RADBOUD UNIVERSITY NIJMEGEN

FACULTY OF SOCIAL SCIENCE

Resolving Misunderstanding

A PRAGMATIC APPROACH

THESIS BSC ARTIFICIAL INTELLIGENCE

Supervisor:
Laura VAN DE BRAAK

Author:
Ron HOMMELSHEIM
s1000522

Co-Supervisor:
Iris VAN ROOIJ

Advisor:
Mark BLOKPOEL

June 2020

Abstract

How Do We Resolve a Misunderstanding? Although two people might have different preferences in choosing a signal for a referent, they seem to be able to understand each other. Even when a misinterpretation occurs, people are able to rectify the conversation and align their understanding. In this Bachelor thesis, Frank, Goodman, and Hawkins' (2017) model, which describes a version of the rational speech act (RSA), was implemented and subsequently used as a base model in order to investigate the resolution of misunderstandings. Both, said implementation and an adjusted model are introduced. The modified model proposes an updating mechanism based on extended conversational output. Although, this mechanism does not resolve a misunderstanding completely, this thesis outlines many assumptions, of the model and of the framework it is used in, which have to be taken into account when addressing this intriguing question of communication.

Contents

1 Introduction

Communication is an extraordinary capacity. It encompasses various subtle and intriguing phenomena which are not yet well understood. How do we understand each other, given that we have different beliefs and associations about the world? How can we reason about the best signal to use, such that a listener will most likely understand what we would like to convey? Likewise, how can we adjust our belief about the others perspective once a misunderstanding has occurred?

Say you have a cat named “The Professor”. At the university you tell your friend Nora: “The Professor has been very tired lately”, to which she responds: “Well, we did just have an exam period.”. Immediately you notice that, clearly, Nora is not showing concern for your cat and misunderstood what you were referring to. This is the problem I am researching in this project: How can we detect and rectify a conversation once a misunderstanding has occurred? To do so, I model this phenomenon and experiment with it in a computational simulation.

1.1 Background

In their 2012 paper, Frank and Goodman introduced the Rational Speech Act (RSA) model [1]. In it, they employ recursive Bayesian inference to model the communication of agents who make use of pragmatic reasoning. The model is based on Grice’s cooperative principle [2], which assumes that agents will act cooperatively, in order to understand each other.

The original publication illuminates the main idea by experimenting with reference games. This is exemplified in figure 1.1. We have three objects, varying on two dimensions: color and shape. Now, a non-pragmatic speaker might refer to the first object by saying either “blue” or “square”, to the second object by saying either “blue” or “circle” and to the third object by saying either “green” or “square”. To be cooperative however, a pragmatic speaker would try to use the most informative signal to refer to the intended object. The first object is neither unique in color, nor in shape. The second object is unique only in shape, and the third object is unique only in color. Therefore, a pragmatic speaker would refer to the first object either by saying “blue” or “square”, since he can refer to the second object uniquely by saying “circle” and to the third by saying “green”. A pragmatic listener then, would reason similarly and deduce that the speaker refers to the first object when he says “blue”. This, in essence, is the idea of the rational speech act. People step into each others perspective to infer the most informative signal and thus to improve their communication. The model successfully imitates experimental observation and therefore is a promising tool, both to model the human capacity and also to use in artificial intelligence applications.

Figure 1.1: A Simple Reference Game

In 2016, the authors extended the model and presented it novel situations [3]. They demonstrated the model's robustness and showed how emergent phenomena such as metaphor and vagueness can be explained. Despite the model's success, further extensions and adaptations were made in order to accommodate for more intricate assumptions. Hawkins et al. [4] showed how an adjusted RSA model can explain how people find the most useful question to ask and how the most appropriate answer can be inferred. The world is ambiguous as there is not a one-to-one mapping between signals and referents. In the example above, there were two objects to which the signal "blue" could refer to. Frank et al. experimented with the model in the face of ambiguity and showed how the model can withstand it [5]. Another obstacle in communication is that people view the world differently from each other and hold individual beliefs about it. Mark Blokpoel et al. [6] adjusted the model and showed that *ambiguity*, that words can refer to multiple things, can be advantageous and therefore favourable in overcoming peoples incongruent beliefs about the world (*asymmetry*). This finding contradicts the long-held assumptions that ambiguity is an obstacle for clear communication.

In the context of RSA, a person's belief is usually modelled using a lexicon. A lexicon is a pairing between signals and referents and can be represented as a matrix. Figure 1.2 gives an example. One fundamental question, is whether people have a personal lexicon which represents a more or less rigid and individual belief about the world or alternatively, hold a probability distribution over all possible lexicons. Perhaps it is something in between. Regardless, researchers are aware of this dichotomy and define their models consequently.

	Blue	Green
Circle	1	0
Square	1	1

Figure 1.2: Associated Lexicon For The Example In Figure 1.1

In 2017 Hawkins, Frank, and Goodman published a model [7], which attempts to explain convention-formation, i.e. how two people converge to use the same set of signal-referent pairs. The authors underline three properties which are detrimental to the model's success. *Arbitrariness*: given two signals that refer to an object equally well, it is arbitrary which one to use - but communicators must agree upon it; *stability*: once a signal successfully refers to an object and is established, it is disadvantageous to change it; *reduction*: longer expressions will be condensed to simpler signals. This model however, makes use of *explicit feedback*. After each turn in a conversation, the speaker is informed of the listeners inferred referent. I argue, that people are able to disambiguate internally, without the need for explicit feedback. Moreover, their model describes a probability distribution over lexicons which I find less plausible than having a personal lexicon. Nonetheless, this model sets a good starting framework in order to investigate how people act, and how the RSA model can be used, when a misunderstanding has occurred. Therefore, I use it as a "base model" but deviate from it to make it align with my own assumptions.

1.2 Tangrams & Fribbles

In subsequent studies, more elaborate tasks have been utilised in order to investigate RSA models. Hawkins, Frank, and Goodman for example, made use of Tangrams which are basic geometrical shapes from which larger structures can be composed (figure 1.3). Here, participants had to work in pairs, in which one of them tries to refer to a specific Tangram while the other has to infer which one was talked about.

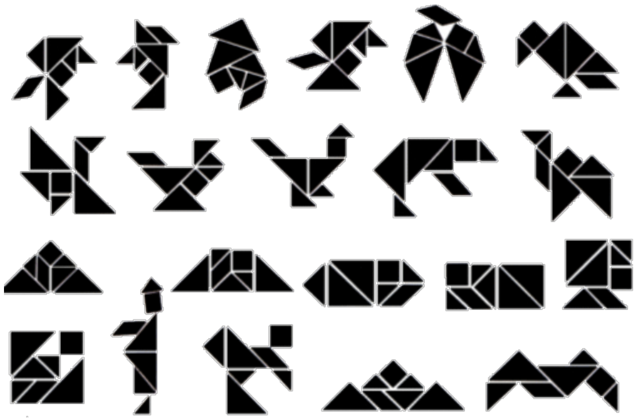

Figure 1.3: Tangrams

Similarly, Barry et al. adapted this experiment and introduced Fribbles [8]. These are a little more complex, three-dimensional objects. The Communicative Alignment in Brain and Behaviour (CABB) lab [9] used these objects in a referral task which was a partial motivation for this research. In the task, two participants stood in front of each other, each with an image of 16 different Fribbles. Again, one of the participants would then have to use different signals to refer to a certain Fribble while the other tries to infer which one was meant. For example, the Fribble in figure 1.4a could be referred to as "the cylindrical one" whereas the Fribble in figure 1.4b could be referred to as "the floating one". This task however, given the amount of Fribbles and their more complex shapes, proved to be non-trivial.

(a) Fribble 1 (b) Fribble 2

Figure 1.4: Two Fribbles

1.3 Relevance

One of the major pitfalls in previous RSA models is the inclusion of explicit feedback. This is analogous to the listener pointing to the object she thinks the speaker refers to in the Fribbles task. Usually in a normal conversation, a listener does not keep mentioning what they think someone is talking about. Whether two people are referring to the same thing in a conversation is subtle and implicit. Therefore, it is necessary to devise a communicative strategy which avoids explicit feedback.

More globally, there are two major reasons to conduct this research. The first one is human-centric. By creating a plausible model which allows us to investigate the intricacies of communication and specifically the resolution of misunderstanding, we can try to compare and extrapolate our findings to our own underlying cognitive mechanisms. In other words, a computational simulation can be a good method to learn about ourselves. The second reason is general. Finding a mechanism in which clear communication can arise from agents with misaligned, incomplete and ambiguous lexicons, can be beneficial for numerous applications such as chat-bots or virtual assistants. In a sense, it boils down to a purely mathematical problem of optimal data generation, -transfer, and signal recovery.

2 Methodology

The resolution of misunderstandings was approached by a two-step process. First, I implemented the model described in Frank, Goodman, and Hawkins [7] in order to use it as a base model. Henceforth, I shall name it "The Pointer Model", since it includes explicit feedback, which is analogous to the listener pointing to the inferred referent. Then, I introduce an original model which is not making use of explicit feedback. I will refer to it as "The Confirmer Model", for reasons which will be clear shortly. The question which I address, necessitates an extended conversation. It is not only the model itself, but also the framework in which a conversation with a misunderstanding can occur, that is paramount to the findings. Therefore, I modelled the required framework which however brought new obstacles to light which all be addressed.

Both models were implemented in Scala [10], a multi-paradigm programming language [1].

2.1 The Pointer Model

Frank, Goodman, and Hawkins initially proposed their model to address the question: "What cognitive mechanisms support the emergence of linguistic conventions from repeated interaction". Their model is an adjusted version of RSA which utilises a probabilistic paradigm: an agent does not posses a personal lexicon which represents ones beliefs of which signals refer to which referents in the world. Instead, an agent always considers all possible lexicons and infers, which lexicon is most likely, given the current situation. This is a crucial and conceptually relevant consideration.

2.1.1 A Monologue

A monologue is an interaction between two agents, say, Adam and Nora, in which Adam has an intention he would like to convey to Nora. To do so, he produces a signal which is inferred by Nora. Here, Adam takes the role of the speaker and Nora takes the role of the listener. This is called a monologue, because only one agent speaks while the other only listens.

2.1.2 Formalisation

A speaker wants to find the most informative signal to use such that a listener will most likely understand it. A listener wants to interpret a signal and find

[1] The full implementation can be found on GitLab [11]

out which referent the signal most likely refers to. The key idea is that either agent puts themselves in the other agents shoes, recursively, until they imitate the literal listener who does not infer but simply consults their lexicon. The order of pragmatic inference implies how often an agent will enter the recursive process, or "step into the others agent shoes" for their reasoning.

Speaker

The speaker creates a signal by considering what the inferred intention of the listener with order $n-1$ would be, given all possible lexicons.

Input: A set of signals $\mathcal{S}$, a set of referents $\mathcal{R}$, an intention $i \in \mathcal{R}$, an order of pragmatic inference $n \in \mathbb{N}$, and a set of observations of previous communication d which is comprised of signal and referent pairs $\{s_i, r_i\}$.

Output: A signal, $s \in \mathcal{S}$ which is most likely to be understood by a listener, given the intended referent i and the observed communication d.

$$S_n(s|i,d) \propto \exp\left(\alpha \log\left(\sum_{\mathcal{L}} P_{S_n}(\mathcal{L}|d) L_{n-1}(i|s,\mathcal{L})\right) - \mathrm{cost}(s)\right)$$

where,

$P_{S_n}(\mathcal{L}|d)$: is the posterior over a lexicon.

$\mathrm{cost}(s)$: is an internal function usually considering the utterance length or frequency of use.

α: is a soft-max optimality parameter controlling the extent to which the speaker maximizes over listener informativity.

Note, that the exponential and the logarithm is only an arithmetic trick to map between spaces. When α is set to 1 and the cost to 0, it can be omitted, which reduces the formula to:

$$S_n(s|i,d) \propto \sum_{\mathcal{L}} P_{S_n}(\mathcal{L}|d) L_{n-1}(i|s,\mathcal{L})$$

The speaker calculates a signal given only a single lexicon and a referent by:

$$S_n(s|r,\mathcal{L}) \propto \exp\left(\alpha \log L_{n-1}(r|s,\mathcal{L}) - \mathrm{cost}(s)\right)$$

Which again, depending on the circumstance, can be reduced to:

$$S_n(s|r,\mathcal{L}) \propto L_{n-1}(r|s,\mathcal{L})$$

Essentially, a speaker will recursively infer the optimal signal until they consider the literal listener (base case).

Listener

A listener infers which referent the speaker's signal most likely refers to by considering the speaker with order n, given all possible lexicons.

Input: A set of signals $\mathcal{S}$, a set of referents $\mathcal{R}$, a specific signal $s \in \mathcal{S}$, an order of pragmatic inference $n \in \mathbb{N}$, and a set of observations of previous communication d which is comprised of signal and referent pairs $\{s_i, r_i\}$.

Output: A referent $r \in \mathcal{R}$ which the speaker most likely refers to, given a signal and observations d.

$$L_n(r|s,d) \propto \sum_{\mathcal{L}} P_{L_n}(\mathcal{L}|d) L_n(r|s,\mathcal{L})$$

where

$$P_{L_n}(\mathcal{L}|d) \propto P(\mathcal{L}) \prod_i S_n\left(r_i|s_i,\mathcal{L}\right)$$

and is used to find the most likely lexicon of the speaker, and

$$L_n(r|s,\mathcal{L}) \propto P(r) S_n(s|r,\mathcal{L})$$

is the probability of a referent given the lexicon under consideration.

$$L_0(r|s,\mathcal{L}) \propto \mathcal{L}(s,r) \cdot P(r)$$

is the literal listener, consulting the given lexicon directly.

2.1.3 Success

A successful monologue is when Nora infers Adam's intended referent correctly, i.e. when $i = r$. Figure 2.1 shows a successful monologue on the left and an unsuccessful monologue on the right.

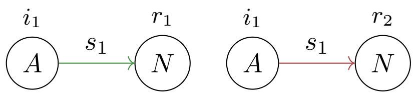

Figure 2.1: In both scenarios, Adam has intention i_1 and produces signal s_1. On the left, Nora infers referent r_1 correctly, which is highlighted by a green arrow. On the right, Nora infers referent r_2 which is incorrect and highlighted by a red arrow.

2.1.4 Implementation

A number of design choices and specifications were made which sometimes deviate from the original paper. Reasons were either because of underspecification in the original model or to avoid a computational explosion. In this implementation it is assumed that both agents have the same sets of signals and referents

$\mathcal{S}$ and $\mathcal{R}$. This is simply to ease the experimentation and aligns with the original paper. Frank, Goodman, and Hawkins specify a continuous lexicon in their paper, i.e. each signal-referent pair has a probability. However, since every possible lexicon has to be considered, I have chosen to use a discrete lexicon. When experimenting with two signals and two referents, which is represented by a 2 x 2 lexicon, there are already $2^4 = 16$ possible lexicons. With three signals and three referents, there are $2^9 = 512$ lexicons. Using continuous lexicons would have made the experimentation unfeasible. Although this is an alteration of the author's model, it is not detrimental to the theory and the results should be similar. The prior over lexicons was chosen to be given by a binomial distribution. This is to accommodate the natural distribution of possible lexicons. For example in a 4 x 4 lexicon there is $\binom{4}{0} = 1$ lexicon with no ones, $\binom{4}{1} = 4$ lexicons with one 1, etc. The α parameter and cost were set to 1 and 0, respectively. These parameters are inessential to the main model. Another assumption made, is that the speaker's produced signal is picked up unambiguously by the listener, i.e. the listener does not have a probability distribution over possibly heard signals. This detail was in fact not clearly specified in the author's paper. Nonetheless, the listener's confidence of hearing the correct signal was not an effect I wanted to test. Instead, only the interpretation of the signal is relevant. A specification which does make a noticeable difference, is the necessity for normalisation after each computation which involves the lexicon. This has the consequence of slowing down the computation immensely, while simultaneously influencing the results to fit the model and the authors own implementation [12]. The implementation further specifies that the agents never switch roles. The speaker stays the speaker and the listener the listener. Switching roles would be unnecessary because they both have access to the conversational observation, i.e. a list of signal-inferred-referent pairs. Within a conversation, a speaker always has the same intention, such as in the Fribbles task. The number of signals and the number of referents can be adjusted. However, due to the computational complexity, I was unable to run simulations with more than three signals or referents. The conversation length can be modified, i.e. how often the speaker tries to convey his intention to the listener and was set to 5. This accommodates the Fribbles task. Since a speaker considers all lexicons and ends up with a probability distribution over signals, the argmax function is used to pick the actual signal that is produced. The alternative would be to choose the signal according to the actual probability distribution the speaker computes, but this is a minor implementational deviation without much impact, especially since both agents do it similarly.

2.2 The Confirmer Model

People are able to communicate without the need for explicit feedback and therefore it will not be included in this model. Remember, Frank, Goodman, and Hawkins' model relies on a shared observation which consists of signal-referent pairs, as if the listener points to her inferred referent. The notion and plausibility of explicit feedback will be further elaborated on in the discussion section.

In this model, instead of computing a probability distribution over all lexicons for each step of the inference, agents possess a discrete, personal lexicon.

Psychologically, this seemed to be more plausible. It appears unlikely that people consider all possible signal-referent pair combinations when trying to either produce or interpret a signal. The amount of computations alone make this paradigm questionable. If there are only 10 signals and 10 referents, then one would have to consider 2^{100} lexicons with each turn in the conversation. Hence, a personal lexicon seemed to be a more suitable decision. One collateral and beneficial consequence is that a personal lexicon eased the computation by a large margin and I could therefore run simulations with much larger lexicons. More consequences of this choice will be discussed in detail in the following sections.

2.2.1 A Dialogue

Detecting a misunderstanding necessitates *some* sort of feedback. That is, presuming the speaker is the one detecting the misunderstanding [2]. Somehow, Adam needs to be able to interpret a signal, a response, to determine or at least reason about Nora's conclusion.

This can be done by the means of a dialogue. Adam produces a signal to refer to his intended referent, let's call it r_A. Nora infers a referent - r_N and will produce another, *different*, signal to refer to it. Now Adam assumes that Nora talks about the same referent and can estimate if Nora's signal coheres with the originally intended referent - r_A. The idea is as follows. Imagine Adam wants to converse about a knife, and hence produces the signal "The sharp thing". Now, Nora can continue the conversation about what she thinks Adam refers to and say "Oh, you mean the pointy thing". This aligns with Adam's idea of a knife, and so Adam assumes that Nora understood him. If however, Nora would answer "Oh, you mean the round thing", Adam would conclude that Nora has not made a correct inference. Although this is the initial example which lead to the naming of the model - The Confirmer Model - as if Nora confirms the conversation by responding to it, this does not always have to be the case. Another way of thinking about it is just as an extended conversation. Recall the initial example in which you initiated a conversation by talking about your cat. Nora responded simply by casually continuing the conversation.

2.2.2 Formalisation

Quite similarly to the "Pointer" model, a speaker tries to find the most informative signal and a listener tries to infer the most likely referent. Besides switching to a personal lexicon paradigm, two significant adjustments were made.

First, since agents conduct themselves in a dialogue instead of a monologue, they can change their behaviour depending on their turn in the conversation. The initiator of the conversation will be referred to as the *active* agent, whereas the responding agent will be referred to as the *reactive* agent. This is merely to avoid confusion. In the example above, Adam is the active agent and Nora the reactive agent.

Second, an agent has a weight matrix over signal-referent pairs which allows for bias and specific modification, both of which are specified below. The weights start off uniformly but can be updated, based on inference. Note, that only the

[2]See the discussion for a more rigorous exploration of ideas, considering the detection and handling of a misunderstanding.

active agent, updates his weights. This is because only he knows the original intention and can compare it to the final inference.

Speaker

The speaker produces a signal by recursively stepping into the other agent's perspective. Here, both agent act similarly.

Input: A set of signals $\mathcal{S}$, a set of referents $\mathcal{R}$, a lexicon $\mathcal{L}$, an intention from the set of referents $i \in \mathcal{R}$, and an order of pragmatic inference $n \in \mathbb{N}$.

Output: A signal, $s \in \mathcal{S}$ which is most likely to be understood by a listener, given the intended referent.

$$S_n(s|i) \propto \frac{L_{n-1}(i|s)}{\sum\limits_{s' \in \mathcal{S}} L_{n-1}(i|s')}$$

The active agent uses argmax to choose the specific signal from the distribution. The reactive agent chooses the second best - the confirming - signal. Essentially, a speaker will recursively infer the optimal signal until they consider the literal listener (base case).

Listener

A listener will adjust its behaviour depending on whether it is the active or reactive agent. Further, to avoid confusion, I will refer to the active agent's signal as the *original* signal - s_o, and to the reactive agent's signal as the *confirming* signal - s_c, to conform with the name of this model.

Input: A set of signals $\mathcal{S}$, a set of referents $\mathcal{R}$, a lexicon $\mathcal{L}$, a signal $s \in \mathcal{S}$, and an order of pragmatic inference $n \in \mathbb{N}$.

Output: The most likely referent $r \in \mathcal{R}$, given the signal. Both agents will infer the most likely referent, given the signal, by:

$$L_n(r|s) \propto \frac{S_n(s|r) \cdot w(s,r) \cdot \lambda}{\sum\limits_{r' \in \mathcal{R}} S_n(s|r') \cdot w(s,r')}$$

Where w is an agents associational weight between signal and referent. Initially, this weight is uniform over all signal-referent pairs. Reasons for not calling this the "prior" will be discussed below. Since the reactive agent never updates her weights which therefore stay uniform, they can be omitted.

Importantly, the active agent will scale his weights by a bias parameter λ whenever he interprets the "confirming" signal. This means, the referent which the active agent referred to, has a higher probability of being inferred:

$$\lambda \in \begin{cases} \mathbb{R}_{>1} & \text{if listener} = \text{active agent} \\ 1 & \text{otherwise} \end{cases}$$

The recursive formula bottoms out at the literal listener:

$$L_0(r|s) \propto \frac{\mathcal{L}(s,r) \cdot w(s,r) \cdot \lambda}{\sum\limits_{r' \in \mathcal{R}} \mathcal{L}(s,r') \cdot w(s,r')}$$

Once the active agent infers the referent, r, from the reactive agent, he will compare it to his own original intention, i, and update his weights accordingly. This is done by:

$$w(s_{o,c}, r) \propto \begin{cases} w(s_{o,c}, i) \cdot \alpha & \text{if } i = r \\ w(s_{o,c}, i) \cdot \beta & \text{otherwise} \end{cases}$$

Where, $\alpha \in \mathbb{R}_{>1}$ and $\beta \in \mathbb{R}_{0<x<1}$. $s_{o,c}$ is the original signal by the active agent and the confirming signal by the reactive agent, respectively.

2.2.3 Success

A dialogue is a little more involved. Therefore, measuring success is a little trickier. A successful dialogue could be operationalised as when agent Nora infers agent Adam's intended referent correctly, and also Adam infers Nora's intended referent correctly. This is more difficult to achieve however, because the second inference depends on the first. Let's make a case distinction to see which scenarios can occur. All scenarios start by Adam having an intention $i_1 = r_1$ and him producing a signal s_1. The coloured arrows indicate whether the intermediate communication was done successfully.

Scenario 1 Everything goes smoothly; this is the ideal case. Nora picks up Adam's signal, interprets r_1 to be Adam's intended referent and produces a secondary signal s_2. Adam picks up the signal and interprets r_1. Since this is equal to Adam's original intention, Adam concludes that Nora has understood her. Note that here, $w(s1, r1)$ and $w(s2, r1)$ get positive reinforcement.

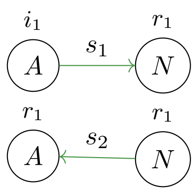

Scenario 2 Nora infers Adam's referent correctly, but when Nora replies, Adam makes an incorrect inference. Adam can not distinguish who has made a mistake. Although actually no misunderstanding occurred, Adam thinks it did and therefore the associated weights will get negative reinforcement.

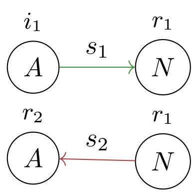

Scenario 3 Nora fails to infer Adam's referent and so sets another referent r_2 as his own intention. Nora produces a signal accordingly and Adam interprets the signal correctly. Again Adam can not distinguish who has made a mistake. This time however, the associated weights will rightly get a negative reinforcement.

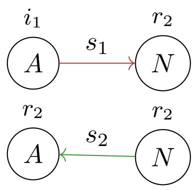

Scenario 4 Both agents fail to infer the intended referent. However, by coincidence Adam interprets Nora's signal as his intended referent. Here, Adam assumes everything went smoothly and cannot distinguish it from scenario 1. Therefore, the weights get an incorrectly assigned positive feedback.

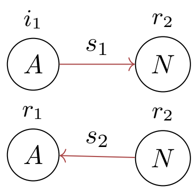

Scenario 5 Both agents fail to infer the intended referent. Adam detects a mistake but cannot attribute it. Rightly, the weights get negatively reinforced.

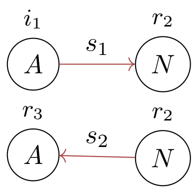

The active agent cannot properly distinguish between a successful and a failed dialogue i.e. he would interpret scenario 4 as a success, although in reality it isn't. Still, a success was operationalised as only when both inferences are made correctly as in scenario 1.

2.2.4 Implementation

In the Confirmer model, both agents have a personal lexicon. A lexicon is discrete (containing only zeros and ones) and is represented as a matrix in which each row is a signal and each column, a referent. The lexicon is produced by creating a random vector consisting of $\approx 70\%$ zero's and $\approx 30\%$ 1's. This gives the lexicon moderate ambiguity ($\approx \frac{1}{3}$), i.e. there is no 1-to-1 mapping between signals and referents [3]. Each referent can be referred to by at least one

[3] See the appendix for a more comprehensive discussion of lexicon creation.

signal. Further, the agents share the same lexicon. Although there certainly is asymmetry between people's lexicons, their beliefs about the world and therefore would have been plausible to include, it is an additionally complicating variable, which I was not primarily interested in. Note, that the Pointer model is defined such that an agent always considers all possible lexicons and thus there is no such measurement as ambiguity and asymmetry. Initially, I have simulated trials with 8x8 lexicons. Conceptually though, I have found it more plausible to use lexicons with more signals than referents and therefore opted for 16x8 lexicons. If we take the Fribbles in figure 1.4 as an example, we can easily see that there are exactly two referents but that one could think of a possibly infinite number of ways to refer to either of them. Crucial in this model, is the assumption that the active and the reactive agent never switch roles. This is to reflect the Fribbles task where one person tries to refer to one of the objects while the other tries to infer which one is meant. Consequentially, only the active agent ever updates his lexicon. Ideas of how to include the reactive agent are outlined in the discussion. To model the notion that Adam assumes that the conversation is still on the same topic or in terms of the Fribbles task, to give the benefit of the doubt that Nora interpreted Adam's signal correctly, the bias λ was set to 1.2. This has the effect of nudging Adam's inference towards his own intention. It can be seen as a slight confirmation bias. Increasing the bias by too much would make Adam oblivious to the signal; he would always choose his own intended referent. Usually, Bayesian inference utilizes the posterior distribution as the new prior distribution. Here however, it proved to be problematic, because I worked with discrete lexicons, rather than with probability distributions, which means that the posterior distribution can be zero. Using this as a prior is of course counter-effective since no new inference can be made (every probability will be zero). Instead, the associated weight matrix has a similar effect. The weight biases α and β were chosen as 1.1 and 0.9, respectively. This was merely done to give the agents a slight positive or negative reinforcement whenever a confirming signal is received and the inference is compared to the original intention. This is slightly tending towards a Hebbian "what fires together wires together".

3 Results

3.1 The Pointer Model

The Pointer model was tested against a number of parameters. The first simulation presented here, was run with 3 signals and 2 referents and with a conversation length of 5, i.e. there are 5 turns/ monologues in a conversation. Per conversation, the speaker keeps trying to convey the same intention. 1000 such conversations were run. The results are shown in figure 3.1. What stands out immediately, is that if a monologue was successful on the first try, the success rate stays at 1, i.e. the agents never make a mistake. This makes sense. If a signal was correctly interpreted by the listener, the speaker has no reason to change it, as if a convention was formed. This aligns with the original paper. However, when the first monologue is unsuccessful, there is always a steady incline in the success rate from that point on. This also makes sense. The probability of getting the referent right by chance is $1/R$ where R is the number of referents, thus 0.5 in this case. So either the listener gets it right straight away, or he gets it wrong. Then the speaker would deem the signal used as unfitting for that particular referent and use another one. This necessarily signals to the listener that the other referent was meant. This is illustrated by the histogram in the same figure, which shows that about half the conversations start off correctly.

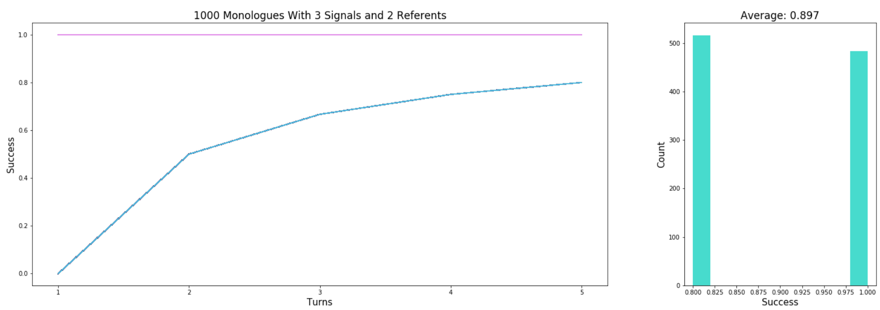

Figure 3.1: Success Rate Over 5 Turns With 3 Signals and 2 Referents

The second simulation was run with similar parameters, with the only difference that 3 referents were used. The results are shown in figure 3.2. Likewise, the agents have a successful monologue, either on the first, second or third turn. The chance level in this configuration is $\frac{1}{3}$ and this is reflected in the outcome. It seems that the binomial prior does not have much of an effect. This makes sense retrospectively, since it is symmetric. The probability of any lexicon's matrix is the same as its inverse. So initially, the signal-referent pairs have to be guessed. It is as if the listener reveals one signal-referent of their lexicon with each turn. Using Bayesian inference, this eliminates a lot of lexicons from the possible choices and thus diminishes their likelihood.

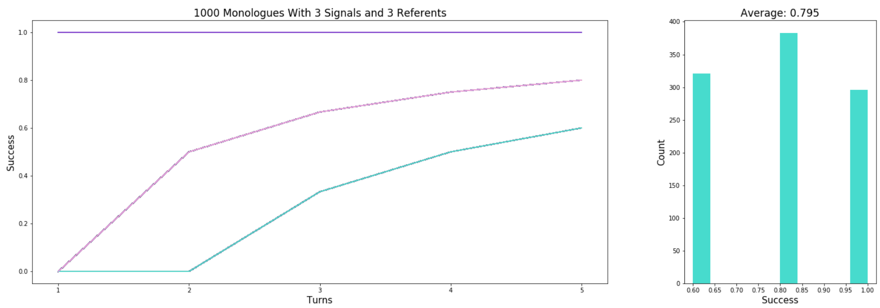

Figure 3.2: Success Rate Over 5 Turns With 3 Signals and 3 Referents

This implementation was tested against Thijme de Valk who implemented the same model and got similar results [13].

3.2 The Confirmer Model

The Confirmer model, along with the dialogue framework was designed in order to investigate the mechanics of the resolution of misunderstandings. Since a dialogue consists of two consecutive monologues, testing the complete model is done incrementally. First, a monologue was run 10.000 times, with a success rate of 0.82 (see figure 5.3 in the appendix). In a monologue Adam tries to convey his intended referent and Nora tries to interpret it. Without explicit feedback, there is no new information which feeds back into the conversation and to the next inference. Therefore, each turn has an equal probability. Then, logically, a plain dialogue, not employing the updating strategy, i.e. $\lambda = \alpha = \beta = 1$, should be $0.8^2 = 0.64$. Another 10.000 simulations revealed that the accuracy of a plain dialogue is pretty close to our prediction with an accuracy of 0.63 (see figure 5.4 in the appendix). Now we know the accuracy of dialogues, when Adam does not adjust his beliefs when Nora gives a confirming signal. There are two distinctive mechanisms to the model. First, the active agent's assumption or, confirmation bias, that the reactive agent continues the conversation. This is reflected by the bias parameter λ. Another 10.000 conversations were simulated, only including this bias parameter. This resulted in an accuracy of 0.72 (see figure 5.5 in the appendix). Now, the actual updating mechanism was included. Another 10.000 simulations showed, that there is no significant benefit over just the bias. Ideally, Adam would choose a better signal after a negative reinforcement and improve the conversation. Figure 5.6 shows however, that virtually any trajectory is possible. The figure shows, that even when there was a success on the first turn, the agents are able to make mistakes afterwards. They are able to improve but they can also deteriorate the conversation. Nonetheless, to analyse the results on a deeper level we can dissect the outcome and individuate the scenarios. The scenarios in which Adam detects a misunderstanding are of importance to us since these are the moments when the conversation can be reconciled. Recall, that in scenario 3 the first monologue is unsuccessful but the second monologue is a success. This corresponds to the original example in the introduction in which you referred to your cat but Nora interpreted your signal as you talking about a professor. When she responded however, you realised this. Figure 3.3 shows all occurrences which started with scenario 3 in the first turn, i.e. in the first dialogue. This seemed to have happened around 800 times out of 10.000 simulations. Ideally, through the negative reinforcement of the signal-referent pairs, the active agent would have corrected himself and would have used a different signal to refer his intention. However, we can see that this is not the case. Although there is some effect, i.e. scenario 2 (first monologue successful, second monologue unsuccessful) and scenario 4 (both monologues unsuccessful, but Adam does not recognise it) are more likely to occur. Especially interesting is that scenario 2 is likely to occur as the conversation continues. This means Adam used a different signal to refer to his intention and Nora interpreted it correctly. So it did actually resolve the misunderstanding. However, when she gives another confirming signal, which may be the same as before, Adam interprets it incorrectly. This might be due to the negative reinforcement of the confirming signal. In any case, more often than not scenario 3 repeats in consecutive turns.

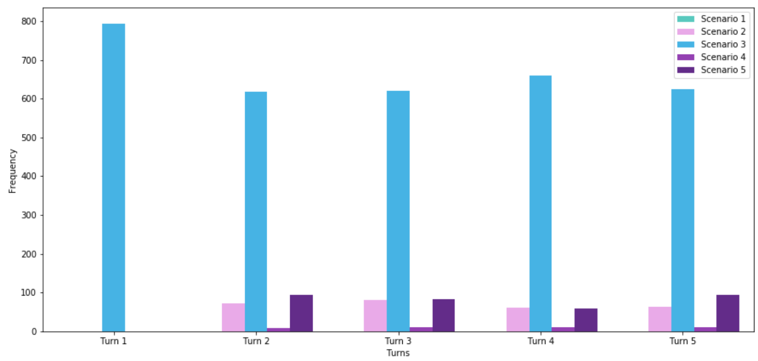

Figure 3.3: All Conversations In Which Scenario 3 Occurred On The First Turn

The other time in which Adam detects a misunderstanding and tries to rectify the conversation is scenario 5. Here, both monologues go awry. We can see in figure 3.4 that this happened around 700 times on the first turn. The subsequent turns show, that Adam has not successfully remedied the conversation. However, we see that either scenario 3 or 5 occur in these following turns. This switching between scenarios might be due to the fact that a single monologue has a probability of 0.8 of being successful. So it is more likely that at least one of the monologues are a success.

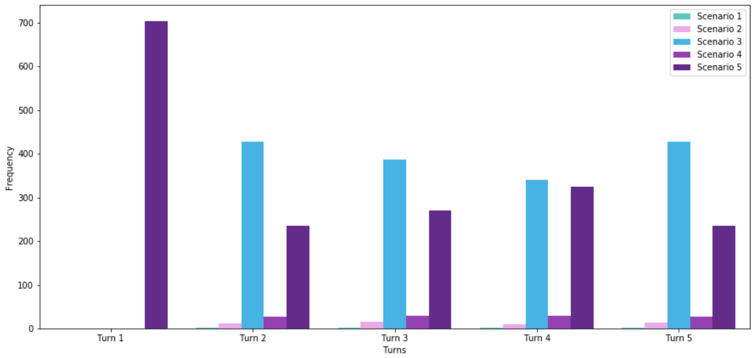

Figure 3.4: All Conversations In Which Scenario 5 Occurred On The First Turn

4 Discussion

4.1 Explicit Feedback

Simulations of monologues using the Pointer model result in a high average of success. Perhaps the most important accomplishment of this method is that both agents "learn", they form lexical conventions. The success rate increases with consecutive turns. This convention-formation though, is only possible due to explicit feedback. This is the unrealistic assumption which is avoided in the Confirmer model. Explicit feedback resembles a game of "Guess Who"[14] in which signal-referent combinations are tested to limit possibilities. Conceptually, the two agents converge to a shared lexicon by using the observed signal-referent pairs to eliminate improbable lexicons. Without choosing a proper prior this can also lead to an arguably unwanted side-effect. Each agent contemplates a new signal-referent pair distribution given prior signal-inferred-referent pairs. Here, it does not matter what intention the speaker has. Whatever referent is inferred by the listener, that signal-referent pair will have formed a "connection". From now on, the speaker will infer a high probability of using the same signal to refer to the same referent and the listener will infer a high probability of inferring the same referent when hearing that particular signal. Controversially, this may not be a bad way of initiating signal-referent pairs but how plausible is it? Imagine a speaker who would like to refer to a banana and signals "The yellow elongated fruit". For some bizarre reason (which may happen in this model unless there are specific priors introduced, which are not specified in the model), the listener infers "apple". From now on, both agents will give the "The yellow elongated fruit"-"apple" pair a high probability and a convention will have been formed. This notion complies with the first factor which they deem detrimental to the model: *arbitrariness.* But I would argue that communication is not arbitrary. Both agents converge to a shared probability distribution over signal-referent pairs. Eventually, the agents would establish a 1-to-1 mapping (second factor *stability*). It would be interesting to investigate the effects of the model with more variations in the number of signals and referents. It would especially be intriguing to inspect the resulting probability distributions over the signal-referent pairs.

4.2 Lexicon

Making use of the probability distribution over all possible lexicons paradigm has various implications. The computational complexity of such a regime, made it impossible for me to run simulations with larger lexicons. It would have been interesting to see whether the results for small lexicons extrapolate to larger ones. Earlier, I have indicated that there seems to be a conceptual conundrum of which regime is cognitively more plausible. Perhaps this result tells us something about the nature of the problem. If it is intractable (which it is not yet proven to be), there must be another way, humans engage in conversation.

The Pointer model certainly helped laying the foundation for the Confirmer model. However, the many adjustments, especially shifting to a personal lexicon paradigm lessened the resemblance. A personal lexicon has many consequences. The size and shape of the lexicon can be detrimental to the agent's success. The Confirmer model requires a referent to be able to be referred to by at least two signals. This is partly the reason why the simulations were run with 16 x 8 lexicons. As mentioned before though, this is plausible and giving the agents more signals would have been too. There were no restrictions on signals in the Fribbles task so this simulation imitates that quite nicely. Moreover, the amount of ambiguity in a lexicon and the asymmetry between two agent's lexicons are additional significant factors, determining an agent's success. Although I haven't experimented with these parameters, it is important to note that these are measures strictly pertaining to a personal lexicon paradigm [1].

4.3 Resolving a Misunderstanding

The purpose of this project was to find how we can resolve a misunderstanding. Investigating the problem lead me to some interesting findings and side-effects of dealing with RSA and related models. A detail, which needs to be highlighted is the distinction between semantically similar, yet different notions. To *resolve* a misunderstanding pertains to the act of fixing, reconciling, or rectifying a miscommunication. This happens after a misunderstanding has already occurred. *Avoiding* a misunderstanding, in contrast, pertains to the circumvention of such a miscommunication. This can happen even before a misunderstanding has occurred. Only the former of the two necessitates the *detection* of a misunderstanding. In scenario 1 for example, in which everything goes smoothly and there is no misunderstanding by either agent, the "confirmation bias" helps avoiding a misunderstanding. It imitates Adam's inkling that Nora is responding to his intended referent. Also, the inferred signal-referent pairs get a positive reinforcement. This further helps avoiding a misunderstanding since the chosen signal is now more likely to be used again to refer to the intended referent, as if a convention was formed. Nonetheless, the updating step also brings unfortunate consequences. There are two scenarios in which the updating is done incorrectly. In scenario 2 no misunderstanding occurred, although Adam assumes it. The weights are negatively reinforced and therefore Adam is less likely to choose the same signal to refer to his intended referent although Nora interpreted it correctly. Also in scenario 4 in which both monologues are unsuccessful and Adam incorrectly assumes they are not, the weights get a positive feedback, making it more likely to refer to the intended referent by the same signal again. Despite the wrong assumptions the active agent makes, these situations which emerge in the face of uncertainty, are plausible. The most relevant results however are of scenario 3 and 5. Here the agent rightfully detects a misunderstanding. Unfortunately, this model does not explain its resolution. Whether the model has to

[1]See more contemplations on this in the appendix.

be minorly tweaked or re-conceptualised is content for future work. The results of scenario 3, i.e. the instances in which scenario 3 occurred on the first turn, offer a little bit of hope in that the consecutive turns tend to scenario 2, at least sometimes. In a way, the misunderstanding is resolved, the problem is only that Adam doesn't know about it. Perhaps indeed only a slight tweaking is necessary to overcome this. Likewise, the results of scenario 5 show that scenario 3 is likely to occur right after. So if the first problem. i.e. going from scenario 3 to scenario 2, is solved, the second problem might follow and solve itself. In any case there are other adjustments which could be made to the model. It would certainly be beneficial to have the reactive agent's weights update as well, i.e. have Nora actively participate in the improvement of the conversation.

The information the agents gain in this model is also different from the Pointer model. The Confirmer model deviates from the three notions: arbitrariness, stability, and reduction. The agents don't comply with stability in that they choose a secondary, different, signal to confirm the conversation. Therefore, they do not converge to 1-to-1 mappings between signals and referents. Instead, they map out multiple attributes of an object in which signal-referent pairs can overlap.

4.4 Alternative Conversation

My intention was to model a conversation. When talking about an event, we respond with a related comment, something that continues the conversation. Also in the Fribbles task, during one conversation, the participants have to communicate about one object. I made the assumption that the active agent is resolving the misunderstanding. Another way to model the conversation would be to let the reactive agent resolve the misunderstanding. An altered version of the example from the introduction might illustrate the idea. Say you are at a dinner party, engaged in a thrilling conversation with the host. They tell you that the professor has been very tired lately, to which you agree. Then, they remark: "He has not been drinking his milk today". After a baffled moment you realise they have been talking about their cat named "The Professor". How did you quickly realise your mistake and shifted your belief to a more likely referent? Here, you take the role of the reactive agent. By simply changing the framework, the possible scenarios differ completely.

4.5 Conclusion

Modelling the resolution of a misunderstanding is challenging. We have seen, that not only the model itself but also the environment in which it resides has to be considered. One must determine, which party (if not both parties) of the conversation detects the misunderstanding and rectifies the conversation. We have also seen, that choosing such a framework implies novel situations from which new difficulties can arise. The Pointer model has provided a good base upon which to extend the model to include new capacities. The Confirmer model makes progress in terms of plausibility of our cognitive mechanisms in that it excludes explicit feedback and shifts to a personal lexicon paradigm. The dialogue framework offers scenarios in which a misunderstanding can occur and such that an active agent is able to detect it. The results show that although the problem of resolving a misunderstanding has not been solved completely, steps in the right direction have been taken. In future work it would be interesting to see whether adjusting the weights help the active agent distinguish between a successful and an unsuccessful monologue. Perhaps, for example, only the original signal- and original intention pair ($w(s_o, r)$) should be reinforced. Another extension to this model would be to involve the reactive agent by making her reconsider her inferences in consecutive turns. Hopefully, this research opens up new avenues to explore and provides a few insights into the problem of resolving a misunderstanding.

Acknowledgement

I would like to thank my supervisors Laura van de Braak and Iris van Rooij for the consistent meetings, feedback, and guidance over the whole course of the bachelor project. I would also like to thank Mark Blokpoel for advice during and before this project which gave me an advantageous insight. Further, I want so say thank you to the teachers of the Radboud University and especially to Franc Grootjen who has supported me and gave me many opportunities throughout the years. Thank you to my friends with whom it was a joy to do this bachelor. I also wish to express my appreciation for my parents who have supported me throughout the bachelor (and all the years prior). And last but not least I want to say thank you to Jess Leondiou who supported and pushed me and convinced me to do this study.

5 Appendix

5.1 Creating a lexicon

The *ambiguity* of a lexicon refers to the amount and distribution of signal-referent pairs. When a signal refers to more than one referent, a listener might not immediately know which referent the speaker is referring to. Ambiguity is operationalised by

$$\frac{1}{R}\sum_{R}\frac{1}{S}\sum_{S}\mathcal{L}(s,r)$$

This is based on the paper by Blokpoel et al. [6]. For example, in the left lexicon in table 5.1, we can sum each column and divide it by the number of signals, so we get $\frac{2}{4}$ for every column. Summing, and dividing by the number of referents gives $\frac{2}{4}\cdot 4 = 2$ and $\frac{2}{4} = \frac{1}{2}$. Therefore, the lexicon has an ambiguity of $\frac{1}{2}$. The *asymmetry* between two lexicons is a measure of the difference of the lexicons. It is operationalised as the number of unequal signal-referent pairs, over all signal-referent pairs.

There are a many ways to create a lexicon. However, creating lexicons for each agent while maintaining control over asymmetry and ambiguity can be quite tricky. To see the problem, consider the two lexicons in figure 5.1. Although, they both have an ambiguity of $\frac{1}{2}$, the left one has evenly distributed signal-referent pairs, while in lexicon on the right every signal refers to $r1$ and there is only one signal referring to $r3$ and $r4$, respectively.

	r_1	r_2	r_3	r_4
s_1	1	0	0	1
s_2	1	1	0	0
s_3	0	1	1	0
s_4	0	0	1	1

	r_1	r_2	r_3	r_4
s_1	1	0	0	0
s_2	1	1	0	0
s_3	1	0	1	0
s_4	1	1	0	1

Figure 5.1: Lexicons With Ambiguity of $\frac{1}{2}$

This problem, of matching signals to referents, can be conceptualised as a bipartite matching problem, for which a maximizing flow algorithm can help to circumvent the issue [15]. Here, it is implemented as follows. The agent's set of signals and referents are used to create a bipartite graph. Say, there are four signals and four referents, then aligning them and attaching a start and target node, would give something like in figure 5.2.

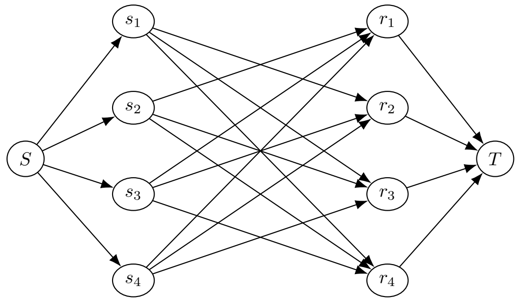

Figure 5.2: A Bipartite Graph For the Creation of a Lexicon

Notice that all nodes are connected, with the exception that there are no edges from s_i to r_j whenever $i = j$. These edges correspond to the diagonal in the matrix and were purposefully left unfilled. To let the agents have at least one signal for each referent in common, it can subsequently be filled. I have found it reasonable to assume that this is the case, even if the agents have to find the common signal first. Each edge from S and to T is initialised with a maximal capacity which is calculated by the formula

$$\frac{\frac{S \cdot R}{a} - min(S, R)}{R}$$

where S and R are the number of signals and referents, respectively. a is a parameter which helps determine the final ambiguity. In this example, a is set to 2 to get a final ambiguity of $\frac{1}{2}$. We get $\frac{\frac{4 \cdot 4}{2} - 4}{4} = 1$. This means four signal-referent pairs will be produced. Remember, that the diagonal with four additional pairs will be filled in. This will give us the desired ambiguity of $\frac{1}{2}$. Now a simple maximizing flow algorithm can be used to find the connecting edges which correspond to the signal-referent pairs. I have made use of the Edmonds-Karp flow algorithm [16] to do this efficiently. The resulting lexicon will be similar to the one on the left one in figure 5.1. Note, that for reasons of time constraints this method of producing lexicons has not been thoroughly tested for generalisation, but only for the simulations done in this research.

5.1.1 How Not To Create A Lexicon

Although the above method does work, it was abandoned. This was done primarily because of time constraints. Further, I noted an interesting side-effect of certain lexicons. When the number of entries in each row equals the number of entries in each column, no inference is possible, assuming a uniform prior. The probability of choosing any signal, e.g. in the left lexicon of figure 5.1, would be $\frac{1}{2}$ even after pragmatic inference, thus not increasing the information. Because this is merely a mathematical side-effect and does not correspond the real world, it was not further considered.

5.2 Additional Results

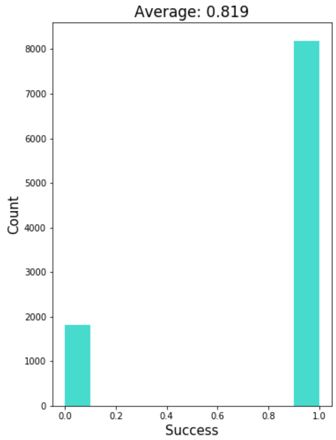

Figure 5.3: Success Rate of the Confirmer Model Over 10.000 Monologues and Using 16 x 8 Lexicons

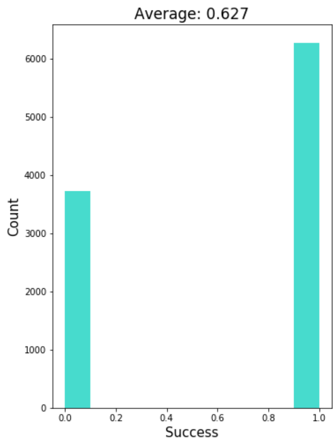

Figure 5.4: Success Rate of the Confirmer Model Over 10.000 Dialogues and Using 16 x 8 Lexicons

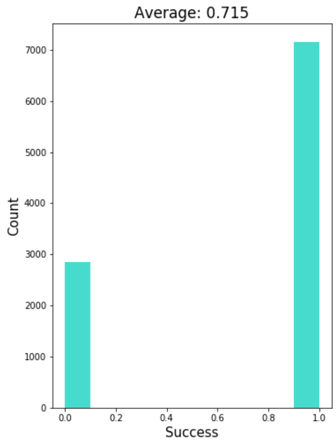

Figure 5.5: Success Rate of the Confirmer Model Over 10.000 Dialogues, Including the Confirmation Bias, and Using 16 x 8 Lexicons

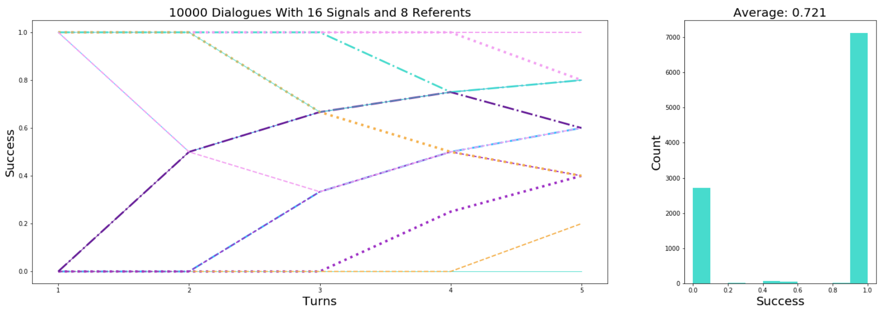

Figure 5.6: Success Rate of the Confirmer Model Over 10.000 Dialogues, Including the Confirmation Bias and the Updating Mechanism, and Using 16 x 8 Lexicons

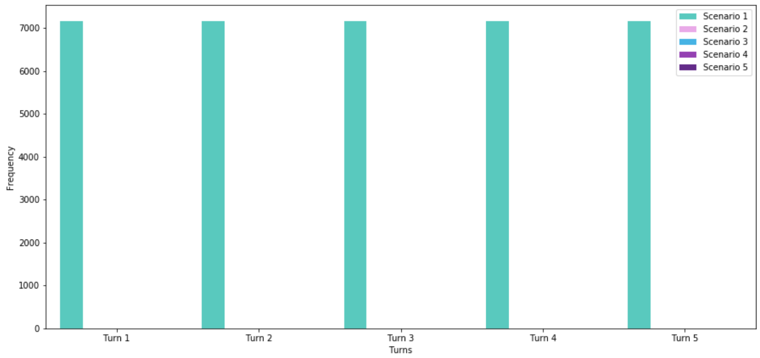

Figure 5.7: All Conversations In Which Scenario 1 Occurred On The First Turn

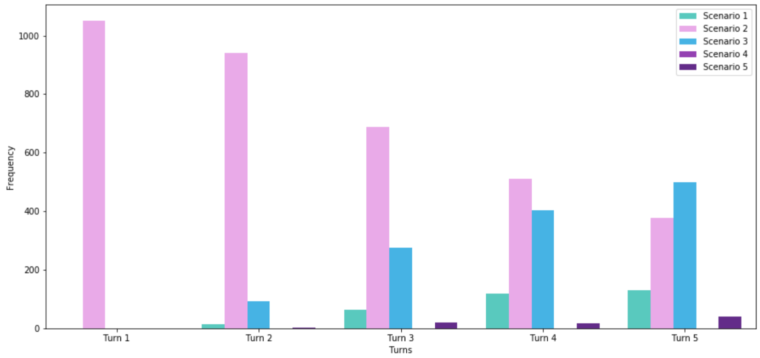

Figure 5.8: All Conversations In Which Scenario 2 Occurred On The First Turn

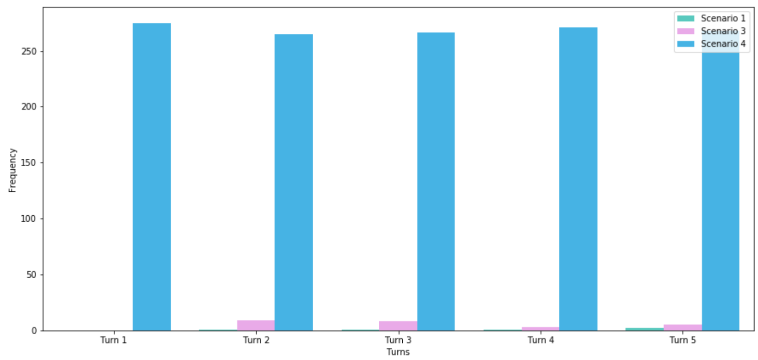

Figure 5.9: All Conversations In Which Scenario 4 Occurred On The First Turn

Bibliography

[1] Michael C Frank and Noah D Goodman. Predicting pragmatic reasoning in language games. *Science*, 336(6084):998–998, 2012.

[2] Herbert P Grice. Logic and conversation. In *Speech acts*, pages 41–58. Brill, 1975.

[3] Noah D Goodman and Michael C Frank. Pragmatic language interpretation as probabilistic inference. *Trends in cognitive sciences*, 20(11):818–829, 2016.

[4] Robert XD Hawkins, Andreas Stuhlmüller, Judith Degen, and Noah D Goodman. Why do you ask? good questions provoke informative answers. In *CogSci*. Citeseer, 2015.

[5] Michael C Frank. Rational speech act models of pragmatic reasoning in reference games. 2016.

[6] Mark Blokpoel, Mark Dingemanse, George Kachergis, Sara Bögels, Ivan Toni, and Iris van Rooij. Pragmatic communicators can overcome asymmetry by exploiting ambiguity. 2019.

[7] Robert XD Hawkins, Mike Frank, and Noah D Goodman. Convention-formation in iterated reference games. In *CogSci*, 2017.

[8] Tom Joseph Barry, James W Griffith, Stephanie De Rossi, and Dirk Hermans. Meet the fribbles: novel stimuli for use within behavioural research. *Frontiers in Psychology*, 5:103, 2014.

[9] The CABB Team. https://www.languageininteraction.nl/bigquestion3.html. `https://www.languageininteraction.nl/BigQuestion3.html`, 2020.

[10] Martin Odersky. Scala language. `https://scala-lang.org/`, 2004.

[11] Ron Hommelsheim. Resolving misunderstanding (remis). `https://gitlab.socsci.ru.nl/R.Hommelsheim/rsa-model`, 2020.

[12] Robert XD Hawkins, Mike Frank, and Noah D Goodman. Conventions. `http://forestdb.org/models/conventions.html`, 2017.

[13] Thijme de Valk. Base model. `https://gitlab.socsci.ru.nl/T.deValk`, 2020.

[14] Ora and Theo Coster. Guess who? `https://en.wikipedia.org/wiki/Guess_Who%3F`, 1979.

[15] Wikipedia. Matching (graph theory). `https://en.wikipedia.org/wiki/Matching_(graph_theory)`, 2020.

[16] Jack Edmonds and Richard M Karp. Theoretical improvements in algorithmic efficiency for network flow problems. *Journal of the ACM (JACM)*, 19(2):248–264, 1972.