

Modelling the Effect of Depression on Working Memory

Maarten van der Velde

Master's thesis

Human-Machine Communication
Department of Artificial Intelligence
University of Groningen

January 2018

Primary supervisor:

Dr. Marieke K. van Vugt

Artificial Intelligence, University of Groningen

Secondary supervisor:

Prof. Dr. Niels A. Taatgen

Artificial Intelligence, University of Groningen

Contents

Abstract	3
1 Introduction	4
Cognitive effects of depression	4
Modelling depression and rumination	5
Problem statement and thesis overview	6
2 Emotional <i>n</i>-back	8
A working memory task with emotional stimuli	8
<i>n</i> -back strategies	9
Separating perceptual bias from WM differences	10
Procedure	10
Conditions	11
Results summary	12
3 Model	13
Approach	13
Modelling strategy and assumptions	13
Architecture	14
0-back model	14
Task strategy	14
Stimuli	15
Mind-wandering	15
Interleaving goals	16
Emotion-specific effects	17
Differentiating a depressed model from a control model	18
2-back model	19
Changes relative to the 0-back model	19
Differentiating a depressed model from a control model	22
4 Results	23
Analysis	23
0-back task	23
Accuracy	24
Response rate	25
Response time	26
2-back	28
Accuracy	29

Response rate	30
Response time	31
5 Discussion	35
General conclusion	35
Future work	36
Concluding remarks	37
References	38
A Model parameters	41
B Resources	42

Abstract

Individuals with depression are prone to engage in rumination, a process in which attention turns inwards to narrowly-focused, negative patterns of thought, at the cost of attending to a task. Depression is furthermore associated with certain cognitive deficits, such as increased difficulty inhibiting information that is no longer relevant and increased difficulty switching between task goals. Previous studies have investigated how these factors affect task performance, but their observations have generally only been supported by verbal theories. Here, we present a computational cognitive model of performance on one such task, an n-back task in which the stimuli are faces with different emotional expressions. Depressed participants were previously found to have a stronger attachment to sad faces than to neutral or happy faces. Compared to faces with other expressions, they integrated sad faces into working memory more quickly, while they spent more time disengaging from sad faces that were no longer relevant. By having a perceptual bias towards sad stimuli, and selectively elaborating on sad items as they are removed from working memory, our model is able to reproduce the reaction time effects observed in the original data, as well as accuracy and response rate metrics. This study demonstrates how a cognitive model can serve as an explicit, testable implementation of a verbal theory, providing a deeper understanding of the cognitive processes underlying behaviour.

Keywords: depression; rumination; mind-wandering; working memory; computational cognitive modelling.

Chapter 1

Introduction

Cognitive effects of depression

A depressed mind is characterised not only by its emotional state, but equally by its cognitive processes. Depression has been associated with a number of cognitive effects, ranging from deficits in inhibition and goal switching to perceptual and attentional biases. Many of these are tied to the occurrence of depressive rumination, a narrowly-focused form of mind-wandering that is a strong indicator for depression (Marchetti, Koster, Klinger & Alloy, 2016).

Smallwood and Schooler (2006) described mind-wandering as a “shift of attention away from a primary task toward internal information, such as memories.” (p. 946). In the case of rumination, this internal information is decidedly negative and focuses largely on the self, e.g. on conversations or arguments that have occurred in the past, or on one’s current dysphoric state (Marchetti et al., 2016). The difference between depressive rumination, which can be quite harmful, and more benign mind-wandering may lie in how thoughts are constrained in these states. Unlike normal mind-wandering, which is relatively unconstrained, rumination is subject to strong automatic constraints, making it more likely that ruminators perseverate in certain narrow tracks of thought from which it is difficult to escape (Christoff, Irving, Fox, Spreng & Andrews-Hanna, 2016; Harel, Tennyson, Fava & Bar, 2016).

The thought patterns that characterise rumination can be found both at short time scales and over longer stretches of time. For example, they are observed in the brief moments of spontaneous thought that occur during lapses in attention to an experimental task. By periodically probing participants’ thoughts as they performed a choice reaction time task, Hoffmann, Banzhaf, Kanske, Bermpohl and Singer (2016) found that participants with depression experienced off-task thoughts more frequently than healthy controls, and that these thoughts were more negative, past-oriented, and focused on the self. In a set of experience sampling studies, which asked participants to complete a small self-report questionnaire 10 times a day over a period of five to six days, depressed participants exhibited a similar pattern: they had a much stronger inclination towards reporting negative thoughts and feelings than their non-depressed counterparts (Wigman et al., 2015).

Those who engage in depressive rumination tend to have more difficulty

inhibiting information that is no longer relevant to their current task goal, compared to non-ruminators. This is most plainly visible in a task-switching paradigm: when the task set switches continuously, depressive rumination leads to weaker inhibition of previous task sets—something which can be beneficial when tasks are cycled, since switching costs will be lower (Whitmer & Banich, 2007). A comparable inhibitory deficit is found in long-term memory: by having participants memorise a list of learning items of which only a subset is subsequently rehearsed, Whitmer and Banich (2010) found that depressive ruminators inhibited unrehearsed items from the original list less strongly than non-ruminators.

Altamirano, Miyake and Whitmer (2010) put forward the trait-inflexibility view, saying that people who are prone to ruminate are characterised by a certain ‘mental inflexibility’, meaning that they emphasise goal maintenance over the ability to shift between goals. They found that, compared to non-ruminators, depressive ruminators were worse at shifting goals in a letter-naming task in which they had to switch between reading letters appearing in one location to reading letters in another location. At the same time, depressive ruminators were better at maintaining their goal in a modified Stroop task that—since it consisted largely of congruent trials—encouraged goal neglect.

Interestingly, some evidence suggests that this mental inflexibility that is associated with depression can largely disappear when a person is not actively engaged in rumination: preventing depressed individuals from ruminating by distracting them reduced their switching costs in a task-switching paradigm (Whitmer & Gotlib, 2012). Inducing rumination in non-depressed controls had no effect on switching cost, however, which implies that it is the combination of rumination and other cognitive factors that changes behaviour in depressed individuals.

The cognitive effects of depression also extend to the domain of perception and attention, where evidence suggests a negativity bias, mirroring the positivity bias that is often found in healthy humans. For example, Gollan, Pane, McCloskey and Coccato (2008) report that, when depressed participants were asked to judge the emotion of a neutral face in a forced-choice task, they saw its emotion as sad more often than happy, while non-depressed participants leaned towards judging the face as happy. This corresponds well with the finding by Surguladze et al. (2004) that depressed individuals exhibit a bias towards classifying faces of varying expressions as more negative than they are. In a review of the literature, Bourke, Douglas and Porter (2010) discuss several studies in which depressed individuals demonstrate selective attention towards sad faces and/or away from happy faces, although they note that the effect has not been widely replicated. When Suslow, Junghanns and Arolt (2001) asked depressed participants to detect deviant facial expressions in a grid of neutral faces, they responded more slowly than non-depressed controls when the deviant face was happy, but not when it was sad, suggesting that their attention to positive stimuli was relatively weak.

Modelling depression and rumination

Gaining a better understanding of how depression influences and interacts with cognition is likely to aid in the development of better diagnostic tools and ef-

fective treatments. Current schemes used to diagnose psychiatric disorders can be overly simplistic in their categorisation, and do not translate well to the clinical level (van Os et al., 1999). Theories of depression tend to be verbal, inferred from empirical observations. This can make them difficult to generalise or integrate with other theories. In addition, strictly verbal theories can afford themselves a degree of flexibility, which, while perhaps making it easier to accommodate new observations, makes these theories harder to falsify.

Taking a mechanistic approach that integrates the ‘what’ with the ‘how’ (i.e., Marr’s computational and algorithmic levels of analysis (Marr, 1982)) helps us create a more precise description of cognition in the depressed mind. As Ashby and Helie (2011) put it, “connections between brain regions do not change from task to task” (p. 276). Theories should therefore be no more flexible than necessary. In recent years, the field of computational psychiatry has demonstrated the benefits of striving towards more rigorous accounts of the role of cognition in mental illness. By implementing a theory in a computational model, we can generate testable predictions on the basis of well-defined computational processes that we suspect may underpin a certain behaviour, thereby improving our understanding of it. This brings with it obvious benefits for the successful diagnosis and treatment of not only depression, but also schizophrenia, anxiety, ADHD, and others (Adams, Huys & Roiser, 2016).

While some computational models focus only on a particular aspect of a task (e.g. perception), cognitive models specified within a cognitive architecture such as ACT-R (Anderson, 2007) or PRIMs (Taatgen, 2013) encompass the entire cognitive process, from an initial task goal to its completion. This makes them well-suited for investigating how changes to each part of a task contribute to changes in performance. In addition, mechanisms that explain differences in behaviour on one task can easily be transferred to a different task.

In a previous computational cognitive model of rumination, van Vugt, van der Velde and ESM-MERGE Investigators (2017) showed how increasing the prominence of negative items in the model’s memory could transform its behaviour from healthy mind-wandering into negatively-focused depressive rumination. When applied to a sustained attention task, the model predicted lower accuracy as a result of the changed quality of its off-task thinking. Following up on earlier models, this model by van Vugt et al. implemented mind-wandering as a process of continuous memory retrievals, competing for resources with the main task process (van Vugt, Taatgen, Sackur & Bastian, 2015; Taatgen, Katidioti, Borst & van Vugt, 2015).

Problem statement and thesis overview

The aim of this thesis is to identify which cognitive processes can account for the behavioural differences that are observed in depression, and to specify how they may do so. To this end, I present a computational cognitive model of performance on a working memory task, that was developed on the basis of existing theories described in this chapter. This model is evaluated against empirical observations of depressed and non-depressed participants.

Chapter 2 introduces the task that will be modelled: the emotional n -back task. Chapter 3 discusses the model. It explains how the task is modelled, and how a depressed version of the model differs from a non-depressed version. In

Chapter 4, the model's performance is compared to empirical data from Levens and Gotlib (2010). These results and their implications are discussed in Chapter 5.

Chapter 2

Emotional *n*-back

A task was selected that incorporates multiple aspects of cognition that are affected by depression. The emotional *n*-back task (Levens & Gotlib, 2010, 2015) involves perception and classification of emotional faces, requires continuous updating of task information in working memory, including the forgetting of information that is no longer relevant, and offers opportunity for rumination to occur. Levens and Gotlib (2010) had depressed and non-depressed participants perform this task. Importantly, they found that depression impaired performance in some, but not all conditions. Both groups were capable of doing the task to a high standard, which means that the differences that were observed between groups did not originate from a general passivity or lack of motivation in the depressed group (Scheurich et al., 2008).

The current chapter describes the emotional *n*-back task, and covers the most important findings of Levens and Gotlib (2010). The following chapter presents a cognitive model of the same task.

A working memory task with emotional stimuli

The emotional *n*-back is a variation on the regular *n*-back paradigm, which is a common method for assessing working memory performance (Owen, McMillan, Laird & Bullmore, 2005). The basic premise of an *n*-back task is that the participant perceives a series of stimuli (visual or auditory), and has to indicate for each stimulus whether it matches a stimulus shown *n* trials earlier. Normally, the stimuli are neutral; they can be letters, numbers, or other symbols that do not carry an emotional meaning.

In an emotional *n*-back task, however, the stimuli are designed to elicit an emotional response in the participant, with the aim of affecting their working memory performance. For example, stimuli can be emotionally charged words (Kopf, Dresler, Reicherts, Herrmann & Reif, 2013) or images of emotional scenes (Marx et al., 2011).

The version of the emotional *n*-back task that is the focus of this thesis uses photographs of faces displaying one of three facial expressions—happy, neutral, or sad—as stimuli (see Figure 2.1). In each trial a face is shown, upon which participants indicate whether its expression matches the expression of the face that was presented *n* trials ago.



Figure 2.1: Examples of the stimuli used by Levens and Gotlib (2010). All images are from the NimStim Face Set (Tottenham et al., 2009) and were presented in greyscale.

***n*-back strategies**

A difficulty of interpreting performance on the n -back task is that it is not necessarily clear how participants approach the task. In some cases, the strategy that participants use can be inferred from their behavioural results (Juvina & Taatgen, 2007), but a more common solution is to simply ask participants what their approach to the task is. Previous research has identified at least two common strategies that people use.

The first, labelled the *activation strategy* (Lovett, Daily & Reder, 2000) or the *low-control strategy* (Juvina & Taatgen, 2007), is best described as a ‘lazy’ approach to the task, since it does not involve any form of rehearsal of recently encountered items. Participants rely on a judgement of familiarity (“Does this item seem familiar?”) to decide whether an item matches the one shown n trials earlier. Juvina and Taatgen (2007) suggest that some basic estimation of recency (e.g. using ‘time tags’) may also contribute to this decision.

The second n -back strategy requires more conscious effort. Known as the *update strategy* (Lovett et al., 2000) or the *high-control strategy* (Juvina & Taatgen, 2007), this method involves actively maintaining a list of the n most recently experienced items in working memory. Participants using this strategy update their mental list every time a new stimulus appears: the new stimulus is added to the front of the list, while the oldest item is discarded, and every other item shifts back one place. Comparing the newest item and the oldest one

then yields the correct response. If time permits, the items in the list may be rehearsed through subvocalisation to keep them accessible.

Because the use of different strategies can be a confounding factor, Levens and Gotlib (2010) specifically instructed participants to only use the second, high-effort strategy.

Separating perceptual bias from WM differences

Assuming the use of a high-control strategy, the process of doing the n -back task can be divided in three distinct components. First, a stimulus that appears on screen has to be perceived and categorised. Here, this means that the participant sees a new face and recognises its expression. Then, the contents of the list in working memory are updated: a representation of the new face and its expression enters the list, while the oldest element is discarded. The item that was previously in the 1-back position now moves to the 2-back position. Following this, the new stimulus can be compared to the stimulus that was shown two trials ago, and a response is initiated.

There are two points where valence-specific biases can emerge: the initial perception and classification of a stimulus, and the subsequent working memory operations involved in processing the stimulus and the ones before it in the rehearsal list. One question that Levens and Gotlib (2010) set out to answer is, if depressed participants do indeed perform differently than healthy controls, should these differences be attributed to a perceptual and/or attentional bias, or to differences in working memory performance (or perhaps both)? As Chapter 1 discussed, previous studies have implicated both perception and working memory in the behavioural differences that emerge from depression.

To separate perception from working memory, Levens and Gotlib (2010) had subjects perform two versions of the task: a 0-back and a 2-back. This 0-back is largely identical to the 2-back, with the difference that participants compare the current stimulus to a target expression, specified at the start of a block, rather than to a previous stimulus. Since the 0-back task requires the same perceptual processes as the 2-back, but does not require participants to maintain a list of stimuli in working memory, it should isolate any bias in attending to and categorising the stimuli from differences in working memory processing. In other words, a difference in performance between depressed and non-depressed participants on the 0-back task would indicate an effect of perception on performance, whereas equal 0-back performance would suggest that any differences that emerge in the 2-back task are attributable to working memory biases.

Procedure

The 0-back task consisted of three blocks, each containing 43 trials. At the start of each block, the target expression for that block was shown. Figure 2.2 shows how trials were structured: a face stimulus appeared on screen for 2 seconds, during which participants could respond by pressing a key labelled “same” (if the face’s expression matched the target) or “diff” (if the face had a different expression than the target). There was a 0.5 s interval between stimulus presentations.

The 2-back task consisted of six blocks, each with 55 trials. Trials in this version of the task were identical to those in the 0-back (see Figure 2.2), except that participants had to compare a face's expression to that of a face shown two trials earlier rather than to a target expression.




Figure 2.2: Timeline showing two trials in the n -back task. Each face stimulus appeared for 2 seconds, during which participants could respond with a button press. The interstimulus interval was 0.5 seconds. In the 0-back condition, participants compared the face to a target expression given at the start of the block. In the 2-back condition, participants compared the face to the face shown two trials ago.

Conditions

The 2-back task requires different mental operations, depending both on the stimulus in the current trial and on the contents of previous trials. In total, we can distinguish four different conditions:

- **Match set:** $[\circ \bullet \circ \bullet]$ ¹ The current stimulus matches the 2-back stimulus.
- **Break set:** $[\bullet \circ \bullet \circ]$ The current stimulus does not match the 2-back stimulus, but the previous trial was a match set trial. That old set has to be 'broken' in order to discard the old 2-back item from the mental list.
- **No set:** $[\circ \circ \circ \circ]$ The current stimulus does not match the 2-back stimulus, nor was there a match in the previous trial.
- **Perseveration set:** $[\bullet \circ \bullet \bullet]$ The current stimulus does not match the 2-back stimulus, but the previous trial was a match set trial, and the valence of the current stimulus is the same as that of the old set.

¹The trial history is symbolised by a list of dots, with the oldest trial on the left and the current trial on the right; filled dots are stimuli that share the same expression.

Results summary

Levens and Gotlib (2010) recorded the response ('same', 'diff', or none) and response time on each trial. Chapter 4 provides a more in-depth look at their results, along with those of the cognitive model, but here a summary is given of similarities and differences between the two groups.

Response rate: In both the 0-back and the 2-back task, depressed participants responded less frequently than those in the control group.

Accuracy: There was no difference in response accuracy between the groups in either of the tasks.

Response time: There was no difference in response time between the groups in the 0-back task. In the 2-back task, depressed participants had slower responses in general. By applying a z-transformation to response times, Levens and Gotlib identified valence-specific differences within and between the groups. Depressed participants responded relatively more quickly than controls to a sad stimulus on screen, than to a happy one. Conversely, non-depressed participants were relatively fast when responding to a happy stimulus. Levens and Gotlib took this to mean that mood-congruent stimuli were easier to integrate into working memory (rather than the effect being due to a perceptual bias, since it would then also have shown up in the 0-back condition). Furthermore, depressed participants responded more slowly when the previous trial had contained a sad set. The authors proposed that this slowdown may be the result of having to break up the old 'set' of sad items, whereby related sad concepts and memories are activated before the oldest member of the set is successfully forgotten. The opposite effect was also observed: in trials during which a happy set had to be broken, participants in the control group were relatively slow to respond. Levens and Gotlib concluded that, while depressed participants seemed to have a tendency to linger on trials where the oldest member of a sad set was discarded, healthy controls must do the same when discarding a happy item.

Chapter 3

Model

Approach

In this chapter, I present two variants of a cognitive model that performs the emotional n -back task of Levens and Gotlib (2010): a ‘regular’ version, representing the control group in that study, and a ‘depressed’ version, representing the depressed participants. The aim of these models is to replicate performance of both groups on the 0-back task and the 2-back task of the original experiment, while offering a mechanistic explanation of observed differences that aligns with existing theoretical insights.

Modelling strategy and assumptions

Because of the large number of trials in the emotional n -back task, and its monotonous nature, it is likely that participants spent a significant portion of their time on off-task thinking. Indeed, the observation that mind-wandering occurs during working memory tasks, and that its intensity increases as time goes on, is a common one (Smallwood, Obonsawin & Heim, 2003; Krinsky, Forster, Llabre & Jha, 2017). Mean response rates of participants in Levens and Gotlib’s experiment ranged from 85% to 98%, supporting the idea that attention was not always directed to the task when a response had to be made. The observation that depressed participants had lower response rates than those in the control group furthermore suggests that depressed participants had a stronger tendency towards off-task thinking—a pattern which would match observations from other experiments involving depressed and non-depressed subjects (Hoffmann et al., 2016), and which is compatible with the idea that depressed participants are more likely to engage in a ‘stickier’ kind of mind-wandering that is harder to break out of (Christoff et al., 2016; Andrews-Hanna et al., 2013)¹. Continuing on previous modelling work (van Vugt et al., 2015; Taatgen et al., 2015), the task is therefore modelled as a competition between goals, each of which is vying for resources: one goal being task-related processing, the other to engage in mind-wandering.

¹An alternative explanation for the difference in response rate between groups could be that depressed participants simply found the task more difficult. There was, however, no indication that depressed participants performed any worse in terms of response accuracy, even in the most demanding conditions, which makes this interpretation unlikely.

Architecture

The model is constructed within the PRIMs cognitive architecture (Taatgen, 2013), an extension to the ACT-R architecture (Anderson, 2007). Like ACT-R, the PRIMs architecture provides a framework for modelling the full set of information processing steps that make up a cognitive task. The architecture consists of a number of modules (declarative memory, working memory, vision, goal state, motor action) that exchange information through buffers. A PRIMs model has one or more goals, each of which has a unique set of operators associated with it. PRIMs operators are comparable to production rules in ACT-R: like a production rule, an operator is styled as an if-then rule requiring certain conditions to be met before a sequence of primitive information processing steps is performed. For example, a goal might be to memorise a letter shown on screen. An operator associated with this goal may first check that there is a letter, read the letter, and then encode it in a new memory chunk in working memory.

While ACT-R makes a distinction between procedural knowledge (operators or productions) and declarative knowledge (memory chunks), PRIMs treats both as regular chunks stored in declarative memory. This means that each goal and operator has an activation value (a measure of memory trace strength) that is boosted by retrievals and decays exponentially over time. As a consequence, a PRIMs model is not constrained to a single goal, but can maintain multiple concurrent, competing goals. The advantage of this approach is that it allows for dynamics that the traditional ACT-R approach does not. Goals spread activation to their operators, which makes operators belonging to more active goals more likely to be selected for execution. This means that, as long as their conditions are met, a PRIMs model can flexibly alternate operators from different goals.

In PRIMs, activation can spread from any chunk to any other chunk, meaning that a chunk in working memory cannot only spread activation to similar chunks in declarative memory, but also to particular operators that it favours.

The additional flexibility offered by PRIMs allows our model to freely switch between operators corresponding to task goals and operators that implement off-task thought.

0-back model

Task strategy

The 0-back model has two concurrent goals, which constantly compete for resources: (i) to perform the task, and (ii) to engage in task-unrelated mind-wandering. Operators belonging to the first goal perform the steps needed to do the task. At the start of each block of trials the target expression (happy, neutral, or sad) is shown on screen. The model stores this information in a slot in its goal buffer. Upon seeing a new face stimulus, the model attempts to recognise the facial expression of the actor in the photograph. If it succeeds, the model responds with a key press: *same* when the expression matches the target, and *diff* when it does not. If the model does not recognise the facial expression, it will respond with a random key press. After responding, the model can keep its attention on the task until the onset of the next trial by subvocalising the target expression or simply by focusing on the task.

Stimuli

Because PRIMs’ model of visual perception and processing is quite simplistic, recognition of facial expressions is modelled as a simple memory retrieval. The stimuli are encoded in text form. Figure 3.1 shows how an image of a particular actor with a happy facial expression can be represented in text form by a chunk.

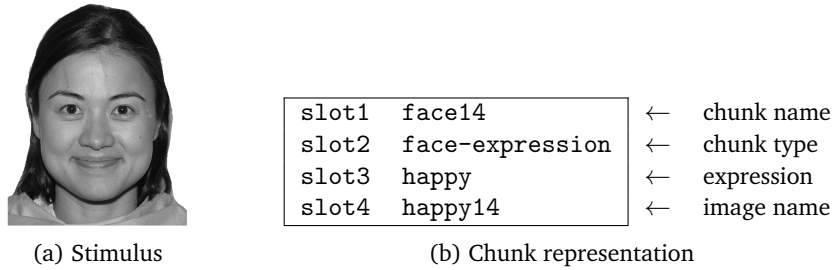


Figure 3.1: Example of a stimulus that was used by Levens and Gotlib (2010) (a), and its representation as a chunk in declarative memory as used in our model (b).

Chunks representing each face are placed in the model’s declarative memory before the task starts. When a stimulus is presented on screen, only the name of the image, stored in the last slot of its corresponding chunk, is used (e.g. happy14 in the example in Figure 3.1). This name serves as a cue for the model to retrieve the full chunk from memory, allowing it to characterise the facial expression.

Mind-wandering

Off-task thinking is modelled through a set of operators under the umbrella of the model’s second goal: mind-wandering. These operators allow the model to retrieve random memories (unrelated to the task) from declarative memory at any time during a trial. Each memory is stored in its own chunk in declarative memory. An example of a chunk encoding a memory is shown in Figure 3.2. Upon retrieving a memory, the model brings it to the forefront of attention by placing it in working memory. At this point, the model can retrieve another memory to prolong the mind-wandering episode. This generates a characteristic ‘thought pump’, in which memories follow one another.

There is no fixed end point for a mind-wandering episode, which means that, in theory, it can go on indefinitely. In practice, however, the stochasticity in the activation of operators means that inevitably an operator related to the task will get selected, at which point the mind-wandering episode ends.

slot1	memory4	←	chunk name
slot2	memory	←	chunk type
slot3	self	←	focus
slot4	memory4	←	memory name

Figure 3.2: Example of a memory chunk retrieved during mind-wandering.

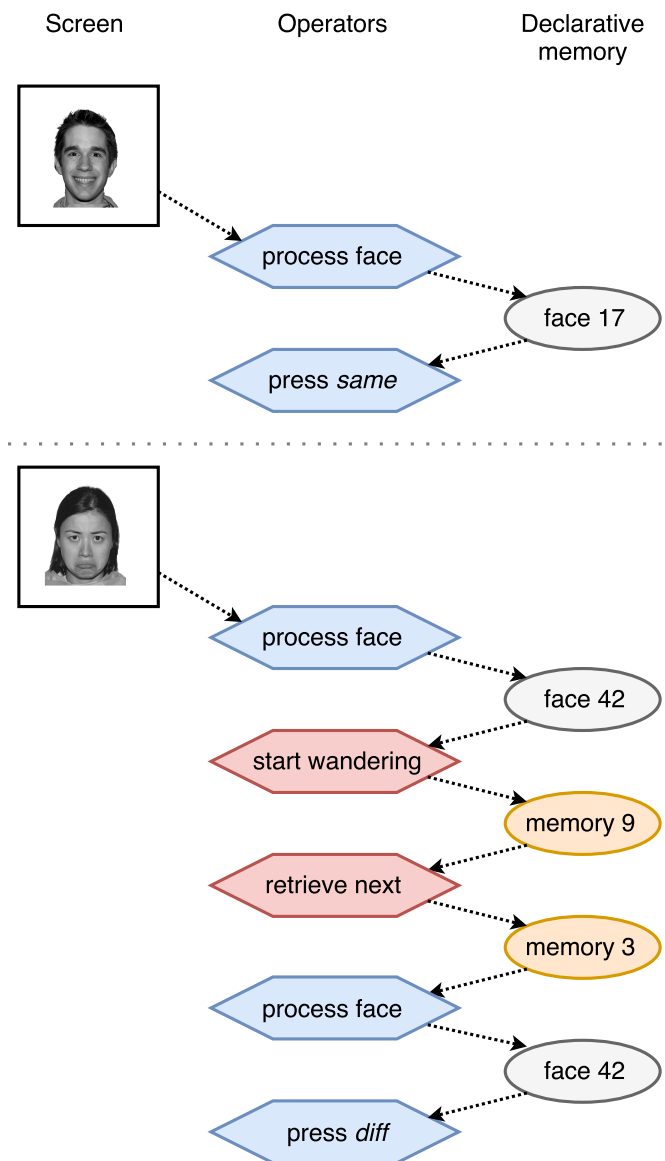


Figure 3.3: Goal interleaving in the 0-back model. Task operators are blue, while mind-wandering operators are red. Two trials are shown: one with only task-related operators, and one in which a series of mind-wandering operators intervene. Depending on the length of the mind-wandering episode, the model may or may not return to the task with enough time to respond.

Interleaving goals

Because of the way the model is structured, there is no set procedure for a trial. At each stage in a trial, operators from both goals can become active. For example, when a stimulus appears, the model can either process it to determine

its expression, or ignore it and start mind-wandering. The same concept applies to dealing with a face chunk that has been retrieved: the model either uses it to respond, or it ignores it. Similarly, in periods between trials, the model can either keep attending to the task, or it can engage in off-task thought. Note that the occurrence of mind-wandering in a trial does not mean that the model cannot engage in the task at all in that trial. As Figure 3.3 illustrates, after a period of mind-wandering the model may still return to the task when a task-related operator gets selected.

Emotion-specific effects

The behavioural data show that task performance was modulated by the emotional content of the stimuli. Specifically, participants were more accurate in their responses to happy faces than to sad and neutral faces, with neutral faces being especially difficult to recognise. In addition, participants responded more quickly to happy faces than to faces from either of the other two categories. These effects were present in both groups.

Partial matching of face chunks

To reproduce the effect of facial expression on response accuracy, the model's process of classifying a face is made to be imperfect. Retrieving a face chunk from memory uses partial matching, which means that on occasion the model will retrieve a chunk that does not perfectly match the face shown on screen, and as a result it will misclassify the facial expression (e.g., thinking a face is sad while it is in fact neutral). This imperfect behaviour mirrors that of humans, who are also prone to making this kind of mistake (Gollan et al., 2008).

To retrieve a face chunk, the stimulus name that is shown on screen (e.g. "happy14") is passed to the retrieval buffer as a query. Since we allow partial matching, all chunks in declarative memory are potential candidates for the subsequent retrieval. Which chunk ultimately gets selected for retrieval depends on activation. In the ACT-R and PRIMs architectures a chunk's activation value (A_i) is comprised of several elements: base-level activation (B_i), spreading activation received from other chunks (S_i), the chunk's mismatch to the retrieval request (M_i), and a random noise term (ϵ):

$$A_i = B_i + S_i + M_i + \epsilon$$

The less a chunk resembles the request, the larger its mismatch term will be, and as a result, the lower its activation.

The mismatch between two face chunks depends on the facial expression encoded in each of the chunks. For example, a sad face is very similar to a different sad face, and is also more similar to a neutral face than to a happy face. The degree to which one face differs from another is specified in a mismatch value. Mismatch values range from 0 (no difference at all, i.e., exactly the same face) to -1 (maximally different). Table 3.1 lists these mismatch values for all possible combinations of face chunks. These values were estimated on the basis of the empirical data and adjusted to provide a reasonable fit. Both versions of the model use the same values.

Table 3.1: Mismatch values of face chunks that do not match the retrieval request. A more negative value represents a greater mismatch.

Expression	Perceived as	Mismatch
happy	happy (different face ¹)	-0.05
	neutral	-0.1
	sad	-0.3
neutral	happy	-0.2
	neutral (different face)	-0.05
	sad	-0.2
sad	happy	-0.3
	neutral	-0.1
	sad (different face)	-0.05

¹A request for one happy face (e.g. *face14*) may return a different happy face (*face18*).

Activation of face chunks

The effect of facial expression on response time is captured in the model through differences in the activation of face chunks in memory. By manipulating the activation of a face chunk, we can make it easier or more difficult to retrieve, and as a result increase or decrease response times.

Following the standard ACT-R formula, the time it takes to retrieve a chunk i with activation A_i from memory is

$$t_{retrieval} = Fe^{-A_i}$$

(where F is the latency factor, a global constant). Thus, assuming sufficient activation for successful retrieval, chunks with a high activation are retrieved more quickly than chunks with a low activation. We can use this property to control how long it takes for the model to respond in trials with faces of a particular mood. In the model, the baselevel activations of face chunks is set individually for each mood. Table 3.2 shows the values that were selected on the basis of their fit.

Table 3.2: Baselevel activation values (in arbitrary units) for face chunks of each mood.

Expression	Activation
happy	1
neutral	0.4
sad	0.4

Differentiating a depressed model from a control model

When it comes to the 0-back task, the depressed participants in Levens and Gotlib (2010) are different from non-depressed controls in two respects: (i) their mean response time was higher, and (ii) their mean response rate was

lower. As Table 3.3 shows, the depressed version of the 0-back model also differs from the control model in two ways to capture these differences. It reproduces the depressed participants' higher response times by taking longer to execute a keypress than the control model. This is consistent with findings of general psychomotor slowing in depressed patients (Schrijvers, Hulstijn & Sabbe, 2008). A lower response rate is achieved by increasing the activation of mind-wandering operators, making the overall proportion of mind-wandering during the task slightly higher and as a consequence reducing the number of responses. This reflects the higher degree of rumination found in depressed patients (Christoff et al., 2016; Andrews-Hanna et al., 2013).

Table 3.3: Parameter differences between the control and depressed versions of the 0-back model. All other parameters kept constant (see Appendix A for a detailed overview of parameter settings).

Parameter	Control	Depressed
Keypress duration	0.4 s	0.5 s
Activation of mind-wandering operators	2	2.25

2-back model

Changes relative to the 0-back model

The 2-back model works the same as the 0-back model when it comes to perceiving and classifying stimuli, making responses, and engaging in off-task thinking. Where the two models differ is in the process of determining the correct response. The 2-back model uses *n-back* chunks to keep track of previous observations. An example of such a chunk is given in Figure 3.4.

slot1	imaginalN42	←	chunk name
slot2	face-observation	←	chunk type
slot3	two-back	←	place in mental list
slot4	happy	←	facial expression
slot5	context12	←	context in which observation was made

Figure 3.4: Example of an *n-back* chunk that is used as part of the model's mental list of recent stimuli. This chunk encodes the fact that the model saw a happy face two trials ago.

Because the response time effects are thought to arise from differences in how quickly the mental list of stimuli is updated under particular circumstances, it is important that the model should capture this updating process in detail. Figure 3.5 gives an overview of the model's task strategy. The mental list of the most recent stimuli is stored in the model's declarative memory, and is updated in every trial. At the start of a trial, the model retrieves from its memory each of the *n-back* observations in turn. Upon retrieving an *n-back* chunk, it creates an updated version in its imaginal buffer that changes its index in the list from *n* to *n* + 1. In the case of the old 2-back chunk, its index is changed to "old",

reflecting the fact that it is no longer part of the list. The updated chunk is pushed back to declarative memory. Once the items on the mental list have been shifted back by one place, the new stimulus that is currently on screen can be added. The model determines its expression in the same way as before, and creates a new n -back chunk that places the observation at the front of the list.

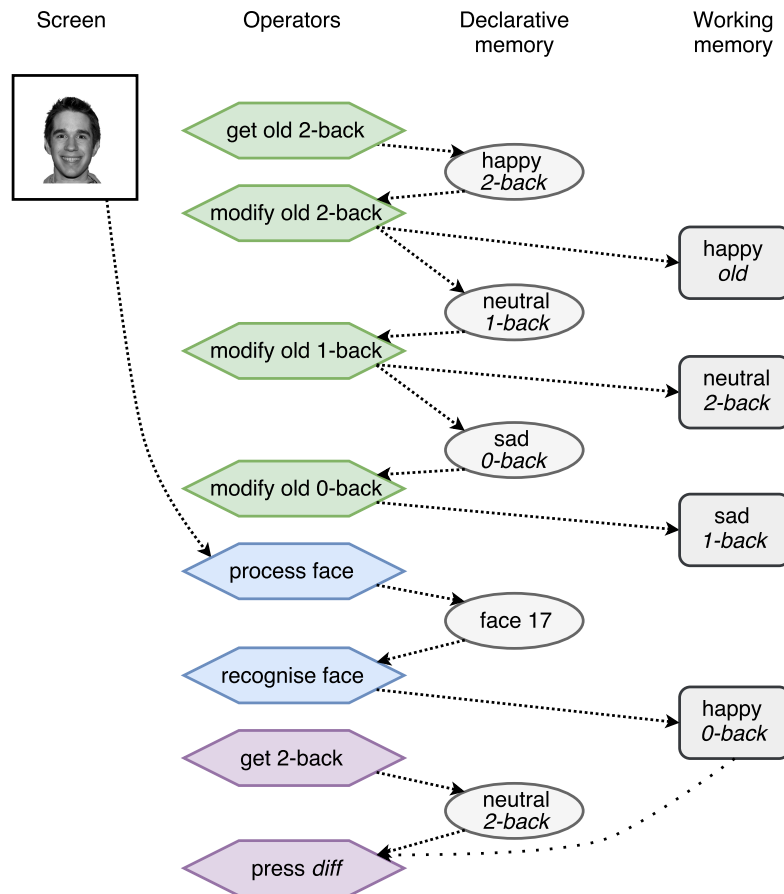


Figure 3.5: Task strategy of the 2-back model. Operators that update the old items on the mental list are green, those that process the new stimulus are blue, and the operators that make the response are purple.

The valence-specific differences in response time that were found in the 2-back task imply that the time it takes to integrate new observations into working memory depends on their valence, and differs between depressed and non-depressed participants. Levens and Gotlib (2010) observed that depressed participants needed less time to integrate a stimulus into their mental list when it was congruent with their mood (i.e., a sad face was integrated more quickly than a happy or neutral face). This was a surprising result, given the lack of any effect of emotion on response time in the 0-back task. Indeed, other studies (including a later one by the same authors using the same task) have shown (formerly) depressed participants to have a negative perceptual bias (Levens

& Gotlib, 2015; Gollan et al., 2008; Bourke et al., 2010), which contradicts what was found here. For this reason, our 2-back model goes along with the consensus view, by implementing such a mood-congruent perceptual bias. This bias was left out of the 0-back model for the sake of reproducing the original data better, but one could argue that both models should include the bias.

The time needed for a stimulus leaving the mental list was also found to depend on its valence, such that mood-congruent stimuli were discarded more slowly (i.e., a depressed participant needed more time to dismiss a sad face than a happy or neutral face, whereas a control participant spent longer on trials in which they had to forget a happy face). It is not certain why some items are more easily purged from working memory than others. Levens and Gotlib (2010) put forth the hypothesis that participants might struggle when breaking up an old mood-congruent ‘set’ of matching stimuli in their mental list, and that this causes them to slow down. While it would account for the observed response time patterns, this explanation is difficult to operationalise, since it is not clear how a ‘set’ would be represented in memory (whether as a single unit, or as two units with a strong connection), nor is it clear what it would mean in practical terms to ‘break’ a set. For this reason, a simpler mechanism is pursued here. Whenever an observation that is pushed out of the mental list is congruent with the model’s mood, there is an opportunity for the model to elaborate on the item before responding, causing a slowdown. Figure 3.6 shows an example of the depressed model lingering on a sad observation.

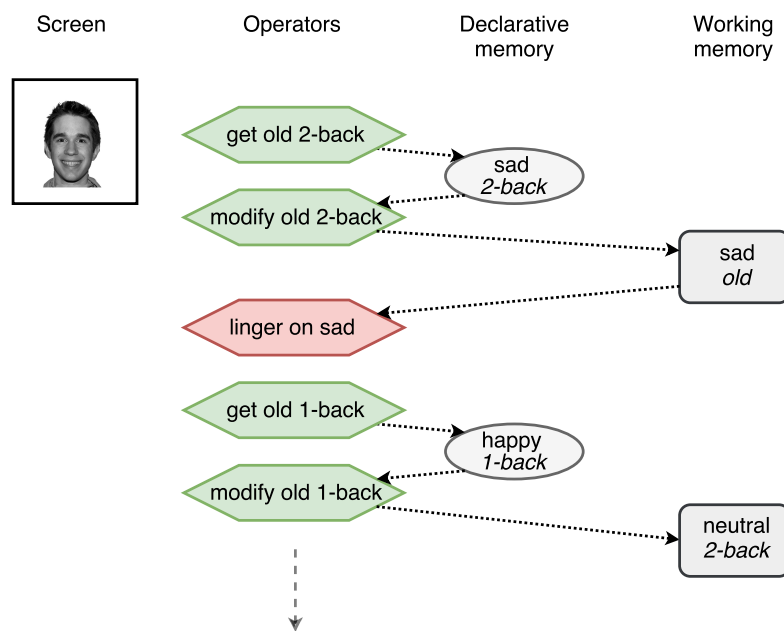


Figure 3.6: Depressed 2-back model elaborating on a mood-congruent item as it is pushed off the mental list, slowing down its response.

Differentiating a depressed model from a control model

The depressed model should perform differently than the control model on several fronts: (a) the overall response time should be higher, (b) the overall response rate should be lower, and (c) there should be valence-specific biases in working memory that are observable in the response time. As in the 0-back model, the higher overall response times in the depressed group are assumed to be the result of psychomotor slowing, which is commonly found in depressed patients (e.g., Schrijvers et al., 2008). Also, as before, the lower response rate is assumed to result from increased mind-wandering in the depressed group. This is encouraged through a higher activation of mind-wandering operators. Finally, the valence-specific biases are implemented through quicker perception of certain faces and elaboration on mood-congruent items as they leave the rehearsal loop. Table 3.4 summarises the differences between the two versions of the 2-back model.

Table 3.4: Parameter differences between the control and depressed versions of the 2-back model. All other parameters kept constant (see Appendix A for a detailed overview of parameter settings).

Parameter	Control	Depressed
Keypress duration	0.35 s	0.5 s
Activation of mind-wandering operators	1.5	1.6
Activation of happy/neutral/sad face chunks	1/0/0	1/0/1.5
Elaboration on discarded items	happy	sad

Quicker perception of sad faces

The depressed version of the 2-back model perceives sad faces more quickly than the nondepressed model. Within the current PRIMs architecture it is not possible to vary, based on the contents of a chunk, how long it takes to create it in working memory, or push it to declarative memory. This means that we can only affect the speed of integration by giving the model a perceptual bias: changing how long it takes to recognise a face’s expression by manipulating the duration of the required retrieval. As described previously, the time required to retrieve a chunk from declarative memory depends on its activation, such that a higher activation leads to a faster retrieval, following the formula: $t_{retrieval} = F e^{-Ai}$. By increasing the activation of sad face chunks from 0 to 1.5, we can reduce the average time needed to retrieve them by a factor of $\frac{e^{-0}}{e^{-1.5}} \approx 4.5$.

Elaborating on mood-congruent items

Both models have the ability to elaborate on old 2-back stimuli as they are removed from the mental list. The valence of the discarded item is important: the depressed model will only elaborate on sad items, while the control model only elaborates on happy items. The time penalty that is incurred by lingering on an old item and subsequently having to refocus on updating the mental list slows down the model’s response in that trial.

Chapter 4

Results

Analysis

Performance of the cognitive models is compared to that of the depressed and healthy participants in Levens and Gotlib (2010). Response rate, accuracy, and response time were recorded in each trial. Following Levens and Gotlib, we removed trials with response times outside 2.5 SD of a model participant's mean response time. Only correct trials were used for calculating mean response time. We analysed the effects of group (*depressed, control*), emotion (*happy, neutral, sad*), and response (*same, diff*) on response rate, accuracy, and response time using repeated measures analysis of variance (RM-ANOVA). Degrees of freedom were adjusted using Greenhouse-Geisser correction wherever Mauchly's test found sphericity to be violated. Effect sizes were estimated using generalised eta-squared statistics. Where theoretically motivated, significant main and interaction effects were further explored through pairwise t-tests with Bonferroni correction of p-values¹.

0-back task

In the 0-back task, participants decided whether a face stimulus had the same facial expression as a target shown at the start of each block. Faces were happy, neutral, or sad. Response times were measured from stimulus onset to the point at which an answer key ("same" or "diff") was pressed. This task only involves the perceptual component of the *n*-back paradigm (perceiving and categorising the stimulus), and leaves out the working memory manipulations that are part of the 2-back task. As such, it allows us to isolate differences in perceptual bias between depressed individuals and healthy controls.

The depressed and control version of the 0-back model were both run 50 times, simulating 50 participants in each group. The models went through 9 training runs (each run containing 3 blocks of 43 trials) to learn the task, before performing a final run. Only this last run was used in the analysis. Table 4.1

¹To maintain a degree of readability, not all effects, interactions, and pairwise comparisons of means are included in the text. Please refer to the online version of the analysis for an exhaustive report: <https://github.com/maartenvandervelde/emotional-n-back/tree/master/analysis>

shows the behavioural fit of the models to the human data. Table 4.2 summarises which of the effects that were observed in the human data are reproduced by our model. The following sections take a more in-depth look at how the 0-back models stack up against depressed and nondepressed humans.

Table 4.1: Behavioural fit of the 0-back model simulations to the human data. RMSE = root mean square error.

Measurement	R^2	RMSE
Accuracy	0.01	0.11
Response rate	0.38	0.04
z-transformed RT	0.71	0.50

Table 4.2: Which 0-back effects in humans were reproduced by the model?

Outcome variable	Effect	Reproduced
Accuracy	<i>Accuracy control</i> = <i>Accuracy depressed</i>	✓
	<i>Accuracy diff</i> > <i>Accuracy same</i>	✓
	<i>Accuracy happy</i> > <i>Accuracy neutral</i>	✓
	<i>Accuracy happy</i> > <i>Accuracy sad</i>	✓
	<i>Accuracy sad</i> > <i>Accuracy neutral</i>	—
Response rate	<i>RR control</i> > <i>RR depressed</i>	✓
	<i>RR diff</i> > <i>RR same</i>	✓
	<i>RR happy</i> > <i>RR sad</i>	—
Response time	<i>RT neutral, same</i> > <i>RT neutral, diff</i>	✓
	<i>RT sad, diff</i> > <i>RT sad, same</i>	—

Accuracy

Figure 4.1 shows the response accuracy by expected response for each of the three stimulus types. In Levens and Gotlib (2010), depressed and nondepressed participants were equally accurate. In addition, both groups showed the same effect of facial expression on response accuracy. Accuracy was lowest when responding to a neutral stimulus, and highest when the stimulus was a happy face. Responses to sad stimuli fell somewhere in the middle in terms of accuracy. The high degree of accuracy in trials with happy stimuli demonstrates that, given a face that is easily categorised, participants in both groups were near-flawless at comparing it to the target and responding appropriately. The drop in accuracy with sad stimuli and particularly with neutral stimuli suggests that it is more difficult to categorise these facial expressions. Furthermore, participants in both groups achieved higher accuracy in trials where the expected response was *different* than in trials where *same* was the correct response.

Our model reproduces all but one of these observations. A three-way RM-ANOVA conducted on the model data shows a significant effect of facial ex-

pression ($F(1.74, 170.21) = 115.43, p < 0.0001, \eta^2 = 0.35$) and expected response ($F(1, 98) = 370.96, p < 0.0001, \eta^2 = 0.51$), as well as their interaction ($F(1.74, 170.21) = 115.43, p < 0.0001, \eta^2 = 0.35$), on accuracy. No other effects were significant. Post-hoc comparisons of means confirmed that both models responded more accurately to happy faces than to neutral faces ($t(196) = 13.999, p < 0.0001$) and sad faces ($t(196) = 12.116, p < 0.0001$). While human participants had more difficulty with neutral faces than with sad faces, the models showed no such effect. The main effect of expected response is due to the models responding more accurately in trials in which the stimulus differs from the target ($t(98) = 19.26, p < 0.0001$); a pattern which was also found in the human data.

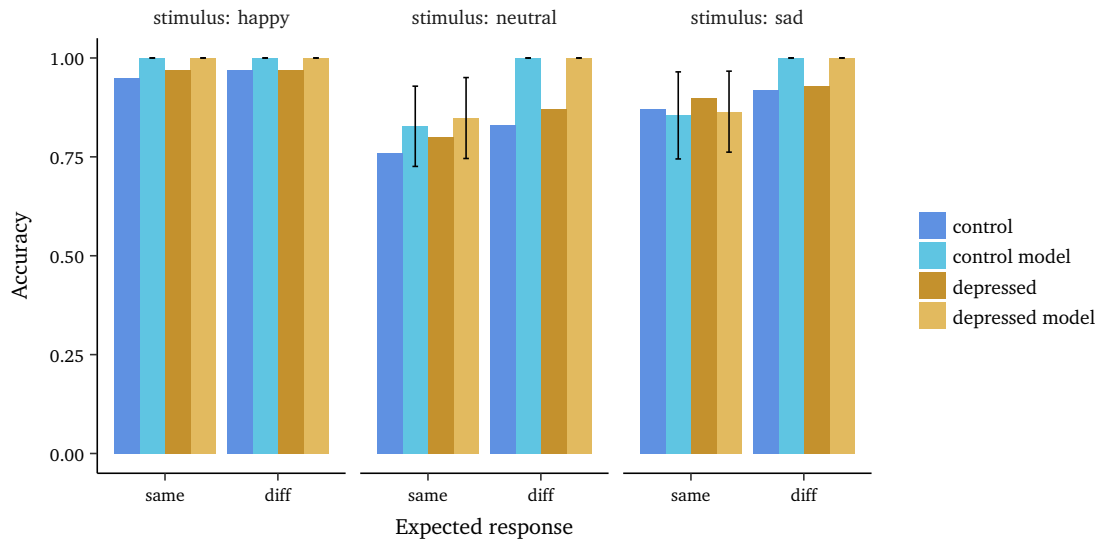


Figure 4.1: Response accuracy of real and modelled depressed and control participants on each condition of the 0-back task. Error bars indicate 1 SD from the mean. Non-response trials are excluded. Human data from Levens and Gotlib (2010).

Response rate

The mean response rate by expected response for each type of stimulus is presented in Figure 4.2. As in the human data, the models' response rates were universally high. However, unlike accuracy, response rates did differ between depressed participants and controls: on average, those with depression responded on about 3% fewer trials. The same is true for the models: a three-way RM-ANOVA indicates a main effect of group on response rate ($F(1, 98) = 81.45, p < 0.0001, \eta^2 = 0.13$). A follow-up test shows that the control model indeed had a higher response rate than the depressed model ($t(98) = 9.025, p < 0.0001$). In addition to group, there is also an effect of the expected response on response rate ($F(1, 98) = 3.80, p = 0.05, \eta^2 = 0.006$): the models tended to respond more

frequently when the expected response was *diff*, although this effect is only marginally significant in a pairwise comparison of means ($t(98) = 1.949, p = 0.0542$). The human data also contained an effect of facial expression: happy faces elicited a higher response rate than sad faces. We do not observe such an effect in the model.




Figure 4.2: Response rate of real and modelled depressed and control participants on each condition of the 0-back task. Error bars indicate 1 SD from the mean. Human data from Levens and Gotlib (2010).

Response time

Figure 4.3 shows mean response times by expected response for each of the three stimulus types. The human data showed a strong effect of group: depressed participants had higher response times across conditions. This effect is also present in the model results. The three-way RM-ANOVA reveals a significant effect of group on RT ($F(1, 98) = 7723.35, p < 0.0001, \eta^2 = 0.92$). As expected, this is because response times of the depressed model are higher ($t(98) = 87.883, p < 0.0001$).

To allow for a direct comparison between groups, despite the observed differences in speed of responding, Levens and Gotlib (2010) z-transformed response times (for every participant, each condition mean RT was subtracted from their overall mean RT and divided by the SD of the condition RT). These z-transformed response times are shown in Figure 4.4. Importantly, they show that when correcting for participants' baseline speed, there is no longer a difference in relative speed between the two groups. Both groups of human participants responded more quickly to happy faces than to sad or neutral ones. The same patterns are present in our model data. A three-way RM-ANOVA confirms that there is no effect of group on z-RTs, while there is an effect of facial




Figure 4.3: Response times of real and modelled depressed and control participants on each condition of the 0-back task. Error bars indicate 1 SD from the mean. Trials with response times deviating more than 2.5 SD from the mean are excluded. Human data from Levens and Gotlib (2010).

expression ($F(1.99, 195.54) = 1694.29, p < 0.0001, \eta^2 = 0.86$). As expected, responses to happy faces were significantly faster than those to neutral faces ($t(196) = 50.826, p < 0.0001$) and sad faces ($t(196) = 49.989, p < 0.0001$). There is also a significant main effect of expected response on z-RT ($F(1, 98) = 714.68, p < 0.0001, \eta^2 = 0.55$): relative response speed was faster on trials in which *diff* was the expected response ($t(98) = 26.734, p < 0.0001$). Finally, the interaction between facial expression and expected response has a significant effect on z-RT ($F(1.91, 187.48) = 21.83, p < 0.0001, \eta^2 = 0.08$). Since all but two of the 15 possible contrasts describe significant differences, we will focus on those comparisons that are motivated by the human data. Levens and Gotlib (2010) found humans to respond more slowly on trials with neutral faces when the correct response was *same* than when it was *diff*. The same is true for our models ($t(292.53) = 17.793, p < 0.0001$). In addition, humans responded more quickly on sad *same* trials than on sad *diff* trials, whereas the opposite is true for our model ($t(292.53) = 16.997, p < 0.0001$).




Figure 4.4: Z-transformed response times of real and modelled depressed and control participants on each condition of the 0-back task. Error bars indicate 1 SD from the mean. Trials with response times deviating more than 2.5 SD from the mean are excluded. Human data from Levens and Gotlib (2010).

2-back

In the 2-back task participants decided whether a face stimulus had the same facial expression as the face shown two trials before. Faces were happy, neutral, or sad. Response times were measured from stimulus onset to the point at which an answer key (“same” or “diff”) was pressed. Along with the perceptual component of the n -back paradigm (perceiving and categorising the stimulus), this version of the task also involves the working memory component (maintaining a representation of recently seen stimuli). In combination with the 0-back version of the task, this allows us to investigate differences in working memory performance between depressed and healthy participants.

Both the depressed and control version of the 2-back model were run 50 times, simulating 50 model participants in each group. As before, the models were trained on the task (6 blocks of 55 trials) for 3 full runs before they performed a final run. Only this final run is analysed. The behavioural fit of the models to the human data is summarised in Table 4.3, while Table 4.4 summarises how well the models capture each of the effects observed in humans by Levens and Gotlib (2010).

The 2-back task requires different mental operations, depending both on the stimulus in the current trial and on the previous trials. Because of this, the results for each of the conditions are analysed separately. The meaning of the condition names is explained on page 11.

Table 4.3: Behavioural fit of the 2-back model simulations to the human data.
RMSE = root mean square error.

Measurement	R^2	RMSE
Accuracy	0.01	0.11
Response rate	0.54	0.06
z-transformed RT	0.10	0.67

Table 4.4: Which 2-back effects in humans were reproduced by the model?

Outcome variable	Effect	Reproduced
Accuracy	<i>Accuracy control = Accuracy depressed</i>	—
Response rate	<i>RR control > RR depressed</i>	✓
	<i>RR perseveration, happy > RR perseveration, sad</i>	—
	<i>RR perseveration, happy > RR perseveration, neutral</i>	✓
Response time	<i>RT depressed > RT control</i>	✓
	<i>RT match, happy > RT match, neutral</i>	✓
	<i>RT match, happy > RT match, sad</i>	✓
	<i>RT depressed, break sad > RT depressed, break neutral</i>	✓
	<i>RT depressed, break sad > RT depressed, break happy</i>	✓
	<i>RT control, break happy > RT control, break neutral</i>	✓
	<i>RT control, break happy > RT control, break sad</i>	✓
	<i>RT control, break happy > RT depressed, break happy</i>	✓
	<i>RT control, break neutral = RT depressed, break neutral</i>	✓
	<i>RT depressed, break sad > RT control, break sad</i>	✓
	<i>RT control, no set neutral > RT control, no set happy</i>	✓
	<i>RT control, no set sad > RT control, no set happy</i>	✓
	<i>RT depressed, no set neutral > RT depressed, no set happy</i>	✓
	<i>RT depressed, no set neutral > RT depressed, no set sad</i>	✓
	<i>RT control, no set happy = RT depressed, no set happy</i>	—
	<i>RT control, no set neutral = RT depressed, no set neutral</i>	—
	<i>RT control, no set sad > RT depressed, no set sad</i>	✓

Accuracy

The human participants in Levens and Gotlib (2010) did not differ significantly in their accuracy on the 2-back task. Accuracy was generally above 80%. As Figure 4.5 shows, our models are unable to replicate this result exactly. Although neither model is perfectly accurate, both perform slightly better on the task than they should; most notably the control model. Contrary to the human data, mean accuracy of the control model appears to be higher than that of the depressed model.

Indeed, the RM-ANOVA reveals that the model's accuracy in match set trials

is predicted by group ($F(1, 98) = 231.83, p < 0.0001, \eta^2 = 0.51$). A follow-up test confirms that the control model is more accurate than the depressed model ($t(98) = 15.226, p < 0.0001$). No other effects are significant. The same is true for break set, no set, and perseveration set trials. In all cases, group membership has a significant effect on response accuracy (break set: $F(1, 98) = 32.82, p < 0.0001, \eta^2 = 0.10$; no set: $F(1, 98) = 202.94, p < 0.0001, \eta^2 = 0.38$; perseveration set $F(1, 98) = 64.38, p < 0.0001, \eta^2 = 0.18$). The control model is always more accurate than the depressed model (break set: $t(98) = 5.729, p < 0.0001$; no set: $t(98) = 14.246, p < 0.0001$; perseveration set: $t(98) = 8.024, p < 0.0001$).




Figure 4.5: Response accuracy of real and modelled depressed and control participants on each condition of the 2-back task. Error bars indicate 1 SD from the mean. Non-response trials are excluded. Human data from Levens and Gotlib (2010).

Response rate

Levens and Gotlib (2010) found a significant effect of group membership on response rate across conditions: participants with depression responded to significantly fewer trials than nondepressed controls. Figure 4.6 suggests that our models reproduce this effect quite well. The models' mean response rates appear to match those of the human participants.

A RM-ANOVA finds that in match set trials, response rate is indeed modulated by group membership ($F(1, 98) = 350.90, p < 0.0001, \eta^2 = 0.65$). A post-hoc comparison of means confirms that the control model responded to significantly more trials than the depressed model ($t(98) = 18.732, p < 0.0001$). No other effects were significant.

The effect of group on response rate is also present in break set trials ($F(1, 98) = 199.64, p < 0.0001, \eta^2 = 0.39$). In addition, the valence of the set that is to be broken has a significant effect on response rate ($F(1.82, 178.81) = 46.95, p < 0.0001, \eta^2 = 0.25$), and so does the interaction between group and valence ($F(1.82, 178.81) = 79.22, p < 0.0001, \eta^2 = 0.36$). Follow-up tests reveal that the control model had a higher response rate than the depressed model ($t(98) = 14.129, p < 0.0001$), and that the main effect of valence is due to lower response rates when the broken set was sad than when it was happy ($t(196) = 7.109, p < 0.0001$) or neutral ($t(196) = 9.257, p < 0.0001$). Further investigation of the interaction between group and valence yields significant differences in 12 of the 15 possible comparisons. These are omitted here.

As in the other conditions, group membership was a significant predictor of response rate in the no set condition ($F(1, 98) = 374.81, p < 0.0001, \eta^2 = 0.63$), owing to the fact that the control model responded more frequently than the depressed model ($t(98) = 19.36, p < 0.0001$). No other effects were significant in this condition.

Finally, response rate in the perseveration set condition was predicted by group membership ($F(1, 98) = 151.05, p < 0.0001, \eta^2 = 0.32$), the valence of the current stimulus (and therefore of the set in the previous trial; $F(1.59, 155.69) = 37.13, p < 0.0001, \eta^2 = 0.21$), as well as the interaction between these two factors ($F(1.59, 155.69) = 46.12, p < 0.0001, \eta^2 = 0.25$). As before, the effect of group is due to higher response rates in the control model ($t(98) = 12.29, p < 0.0001$). The effect of valence comes from lower response rates in trials that perseverate on a sad set compared to trials that perseverate on a happy set ($t(196) = 7.389, p < 0.0001$) or a neutral set ($t(196) = 7.535, p < 0.0001$). The interaction between the two main effects yields significant differences in 9 of the 15 possible comparisons. For this reason, it is not explored further here.

In summary, the depressed model's response rate was found to be lower than that of the control model in all four conditions, which matches the empirical observations of Levens and Gotlib (2010). In conditions where a previous set had to be broken (break set and perseveration set) both models responded less frequently when this set was sad than when it was happy or neutral. Levens and Gotlib (2010) reported a similar effect, albeit only in the perseveration set condition: response rate was higher in happy perseveration set trials than in neutral or sad ones. Both cases are compatible with the notion that a happy set may be easier to 'break up and forget' than a sad one.

Response time

Response times of both models are shown alongside human response times in Figure 4.7. There are several effects in the human data that we are interested in replicating. Firstly, depressed participants responded more slowly than nondepressed controls in all conditions, an effect that is likely due to psychomotor slowing. In addition, there are some condition-specific effects. In the match set condition, both groups responded more quickly when the stimulus was a happy one. In the break set condition, depressed participants took longer to respond when breaking up a sad set than they did when breaking happy or neutral sets, whereas nondepressed control participants had the opposite tendency, responding more slowly when breaking a happy set than when breaking a neutral or sad set. This effect can be explained through a tendency to linger on

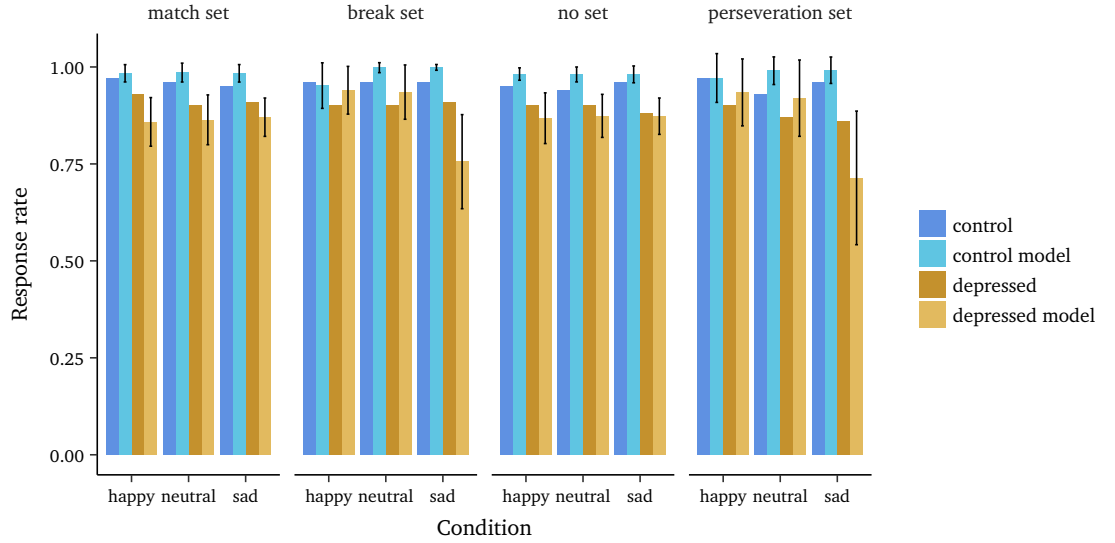


Figure 4.6: Response rate of real and modelled depressed and control participants on each condition of the 2-back task. Error bars indicate 1 SD from the mean. Human data from Levens and Gotlib (2010).

old mood-congruent items as they are discarded. Finally, depressed participants had relatively fast responses to sad stimuli in the no set condition, compared to happy or neutral stimuli, which may reflect facilitated processing of sad stimuli.

RM-ANOVAs confirm that the depressed model's response times are indeed slower in each of the four conditions (match set: $F(1, 98) = 2244.53, p < 0.0001, \eta^2 = 0.91$; $t(98) = 47.377, p < 0.0001$; break set: $F(1, 98) = 1319.48, p < 0.0001, \eta^2 = 0.85$; $t(98) = 36.325, p < 0.0001$; no set: $F(1, 98) = 3083.83, p < 0.0001, \eta^2 = 0.93$; $t(98) = -55.532, p < 0.0001$; perseveration set: $F(1, 98) = 570.67, p < 0.0001, \eta^2 = 0.68$; $t(98) = 23.889, p < 0.0001$).

The RM-ANOVAs reveal a large effect size of group membership. Since the relatively large difference in response time between groups is likely to obscure the other effects, response times are once again z-transformed to investigate the other effects. Figure 4.8 depicts these transformed RTs. It allows us to see the relative speed of responding, taking into account overall differences in response speed.

A RM-ANOVA conducted on match set trials yields significant effects of group ($F(1, 98) = 2366.71, p < 0.0001, \eta^2 = 0.90$), stimulus valence ($F(1.97, 192.68) = 187.26, p < 0.0001, \eta^2 = 0.54$), and the interaction between group and valence ($F(1.97, 192.68) = 79.07, p < 0.0001, \eta^2 = 0.33$) on z-transformed RT. Most interestingly, the effect of valence is due to both models responding relatively more quickly to happy trials than to neutral trials ($t(196) = 19.326, p < 0.0001$) or sad trials ($t(196) = 8.784, p < 0.0001$), in addition to responding more quickly to sad trials than to neutral trials ($t(196) = 10.542, p < 0.0001$). This result matches the pattern in the human data, where both groups had the most rapid responses to happy match set trials.

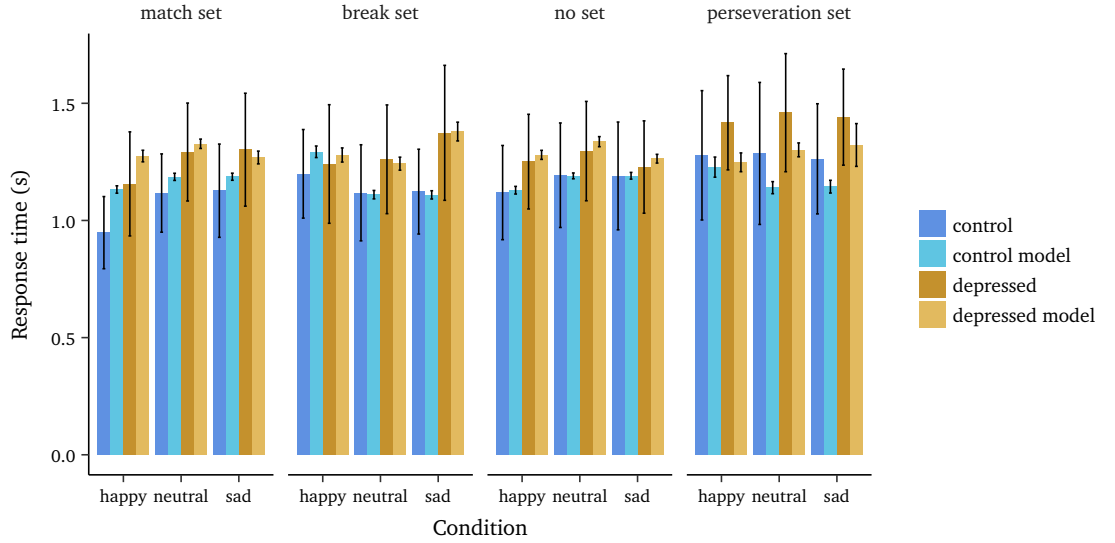


Figure 4.7: Response times of real and modelled depressed and control participants on each condition of the 2-back task. Error bars indicate 1 SD from the mean. Trials with response times deviating more than 2.5 SD from the mean are excluded. Human data from Levens and Gotlib (2010).

In break set trials, the valence of the set that is broken has a significant effect on z-RT ($F(1.77, 173.82) = 294.65, p < 0.0001, \eta^2 = 0.65$), which is qualified by an interaction effect of group and valence ($F(1.77, 173.82) = 398.05, p < 0.0001, \eta^2 = 0.72$). 12 out of 15 post-hoc comparisons of means yield a significant difference, but here we will focus on the comparisons that are relevant to our modelling goals. As expected, we observe a mood-congruent slowing effect in both groups. The control model responded more slowly in trials in which a happy set was broken than in trials in which it broke a neutral set ($t(196) = 28.884, p < 0.0001$) or a sad set ($t(196) = 28.850, p < 0.0001$). Similarly, the depressed model made slower responses when it broke a sad set than when it broke a happy set ($t(196) = 10.834, p < 0.0001$) or a neutral set ($t(196) = 16.267, p < 0.0001$). In addition, in trials in which a sad set was broken, the control model responded more quickly than the depressed model ($t(291.25) = 18.561, p < 0.0001$), whereas the depressed model was quicker to respond than the control model when a happy set was broken ($t(291.25) = 19.736, p < 0.0001$). The two models did not differ in their response time on trials in which they disengaged from a neutral set ($t(291.25) = 2.896, p = 0.0609$). These findings match the mood-congruent slowing effects found by Levens and Gotlib (2010).

In no set trials, a RM-ANOVA yields a significant effect of group ($F(1, 98) = 3145.81, p < 0.0001, \eta^2 = 0.93$), stimulus valence ($F(2, 195.66) = 322.50, p < 0.0001, \eta^2 = 0.66$), and the interaction between group and valence ($F(2, 195.66) = 175.30, p < 0.0001, \eta^2 = 0.52$). Follow-up comparisons of means reveal significant differences for 14 out of 15 possible comparisons. Here, too, we are primar-

ily interested in comparisons that are relevant to the effects we are modelling. In line with our expectations, the depressed model responded more slowly to trials with neutral stimuli than to trials with sad stimuli ($t(196) = 22.024, p < 0.0001$) or happy stimuli ($t(196) = 17.068, p < 0.0001$). In addition, the control model was faster in its responses to happy stimuli than to neutral stimuli ($t(196) = 18.538, p < 0.0001$) or sad stimuli ($t(196) = 18.675, p < 0.0001$). The depressed model was quicker at responding to sad stimuli than the control model ($t(288.35) = 20.953, p < 0.0001$). These results align perfectly with our expectations based on the human data. Two results are unexpected: both models were expected to respond equally quickly to happy and neutral stimuli, but in both cases the depressed model is significantly slower ($t(288.35) = 43.385, p < 0.0001$; and $t(288.35) = 41.989, p < 0.0001$, respectively).

We find the biggest discrepancy between the model's response times and the human response times in the perseveration set condition. In general, both models appear to be too fast in this condition, whereas humans tend to be slowest. Since this condition is not directly related to our modelling goals (Levens and Gotlib (2010) did not find any significant effects of group, valence, or their interaction), we will not explore it further here.

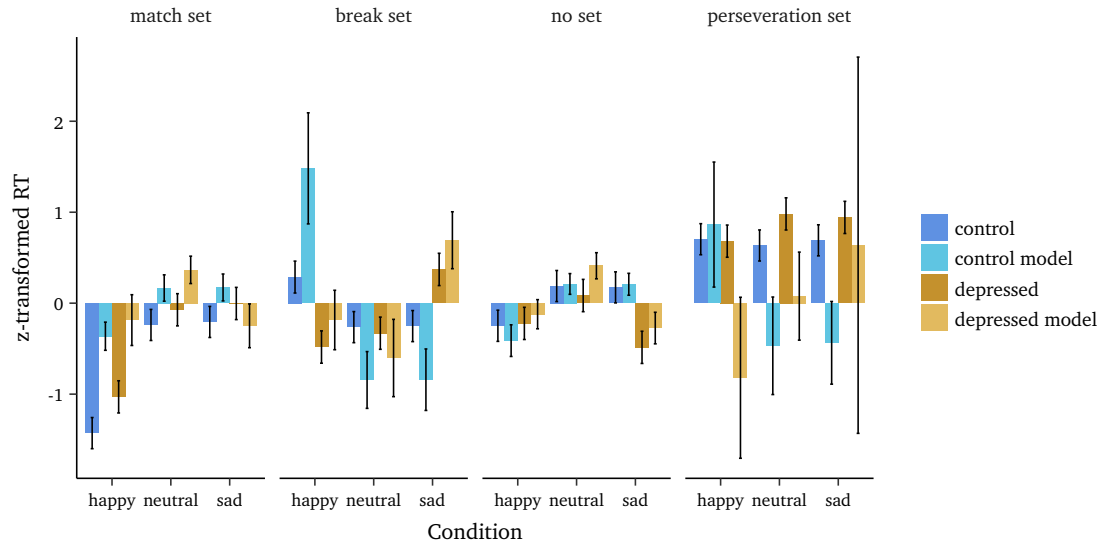


Figure 4.8: Z-transformed response times of real and modelled depressed and control participants on each condition of the 2-back task. Error bars indicate 1 SD from the mean. Trials with response times deviating more than 2.5 SD from the mean are excluded. Human data from Levens and Gotlib (2010).

Chapter 5

Discussion

General conclusion

The effects of depression on cognition have long been of interest to researchers. Experimental studies have highlighted cognitive impairments related to inhibition of information that is no longer relevant, and switching between task goals, as well as attentional and processing biases towards information with negative valence. The occurrence of excessively stable, negatively-coloured forms of mind-wandering has been demonstrated on short time scales and, through real-world experience sampling studies, over a much longer term. Together, these factors cause subtle changes in the depressed mind that are observable through behaviour.

This study has combined theoretical accounts of depression's effect on cognition and applied them in a cognitive model that performs a working memory task. I have shown how manipulating this model can change its behaviour to match either that of a nondepressed control subject, or that of a subject with depression. As such, the model can provide an explicit account of how differences in cognition may bring about observed differences in behaviour.

Rather than attributing behavioural abnormalities in depression to a single factor or model parameter, this modelling approach has pointed us towards a more complex set of differences. Through the model, we can evaluate to what extent the cognitive factors described in the depression literature can explain how behaviour changes. For instance, by increasing the model's tendency to engage in off-task thinking, we can cause its response rate in the task to decrease. This fits well with empirical findings that link depression to higher rates of rumination, a negative, narrowly focussed form of mind-wandering from which it is particularly difficult to disengage (Wigman et al., 2015; Marchetti et al., 2016; Hoffmann et al., 2016). Slower responses in trials during which a previously seen sad item must be forgotten are captured in the model through a process of selective elaboration. By lingering on a mood-congruent stimulus as it is removed from the focus of attention, the model implements a mechanism that can explain inhibitory impairments observed in experimental studies (Whitmer & Banich, 2007; De Lissnyder, Derakshan, De Raedt & Koster, 2011; Whitmer & Gotlib, 2012). And giving the model a perceptual bias toward negative stimuli can explain why a depressed participant responds faster to a sad face than

to a neutral one, in line with previous findings of a negativity bias in perception (Bourke et al., 2010).

Future work

The model of depression that is the topic of this thesis is applicable beyond the *n*-back task that is reported here. The mental ‘operators’ that implement off-task thought and undirected elaboration on particular kinds of task information can be transferred to models of other tasks in which participants are prone to disengage from task-related processing. An example of this is a complex working memory (CWM) task (Huijser, van Vugt & Taatgen, 2017), which mixes an attention task with triggers for self-referential processing (i.e., having the participant think about questions such as “Am I kind?”). The occurrence of self-referential processing makes participants more likely to engage in mind-wandering during the task, with a negative effect on performance. Using the model of depression presented here, one could make predictions of performance on the CWM task by depressed participants, which can in turn be evaluated empirically. Through this process, the generalisability of the current model can be put to the test: which elements of the model are specific to the emotional *n*-back task, and which elements are shared with other tasks?

Improvements in our understanding of how depression affects cognition may benefit the development of novel treatments and diagnostic tools for depression. Current diagnostic tools tend to rely on placing symptoms into discrete, artificial categories, an approach that does not do justice to the high dimensionality of a condition like depression, which can express itself through a variety of mood symptoms, physical symptoms, thought patterns, and cognitive impairments (Adams et al., 2016). Identifying which aspects of behaviour are attributable to ongoing thought (e.g. rumination) and which are due to differences in the cognitive ‘hardware’ can also be helpful in shaping how depression is treated. In a study that repeated the emotional *n*-back task with participants who had recovered from depression, Levens and Gotlib (2015) found that the working memory biases that characterised depressed participants were still present in these subjects—a finding that suggests that these negativity biases are not merely the product of ongoing rumination, but may be hard-wired into the brain. While further work is needed, this may turn out to be a useful tool for proactively identifying whether someone has a heightened risk of developing depression later in life.

As this thesis has demonstrated, a cognitive architecture can serve a useful purpose in evaluating verbal theories of cognitive processes; it forces them to be operationalised within a well-defined set of constraints, and to integrate within a larger whole. While the model in its current form has separate mechanisms that each account for a part of the observed behavioural differences, future work should aim towards a more parsimonious approach. This may require adapting the structure of the architecture in which it is constructed. For example, a factor such as cognitive slowing may be implemented more deeply into the architecture, affecting not just the time it takes to press a key, but also how quickly a production rule fires, or how quickly information moves between buffers.

The current model also does not specify how its working memory biases

interact with its tendency to ruminate. Modelling the contents of rumination in more detail, as was done by van Vugt et al. (2017), may allow for stronger integration between working memory biases and the ruminative process. For instance, a model with a bias towards negativity in its working memory updating may extend this bias through spreading activation towards similarly negative items in declarative memory during off-task thinking. In this way, a bias in working memory places a top-down constraint on the contents of mind-wandering. Alternatively, highly active concepts in memory may interrupt task-directed cognition in a bottom-up fashion when they are thematically related to an item in the focus of attention. Finally, it may be possible to have a general, architecture-wide influence of mood on processing to account for these biases, e.g. as in ACT-R Φ , which extends ACT-R with an affective system (Dancy, 2013).

Concluding remarks

Computational methods, including cognitive modelling, can help us uncover the cognitive mechanisms that shape behaviour in depression and other psychiatric disorders. In this thesis, I have shown how a cognitive model's performance on a working memory task can be transformed from that of a healthy, non-depressed participant, into that of a participant with depression, using mechanisms that were motivated by existing theory. Through the concrete implementation of verbal theories of depression, and by comparing the resulting predictions against human behaviour, we can deepen our understanding of the depressed mind.

References

- Adams, R. A., Huys, Q. J. M. & Roiser, J. P. (2016). Computational Psychiatry: Towards a mathematically informed understanding of mental illness. *J Neurol Neurosurg Psychiatry*, 87(1), 53–63. doi: 10.1136/jnnp-2015-310737
- Altamirano, L. J., Miyake, A. & Whitmer, A. J. (2010). When Mental Inflexibility Facilitates Executive Control: Beneficial Side Effects of Ruminative Tendencies on Goal Maintenance. *Psychological Science*, 21(10), 1377–1382. doi: 10.1177/0956797610381505
- Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* Oxford, United Kingdom: Oxford University Press.
- Andrews-Hanna, J. R., Kaiser, R. H., Turner, A. E. J., Reineberg, A. E., Godinez, D., Dimidjian, S. & Banich, M. T. (2013). A penny for your thoughts: Dimensions of self-generated thought content and relationships with individual differences in emotional wellbeing. *Frontiers in Psychology*, 4. doi: 10.3389/fpsyg.2013.00900
- Ashby, F. G. & Helie, S. (2011). A tutorial on computational cognitive neuroscience: Modeling the neurodynamics of cognition. *Journal of Mathematical Psychology*, 55(4), 273–289. doi: 10.1016/j.jmp.2011.04.003
- Bourke, C., Douglas, K. & Porter, R. (2010). Processing of facial emotion expression in major depression: A review. *Australian and New Zealand Journal of Psychiatry*, 44(8), 681–696.
- Christoff, K., Irving, Z. C., Fox, K. C. R., Spreng, R. N. & Andrews-Hanna, J. R. (2016). Mind-wandering as spontaneous thought: A dynamic framework. *Nature Reviews Neuroscience*, 17(11), 718–731. doi: 10.1038/nrn.2016.113
- Dancy, C. L. (2013). ACT-RΦ: A cognitive architecture with physiology and affect. *Biologically Inspired Cognitive Architectures*, 6, 40–45.
- De Lissnyder, E., Derakshan, N., De Raedt, R. & Koster, E. H. W. (2011). Depressive symptoms and cognitive control in a mixed antisaccade task: Specific effects of depressive rumination. *Cognition & Emotion*, 25(5), 886–897. doi: 10.1080/02699931.2010.514711
- Gollan, J. K., Pane, H. T., McCloskey, M. S. & Coccaro, E. F. (2008). Identifying differences in biased affective information processing in major depression. *Psychiatry Research*, 159(1), 18–24. doi: 10.1016/j.psychres.2007.06.011
- Harel, E. V., Tennyson, R. L., Fava, M. & Bar, M. (2016). Linking major depression and the neural substrates of associative processing. *Cognitive, Affective, & Behavioral Neuroscience*, 16(6), 1017–1026. doi: 10.3758/s13415-016-0449-9

- Hoffmann, F., Banzhaf, C., Kanske, P., Bermpohl, F. & Singer, T. (2016). Where the depressed mind wanders: Self-generated thought patterns as assessed through experience sampling as a state marker of depression. *Journal of Affective Disorders*, 198, 127–134. doi: 10.1016/j.jad.2016.03.005
- Huijser, S., van Vugt, M. K. & Taatgen, N. A. (2017). The wandering self: Tracking distracting self-generated thought in a cognitively demanding context. *Consciousness and Cognition*. doi: 10.1016/j.concog.2017.12.004
- Juvina, I. & Taatgen, N. A. (2007). Modeling control strategies in the n-back task. In *Proceedings of the 8th international conference on cognitive modeling* (pp. 73–78).
- Kopf, J., Dresler, T., Reicherts, P., Herrmann, M. J. & Reif, A. (2013). The Effect of Emotional Content on Brain Activation and the Late Positive Potential in a Word n-back Task. *PLOS ONE*, 8(9), e75598. doi: 10.1371/journal.pone.0075598
- Krimsky, M., Forster, D. E., Llabre, M. M. & Jha, A. P. (2017). The influence of time on task on mind wandering and visual working memory. *Cognition*, 169, 84–90. doi: 10.1016/j.cognition.2017.08.006
- Levens, S. M. & Gotlib, I. H. (2010). Updating positive and negative stimuli in working memory in depression. *Journal of Experimental Psychology: General*, 139(4), 654–664. doi: 10.1037/a0020283
- Levens, S. M. & Gotlib, I. H. (2015). Updating emotional content in recovered depressed individuals: Evaluating deficits in emotion processing following a depressive episode. *Journal of Behavior Therapy and Experimental Psychiatry*, 48(Supplement C), 156–163. doi: 10.1016/j.jbtep.2015.03.009
- Lovett, M. C., Daily, L. Z. & Reder, L. M. (2000). A source activation theory of working memory: Cross-task prediction of performance in ACT-R. *Cognitive Systems Research*, 1(2), 99–118.
- Marchetti, I., Koster, E. H., Klinger, E. & Alloy, L. B. (2016). Spontaneous Thought and Vulnerability to Mood Disorders The Dark Side of the Wandering Mind. *Clinical Psychological Science*, 835 – 857.
- Marr, D. (1982). *A computational investigation into the human representation and processing of visual information*. San Francisco: Freeman and Company.
- Marx, I., Domes, G., Havenstein, C., Berger, C., Schulze, L. & Herpertz, S. C. (2011). Enhanced emotional interference on working memory performance in adults with ADHD. *The World Journal of Biological Psychiatry*, 12(sup1), 70–75. doi: 10.3109/15622975.2011.599213
- Owen, A. M., McMillan, K. M., Laird, A. R. & Bullmore, E. (2005). N-back working memory paradigm: A meta-analysis of normative functional neuroimaging studies. *Human Brain Mapping*, 25(1), 46–59. doi: 10.1002/hbm.20131
- Scheurich, A., Fellgiebel, A., Schermuly, I., Bauer, S., Wölfiges, R. & Müller, M. J. (2008). Experimental evidence for a motivational origin of cognitive impairment in major depression. *Psychological Medicine*, 38(2), 237–246. doi: 10.1017/S0033291707002206
- Schrijvers, D., Hulstijn, W. & Sabbe, B. G. (2008). Psychomotor symptoms in depression: A diagnostic, pathophysiological and therapeutic tool. *Journal of Affective Disorders*, 109(1-2), 1–20. doi: 10.1016/j.jad.2007.10.019
- Smallwood, J., Obonsawin, M. & Heim, D. (2003). Task unrelated thought: The role of distributed processing. *Consciousness and Cognition*, 12(2),

- 169–189. doi: 10.1016/S1053-8100(02)00003-X
- Smallwood, J. & Schooler, J. W. (2006). The Restless Mind. *Psychological Bulletin*, 132(6), 946–958. doi: 10.1037/0033-2909.132.6.946
- Surguladze, S., W Young, A., Senior, C., Brébion, G., J Travis, M. & L Phillips, M. (2004). Recognition Accuracy and Response Bias to Happy and Sad Facial Expressions in Patients With Major Depression. *Neuropsychology*, 18, 212–8. doi: 10.1037/0894-4105.18.2.212
- Suslow, T., Junghanns, K. & Arolt, V. (2001). Detection of Facial Expressions of Emotions in Depression. *Perceptual and Motor Skills*, 92(3), 857–868. doi: 10.2466/pms.2001.92.3.857
- Taatgen, N. A. (2013). The nature and transfer of cognitive skills. *Psychological Review*, 120(3), 439–471. doi: 10.1037/a0033138
- Taatgen, N. A., Katidioti, I., Borst, J. & van Vugt, M. (2015). A model of distraction using new architectural mechanisms to manage multiple goals. In *Proceedings of the International Conference on Cognitive Modeling* (pp. 264–269).
- Tottenham, N., Tanaka, J. W., Leon, A. C., McCarry, T., Nurse, M., Hare, T. A., ... Nelson, C. (2009). The NimStim set of facial expressions: Judgments from untrained research participants. *Psychiatry Research*, 168(3), 242–249. doi: 10.1016/j.psychres.2008.05.006
- van Os, J., Gilvarry, C., Bale, R., Horn, E. V., Tattan, T., White, I. & Group, R. M. o. B. o. t. U. (1999). A comparison of the utility of dimensional and categorical representations of psychosis. *Psychological Medicine*, 29(3), 595–606.
- van Vugt, M. K., Taatgen, N. A., Sackur, J. & Bastian, M. (2015). Modeling mind-wandering: A tool to better understand distraction. In *International Conference on Cognitive Modeling (ICCM)*.
- van Vugt, M. K., van der Velde, M. & ESM-MERGE Investigators. (2017). How does rumination impact cognition? A first mechanistic model. In *Proceedings of the 15th International Conference on Cognitive Modeling* (pp. 25–30). Coventry, UK: University of Warwick, Inclusive Technology Ltd.
- Whitmer, A. J. & Banich, M. T. (2007). Inhibition versus switching deficits in different forms of rumination. *Psychological science*, 18(6), 546–553.
- Whitmer, A. J. & Banich, M. T. (2010). Trait rumination and inhibitory deficits in long-term memory. *Cognition & Emotion*, 24(1), 168–179. doi: 10.1080/02699930802645762
- Whitmer, A. J. & Gotlib, I. H. (2012). Switching and backward inhibition in major depressive disorder: The role of rumination. *Journal of Abnormal Psychology*, 121(3), 570–578. doi: 10.1037/a0027474
- Wigman, J. T. W., van Os, J., Borsboom, D., Wardenaar, K. J., Epskamp, S., Klippel, A., ... Wichers, M. (2015). Exploring the underlying structure of mental disorders: Cross-diagnostic differences and similarities from a network perspective using both a top-down and a bottom-up approach. *Psychological Medicine*, 45(11), 2375–2387. doi: 10.1017/S0033291715000331

Appendix A

Model parameters

The set of architectural parameters comprises those parameters that are set to a fixed value and not changed between depressed and non-depressed variants of the model. Where possible, these parameters are set to their respective default values. The parameters of the 0-back model whose value is different from the default are listed in Table A.1. Those of the 2-back model are given in Table A.2.

Table A.1: Non-standard parameter settings of the 0-back model.

Parameter	Value	Default	Explanation
default-activation	2.0	none	lower bound on chunk activation
ans	0.05	0.2	activation noise
rt	-5.0	-2.0	retrieval threshold
default-inter-operator-assoc	2.0	1.0	associative strength between operators of same goal
default-operator-self-assoc	-3.0	-1.0	spreading activation from each operator to itself
pm	t	nil	partial matching of retrievals
subvocalise	0.3 ± 0.1 s	none	duration of subvocalising the target
focus	0.15 ± 0.05 s	none	duration of focusing on the task

Table A.2: Non-standard parameter settings of the 2-back model.

Parameter	Value	Default	Explanation
default-activation	0.0	none	lower bound on chunk activation
ans	0.05	0.2	activation noise
rt	-5.0	-2.0	retrieval threshold
lf	0.1	0.2	latency factor
default-operator-self-assoc	-3.0	-1.0	spreading activation from each operator to itself
production-prim-latency	0.0	0.3	time to fire subsequent productions of an operator
imaginal-delay	0.0	0.2	time to create an imaginal chunk
pm	t	nil	partial matching of retrievals
mp	10	5	mismatch penalty applied to partial matches
focus	0.15 ± 0.05 s	none	duration of focusing on the task

Appendix B

Resources

The model files, simulation data, and the scripts used for the analysis (along with instructions and additional analyses) can be downloaded from <https://github.com/maartenvandervelde/emotional-n-back>

A custom build of the PRIMs cognitive architecture is needed to run the models (requires macOS 10.12 or newer). Its source code can be downloaded from <https://github.com/maartenvandervelde/ACTransfer/tree/Maarten>