Modelling the Effect of Depression on Working Memory

van der Velde, Maarten; van Vugt, Marieke; Taatgen, Niels

Modelling the Effect of Depression on Working Memory

Maarten A. van der Velde (m.a.van.der.velde@rug.nl)
Marieke K. van Vugt (m.k.van.vugt@rug.nl)
Niels A. Taatgen (n.a.taatgen@rug.nl)
Department of Artificial Intelligence, University of Groningen
Nijenborgh 9, 9747 AG Groningen, The Netherlands

Abstract

Individuals with depression are prone to engaging in rumination, a process in which attention turns inwards to narrowly-focused, negative patterns of thought, at the cost of attending to a task. Other core deficits associated with depression are weaker inhibition of information that is no longer relevant, and a negative perceptual bias. Here, we present a computational cognitive model that uses these mechanisms to explain performance on an n-back task in which the stimuli are faces with different emotional expressions, and in which depressed participants exhibit specific impairments. These impairments are explained by assuming that depressed participants selectively elaborate on sad items as they are removed from working memory, and that they have a perceptual bias towards sad faces. In this way, by specifying a mechanism by which performance impairments come about, the model helps to provide a deeper understanding of the cognitive processes underlying behaviour.

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

Introduction

Depression is an important mental health issue, but much is still unknown about its cognitive mechanisms. A major cognitive component of depression is thought to be rumination, which can be conceptualised as a maladaptive form of mind-wandering (Marchetti, Koster, Klinger, & Alloy, 2016). Mind-wandering, a process in which attention is directed away from a task towards internal thoughts and memories, can take up as much as 50% of our waking hours, and can have both positive and negative effects on cognitive performance (Mooneyham & Schooler, 2013). What sets rumination apart from other forms of mind-wandering is that it is characterised by a thematic narrowness, focusing primarily on negative memories and thoughts related to one's current dysphoric state. Ruminators tend to repeatedly engage in certain narrow trains of thought from which it can be difficult to escape (Christoff, Irving, Fox, Spreng, & Andrews-Hanna, 2016). While both depressed and non-depressed individuals have been observed to engage in off-task thinking during laboratory tasks (e.g., Smallwood, Obonsawin, & Heim, 2003), participants with depression spend more time mind-wandering than their non-depressed counterparts (Hoffmann, Banzhaf, Kanske, Bermpohl, & Singer, 2016).

Rumination has been associated with several cognitive effects, including inhibitory and switching deficits. Individuals who ruminate tend to have more difficulty inhibiting information that is no longer relevant to their current situation, compared to non-ruminators. A study by Whitmer and Banich (2010), in which participants memorised a list of study items and subsequently rehearsed only a subset of this list, found that those with a tendency for rumination inhibited unrehearsed items from the original list less strongly than non-ruminators. This inability to effectively inhibit outdated information also affects working memory. In a task-switching paradigm, the occurrence of depressive rumination is associated with a weaker inhibition of previous task sets (Whitmer & Banich, 2007). Depressed individuals find it particularly difficult to inhibit obsolete information when it has a negative valence (Joormann & Gotlib, 2008), whereas non-depressed individuals have similar difficulty inhibiting positive stimuli (Deveney & Deldin, 2006), suggesting a mood-congruency effect.

Aside from the inhibitory deficits associated with rumination, depression can also influence visual perception and attention. Whereas healthy humans commonly exhibit a perceptual bias towards positive stimuli, perception in those with depression tends to be coloured by a negativity bias. When Gollan, Pane, McCloskey, and Coccaro (2008) asked participants to judge the emotion of a neutral face in a forced-choice task, non-depressed participants saw its emotion as happy more often than sad, while depressed participants were more inclined to judge the same face as sad. Visual attention is also not immune to bias: depressed individuals tend to demonstrate selective attention towards sad faces and/or away from happy faces (Bourke, Douglas, & Porter, 2010; Suslow, Junghanns, & Arolt, 2001). Indeed, depression may be linked to a general negativity bias, which also reveals itself in, e.g., memory recall (Dalgleish & Watts, 1990).

Gaining a better understanding of how depression influences and interacts with cognition could potentially aid the development of better diagnostic tools and more effective treatments. Schemes that are currently used to diagnose psychiatric disorders tend to be overly simplistic in their categorisation, and do not translate well to the clinical level (van Os et al., 1999). Existing cognitive theories of depression, generally formulated as verbal accounts, can be difficult to generalise or integrate with other theories. 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 (Adams, Huys, & Roiser, 2016).

We have been developing such a computational model in the ACT-R cognitive architecture (Anderson, 2007), which is well-suited to translating verbal theories into quantitative predictions. Increasing the prominence of negative items in the model's declarative memory transforms its regular mind-wandering behaviour into more persistent, negatively-coloured depressive rumination (van Vugt, van der Velde, &

ESM-MERGE Investigators, 2018). When applied to a sustained attention task, the model predicts lower response accuracy as a result of this change.

Here, we extend this model to a more complex task, the *emotional n-back* task (Levens & Gotlib, 2010). Modelling this task requires the addition of some mechanisms to our existing model. The task, a modified version of the *n*-back, uses images of happy, neutral, and sad faces as its stimuli. It involves complex working memory operations, as participants are required to maintain an up-to-date mental list of recently seen items in their memory in order to respond correctly (Lovett, Daily, & Reder, 2000; Juvina & Taatgen, 2007). Because of its use of emotional face stimuli and its reliance on working memory, this task is expected to elicit many of the cognitive effects of depression.

Levens and Gotlib (2010) had depressed and non-depressed participants perform the emotional *n*-back task, and found that, while depressed participants responded more slowly and slightly less frequently across conditions, they were as accurate as healthy controls at identifying whether a face had the same expression as one shown two trials earlier. Response times did reveal differences in working memory updating: depressed participants responded relatively quickly in trials with a sad stimulus, compared to happy or neutral trials, indicating easier integration of a sad face into their mental list. Furthermore, depressed participants were relatively slow when breaking a former set of matching faces with sad expressions, suggesting a tendency to linger on negative items as they were pushed off the mental list, possibly because these items activated related concepts and memories more easily. For instance, a sad face might trigger ruminative thoughts in a participant about a negative experience or their current dysphoria. Such an elaboration on mood-congruent items could also account for the observation that non-depressed control participants responded relatively slowly in trials that involved breaking a set of happy faces, as these items might trigger positive elaboration in healthy participants. Surprisingly, there was no evidence of a perceptual bias: in a 0-back control condition, in which participants compared each stimulus to a target expression given earlier, depressed and non-depressed participants performed identically.

We present a computational cognitive model that can explain the observed differences in performance between depressed participants and healthy controls, on the basis of simple assumptions about mental elaboration on mood-congruent items as they are removed from working memory, which can account for the behavioural differences between participants with and without depression. Our model also suggests that the existence of a perceptual bias is necessary for capturing depressed participants' faster responses to sad stimuli in the 2-back task.

Model

Following our previous model of rumination, performance on the emotional *n*-back task is modelled as a competition for cognitive resources between task-directed thought and off-task thought (i.e., mind-wandering or rumination).

The model is implemented in the PRIMs cognitive architecture (Taatgen, 2013), an extension of the ACT-R architecture (Anderson, 2007). Like ACT-R, PRIMs provides a framework for modelling an entire task, from perceptual input to motor output, and everything in between. 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: if-then rules requiring certain conditions to be met before a sequence of 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.

Unlike ACT-R, the PRIMs architecture makes no theoretical distinction between procedural knowledge (operators or productions) and declarative knowledge (facts): both are treated as regular chunks stored in declarative memory. This means that, like declarative facts, each operator has an activation value (representing its memory trace strength, which is boosted by retrievals and decays over time), and is subject to spreading activation. Goals spread activation to their operators, making operators belonging to more active goals more likely to be selected for execution. Operators can also spread activation to each other, thereby increasing the likelihood that operator execution follows a certain favoured sequence. Chunks in other buffers, such as the imaginal buffer or the retrieval buffer, can furthermore influence operator selection through spreading activation. Because of these dynamics, a PRIMs model can freely alternate operators from different goals. Our emotional *n*-back model switches flexibly between operators that fall under the umbrella of the task goal, and operators that enact the mind-wandering goal.

The model also builds on previous ACT-R models of the 2-back task. Lovett et al. (2000) and Juvina and Taatgen (2007) implemented multiple 2-back strategies. Here, we implement their high-effort strategy, since participants were explicitly instructed to use this method. This means that the model holds a list of the most recently seen stimuli in its declarative memory that is updated every trial. Each item on this list is a chunk, encoding the facial expression of the stimulus (*happy*, *neutral*, or *sad*) as well as its position in the list (*0-back*, *1-back*, or *2-back*). At the start of a trial, the model retrieves each item on the list in turn, increments its list position, and pushes the modified item back to declarative memory. 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 new stimulus that is currently on screen is then added to the front of the list. To decide its response, the model retrieves the 2-back item from memory and compares it to the current stimulus.

The recognition of facial expressions is modelled as a memory retrieval: the model sees a string representing a face, which it uses as a cue to retrieve the corresponding expression from its declarative memory. This mechanism allows us to model differences in recognition speed by varying the activation of face chunks of different emotional valence. Levens and Gotlib (2010) found that both depressed and non-depressed participants responded more quickly to happy faces than to sad or neutral faces, so chunks representing happy faces are given a higher activation, making them faster to retrieve. Since humans are not perfect at recognising facial expressions, we use partial matching to allow the occasional retrieval of an incorrect facial expression. Participants in both groups were more accurate in their responses to happy faces than to sad or neutral faces, with neutral faces being the most difficult to categorise (Levens & Gotlib, 2010). This effect is captured in the model by varying the mismatch penalty that is applied to a non-matching face on the basis of its expression.

The competing mind-wandering goal is implemented as a process of memory retrieval, following earlier models of mind-wandering (van Vugt, Taatgen, Sackur, & Bastian, 2015; van Vugt et al., 2018). At any point in the task, operators from this goal can initiate the retrieval of a sequence of random items from declarative memory. Once a memory is retrieved, it is placed in the model's working memory to bring it to the forefront of attention, and a new memory is retrieved. This process continues until a task operator is selected.

In addition to the 2-back condition, we also model the 0-back condition, which Levens and Gotlib (2010) used to show that depressed participants did not have a different perceptual bias than non-depressed controls. The 0-back model works the same as the 2-back model when it comes to perceiving and classifying stimuli, but differs in how it determines the correct response. In the 0-back condition, each stimulus has to be compared to a target expression, which eliminates the need for maintaining a history of recent stimuli in memory. Instead, once the model has determined the facial expression of a stimulus, it directly compares the expression to the target stored in its goal buffer, and responds accordingly.

To capture the differences between depressed and non-depressed participants, we manipulated several aspects of the model. Firstly, since depression is associated with more pervasive mind-wandering, the activation of mind-wandering operators was increased in the depressed version of the model, making them more likely to be selected. In addition, we recreated the general psychomotor slowing found in those with depression (Schrijvers, Hulstijn, & Sabbe, 2008) by increasing the time needed to execute a keypress in the depressed model. These two manipulations were expected to lower the depressed model's response rate on the *n*-back, while increasing its average response time.

Depression is linked to difficulties inhibiting negative emotional information that is no longer relevant, so the model includes a mechanism by which it can elaborate on sad faces as they are discarded. Whenever a sad face is pushed off the model's mental list of recent stimuli, the task process can be briefly interrupted by an elaboration operator from the mind-wandering goal. Although it is simply implemented as a delay in the current model, this operator represents the activation of related concepts and memories that are triggered by the sad face. The time penalty that is incurred by lingering on old information and subsequently having to refocus on updating the mental list will slow down the model's response in the affected trials. Since non-depressed individuals have been found to exhibit a difficulty letting go of positive emotional information, a similar mechanism is implemented in the non-depressed version of the model. Rather than elaborating on negative items, however, the non-depressed model has an operator that can elaborate on happy faces as they are removed from the mental list.

Finally, the depressed version of the 2-back model includes a negative perceptual bias: sad face chunks are given a higher activation than in the non-depressed version of the model, which makes them faster to retrieve. Since Levens and Gotlib (2010) did not find evidence of such a bias in the 0-back condition, it is left out of the 0-back model.

Parameters that distinguish the depressed model from the control model (activation of mind-wandering operators, keypress duration, and activation of face chunks) were initially set to reasonable values and were subsequently adjusted to better fit the behavioural results. The full code for depressed and control versions of the model is available at github.com/maartenvandervelde/emotional-n-back/.

Methods

0-back
The depressed and control version of the 0-back model were each run 50 times, simulating 50 participants per group (similar to the 29 participants per group in the empirical data). Both models went through 9 full training runs of 3 blocks with 43 trials each to learn the task, before they performed a final run. Only the final run of each model was analysed.

2-back
Both versions of the 2-back model were run 50 times, simulating 50 depressed participants and 50 control participants. The models received a similar amount of training as before (3 full runs containing 6 blocks of 55 trials), before performing a final run. Once again, only this last run was analysed.

Analysis
Performance of the depressed and control versions of the emotional *n*-back model was compared to that of the depressed and healthy participants in Levens and Gotlib (2010). 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. Trials in which the model did not respond within 2 seconds after stimulus onset were considered non-response trials. These trials were excluded when calculating response accuracy. To allow for further comparison of
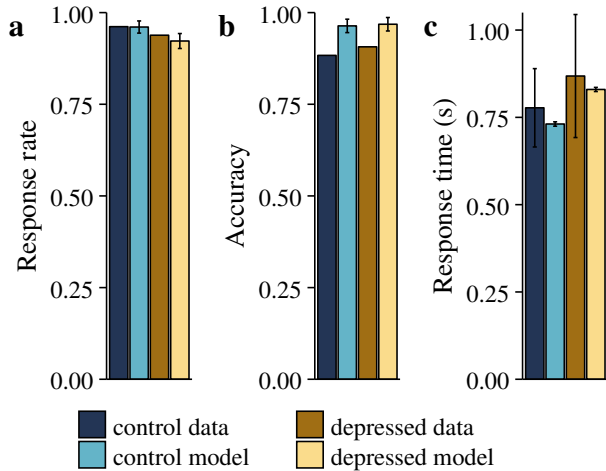
Figure 1: 0-back model performance compared to that of real participants. The model replicates the slightly lower response rate of depressed participants relative to controls (**a**), as well as the lack of a difference in accuracy between groups (**b**). As in humans, the depressed version of the model has longer response times (**c**). Error bars indicate 1 SD from the mean (SD of the human response rate and accuracy was unavailable).
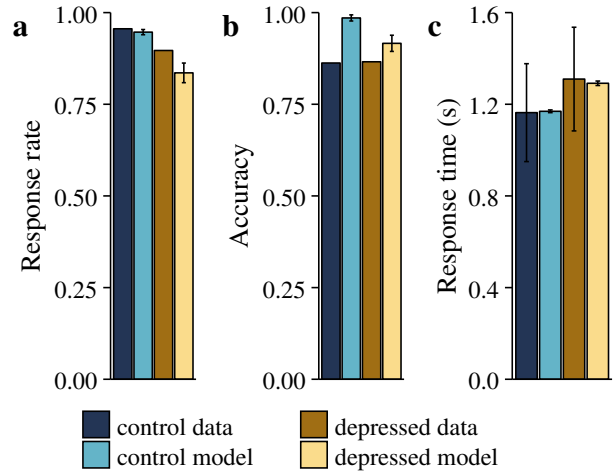

Figure 2: 2-back model performance compared to that of real participants. The model replicates the lower response rate of depressed participants relative to controls (**a**), but does not capture the similar response accuracy of both groups (**b**). As in humans, the depressed version of the model has longer response times (**c**). Error bars indicate 1 SD from the mean (SD of the human response rate and accuracy was unavailable).

relative response speed between groups, response times were normalised using a z-transformation (for every participant, each condition mean RT was subtracted from their overall mean RT and divided by the SD of the condition RT).

Results

0-back

Figure 1 summarises the 0-back model's behavioural fit to the human data[1]. It shows that, by increasing the model's tendency to engage in mind-wandering and by slowing down its motor response, we can reproduce the lower response rate and higher response times of depressed participants, relative to their non-depressed counterparts (although the model's response times are less variable than those of human participants). As in humans, depression does not affect response accuracy; the depressed and non-depressed versions of the model have equal accuracy (albeit slightly higher than the accuracy of human participants). Table 1 confirms that the model fits the data satisfactorily.

2-back

Figure 2 shows the fit of the 2-back model to the human data. As before, the higher activation of mind-wandering operators in the depressed version, together with a longer keypress time, causes it to correctly predict a lower response rate and higher response times relative to the control model. However,

while participants in both groups responded equally accurately, the control model is more accurate than the depressed model. Both models' accuracy is slightly too high, but otherwise they capture the differences between groups (Table 1).

Z-transformed response times in two relevant 2-back conditions are shown in Figure 3. As Figure 3a shows, the model reproduces the relatively slow responses that occur in trials which require a mood-congruent face to be removed from the mental list. By elaborating on sad faces as they are being discarded, the depressed version of the model takes relatively long to respond in trials in which a previous set of sad faces is broken. In contrast, the control model is slower than usual in trials that require the breaking of an old set of happy faces, as it spends some time elaborating on the happy face being discarded. When breaking a neutral set, neither model elaborates on the discarded face. As a result, both are equally fast (accounting for differences in baseline speed), matching the

---

[1]The output data of the model, as well as additional analyses of the results, including a more detailed breakdown by condition and statistical comparisons, are available at github.com/maartenvandervelde/emotional-n-back/.

Table 1: Model fit to data. RMSE = root-mean-square error.

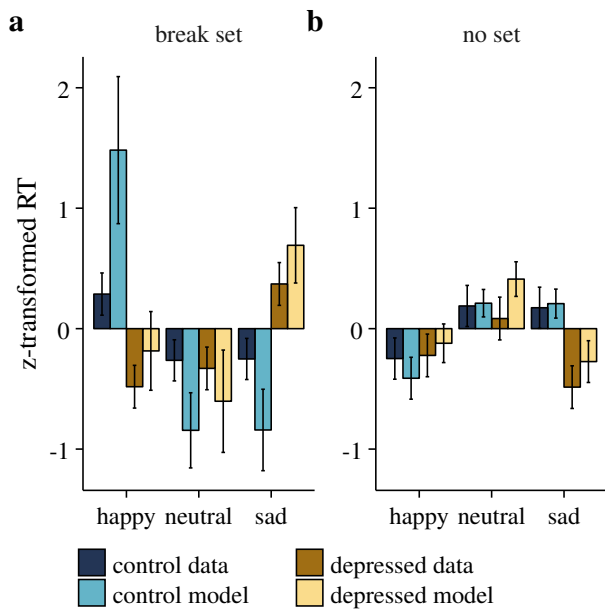| Task	Measure	RMSE
0-back	Response rate	.01
	Accuracy	.07
	Response time (s)	.04
2-back	Response rate	.04
	Accuracy	.09
	Response time (s)	.01
	Z-transformed RT	.46

Figure 3: Z-transformed response times of the model compared to those of human participants on two conditions of the 2-back task. **(a)** When breaking a mental set (i.e., discarding one of the elements that previously made up a set from the mental list), the model reproduces the mood-congruent slowing effect seen in humans: the control model responds relatively slowly when breaking a happy set, while the depressed model is slow when breaking a sad set. Like humans, both models are equally fast when breaking a neutral set. **(b)** In non-match trials without a set to break, the depressed model responds more quickly to a sad stimulus than the control model, matching the pattern found in humans. Error bars indicate 1 SD from the mean.

pattern found in human participants.

Figure 3b further shows that, by manipulating the speed at which the model recognises a particular facial expression, we can recreate the human response time pattern observed in trials without a previous set to break or a current match. In particular, the higher activation of sad face chunks in the depressed model allows it to respond more quickly in trials with a sad stimulus, compared both to its own response times in happy or neutral trials and to the control model's responses in sad trials.

Discussion

In this paper we have presented a model of the cognitive effects of depression on performance in an emotional working memory task. Through the model, we have identified several key factors that differ between depressed and non-depressed participants. Depressed patients have a stronger tendency for mind-wandering, at the cost of attending to the task, which explains a lower response rate. Additionally, a slower motor response accounts for a higher average response time. Since

individuals with depression were previously found to have difficulty inhibiting negatively-valenced information once it was no longer relevant, we implemented a mechanism for elaborating on negative items as they were discarded. With this mechanism, the model reproduces the observed response time pattern in the human data. The same mechanism, but focused on positive instead of negative items, furthermore captures the opposite pattern that was present in the non-depressed control group.

A potential criticism of our model is that it fails to provide a unified explanation for depressed participants' lack of a perceptual bias in the 0-back on one hand, and their faster responses to sad stimuli in the 2-back on the other. Depressed 0-back performance is captured without a perceptual bias relative to the control model. Yet, the depressed 2-back model requires faster perception of sad faces to reproduce depressed participants' relatively fast responses to sad faces.

There are two responses to this. Firstly, it is possible that the absence of perceptual bias in the 0-back task was a chance finding—indeed, it is inconsistent with earlier studies associating depression with a negative perceptual bias, and a study by the same authors in which participants who had recovered from depression performed the same task did find a sadness bias in the 0-back (Levens & Gotlib, 2015). If so, we could use the same response bias in the 0-back and 2-back models.

Alternatively, the inconsistency is due to an architectural limitation. It can be argued that the lower response times to sad stimuli in the 2-back task are not the result of faster perception, but rather of faster integration into working memory. Currently, PRIMs (and ACT-R) assumes that creating a new chunk in working memory takes a fixed length of time, irrespective of its contents. If it is indeed the case that depressed individuals integrate sad faces more quickly into working memory than neutral or happy faces, the architecture would have to be extended to support a variable working memory integration time that is dependent on an item's content.

An additional limitation of the model is that it assumes there to be little interaction between perceptual bias, working memory bias, and the mind-wandering process. Instead, these components are modelled as separate mechanisms. In an earlier version of the model we unsuccessfully used spreading activation in an attempt to influence task performance through the occurrence of mind-wandering. In that model, mind-wandering memories had an emotional valence (*happy*, *neutral*, or *sad*) and could spread activation to face chunks with their respective emotion, such that negatively-themed mind-wandering would make recognition of sad faces faster, and the processing of sad faces would more easily trigger mind-wandering. This approach failed because spreading activation is too fleeting, disappearing as soon as its source leaves its buffer. To make it work would require having a longer-lasting effect of spreading activation. Alternatively, affect-related spreading activation should come from an external source (such as physiology, see, e.g., ACT-R Φ (Dancy, 2013), which would also allow one to model some of the biologi-

cal factors underlying depression). A further option would be to model the tendency to elaborate on mood-congruent information as a learned association. Future work should examine whether these mechanisms could achieve a similar fit to the data while requiring fewer parameters.

In summary, we have shown how computational cognitive modelling can help us uncover the cognitive mechanisms that shape behaviour in depression and other psychiatric disorders. With a stronger propensity for mind-wandering, slower motor actions, a negative perceptual bias, and selective elaboration on mood-congruent items in working memory, our model reproduces specific behavioural deficits in working memory observed in humans. Through the implementation of existing theory, 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.

Anderson, J. R. (2007). *How Can the Human Mind Occur in the Physical Universe?* Oxford, United Kingdom: OUP.

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.

Dalgleish, T., & Watts, F. N. (1990). Biases of attention and memory in disorders of anxiety and depression. *Clinical Psychology Review*, *10*(5), 589–604.

Dancy, C. L. (2013). ACT-RΦ: A cognitive architecture with physiology and affect. *Biologically Inspired Cognitive Architectures*, *6*, 40–45.

Deveney, C., & Deldin, P. (2006). A preliminary investigation of cognitive flexibility for emotional information in major depressive disorder and nonpsychiatric controls. *Emotion*, *6*, 429–37.

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.

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.

Joormann, J., & Gotlib, I. H. (2008). Updating the contents of working memory in depression: Interference from irrelevant negative material. *Journal of Abnormal Psychology*, *117*(1), 182–192.

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).

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.

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.

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.

Mooneyham, B. W., & Schooler, J. W. (2013). The costs and benefits of mind-wandering: A review. *Canadian Journal of Experimental Psychology/Revue canadienne de psychologie expérimentale*, *67*(1), 11–18.

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.

Smallwood, J., Obonsawin, M., & Heim, D. (2003). Task unrelated thought: The role of distributed processing. *Consciousness and Cognition*, *12*(2), 169–189.

Suslow, T., Junghanns, K., & Arolt, V. (2001). Detection of Facial Expressions of Emotions in Depression. *Perceptual and Motor Skills*, *92*(3), 857–868.

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

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 *Proceedings of the 13th International Conference on Cognitive Modeling* (pp. 252–257).

van Vugt, M. K., van der Velde, M., & ESM-MERGE Investigators. (2018). How Does Rumination Impact Cognition? A First Mechanistic Model. *Topics in Cognitive Science*, *10*(1), 175–191.

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.