The Metronome Counting Task for Measuring Meta-awareness

Abstract

How can we measure the absence of awareness? Attention research has developed tools for measuring self-caught meta-awareness restoration and behavioural mind-wandering, but we lack a way to dynamically track the loss of meta-awareness. The present pre-registered study sought to bring together three extant paradigms into one tool designed to dynamically measure meta-awareness: the Metronome Counting Task (MCT). The MCT is a continuous performance task wherein participants tap along to a steady beat while counting to twenty, indicating the final count by a special button press. This sample ($N = 74$) provides evidence that participants could self-catch their failures in the task, that a response variability metric measuring mind-wandering depth was successfully recreated in this new tool, and that dynamic performance changes may be useful for detecting meta-awareness loss before participants become internally aware of the loss or are caught by external errors. The MCT was conceived as a tool that will support neuroimaging models of dynamic fluctuations during sustained attention, providing a link between the phenomenology of meta-awareness, the behaviour measured by a replicable index of task engagement, and a continuous performance task on time-scales relevant for MRI. We discuss the possibility that meta-awareness may exist on a continuum and that conceptions of mind-wandering as attention failures may plausibly be reconceived as changes in goal priority manifesting as shifting task engagement.

Keywords: attention; meta-awareness; continuous performance task; mind-wandering; task-engagement;

Glossary of Abbreviations

ARCES: Attention Related Cognitive Errors Scale

COG: Constraints on Generality statement

I/O: Input / Output

MA: Meta-awareness

MCT: Metronome Counting Task

MRT: Metronome Response Task

MW: Mind-wandering

OSF: Open Science Framework

RRTv: Rhythmic Response Time variance

SART: Sustained Attention to Response Task

1. Introduction

How do you find something you don't know you've lost? Attention drifts away from our activities despite our intentions yet, after some time, we notice that we have lost focus. This "noticing" is the core of meta-awareness (MA): the process of attending to and appraising the current contents of consciousness (Chin & Schooler, 2010). Critically, MA involves noticing when focal behaviors are inconsistent with broader intentions, allowing resource engagement to overcome the inconsistency.

Meditation provides perhaps the clearest example of MA and cultivation of MA is central to mechanistic theories of meditation (Lutz, Jha, Dunne, & Saron, 2015). In most styles of meditation, the broader intention is to focus on an anchor of attention (e.g. the breath) (Anderson & Farb, 2018). The meditator pays attention to the anchor, eventually becomes lost in thought, subsequently becomes meta-aware that their mind is wandering, and shifts back to focusing on the anchor (Hasenkamp, Wilson-Mendenhall, Duncan, & Barsalou, 2012). The phenomenological experience of MA in anchor- or intention-based meditation is especially clear because the intention and task are so well-defined yet distraction is so inevitable despite a lack of outside interference. This makes noticing the inevitable misalignment between behaviour and intention extremely salient.

MA is also central to therapy (Hayes, 2004). Without MA, vital self-reflection is lacking. MA deficits underlie numerous psychopathologies (Goodkind et al., 2015) associated with grey-matter reductions in the salience network (Menon, 2011), neural areas related to noticing unexpected events (i.e. bilateral insula and dorsal anterior cingulate cortex). Research has also shown that MA activity correlates with salience network activity (Hasenkamp et al., 2012), but no extant measurement tools can detect the loss of MA. Indeed, participants cannot report MA loss as one would need MA to report losing it: a catch-22.

While MA is an inherently subjective phenomenon, we propose that one objective, measurable marker of MA loss is misalignment between intention and behaviour in feasible tasks [Footnote]. Tasks

are here conceived to be behavioural patterns pursuant to intentions. A task is feasible insofar as intended behaviours are readily accomplishable, i.e., expertise is present or not required and there is no environmental interference preventing the intended behaviour. According to this framework, MA loss could occur in any intended task —reading, driving, listening to a colleague, counting— and task-disengagement could be to anything —current concerns, flights of fancy, loved ones, etc. So long as the task is feasible, MA loss can be inferred any time behavior deviates from intention because if MA was maintained, the alignment between intention and behaviour would be maintained. Crucially, these definitions allow for the existence of intentional disengagement, e.g. intentional mind-wandering (Seli, Risko, Smilek, & Schacter, 2016). Intention-behaviour alignment can be assessed by task-performance thus MA involves understanding the relationship between broader intentions and focal task-based behavioural performance, two constructs that are rarely explicitly distinguished.

[Footnote: One way to think of MA loss is that it should be reflected behaviourally as the loss of a broader intention. The relation is one of *modus tollens*: in the absence of external influences or skill-based constraints, if one retains MA of an intention then one will behave in alignment with that intention; as such, if one does not behave in alignment with an intention, one must not have retained MA of that intention. The converse is not necessarily true, that is, if one does not retain MA of the intention, behaviour may remain in alignment with the intention despite the loss of MA. The common ideas of “autopilot” and well-trained behaviours are useful here. For example, one might walk around without MA of the intention of placing one foot in front of the other and yet succeed in walking without bumping into anything (Harms, van Dijken, Brookhuis, & de Waard, 2019).]

The present study aims to introduce a task that begins to measure MA loss more explicitly by distinguishing intention and behaviour within one task. Intentions operate in wider temporal windows than focal behaviour; in this way, intentions can be understood as short-term prospective goals for behaviour. A higher-order prospective memory measure could thus indicate MA loss through a failure to behaviourally comply with the broader intention. A continuous performance task was selected to avoid

speed-accuracy trade-offs inherent to sustained attention tasks that require response inhibition. Continuous performance tasks allow moment-to-moment variability in task-performance to be used to index task-engagement.

Recently, Seli et al. (Seli, Smilek, Ralph, & Schacter, 2018) administered a time-based prospective memory task wherein participants indicated whether on-screen words represented natural or artificial objects, meanwhile also being required to indicate when one minute had passed; participants were able to check a timer to enhance their accuracy and the number of checks was used as a measure of MA “checking” (or “polling”, see Future Directions). They then administered the Sustained Attention to Response Task (SART) in which participants self-caught their own mind-wandering; as discussed below, self-caught mind-wandering is a measure of MA restoration. They found that the rate at which participants self-caught and the rate of timer-checking was correlated at $r = 0.2$ (Seli, Smilek, et al., 2018). While the aforementioned study is both innovative and informative, we believe the principle metrics of multiple tasks and paradigms can be combined into one tool.

The present study introduces a novel tool: the Metronome Counting Task (MCT). This continuous performance task with built-in prospective memory task is intended to measure meta-awareness and is based on three extant paradigms: a self-caught paradigm that measures meta-awareness restoration (adapted from (Hasenkamp et al., 2012)), a focal versus broad task monitoring paradigm to rapidly indicate MA loss (adapted from (Levinson, Stoll, Kindy, Merry, & Davidson, 2014)), and a response variability paradigm that tracks task-engagement (adapted from (Anderson, Lin, & Petranker, 2017)). Each of these paradigms will be reviewed in turn.

1.1 The Self-Caught Paradigm

Determining when MA is restored is relatively straightforward: a participant presses a button when they catch their mind wandering from an experimenter-defined task (2012). Catching one’s mind wandering is perhaps the most emblematic example of MA restoration. This simple and intuitive paradigm has been used extensively in mind-wandering and meditation research. For example,

Hasenkamp et al. (Hasenkamp et al., 2012) had meditators in an MRI indicate self-caught mind-wandering (i.e. spontaneous MA restoration) with a button-press. The MCT incorporates this paradigm by allowing participants to indicate when they notice that they have lost track of the broader intention.

1.2 The Focal versus Broad Task Monitoring Paradigm

The self-caught paradigm is limited insofar as it can only detect the spontaneous restoration of MA, not track the loss of MA or any dynamic fluctuations in MA. Measuring MA loss is far less intuitive as to consciously report MA loss, the participant would need MA of losing MA. One way of assessing whether participants have lost MA is by distinguishing between a focal behavior and a broader intention within a task: if participants lose track of the broader intention, they have lost MA of that intention. For example, in a breath counting task (Levinson et al., 2014), participants counted their breaths and indicated each breath with a button-press (focal behavior), but they indicated the ninth breath with a different button-press (broad intention). An incorrect button-press thus signaled that the broader intention was overlooked, i.e. MA was lost. This task has been adapted in the MCT by having participants simultaneously engage a focal task (tapping synchronously with the metronome) and a broad intention (counting to twenty).

1.3 The Response Variability Paradigm

The self-caught and focal versus broad task monitoring paradigms are limited in that they cast MA as a binary phenomenon. We prefer to allow for the possibility that there may be dynamic fluctuations in MA and task-engagement, that participants may engage tasks other than the experimenter-defined task, and that participants may engage more than one task simultaneously. It may be that one keeps several goals in mind, using meta-awareness to dynamically monitor performance and reprioritize goals, redistributing resources so acceptable performance levels can be maintained. Other times, meta-awareness may fail to monitor multiple goals adequately resulting in suboptimal resource allocation resulting in errors. It is also a practical reality that MA may be lost relatively infrequently by task-experts

or in certain tasks, which can result in inefficient assessment, like the bulk of expensive neuroimaging time going underused (2012).

Effective dynamic detection of MA may be possible through continuously measuring task-engagement. Task-engagement has been measured as an inverse proxy for mind-wandering (MW), which is often (imperfectly) assumed to reflect the absence of MA (Chin & Schooler, 2010). A number of behavioral tasks assess engagement and MW (see Fortenbaugh et al., 2017 for an extended treatment). The two most common ways to measure MW are by response variability and response-inhibition, though response inhibition tasks face a number of theoretical issues, i.e. speed-accuracy trade-offs and response-strategy confounding interpretation (Seli, Jonker, Cheyne, & Smilek, 2013). To overcome these issues, the Metronome Response Task (MRT) (Seli, Cheyne, & Smilek, 2013) was created.

The MRT is a continuous performance task wherein response variability dynamically indexes MW (Anderson et al., 2017; Laflamme, Seli, & Smilek, 2018; Meier, 2018; Seli, Cheyne, et al., 2013; Seli, Cheyne, Xu, Purdon, & Smilek, 2015; Seli, Jonker, Cheyne, Cortes, & Smilek, 2015). The experimenter-defined task has participants tap along with a steady metronome beat as synchronously as they can; unlike response inhibition tasks, the MRT creates no speed-accuracy trade-off as accuracy is defined by temporal synchrony with the metronome. The MRT employs probes to measure attentional states and self-reported MW is associated with greater variability in tapping (Anderson et al., 2017; Laflamme et al., 2018; Meier, 2018; Seli, Cheyne, et al., 2013, 2015; Seli, Jonker, et al., 2015). This variability is called Rhythmic Response Time variance (RRTv), which is calculated as the natural-log transformed response variance in the five trials before a probe. While some moderator and mediator analyses have failed to replicate, the link between MW and RRTv has been successfully replicated independently of the original authors' numerous self-replications (Anderson et al., 2017; Meier, 2018). This task has been adapted in the MCT to measure task-engagement by having participants tap synchronously with a metronome beat (Figure 1).

Figure 1. Diagrammatic representation of study procedure. In the experiment, participants count in cycles of twenty by pressing the <down arrow> key synchronously with a metronome beat; after a complete cycle, participants begin again at 1. To correctly complete a cycle, participants indicate twenty by pressing the <left arrow> key (Correct cycle). Participants could incorrectly complete a cycle in three ways: indicating they reached twenty before or after twenty (Incorrect: Target), continuing long past twenty to the programmed limit of thirty presses (Incorrect: Limit), or indicating losing count using the <up arrow> key at any time (Incorrect: Self-catch). Self-Catching thus reflects spontaneous MA restoration whereas Target and Limit errors reflect loss of MA of the broader counting goal.

1.4 Hypotheses

The MCT incorporates the three paradigms outlined above and the present study aims to validate this new tool. Failure to comply with the broad intention of counting to twenty should reflect the loss of MA while variance in tapping should index momentary task-engagement. Performance decline at the focal behavioural level is interpreted as task-disengagement whereas failure in the higher-order counting task is interpreted as noncompliance with intention, i.e. the behavioural marker of MA loss. The following hypotheses were registered on the Open Science Framework (Anderson, 2019):

1.4.1 Task Focus

In the present study, there was one way to be on-task (i.e. correctly counting to twenty) and three ways to fail to be on-task (i.e. counting incorrectly without realizing, reaching an upper limit without realizing, or self-catching one's miscount). Based on previous findings from various iterations of the MRT (Anderson et al., 2017; Meier, 2018; 2013, 2015; Seli, Jonker, et al., 2015)—which consistently found lower RRTv for on-task reports versus MW reports of various types— we predicted that cycles wherein participants remained on-task (i.e. correctly counting) would have lower RRTv than cycles in which participants lost count at some point. We registered the following:

H1: We hypothesize that Correct cycles will have lower RRTv than Incorrect, Limit, and Self-Catch cycles.

We also registered that we would perform follow-up analyses exploring these different RRTv rates for each type of cycle, i.e. follow-up t-tests to explore differences between Correct, Incorrect, Limit, and Self-Catch cycles.

1.4.2 Performance Decrement

Performance and vigilance on tasks declines over time (Mackworth, 1964; Warm, Parasuraman, & Matthews, 2008). This finding was present in both the original MRT study (2013) and our replication (Anderson et al., 2017) such that higher RRTv was found in trials later in the experiment. We registered the following:

H2: We hypothesize that trials later in the experiment will have higher RRTv than trials earlier in the experiment.

1.4.3 Trait-level Control Variables

We registered two participant-level covariate hypotheses based on previous findings (Anderson et al., 2017). The MCT involves rhythmic tapping thus experience with rhythm-keeping was predicted to

precipitate superior performance, which was true of the MRT (Anderson et al., 2017). We predicted participants with musical experience, rhythm-game experience (e.g. Rock Band, Guitar Hero), and experience with tasks involving internal counting (e.g. counting reps in cross-fit) would have lower RRTv given their existing proficiencies. We registered the following hypothesis:

H3: We hypothesize that participants with more Timing Experience will have lower RRTv

As RRTv indexes attention we predicted that a trait-measure of attention-lapses, the Attention Related Cognitive Errors Scale (ARCES; (Cheyne, Carriere, & Smilek, 2006)), would be predictive of RRTv. This finding has been shown with the Sustained Attention to Response Task (SART; (Cheyne et al., 2006; Smilek, Carriere, & Cheyne, 2010)) and in our previous work with the MRT (Anderson et al., 2017). We registered the following hypothesis:

H4: We hypothesize that participants with higher Attention Related Cognitive Errors (ARCES) will have higher RRTv

1.4.4 Meta-Accuracy through Self-Reported Performance

We predicted that participants would have an accurate sense of their performance relative to their peers. This prediction was based on findings to this effect in our previous work (Anderson et al., 2017). Based on these findings, we registered the following hypothesis:

H5 We hypothesize that self-rated performance and Mean RRTv will negatively correlate, i.e. higher self-rated performance means lower Mean RRTv, i.e. participants will be able to rate their performance accurately

2. Methods

2.1 Participants

This experiment was conducted under informed consent in accordance with the Declaration of Helsinki under the Social Sciences, Humanities and Education Research Ethics Board at the University of Toronto, Mississauga campus. Undergraduates from the University of Toronto Mississauga campus participated in exchange for course-credit. Participants were recruited for the January to April 2019 term; the final sample was $N = 74$ after data exclusion. Participants were young-adult undergraduates (mean age: 19, SD: 0.72, range 18–22; 23 male).

As registered, participants were removed for missing 10% or more trials ($n = 0$) or reporting disengaging from the experiment (e.g. using their cellphone or removing the headphones, $n = 12$). Hard-drive failure resulted in loss of task-data for 77 participants; this was due to a fault in one of the testing computers such that task-data was never saved to the hard-drive. One participant asked that their data be removed.

As registered, we use an alpha of 0.01 for good evidence against the null and the common alpha of 0.05 for modest evidence against the null. The final sample size is consistent with prior work using the MRT (Laflamme et al., 2018; Seli, Cheyne, et al., 2013, 2015; Seli, Jonker, et al., 2015).

2. 2 Procedure

2.2.1 Metronome Count Task (MCT)

Instructions used in the MCT are available on the OSF. In the MCT, participants tap along to a steady aural beat (Seli, Cheyne, et al., 2013), as in Seli et al.'s (2013) MRT. Trials are 575 ms of silence, then a metronome tone lasting approximately 75 ms, then another 650 ms of silence, resulting in a total trial duration of 1300 ms. Participants were instructed to use the arrow-keys to tap synchronously with the

onset of the tone. Participants completed 30 practice trials before completing 900 experimental trials in two blocks of 450 trials (~10 minutes each) with a short break between.

In addition to tapping along to the beat, participants were instructed to count the beats. Participants were instructed to press the <down arrow> key in synchrony with the metronome and to press the <right arrow> every twentieth beat, also in synchrony with the metronome. Participants were also instructed that if they noticed they lost count, they were to press the <up arrow> to reset the count at 1. If participants hit the 30th beat without indicating they passed the 20th beat and without self-catching, the program alerted them and reset the count. These instructions were framed for participants as manipulating points they would gain and lose during the experiment; these points were never shown to participants and did not alter any part of the procedure (see full instructions for details).

2.2.2 Cycle Type

A “Cycle” is defined as a sequential series of trials counting from 1 to any of four possible end-points (Figure 1). There is one correct end-point: correctly identifying the twentieth beat. There are three incorrect end-points: incorrectly indicating a beat as the twentieth beat, self-catching lost counts, and hitting the programmed limit (i.e. the thirtieth beat).

Following incorrect cycles, participants were asked to indicate which of the following best characterized why they lost count: (1) Mind-Wandering: I got distracted by my own internal thoughts, (2) Environmental: I got distracted by something in my environment, (3) Intentional: I was not actually trying this time (Be honest! This answer is legitimate!), (4) Accidental: I was counting properly but my finger accidentally hit the wrong key, (5) Unclear: I thought I was counting correctly. I was not aware of any mind-wandering or other distractions, (6) Other Reason: (You can elaborate after the experiment). These questions were included for exploratory and task-development purposes and are not reported here.

2.2.3 Disengagement and Exclusion

After the MCT, participants indicated whether they disengaged from the experiment. As registered, this question was used to exclude participants if they disengaged in a way that would invalidate their responses. Participants indicating any of the following were removed: Reading or sending text messages, Reading or sending emails, Taking or making phone calls, Removing the headphones, I didn't understand the instructions. Participants also indicated instruction clarity ("How clear did you find the instructions?") on a 0–100 slider scale and were excluded if they indicated under 50 on this scale, i.e. were treated as if they indicated "I didn't understand the instructions". Participants indicating "Other" were screened based on their supplied disengagement.

Participants indicating the following were not removed based on these responses: Checking my phone, Stopped trying (too tired), Stopped trying (too boring), Mind Wandering, I was distracted by something in the environment, None: I did not disengage.

2.2.4 Control Questions

After the disengagement question, participants completed the Attention Related Cognitive Errors scale (Cheyne et al., 2006) on 0–100 slider scales with nominal end-point anchors of "Never" and "Very Often". They also indicated their level of experience with musical instruments, rhythm games, and other internal counting (Do you have any experience with musical instruments (e.g. drums, bass guitar, piano)? Do you have any experience with Dance Dance Revolution, Rock Band, Guitar Hero, or other rhythm-based games? Do you have any experience with activities that involve a great deal of internal counting (e.g. counting reps in cross-fit)?), which were summed according to this rubric: No, none = 0, Yes, a little = 1, Yes, some = 2, Yes, a lot = 3 (total range 0–9).

2.2.5 Meta-Accuracy through Self-Rated Performance

Participants rated their own Self-rated Performance (“How well do you think you did on the task compared to other people?”) on a 0–100 slider scale under an image of a standard normal distribution.

2.2.6 Additional Questions

After completing the MCT participants indicated their Motivation (“How motivated were you to do well on the task?”) on 0–100 slider scale. Participants also indicated their primary motivator from the following list: Course credit, Hitting timing, Keeping count, Finishing fast, Other, I was not motivated. Finally, participants reported demographic variables and were allowed to give open-ended feedback about the MCT and experiment as a whole. These questions were included for exploratory and task-development purposes and are not reported here.

2.3 Deviations from Registered Hypotheses

The registration did not indicate that the first cycle and first cycle following the break would be removed. This has been done as the first cycle has trials in which RRTvs cannot be calculated. The last cycle and last cycle preceding the break have also been removed if the cycle was incomplete. That is, the break occurred on the 450th trial and the end on the 900th, which were not always cycle termini. For example, if trial 450 coincided with a count of 15 in the cycle, the cycle would be incomplete and could not be classified as correct, incorrect, self-caught, or limit-caught. These incomplete cycles were removed from analyses.

2.4 Data Analysis

The R Language (R Core Team, 2014) was used for statistical analyses using the nlme package for modelling (Pinheiro, Bates, DebRoy, Sarkar, & R Core Team, 2018), the housekeep package for reporting (Lin, 2019), and the ggplot2 package for data visualization (Wickham, 2016, p. 2). As

registered, we use an alpha of 0.01 for good evidence against the null and alpha of 0.05 for modest evidence against the null.

3. Results

3.1 Task Focus, Performance Decrement, and Trait-level Covariates

To test our hypotheses we regressed RRTv onto the standardized within-participant fixed-effects of cycle type (correct vs incorrect), trial (number from 1 to 900), the trial by cycle interaction, and the between-participants fixed-effects of timing experience and ARCES scores. Participant intercepts were modelled as random effects.

Supporting the hypothesized effect of task-focus, there was a significant main effect of cycle type (H1: $b = 0.23$, $SE = 0.01$, $t(61067) = 16.70$, $p < .001$, $r = 0.07$) such that incorrect-ending cycles had higher RRTv than correct cycles. There was also support for the hypothesized performance decrement, i.e. there was a significant main effect of trial (H2: $b = 0.04$, $SE = 0.005$, $t(61067) = 6.85$, $p < .001$, $r = 0.03$), and there was support for the predicted effect of timing experience (H3: $b = -0.27$, $SE = 0.06$, $t(71) = -4.58$, $p < .001$, $r = 0.48$). The hypothesized effect of ARCES scores was non-significant (H4: $b = -0.08$, $SE = 0.06$, $t(71) = -1.34$, $p = .183$, $r = 0.16$). The trial by cycle interaction was also non-significant ($b = 0.005$, $SE = 0.01$, $t(61067) = 0.37$, $p = .708$, $r = 0.002$).

3.1.1 Performance Decrement within and across Blocks - Exploratory Follow-up

The lack of interaction between trial and cycle was unexpected. As an exploration, we considered whether this may be because there was a break in the middle of the MCT, which differs from the MRT. To test this possibility, we separated trials into pre- and post-break blocks and re-ran the above analyses with the fixed-effect of block (see Figure 2).

The significant main effects in the registered analyses remained (cycle type: $b = 0.23$, $SE = 0.01$, $t(61066) = 16.49$, $p < .001$, $r = 0.07$; trial: $b = 0.04$, $SE = 0.005$, $t(61066) = 8.17$, $p < .001$, $r = 0.03$; timing experience: $b = -0.27$, $SE = 0.06$, $t(71) = -4.57$, $p < .001$, $r = 0.48$). The non-significant ARCES results also remained ($b = -0.08$, $SE = 0.06$, $t(71) = -1.35$, $p = .182$, $r = 0.16$). Under this analysis, however, the trial by cycle interaction was significant ($b = 0.07$, $SE = 0.01$, $t(61066) = 5.21$, $p < .001$, $r = 0.02$) as was the main effect of block ($b = 0.02$, $SE = 0.009$, $t(61066) = 2.61$, $p = .009$, $r = 0.01$). A model comparison between the previous model and this model shows that this exploratory model is superior ($\chi^2(1) = 98.7$, $p < 0.0001$). The correlation between blocks also shows excellent reliability in RRTv ($r(72) = 0.88$, $p < .001$).

Figure 2. Linear regressions showing main effects of cycle, block, and trial. Shaded area reflects 95% confidence interval around fitted model.

3.2 Meta-Accuracy through Self-Reported Performance

To test our registered hypothesis that self-rated performance and Mean RRTv would negatively correlate we conducted a participant-level Pearson's R correlation (one-sided, as registered) with grand-mean centred Self-Rated Performance and Mean RRTv. Counter to H5, this correlation was non-significant ($r(72) = -0.01$, $p = .450$, one-sided).

3.3 The Four Types of Cycles - Registered Follow-up

Correct cycles made up 83.0% (SD: 14.9%) of responses on average. Incorrectly indicating a target on a non-target trial was the most common way to be incorrect (M: 10.9%, SD: 10.9%) and hitting the limit was the least common (M: 1.5%, SD: 2.6%). Self-catching was relatively rare (M: 3.2%, SD: 4.0%). After an incorrect cycle-end, participants indicated why they lost count: the most common reasons were mind-wandering (43.8%), unclear (26.5%, i.e. "I thought I was counting correctly. I was not aware of any mind-wandering or other distractions"), and accidental (18.3%). Uncommon reasons were environmental distraction (7.0%), intentional disengagement (2.2%), and other reasons (2.2%). Four participants had 100% correct performance, 11 never indicated an incorrect target, 50 never hit the limit, and 28 never self-caught.

We explored within-participant scaled RRTv split by the four cycle types (Figure 3). Exploratory t-tests show that incorrect (target) ($M=0.13$, $SD=1.08$) cycles had higher RRTv than both correct ($M=-0.03$, $SD=0.96$, $p_{\text{adj}} < .0001$) and incorrect (limit) ($M=-0.02$, $SD=1.06$, $p_{\text{adj}} = .0001$) cycles. There was no difference between correct and incorrect (limit) cycles ($p = 0.97$). Self-caught miscounts ($M=0.2$, $SD=1.13$) had higher RRTv than all other cycles (correct: $p_{\text{adj}} < .0001$, incorrect (limit): $p_{\text{adj}} < .0001$, incorrect (target): $p_{\text{adj}} = .025$), which may be because self-catching was not necessarily synchronous with the metronome beat. Upon exploratory reanalysis using only the down-arrow button presses, i.e. dropping off-beat self-catch trials and target-indication trials, presses, i.e. dropping off-beat self-catch trials and target-indication trials, the pattern of results was identical, but less extreme, especially for self-caught; the only significance-change was that self-catch was no longer statistically higher than incorrect (target).

Figure 3. Within-participant mean RRTvs for each type of cycle (error bars indicate bootstrapped 95% CI).

3.4 Performance Divergence - Exploratory Analysis

Given that RRTv differs as a function of cycle type, we explored whether divergence occurred at a specific time-point within a cycle via one-tailed t-tests and an alpha of 0.01 (Figure 4 and Supplemental Materials). Cycles that eventually end in incorrect responses were divergent during the first five trials of the cycle (all p s ≤ 0.002). Differences were then non-significant until the 15th trial, after which all differences were significant (all p s ≤ 0.004). Worth noting is that calculating the RRTv in the first five trials of each cycle necessarily includes trials a previous cycle as RRTv is based on the preceding 5 trials; calculating RRTv for trials 15–20 includes only trials from the present cycle.

Figure 4. Within-participant means of RRTvs for trials counting up to 20 for cycles ending as Correct (blue) and Incorrect (red) trials (error bars indicate bootstrapped 95% CI). Note the observable divergence on trials 15–20, which are statistically significant (see Supplemental Materials).

4. Discussion

The Metronome Counting Task (MCT) is a continuous performance task intended for measuring meta-awareness based on three extant paradigms: the self-caught paradigm for indicating spontaneous MA restoration, focal versus broad task monitoring for objectively indicating MA loss, and response variability for dynamically tracking task-engagement. First, the MCT allows participants to report losing track of the experimenter-defined task's broader intention (i.e. losing count when counting to twenty). These self-caught events embedded within the context of the ongoing focal task (i.e. synchronously tapping the metronome beat) provide a behaviour-phenomenology link from temporally relevant task-based indices of MW to participant-indicated moments of MA restoration. Second, the MCT implements both focal and broad task monitoring to maintain objective and temporally relevant behavioural metrics of focal task-engagement and MA loss through noncompliance with the broader intention. Focal monitoring of trial-by-trial key-presses indexes task-engagement on the timescale of trial duration (1300ms) whereas the broader intention allows for objective classification of task performance on timescales that ensure time is not wasted (20–30 trials, 26–39 seconds), especially useful during expensive neuroimaging. Third, the MCT's RRTv measure continuously indexes task-engagement and thus allows investigation of dynamic changes in attention on the above timescales. Dynamic performance measures are critical for aligning to naturally dynamic neuroimaging data, establishing the link between behaviour and neural activity. Additionally, because of the nature of the task, we know that differences are based in attention or other factors, not in skill (i.e. everyone can count to twenty). Together, the present study validates the MCT as successfully integrating multiple MA measurement constructs into a single efficient tool.

To confirm the utility of the MCT we tested a number of registered hypotheses. First, our findings supported the critical hypothesis that RRTv would differ between correct and incorrect cycles. This central finding instills confidence that RRTv in the MCT tracks task-engagement as it does in the MRT. Exploratory analyses revealed that performance divergence reached statistical significance as early as count 15 of 20, suggesting that MA loss could be objectively detectable in advance of participants'

subjective MA restoration and, critically, in advance of task errors. If these divergence patterns remain consistent in confirmatory studies they will allow us to probe the neural mechanisms of MA loss by using this objective behavioural index to event-lock neuroimaging data.

As predicted, we found the classic vigilance decrement affected the MCT as performance declined over time in the experiment. Exploratory follow-up suggests that performance declines both within and across blocks, consistent with Ralph et al. (Ralph, Onderwater, Thomson, & Smilek, 2017). The relationship between trial and block effects is currently being investigated in an extended 40 minute version of the MCT. Future research can examine whether real-world performance or health indicators can be predicted by block-wide vigilance decrements in the MCT. Block-wide performance may be developed into a measure of sustained attention and self-monitoring over a longer period (10 minutes), which may prove useful to provide objective metrics for clinical assessments involving attention-related disorders, e.g. ADHD.

A trait-level prediction relating performance to timing experience was corroborated and the hypothesized predictive capacity of ARCES scores was not. The lack of significant predictive value of ARCES may be due to this study being underpowered to detect the very small effect that was found for the MRT in a multisite sample of 300 participants (Anderson et al., 2017). If this is the case, the effect may be small enough that it is of no great interest in the context of the larger effects found here. Given the large effect size of timing experience in this task, and given that it was one of the larger effects when modelling the MRT (Anderson et al., 2017), we recommend this control variable be measured and refined. Finer grained response options than the limited 4-point ordinal scale used here may provide a higher signal-to-noise ratio in future studies.

Counter to our hypothesis participants did not have an accurate sense of how they performed relative to their peers. We can hypothesize that the MCT is a novel enough task that participants do not have a well-informed sense of how they or most people would perform. Counting to twenty is trivially

easy for any healthy adult so, in principle, everyone should perform with perfect accuracy. Only four participants performed the task without error, however, and failures in the MCT often came as a surprise: nearly half of all reasons given for incorrect responses indicated that participants were unaware of any distraction or that the mistake was accidental, indicative of their lack of MA. Unsurprisingly, the remaining mistakes were mostly attributed to MW.

5. Constraints on Generality

Participants were young-adult undergraduates with diverse ethnic and religious backgrounds. Age effects may emerge if the MCT is tested with older or younger participants. This study was performed with healthy normals; performance may interact with psychopathologies involving meta-awareness or attention-related deficits (Goodkind et al., 2015; Menon, 2011). Conversely, individuals with advanced attention or rhythm training, (e.g. expert meditators or musicians) may demonstrate altered MA profiles. This study rewarded participants with course credit; results may differ if the MCT is used with other incentives (e.g. financial), especially if rewards are performance-based.

The present findings are particular to a broad intention of counting to twenty, a number selected to allow for MA failures while still maximizing the number of cycles observed. Other counting goals would likely result in different rates of correct and incorrect cycles and different performance divergence patterns. With changing counting goals, we would still predict a performance decrement and correct cycles would still be predicted to have lower RRTv. Different trial durations may also affect the proportion of correct vs. incorrect cycles as broad intentions must be maintained in working memory over the cycle period, but RRTv patterns should remain consistent unless delays between metronome beats reach extreme values.

Calibration may be needed to overcome processing and input-delay so participants' tapping is subjectively aligned with metronome onset; without proper calibration, participants may learn to press before they subjectively hear the metronome to get the best objective synchrony (i.e. anticipatory

pressing). We have no reason to believe that the results depend on other characteristics of the participants, materials, or context. This Constraints on Generality (COG) statement is made in accordance with Simons, Shoda, & Lindsay (2017).

6. Future Directions

The MCT sets task-engagement and meta-awareness on continua and measures focal behavioural performance simultaneously to compliance with broader intentions. The result is a task wherein MA is frequently needed to keep the broader intention in mind: if attention strays, the count is quickly lost. It may be that participants devote resources to tasks according to the priority of their strategic performance goals. Maximally prioritizing focal performance implies immediate task-engagement, devoting resources to accurately synchronizing with the metronome. However, devoting too many resources to focal accuracy could compromise broader MA of counting to twenty. Conversely, too little attention to the focal goal in favour of the broad intention may result in poorer focal performance. Meanwhile, as time in the MCT progresses, focal task-performance decreases (the expected performance decrement (Warm et al., 2008)).

Despite the success of the MCT, MRT, and related tasks, using MW as a proxy for MA presents major unresolved issues, especially as concerns “intentional mind-wandering” and “mind-wandering with awareness” (Seli, Carriere, Levene, & Smilek, 2013; Seli et al., 2017; Seli, Risko, & Smilek, 2016; Seli, Risko, Smilek, et al., 2016). Indeed, participants often consciously decide to disengage from experimenter-defined tasks, preferring instead to engage task-unrelated thoughts, fully meta-aware that they are doing so (Anderson et al., 2017; Seli et al., 2017). In light of these and other dimensions of MW (e.g. guided vs. unguided thoughts), MW is not necessarily best thought of as a unitary construct (Christoff, Irving, Fox, Spreng, & Andrews-Hanna, 2016; Seli, Kane, et al., 2018). Indeed, one could think of MW as a task in its own right, not merely a proxy for disengagement from experimenter-defined tasks. Viewing MW as a participant-defined task may aide in framing how both MW and MA relate to the

sustained attention cycle, possibly providing insights that could eventually overturn the very idea that MW represents a failure of sustained attention. Indeed, it may be that dynamically shifting priorities inform which tasks most meet the currently motivated goals of participants. For example, a participant may disengage from an experimenter-defined task of subjectively low value in order to engage in a participant-defined task of greater value; this disengagement could be seen not as a “failure”, but as an optimization of the participant’s goals. Such a framework of attention would suggest a more fluid goal-oriented resource-allocation process that would closely link motivation, attention, and behaviour. This framework could contrast against models of attention based on executive-control failures or on notions of resource scarcity and could be seen as complimentary to the natural fluctuations hypothesis (Seli, Smilek, et al., 2018).

Focal versus broad goal-prioritization has plausible neurological foundations (Pezzulo, Rigoli, & Friston, 2018) and it seems likely that focal tapping and broader counting involve both shared and non-overlapping neural processes (Grahn & Manly, 2012). Performance deterioration likely has its own biomarkers, as demonstrated by existing fMRI evidence of performance decrement related neural modulation (Hilti et al., 2013; Langner & Eickhoff, 2013) and EEG evidence of the same (Reteig, van den Brink, Prinssen, Cohen, & Slagter, 2019). There may be a third neural process that closely correlates with the subjective state of meta-awareness, with the interplay between the dorsal anterior cingulate cortex and locus coeruleus being a major candidate (Melnychuk et al., 2018; Mittner, Hawkins, Boekel, & Forstmann, 2016). By measuring variation in focal behaviours under the constraints of broader intentions, the MCT may aid in modelling the neural processes that lie at the heart of MA, answering the call for further investigation into possible general consciousness-monitoring systems in the brain (Seli, Smilek, et al., 2018).

It could also be that the brain manifests two MA systems analogous to two input/output (I/O) methods in computer science: polling and interrupts. Generally, I/O devices operate much more slowly than internal processing (e.g. a hard-disk drive reads and writes data orders of magnitude slower than a

CPU processes instructions) and an analogy could be made to the human body: one can think faster than one can move one's arm. These processing speed differences would make it extremely inefficient for the internal system to wait for the I/O device to finish its operations before the internal system were allowed to process more information. Two ways of dealing with this processing difference are polling and interrupts. Polling is a system by which an internal system actively monitors an I/O device by regularly checking whether the I/O device is ready for operation; the human analogy would be that there may be a system that periodically "checks" the contents of consciousness to appraise whether behaviours are aligned with intentions. Such a system could plausibly be provided by phasic bursts of locus coeruleus activity or some function of the claustrum (Torgerson, Irimia, Goh, & Van Horn, 2015). In contrast, interrupts are a system by which the I/O device, when ready for operation, signals the internal system to stop its current operations and handle the I/O device's request. The human analogy of interrupts could plausibly be detection of predictive-coding errors processed by the dorsal anterior cingulate cortex, which then feed into the locus coeruleus and produce an event-based burst of activity. This two system model could explain the dual experiences of MA in its wavering, cyclic manifestation—as in the natural fluctuations hypothesis (Seli, Smilek, et al., 2018) and the sustained attention cycle model (Hasenkamp et al., 2012)—and in its immediate, salience-based manifestation during surprise and behavioural or psychological error-detection—as in the cocktail party effect (Cherry, 1953) or when stubbing one's toe. This conjecture is speculative and based on analogy thus more research is needed to test the hypothesized neural activation patterns predicted by such a dual model of MA.

To further probe mechanistic questions and search for the neural markers of MA, the MCT is being deployed using functional near-infrared spectroscopy (fNIRS), electroencephalography (EEG), and functional magnetic resonance imaging (fMRI). Longer versions of the MCT will also be validated to uncover whether block effects replicate and extend to longer experiment durations. We propose that these non-dichotomous measures of meta-awareness are needed to more completely understand the subtleties of sustained attention and the instability of task-engagement.

7. Conclusion

The MCT appears to be a useful tool to help objectively quantify meta-awareness. Response-variability, which has previously been linked with phenomenological reports in a structured and replicable manner, was here linked with objective markers of MA loss through behavioural non-compliance with a broader intention. The MCT is easy for participants to understand and participants perform adequately. The continuous performance indices afforded by the MCT offers insight into task-engagement on small time-scales and will be useful for maximizing data usability in neuroimaging as every moment returns meaningful task-data. By continuing to build on the MCT paradigm, more sensitive models of meta-awareness, task-engagement, and sustained attention may soon emerge.

8. Funding

This work was supported by the Mind & Life Institute's Varela Grant program [grant number A-84847481] and the Canada Graduate Doctoral Scholarships program the Natural Sciences and Engineering Research Council of Canada (NSERC) [grant number PGSD2 - 518591].

9. Data Availability Statement

The datasets generated and analysed for the current study are available on the Open Science Framework (Anderson, 2019).

References

- Anderson, T. (2019). *Metronome Counting Task (MCT)*. Retrieved from <https://osf.io/69q5s>
- Anderson, T., & Farb, N. A. S. (2018). Personalising Practice Using Preferences for Meditation Anchor Modality. *Frontiers in Psychology*, 9. <https://doi.org/10.3389/fpsyg.2018.02521>
- Anderson, T., Lin, H., & Petranker, R. (2017, May 22). Replication of Seli et al 2013 ‘Wandering Minds and Wavering Rhythms’. Retrieved 29 August 2018, from Open Science Framework website: <https://osf.io/m6gtw/>
- Cherry, E. C. (1953). Some Experiments on the Recognition of Speech, with One and with Two Ears. *The Journal of the Acoustical Society of America*, 25(5), 975–979. <https://doi.org/10.1121/1.1907229>
- Cheyne, J. A., Carriere, J. S. A., & Smilek, D. (2006). Absent-mindedness: Lapses of conscious awareness and everyday cognitive failures. *Consciousness and Cognition*, 15(3), 578–592. <https://doi.org/10.1016/j.concog.2005.11.009>
- Chin, J. M., & Schooler, J. W. (2010). Meta-Awareness. In *Encyclopedia of Consciousness* (pp. 33–41). Retrieved from <http://www.scopus.com/inward/record.url?eid=2-s2.0-79959874929&partnerID=40&md5=f171dc0f7e18fbf691c84d9ec95145c5>
- 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. <https://doi.org/10.1038/nrn.2016.113>
- Fortenbaugh, F. C., DeGutis, J., & Esterman, M. (2017). Recent theoretical, neural, and clinical advances in sustained attention research. *Annals of the New York Academy of Sciences*, 1396(1), 70–91. <https://doi.org/10.1111/nyas.13318>
- Goodkind, M., Eickhoff, S. B., Oathes, D. J., Jiang, Y., Chang, A., Jones-Hagata, L. B., ... Etkin, A. (2015). Identification of a Common Neurobiological Substrate for Mental Illness. *JAMA Psychiatry*, 72(4), 305–315. <https://doi.org/10.1001/jamapsychiatry.2014.2206>

- Grahn, J. A., & Manly, T. (2012). Common Neural Recruitment across Diverse Sustained Attention Tasks. *PLOS ONE*, 7(11), e49556. <https://doi.org/10.1371/journal.pone.0049556>
- Harms, I. M., van Dijken, J. H., Brookhuis, K. A., & de Waard, D. (2019). Walking Without Awareness. *Frontiers in Psychology*, 10. <https://doi.org/10.3389/fpsyg.2019.01846>
- Hasenkamp, W., Wilson-Mendenhall, C. D., Duncan, E., & Barsalou, L. W. (2012). Mind wandering and attention during focused meditation: A fine-grained temporal analysis of fluctuating cognitive states. *NeuroImage*, 59(1), 750–760. <https://doi.org/10.1016/j.neuroimage.2011.07.008>
- Hayes, S. C. (2004). Acceptance and commitment therapy, relational frame theory, and the third wave of behavioral and cognitive therapies. *Behavior Therapy*, 35(4), 639–665. [https://doi.org/10.1016/S0005-7894\(04\)80013-3](https://doi.org/10.1016/S0005-7894(04)80013-3)
- Hilti, C. C., Jann, K., Heinemann, D., Federspiel, A., Dierks, T., Seifritz, E., & Cattapan-Ludewig, K. (2013). Evidence for a cognitive control network for goal-directed attention in simple sustained attention. *Brain and Cognition*, 81(2), 193–202. <https://doi.org/10.1016/j.bandc.2012.10.013>
- Laflamme, P., Seli, P., & Smilek, D. (2018). Validating a visual version of the metronome response task. *Behavior Research Methods*, 1–12. <https://doi.org/10.3758/s13428-018-1020-0>
- Langner, R., & Eickhoff, S. B. (2013). Sustaining Attention to Simple Tasks: A Meta-Analytic Review of the Neural Mechanisms of Vigilant Attention. *Psychological Bulletin*, 139(4), 870–900. <https://doi.org/10.1037/a0030694>
- Levinson, D. B., Stoll, E. L., Kindy, S. D., Merry, H. L., & Davidson, R. J. (2014). A mind you can count on: Validating breath counting as a behavioral measure of mindfulness. *Frontiers in Psychology*, 5. <https://doi.org/10.3389/fpsyg.2014.01202>
- Lin, H. (2019). *housekeep*: Miscellaneous functions for research and housekeeping (Version 0.0.0.9001). <https://doi.org/10.5281/zenodo.2557034>

- Lutz, A., Jha, A. P., Dunne, J. D., & Saron, C. D. (2015). Investigating the phenomenological matrix of mindfulness-related practices from a neurocognitive perspective. *American Psychologist*, 70(7), 632–658. <http://dx.doi.org.myaccess.library.utoronto.ca/10.1037/a0039585>
- Mackworth, J. F. (1964). Performance decrement in vigilance, threshold, and high-speed perceptual motor tasks. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, 18(3), 209–223. <http://dx.doi.org.myaccess.library.utoronto.ca/10.1037/h0083302>
- Meier, M. E. (2018). Can research participants comment authoritatively on the validity of their self-reports of mind wandering and task engagement? A replication and extension of Seli, Jonker, Cheyne, Cortes, and Smilek (2015). *Journal of Experimental Psychology: Human Perception and Performance*. <https://dx.doi.org/10.1037/xhp0000556>
- Melnychuk, M. C., Dockree, P. M., O’Connell, R. G., Murphy, P. R., Balsters, J. H., & Robertson, I. H. (2018). Coupling of respiration and attention via the locus coeruleus: Effects of meditation and pranayama. *Psychophysiology*, 55(9), e13091. <https://doi.org/10.1111/psyp.13091>
- Menon, V. (2011). Large-scale brain networks and psychopathology: A unifying triple network model. *Trends in Cognitive Sciences*, 15(10), 483–506. <https://doi.org/10.1016/j.tics.2011.08.003>
- Mittner, M., Hawkins, G. E., Boekel, W., & Forstmann, B. U. (2016). A Neural Model of Mind Wandering. *Trends in Cognitive Sciences*, 20(8), 570–578. <https://doi.org/10.1016/j.tics.2016.06.004>
- Pezzulo, G., Rigoli, F., & Friston, K. J. (2018). Hierarchical Active Inference: A Theory of Motivated Control. *Trends in Cognitive Sciences*, 22(4), 294–306. <https://doi.org/10.1016/j.tics.2018.01.009>
- Pinheiro, J., Bates, D., DebRoy, S., Sarkar, D., & R Core Team. (2018). nlme: Linear and Nonlinear Mixed Effects Models (Version 3.1-137). Retrieved from <https://CRAN.R-project.org/package=nlme>
- R Core Team. (2014). *R: A Language and Environment for Statistical Computing*. Retrieved from <http://www.r-project.org/>

- Ralph, B. C. W., Onderwater, K., Thomson, D. R., & Smilek, D. (2017). Disrupting monotony while increasing demand: Benefits of rest and intervening tasks on vigilance. *Psychological Research*, 81(2), 432–444. <https://doi.org/10.1007/s00426-016-0752-7>
- Reteig, L. C., van den Brink, R. L., Prinssen, S., Cohen, M. X., & Slagter, H. A. (2019). Sustaining attention for a prolonged period of time increases temporal variability in cortical responses. *Cortex*. <https://doi.org/10.1016/j.cortex.2019.02.016>
- Seli, P., Carriere, J. S. A., Levene, M., & Smilek, D. (2013). How few and far between? Examining the effects of probe rate on self-reported mind wandering. *Frontiers in Psychology*, 4(JUL). <https://doi.org/10.3389/fpsyg.2013.00430>
- Seli, P., Cheyne, J. A., & Smilek, D. (2013). Wandering minds and wavering rhythms: Linking mind wandering and behavioral variability. *Journal of Experimental Psychology: Human Perception and Performance*, 39(1), 1–5. <https://doi.org/10.1037/a0030954>
- Seli, P., Cheyne, J. A., Xu, M., Purdon, C., & Smilek, D. (2015). Motivation, intentionality, and mind wandering: Implications for assessments of task-unrelated thought. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(5), 1417–1425. <http://dx.doi.org/10.1037/xlm0000116>
- Seli, P., Jonker, T. R., Cheyne, J. A., Cortes, K., & Smilek, D. (2015). Can research participants comment authoritatively on the validity of their self-reports of mind wandering and task engagement? *Journal of Experimental Psychology: Human Perception and Performance*, 41(3), 703–709. <https://doi.org/10.1037/xhp0000029>
- Seli, P., Jonker, T. R., Cheyne, J. A., & Smilek, D. (2013). Enhancing SART Validity by Statistically Controlling Speed-Accuracy Trade-Offs. *Frontiers in Psychology*, 4, 265. <https://doi.org/10.3389/fpsyg.2013.00265>

- Seli, P., Kane, M. J., Smallwood, J., Schacter, D. L., Maillet, D., Schooler, J. W., & Smilek, D. (2018). Mind-Wandering as a Natural Kind: A Family-Resemblances View. *Trends in Cognitive Sciences*, 22(6), 479–490. <https://doi.org/10.1016/j.tics.2018.03.010>
- Seli, P., Ralph, B. C. W., Risko, E. F., Schooler, J. W., Schacter, D. L., & Smilek, D. (2017). Intentionality and meta-awareness of mind wandering: Are they one and the same, or distinct dimensions? *Psychonomic Bulletin & Review*, 24(6), 1808–1818. <https://doi.org/10.3758/s13423-017-1249-0>
- Seli, P., Risko, E. F., & Smilek, D. (2016). On the Necessity of Distinguishing Between Unintentional and Intentional Mind Wandering. *Psychological Science*, 27(5), 685–691. <https://doi.org/10.1177/0956797616634068>
- Seli, P., Risko, E. F., Smilek, D., & Schacter, D. L. (2016). Mind-Wandering With and Without Intention. *Trends in Cognitive Sciences*, 20(8), 605–617. <https://doi.org/10.1016/j.tics.2016.05.010>
- Seli, P., Smilek, D., Ralph, B. C. W., & Schacter, D. L. (2018). The awakening of the attention: Evidence for a link between the monitoring of mind wandering and prospective goals. *Journal of Experimental Psychology. General*, 147(3), 431–443. <https://doi.org/10.1037/xge0000385>
- Simons, D. J., Shoda, Y., & Lindsay, D. S. (2017). Constraints on Generality (COG): A Proposed Addition to All Empirical Papers. *Perspectives on Psychological Science*, 12(6), 1123–1128. <https://doi.org/10.1177/1745691617708630>
- Smilek, D., Carriere, J. S. A., & Cheyne, J. A. (2010). Failures of sustained attention in life, lab, and brain: Ecological validity of the SART. *Neuropsychologia*, 48(9), 2564–2570. <https://doi.org/10.1016/j.neuropsychologia.2010.05.002>
- Torgerson, C. M., Irimia, A., Goh, S. Y. M., & Van Horn, J. D. (2015). The DTI connectivity of the human claustrum. *Human Brain Mapping*, 36(3), 827–838. <https://doi.org/10.1002/hbm.22667>
- Warm, J. S., Parasuraman, R., & Matthews, G. (2008). Vigilance Requires Hard Mental Work and Is Stressful. *Human Factors*, 50(3), 433–441. <https://doi.org/10.1518/001872008X312152>

Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Retrieved from <http://ggplot2.org>