

Pupil Diameter Tracks the Exploration–Exploitation Trade-off during Analogical Reasoning and Explains Individual Differences in Fluid Intelligence

Taylor R. Hayes and Alexander A. Petrov

Abstract

■ The ability to adaptively shift between exploration and exploitation control states is critical for optimizing behavioral performance. Converging evidence from primate electrophysiology and computational neural modeling has suggested that this ability may be mediated by the broad norepinephrine projections emanating from the locus coeruleus (LC) [Aston-Jones, G., & Cohen, J. D. An integrative theory of locus coeruleus-norepinephrine function: Adaptive gain and optimal performance. *Annual Review of Neuroscience*, 28, 403–450, 2005]. There is also evidence that pupil diameter covaries systematically with LC activity. Although imperfect and indirect, this link makes pupillometry a useful tool for studying the locus coeruleus norepinephrine system in humans and in high-level tasks. Here, we present a novel paradigm that examines how the pupillary response during exploration and exploitation covaries with individual differences in fluid intelligence during analogical reasoning on Raven's Advanced Progressive

Matrices. Pupillometry was used as a noninvasive proxy for LC activity, and concurrent think-aloud verbal protocols were used to identify exploratory and exploitative solution periods. This novel combination of pupillometry and verbal protocols from 40 participants revealed a decrease in pupil diameter during exploitation and an increase during exploration. The temporal dynamics of the pupillary response was characterized by a steep increase during the transition to exploratory periods, sustained dilation for many seconds afterward, and followed by gradual return to baseline. Moreover, the individual differences in the relative magnitude of pupillary dilation accounted for 16% of the variance in Advanced Progressive Matrices scores. Assuming that pupil diameter is a valid index of LC activity, these results establish promising preliminary connections between the literature on locus coeruleus norepinephrine-mediated cognitive control and the literature on analogical reasoning and fluid intelligence. ■

INTRODUCTION

The ability to adaptively regulate the balance between exploration and exploitation is critical for optimizing behavior in the diverse, dynamic environments we encounter on a daily basis. Despite the ubiquity of the exploration–exploitation trade-off and its broad importance in understanding executive control, the neural mechanisms involved are still not well understood (Cohen, McClure, & Yu, 2007; Berridge & Waterhouse, 2003). Recent animal and human studies suggest that the locus coeruleus norepinephrine (LC-NE) system may support the exploration–exploitation trade-off, but the work is limited by the use of low-level tasks (Jepma & Nieuwenhuis, 2011; Gilzenrat, Nieuwenhuis, Jepma, & Cohen, 2010; Aston-Jones & Cohen, 2005). Here, we present a unique paradigm that examines how the pupillary response during exploration and exploitation covaries with individual differences in fluid intelligence (Gf) by combining pupillometry and verbal protocol analysis during analogical reasoning on Raven's Advanced Progressive Matrices (APM; Raven, Raven, & Court, 1998). Expanding the study of the exploration–exploitation trade-

off to a high-level analogical reasoning task and employing an individual differences approach provided novel insights into the relationship between the exploration–exploitation trade-off, noradrenergic function, and individual differences in Gf.

Converging evidence suggests that the LC-NE system plays a central role in mediating the exploration–exploitation trade-off (Cohen et al., 2007; Aston-Jones & Cohen, 2005). Specifically, the LC-NE system is thought to monitor for unexpected uncertainty and actively mediate the shift between exploration and exploitation in response to reward history (Aston-Jones & Cohen, 2005). Much of the current theory of LC-NE function is based on monkey electrophysiological recordings (see Aston-Jones & Cohen, 2005, for a review) and computational models thereof (Brown, Gilzenrat, & Cohen, 2005; Usher, Cohen, Servan-Schreiber, Rajkowski, & Aston-Jones, 1999). The electrophysiological data suggested that locus coeruleus (LC) neurons in monkeys exhibit distinct firing patterns that lie along a continuum of task performance from offline to phasic to tonic modes.

Offline mode occurs when the animal is drowsy or actively sedated and is characterized by low levels of LC activity and poor task performance. In phasic mode, the

baseline firing rate remains low, but the LC neurons fire phasic bursts of activity synchronized to task-relevant events; the animals exhibit high task performance. Finally, the tonic mode is characterized by high baseline rates of LC firing, poor task performance, exploratory behaviors, and indiscriminate sensitivity to both task-related and task-unrelated stimuli. These electrophysiological findings have been incorporated into a broader theory of LC function in which the LC mediates the exploration–exploitation trade-off in response to online assessments of task utility (adaptive gain theory, AGT; Aston-Jones & Cohen, 2005). AGT postulates that the LC actively mediates the gain of cortical units through the release of norepinephrine to promote exploitation via phasic mode or exploration via tonic mode (Aston-Jones & Cohen, 2005). Although AGT provides an elegant account of the existing data, it is limited by a paucity of corroborative human studies to validate and test the theorized link between LC function and the exploration–exploitation trade-off.

One major obstacle to studying LC function in humans is identifying a noninvasive method for measuring LC activity. Recently, pupil diameter has emerged as one promising noninvasive measure for LC activity and is being increasingly employed for this purpose (e.g., Cheadle et al., 2014; Eldar, Cohen, & Niv, 2013; Jepma & Nieuwenhuis, 2011; Einhäuser, Koch, & Carter, 2010; Gilzenrat et al., 2010; Einhäuser, Stout, Koch, & Carter, 2008). Neuroimaging work has shown that the pupil diameter covaries with fMRI BOLD activity in the LC (Murphy, O’Connell, O’Sullivan, Robertson, & Balsters, 2014) and that both P3 ERP and pupil diameter are sensitive to LC-NE modes of task engagement (Cheadle et al., 2014; Murphy, Robertson, Balsters, & O’Connell, 2011). Converging evidence from electrophysiology (e.g., Rajkowski, Kubiak, & Aston-Jones, 1994) and pharmacology (e.g., Phillips, Szabadi, & Bradshaw, 2000; Koss, 1986) also suggests that pupil diameter correlates with LC activity in animals. The anatomical pathways linking the LC and the pupil are a topic of ongoing research but probably involve α_2 -adrenoreceptor-mediated inhibition of the parasympathetic Edinger–Westphal nucleus responsible for pupil constriction (Samuels & Szabadi, 2008).

Here, we extend the study of the exploration–exploitation trade-off and LC function by tracking real-time shifts in exploration and exploitation in a rich, temporally extended analogical reasoning task, Raven’s APM (Raven et al., 1998). The APM is a geometric analogy test with excellent psychometric properties (Brouwers, Van de Viver, & Van Hemert, 2009) that has been a popular and trusted instrument in psychology for 70 years (e.g., Hayes, Petrov, & Sederberg, 2011, 2015; Gray, Chabris, & Braver, 2003; Carpenter, Just, & Shell, 1990). Raven’s APM is an excellent environment to induce strong shifts between exploration and exploitation because it repeatedly places participants in an unfamiliar relational environment in which they must engage in relational foraging to attempt to construct and/or identify the correct answer.

Although the term “relational foraging” is not commonly used in the literature on relational reasoning and Gf, it is pertinent to geometric analogy problems such as Raven’s APM (Figure 1, left). The problem-solving process is a search of an abstract problem space and involves the formulation of hypotheses and subgoals that either succeed or fail (e.g., Taatgen, Huss, Dickison, & Anderson, 2008; Newell & Simon, 1976). There are deep structural parallels, although not strict isomorphism, with reinforcement learning concepts such as reward, punishment, exploration, and exploitation. Indeed, reinforcement learning algorithms are commonly used to learn the utilities of production rules in production systems (e.g., Taatgen, 2013; Anderson et al., 2004). These structural similarities provide connections with independently developed theories of LC-NE function.

Specifically in Raven’s APM, the relational foraging process consists of an extended series of goals and subgoals, in which hypothesized patterns are extracted from one part of the problem matrix and tested on others. If the hypothesized pattern generalizes to another part of the matrix, the current pattern receives reinforcement. If the test fails, then a new pattern hypothesis must be generated (Carpenter et al., 1990). This iterative strategy is used to extract all the relations contained in a given Raven item or as many relations as needed to narrow the number of possible responses. This process has been formalized in models of matrix reasoning (Lovett, Tomai, Forbus, & Usher, 2009; Carpenter et al., 1990).

Importantly, for our present purposes, the relational foraging process maps onto key aspects of AGT. Each Raven problem is a miniature environment with a precise definition of optimal performance (i.e., pinpointing the item that best completes the relational pattern), fluctuations in task utility over time that result from testing hypothesized patterns, and implicit reinforcement received when hypothesized patterns generalize or fail to generalize. Last but not least, the use of Raven’s APM instead of simpler reinforcement learning tasks is methodologically beneficial for a pupillometric study because it produces more extended periods of exploration and exploitation, which are better suited to the relatively low temporal resolution of the pupillary response.

In the current study, pupil diameter was recorded as an indirect proxy for LC activity on each trial. Periods of exploration and exploitation were identified with the aid of think-aloud verbal protocols (Ericsson & Simon, 1993) that were collected while the participants solved each Raven problem. The results revealed a decrease in pupillary response during exploitative periods and a significant increase during exploratory periods.

This pattern is consistent with prominent theories of LC-NE system function and provides the first evidence that this system may be involved in cognitive control of the exploration–exploitation trade-off during analogical reasoning. Moreover, the individual differences in the

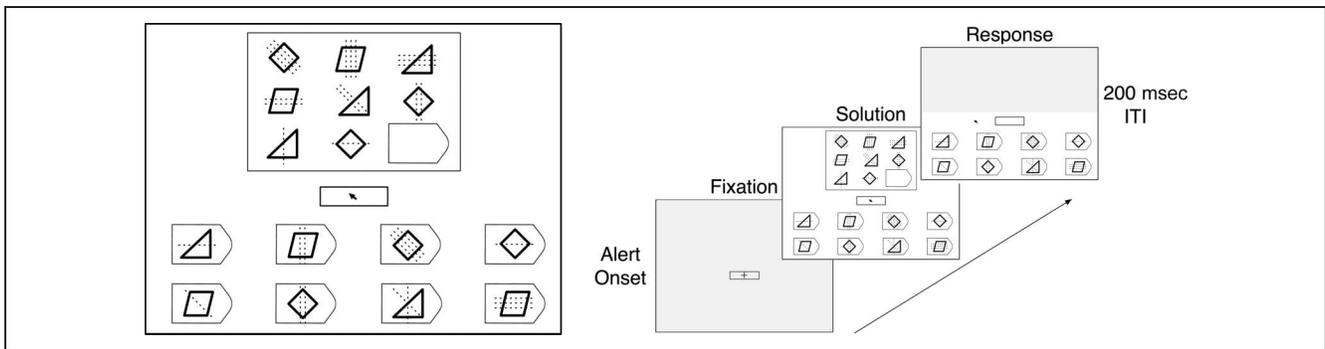


Figure 1. Raven problem format and trial sequence. (Left) The problem matrix and the eight response alternatives are shown with solid lines. The height of the rectangular box around the matrix subtended 9° of visual angle. This example item (generated by the authors) contains three relations that must be extracted: distribution of three shapes (diamond, triangle, parallelogram), distribution of three line orientations (0° , 45° , 90°), and decreasing line number down columns ($3 \rightarrow 2 \rightarrow 1$). (Right) Each trial had three phases: fixation, solution, and response. Participants fixated for 1 sec. Eye movements and concurrent think-aloud verbal protocols were collected during the solution phase. Moving the mouse cursor out of the fixation box triggered the response phase, during which the problem matrix was masked, and the participant clicked on their chosen answer. The intertrial interval (ITI) was 200 msec.

exploratory pupillary dilation could account for 16% of the variation in APM scores across participants.

METHODS

This study was conducted in a larger context of several related experiments that combined think-aloud verbal protocols and eye tracking to investigate the role of strategic cognitive control during visual relational reasoning on Raven's APM (Hayes, 2015; Hayes et al., 2011, 2015). These experiments involved multiple sessions and various manipulations from Session 2 onward, but the first session was always the same: To establish a common baseline, eye-tracking and think-aloud protocols were collected while the participants worked on 14 Raven problems as detailed below. This study is based exclusively on data from this common baseline session. Although other aspects of this large and multifaceted data set have been published elsewhere (Hayes et al., 2011, 2015), the pupillometric and verbal protocol aspects are reported here for the first time.

Participants

One hundred thirty-six students at the Ohio State University participated in the experiments outlined above (Hayes, 2015). They responded to recruitment flyers posted in the Ohio State University Psychology Building and were paid \$6 per hour plus \$1 bonus for each correct answer to a Raven problem. Sixteen participants did not consistently provide think-aloud protocols throughout each trial and were excluded from further consideration for the present purposes. Because of the labor-intensive nature of verbal protocol preprocessing, only 40 sessions' worth of verbal protocols were coded and analyzed. Thus, all results reported below are based on a random stratified sample of 40 participants (20 women and 20 men).

The distribution of Gf scores in the large sample ($N = 120$) was partitioned into four ability groups as follows: high (APM scores of 13–14), medium-high (scores of 11–12), medium-low (scores of 8–10), and low (scores of ≤ 7) abilities. Ten participants were then drawn at random from each ability group, and the verbal protocols and pupillometric data from their first (baseline) session were processed.

Stimuli

The participants completed a short-form test from Raven's APM Set II (Raven et al., 1998). Participants either completed Items 2, 4, 6, 9, 10, 11, 16, 17, 19, 21, 23, 24, 26, and 29 or Items 1, 3, 5, 7, 12, 13, 14, 15, 18, 20, 22, 25, 27, and 28. The participant instructions followed the *Raven APM Manual* guidelines for untimed individual test administration (Raven et al., 1998). The two 14-item subsets of the complete (36-item) APM were chosen to be approximately matched for difficulty on the basis of their psychometric characteristics published in the manual (Raven et al., 1998). There were no statistically significant differences in the respective distributions of scores in our sample (Hayes, 2015).

Apparatus

The Raven items were presented on a 21-in. NEC AccuSync 120 color CRT using Experiment Builder (SR Research, Mississauga, Canada). Participants viewed the items binocularly from a chin-and-forehead rest located 935 mm away. Gaze position and pupil response data were recorded from the left eye using an EyeLink 1000 desktop eye tracker (SR Research) at a sampling rate of 250 Hz. The experimental room had a constant ambient illuminance with 25 lux incident at participants' eyes to control for the pupillary light reflex. Image analysis of the Raven APM items revealed high luminance consistency across the 28 Raven test items (grayscale intensity: $M = 0.96$, $SD = 0.02$) and

across the individual matrix and response cells within each item (grayscale intensity: $M = 0.90$, $SD = 0.04$). Therefore, we did not alter the luminance properties of the original Raven APM test images, preserving their original psychometric properties.

Verbal protocols were recorded for each Raven item using a Shure Beta 58A supercardioid dynamic microphone (Shure, Inc., Niles, IL) and E-MU 0202 audio interface (E-MU, Scotts Valley, CA) controlled via Experiment Builder (SR Research). The microphone was placed close to the participant's mouth (≈ 5 cm) using a telescoping boom tripod microphone stand to provide clear audio recordings. The concurrent think-aloud verbal protocols were collected according to standard think-aloud procedures (Ericsson & Simon, 1993). After the participants received instructions on thinking aloud, they practiced it on unrelated items such as multiplication problems until the experimenter was confident they understood the instructions.

Procedure

Before the study, participants completed the EyeLink 1000 9-point calibration procedure. Each Raven item was preceded by a beep and fixation cross (similar to the EyeLink 1000 9-point calibration procedure) that appeared in the middle of the screen (Figure 1, right). The fixation screen was equal in luminance to the subsequent Raven item to avoid luminance changes at stimulus onset. After the participant fixated for 1 sec, which allowed for equipment recalibration, the Raven problem appeared, and the participant had unlimited time to work on it. Once participants had chosen an answer, they used the mouse to click on one of the eight responses, thereby ending the trial. Moving the mouse out of the fixation box triggered an isoluminant mask to be drawn over the problem matrix, which delineated solution and response phases. No accuracy feedback was provided until the very end of the experimental session to avoid feedback-induced pupillary dilations. Accuracy and solution time data were collected for each trial. Accuracy was defined as the total number of Raven items answered correctly, and solution time was measured from stimulus onset until response selection. Pupil diameter and gaze position were recorded throughout (i.e., from pretrial fixation through the end of the inter-trial interval).

Pupil Data Preprocessing

Before analysis, the pupillary data were corrected for blink artifacts and pupil foreshortening error. Following standard procedures, pupillary measurements were first filtered for blink artifacts, linearly interpolated, and then smoothed for measurement noise (Klingner, 2010; Beatty & Lucero-Wagoner, 2000). In addition, the pupil data were corrected to account for pupil foreshortening error—the systematic foreshortening of the pupil image as the eye

rotates away from the eye-tracking camera (Hayes & Petrov, 2015). Pupil foreshortening error must be corrected before analysis because solving Raven items requires the participants to freely scan the screen. The pupil foreshortening error correction described in detail in Hayes and Petrov (2015) fits a geometric model that expresses the pupil foreshortening as a function of the cosine of the angle between the eye-to-camera axis and the eye-to-stimulus axis. In calibration studies with artificial eyes with known, fixed pupil diameters (Hayes & Petrov, 2015), the geometric correction successfully reduced the root mean squared error in pupil diameter estimates by 97.5% when the model parameters were optimized to fit the empirical error surface. The calibration results strongly indicated that the pupil foreshortening error is invariant across changes in pupil size and systematically varies as a function of the orientation of the eye with respect to the camera. The results corresponded well with previous empirical measurements of pupil foreshortening error in biological human eyes (Mathur, Gehrmann, & Atchison, 2013; Jennings & Charman, 1978; Jay, 1962; Spring & Stiles, 1948). Together, these findings suggest the geometric correction can be used to virtually eliminate pupil foreshortening error. In this study, we first performed artifact correction to measure the apparent pupil diameter and then applied the optimized geometric model correction from Hayes and Petrov (2015) to estimate the true pupil diameter.

Verbal Protocol Coding

Concurrent think-aloud verbal protocols were used to segment each trial into exploration and exploitation solution periods. A broad coding scheme was developed to assist the coder in identifying exploration and exploitation periods during each Raven trial. Exploration periods were indicated by utterances that described isolated Raven image features (e.g., “Alright, it looks like we have a bunch of circles and squares....”) or expressed uncertainty (“Not sure what is going on here...I don't see any patterns yet.”). Exploitation periods were indicated by utterances that described a specific pattern within the Raven item (e.g., “In each line it looks like we have circle and diamond and square.... So on this we have square... circle...So the bottom should be a diamond.”). Many transitions from exploration to exploitation were signaled by insight language (“Oh, I see!”). Early Raven items that were easier to solve often only contained one exploration-to-exploitation shift, whereas later more difficult Raven items contained multiple transitions between exploring and exploiting. On these more difficult items, subsequent transitions from exploiting back to exploring were preceded by participants realizing either that a pattern extracted on one row of the Raven problem matrix did not generalize to the subsequent rows or that there was no response option that matched the final solution they had in mind. In these cases, the return to exploration

was signaled by failure utterances (e.g., “But that doesn’t match the second row” or “ok, it doesn’t look like that is even one of the possible options”) followed by a transition back to uncertainty utterances and/or isolated feature descriptions.

A semiautomated coding routine was developed in MATLAB (The MathWorks, Natick, MA) and was used to code all verbal protocol data. In this routine, for each trial, the human expert coder would be presented with an image of the relevant APM item while the recorded verbal protocol audio was played back in real time. The coder served as an “exploration detector” pressing one key to indicate the beginning of an exploratory period and another key to indicate the end of an exploratory period and the beginning of an exploitative period. The beginning of a trial was coded as neutral before any key presses. The MATLAB routine would then convert the time-stamped key presses into a code stream that contained the neutral (0), exploratory (+1), and exploitative (−1) codes for that verbal protocol, each sampled at 250 Hz. This procedure was completed for all participants ($n = 40$) and trials ($n = 14$), resulting in 560 individual protocol code streams.

Recent studies that have examined the effect of thinking aloud relative to silent control conditions have not shown any pupillary effect of vocalization (Hertzum & Holmegaard, 2013; Kammerer & Gerjets, 2013). Therefore, no distinction was made in the verbal protocol coding between periods of vocalization and gaps in vocalization.

All coding was performed by the first author (T. R. H.). He did not have access to any pupillometric data while he was coding the verbal protocols. Coder reliability was assessed by coding the data from five randomly sampled participants twice. The recoding was done approximately 1 full year after the original coding. The intrarater reliability for T. R. H. across the two coding sessions was high (mean % agreement = 82.16, 95% CI [80.17, 84.15]). This suggests that the coding scheme was applied consistently.

Synchronizing Pupil and Verbal Protocol Streams

To synchronize the pupillary response stream with the verbal protocol code stream, three sources of latency were considered: participant latency, coder latency, and pupillary response/LC latency. Participant latency refers to the latency that occurs because of a participant processing the APM item information and transforming it into an utterance. Participant latency unfortunately cannot be accounted for in our study because it is known to vary across individuals and types of processing steps and, therefore, will invariably add some noise to our data (Ericsson & Simon, 1993). In contrast, coder latency and LC pupillary response can and were accounted for before analysis. Coder response latency refers to the processing time it takes for the verbal protocol coder to process what they are hearing, make the decision to switch codes, and then actually press the key on the keyboard. To estimate this value, a random sample of 50 trials was

used to compare the coder key RT stamps to the original audio time series using audio editing software (Apple, Cupertino, CA). The results showed a coder response latency of approximately 1 sec ($M = 1014$ msec, $SD = 198$ msec). Finally, we considered the documented lag between LC activity and the pupillary response. Single-cell studies of LC neurons show that LC activity is tightly linked to stimulus onset, with a lag of only ≈ 200 msec (Clayton, Rajkowski, Cohen, & Aston-Jones, 2004; Rajkowski, Majczynski, Clayton, & Aston-Jones, 2004). However, the temporal resolution of the pupillary response is much lower than that of LC neurons. The pupil acts as a low-pass filter of LC activity with a lag of approximately 1 sec after stimulus onset (Hayes & Petrov, submitted; van Steenbergen & Band, 2013; Gagl, Hawelka, & Hutzler, 2011). As the coder and pupillary response latencies were approximately equivalent (each about 1 sec), no additional pre-processing was necessary to synchronize the pupil and code streams before analysis.

Segmentation of the Pupillary Data

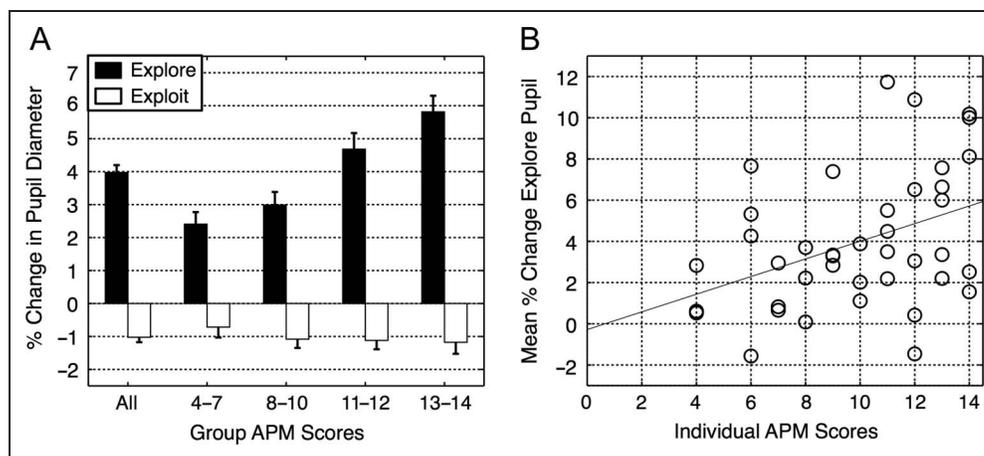
Finally, the pupillary data were segmented according to the exploratory and exploitative periods obtained from the verbal protocols. First, a baseline pupil diameter was calculated for each segment as follows: The baseline for the first nonneutral (exploratory) segment at the beginning of each trial was computed as the average pupil diameter during the first 500 msec of that segment to provide a more accurate baseline estimate as participants began the trial. The baseline for all subsequent segments until the end of the trial was computed as the average pupil diameter during the last 1000 msec of the immediately preceding segment.

Our main dependent variable is the percent change in pupil diameter (PCPD) relative to the relevant (most recent) baseline. The PCPD is measured in dimensionless units and is invariant with respect to the considerable individual differences in absolute pupil diameter as well as to slow drifts in pupillary tone. Specifically, the PCPD was computed as the task-evoked diameter minus the baseline diameter, divided by the baseline diameter. The mean PCPD was calculated by averaging the PCPD time series within each exploratory or exploitative segment for each participant on each trial (Beatty & Lucero-Wagoner, 2000).

RESULTS

The accuracy and RT data replicated well-documented patterns in the literature on Raven’s APM (e.g., Bors & Vigneau, 2003; Carpenter et al., 1990). There were substantial individual differences in overall APM scores, and the trial-by-trial accuracy decreased whereas RTs increased for the later, more difficult problems on the test. The verbal protocols indicated a slightly greater number of exploration than exploitation periods (990 explore,

Figure 2. Comparison of group-averaged PCPD for exploratory and exploitative periods by APM score, and a scatterplot of exploratory PCPD by individual APM score. (Left) Mean PCPD from baseline for exploratory and exploitative periods, averaged across all 40 participants and/or for subgroups at four ability levels ($n = 10$ for each subgroup). The group averages revealed a decrease in PCPD during exploitative periods and an increase during exploratory periods. The latter increase was significantly greater in the higher ability subgroups. The error bars represent ± 1 SEM. (Right) Individual differences in APM score were correlated with individual differences in mean PCPD during the exploratory periods.



945 exploit, 560 neutral). This is because the first non-neutral period on a trial was always exploratory, the two types alternated thereafter, and some trials ended in exploration mode. However, the exploitation periods were, on average, longer in duration (exploit: $M = 25.9$ sec, $SD = 25.2$ sec; explore: $M = 18.7$ sec, $SD = 14.4$ sec; neutral: $M = 1.75$ sec, $SD = 0.9$ sec).

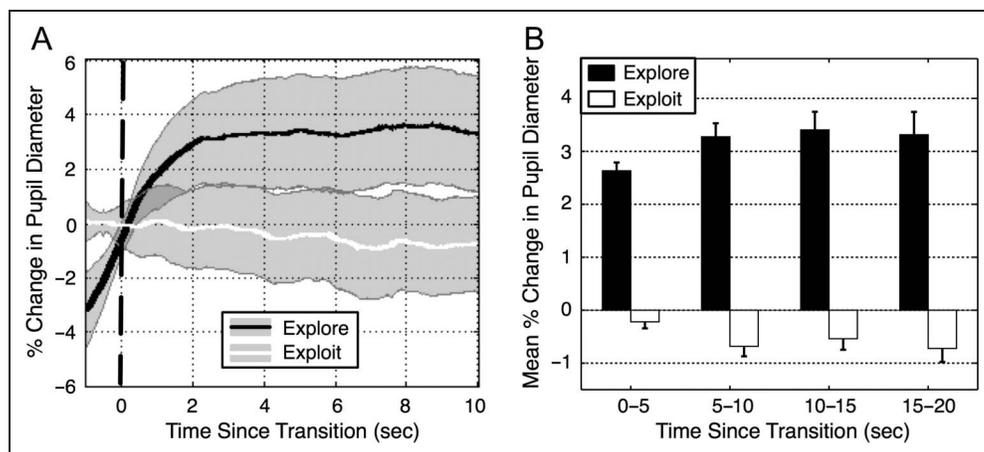
A significant boost in mean PCPD was observed during exploration periods relative to exploitation periods (Figure 2, left). A repeated-measures ANOVA with Segment type (explore vs. exploit) as a fixed factor and Participant as a random factor confirmed a strong exploration/exploitation effect on the PCPD ($F(1, 39) = 71.9, p < .001, \eta_p^2 = 0.65$). A pair of one-tailed t tests confirmed that the exploration effect was significantly greater than zero ($t(39) = 7.59, p < .001, r^2 = .59$) and the exploitation effect was significantly less than zero ($t(39) = -4.90, p < .001, r^2 = .38$). Under the linking hypothesis that PCPD is a valid index of LC-NE function, which in turn mediates the exploration–exploitation trade-off, this novel finding

suggests that this LC-NE mediation also operates during high-level analogical reasoning.

Furthermore, the mean exploratory PCPD increased linearly as a function of fluid reasoning ability as indexed by the APM. The steady increase is evident both in the group level (Figure 2, left) and individual level (Figure 2, right) data. A linear regression with mean exploratory PCPD as the sole predictor accounted for 16% of the variance in individual APM scores ($F(1, 38) = 7.05, p = .01, r^2 = .16$). Under the linking hypothesis outlined above, this additional novel finding suggests that individual differences in the mediation of the exploration–exploitation trade-off may contribute to individual differences in Gf. By contrast, no significant trends were observed in the exploitative pupillary response as a function of fluid ability.

The underlying temporal dynamics of the pupillary response revealed that these patterns in the averaged PCPD data were driven by sustained rather than momentary changes in pupil diameter (Figure 3). Figure 3A shows

Figure 3. Temporal dynamics of the pupillary response near segment transition boundaries. (Left) Group-averaged PCPD, time-locked to exploration and exploitation onset. The dashed vertical line at 0 sec indicates the onset of the new segment as identified in the verbal protocols. Transitions to exploration (black line) were accompanied with steep, sustained pupillary dilation, whereas transitions to exploitation (white line) showed only slow steady decrease in pupil diameter. The semitransparent gray error bands delineate 95% within-subject confidence limits. (Right) PCPD averaged within 5-sec-wide bins as a function of time spent exploring or exploiting. The error bars represent ± 1 SEM.



the grand-mean PCPD time-locked to exploration and exploitation onset.¹ The transitions from exploratory to exploitative segments were not associated (on average) with steep changes in pupil diameter. Rather, the exploitative segments exhibited (on average) a slow steady decrease in pupil diameter, depicted by the white line in Figure 3A. By contrast, the transitions from neutral/exploitative to exploratory segments were accompanied with steep increases in pupil diameter (black line) that began before exploratory language became manifest in the verbal protocols. This suggests that exploration likely preceded verbalization (on average) as can be seen by the positive slope of the black line near the transition boundary in Figure 3A. As this same period was taken as the baseline for subsequent change in pupil diameter, the PCPD values used in the statistical analyses may underestimate the magnitude of the pupillary dilation during exploration.

Furthermore, the exploratory pupil dilation seemed to persist for many seconds into the exploration period, as depicted in Figure 3 (right). We interpret this sustained pupillary dilation as a marker of a temporally extended exploratory state as opposed to a transient event.

In the behavioral data, we expected that both mean error rate and mean solution time would increase as a function of trial number according to the progressive nature of Raven's test and in agreement with previous findings (e.g., Bors & Vigneau, 2003; Carpenter et al., 1990). A one-tailed Pearson's product-moment correlation test confirmed that trial number accounted for a significant amount of variance in both error rate ($t(12) = 8.42, p < .001, r^2 = .85$) and solution time ($t(12) = 9.13, p < .001, r^2 = .87$).

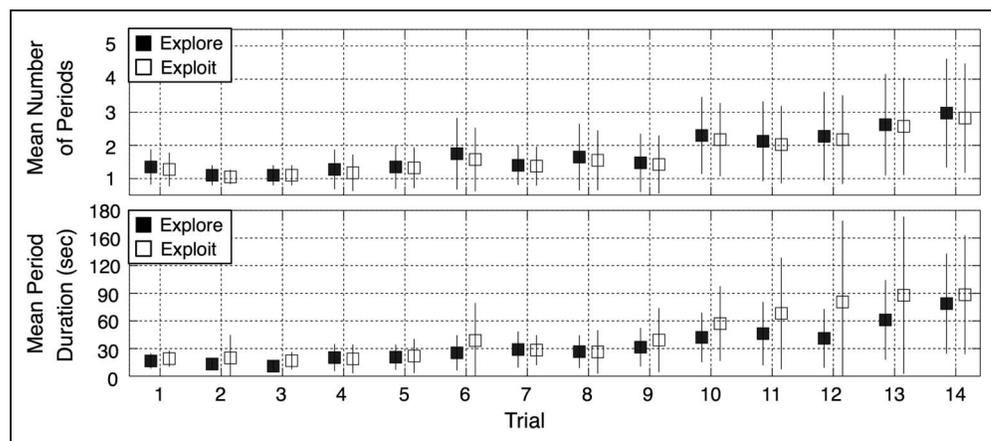
Given the strong trial effect on error rate and solution time, we tested for a trial difficulty effect on the pupillary response during exploration and exploitation. A trend analysis revealed a significant linear decrease in mean PCPD as a function of within-subject trial number during the exploration periods ($F(1, 507) = 41.01, p < .001, \eta_p^2 = .08$), whereas no statistically significant trend was de-

tected during the exploitation periods ($F(1, 507) = 1.08, p = .299$). This negative relationship between the magnitude of the exploratory pupillary dilation, on the one hand, and trial number, on the other, can be attributed to the much longer solution times on later, more difficult trials. Although the exploratory dilation could be sustained, on average, for at least 20 sec (Figure 3B), many exploratory periods were quite longer on difficult trials, eventually diluting the exploratory PCPD increase.

To investigate this further, Figure 4 presents some basic descriptive statistics about the number and duration of exploratory and exploitative segments. Both quantities increased as a function of trial number (and trial difficulty). The earliest trials typically exhibited only one brief exploration period followed by a single brief exploitation period. On the more difficult middle and late items, however, the participants tended to alternate multiple times between exploring and exploiting. Trend analyses revealed significant linear ($F(1, 507) = 193.38, p < .001, \eta_p^2 = .27$) and quadratic ($F(1, 507) = 21.02, p < .001, \eta_p^2 = .04$) trends in the total number of transitions between exploration and exploitation as a function of Trial. Analogous analyses also revealed significant linear and quadratic trends in exploration duration (linear: $F(1, 468) = 32.72, p < .001, \eta_p^2 = .06$; quadratic: $F(1, 468) = 46.82, p < .001, \eta_p^2 = .09$) and exploitation duration (linear: $F(1, 468) = 18.21, p < .001, \eta_p^2 = .04$; quadratic: $F(1, 468) = 26.47, p < .001, \eta_p^2 = .05$).

Furthermore, there was evidence for interactions between the difficulty of the test items and the Gf of the participants as indexed by their APM scores. Recall that the participants were sampled from four ability groups. Repeated-measures ANOVAs with Group as a between-subject factor and Trial as a within-subject factor showed significant Trial \times Group interactions for the number of transitions ($F(39, 468) = 1.65, p < .01, \eta_p^2 = .12$), exploration duration ($F(39, 468) = 1.83, p < .01, \eta_p^2 = .13$), exploitation duration ($F(39, 468) = 2.22, p < .01, \eta_p^2 = .16$), and total solution time ($F(39, 468) = 2.37, p < .001, \eta_p^2 = .16$). These Ability \times Difficulty interactions reflected

Figure 4. Mean number and duration of exploration and exploitation periods as a function of trial number. (Top) The mean number of exploration and exploitation periods increased as the Raven problems got progressively more difficult. (Bottom) The mean period duration also increased as trial difficulty increased, reflecting the increasingly complex figural elements and relations characteristic of the most difficult problems. The error bars on both panels represent $\pm 1 SD$.



the differential ability of the participants to engage with the most difficult items (Trials 11–14). High-ability participants would struggle yet work through those difficult items over a lengthy sequence of alternating exploration and exploitation periods, whereas lower ability participants were prone to become overwhelmed, take a guess, and terminate the trial after a comparatively short effort.

These Ability \times Difficulty interactions raise a possible alternative explanation for the correlation between exploratory PCPD and ability group depicted in Figure 2A above. It is possible that this correlation might simply be driven by the difference in performance between high- and low-ability participants on the most difficult items. However, the following analysis suggests that this is unlikely. When the mean exploratory PCPD, taken across Trials 1–10 only, excluding the most difficult trials (Trials 11–14), was used as a predictor in a linear regression, it accounted for 20% of the variance in individual APM scores ($F(1, 38) = 9.28, p < .01, r^2 = .20$). Recall that the mean exploratory PCPD across all 14 trials accounted for 16% of this variance. Therefore, when the high- and low-ability participants spent the same amount of time exploring, the correlation between exploratory PCPD and APM scores actually increased. This rules out the late trials as a potential confounding factor. If anything, the random guessing on the most difficult trials in the lower ability groups probably adds nonsystematic variance that degrades the correlation.

We also checked the so-called time-on-task effect as a potential confound. Prior studies have found that pupil diameter can decrease systematically during the experimental session (Hayes & Petrov, submitted; Beatty, 1982; Kahneman & Beatty, 1967). It should be noted that these studies used low-level perceptual tasks with subsecond RTs (e.g., vigilance, auditory discrimination, visual-motion discrimination). These monotonous simple tasks are quite different than Raven's APM, which is designed to vary the figural material constantly to measure fluid (as opposed to crystallized) intelligence. It has been suggested that the decreasing pupil size in earlier studies may be a result of decreasing arousal as participants get bored with the task (Laeng, Sirois, & Gredeback, 2012; Beatty & Lucero-Wagoner, 2000). Some participants in our study worked for over 5 min on some of the (difficult) Raven problems, which raises the concern of a within-trial time-on-task effect as a potentially confounding factor in our data. To estimate the magnitude of the time-on-task effect over the course of individual trials, we performed a series of robust linear regressions² on the pupillary diameter as a function of the time since each stimulus onset. This produced one slope-parameter estimate per trial. These were averaged across trials to produce one aggregate slope estimate per participant. The latter estimates did not differ significantly from zero ($t(39) = 1.18$, two-tailed $p = .25$). In addition, recall that the exclusion of the four longest trials from the analysis only strengthened the pattern in Figure 2A. Being the longest, these trials should be the most vulnerable to a possible time-on-task effect. Overall,

this confound does not seem a viable explanation of our results. Apparently, the consistently novel and challenging nature of the Raven task kept our participants engaged throughout each trial and throughout the session as a whole.

Finally, we checked whether the exploration and exploitation periods differed in terms of missing values in the pupillometric time series and in terms of saccade frequency. Missing values occur when the eye tracker temporarily loses the pupil, for example, because of blink artifacts. Such artifacts were rejected during preprocessing. The artifact frequency in the raw pupil data was similar for exploration ($M = 11.18, SD = 12.42$ percent of period) and exploitation ($M = 10.96, SD = 12.47, t(39) = 0.50, p = .62$). Saccade frequency was significantly lower during exploration ($M = 10.03, SD = 5.11$ percent of period) compared with exploitation ($M = 11.19, SD = 5.58, t(39) = 6.28, p < .001$). However, we are not aware of any studies showing a systematic effect of saccade frequency on the pupillary response, and the 1% difference is not likely to account for the large exploration/exploitation effect in our data. Saccades produce a known risk of pupil foreshortening error as they change the gaze position, but this source of systematic error was corrected during preprocessing (Hayes & Petrov, 2015).

DISCUSSION

A novel combination of pupillometry and verbal protocol analysis was used to compare changes in pupil diameter during exploration and exploitation control states during visual analogy making. The analysis revealed a significant increase in pupil diameter during exploration and decrease during exploitation. This broad finding is the first to generalize theories of the LC-NE system's role in the exploration–exploitation trade-off to a high-level analogical reasoning task such as Raven's APM. More importantly, individual differences in the relative magnitude of exploratory pupillary dilation accounted for 16% of the variance in APM scores. This novel result suggests that individual differences in general Gf may be related to underlying differences in noradrenergic function.

Our findings build upon and are consistent with previous studies of the exploration–exploitation trade-off that monitored the pupillary response as a noninvasive index of the LC-NE system (Jepma & Nieuwenhuis, 2011; Gilzenrat et al., 2010). Gilzenrat et al. (2010) had participants complete an auditory pitch discrimination task in which reward increased as the pitch discrimination difficulty increased, until the discrimination eventually became impossible. Importantly, their participants were allowed the option to escape before each trial. Escaping would reset the reference tone, difficulty, and reward levels. They found that baseline pupil diameter increased leading up to escape trials and decreased afterward. This is consistent with our observation of a decrease during exploitation and increase during exploration. However, the effect size

and increased diameter during exploration, consistent with prominent theories of LC-NE function. Importantly, one sixth of the variance in Raven scores was accounted for by individual differences in exploratory pupillary dilation. These findings shed new light on the relationship between the exploration–exploitation trade-off, noradrenergic function, and individual differences in Gf.

Acknowledgments

This research was supported by the National Eye Institute (R21 EY022745).

Reprint requests should be sent to Taylor R. Hayes, Center for Mind and Brain, University of California, Davis, CA 95618, or via e-mail: taylor.r.hayes@gmail.com.

Notes

- Note that the pupil baseline is applied retroactively to the data points before the segment transition boundary in Figure 3A. This is for plotting purposes only. In the statistical analyses, these data points were incorporated into the preceding segment.
- Each regression used iteratively reweighted least squares with a bisquare weighting function. Ordinary regressions yielded similar results.

REFERENCES

- Ahern, S. K., & Beatty, J. (1979). Pupillary responses during information processing vary with scholastic aptitude test scores. *Science*, *205*, 1289–1292.
- Ahern, S. K., & Beatty, J. (1981). Physiological evidence that demand for processing capacity varies with intelligence. In M. P. Friedman, J. P. Das, & N. O'Connor (Eds.), *Intelligence and learning* (pp. 121–128). New York: Plenum.
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, *111*, 1036–1060.
- Aston-Jones, G., & Cohen, J. D. (2005). An integrative theory of locus coeruleus-norepinephrine function: Adaptive gain and optimal performance. *Annual Review of Neuroscience*, *28*, 403–450.
- Beatty, J. (1982). Phasic not tonic pupillary responses vary with auditory vigilance performance. *Psychophysiology*, *19*, 167–172.
- Beatty, J., & Lucero-Wagoner, B. (2000). The pupillary system. In J. T. Cacioppo, L. G. Tassinary, & G. G. Berntson (Eds.), *Handbook of psychophysiology* (2nd ed., pp. 142–162). Cambridge: Cambridge University Press.
- Berridge, C. W., & Waterhouse, B. D. (2003). The locus coeruleus-noradrenergic system: Modulation of behavioral state and state-dependent cognitive processes. *Brain Research Reviews*, *42*, 33–84.
- Bors, D. A., & Vigneau, F. (2003). The effect of practice on Raven's advanced progressive matrices. *Learning and Individual Differences*, *13*, 291–312.
- Brouwers, S. A., Van de Viver, F. J. R., & Van Hemert, D. A. (2009). Variation in Raven's progressive matrices scores across time and place. *Learning and Individual Differences*, *19*, 330–338.
- Brown, E. T., Gilzenrat, M. S., & Cohen, J. D. (2005). *The locus coeruleus, adaptive gain, and the optimization of simple decision tasks* (Technical Report). Princeton, NJ: Princeton University.
- Carpenter, P. A., Just, M. A., & Shell, P. (1990). What one intelligence test measures: A theoretical account of the processing in the Raven Progressive Matrices test. *Psychological Review*, *97*, 404–431.
- Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., de Gardelle, V., Castanon, S. H., et al. (2014). Adaptive gain control during human perceptual choice. *Neuron*, *81*, 1429–1441.
- Clayton, E. C., Rajkowski, J., Cohen, J. D., & Aston-Jones, G. (2004). Phasic activation of monkey locus coeruleus neurons by simple decisions in a forced-choice task. *Journal of Neuroscience*, *24*, 9914–9920.
- Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society, Series B, Biological Sciences*, *362*, 933–942.
- Einhäuser, W., Koch, C., & Carter, O. L. (2010). Pupil dilation betrays the timing of decisions. *Frontiers in Human Neuroscience*, *4*, 1–9.
- Einhäuser, W., Stout, J., Koch, C., & Carter, O. (2008). Pupil dilation reflects perceptual selection and predicts subsequent stability in perceptual rivalry. *Proceedings of the National Academy of Sciences, U.S.A.*, *105*, 1704–1709.
- Eldar, E., Cohen, J. D., & Niv, Y. (2013). The effects of neural gain on attention and learning. *Nature Neuroscience*, *16*, 1146–1153.
- Ericsson, K. A., & Simon, H. A. (1993). *Protocol analysis: Verbal reports as data* (Rev. ed.). Cambridge, MA: MIT Press.
- Gagl, B., Hawelka, S., & Hutzler, F. (2011). Systematic influence of gaze position on pupil size measurement: Analysis and correction. *Behavior Research Methods*, *43*, 1171–1181.
- Gilzenrat, M. S., Nieuwenhuis, S., Jepma, M., & Cohen, J. D. (2010). Pupil diameter tracks changes in control state predicted by the adaptive gain theory of locus coeruleus function. *Cognitive, Affective & Behavioral Neuroscience*, *10*, 252–269.
- Gray, J. R., Chabris, C. F., & Braver, T. S. (2003). Neural mechanisms of general fluid intelligence. *Nature Neuroscience*, *6*, 316–322.
- Hayes, T. R. (2015). *Mechanisms of visual relational reasoning* (Unpublished doctoral dissertation). The Ohio State University, Columbus, OH.
- Hayes, T. R., & Petrov, A. A. (2015). Mapping and correcting the influence of gaze position on pupil size measurements. *Behavior Research Methods*, 1–18. Advance online publication. doi:10.3758/s13428-015-0588-x.
- Hayes, T. R., & Petrov, A. A. (submitted). Learning is in the eye of the beholder: Phasic pupil diameter decreases during perceptual learning.
- Hayes, T. R., Petrov, A. A., & Sederberg, P. B. (2011). A novel method for analyzing sequential eye movements reveals strategic influence on Raven's Advanced Progressive Matrices. *Journal of Vision*, *11*, 1–11.
- Hayes, T. R., Petrov, A. A., & Sederberg, P. B. (2015). Do we really become smarter when our fluid-intelligence scores improve? *Intelligence*, *48*, 1–14.
- Hertzum, M., & Holmegaard, K. D. (2013). Thinking aloud in the presence of interruptions and time constraints. *International Journal of Human-Computer Interaction*, *29*, 351–364.
- Jay, B. S. (1962). The effective pupillary area at varying perimetric angles. *Vision Research*, *1*, 418–424.
- Jennings, J. A., & Charman, W. N. (1978). Optical image quality in the peripheral retina. *American Journal of Optometry and Physiological Optics*, *55*, 582–590.
- Jepma, M., & Nieuwenhuis, S. (2011). Pupil diameter predicts changes in exploration–exploitation trade-off: Evidence for the adaptive gain theory. *Journal of Cognitive Neuroscience*, *23*, 1587–1596.

- Kahneman, D., & Beatty, J. (1967). Pupillary responses in a pitch-discrimination task. *Perception & Psychophysics*, *2*, 101–105.
- Kammerer, Y., & Gerjets, P. (2013). The role of thinking-aloud instructions and prior domain knowledge in information processing and source evaluation during Web search. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 716–721). Austin, TX: Cognitive Science Society.
- Klingner, J. (2010). *Measuring cognitive load during visual tasks by combining pupillometry and eye tracking* (Unpublished doctoral dissertation). Stanford University, Stanford, CA.
- Koss, M. (1986). Pupillary dilation as an index of central nervous system α_2 -adrenoceptor activation. *Journal of Pharmacological Methods*, *15*, 1–19.
- Laeng, B., Sirois, S., & Gredeback, G. (2012). Pupillometry: A window to the preconscious? *Perspectives on Psychological Science*, *7*, 18–27.
- Lovett, A., Tomai, E., Forbus, K., & Usher, J. (2009). Solving geometric analogy problems through two-stage analogical mapping. *Cognitive Science*, *33*, 1192–1231.
- Mathur, A., Gehrman, J., & Atchison, D. A. (2013). Pupil shape as viewed along the horizontal visual field. *Journal of Vision*, *13*, 1–8.
- Murphy, P. R., O'Connell, R. G., O'Sullivan, M., Robertson, I. H., & Balsters, J. H. (2014). Pupil diameter covaries with BOLD activity in human locus coeruleus. *Human Brain Mapping*, *35*, 4140–4154.
- Murphy, P. R., Robertson, I. H., Balsters, J. H., & O'Connell, R. G. (2011). Pupillometry and P3 index of locus coeruleus noradrenergic arousal function in humans. *Psychophysiology*, *48*, 1531–1542.
- Newell, A., & Simon, H. A. (1976). Computer science as empirical enquiry: Symbols and search. *Communications of the Association of Computing Machinery*, *19*, 113–126.
- Phillips, M. A., Szabadi, E., & Bradshaw, C. M. (2000). Comparison of the effects of clonidine and yohimbine on spontaneous pupillary fluctuations in healthy human volunteers. *Psychopharmacology*, *150*, 85–89.
- Rajkowski, J., Kubiak, P., & Aston-Jones, G. (1994). Locus coeruleus activity in monkey: Phasic and tonic changes are associated with altered vigilance. *Brain Research Bulletin*, *35*, 607–616.
- Rajkowski, J., Majczynski, H., Clayton, E., & Aston-Jones, G. (2004). Activation of monkey locus coeruleus neurons varies with difficulty and performance in a target detection task. *Journal of Neurophysiology*, *92*, 361–371.
- Raven, J. C., Raven, J., & Court, J. H. (1998). *Manual for Raven's progressive matrices and vocabulary scales. Section 4: Advanced progressive matrices*. San Antonio, TX: Pearson.
- Samuels, E. R., & Szabadi, E. (2008). Functional neuroanatomy of the noradrenergic locus coeruleus: Its roles in the regulation of arousal and autonomic function part II: Physiological and pharmacological manipulations and pathological alterations of locus coeruleus activity in humans. *Current Neuropharmacology*, *6*, 254–285.
- Spring, K. H., & Stiles, W. S. (1948). Apparent shape and size of the pupil viewed obliquely. *British Journal of Ophthalmology*, *32*, 347–354.
- Taatgen, N. A. (2013). The nature and transfer of cognitive skills. *Psychological Review*, *120*, 439–471.
- Taatgen, N. A., Huss, D., Dickison, D., & Anderson, J. R. (2008). The acquisition of robust and flexible cognitive skills. *Journal of Experimental Psychology: General*, *137*, 548–565.
- Usher, M., Cohen, J. D., Servan-Schreiber, D., Rajkowski, J., & Aston-Jones, G. (1999). The role of locus coeruleus in the regulation of cognitive control. *Science*, *283*, 549–554.
- Van Der Meer, E., Beyer, R., Horn, J., Foth, M., Bornemann, B., Ries, J., et al. (2010). Resource allocation and fluid intelligence: Insights from pupillometry. *Psychophysiology*, *47*, 158–169.
- van Steenbergen, H., & Band, G. P. H. (2013). Pupil dilation in the Simon task as a marker of conflict processing. *Frontiers in Human Neuroscience*, *7*, 1–11.