UNIVERSITY OF GRONINGEN

MASTER THESIS

# Modelling Implicit Learning of Temporal Relations

*Author:*
Aaron Robert THIEL
3500101

*Supervisors:*
Dr. Jacolien VAN RIJ
Prof. Dr. Niels TAATGEN

*A thesis submitted for the degree of Master of Science in*

Computational Cognitive Science

August 25, 2023

# *Abstract*

Science and Engineering
Computational Cognitive Science

Master of Science

**Modelling Implicit Learning of Temporal Relations**

by Aaron Robert THIEL
3500101

Awareness of time passing between events informs our perception of our environment. While much research exists on how it guides daily decisions, it is not as clear how humans *learn* temporal relations in the first place. As such, this thesis focuses on integrating theories on temporal relations with existing models of implicit learning. The aim is to implement a formal cognitive model that includes temporal cues and allows for discriminate learning and transference. A first attempt at such a model is made, and an experiment is conducted to collect additional required data. The experiment involved participants playing an anticipation-based video game to capture under which conditions awareness of temporal relations leads to increased performance over time. The insight gained from the experiment is limited, but potential problems and future steps are discussed.

# *Acknowledgements*

I want to thank my supervisor Jacolien van Rij for her engagement and encouragement during our weekly meetings and for helping me to stay focused during the more hectic periods of the project.

I also thank my second supervisor Niels Taatgen for showing interest in the project at all our meetings and offering quick and precise expert knowledge. I also thank the members of the Cognitive Modelling research group for providing feedback and welcoming me to their meetings which greatly shaped my approach to research and writing.

Lastly, I would like to thank my parents for supporting me throughout my thesis, my degree and my education in general. Thanks also to my fellow RUG students, namely Gilles Lijnzaad, Leonidas Zotos, Gonçalo Carvalho and Elisabeth Blyumin.

# Contents

**Abstract** **iii**

**Acknowledgements** **v**

**1 Introduction** **1**
1.1 Implicit Learning . . . . . 1
1.1.1 Existing Theoretical Frameworks of Implicit Learning . . . . . 1
1.1.2 Modelling of Implicit Learning . . . . . 3
1.1.3 Limitations regarding Time . . . . . 4
1.2 Existing Theoretical Frameworks of Temporal Relations . . . . . 6
1.2.1 Contiguity Principle . . . . . 6
1.2.2 Temporal Coding Hypothesis . . . . . 6
1.3 Formalising the Temporal Coding Hypothesis . . . . . 7
1.3.1 Cue-Blocking . . . . . 7
1.3.2 Cue-Blocking involving different Temporal Relations . . . . . 8
1.3.3 Transference . . . . . 10

**2 Simulation Study** **11**
2.1 Base-EDL . . . . . 11
2.1.1 Simultaneous cues . . . . . 11
Why not Outcome-based encoding? . . . . . 13
2.1.2 Non simultaneous cues . . . . . 15
2.2 Adding Transference . . . . . 16
2.2.1 Method . . . . . 16
2.2.2 Results for replicating Amundson and Miller (2008) . . . . . 17
2.3 Discussion . . . . . 18

**3 Human Study - Methods** **21**
3.1 Experiment Design . . . . . 21
3.1.1 Phases and Conditions . . . . . 22
3.1.2 Timing and Intervals . . . . . 23
3.1.3 Alien Language . . . . . 24
3.1.4 Structure of Sequence . . . . . 25
3.1.5 Second Level . . . . . 25
3.2 Procedure . . . . . 26
3.3 Planned Analysis . . . . . 27
3.3.1 Expected Results . . . . . 28

**4 Human Study - Results** **29**
4.1 Non-game data . . . . . 29
4.1.1 Thought probes . . . . . 29
4.1.2 Memory Test . . . . . 29
4.1.3 Interviews . . . . . 30

4.2 Game data . . . 31
4.2.1 Elemental Phase . . . 32
4.2.2 Compound Phase . . . 33
4.2.3 Test Phase . . . 33

**5 Discussion** **37**
5.1 Limitations of Software . . . 37
5.2 Limitations of Experiment Design . . . 38
5.3 Interpreting Results . . . 40
5.4 Future Research . . . 40
5.4.1 PRIMs . . . 41
5.5 A waste of time? . . . 42
5.6 Conclusion . . . 42

**A Model Output** **43**
A.1 Elemental Phase . . . 43
A.1.1 Level 1 . . . 43
A.1.2 Level 2 . . . 45
A.2 Compound Phase . . . 47
A.2.1 Level 1 . . . 47
A.2.2 Level 2 . . . 49
A.3 Test Phase . . . 51
A.3.1 Level 1 . . . 51
A.3.2 Level 2 . . . 53

**Bibliography** **55**

# 1 Introduction

Most people use knowledge of temporal relations daily, even if they do not realise it. For instance, when powering up a personal computer, a user knows that the desired result will occur not immediately, but only after a certain time. Similar knowledge is used when waiting for a response after ringing a doorbell or filling a bucket with water; In each case, a person has an expectation of how long the process should last. We primarily notice these expectations when they are not met, such as if the computer screen is still black a minute after the user presses the power button, no one opens the door for us, or the bucket never fills up because it has a hole.

While this ability to make use of time estimation has been studied (e.g. Ramscar et al., 2010; Van Rijn & Taatgen, 2008), it is less clear how we learn such connections in the first place. They are generally not explicitly taught to us, nor can we actively recall the knowledge they represent at will in most cases (e.g. knowing how many milliseconds one would wait for someone to answer the door.). This thesis aims to investigate how humans might form expectations of temporal relations using implicit learning. To this end, the following section will outline how learning without awareness occurs, how we may apply this process to time passing between events, and how to create a formalised theory that explains our implicit knowledge of time.

## 1.1 Implicit Learning

### 1.1.1 Existing Theoretical Frameworks of Implicit Learning

In general, *implicit learning* describes the process of learning complex information about the environment without the learner being aware of the process or being able to stop it. For instance, when learning to ride a bike, the learner has to remember a sequence of bodily movements to keep the bike moving and upright. However, from the learner's point of view, the only observable change may be that their general ability to ride a bike increases with each attempt.

A special case of implicit learning is *classical conditioning* as proposed by Pavlov (1927), which involves learning of cause and effect relations. According to this theory, a learner may form an association between two stimuli after repeated exposure to co-occurrences, such that responses previously triggered by only one stimulus will eventually also be elicited by the other. For example, in the experiments conducted by Pavlov, a dog was repeatedly exposed to co-occurrences of a ringing bell and the arrival of food. Pavlov inferred that an implicit association had been formed in the dog's mind after training, because the dog showed signs of anticipating food even if only the bell was presented (Pavlov, 1927).

In the terminology used by Pavlov, the smell of food is called the *unconditioned stimulus* (US), which triggers an *unconditioned response* (UR) of salivation. The same response is called a *conditioned response* (CR) if it is triggered by a *conditioned stimulus* (CS) created through training, such as the bell. An important assumption here is that the response is originally a reflex, meaning that the dog can not make itself salivate through any explicit reasoning it might have. As such, it is assumed that the observed anticipation is part of the automatic process of implicit learning, rather than the dog applying conscious reasoning to infer that it is about to receive food.

One approach to explaining how implicit learning works is *associative learning*, which assumes that the strength of an association between two stimuli is defined by the cumulative number of previous co-occurrences (Barto et al., 1981; Tesauro, 1995). Ideas similar to this approach to learning were theorised as early as 1898 in what would later be known as Thorndike's *Law of Effect* (Thorndike, 1898; Thorndike, 1927), laying part of the groundwork for modern reinforcement learning as well as machine learning.

However, while the associative approach captures the basic aspects of implicit learning, it has limitations when explaining certain learning phenomena found in animals and humans. For example, learners are usually able to ignore constantly occurring background stimuli, unlearn associations when presented with conflicting data or learn from currently absent stimuli (Nixon, 2020; Ramscar et al., 2013; Rescorla, 1988). As discussed by Rescorla (1988), learning theories based on simple co-occurrence do not provide mechanics to explain these more complex learning phenomena and an alternate, interference-based approach such as *discriminative learning* is needed. The approaches proposed by Rescorla and Wagner (1972) or Widrow and Hoff (1960), for instance, are able to explain the mentioned learning phenomena more efficiently, because they assume that the strength of associations between stimuli is dependent not just on a high number of co-occurrences, but also on if the information provided will be useful for minimising prediction errors in the future. Essentially — while with associative learning, the count usually increases the same amount for each learning event — under a discriminative approach, associations will change more drastically if there is a large difference between the expected and observed outcomes.

Since it's conception, discriminative learning has been successfully used to study several different forms of human learning such as category learning, interference reasoning and social psychology (reviewed in Siegel & Allan, 1996). Even neuroscientific evidence for the role of recognising error in predictions has been found (reviewed in Schultz, 1998), further giving credit to the theory within the context of cognitive modelling.

To give a practical example of how the discriminative approach differs from an associative one: If a conditioned stimulus $A$ occurs with an unconditioned stimulus $B$, then under both approaches, the connection $AB$ will become stronger — meaning that both cues start to produce the response. However, the approaches will predict different behaviours if $A$ is then also shown to occur together with a new stimulus $C$. Under an associative approach, the new association $AC$ will form incrementally, with no change to the already established association $AB$ (though, depending on implementation, $AB$ may decrease over time if there are no further co-occurrences of $A$ and $B$). In other words, there is now an expectation that $A$ should occur with $B$ or $C$. In the discriminative approach, on the other hand, once $C$ is introduced, the connection $AB$ will be immediately unlearned because of the large difference between what was expected ($A$ always occurring with $B$) and what was encountered ($A$ now also occurring without $B$). By the end of the learning process, the learner should only expect $A$ to occur with $C$ (for clarity, see Figure 1.1 for a visual representation of this difference in outcomes).

To sum up, the ability of the discriminative approach to quickly respond to perceived mistakes in expectations is the key difference which sets it apart from the associative approach, when explaining implicit learning phenomena beyond classical conditioning.

| | |
|---|---|
| $\boxed{A \rightarrow B}$ | $\boxed{A \rightarrow B}$ |
| $\downarrow$ | $\downarrow$ |
| Knowledge:<br>*A predicts B* | Knowledge:<br>*A predicts B* |
| $\downarrow$ | $\downarrow$ |
| $\boxed{A \rightarrow C}$ | $\boxed{A \rightarrow C}$ |
| $\downarrow$ | $\downarrow$ |
| Knowledge:<br>*A predicts B **and** C* | Knowledge:<br>*A predicts **only** C* |
| associative learning | discriminative learning |

Figure 1.1: Simplified comparison of differences in learning outcomes for the two approaches. In the associative approach, both outcomes are learned, while in the discriminative approach the original connection is unlearned.

### 1.1.2 Modelling of Implicit Learning

Both approaches to implicit learning can be mathematically formalised to create cognitive models using a simple two-layer neural network (Baayen et al., 2011; Hoppe et al., 2022; Klopf, 1988). In such implementations, the first layer represents all possible cues and the second layer represents all possible outcomes. As a first modelling step of our aim to model the learning of temporal relations, a network implementation is preferred over a more holistic symbolic representations such as ACT-R. This is because we are interested in modelling the sub-symbolic aspects of the learning process namely the dynamic created by the presence of certain cues and expectations. A model of how the learned information fits into the larger cognitive process would be unnecessarily complex for this purpose.

As previously described, the difference between associative and discriminative models is in the way the associations (i.e. network weights) change after each learning event. For all models, we assume that after each learning event, the weight $V_{ij}^t$ between cue $i$ and outcome $j$ is updated by a change $\Delta V_{ij}^t$, as defined by some learning rule:

$$V_{ij}^{t+1} = V_{ij}^t + \Delta V_{ij}^t \tag{1.1}$$

In a basic associative model, if a learning event includes cue $i$ leading to outcome $j$, then the network's connection weight $V_{ij}$ will increase by some set value, which we will call the learning rate $\eta$ (sometimes also called $\alpha$). Such a model would capture the counting of co-occurrences for every possible combination of two items, but nothing more; No information is gathered (or rather, no change in knowledge occurs) regarding any items that are not explicitly shown to relate to another.

A more advanced version of this approach would be to also include a negative reward whenever a prediction is wrong, i.e. is when two items do *not* co-occur. Such an implementation offers a simple version of *reinforcement learning* (Farley & Clark, 1954; Thorndike, 1911), with a simple positive and negative reward. In these models, a learning event involving both $i$ and $j$ would increase $V_{ij}$ by the learning rate $\eta$, while in a learning event including only $i$ without $j$, the change to $V_ij$ would be negative $\eta$. Figure 1.2 shows a visualisation of this difference in weight changes.

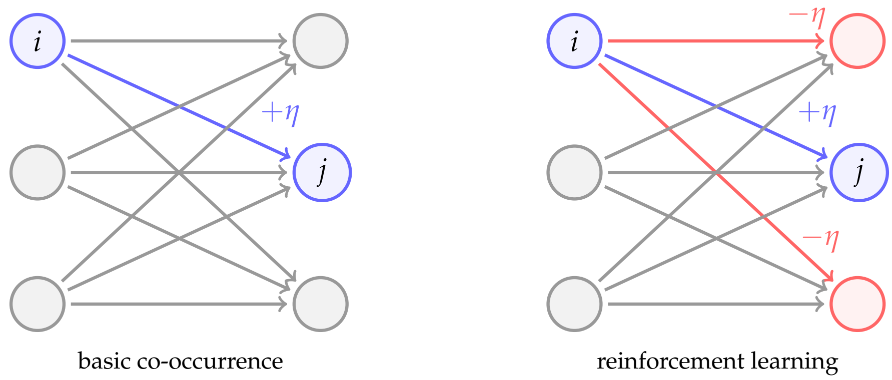


Figure 1.2: Visual representation of differences in the two learning algorithms for a learning event, including $i$ and $j$; Each arrow represents the association between two nodes. Blue and red arrows represent strengthened and weakened connections, respectively. Grey arrows represent unchanged connections.

A further adjustment to this learning rule, is to adjust the weight not by a fixed amount, but by different values depending on each specific learning event. For instance, the learning rule used within the cognitive architecture *ACT-R* uses the elapsed time since the last occurrence of an item to adjust the magnitude of change on the learning weights (Anderson et al., 1997). In this thesis, we will be using *Error-Driven Learning* which implements a form of discriminative learning by changing the connection weights based on positive or negative evidence, *and* with the magnitude of this change being determined by the size of error. Within these models, a learning event involving both $i$ and $j$ will increase $V_{ij}$ by a large amount if the co-occurrence of $i$ and $j$ was not already expected by the model, with the same being true for decreases to $V_{ij}$.

This dynamic is usually formalised based on the *delta-rule* as discussed by Widrow and Hoff (1960). An example implementation of this rule for a network with discrete cues and outcomes and activation between zero and one looks as follows:

$$\Delta V_{ij}^{t} = \begin{cases} 0 & \text{, cue } i \text{ absent} \\ \eta(1 - act_j^t) & \text{, cue } i \text{ and outcome } j \text{ present} \\ \eta(0 - act_j^t) & \text{, cue } i \text{ present but outcome } j \text{ absent} \end{cases} \tag{1.2}$$

Here $\Delta V_{ij}^{t}$ is the change in weights between cue $i$ and outcome $j$, which is calculated differently depending on which of either items is present in the learning event. The part of the equation $act_j^t$ is the overall activation of outcome $j$ at point in time $t$ which is given as the sum of all weights leading from any of the present cues to $j$:

$$act_j^t = net_j^t = \sum_{x \in cues(t)} v_{xj}^t \tag{1.3}$$

Further parameters may be added to these equations for more complex models. However, the basic dynamic of considering the existing activation, that is, the previous expectation, during learning remains the same in all discriminative models.

### 1.1.3 Limitations regarding Time

Having outlined these different approaches to network-based models, it may seem that one or even a number of them offer a sufficiently accurate way of modelling implicit learning. However, by their nature, network models do not consider the dimension of *time* during learning. Instead, learning instances are encoded as abstract events consisting only of a singular moment, allowing for no distinction between cases of an outcome following a cue immediately and those where it occurs after some time has passed; The changes in connection weights are the same in either scenario.

Naturally, this is an abstraction from real life, where learning does not happen as instantaneous events, but rather as an ongoing process. However, our goal is more complex than just zooming in to a deeper level of detail. Returning, for example, to Pavlov's experiments, if we imagine a scenario where the experimenter plays not one but two auditory cues before presenting the food, it must be decided whether it is more accurate to treat this as a single learning event or two which overlap. An ideal model would be able to correctly handle either case, without the need for special hard-coding on the part of it's creator or the user.

Time-decay-based models such as the already mentioned ACT-R are less discreet than a network representation and offer ways of handling such overlapping events. However, these systems primarily use time to *inform* learning rather than something that is in itself learned. Temporal modules can and have been added to ACT-R (Taatgen & van Rijn, 2011), but, as outlined in Section 1.1.2, these models do not allow us to study the specific changes to individual weights in the same way that network models do. In short, there is value in incorporating time into a network representation itself. Indeed, associative network models that attempt to do so already exist. Single recurrent networks, for example, capture learning phenomena involving sequencing and memory (Elman, 1990) and comparable implementations using a discriminative model are possible (Thiel & Spenader, 2022). However, models such as these merely capture that something has occurred at some point in the past but not the exact distance in time between events; They still function on a level of abstraction that cannot capture specific nuances involved when learning specific rhythms or time intervals. In spoken conversation, for example, two people can usually estimate when it is their turn to speak based not just on what the other person is saying but also partly on the length of pauses between words. Responding to a question immediately or with hesitation can convey different meanings. It is easy to imagine further examples with similar dynamics, such as when two people simultaneously reach for the same object or sing a melody together. All these examples involve knowing not just *if* something has already happened but also *when* and how long ago it occurred. Therefore, we must only further consider models which can use these relational forms of temporal knowledge as well.

Out of all of these, the most promising approach, perhaps, would be the cognitive architecture *PRIMs* (Taatgen, 2013), which simulates learning as an ongoing process rather than a sequence of self-contained events and can simulate both implicit and discriminative learning, when using a modified learning rule. Section 5 will return to the topic of PRIMs and how we may turn it into a discriminative model, but for now, we will stay on the smaller scale of network representations. To start with this, the next section will discuss how exactly to describe temporal relations.

## 1.2 Existing Theoretical Frameworks of Temporal Relations

### 1.2.1 Contiguity Principle

When Pavlov initially proposed his theory of classical conditioning, he already attempted to explain how learned associations might be affected by different temporal relations. He proposed the *contiguity principle*, which assumes that the closer two stimuli are in time, the stronger their (non-temporal) association will be. This framing of temporal proximity as a multiplier of the regular conditioning enables distinction between cases where a CS is followed by a US after just one second and those where the elapsed time before the US appears is ten seconds.

However, if only the contiguity principle were relevant, it would make no difference if the CS occurs followed two seconds later by the US or if the US goes first and then the CS, as long as the temporal proximity between the items is kept the same in both cases. Similarly, we would expect that if the CS and US appear simultaneously (i.e. have the highest possible temporal proximity), this association would be learned faster than all others. However, for both of these assumptions, Pavlov himself reported contradictory findings, discrediting his theory (Molet & Miller, 2014; Pavlov, 1927).

### 1.2.2 Temporal Coding Hypothesis

A better-suited theoretical framework of how temporal relations influence implicit learning is the *Temporal Coding Hypothesis* (TCH; see Matzel et al., 1988; Molet & Miller, 2014; Savastano & Miller, 1998). Instead of using the temporal relation between two stimuli as an amplifier of their association, under the TCH, this relation is part of the association itself. Essentially, the time interval between cue and outcome becomes another cue during learning, with the ability to hold useful or redundant information when distinguishing between different learning events. Under this principle, two learning events which involve the same stimuli, but differ in the interval between them, can easily be distinguished. For example, Pavlov's dog would be able to distinguish a scenario where the bell rings and food is presented after ten seconds and already expecting food just three seconds after the sound of the bell. The TCH would even capture compound effects, where two cues hold information that neither holds separately. An example of this could be Pavlov's dog only reacting to the sound of a bell after ten seconds, but not the sound of his owner unlocking the door with the same temporal interval. Beyond these dynamics, the TCH also correctly models both cases that Pavlov's contiguity principle fell short on, as the dynamics of changed order and simultaneous occurrence are correctly captured.

It is worth noting at this point that surrounding the idea of encoding time, further alternatives and variations of the TCH exist, such as for example the concept of *temporal maps* (e.g. Honig, 1981). Indeed if we widen our scope to general cognitive models of temporal relations, a multitude of alternative views opens up (for an extensive review see Block, 2014). However, these other approaches are often less developed, with the TCH appearing both as the most researched and as the most compatible with existing cognitive modelling ideas, based on its assumptions. The TCH assumes that associations between two events form based on contiguity (i.e. co-occurence), that the temporal relation between the same two events is somehow stored with or as part of the association, and that this information about temporal relations will influence the magnitude or timing of responses at any future presentations of either event. Beyond this, it is also assumed that people are able to combine or transfer knowledge about different temporal relationships regardless of if the knowledge was acquired separately. So for example, if two items both predict the same event at the same time, then this will lead to a stronger prediction overall. Similarly, if two otherwise identical learning events have different temporal relations, this difference will be reflected in the type and strength of response. What is also meant by this assumption, is that learners may make inferences about temporal relations they were never presented with. We will return to these more advanced ideas later, but for now, we can summarise that assumptions such as temporal relations acting as cues with predictive power, make the TCH a natural fit for discriminative models, which already distinguish between different learning events in a similar way (Molet & Miller, 2014). Crucially however, the TCH is just that — a set of assumptions, and not a formalised cognitive model. As we shall see in the following section, there is insufficient data to combine the two ideas seamlessly.

## 1.3 Formalising the Temporal Coding Hypothesis

Suppose we follow the assumption laid out by the TCH, that temporal relations are part of the cue-set of a learning event and can therefore act as a marker to distinguish one learning event from another. If this is the case, changes in temporal relations should affect other learning phenomena that rely on the presence of specific cues. As outlined in Section 1.1.1, the discriminative approach can explain many of the more complex learning phenomena involving implicit learning. What a number of these phenomena have in common is that they involve learners seeming to ignore frequently occurring pairings because they do not provide useful information to the model or do not contradict anything previously learned. Examining the exact mechanics behind one such learning phenomenon will show where implementing the Temporal Coding Hypothesis becomes more difficult.

### 1.3.1 Cue-Blocking

One learning phenomenon shown to occur during implicit learning for humans and animals is *cue-blocking*. This phenomenon involves an already strong established cue-to-outcome relation (between a CS and a US) blocking any new cues from becoming seen as reliable predictors (Kamin, 1967a, 1967b; Nixon, 2020).

For instance, we return once again to Pavlov's dog who has already gone through the so-called elemental phase of being trained with the sound of a bell as a reliable predictor for the arrival of food. Now the dog enters a second learning phase, called the compound phase, where a new cue, such as a blinking lamp, is shown simultaneously with the initial cue of the bell. If cue-blocking occurs, the dog should not form an association between the lamp and the food, meaning that, by the end of the learning process, the dog would not show signs of anticipating food when presented with the lamp alone. If the bell sound is present (either by itself or with the lamp) the dog will still react as before, making clear that it is not the case that the dog has become incapable of anticipation.

In reality, of course, it is also possible that partial blocking occurs so that by the end of the training, both cues elicit some level of the expected conditioned response. A discriminative approach explains this partial blocking by proposing that at the end of each learning step, there is some amount of uncertainty left which can be filled by further information. If more uncertainty is left, new information can lead to a more drastic reshaping of expectations. Figure 1.3 shows a visualisation of the level of certainty when predicting two different cues over time. The main point shown here is that the maximum certainty reached by the new cue introduced in the compound phase, is the same as the amount of uncertainty still left by the first cue in the previous phase.

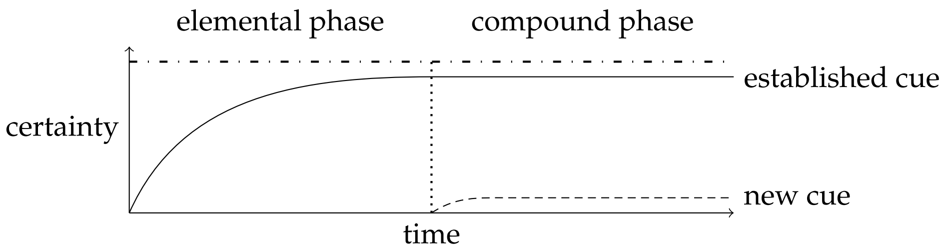


Figure 1.3: Example of the cue-blocking phenomenon. In the elemental phase, only one cue is presented, allowing a strong connection to form. During the compound phase, two cues are presented together, but only a small connection is formed for the new cue. The horizontal dashed line represents the maximum available certainty for all cues, showcasing how the certainty available to the new cue is equal to the remaining uncertainty of the previously established cue.

What is really meant by uncertainty here is the expected error based on the current network weights. If the weight between two items is already very high, it will only change if it is unambiguously at odds with newly encountered information. It should be noted that all this naturally emerges from the learning rules as previously described in Section 1.1.2; After each learning event, the weights in the model are increased depending on the amount of expected activation. Any learning event which includes an established cue to outcome relation will lead to a high activation and a small error, regardless of the presence of other cues. Because the error of the learning event is small, the changes to the learning weights for any of the present cues is small as well. Therefore, the weight of the new cue will only ever increase to the point till the maximum activation is reached, at which point the error will become zero and no more improvement of weights is possible.

### 1.3.2 Cue-Blocking involving different Temporal Relations

As outlined in Section 1.1.3, assuming that two stimuli always occur completely simultaneously is a limiting view. If we assume the TCH is correct, we would expect different levels of blocking for different combinations of temporal relations, depending on the placement of the new cue in relation to the already established one.

This prediction is confirmed in experiments by Amundson and Miller (2008), which show that differences in the temporal relations between stimuli lessen the impact of cue blocking. Using rats conditioned to auditory stimuli, Amundson and Miller showed that an established association of a cue $A$ to the outcome US will only block a new cue $X$ from also becoming a predictor of the US, if the time passing between $A$ and the US is kept consistent throughout all learning phases (see Table 1.1 for cases tested). These findings provide evidence for the assumption of the TCH that temporal relations are cues that carry either useful or distracting information and that learners use this information to distinguish otherwise identical learning events and produce different learning outcomes (Amundson & Miller, 2008).

Unexpectedly, Amundson and Miller also found that if they place the new cue $X$ *before* the already established sequence of $A$ and the US, it will not be (fully) blocked

| Group | Phase 1 | Phase 2 | Test X Expected |
|---|---|---|---|
| 15–15 Block | A US | A X US | cr |
| 15–15 No Block | B US | A X US | CR |
| 3–15 Block | A US | A X US | Cr |
| 3–15 No Block | B US | A X US | CR |

Table 1.1: Examples of different cases tested by Amundson and Miller (2008), in regards to consistency of temporal relations. The letters cr, Cr and CR indicate the expected magnitude of the conditioned response. Note how, when temporal relations are inconsistent, the response is stronger, i.e. blocking is attenuated. Table adapted from Amundson and Miller (2008).

and will produce similar CR as previously only elicited through the presentation of $A$ or the US (see Table 1.2 for cases tested).

| Group | Phase 1 | Phase 2 | Test X Expected |
|---|---|---|---|
| 15–3 Block | A US | A X US | Cr |
| 15–3 No Block | B US | A X US | CR |
| 3–15 Block | A US | A X US | Cr * |
| 3–15 No Block | B US | A X US | CR |

Table 1.2: Examples of different cases tested by Amundson and Miller (2008), in regards to variation of placement of a new cue $X$ in relation to an established sequence. The letters cr, Cr and CR indicate the expected magnitude of the conditioned response. Note that the marked case "3-15 Block" did not result in blocking as expected.

At first glance this may seem intuitive, as it seems correct that if stimulus $A$ occurs *after* $X$, it should have no influence over whether a connection between $X$ and the US is formed. However, these findings are surprising under our current understanding of blocking as modelled by a discriminative model. From a network-modelling standpoint, we would expect that it makes no difference in which order two events occur, just that they share features and as such interfere with each other. However, this mismatch of model and data is precisely because of the discrete representation of learning events. The network-based approach allows only for a limited representation of sequence learning: Things are either cues or outcomes — before or after. Anything that is seen as a cue is grouped together by the model, making the two cases of placement of a new cue identical to each other. This means that in either case — regardless of where the new cue is placed — the new cue should not provide information that contradicts the association formed in the previous learning stage. But, as the data collected by Amundson and Miller show, this is not that case if the cue is placed before the established sequence, and some amount of sequential information is, in fact, available during recall.

### 1.3.3 Transference

A possible explanation for how sequential information could make it into the learner's mind is the aforementioned concept of *transference*. The basic idea is that given knowledge of $A$ predicting $B$ and $B$ predicting a third stimulus $C$, the temporal relationship between $A$ and $C$ can be (implicitly) inferred by the learner (Molet & Miller, 2014). Such transference of learned associations would solve the problem outlined in the previous section in that a learner who already knows that $A$ predicts the US and who then learns that the newly introduced cue $X$ predicts $A$ could then infer using transference that $X$ predicts the US as well. Noteworthy is that while we would expect transference to occur in all cases, it would only have a large effect if there is information to transfer. In the example of the new cue $X$ being introduced *after* $A$, we don't expect much if any information to be transferred, because $X$ is being blocked from becoming associated with the US.

Of course, there may be alternate explanations. For instance, it may be the case that the new cue $X$ is not seen as linked to the co-occurrence of $A$ and the US. We recall that blocking essentially functions on the principle of new information being redundant. In the situation where $X$ is placed after $A$, the new cue can only provide redundant information, as the learner already knows that the US is about to occur. Alternatively, when $X$ occurs first, it could be argued that at the moment when $X$ is encountered, it is not yet certain if $A$ will occur, so $X$ holds some predictive power, allowing it to avoid being blocked. This is the explanation discussed in Amundson and Miller (2008).

For clarity, it should be highlighted perhaps in what way this information theory viewpoint is different from the concept of transference: The former assumes that in the described scenario $X$ is not blocked at all, while the latter assumes that while blocking occurs, the gaps are filled in afterwards, with the end result appearing the same. The crucial difference is that transference does not create "real" learning of sequences, but only offers an alternative path to infer missing information. In a review of transference, Molet and Miller (2014) outline that research on transference as a cognitive process (rather than just a theoretical concept) shows that it only occurs during the explicit recall of information. This is measurable, for example, in that transference captures only the general order of cues rather than the exact temporal relations. So, in the cue-blocking example, it would be possible only to infer that the US will occur *at some point* after $X$ but not exactly when, which may be a useful marker for examining if transference is used or not.

To sum up, this thesis is interested in the question of how temporal relations interact with blocking effects with the larger goal of investigating how to include temporal relations inside of a discriminative model. For this we examine the differences in responses related to the placement of a new cue in an established sequence, with the initial theory being that learners can use transference to restore blocked information. For this purpose the next section will outline a simulation study, to further identify questions regarding formalisation of this process.

# 2 Simulation Study

The aim of this simulation study is to model cues holding information about temporal relations as part of a network model using a discriminative learning rule, to compare different ways of encoding, and to simulate the effect of transference. We will first explore a direct implementation of temporal relations as cues and then highlight the benefits of adding simulated transference on top.

## 2.1 Base-EDL

The *edl* package developed by van Rij and Hoppe (2021), which runs in *R* (R Core Team, 2021), provides a good base to calculate changes to a weight matrix based on specified learning events. Within this implementation, strings define discrete cues and outcomes of each learning event. An example of a string representing the cue-set of two cues $A$ and $X$ would be `A_X`. This structure can be expanded to include cues representing the temporal relation between each cue and the outcome of each learning event (not to be confused with the temporal relation between each of the cues in the cue-set). For instance. it may be encoded that cue $A$ occurs 10 seconds before an outcome or that cue $X$ occurs 3 seconds before a different outcome.

Ahead of discussing the exact nature of encoding, it should be mentioned our model will work on the assumptions that temporal relations are encoded in seconds and that each time interval represents a unique cue with no relevance to how similar or dissimilar two different time intervals are to each other. In practice, the similarity between 10 and 11 seconds is of course larger than between 10 and 100 seconds. In addition, human time perception is logarithmic and people treat small and large intervals differently, to the point of showing different levels of accuracy when having to make estimates (Taatgen & van Rijn, 2011). However, it is not straightforward to include such a continuous notion of time into the EDL framework, which assumes discrete cues. Therefore, we will ignore this knowledge for now and assume that the time intervals used are possible to distinguish and are estimated to the nearest second. Let us now turn to how to encode this information in our model, starting with the most straight forward cases discussed in Amundson and Miller (2008).

### 2.1.1 Simultaneous cues

The simplest way of encoding temporal information into a cue-set is to define each cue as representing both the sensory aspect of the stimulus as well as the temporal relation. For example, the cue-set of $A$ and $X$ may be written as `A-3s_X-3s`, representing a learning event where both $A$ and $X$ occur three seconds before the same outcome (this is the first case tested in Amundson and Miller (2008), which was found to result in blocking). Figure 2.1 shows that for this compound-cue setup, there is no increase in learning weight for the new cue $X$ once it is introduced and that the only activation visible is carried by the general background cue (BG), which represents information that is always present. From this we can conclude that the cue $X$ is blocked, as it was in the real experiment.

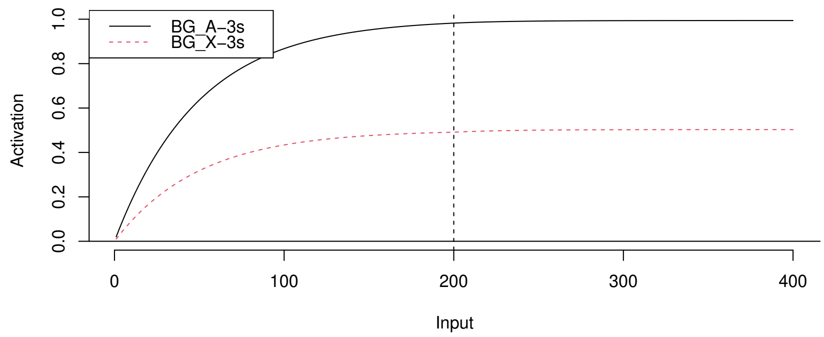


Figure 2.1: Learning trajectory for a 3-3 blocking paradigm, as given by a model using compound-cues. Initially only cue $A$ is presented, with additional cue $X$ introduced after the dotted line. Both cues occur three seconds before the US. No change to learning weights of $X$ occurs after it is introduced, though the background cue (BG) provides some activation.

An alternative implementation to the compound-cue method would be to treat the sensory content and temporal relation as two separate cues. For the same learning event as described above, this would give us the cue-set `A_3s_B_3s`. Figure 2.2 shows the learning trajectory of the same scenario, using this different implementation. Notably, while blocking of the new cue occurs, there is some unlearning for both learned cues. This is because in this implementation, the cue `3s` looses predictive power. Amundson and Miller do not report anything that would suggest this occurring, though they don't explicitly check for it either. In general, it can be said that this implementation is unlikely to be accurate, as previous research already indicates that temporal relations are not stored without the event they are anchored with (Molet & Miller, 2014).

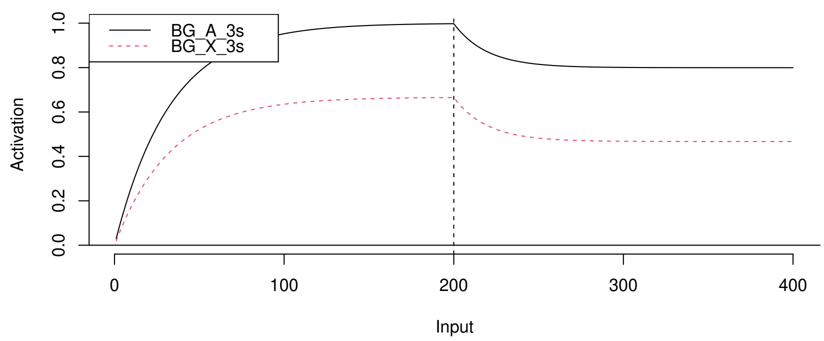


Figure 2.2: Learning trajectory for a 3-3 blocking paradigm, as given by a model using individual cues. Both cues occur three seconds before the US. While blocking occurs, both $A$ and $X$ are less strongly predicted once the second cue $X$ is introduced.

An advantage of the implementation using two cues, would have been that it would be possible for learners to ignore the information about temporal relations when it is not useful, while still gaining knowledge about the sensory aspect. To maintain this feature, we try a third approach, a combination of the previous two (cf.

Ramscar et al., 2011): A cue of just sensory aspect combined with the compound cue. In practice for our example this would be a string of four cues: `A_A-3s_B_B-3s`. Figure 2.3 shows that in this implementation there is no unlearning taking place, and the blocking of $X$ is equally as strong as before.

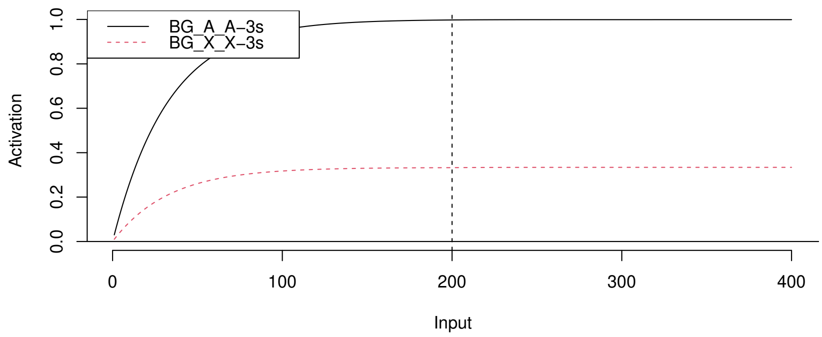


Figure 2.3: Learning trajectory for a 3-3 blocking paradigm, as given by a model using compound-cues and additional sensory cues. No change to learning weights of $X$ occurs after it is introduced.

We recall that one of the main findings of Amundson and Miller is that blocking is attenuated if temporal relations are inconsistent through out the learning process. For example, if the interval length changes from 15 seconds in the compound phase to 3 seconds in the elemental phase, the amount of blocking is equal to control conditions with no blocking at all. Figure 2.4 shows the corresponding simulations of this experiment for the two ways of cue-based encoding. In the pure compound cue implementation, there is no blocking whatsoever, while the implementation using separate sensory cues has some blocking occurring. This makes sense, as with the split cues, the sensory cue $A$ is a good predictor even if the temporal interval changes. This is exactly the dynamic described in the previous paragraph, of unreliable temporal relations being simply ignored. Which of the two implementations is more accurate for this case is hard to say, since this is something that would lead to changes for cue $A$, which Amundson and Miller, did not measure or report.

### Why not Outcome-based encoding?

At this point it is good to briefly explain why the temporal relation is not encoded as part of the output. To clarify what is meant by this, in the so far outlined approach the learned information is that a stimulus and a pause of a certain length predict an outcome rather than a stimulus being seen as a predictor of an outcome after a pause has passed.

While the alternative implementation is possible, if the temporal relation is stored with the outcome (e.g. `US-15s`), there is nothing with in the cue-set itself that distinguishes the different learning phases from another. Instead, the semantic cue comes to be linked and unlinked with different information. Figure 2.5 shows such a simulation for the same case as before. The main difference is that information from the elemental phase is now more drastically unlearned. This means that if we now present again the cue $A$ with the old interval of 15 seconds, people should show no signs of having previously learned this information. For this thesis, the possible

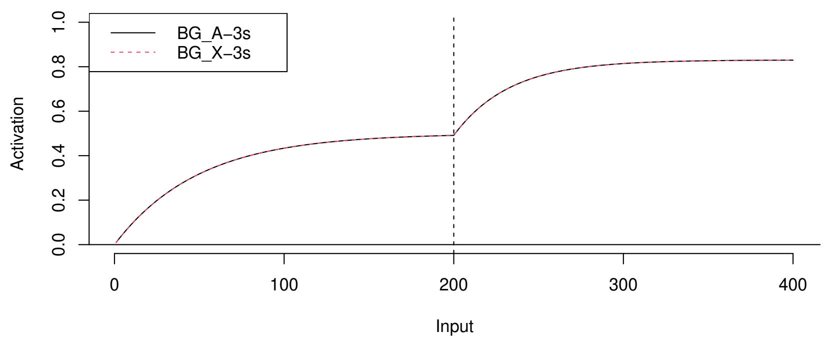


(a) Compound cues

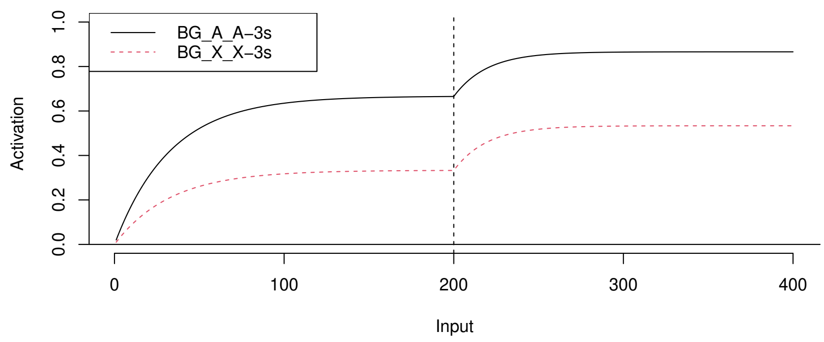


(b) Sensory cues and Compound cues

Figure 2.4: Learning trajectories for an unreliable blocking paradigm, as given by a model using compound-cues only and one with additional sensory cues. The former shows no difference between the two cues, while the latter shows that the new cue $X$ is still partially blocked.

validity of this theory is acknowledged, but a targeted exploration is left to future research.

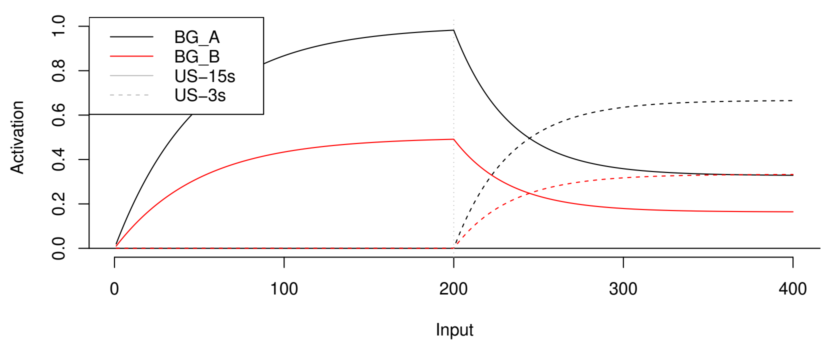


Figure 2.5: Learning trajectory for an unreliable blocking paradigm, as given by a model where temporal relation is encoded in the outcome. Full lines show activation for US after 15 seconds, dashed lines show activation for US after 3 seconds. Information of first learning phase is partially unlearned.

### 2.1.2 Non simultaneous cues

Let us now focus on the experiments of Amundson and Miller involving cues with different temporal relations to the outcome US.

One experiment described (experiment 1 in Amundson & Miller, 2008) involves an already established cue $A$ occurring 15 seconds before the US, and the new cue $X$ occurring just three seconds before the US. Amundson and Miller found that in such cases blocking does occur, as was in line with their predictions. Going back to our compound-cue-based implementation, we can implement this learning event as the cue-set `A-15s_X-3s` and the outcome `US`. However, we must now also include a second learning event which happens right before this one, which is that $A$ is followed by $X$ after 12 seconds. For this secondary learning event — secondary in the sense that it is the less interesting one, though it occurs first, *before* the main one — we have the compound-cue `A-12s` and outcome `X`. Figure 2.6 shows the model results for experiment for either of the usual implementations. There is some slight unlearning in both simulations, but the results are still mainly in line with the animal data, in that $X$ gets blocked.[1]

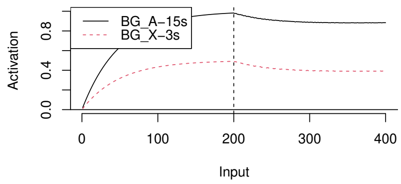


(a) Compound cues

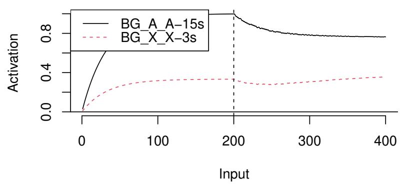


(b) Sensory cues and Compound cues

Figure 2.6: Learning trajectories for a **15-3 blocking paradigm**, as given by a model using compound-cues only and one with additional sensory cues.

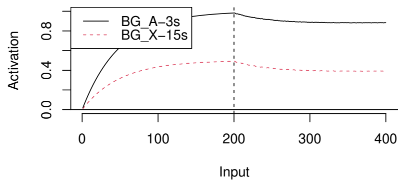


(a) Compound cues

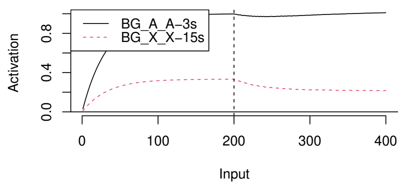


(b) Sensory cues and Compound cues

Figure 2.7: Learning trajectories for a **3-15 blocking paradigm**, as given by a model using compound-cues only and one with additional sensory cues.

The more interesting case from Amundson and Miller (2008) is of course when the new cue $X$ is placed *before* $A$. In this case Amundson and Miller unexpectedly found that no blocking occurs. However, as Figure 2.7 shows, the results are more or less the same as in the previous set of simulations, where $X$ occurred *after* $A$. This is because — from a modelling standpoint — there is little distinguishing these cases. This again illustrates the current limit of the discreet network model. Therefore, we now turn to including transference in our model.

[1] The implementation completely splitting sensory cue and temporal relation is not shown here as it was already shown to not be useful. It produces similar results to those shown in Figure 2.6b.

## 2.2 Adding Transference

### 2.2.1 Method

As outlined in Section 1.3.3, transference is a process which occurs during the recall of information rather than the actual learning process. However, the EDL model, can only model the effect of learning events on a weight matrix. It is possible, though, to add a simulation of transference after the fact.

For this, we adapt the equation for calculating the activation of an outcome: Let us assume we have a learning event at time $t$. We recall, that, if we want to know the activation of outcome $j$ given a specific cue $i$ at moment $t$, we must retrieve the weight $v_{ij}^{t}$. As there may be more than just one cue in our cue-set, we take the sum of all weights leading from cue $a$ to outcome $j$, where $a$ is any cue present during moment $t$.

$$act_j^t = net_j^t = \sum_{a \in cues(t)} v_{aj}^t \tag{2.1}$$

This is what the base-EDL system discussed in the previous section already does. Now let us consider the situation that there are three cues occurring in a row, $A$ followed by $B$, which itself is followed by $C$. We are interested in finding a formula that given just the cue $A$ gives us the activation of $C$ by using another cue (in this case $B$) as a 'pit stop'.

To do this we will combine the weights $v_{AB}^t$ and $v_{BC}^t$ in a way that we get a positive activation. However, for our method to be realistic, it should only produce a positive activation if the learner indeed posses the two pieces of knowledge to create this alternate route, namely that $A$ predicts $B$ and $B$ predicts $C$. Therefore, a good way to combine the two weights is to multiply them together. This way, we only get a positive activation at the end, if both the weight leading to and from the 'pit stop' are positive. Further, if the learner has only partially learned one of the linking connections, they will also only be able to produce a small amount of transferred information.

For our simulation, we will also use a *Rectified Linear Unit* (ReLu) function on the 'pit stop' weights, so that we will not get a negative activation. This is done to prevent transference being used to infer the absence of a cue, which none of the literature suggests is the case.

Returning to formalisation, we add to our activation function the product of all possible combinations of two (above-zero) weights. First, we take the weight $V_{ab}^t$, which leads from cue $a$ (which may be any cue present at moment $t$) to outcome $b$ (which can be any outcome part of $V_i(t)$, meaning any outcome already in the weight matrix at moment $t$). We multiply this with the weight between any matched cue $b$ and our goal outcome $j$, simulating a chain from $a$ to $b$ to $j$:

$$act_j^t = net_j^t = \sum_{a \in cues(t)} v_{aj}^t + \sum_{b \in V_i(t)} ReLu(v_{ab}^t) * ReLu(v_{bj}^t) \tag{2.2}$$

To show a practical example, let us assume we have a weight matrix at moment $t = 6$, representing a situation where the learner has partially formed associations between $A$ and $B$ as well as $B$ and $C$, but nothing else. We let $V_{AB}^6 = 0.8$ and

$V^6_{BC} = 0.9$ and all other weights be zero.[2]. To calculate the activation of outcome $C$ given just the cue $A$ we get:

$$\begin{aligned} act^6_C = net^6_C &= \sum_{a \in \{A\}} v^6_{aC} + \sum_{b \in \{A,B\}} ReLu(v^6_{ab}) * ReLu(v^6_{bC}) \\ &= v^6_{AC} + \sum_{b \in \{A,B\}} ReLu(v^6_{Ab}) * ReLu(v^6_{bC}) \\ &= 0 + ReLu(v^6_{AA}) * ReLu(v^6_{AC}) + ReLu(v^6_{AB}) * ReLu(v^6_{BC}) \\ &= 0 + (0 * 0) + (0.8 * 0.9) \\ &= 0.72 \end{aligned} \tag{2.3}$$

Importantly, $t$ here refers always to the time of recall of information, rather than when information was learned; Even if, for instance, the learner encounters first an occurrence of $A$ predicting $B$ and then later is exposed to $B$ predicting $C$, the only thing that matters regarding transference, is what the weights are at the time of recall. For this reason, we use the same value for $t$ throughout the equation. It is noted also that, for the process of changing the weights during learning, the unaltered equation should be used. Further, before moving on to the next section, it must be stressed that this study does not claim that this mathematical approach to what ultimately appears to be a reasoning process is in any way based on cognitive research. It is simply an approximation of the result of transference, rather than the actual process itself. Still, this simple example shows that with this new method, even though a direct connection between $A$ and $C$ was not present in the weight matrix, it is possible to simulate a positive activation reached through transference.

### 2.2.2 Results for replicating Amundson and Miller (2008)

We can use our new equation to examine how adding transference affects the outcome of the simulations discussed in Section 2.1. We can first say, that in cases where cues occur simultaneous, there is absolutely no transference at all. This is because there is simply nothing to transfer, as the weight matrix only ever has one possible outcome — the unconditioned stimulus. No associations between $A$ and $X$ are formed, and as such there is nothing to transfer.

More interesting then are the cases where cues are non simultaneous. Returning for example to the scenario where the new cue $X$ is placed inside the established sequence $A$ to the US, we can see in Figure 2.8 that in both of our implementation methods, transference eventually counters the initial unlearning that occurred after the introduction of $X$. Overall though, transference does not change the situation that $X$ is blocked from becoming associated, staying in line with the experiment data collected by Amundson and Miller.

As we turn now to the case of $X$ being placed *outside* the established sequence, we recall that we were not able to correctly simulate this experiment when using just the base EDL-system. Figure 2.9 shows now that in both of the simulations, transference pulls up the learning trajectory for the $X$ cue significantly. Interesting to note is that even though the weight for $X$ is initially much lower in the model

[2]The simulations replicating Amundson and Miller (2008) will use a background cue, but for now, to keep the mathematics simpler, it is omitted.

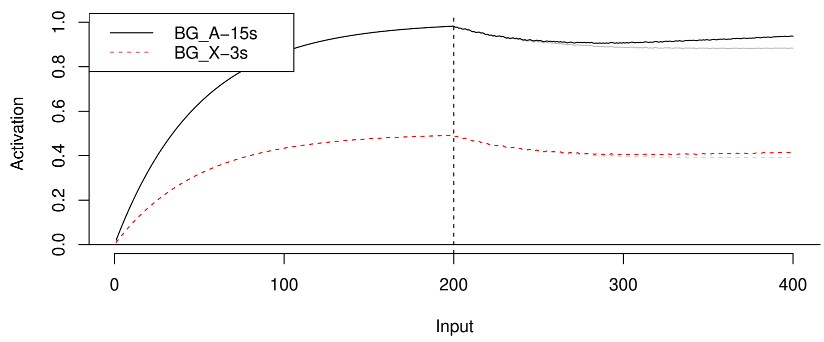


(a) Compound cues

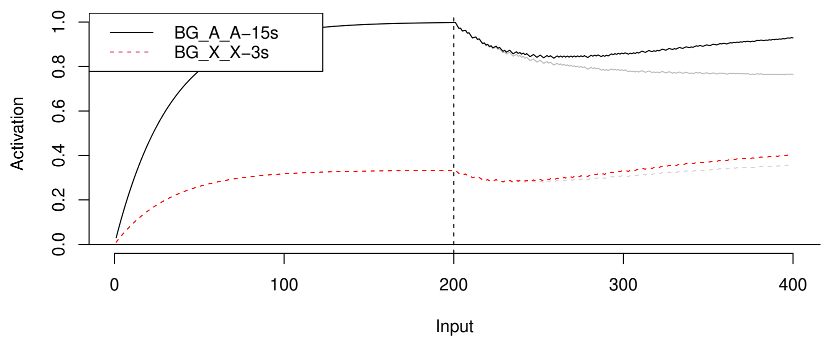


(b) Sensory cues and Compound cues

Figure 2.8: Learning trajectories for 15-3 blocking paradigm, as given by a model using compound-cues only and one with additional sensory cues. Coloured lines show the new trajectories with added transference, with old trajectories in grey for comparison.

using a separate sensory cue, by the end of the learning process, it is on the same level as in the implementation using just a single compound cue.

## 2.3 Discussion

The aim of this simulation study was to test different approaches to encoding temporal relations and the results of combining them with transference to identify the most accurate approach and identify remaining ambiguity. From the obtained results, it seems that encoding temporal relations and sensory aspects as a compound cue is effective, and that results are even better if an additional cue of just the sensory aspects is added.

We could also see that the chosen method of simulating transference can have a significant affect on getting the network model in line with the results found by Amundson and Miller (2008). Importantly, the method only causes big shifts in learning trajectories when it would reasonably make sense for transference to be used by a learner. As shown, it does not magically fix situations where the reason that blocking was attenuated is inconsistent temporal relations, but only changes the outcome in the specific case of a cue being placed outside of an established sequence.

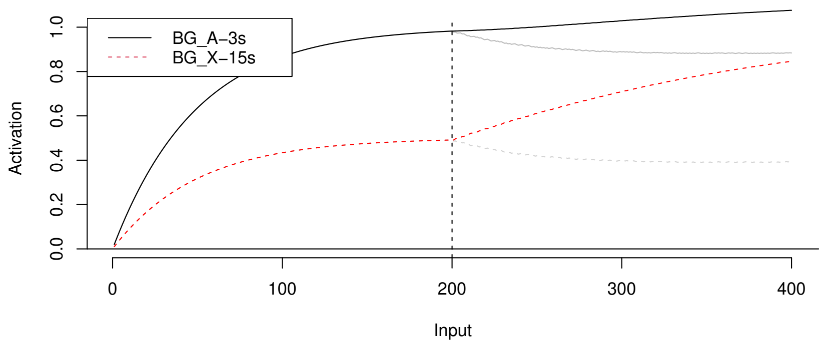


(a) Compound cues

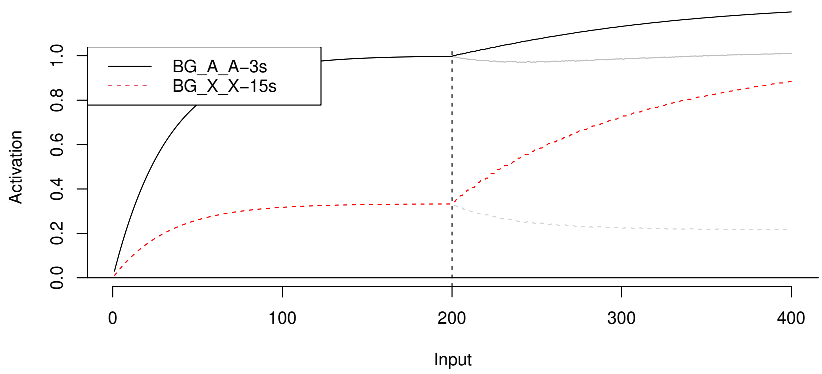


(b) Sensory cues and Compound cues

Figure 2.9: Learning trajectories for 3-15 blocking paradigm, as given by a model using compound-cues only and one with additional sensory cues. Coloured lines show the new trajectories with added transference, with old trajectories in grey for comparison.

However, we can identify some issues. First of all, it remains to be tested whether the findings of Amundson and Miller are consistent when repeated using human participants. Secondly, we should investigate if the original sequence is unaffected in cases where blocking occurs. As we have seen, for example, in the issue of encoding the temporal relation as it's own cue, once two learning events share the same temporal relation, there would be some ambiguity introduced which may result in the initial connection becoming weaker. Lastly, when comparing cases that we suspect to rely on transference with those that do not, it should be examined if there are differences in the accuracy of predictions. This would give us indication if transference is really the right idea, or if we might rather use a different explanation.

To summarise, the remaining questions are:

- Can we replicate the findings on the attenuation of blocking in relation to cue placement with human participants?
- Is the knowledge, which is learned in the elemental phase, unlearned, when a new cue is introduced?
- Is the accuracy of temporal estimations higher, if the new cue is placed inside the sequence rather than outside?

Answering these questions should help us determine the accuracy of our cognitive model and provide more insight into the validity of the model's assumption. The following sections will outline an experiment to collect the required human data to answer them.

# 3 Human Study - Methods

## 3.1 Experiment Design

Human participants are tested to gather reaction time data across different conditions. The recorded data should capture accuracy, speed and precision, i.e. if people can respond quickly and during a specific time interval. For this purpose, participants look at stimuli on a computer screen and give a response as soon as a specified stimulus appears.

The experiment is in the form of a video game to discourage participants from analysing the underlying rules and instead focus on reacting as quickly as possible. The game's story is that aliens are attacking a city using a laser. The participant has to block these attacks with a shield which they activate by pressing the space bar. The participant can see a sequence of symbols displayed one at a time on the screen, which they are told represents the aliens' communication with each other. Before the start of the game, participants are shown one of these symbols and are told that when the aliens use this symbol, an attack is imminent. The participants are told nothing about any of the other symbols. The game's goal is to deflect all attacks by reacting quickly to the trigger symbol. Participants are informed that each blocked attack earns them 50 points, failure to block will result in a loss of 50 points, and unnecessary uses of the shield will result in a loss of 10 points. The game explicitly instructs participants to get as high a score as possible.

The game consists of continuous trials involving multiple presentations of the trigger symbol, each reliably followed by an attack serving as feedback to the participant. A failure to block an attack is visualised as the alien laser hitting the city and the screen turning red for 300 ms. A successfully blocked attack is visualised as the alien laser reaching the shield and bouncing off from it as sparks, shown for the same time interval. This visual feedback functions to keep participants engaged in the game and to encourage them to avoid the negative affect of the sudden red screen and loss of points. However, the game is designed in a cartoon style to keep the caused distress from exceeding the required amount. Figure 3.1 and 3.2 show examples of the visual content of the game, which were also shown to and approved by the research ethics committee *Commissie Ethische Toetsing Onderzoek* (CETO).

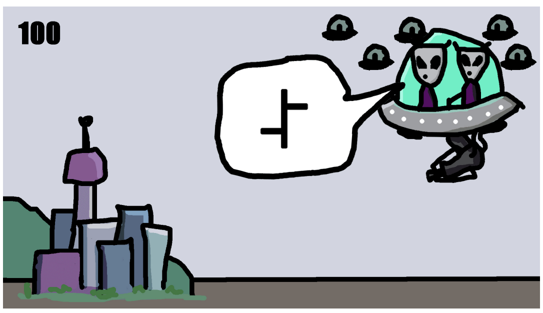


Figure 3.1: Example screenshot of the invading aliens game, showcasing the display of a stimulus.

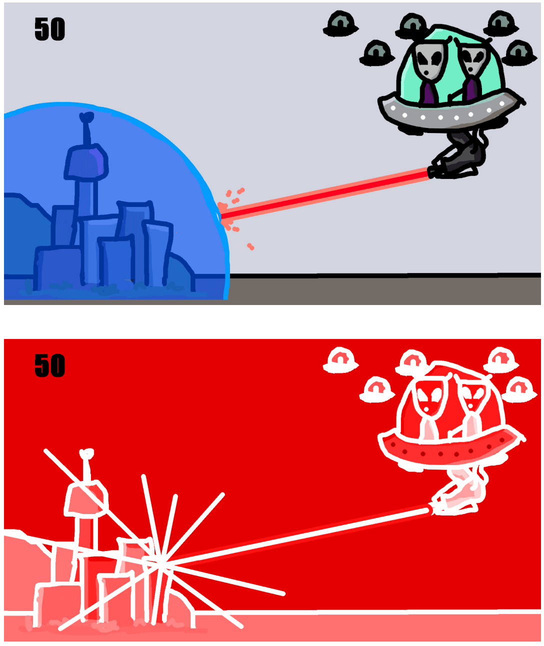


Figure 3.2: Example screenshots of the invading aliens game showcasing the two feedback screens for a successfully blocked attack (top) and failure to block (bottom).

Within the game, the visual feedback of the attacks acts as the unconditioned stimulus (US), with the resulting negative affect of failing to block an attack being the unconditioned response (UR). Once the participant starts performing well at the game, they will experience a positive affect in the form of relief as the negative affect subsides. Through backward conditioning, participants will connect this positive affect to the conditioned response of pressing the response key at the right time. The reward system of points within the game serves to further this connection.

Unbeknown to the participants, the stimulus immediately preceding the trigger stimulus is consistent across trials. These stimuli serve as a hidden target sequence, with the idea being that participants may form a secondary implicit association between the presentation of the hidden stimuli and the positive affect of avoiding punishment. To keep consistency with the descriptions used in Amundson and Miller (2008), we will refer to the hidden stimulus as the $A$ stimulus. An association is considered formed if participants react faster to attacks when presented with the $A$ stimulus than in other cases.

### 3.1.1 Phases and Conditions

As described in Section 1.3.1, for tests of blocking, the learner usually goes through an elemental learning phase, a compound learning phase and a test phase. Similarly, in this experiment, participants pass through three phases of continuous trials, with breaks in between each one. Each phase of the game is presented as a day in the fight against the aliens.

In the first phase, the target sequence consists of the $A$ stimulus followed by the trigger stimulus, with the rest of the trials as a random sequence of filler stimuli so that only the target sequence provides consistent information to establish it as reliable in the participant's mind. The filler cues also serve to obscure the presence of the hidden sequence. One constraint on the trials is that the target sequence is never the first or last thing participants see to ensure that measurements are not affected by participants getting used to the game and to prevent special attention being drawn to the sequence by ending a trial with it.

The second phase introduces a new cue $X$ to the sequence. This phase could show participants unlearning the previously learned sequence in favour of a new sequence or ignoring the new cue completely. The placement of the new cue is either just before $A$ or between $A$ and the trigger symbol, depending on the condition. These conditions are comparable to the non-simultaneous experiments discussed in Section 2.1.2 and will be similarly referred to as "Inside" and "Outside" conditions respectively.

The experiment also includes two control conditions to allow making inferences about whether blocking took place, which will be referred to as "Inside (control)" and "Outside (control)". These conditions differ from the previously described phases primarily in the elemental phase; Participants are not shown a sequence of $A$ preceding the US but an alternate sequence with an unrelated masking cue $B$ instead of $A$. This new cue $B$ is only used in this condition and not part of the set of filler cues. In keeping with Amundson and Miller (2008), the control condition serves as a comparison showing that any differences in responses in the non-control conditions must be due to the fact that participants' previous knowledge affected the forming of new associations i.e. due to the presence of blocking effects.

Lastly, a test phase closes out a run of the experiment. In this phase, the target sequence consists of only one symbol, *A or X*, presented before the trigger symbol. This phase should test which stimulus, if any, a participant relies on to anticipate an attack once they have undergone both training phases. Individual participants are only tested on one stimulus, as only the association based on the training phases is wanted, and the testing phase itself may provide new information. For this purpose, half of the participants are assigned to be tested on $A$, with the other half tested on $X$. However, it is stressed that the procedure in phases before the test phase is the same, with differences in response only being expected within the test phase itself.

Combining the four training conditions with the two possible test conditions gives eight possible condition combinations, as shown in Table 3.1. Important also is that across one condition, the time interval between two stimuli never changes, e.g. the time passing between the onset of $A$ and the onset of the trigger symbol US is the same during the elemental, compound and testing phase, regardless of the presence of $X$ in between or before. This fixed length ensures that the temporal relation can be used as a predictor and that if blocking effects are weaker across two conditions, it is not caused by inconsistent temporal relations introducing noise.

### 3.1.2 Timing and Intervals

Each trial of the experiment begins with the onset of the respective stimulus. A stimulus always appears for 400 ms. After this, the stimulus disappears until the trial ends and the next stimulus appears. There is a distinction between short and long *inter-stimulus intervals* (ISI) for the target trials. A short ISI is always 300 ms, while a long ISI is always 1000 ms. The length of these intervals was selected so that two cues with a short ISI take up the same time as one presentation of a cue with a long ISI. The length of the short ISI is similar to the reaction window used in *Space Fortress*, a game used to study reaction times involving variation in time intervals between cue (Mané & Donchin, 1989; Moon & Anderson, 2012).

Regarding the trials which are filler trials, the ISI is randomly sampled from a uniform distribution with the range of 200 ms up to 1100 ms (an exception to this are filler stimuli which preceded the target sequence, which are sampled from a

| Conditions Training | Test | elemental | Phase compound | test | number of participants |
|---|---|---|---|---|---|
| Inside | A | A - - - US | A - X - US | A - - - US | 6 |
| | X | | | X - US | 6 |
| Outside | A | A - US | X - A - US | A - US | 6 |
| | X | | | X - - - US | 6 |
| Inside (control) | A | B - - - US | A - X - US | A - - - US | 6 |
| | X | | | X - US | 6 |
| Outside (control) | A | B - US | X - A - US | A - US | 6 |
| | X | | | X - - - US | 6 |

Table 3.1: Table of the eighth possible condition combinations, showing target sequences across experiment phases. The length of the dashed lines indicates the relative length of time between stimuli onsets (long or short). Two short intervals equal one long interval.

range with a minimum ISI of 300 ms). The shortest time between the onset of two concurrent stimuli is 600 ms, while the longest is 1500 ms.

Trials in which an attack occurs (i.e. those involving the trigger symbol) have a total length of 1000 ms, with the attack appearing 700 ms after the initial stimulus onset. The visual feedback of the attack always lasts 300 ms and is counted as part of that trial. The resulting distribution of trial lengths across the whole experiment run is shown in Figure 3.3.

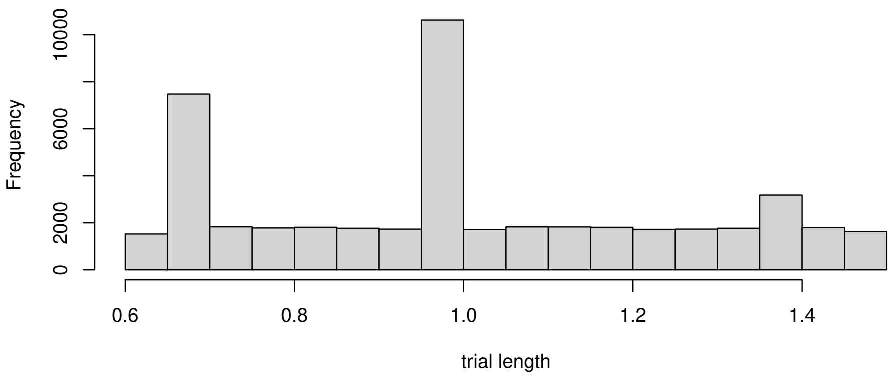


Figure 3.3: Histogram showing the distribution of trial lengths in trial sequences across all participants

After a key press, the shield stays switched on for 1000 ms, making it possible for participants to press up to 300 ms before the onset of the trigger stimulus and still be rewarded. This allows participants to use any knowledge they have about the temporal relations of the target sequence to time their key presses. However, participants received the same reward regardless of how accurately they could estimate how soon to press, as long as the shield was up during the start of the attack.

### 3.1.3 Alien Language

For this experiment, the stimuli are only visual, instead of auditory as they were in Amundson and Miller (2008). This is because not all auditory stimuli produce equally strong blocking effects (Nixon, 2020) and that working with auditory stimuli is less practical regarding creation and presentation during the experiment.

The experiment uses three different stimuli sets to ensure that the difficulty of stimuli does not influence results. The symbols are taken from the *Brussels Artificial Character Sets* (BACS) developed by Vidal et al. (2016). These symbols were specifically developed for use in studies and are based on existing alphabets such as Chinese, Asomtavruli, Cyrillic, Oriya, and Latin. The ten most unique characters were selected from each of the three available character sets using a similarity matrix provided by Vidal et al. The symbols were randomly grouped into six filler stimuli and four target stimuli per set. The stimuli are shown in Figure 3.4.

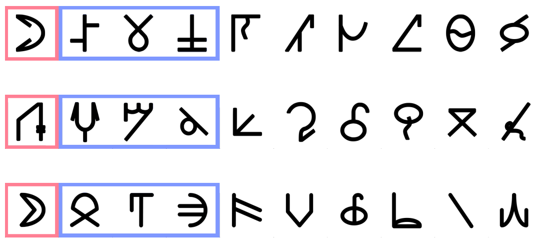

Figure 3.4: The three stimuli sets used in the experiment. Marked in red is the respective trigger symbol. Marked in blue are symbols used as part of the target sequence of hidden symbols. The remaining symbols are used as fillers.

### 3.1.4 Structure of Sequence

To ensure that the occurrences of attacks were evenly spread out throughout the attacks, the sequence of symbols in each phase of the experiment was semi-randomised: A phase was divided into a certain number of blocks (four for the learning phases, one for the test phase) each containing six bins, with one bin containing one of each filler stimuli and the target sequence. The placement of the target sequence inside a bin and the order of the surrounding filler stimuli is random. Overall, participants encountered the trigger symbol 24 times in one experiment run, plus three additional times in a preceding practice round. The practice round includes nine presentations of filler stimuli and three presentations of the trigger stimulus, with none of the 'hidden' stimuli present.

### 3.1.5 Second Level

A second level of the game is included as well, which participants play once they have finished a complete run of the experiment. This level is another experiment run, but in this version the aliens use a shield and the trigger symbol signals when they are about to turn it off temporarily. At those times, participants can press the space bar to shoot a laser at the aliens and receive points. The alien gun was removed from the spaceship to make clear to participants that they did not have to worry about being attacked (see Figure 3.5).

The changes in the second level are primarily visual, with the underlying task still consisting of participants having to react with a key press after the trigger symbol appears. The point reward system is unchanged in that pressing at the right time results in gaining 50 points, failing to press results in losing 50 points, and unnecessary presses result in losing 10 points. However, the second level is slightly more difficult as the window when participants should press to score points is much smaller, with the laser only staying activated for 200 ms, compared to the 1000 ms of

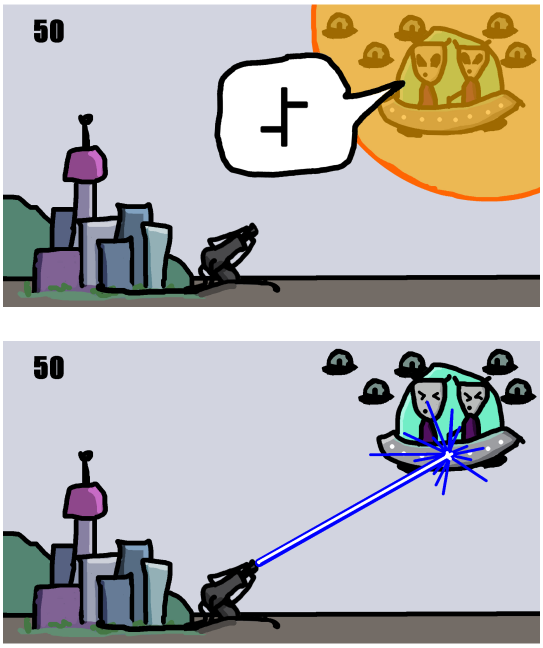


Figure 3.5: Example screenshots of the second level of the invading aliens game showcasing the display of a stimulus (top) and a successful attack from the player on the aliens (bottom).

the shield. As before, participants are rewarded as long as the laser was active when the aliens lowered their shield, meaning that the earliest possible time a participant could press was 500 ms after the onset of the trigger symbol. Stimuli sets and conditions are counterbalanced across participants and levels to prevent order effects from affecting the results.

## 3.2 Procedure

The experiment is coded in *Python* and run using the experiment builder software *PsychoPy* (Peirce et al., 2019). The experiment is conducted under controlled conditions at the University of Groningen. The present instructor notes any visual or auditory disturbances occurring during the experiment to allow referencing in cases of outlier data during analysis. The experiment is conducted over multiple days between 11 a.m. and 5 p.m., with participants able to choose their timeslot. Participants knowingly provide some personal information for storage in a separate payment system, but no personal data is recorded with the experiment data.

Before the start of the experiment, participants are given a printed information letter and asked to sign an informed consent form. The information letter described the study's aim as investigating "how humans can learn over time to give faster responses to a specific signal", which is an intentionally ambiguous description, to avoid participants knowing the full extent of the experiment. Participants are informed that they may withdraw their consent to participate in the study at any point during the experiment and up to 24 hours afterwards. The information letter, consent form, and the necessary level of deception were evaluated as acceptable with no objections by the ethics committee CETO.

Once participants have signed the consent form, the experiment starts with instructions about the game appearing on the screen. Participants are shown the trigger symbol before and after the practice round, after which they enter the first "day" of the game. After completing this first section (i.e., the first learning phase), the participants can take a short break of their chosen length. During this break, they are asked to respond to a thought probe with the following options:

1. *I was focused on responding in time*
2. *I was thinking about my score*
3. *I was thinking about aspects of the game (e.g. the alien language or the level of difficulty)*
4. *I was distracted by my environment*
5. *I was daydreaming about task-unrelated things*
6. *I was not paying attention, but I did not think about anything specific*

If they feel multiple answers apply, participants are instructed to answer with what most applies to the part of the game they just played right before the break. There are a total of two thought probes per level of the game.

Once participants complete the testing phase of the first level, the game instructs them to take a longer break, with the exact length up to the participant. Participants are instructed to stand up or look away from the computer. Once participants indicate they are ready to continue, the screen shows instructions explaining the changes in the second level. Participants are again shown the new target symbol and given a practice round.

After participants finish the second level, they are asked to complete a memory test. In this test, participants see two stimuli on screen, one of which they have previously seen and one which is novel. Participants must select which symbol they have seen before, though there is no time limit and no points to be earned.

Once participants have completed the memory test, they are debriefed using a pre-written text. The present instructor then conducts a short semi-structured interview asking if participants noticed the sequence, if they applied any strategies, and if they found the game engaging. Responses are recorded paraphrased, as are any further comments that the instructor may deem relevant. Participants are also given the option to ask any questions they might have. Participants are paid between eight and ten euros, a few days after the experiment, based on their final score.

## 3.3 Planned Analysis

Any participant data may be excluded if the participant withdraws their consent or if their data is deemed unusable. Data is considered suspicious if the experiment log, the thought probes, or the memory test indicate that the participant did not complete the experiment correctly. If multiple factors point to data being suspicious, it is considered unusable and excluded. Further, a whole stimulus set may be excluded if data from the main game or the memory test indicate it as significantly more difficult to parse than the others.

The data analysis will be conducted in R version 4.1.2 (R Core Team, 2021). A General Additive Mixed Model (Hastie & Tibshirani, 1990; Wood, 2017) will be fit to the data to determine the presence of any fixed effects and data patterns.

### 3.3.1 Expected Results

Random effects for stimuli set and participant differences are expected. Fixed effects for condition and encounter of the trigger stimulus are expected, with a possible interaction. Based on the previous research, no differences in reaction time (RT) across conditions are expected in the elemental phase. However, data may show some speed-up taking place, due to learning, regardless of condition. In the compound phase, participants in the control condition should slow down at the start of the phase as they have to adjust to the new sequence. Participants in other conditions may also show a slight readjustment.

The main findings are expected in the test phase. It is predicted that participants in the “Inside" condition will react fast if tested on $A$ and slow if tested on $X$. No significant difference is expected for all other conditions, though there may be some slight differences. Figure 3.6 approximates the expected reaction speeds in the test phase. It is noted that these predictions are formed based on the original findings of Amundson and Miller (2008), not the simulation study described in Section 2.

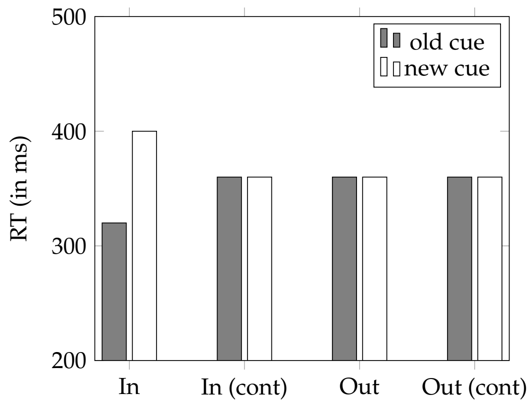


Figure 3.6: Predicted reaction time values for test phase of the experiment.

Lastly, people are expected to adjust their behaviour slightly when moving to the second level, resulting in a slightly more delayed reaction overall to match the smaller target window.

# 4 Human Study - Results

In total, 48 participants were tested (26 males / 22 females, aged 18-26). One participant's data had to be excluded due to technical difficulties. The highest score possible to achieve in the game was 5800, which one participant achieved. The lowest score achieved by a participant was 5410 points. The median score across all participants was 5765.

Data was taken from the log files of each experiment run as they were less affected by noise than the other timer values stored by PsychoPy. Some noise is still retained in these timers, as visible, for instance, in Figure 4.1, which shows that most target trials deviated from the supposed trial length of 1000 ms and that this deviation was not consistent from trial to trial. Noteworthy also, is that the magnitude of these effects differs between the two levels. The relevance of this noise is further discussed in Section 5.

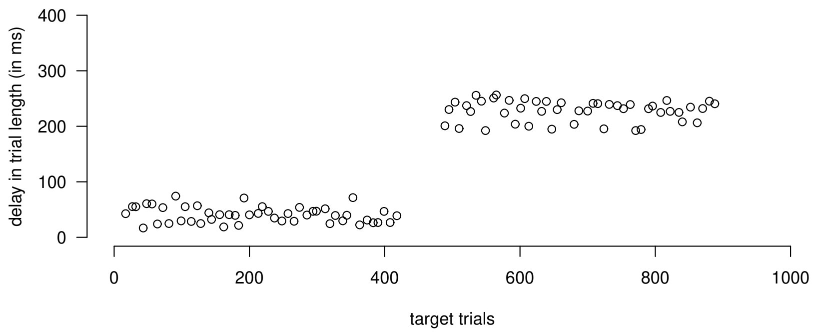


Figure 4.1: Delay of recorded trial length when compared to supposed length of 1000 ms. Plot shows only one experiment run. The X axis shows stimID value which can be taken the progression in the experiment. Clearly visible is a shift from first level to second level.

## 4.1 Non-game data

### 4.1.1 Thought probes

Figure 4.2 shows the distribution of thought probe responses across the two levels. It can be summarised that as participants move to the second level, they slightly lose their focus, but overall there are no big changes.

### 4.1.2 Memory Test

Regarding accuracy in the memory test, almost all symbols scored a mean accuracy above 90 per cent, with the exception of one of the filler symbols in the first stimuli

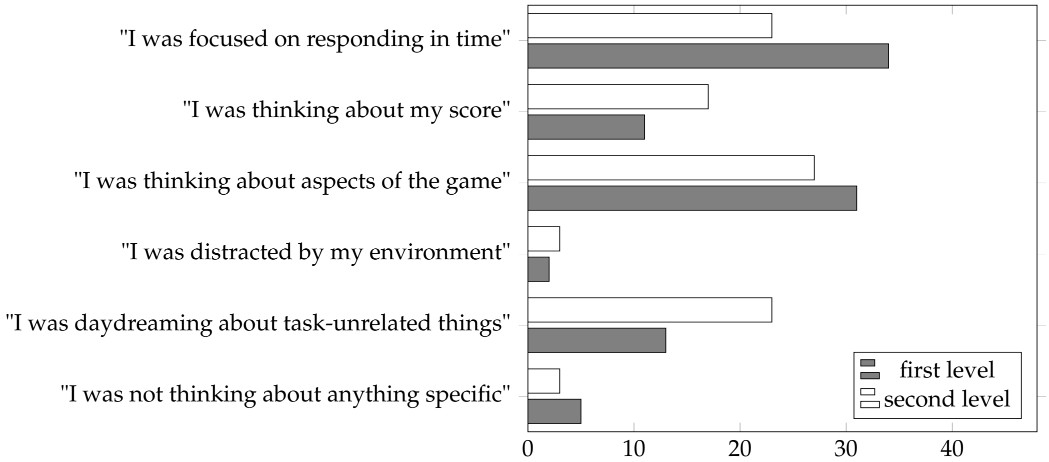


Figure 4.2: Responses to the thought probes across levels

set, which had a mean accuracy of just 63 per cent. This lower value is most likely caused by the fact that, when participants had to choose between this symbol and a foil symbol as part of the memory test, the foil symbol looked very similar to a different target symbol. This oversight in the design of the memory test should not affect the experiment's results. The affected stimulus was also the only one for which participants took significantly longer to decide on an answer (4 seconds on average), than other symbols (between 1 and 3 seconds on average). No further relevant differences were found, neither regarding accuracy nor reaction speed. There were no visible differences between filler and target stimuli.

### 4.1.3 Interviews

In the post-experiment interview, 25 participants indicated being aware of the sequence before the start of the second level, with an additional seven claiming they became aware during the second level.

All participants indicated that they found it easy to complete the memory test. However, some did admit that they found certain symbols easier than others, which aligns with the recorded data. Regarding the main section of the experiment, most participants responded that they did not use any strategies beyond simply trying to press quickly after the trigger symbol appeared. Some noteworthy exceptions were individual instances of participants saying they were trying to count the number of symbols between appearances of the trigger symbol. However, those participants also mentioned that they quickly abandoned this strategy once they noticed that it was not effective. One participant noted that they were focusing only on parts of the trigger symbol (e.g. a curve or wiggle), rather than the entire symbol. Similarly, another participant noted that they were focused only on the area of the screen where the stimuli appeared, to the point that they did not recall things such as seeing the aliens react to being hit by a laser or the score changing.

Two participants noted that they experienced an urge to press right before the trigger symbol appeared but restrained themselves from pressing until they felt it was safe. It is possible that more participants experienced this but did not bring it up in the interview, as they were not explicitly asked about it.

Overall, participants indicated they enjoyed the game and found it easy to stay focused. Some participants went so far as to say they would gladly play the game in their free time.

## 4.2 Game data

For analysis of reaction time during the game parts of the experiment, only the target stimuli were looked at, with stimuli of one sequence grouped together as one data point. Reaction times were calculated as key presses relative to the onset of the trigger symbol, meaning any key press before the trigger symbol would give a negative reaction time. Only the first key press was considered for further analysis for each group. Out of all data points, only 34 were below zero, meaning that in the majority of encounters, participants pressed only after the known target symbol had appeared.

The initially planned analysis regarding participants' ability to press in the expected time window was not possible from the data obtained due to issues with the experiment implementation (further discussed in Section 5). All analysis will therefore focus only on relative speed-ups in reaction time inside one phase of the experiment rather than the specific times obtained.

For the sake of completeness, it can still be mentioned that the average reaction time in the first level was 416.3 ms (with a minimum of 0.3 ms and a maximum of 1485.0 ms), and the average reaction time in the second level was 452.5 ms (with a minimum of 0.3 ms and a maximum of 1501.3 ms). Figure 4.3 shows the average time of key presses in each of the four conditions for both levels. Notable is that it appears that participants responded faster in later blocks of the experiment but overall slower in the second level (represented by the blue lines in Figure 4.3). However, it is not yet clear if these differences are significant and if there are any differences between the elemental learning phase (taking place in blocks one through four) and the compound learning phase (taking place in blocks five through eight).

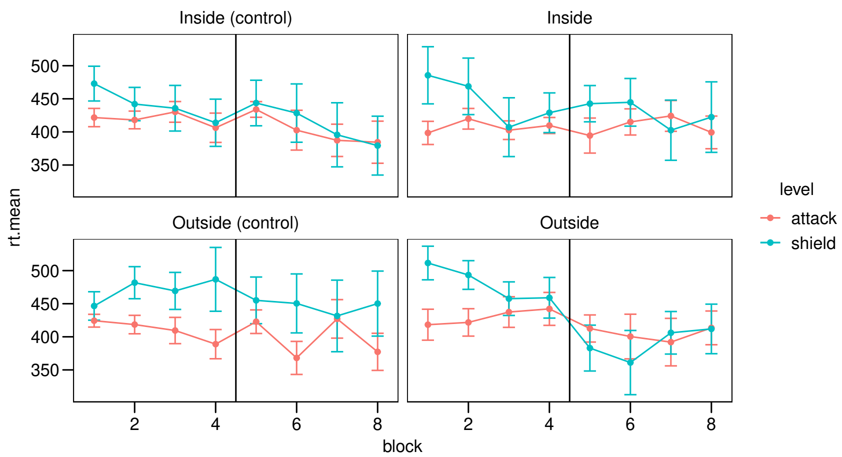


Figure 4.3: Averaged reaction time over experiment blocks divided by experiment conditions and levels. Error bars indicate standard error. Blocks 1-4 correspond to the elemental learning phase and blocks 5-8 correspond to the compound learning phase.

To gain more insight into the obtained data, a statistical analysis involving General Additive Mixed Models was performed in R. Data from each experiment phase was fit to a separate model for better comparison across phases. Models were fit using the *Maximum Likelihood* (ML) method and logRTs with values below zero excluded.

Random effects were assumed for all models to be a random intercept for participants and stimuli set. When the change in RT over the course of the learning phases was modelled, a random factor smooth for the number of times each participant had already encountered the target sequence was included. This aimed to improve model fit and reduce possible auto-correlation in the model's residuals. The fixed effects for each model were fitted using forward fitting, the process and results of which will be outlined in the following sections.

For the first level, a separate model was fit using only participants who reported being unaware of the sequence. This secondary analysis was not possible for the second level, as the number of remaining unaware participants was not high enough across all conditions. The output for all models listed is given in Appendix A.

### 4.2.1 Elemental Phase

For the RTs collected during the first day (i.e. the first learning phase) of the first level, a base model with no fixed effects beyond encounters explained 35.0 per cent of the maximum deviance to explain. A model with an added fixed effect for the condition explained the same deviance and did not significantly improve the model fit ($\chi^2(3.00) = 0.958, p > .1$). Similarly, adding an interaction between condition and encounter decreases explained deviance to 34.6 per cent and did not significantly improve upon the previous model ($\chi^2(8.00) = 4.081, p > .1$). Removing participants who reported being aware of the sequence did not improve the model's fit, leading to an explained deviance of just 28.3 per cent, much lower than any of the models using all data.

Importantly, no significant difference across conditions was found in any of the models. Figure 4.4 shows the overall trend of reaction times over encounters, highlighting that there appears to be no underlying speed-up found in this learning phase.

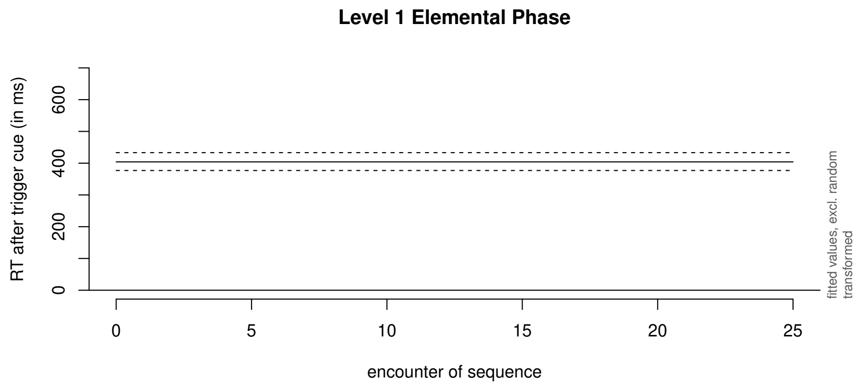


Figure 4.4: Fitted values of GAM model for elemental phase in the first level. All random effects are removed. Dotted lines indicate 95 per cent confident intervals.

Repeating the same analysis for the second level, a base model explains 37.1 per cent of the maximum deviance. This model is not significantly improved by neither a fixed effect for condition ($\chi^2(3.00) = 1.469, p > .1$) nor a further interaction between condition and encounters ($\chi^2(8.00) = 6.602, p > .1$). As in the first level, no general slope remains once random effects are removed, suggesting there was no general speed-up in this learning phase.

### 4.2.2 Compound Phase

For the second learning phase of the first level, a base model with no fixed effects beyond encounters explained 31.5 per cent of the maximum deviance to explain. As before, a model with an added fixed effect for condition explained the same deviance and did not significantly improve the model fit ($\chi^2(3.00) = 1.197, p > .1$). Adding an interaction between condition and encounter decreases explained deviance slightly to 31.1 per cent and does not significantly improve the model ($\chi^2(8.00) = 2.051, p > .1$). Removing participants who reported being aware of the sequence does not improve the model's fit, leading to an explained deviance of just 21.3 per cent.

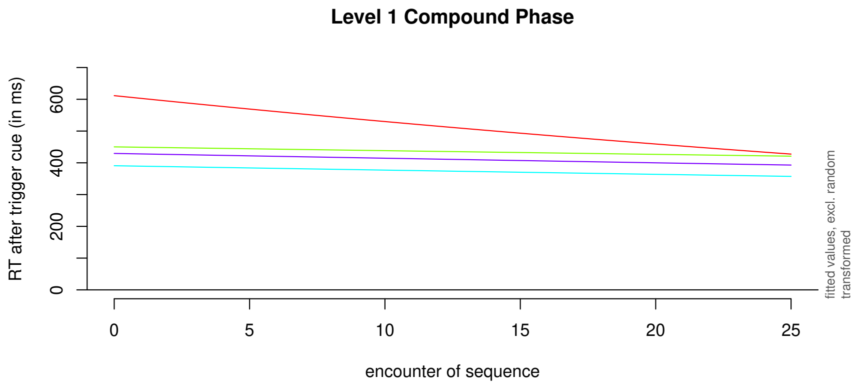


Figure 4.5: Fitted values of GAM model for compound phase in the first level. All random effects are removed. Colours indicate different conditions: *Inside* in green, *Inside (control)* in red, *Outside* in purple and *Outside (control)* in cyan.

Although the interaction model was not preferred, it is still interesting to look at the different slopes across conditions. As Figure 4.5 shows, the fitted value for the *Inside (control)* condition is high at the start of the compound phase. However, this is not occurring when looking at data of only unaware participants (see Figure 4.6). This suggests that this increase was caused by participants who were already aware of the hidden cues and were caught off guard by the change to what they had consciously picked to use as a strategy.

Repeating the same analysis for Level 2, a base model explains 32.5 per cent of the maximum deviance. This model is again not significantly improved by a fixed effect for condition ($\chi^2(3.00) = 0.289, p > .1$). A model including a further interaction between condition and encounters puts the explained deviance to just 31.8 per cent and does not improve the model fit ($\chi^2(8.00) = 4.237, p > .1$). As in the first level, a general slope and difference is visible between some conditions, (see Figure 4.7), but it is not statistically significant.

### 4.2.3 Test Phase

For the test phase, the analysis mainly focuses on changes in the average response time during the whole period, as the length of just six measurements is insufficient to find any slopes. However, all models are still fit with a factor smooth for the

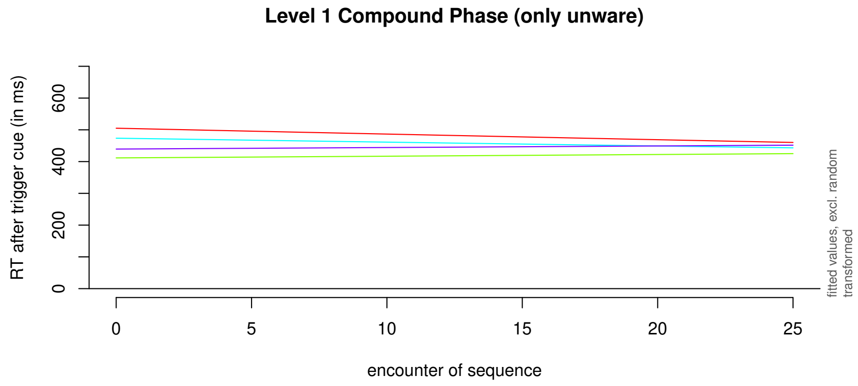


Figure 4.6: Fitted values of GAM model for compound phase in the first level, but using only data from unaware participants. All random effects are removed. Colours indicate different conditions: *Inside* in green, *Inside (control)* in red, *Outside* in purple and *Outside (control)* in cyan.

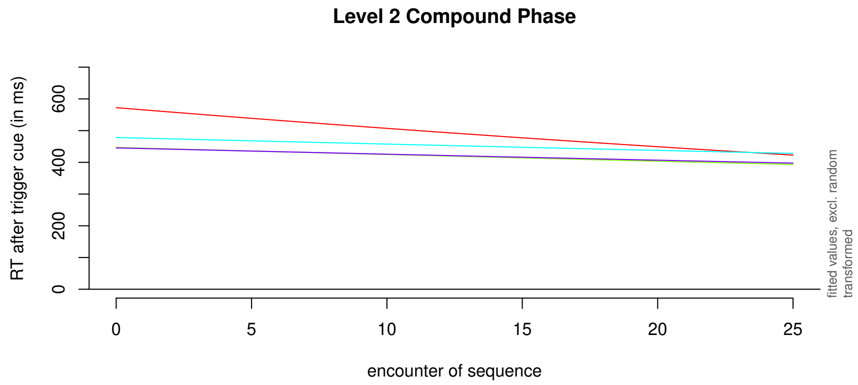


Figure 4.7: Fitted values of GAM model for compound phase in the second level. All random effects are removed. Colours indicate different conditions: *Inside* in green, *Inside (control)* in red, *Outside* in purple and *Outside (control)* in cyan.

encounter variable as there may still be individual differences caused by participants being distracted or otherwise delayed during a trial.

For the test phase in the first level, the base model includes a fixed effect between test conditions (i.e. which cue people were tested on). This model explains 65 per cent of the deviance to be explained. However, the intercept adjustment between the two test conditions is insignificant ($p > .1$). Adding a distinction by condition does not improve the model fit ($\chi^2(3.00) = 0.320, p > .1$), and neither does adding an interaction between test and condition ($\chi^2(3.00) = 1.613, p > .1$). Removing participants who reported being unaware did not change the model fit, with the intercept adjustment for test conditions still being insignificant. Figure 4.8 shows the fitted results across conditions, highlighting how neither learning nor test condition significantly affected the reaction speed in the test phase.

Repeating the same procedure for the second level, the base model using just the test variable explains 63.3 per cent of deviance. However, as before, the intercept adjustment for test conditions is non-significant. The model with a further fixed effect for learning conditions is also not significantly better than the base

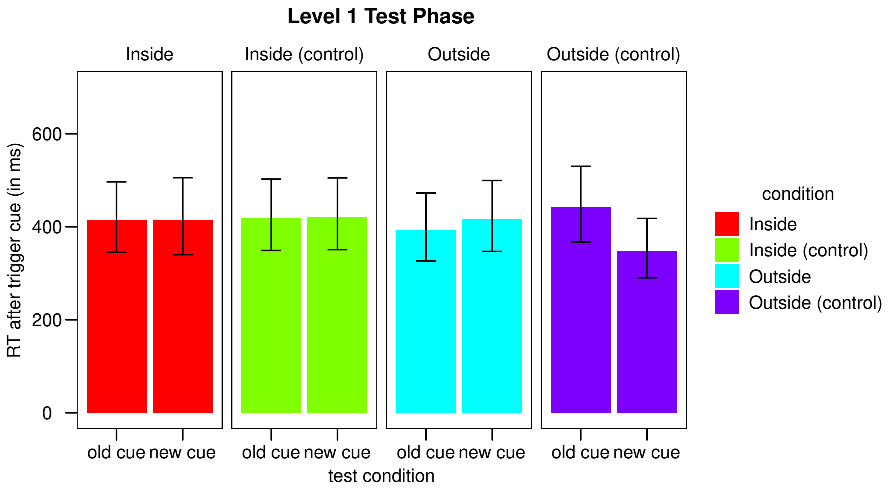


Figure 4.8: Fitted values of GAM model for test phase in the first level. All random effects are removed. Error bars indicate 95 per cent confident intervals.

model ($\chi^2(3.00) = 0.431, p > .1$). The model with an added interaction is not significantly better than the one without ($\chi^2(3.00) = 1.210, p > .1$).

Lastly, some data on the general fit of the discussed models: Figure 4.9 shows the distribution of residuals for the model used in the analysis of the test phase in the first level, which explained 65 per cent of deviance. We can see that the residuals have heavy tails on both ends, suggesting that the residuals are not normally distributed. This finding extend to all other models.

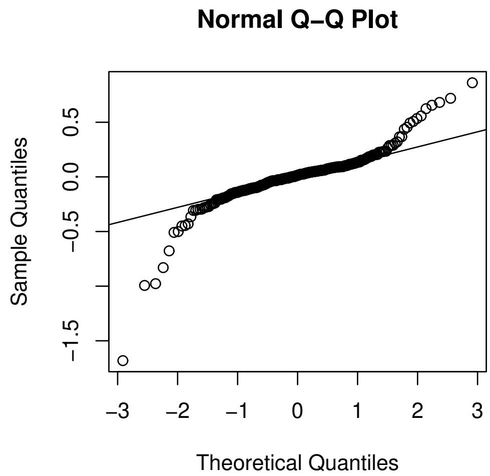


Figure 4.9: Distribution of residuals for model of test phase in the first level

# 5 Discussion

The research question of this thesis was to investigate differences in reaction times when comparing blocking paradigms involving different temporal relations. The simulation study aimed to identify open issues regarding the cognitive modelling of such differences. The human experiment aimed to collect more data on the identified issues, namely the change in blocking depending on the order placement of a novel cue in or before an established sequence. Further, we were interested in whether the original sequence would be unlearned and if prediction accuracy varied in different conditions.

The results of the experiment indicate no significant differences in behaviour regardless of where a new cue is placed. There was no change in performance within the learning phases and no differences in performance during the test phase. This suggests that in our experiment, participants did not use the provided hidden cues and instead focused only on the explicit signal of the trigger cue for their responses. Data regarding the accuracy of estimates made by participants was not able to be obtained.

Before interpreting these findings in the context of cognitive theory, the following section will first discuss the possible limitations of the experiment in both implementation and design, which will inform whether any conclusions may be drawn.

## 5.1 Limitations of Software

As mentioned in Section 4, the data used for analysis included large amount of noise. This noise was found most clearly in the recorded length of each trial, which deviated considerably from the supposed length. This suggests that in this implementation, the experiment was either unable to record accurate values or unable to produce them in the first place. It is likely that this also extended to the onset of cues, the onset of the feedback and the length of timers for the shield and laser in the respective levels.

This issue is likely caused by the fact that PsychoPy uses a timer-based implementation where each event is given a start condition checked at every frame. It is possible that, because each learning phase was coded as one long continuous trial and because the game involved loading a lot of large images, PsychoPy could not accurately perform these checks. This hypothesis is also supported by the fact that the deviation from the supposed trial length as shown in Figure 4.1 is different in the two levels, suggesting that the amount of code to run through resulted in different deviations from the supposed time intervals.

An implication of this inconsistency is that, if the inter stimuli intervals were not as intended, it may be that the temporal relations of the target stimuli were inconsistent during the experiment. As outlined in Section 1.3.2 and confirmed by the simulations in Section 2.2.2, if temporal relations are inconsistent, they do not act as an informative cue, and no blocking can take place. This could partially explain the lack of effect the presence of these cues had on participants' responses.

Further, the failure to produce accurate timings also leads to unreliable information about reaction times. For instance, the average key press time in the second level was 452.5 ms, which is 50 ms before a press would be correct. This would suggest that, on average, all first presses in the second level were incorrect. While this is possible in theory, it seems unlikely as participants would most likely adjust their behaviour if they kept pressing too early for their response to be rewarded. Indeed, the high scores obtained by participants indicate that few mistakes were made. Even more so, as one participant received the perfect score even though all their presses were in the 450 ms range, we can rule out bad performance and must instead conclude fault in the obtained data.

For the sake of having any data for analysis, it was assumed that the relative speed-up of values to each other is still accurate (i.e., it could be observed if participants sped up over time). However, it should be stressed that there is no certain way of knowing even this. The reason it was nonetheless assumed that the data on speed-ups were accurate is because there are visible changes in the data at moments when people reported becoming aware of the sequence. Therefore, any changes in behaviour caused by implicit learning of the sequence would also have been visible, had they occurred.

## 5.2 Limitations of Experiment Design

As mentioned in the analysis, the experiment is limited by the possibility of people becoming aware of the hidden sequence at some point during the task. This seems inevitable as the game is simple and played for around 30 minutes. However, as we saw in Section 4, the result does not improve when looking at only unaware participants, at least not to the point of results becoming significant. This suggests that people rarely picked up on the sequence implicitly.

Alternatively, some participants reported suppressing the urge to respond before the trigger symbol appeared. This suggests that though they had formed a reflex-like response through implicit learning, their executive control intervened. Similarly, some participants noted that they were trying to ‘experiment’ with the timing of their responses, but noticed that this often led to a loss of points, when straying too far from the right time. At the same time, due to the binary nature of the feedback there is no incentive for participants to optimise their responses, as a response which is just good enough earns as many points as one which is perfect. Because of this, many participants took a conservative approach with their presses and only pressed when they were sure it was safe to do so, i.e. once the trigger symbol had appeared. This again suggests that explicit reasoning about the game’s rules intervened with any implicit knowledge that may have been picked up. Another indicator for this is also the high number of responses to the thought probes regarding participants thinking about meta aspects of the game. A further explanation for the lack of implicit learning could be the number of participants who reported that they were not paying much attention to anything except the trigger cue itself or even just parts of the shapes that made up the trigger cue, which may have limited the amount of implicit information they were able to pick up.

In summary, while some participants were suppressing what they had implicitly learned, others were not approaching the game implicitly at all. The answer to both these problems might be to make the game less strict. Participants seemed scared to press at the wrong time, resulting in either conservative behaviour or extreme concentration. Therefore, it would be better to make the feedback given to participants less binary. For instance, giving feedback on a scale would allow more room for responses that are too late but close to being rewarded, while also communicating to participants that there is an optimal time to react, where the reward will be highest.

Similar to the issue of binary feedback, the present game may have also involved some ambiguity as to when exactly the ideal time to press was. Particularly in the second level, it often occurred that the participants would shoot their laser slightly too early, causing it to hit the alien shield at first, but then appearing to 'break-through' as the aliens dropped their shield. Participants may have needed clarification on if it was them who caused the aliens to drop their shield at that time or if it would have happened regardless. Again, introducing different levels of feedback might convey better that participants had pressed slightly too early.

It may also be considered to change the design of the experiment altogether, namely regarding when feedback is presented. In Amundson and Miller (2008), the feedback was given more directly regarding when the rat made a prediction. If the mouse rat too early, it would have to wait longer for its reward than if it was on time. A similar idea could be implemented in the alien game, for instance, by having participants have to hold the button down for the shield to go up and loosing points with each second. This would mean that there would be an incentive to learn the timing of the attacks, and to be efficient with the use of the shield. This might place more attention on the temporal relation between cues than in the current implementation.

This brings us to the last limitation of the design, the predictive power of the cues used. As outlined in Section 3, the experiment used filler cues to obscure the presence of the hidden cues. Additionally, the time intervals between cue onsets were kept much shorter than in the experiments by Amundson and Miller, where the time between cues could be up to 15 seconds. It can therefore be summarised that in the present experiment, there is an increased pool of possible stimuli and a decreased range of time intervals. If we recall the principles of discriminative learning and the temporal coding hypothesis, namely the importance of predictive power, we may suspect that in our experiment, each cue held significantly less predictive power, so there were few associations for participants to learn.

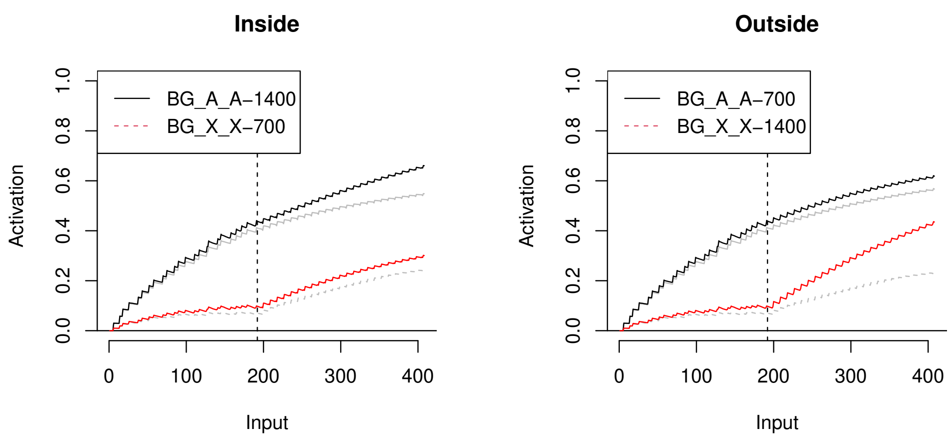


Figure 5.1: Simulated learning trajectories for placement of a new cue *Inside* or *Outside* an established sequence, when each cue is less powerful due to presence of filer cues. Coloured lines show trajectories with transference, compare to trajectories without in grey. Dotted line marks the switch from first learning phase to second.

Figure 5.1 shows the results of a simulation similar to the one discussed in Section 2.2.2, with the difference being the inclusion of filler stimuli and time intervals being equivalent to those used in the present experiment. Note that the number of exposures within one learning phase is also just 24 times per stimulus, as it was for the human participants. In this setup, activation values are generally much lower and do not increase as quickly. This is because, as the simulation is using an interference-based model, if there is more information, there is also more interference and, therefore, each cue's impact is much smaller. From this simulation, we may conclude that the inclusion of filler cues in the conducted experiment lessens the impact of the target cues, partially explaining the lack of findings in the human experiment.

## 5.3 Interpreting Results

Putting these limitations aside, let us now interpret what the experiment showed with respect to what it tells us about learning temporal relations.

First of all, we do see some adjustment when switching from the first level to the second; this suggests that participants are aware of the temporal relations overall, as they can use them to make a broad prediction of having to wait a few moments before responding, similar to the various examples outlined at the very beginning of this thesis. However, it appears that the only temporal relation people learned is between the trigger cue and the feedback (i.e. the aliens attacking or the aliens lowering their shield). Focusing on the temporal relation after the trigger symbol is a reasonable strategy, as it helps participants automate their response and only leaves them waiting to anticipate the trigger symbol, which could be accomplished via the cue symbols. However, it appears that participants did not use whatever information was provided by the hidden cues for one reason or another. It could be that the cues were not perceived as reliable by participants, and that once the sequence changed, people abandoned the idea of using the hidden cue to guide them. On the other hand, it is more likely that participants never made use of the cues at all, as even in the very first learning phase, the data shows no significant changes in reaction speed, showing that people did not or were not able to improve their ability to anticipate the trigger symbol any further.

This lack of learning in the elemental and compound phase is also relevant for reflecting on the results of the testing phase, where no significant differences were found regardless of which cue was used as a hidden cue. The data might be taken to suggest that participants need both cues to be present to make use of learned information, but considering the apparent lack of performance changes in previous phases, it is more plausible that this is not the case, and instead, participants did not learn anything about the hidden sequence at all.

## 5.4 Future Research

As participants did not learn any associations of the hidden cues temporal relations in our experiment, it is worth first determining when and how participants implicitly learn such relations. For this, it might be worth taking a step back and re-examining some of the other experiments done by Amundson and Miller. For example, it may be worth-wile to replicate the experiments using simultaneous cues regarding reliable and unreliable temporal relations in blocking paradigms, in order to verify that temporal relations inform learning in any way at all.

Similarly, if we want to show transference being used, we should first determine what happens to transferred information if a cue disappears unexpectedly. For example, one might create an experiment using a simple cue sequence $A - B - C$ and investigate if $A$ on its own will lead to anticipation of $C$, and if this changes if we introduce information that should prompt unlearning of the $B - C$ connection. This would show how much of the $A - C$ connection is created by transference. Knowing this would allow to more accurately test for transference in scenarios such as the one in the performed experiment.

Zooming out further, we return to the question of network representations being a good approach to modelling temporal relations. So far, we have only examined the distance between two cues as relevant. However, we could just as well look at other time-related variables, such as frequency rate, duration or order. All of these have equal claim to being encoded as a cue in our network. However, doing so would lead to an explosion in model complexity and it is worth asking if there is a better way to go about modelling time, that requires less hard-wiring of cues and attributes.

### 5.4.1 PRIMs

The previously mentioned cognitive architecture PRIMs (Taatgen, 2013) can model the whole cognitive process from the act of perceiving stimuli to the response being given, with each step, like retrieval of information from memory, being performed and recorded. PRIMs can also model implicit learning using so-called *context operator learning*. Within PRIMs, an operator is an action that the model can take, such as, for example, taking whatever information was retrieved from memory and saying it out loud. Normally these operators are executed when specific conditions are met, but, in context-operator learning, certain operators are more likely to activate depending on if the model is in a specific state. This method has been used to model learning biases, where the model is more likely to respond a certain way if exposed to a certain type of word (Toth et al., 2022). Using the context at the moment in time to adapt the behaviour could simulate a learner anticipating the appearance of a trigger symbol if certain stimuli are still in working memory. The learning rule used to establish the connections between model context and operators could be changed to create discriminative learning, which should allow for (at the very least) the modelling of blocking effects in situations where both cues are simultaneous.

Regarding encoding temporal relations, PRIMs can explicitly store and retrieve information about time intervals. A model could be used to form estimates of the interval between the last seen cue and the current cue and then use this interval as another cue that is part of the background context involved in context-operator learning. The exact details of encoding would of course be specific to PRIMs as the cognitive architecture used, in the same way as the encoding used in the simulation discussed in Section 2 was specific to the edl package developed by van Rij and Hoppe. As such, there may be new issues that arise when using PRIMs, but the exploration of these is left to future research.

The advantage of using PRIMS would be, that by going beyond a network-based model, it would allow behaviours to emerge naturally. For instance, we recall the alternative information-theoretical explanation for the findings of Amundson and Miller (2008), that the cue $X$ is not blocked when it occurs outside of an established sequence, because it still holds some predictive power. It may be that a less discrete model like PRIMs can replicate the findings of Amundson and Miller without ever needing to rely on explicit strategies like transference.

## 5.5 A waste of time?

It is hopefully clear by now that time is an important factor when considering the learning of information. As the work by Amundson and Miller as well as the experiments of this thesis show, introducing changes in time to otherwise well-established learning phenomena can create unexpected results. Adding the dimension of time to an otherwise simple experiment can shift what participants pay attention to, which information is useful to them and how they use it. The fact that just changing the temporal placement of a single item can completely break an established theory means that a reliable way of predicting the effect of temporal relations on learning is needed. In real life, information is rarely presented simultaneously or occurs in the same order each time. Increasing our understanding of time and creating a model that can predict its effect on other learning processes would be useful for having the means to rule out time as a factor causing unexpected results.

Another reason why time is worth investigating becomes clear when we return to the differences between discriminative models and other ways to model aspects of implicit learning. Features like blocking or unlearning require adjustments and special treatments in other models, while in discriminative models, they emerge naturally. Discriminative models can explain more with less and models which treat events occurring in sequential time, rather than in discrete moments, may offer a similar advantage. Things which currently seem like exceptions to rules might be perfectly in line with what we would expect to happen under the right theory. As discussed in previous sections, even the need for a concept like transference might evaporate once we start to expand our theories of what it means for a cue to hold predictive power.

## 5.6 Conclusion

This study was concerned with studying how changes in temporal relations between items affect the learning phenomena of blocking, using a discriminative approach to learning. A simulation study was conducted to identify what is and is not possible using network-based models of such learning to show how the concept of transference could play a role in filling these gaps. An experiment using human participants was conducted to confirm these findings. However, it was found that in the used set-up, participants did not adjust their behaviour in a way that suggested the learning of information. These findings were interpreted as the result of limitations specific to the experiment's method and design. It is recommended that the study is replicated with more attention paid to what might affect how participants engage with the stimuli and how more attention could be drawn to the predictive power of the time intervals between cues.

# A Model Output

## A.1 Elemental Phase

### A.1.1 Level 1

**Formula**

$$\begin{aligned} logRT \quad & \sim s(encounter, pID, bs = "fs", m = 1) \\ & + s(stimSet, bs = "re") \end{aligned}$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 6.00154 | 0.03545 | 169.3 | <2e-16 | *** |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 105.06 | 422 | 1.088 | <2e-16 | *** |
| s(stimSet) | 1.01 | 2 | 1.517 | <2e-16 | *** |

Table A.1: Model output for a model of the elemental phase in the first level with fixed effects for encounter and random effects for participant and item differences

**Formula**

$$\begin{aligned} logRT \quad & \sim condition \\ & + s(encounter, pID, bs = "fs", m = 1) \\ & + s(stimSet, bs = "re") \end{aligned}$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 6.01910 | 0.05466 | 110.114 | <2e-16 | *** |
| condition: Inside | -0.02227 | 0.06987 | -0.319 | 0.750 | |
| condition: Outside (cont) | -0.06968 | 0.06830 | -1.020 | 0.308 | |
| condition: Outside | 0.02142 | 0.06829 | 0.314 | 0.754 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 102.215 | 419 | 1.054 | <2e-16 | *** |
| s(stimSet) | 1.048 | 2 | 1.633 | <2e-16 | *** |

Table A.2: Model output for a model of the elemental phase in the first level with fixed effects for encounter and condition and random effects for participant and item differences

**Formula**

$$\begin{aligned} logRT &\sim condition \\ &+ s(encounter, by = condition) \\ &+ s(encounter, pID, bs = "fs", m = 1) \\ &+ s(stimSet, bs = "re") \end{aligned}$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 6.01909 | 0.05471 | 110.019 | <2e-16 | *** |
| condition: Inside | -0.02222 | 0.06995 | -0.318 | 0.751 | |
| condition: Outside (cont) | -0.06964 | 0.06837 | -1.019 | 0.309 | |
| condition: Outside | 0.02137 | 0.06836 | 0.313 | 0.755 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter):Inside (cont) | 1.000 | 1 | 0.022 | 0.88093 | |
| s(encounter):Inside | 1.000 | 1 | 0.150 | 0.69878 | |
| s(encounter):Outside (cont) | 1.000 | 1 | 8.379 | 0.00388 | ** |
| s(encounter):Outside | 1.000 | 1 | 0.019 | 0.89135 | |
| s(encounter,pID) | 94.047 | 419 | 0.991 | 8.57e-07 | *** |
| s(stimSet) | 1.047 | 2 | 1.628 | <2e-16 | *** |

Table A.3: Model output for a model of the elemental phase in the first level with fixed effects for encounter, condition and their interaction and random effects for participant and item differences

**Formula**

$$\begin{aligned} logRT &\sim condition \\ &+ s(encounter, by = condition) \\ &+ s(encounter, pID, bs = "fs", m = 1) \\ &+ s(stimSet, bs = "re") \end{aligned}$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 6.073597 | 0.021904 | 277.277 | <2e-16 | *** |
| condition: Inside | -0.087186 | 0.029879 | -2.918 | 0.00368 | ** |
| condition: Outside (cont) | 0.003213 | 0.037997 | 0.085 | 0.93264 | |
| condition: Outside | 0.083810 | 0.030977 | 2.706 | 0.00705 | ** |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter):Inside (cont) | 1.000 | 1 | 1.655 | 0.199 | |
| s(encounter):Inside | 1.000 | 1 | 0.397 | 0.529 | |
| s(encounter):Outside (cont) | 1.000 | 1 | 0.828 | 0.363 | ** |
| s(encounter):Outside | 1.000 | 1 | 2.373 | 0.124 | |
| s(encounter,pID) | 2.123e+01 | 194 | 0.307 | <2e-16 | *** |
| s(stimSet) | 3.330e-05 | 2 | 0.000 | 0.259 | |

Table A.4: Model output for a model of the elemental phase in the first level with fixed effects for encounter, condition and their interaction and random effects for participant and item differences **with only unaware participants**

### A.1.2 Level 2

**Formula**

$$logRT \sim s(encounter, pID, bs = "fs", m = 1) + s(stimSet, bs = "re")$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 6.08191 | 0.02882 | 211 | <2e-16 | *** |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 1.084e+02 | 422 | 1.151 | <2e-16 | *** |
| s(stimSet) | 1.625e-04 | 2 | 0.000 | 0.109 | |

Table A.5: Model output for a model of the elemental phase in the second level with fixed effects for encounter and random effects for participant and item differences

**Formula**

$$logRT \sim condition + s(encounter, pID, bs = "fs", m = 1) + s(stimSet, bs = "re")$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 6.036447 | 0.055246 | 109.265 | <2e-16 | *** |
| condition: Inside | -0.003434 | 0.078186 | -0.044 | 0.965 | |
| condition: Outside (cont) | 0.085093 | 0.079930 | 1.065 | 0.287 | |
| condition: Outside | 0.103473 | 0.078147 | 1.324 | 0.186 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 1.053e+02 | 419 | 1.101 | <2e-16 | *** |
| s(stimSet) | 1.701e-04 | 2 | 0.000 | 0.132 | |

Table A.6: Model output for a model of the elemental phase in the second level with fixed effects for encounter and condition and random effects for participant and item differences

**Formula**

$$\begin{aligned} logRT \quad & \sim condition \\ & + s(encounter, by = condition) \\ & + s(encounter, pID, bs = "fs", m = 1) \\ & + s(stimSet, bs = "re") \end{aligned}$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 6.036468 | 0.055264 | 109.230 | <2e-16 | *** |
| condition: Inside | -0.003597 | 0.078212 | -0.046 | 0.963 | |
| condition: Outside (cont) | 0.085028 | 0.079956 | 1.063 | 0.288 | |
| condition: Outside | 0.103448 | 0.078173 | 1.323 | 0.186 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter):Inside (cont) | 1.000 | 1 | 2.956 | 0.08586 | . |
| s(encounter):Inside | 1.000 | 1 | 8.969 | 0.00281 | ** |
| s(encounter):Outside (cont) | 1.000 | 1 | 0.002 | 0.96041 | |
| s(encounter):Outside | 1.000 | 1 | 3.164 | 0.07559 | . |
| s(encounter,pID) | 8.781e+01 | 419 | 0.973 | <2e-16 | *** |
| s(stimSet) | 4.326e-05 | 2 | 0.000 | 0.13516 | |

Table A.7: Model output for a model of the elemental phase in the second level with fixed effects for encounter, condition and their interaction and random effects for participant and item differences

## A.2 Compound Phase

### A.2.1 Level 1

**Formula**

$$\begin{aligned} logRT \quad & \sim s(encounter, pID, bs = "fs", m = 1) \\ & + s(stimSet, bs = "re") \end{aligned}$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 5.91664 | 0.04088 | 144.7 | <2e-16 | *** |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 1.036e+02 | 422 | 0.861 | <2e-16 | *** |
| s(stimSet) | 2.156e-03 | 2 | 0.002 | 0.000193 | *** |

Table A.8: Model output for a model of the compound phase in the first level with fixed effects for encounter and random effects for participant and item differences

**Formula**

$$\begin{aligned} logRT \quad & \sim condition \\ & + s(encounter, pID, bs = "fs", m = 1) \\ & + s(stimSet, bs = "re") \end{aligned}$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 5.89327 | 0.07944 | 74.188 | <2e-16 | *** |
| condition: Inside | 0.11830 | 0.11373 | 1.040 | 0.299 | |
| condition: Outside (cont) | -0.05583 | 0.11123 | -0.502 | 0.616 | |
| condition: Outside | 0.03950 | 0.11120 | 0.355 | 0.722 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 100.5516 | 419 | 0.839 | <2e-16 | *** |
| s(stimSet) | 0.1518 | 2 | 0.123 | 9.55e-05 | *** |

Table A.9: Model output for a model of the compound phase in the first level with fixed effects for encounter and condition and random effects for participant and item differences

**Formula**

$$\begin{aligned} logRT \quad & \sim condition \\ & + s(encounter, by = condition) \\ & + s(encounter, pID, bs = "fs", m = 1) \\ & + s(stimSet, bs = "re") \end{aligned}$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 5.89329 | 0.07944 | 74.185 | <2e-16 | *** |
| condition: Inside | 0.11825 | 0.11383 | 1.039 | 0.299 | |
| condition: Outside (cont) | -0.05587 | 0.11133 | -0.502 | 0.616 | |
| condition: Outside | 0.03946 | 0.11129 | 0.355 | 0.723 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter):Inside (cont) | 1.0001 | 1 | 3.621 | 0.0573 | . |
| s(encounter):Inside | 1.0001 | 1 | 0.116 | 0.7330 | |
| s(encounter):Outside (cont) | 1.0001 | 1 | 0.226 | 0.6345 | |
| s(encounter):Outside | 1.0001 | 1 | 0.223 | 0.6366 | |
| s(encounter,pID) | 94.3232 | 419 | 0.808 | <2e-16 | *** |
| s(stimSet) | 0.1402 | 2 | 0.113 | 0.0001 | *** |

Table A.10: Model output for a model of the compound phase in the first level with fixed effects for encounter, condition and their interaction and random effects for participant and item differences

**Formula**

$$\begin{aligned} logRT \quad & \sim condition \\ & + s(encounter, by = condition) \\ & + s(encounter, pID, bs = "fs", m = 1) \\ & + s(stimSet, bs = "re") \end{aligned}$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 6.08908 | 0.03631 | 167.720 | <2e-16 | *** |
| condition: Inside | -0.02219 | 0.03940 | -0.563 | 0.573 | |
| condition: Outside (cont) | -0.02535 | 0.05105 | -0.497 | 0.620 | |
| condition: Outside | 0.03655 | 0.04175 | 0.875 | 0.382 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter):Inside (cont) | 1.000 | 1 | 3.422 | 0.0649 | . |
| s(encounter):Inside | 1.000 | 1 | 0.467 | 0.4945 | |
| s(encounter):Outside (cont) | 1.000 | 1 | 0.851 | 0.3567 | |
| s(encounter):Outside | 1.000 | 1 | 0.300 | 0.5842 | |
| s(encounter,pID) | 12.750 | 194 | 0.367 | 0.0015 | ** |
| s(stimSet) | 1.287 | 2 | 2.393 | 1.04e-06 | *** |

Table A.11: Model output for a model of the compound phase in the first level with fixed effects for encounter, condition and their interaction and random effects for participant and item differences **with only unaware participants**

### A.2.2 Level 2

**Formula**

$$logRT \sim s(encounter, pID, bs = "fs", m = 1) + s(stimSet, bs = "re")$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 5.94017 | 0.05133 | 115.7 | <2e-16 | *** |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 5.722e+01 | 422 | 1.063 | <2e-16 | *** |
| s(stimSet) | 4.838e-04 | 2 | 0.000 | 0.349 | |

Table A.12: Model output for a model of the compound phase in the second level with fixed effects for encounter and random effects for participant and item differences

**Formula**

$$logRT \sim condition + s(encounter, pID, bs = "fs", m = 1) + s(stimSet, bs = "re")$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 5.907820 | 0.101007 | 58.489 | <2e-16 | *** |
| condition: Inside | 0.008805 | 0.142788 | 0.062 | 0.951 | |
| condition: Outside (cont) | 0.100587 | 0.145982 | 0.689 | 0.491 | |
| condition: Outside | 0.025600 | 0.142885 | 0.179 | 0.858 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 5.446e+01 | 419 | 1.057 | <2e-16 | *** |
| s(stimSet) | 4.728e-04 | 2 | 0.000 | 0.297 | |

Table A.13: Model output for a model of the compound phase in the second level with fixed effects for encounter and condition and random effects for participant and item differences

**Formula**

$$
\begin{aligned}
logRT \quad & \sim condition \\
& + s(encounter, by = condition) \\
& + s(encounter, pID, bs = "fs", m = 1) \\
& + s(stimSet, bs = "re")
\end{aligned}
$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 5.908080 | 0.101175 | 58.395 | <2e-16 | *** |
| condition: Inside | 0.008374 | 0.143030 | 0.059 | 0.953 | |
| condition: Outside (cont) | 0.100385 | 0.146230 | 0.686 | 0.493 | |
| condition: Outside | 0.025185 | 0.143127 | 0.176 | 0.860 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter):Inside (cont) | 1.0001 | 1 | 6.178 | 0.0131 | * |
| s(encounter):Inside | 1.0001 | 1 | 1.119 | 0.2904 | |
| s(encounter):Outside (cont) | 1.0001 | 1 | 0.779 | 0.3775 | |
| s(encounter):Outside | 1.0001 | 1 | 0.868 | 0.3518 | |
| s(encounter,pID) | 4.503e+01 | 419 | 1.024 | <2e-16 | *** |
| s(stimSet) | 7.633e-05 | 2 | 0.000 | 0.2985 | |

Table A.14: Model output for a model of the compound phase in the second level with fixed effects for encounter, condition and their interaction and random effects for participant and item differences

## A.3 Test Phase

### A.3.1 Level 1

**Formula**

$$\begin{aligned} logRT \quad & \sim testCondition \\ & + s(encounter, pID, bs = "fs", m = 1, k = 5) \\ & + s(stimSet, bs = "re") \end{aligned}$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 6.03180 | 0.05322 | 113.341 | <2e-16 | *** |
| testCondition: New Cue | -0.04509 | 0.06567 | -0.687 | 0.493 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 92.1489 | 232 | 1.049 | <2e-16 | *** |
| s(stimSet) | 0.7952 | 2 | 0.981 | 0.000226 | *** |

Table A.15: Model output for a model of the test phase in the first level with fixed effects for the test condition and encounter and random effects for participant and item differences

**Formula**

$$\begin{aligned} logRT \quad & \sim testCondition + condition \\ & + s(encounter, pID, bs = "fs", m = 1, k = 5) \\ & + s(stimSet, bs = "re") \end{aligned}$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 6.06221 | 0.07731 | 78.414 | <2e-16 | *** |
| testCondition: New Cue | -0.04417 | 0.06536 | -0.676 | 0.500 | |
| condition: Inside | -0.01611 | 0.09325 | -0.173 | 0.863 | |
| condition: Outside (cont) | -0.06938 | 0.09127 | -0.760 | 0.448 | |
| condition: Outside | -0.03689 | 0.09141 | -0.404 | 0.687 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 89.8502 | 229 | 1.047 | <2e-16 | *** |
| s(stimSet) | 0.8304 | 2 | 1.051 | 0.00021 | *** |

Table A.16: Model output for a model of the test phase in the first level with fixed effects for the test condition, experiment condition and encounter and random effects for participant and item differences

**Formula**

$$
\begin{aligned}
logRT \quad & \sim testCondition * condition \\
& + s(encounter, pID, bs = "fs", m = 1, k = 5) \\
& + s(stimSet, bs = "re")
\end{aligned}
$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 6.037414 | 0.092959 | 64.947 | <2e-16 | *** |
| testCondition: New Cue | 0.005415 | 0.124547 | 0.043 | 0.965 | |
| condition: Inside | -0.011655 | 0.124547 | -0.094 | 0.926 | |
| condition: Outside (cont) | 0.052038 | 0.125014 | 0.416 | 0.678 | |
| condition: Outside | -0.063858 | 0.125419 | -0.509 | 0.611 | |
| New:Inside | -0.003619 | 0.180590 | -0.020 | 0.984 | |
| New:Outside (cont) | -0.241703 | 0.176467 | -1.370 | 0.172 | |
| New:Outside | 0.052529 | 0.176753 | 0.297 | 0.767 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 86.6330 | 226 | 0.998 | <2e-16 | *** |
| s(stimSet) | 0.9391 | 2 | 1.302 | 0.000182 | *** |

Table A.17: Model output for a model of the test phase in the first level with fixed effects for the test condition, experiment condition, their interaction and encounter and random effects for participant and item differences

**Formula**

$$
\begin{aligned}
logRT \quad & \sim testCondition * condition \\
& + s(encounter, pID, bs = "fs", m = 1, k = 5) \\
& + s(stimSet, bs = "re")
\end{aligned}
$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 6.06060 | 0.05729 | 105.793 | <2e-16 | *** |
| testCondition: New Cue | 0.04118 | 0.05883 | 0.700 | 0.485 | |
| condition: Inside | -0.05694 | 0.06092 | -0.935 | 0.352 | |
| condition: Outside (cont) | 0.01897 | 0.08303 | 0.228 | 0.820 | |
| condition: Outside | 0.01156 | 0.06248 | 0.185 | 0.854 | |
| New:Inside | 0.06405 | 0.07744 | 0.827 | 0.410 | |
| New:Outside (cont) | -0.10076 | 0.09935 | -1.014 | 0.313 | |
| New:Outside | 0.05156 | 0.07891 | 0.653 | 0.515 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 9.543 | 102 | 0.162 | 0.0163 | * |
| s(stimSet) | 1.578 | 2 | 5.081 | 9.02e-05 | *** |

Table A.18: Model output for a model of the test phase in the first level with fixed effects for the test condition, experiment condition, their interaction and encounter and random effects for participant and item differences **with only unaware participants**

### A.3.2 Level 2

**Formula**

$$\begin{aligned} logRT &\sim testCondition \\ &+ s(encounter, pID, bs = "fs", m = 1, k = 5) \\ &+ s(stimSet, bs = "re") \end{aligned}$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 6.01387 | 0.06145 | 97.86 | <2e-16 | *** |
| testCondition: New Cue | 0.07389 | 0.08796 | 0.84 | 0.402 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 6.744e+01 | 231 | 1.217 | <2e-16 | *** |
| s(stimSet) | 6.994e-05 | 2 | 0.000 | 0.0862 | . |

Table A.19: Model output for a model of the test phase in the second level with fixed effects for the test condition and encounter and random effects for participant and item differences

**Formula**

$$\begin{aligned} logRT &\sim testCondition + condition \\ &+ s(encounter, pID, bs = "fs", m = 1, k = 5) \\ &+ s(stimSet, bs = "re") \end{aligned}$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 5.95260 | 0.09666 | 61.586 | <2e-16 | *** |
| testCondition: New Cue | 0.07358 | 0.08734 | 0.842 | 0.401 | |
| condition: Inside | 0.07415 | 0.12238 | 0.606 | 0.545 | |
| condition: Outside (cont) | 0.05972 | 0.12484 | 0.478 | 0.633 | |
| condition: Outside | 0.11147 | 0.12195 | 0.914 | 0.362 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 6.481e+01 | 228 | 1.208 | <2e-16 | *** |
| s(stimSet) | 6.773e-05 | 2 | 0.000 | 0.0862 | . |

Table A.20: Model output for a model of the test phase in the second level with fixed effects for the test condition, experiment condition and encounter and random effects for participant and item differences

**Formula**

$$logRT \sim testCondition * condition + s(encounter, pID, bs = "fs", m = 1, k = 5) + s(stimSet, bs = "re")$$

**Parametric coefficients**

| | Estimate | Std. Error | t value | Pr(>\|t\|) | |
|---|---|---|---|---|---|
| (Intercept) | 5.87802 | 0.11863 | 49.549 | <2e-16 | *** |
| testCondition: New Cue | 0.22404 | 0.16842 | 1.330 | 0.185 | |
| condition: Inside | 0.25012 | 0.16843 | 1.485 | 0.139 | |
| condition: Outside (cont) | 0.14787 | 0.16777 | 0.881 | 0.379 | |
| condition: Outside | 0.14741 | 0.16777 | 0.879 | 0.381 | |
| New:Inside | -0.35284 | 0.23861 | -1.479 | 0.141 | |
| New:Outside (cont) | -0.18045 | 0.24383 | -0.740 | 0.460 | |
| New:Outside | -0.07319 | 0.23773 | -0.308 | 0.758 | |

**Approximate significance of smooth terms**

| | edf | Ref.df | F | p-value | |
|---|---|---|---|---|---|
| s(encounter,pID) | 6.183e+01 | 225 | 1.146 | <2e-16 | *** |
| s(stimSet) | 7.111e-05 | 2 | 0.000 | 0.0963 | . |

Table A.21: Model output for a model of the test phase in the first level with fixed effects for the test condition, experiment condition, their interaction and encounter and random effects for participant and item differences

# Bibliography


Amundson, J. C., & Miller, R. R. (2008). CS-US temporal relations in blocking. *Learning & Behavior*, *36*(2), 92–103.

Anderson, J. R., Matessa, M., & Lebiere, C. (1997). ACT-R: A theory of higher level cognition and its relation to visual attention. *Human–Computer Interaction*, *12*(4), 439–462.

Baayen, R. H., Milin, P., Đurđević, D. F., Hendrix, P., & Marelli, M. (2011). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological review*, *118*(3), 438.

Barto, A. G., Sutton, R. S., & Brouwer, P. S. (1981). Associative search network: A reinforcement learning associative memory. *Biological cybernetics*, *40*, 201–211.

Block, R. A. (2014). *Cognitive models of psychological time*. Psychology Press.

Elman, J. L. (1990). Finding structure in time. *Cognitive science*, *14*(2), 179–211.

Farley, B., & Clark, W. (1954). Simulation of self-organizing systems by digital computer. *Transactions of the IRE Professional Group on Information Theory*, *4*(4), 76–84.

Hastie, T. J., & Tibshirani, R. J. (1990). *Generalized additive models* (Vol. 43). CRC press.

Honig, W. K. (1981). Working memory and the temporal map. In N. E. Spear & R. R. Miller (Eds.), *Information processing in animals: Memory mechanisms* (pp. 177–208). Psychology Press.

Hoppe, D. B., Hendriks, P., Ramscar, M., & van Rij, J. (2022). An exploration of error-driven learning in simple two-layer networks from a discriminative learning perspective. *Behavior Research Methods*, *54*(5), 2221–2251.

Kamin, L. J. 'Attention-like' processes in classical conditioning (M. R. Jones, Ed.). In: *Miami symposium on the prediction of behavior* (M. R. Jones, Ed.). Ed. by Jones, M R. Miami: University of Miami Press, 1967.

Kamin, L. J. Predictability, surprise, attention, and conditioning (R. M Campbell B. A. Church, Ed.). In: *Miami symposium on the prediction of behavior* (R. M Campbell B. A. Church, Ed.). Ed. by Campbell B. A. Church, R. M. Miami: University of Miami Press, 1967, 279–296.

Klopf, A. H. (1988). A neuronal model of classical conditioning. *Psychobiology*, *16*(2), 85–125.

Mané, A., & Donchin, E. (1989). The space fortress game. *Acta psychologica*, *71*(1-3), 17–22.

Matzel, L. D., Held, F. P., & Miller, R. P. (1988). Information and expression of simultaneous and backward associations: Implications for contiguity theory. *Learning and Motivation*, *19*(4), 317–344.

Molet, M., & Miller, R. R. (2014). Timing: An attribute of associative learning. *Behavioural Processes*, *101*, 4–14.

Moon, J., & Anderson, J. (2012). Modeling millisecond time interval estimation in space fortress game. *Proceedings of the annual meeting of the cognitive science society*, *34*.

Nixon, J. S. (2020). Of mice and men: Speech sound acquisition as discriminative learning from prediction error, not just statistical tracking. *Cognition*, *197*, 104081.

Pavlov, I. P. (1927). *Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex.* Oxford University Press.

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*(1), 195–203.

R Core Team. (2021). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. Vienna, Austria.

Ramscar, M., Dye, M., & McCauley, S. M. (2013). Error and expectation in language learning: The curious absence of "mouses" in adult speech. *Language*, 760–793.

Ramscar, M., Dye, M., Popick, H. M., & O'Donnell-McCarthy, F. (2011). The enigma of number: Why children find the meanings of even small number words hard to learn and how we can help them do better (D. Bishop, Ed.). *PLoS ONE*, *6*(7), e22501.

Ramscar, M., Matlock, T., & Dye, M. (2010). Running down the clock: The role of expectation in our understanding of time and motion. *Language and Cognitive Processes*, *25*(5), 589–615.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical conditioning II: Current research and theory*, 64–99.

Rescorla, R. A. (1988). Pavlovian conditioning: It's not what you think it is. *American Psychologist*, *43*(3), 151–160.

Savastano, H. I., & Miller, R. R. (1998). Time as content in pavlovian conditioning. *Behavioural Processes*, *44*(2), 147–162.

Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of neurophysiology*.

Siegel, S., & Allan, L. G. (1996). The widespread influence of the rescorla-wagner model. *Psychonomic Bulletin & Review*, *3*(3), 314–321.

Taatgen, N. A. (2013). The nature and transfer of cognitive skills. *Psychological Review*, *120*(3), 439–471.

Taatgen, N. A., & van Rijn, H. (2011). Traces of times past: Representations of temporal intervals in memory. *Memory & cognition*, *39*, 1546–1560.

Tesauro, G. (1995). Td-gammon: A self-teaching backgammon program. In A. F. Murray (Ed.), *Applications of neural networks* (pp. 267–285). Springer.

Thiel, A., & Spenader, J. (2022, June). *Exploration of an EDL approach in recreating Mirman, Graf Estes, and Magnuson (2010)* [unpublished research project], University of Groningen.

Thorndike, E. L. (1898). Animal intelligence: An experimental study of the associative processes in animals. *The Psychological Review: Monograph Supplements*, *2*(4), i.

Thorndike, E. L. (1927). The law of effect. *The American Journal of Psychology*, *39*(1/4), 212.

Thorndike, E. L. (1911). *Animal intelligence: Experimental studies*. Transaction Publishers.

Toth, A. G., Hendriks, P., Taatgen, N. A., & van Rij, J. (2022). A cognitive modeling approach to learning and using reference biases in language. *Frontiers in Artificial Intelligence*, *5*, 933504.

van Rij, J, & Hoppe, D. (2021). Edl: Toolbox for error-driven learning simulations with two-layer networks. *R package version*, *1*.

Van Rijn, H., & Taatgen, N. A. (2008). Timing of multiple overlapping intervals: How many clocks do we have? *Acta Psychologica*, *129*(3), 365–375.

Vidal, C., Chetail, F., & Content, A. (2016). BACS: The Brussels Artificial Characters Sets for studies in cognitive psychology and neuroscience.

Widrow, B., & Hoff, M. E. (1960). *Adaptive switching circuits* (tech. rep.). Stanford Univ Ca Stanford Electronics Labs.

Wood, S. N. (2017). *Generalized additive models: an introduction with R*. CRC press.