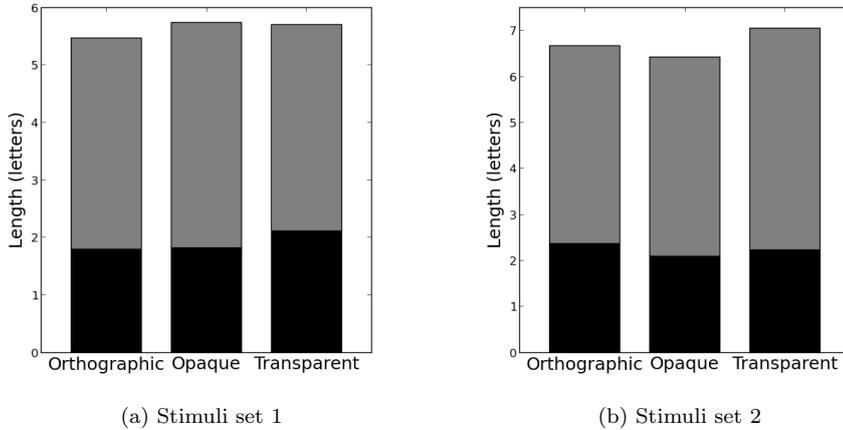


Figure 5.2: Stimuli length per condition. Gray is prime length, black is affix length, total height is target length (prime plus affix)

prime and target is large, the proportion of sublexical nodes that activate the lexical prime node during target presentation is smaller. Therefore, the prime node is lower in activation, which causes the inhibition from prime to target node to be less strong. This increases the target node activation, so that the target can be recognized earlier and response times are faster.

The simulation data show that response time indeed differs across affix lengths (see Figure 5.1). When a linear model is fitted on the response times of the first simulation, with repeated contrasts between the affix lengths, response time was significantly faster when affix length was longer (1 vs. 2: $t = 5.89, p < 0.001$, 2 vs. 3: $t = 3.46, p < 0.001$). The differences between an affix length of 3 and 4 wasn't tested, since only one prime had an affix length of 4.

The average affix length of the transparent conditions is slightly higher than in the two other conditions in this first set of stimuli (see Figure 5.2). Note that the article reports that the stimuli in the conditions were matched on prime length, target length and orthographic overlap, which was calculated by dividing the prime length by target length. There was some discrepancy of these variables between conditions, despite the matching, which is understandable, because of the many factors that stimuli must be matched on. No post-hoc analyses of these variables were reported, in order to check their influence on the data.

A post-hoc analysis was performed on the data from the first simulation without affix stripping to investigate whether affix length can cause the differences in priming means when both affix stripping mechanisms are switched off. A linear model with simple contrasts between conditions was fitted on the response times and showed that the difference between the orthographic and transparent condition vanished if the responses were controlled for affix length (from $t = 2.17, p = 0.032$ to $t = 0.60, p = 0.549$). This means that the effect from the transparent condition could be fully explained by affix length.

The contribution of the transparent condition to this model is significant ($t = 2.17, p = 0.032$), whereas the opaque contribution is not. If we add affix

length to the model, the contribution from affix length is similar to the model without experimental conditions and the contribution from the transparent condition disappears ($t = 0.60, p = 0.549$). This means that the difference between conditions in the simulation without affix stripping disappears when controlled for affix length.

The fact that semantic priming isn't a very robust effect, illustrated by the comparison of the studies that were used for the IAAS simulations, raises the question which confounding variables might influence the semantic priming effect. Because Beyersmann et al. (2012a) did find a semantic priming effect and this difference in the IAAS simulation can be explained by the differences in affix length, and because neither Beyersmann et al. (2012b), nor the IAAS simulation found a semantic priming effect for the second set of stimuli, and no large difference in affix length was present between these conditions, affix length might also explain part of the effect in the experimental data.

The recommendation for further research is to investigate whether affix length, or the general relation between prime and target length, also influences human behaviour. If it does, future MMP experiments should control for the difference in length between prime and target. Since stimulus selection can become too complicated if too many variables need to be controlled for, the effects could be controlled for statistically. If affix length doesn't influence behaviour, this aspect of the model should be reconsidered in order to make future simulations more realistic.

5.3 Stimulus Selection

As pointed out in the previous section, selecting stimuli can be tricky. Many variables might bias results if no proper matching has taken place. However, the number of words that could be used is finite and decreases with each additional matching criterion. Nevertheless, we put forward one additional factor that might bias priming effects.

In selecting the stimuli for the experiments, the condition for a prime to contain a (pseudo-)affix was whether the letters that remained after the target was stripped from the prime did form an affix. However, as described in Section 3.6.1, this raises the question whether the orthographic priming effect is dependent on this subtraction between prime and target or on the prime only. The IAAS model supports the last possibility. From primes that were included as non-affixed, an affix was stripped off on an orthographic basis, such as *twinkle-twin* and *comet-come*. Because these primes were regarded as pseudo-affixed by the model, the difference in results between the orthographic and the other conditions decreased.

Selecting stimuli carefully not only reduces variance and decreases bias, it also makes modelling easier. For example, word frequency is incorporated in many IA models, since it is known to influence speed of recognition. This is done by slightly increasing the resting level if a word is high frequent. However, this increases the number of factors that influences response times, which makes the sources of the effects less transparent. Fortunately, the stimuli selected in the experiments were corrected for word frequencies and the number of orthographic neighbours, so that no bias in the results is expected from these variables. Nevertheless, many factors needed to be incorporated in the model

and therefore, many decision needed to be made.

5.4 Interactive Activation

The decision to use the interactive activation model as basis for the model has advantages and disadvantages. The main advantages of the model are that it is widely used in previous research in which numerous human behaviours have been successfully modelled, and that it has been used repeatedly as a framework in which to explain morphological processing. Using a similar framework connects the implemented model with the proposed IA models in the literature, which haven't been implemented, and in that way bridges the gap between these theoretical models and implemented IA models that successfully simulate other human effects in visual word recognition.

The main disadvantage of the IA model is that it is a model with many free parameters and many connections that are influenced by these parameters. Activation patterns can become complex, because of the large number of connections between nodes and the way they interact. Keeping all connections that are present in the original IA model was therefore too complicated for these simulations. Excitatory top-down connections, as well as the feature level, were removed. Nevertheless, many free parameters remained.

5.5 Parameter Setting

Parameter values were as much as possible taken over from other models, such as the original IA model (McClelland and Rumelhart, 1981), an adaption of this model that incorporates words of different lengths (Grainger and Van Heuven, 2003) and a model that makes use of masked field weighting (Davis, 2010). If a parameter value was problematic for a successful simulation, alternative values were systematically tested, in order to find the optimal value. This was sometimes done by looking at activation curves and how they are influenced by a certain parameter, which increases the understanding of the model but also increases the probability that the model is over-adjusted to a specific trial. Therefore, the model was also tested with a larger set of stimuli, and only after doing runs with different parameter values, results were compared.

As an example of finding an optimal value, the criteria are given of the optimal weight for the inhibition from an affix node to its corresponding sublexical nodes. The priming effect can be increased by increasing this parameter. However, an increase in the priming effect leads in some cases to incorrect word recognition. Some morphologically simple target words end with a letter combination that forms a legitimate affix, such as *sing*, *manic*, *invent* and *ether*. If the pseudo-affixes of these short words are inhibited too strongly, the remaining letters don't have enough power to activate the corresponding target node and in that way to inhibit the prime node. For example, upon presentation of the target *sing*, after the prime *singular*, I-3, N-2 and G-1 are very strongly inhibited by the affix node -ING. The nodes that remain activated, S+1, I+2, N+3, G+4 and S-4, don't have the power to activate the target node SING enough to inhibit the prime node SINGULAR. This leads to recognition of the prime instead of the target or in no recognition at all. Therefore, the optimal value of

the parameter is the highest value of inhibition for which all targets and other words are recognized correctly. However, when another parameter changes, such as the lateral inhibition between words, the optimal value for affix-to-sublexical node inhibition might change as well.

This example shows that optimal parameter values depend on other parameters. If one value changes, the whole activation flow through the model might be different, so that other parameters might have new optimal values. The number of parameters is too high to test each parameter in combination with every other parameter adjustment. Therefore, adjustments were ended when an optimal value was found and only reconsidered if there were indications that the model did not work properly because of this parameter. Nevertheless, parameter adjustment took a long time and could be frustrating, because it often wasn't clear whether and how the model should be improved.

An alternative for such a hand-wired approach is learning. Learning can be incorporated in IA, by adjusting the weights of connections each time units are activated together, so that the connections between units that are active together become stronger, resembling Hebbian learning (Hebb, 1949). However, implementing a learning algorithm and training the model was not the purpose of this model, because it makes it harder to determine whether effects are caused by affix stripping or by other factors that originate from adjusted weights.

Related to the difficulty of testing a large number of parameters is the computational efficiency of the model and processing speed of the computer. If the model can be run faster, more parameter values can be tested. In addition, a simulation is more realistic with a large lexicon. However, the running time of the model was quadratically related to the lexicon size and therefore increased rapidly with a larger number of words in the lexicon. Although efficiency has been increased by removing redundant computational steps, running an experiment would take several hours. Therefore, a suggestion for future versions of the model is to try to improve computational speed.

5.6 Behavioural Plausibility

Neural network models share many characteristics with human brain cells, or groups of cells. Therefore, the IA model is neuronally more plausible than alternatives that aren't neural network models. The similarity between the brain and the IA model is of course limited, and the IAAS model, for which the foundation is a simplified version of the IA model, without feature level and almost without top-down connections, might resemble the brain even less. Although neuronal constraints were incorporated in the design as much as possible, the main priority was to simulate the behavioural results, which were already present at the start of implementation. The design and implementation of the IAAS model were therefore data-driven. The behavioural and simulation results were compared repeatedly between implementation iterations and the model was adjusted according to these evaluations. This drew attention away from the neuronal constraints, but toward behavioural plausibility.

Some predictions can be derived from the IAAS model, that can be tested with human participants, and in that way increase or reduce its plausibility. One of these predictions that has already been mentioned in Section 5.2 is that the priming effect in an MMP task is stronger when the difference in length

between prime and target, the affix length, is larger.

Another prediction is that pseudo-affixed words are in general recognized slower than non-affixed words, because the inhibition from affix nodes to sublexical nodes during affix stripping reduces bottom-up excitation to the corresponding lexical node.

This might seem counter-intuitive, because affix stripping slows down processing instead of increasing processing efficiency. However, while affix stripping would slow down processing of pseudo-affixed words, it could facilitate processing of real affixed words, when semantic priming is incorporated in the model. Semantic priming has been found between prime-target word pairs that are semantically related words and can be either syntactically related or unrelated (Perea and Gotor, 1997). Semantic priming is therefore independent of orthographic overlap. In the IA model, semantic priming could be reflected by positive feedback from a shared node in the semantical level (e.g. “farm”, connected to FARM and FARMER).

The advantage of affix stripping could be explained as follows: if a lexical word stem is better able to activate its morpho-semantical representation than the derived affixed form is, activation of the derived form could be faster by top-down excitation from the morpho-semantical node to the lexical node of the affixed word than by the direct bottom-up excitation from sublexical nodes. The purpose of affix stripping would then be that it increases the activation of the lexical stem node, so that the faster indirect path of activation of the affixed word node is stronger. Pseudo-affixed words do not benefit from the positive effect between semantically related lexical nodes, so that affix stripping is disadvantageous for these words. Real affixed words are far more common than pseudo-affixed words in everyday language, so the overall effect would be a gain in recognition efficiency. This conception suggests that real affixed words are recognized faster than pseudo-affixed words and than non-affixed morphologically complex words, during single word presentation tasks.

This hypothetical purpose of affix stripping has the additional advantage of being an alternative explanation for the semantic priming effect. It could serve as a replacement for morpho-semantic affix stripping, which couldn't be successfully implemented in the IAAS model. In future research, both options could be compared with brain imaging results (e.g. Lavric et al. (2012)), to look which option fits best with the timing and location of brain activity.

This section has made some explicit behavioural predictions, which can be tested in order to make conclusions about the model's behavioural plausibility. As mentioned in previous chapters, previous IA models that were proposed to explain morphological processing aren't implemented, which causes them to remain vague on their predictions. In addition, these models seem to be designed in parallel with or after the experimental research, so that the design is shaped by the latest behavioural results. Like the IAAS model, their design is data-driven. Although the latest findings are constraining the model, the theoretical implications of the model are never tested. In this way, no explicit reasons can be determined to favour one model over another.

In order to test the behavioural plausibility of a model, not only data-driven modelling should be performed, but also model-driven behavioural research. Therefore, a final advice for future studies that want to incorporate a model, is to change the aim of research from comparing experimental conditions and designing a model that explains the results, into designing a model or taking

a previously proposed model, implementing it and then testing the model's predictions with behavioural research.

5.7 Conclusions

The IAAS model has successfully simulated stronger transparent and opaque priming than orthographic priming in the masked morphological priming task. Therefore, morpho-orthographic affix stripping can be used as a mechanism to explain this priming effect. Furthermore, the model has not simulated stronger transparent priming than opaque priming. A possible reason for this is that the time-window in which affix stripping should take place is too small to differentiate between both priming effects.

Therefore, we can conclude that the answer to the main research question of this thesis - whether morpho-orthographic and morpho-semantic effects in human processing of (pseudo-)affixed words can be simulated with an affix stripping mechanism in an IA model - is positive with respect to morpho-orthographic effects. However, the presented simulations with the IAAS model are inconclusive on morpho-semantic priming effects.

The answer to the additional research question, whether precise positional information is necessary for affix detection, is positive. Affix stripping couldn't be implemented with coarse-grained open-bigram coding, whereas simulations were successful with fine-grained both-end position coding.

Future research should investigate whether a mechanism like orthographic affix stripping is used in the brain. In addition, research should investigate whether improvements in the model could be made to simulate morpho-semantic affix stripping, or whether alternative accounts, such as semantic priming, are more favourable explanations. Finally, investigations might clarify which factors contribute to the robustness of the priming effects, for example by testing the influence of possible confounding variables, like the length of the affix.

Bibliography

- Baayen, R. H., Milin, P., Durdevic, D. F., Hendrix, P., and Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological Review*, 118(3):438–481.
- Beyersmann, E., Castles, A., and Coltheart, M. (2012a). Morphological processing during visual word recognition in developing readers: Evidence from masked priming. *The Quarterly Journal of Experimental Psychology*, 65(7):1306–1326. PMID: 22530670.
- Beyersmann, E., Coltheart, M., and Castles, A. (2012b). Parallel processing of whole words and morphemes in visual word recognition. *The Quarterly Journal of Experimental Psychology*, 65(9):1798–1819. PMID: 22530670.
- Caramazza, A. and Hillis, A. E. (1990). Levels of representation, co-ordinate frames, and unilateral neglect. *Cognitive Neuropsychology*, 7(5-6):391–445.
- Cohen, P. R. (1995). *Empirical methods for artificial intelligence*. MIT press Cambridge.
- Coltheart, M., Curtis, B., Atkins, P., and Halter, M. (1993). Models of reading aloud: Dual-route and parallel-distributed-processing approaches. *Psychological Review*, 100:589–608.
- Davis, C. J. (2010). The Spatial Coding Model of Visual Word Identification. *Psychological Review*, 117:713–758.
- Diependaele, K., Sandra, D., and Grainger, J. (2005). Masked cross-modal morphological priming: Unravelling morpho-orthographic and morpho-semantic influences in early word recognition. *Language and Cognitive Processes*, 20(1-2):75–114.
- Diependaele, K., Sandra, D., and Grainger, J. (2009). Semantic transparency and masked morphological priming: The case of prefixed words. *Memory & cognition*, 37(6):895–908.
- Ellis, R. and Humphreys, G. (1999). *Connectionist Psychology: A Text With Readings*. Psychology Press.
- Fischer-Baum, S., Charny, J., and McCloskey, M. (2011). Both-edges representation of letter position in reading. *Psychonomic Bulletin & Review*, 18:1083–1089.

- Fischer-Baum, S., McCloskey, M., and Rapp, B. (2010). Representation of letter position in spelling: Evidence from acquired dysgraphia. *Cognition*, 115(3):466 – 490.
- Forster, K. I. and Davis, C. (1984). Repetition priming and frequency attenuation in lexical access. *Journal Of Experimental Psychology. Learning Memory And Cognition*, 10(4):680–698.
- Forster, K. I., Davis, C., Schoknecht, C., and Carter, R. (1987). Masked priming with graphemically related forms: Repetition or partial activation? *The Quarterly Journal of Experimental Psychology Section A*, 39(2):211–251.
- Glasspool, D. W. and Houghton, G. (2005). Serial order and consonantvowel structure in a graphemic output buffer model. *Brain and Language*, 94(3):304 – 330.
- Grainger, J., Colé, P., and Segui, J. (1991). Masked morphological priming in visual word recognition. *Journal of memory and language*, 30(3):370–384.
- Grainger, J. and Jacobs, A. M. (1996). Orthographic processing in visual word recognition: A multiple read-out model. *Psychological review*, 103(3):518–565.
- Grainger, J. and Van Heuven, W. (2003). Modeling letter position coding in printed word perception. *The mental lexicon*, pages 1–24.
- Grainger, J. and Ziegler, J. (2011). A dual-route approach to orthographic processing. *Frontiers in psychology*, 2.
- Hebb, D. (1949). The organization of behavior; a neuropsychological theory.
- Heide, J., Lorenz, A., Meinunger, A., and Burchert, F. (2010). The influence of morphological structure on the processing of german prefixed verbs. In Onysko, A. and Michel, S., editors, *Cognitive perspectives on word formation*, pages 375–398. De Gruyter Mouton.
- Hubel, D. H. and Wiesel, T. N. (1959). Receptive fields of single neurones in the cat’s striate cortex. *The Journal of physiology*, 148(3):574–591.
- Hubel, D. H. and Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, 195(1):215–243.
- Jacobs, A. M., Rey, A., Ziegler, J. C., and Grainger, J. (1998). Mrom-p: An interactive activation, multiple readout model of orthographic and phonological processes in visual word recognition. In Grainger, J. and Jacobs, A., editors, *Localist Connectionist Approaches to Human Cognition*, Scientific Psychology Series, pages 147–188. Lawrence Erlbaum Associates, Incorporated.
- Kazanina, N., Dukova-Zheleva, G., Geber, D., Kharlamov, V., and Tonciulescu, K. (2008). Decomposition into multiple morphemes during lexical access: A masked priming study of russian nouns. *Language and Cognitive Processes*, 23(6):800–823.
- Koskenniemi, K. (1984). A general computational model for word-form recognition and production. In *Proceedings of the 10th international conference on Computational linguistics*, pages 178–181. Association for Computational Linguistics.

- Lavric, A., Elchlepp, H., and Rastle, K. (2012). Tracking hierarchical processing in morphological decomposition with brain potentials.
- Lété, B. and Fayol, M. (2013). Substituted-letter and transposed-letter effects in a masked priming paradigm with french developing readers and dyslexics. *Journal of Experimental Child Psychology*, 114(1):47 – 62.
- Longtin, C.-M., Segui, J., and Halle, P. A. (2003). Morphological priming without morphological relationship. *Language and cognitive processes*, 18(3):313–334.
- Max Planck Institute for Psycholinguistics (2011). Webcelex. Retrieved 9 January 2013 from <http://celex.mpi.nl>.
- McClelland, J. and Rumelhart, D. (1981). An interactive activation model of context effects in letter perception: I. an account of basic findings. *Psychological Review; Psychological Review*, 88(5):375.
- Nooteboom, S., Weerman, F., and Wijnen, F., editors (2002). *Storage and Computation in the Language Faculty*. Kluwer Academic Publishers.
- Perea, M. and Gotor, A. (1997). Associative and semantic priming effects occur at very short stimulus-onset asynchronies in lexical decision and naming. *Cognition*, 62(2):223 – 240.
- Rastle, K. and Davis, M. H. (2008). Morphological decomposition based on the analysis of orthography. *Language and Cognitive Processes*, 23(7-8):942–971.
- Rastle, K., Davis, M. H., Marslen-Wilson, W. D., and Tyler, L. K. (2000). Morphological and semantic effects in visual word recognition: A time-course study. *Language and Cognitive Processes*, 15(4-5):507–537.
- Reicher, G. (1969). Perceptual recognition as a function of meaningfulness of stimulus material. *Journal of Experimental Psychology*, 81(Vol 81(2)):275–280.
- Sandra, D. (1994). The morphology of the mental lexicon: Internal word structure viewed from a psycholinguistic perspective. *Language and cognitive processes*, 9(3):227–269.
- Seidenberg, M. S. and McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological review*, 96(4):523–568.
- Sprenger, S. A. and van Rijn, H. (2013). Its time to do the math: Computation and retrieval in phrase production. *The Mental Lexicon*, 8(1):1–25.
- Taft, M. (1994). Interactive-activation as a framework for understanding morphological processing. *Language and Cognitive processes*, 9(3):271–294.
- Taft, M. and Forster, K. I. (1975). Lexical storage and retrieval of prefixed words. *Journal of Verbal Learning and Verbal Behavior*, 14(6):638 – 647.
- Wheeler, D. D. (1970). Processes in word recognition. *Cognitive Psychology*, 1(1):59 – 85.

- Whitney, C. (2001). How the brain encodes the order of letters in a printed word: The serial model and selective literature review. *Psychonomic Bulletin & Review*, 8(2):221–243.