Pupillary manifolds: uncovering the latent geometrical structures behind phasic changes in pupil size

***Elvio Blini, Roberto Arrighi, and Giovanni Anobile**

Department of Neuroscience, Psychology, Pharmacology and Child Health, University of Florence, Florence, Italy

*** Corresponding author:**

Elvio Blini

elvioadalberto.blini (at) unifi.it

Firenze, 50136, Via di San Salvi 12, Building 26

Abstract

The size of the pupils reflects directly the balance of different branches of the autonomic nervous system. This measure is inexpensive, non-invasive, and has provided invaluable insights on a wide range of mental processes, from attention to emotion and executive functions. Two outstanding limitations of current pupillometry research are the lack of consensus in the analytical approaches, which vary wildly across research groups and disciplines, and the fact that, unlike other neuroimaging techniques, pupillometry lacks the dimensionality to shed light on the different sources of the observed effects. In other words, pupillometry provides an integrated readout of several distinct networks, but it is unclear whether each has a specific fingerprint, stemming from its function or physiological substrate. Here we show that phasic changes in pupil size are inherently low-dimensional, with modes that are highly consistent across behavioral tasks of very different nature, suggesting that these changes occur along pupillary manifolds that are highly constrained by the underlying physiological structures rather than functions. These results provide not only a unified approach to analyze pupillary data, but also the opportunity for physiology and psychology to refer to the same processes by tracing the sources of the reported changes in pupil size in the underlying biology.

Keywords

Pupillometry; Cognitive load; Light reflex; Dimensionality reduction; Pupillary manifolds; Autonomic nervous system

Graphical abstract

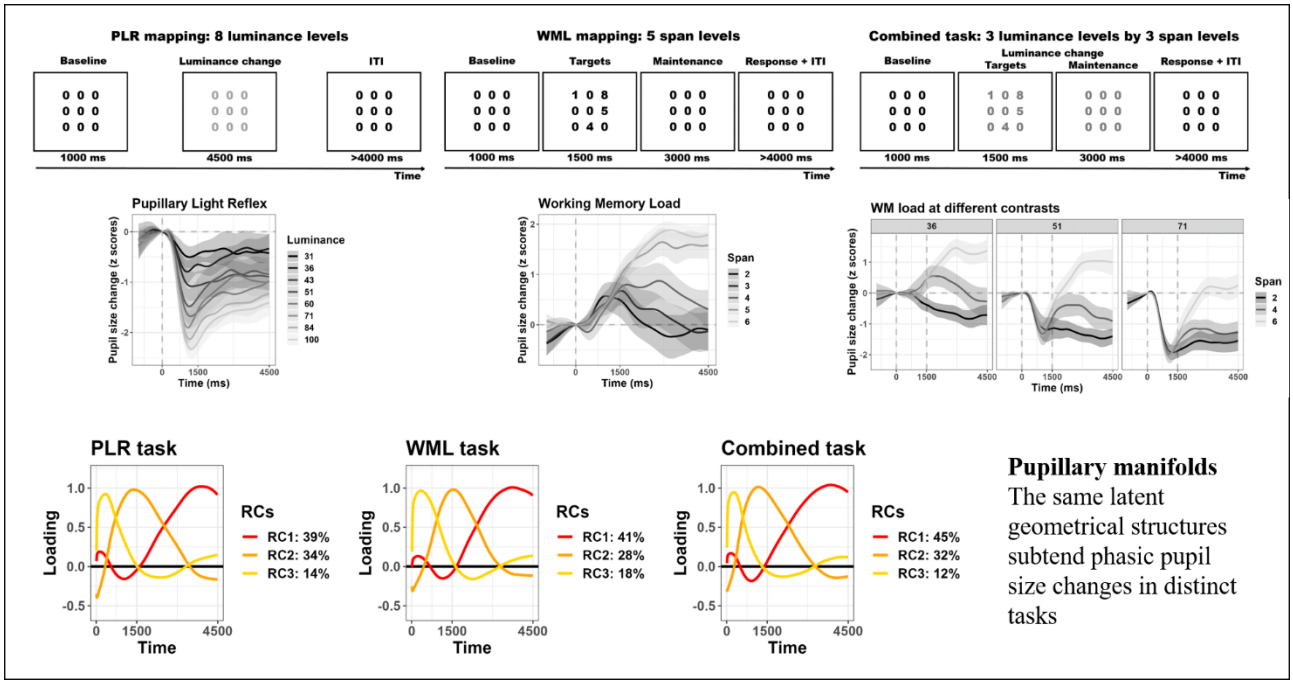

Significance statement

Phasic changes in pupil size are thought to reflect dynamic shifts between attentional states as instantiated by the locus-coeruleus noradrenaline system, and are crucial for adaptive behaviors. We found that the latent space of these changes is low-dimensional and remarkably similar across very different tasks, involving distinct cognitive processes. We therefore introduce the notion of pupillary manifolds as latent spaces that subtend the generative processes behind these changes. We suggest that manifolds arise due to hard constraints in the underlying physiological substrate – the relative balance between sympathetic and parasympathetic activity. In the framework outlined here, these mechanisms can be accessed and described directly, with only a handful of parameters, thus better informing computational modelling.

Introduction

The pupils evolved primarily as a key tool for vision. Their core task is to manage the amount of light reaching the retina at any given moment, depending on the environment, as to optimize visual acuity [1–3]. The nature of pupillary responses to light, i.e. pupil constriction, is therefore largely reflexive. However, pupil dynamics that do not have a strict environmental explanation (in terms of light levels) also exist; decades of research have ascribed them to a plethora of distinct cognitive processes spanning attention, emotion, working memory load, and executive functions more generally [3–11]. Even the most fundamental **Pupillary Light Reflex (PLR)** is not completely impervious to cognitive, top-down modulations [12]. For example, the PLR is increased, and the pupils constrict more, whenever stimuli are attended, even if covertly [13–16]. The pupils can dilate, instead, for all sorts of arousing, demanding tasks and stimuli [17–21], often included under the same umbrella term of "psychosensory" modulators [3]. The dilation response is notoriously sluggish, and is only seen several hundreds of milliseconds after the trigger event and the corresponding early orienting reflex, which also yields effects that are protracted in time. This means that pupillary responses that are functionally distinct typically overlap, on the one hand, and that searching for timepoints with significant differences in pupil size between experimental conditions is questionable at best, on the other hand. Regardless of the chosen analytical procedure, it is clear that as soon as one timepoint shows a significant difference this is the result of fully-blown latent processes and their interaction. Here we therefore turned to dimensionality reduction techniques to unveil these latent, unobserved processes and characterize them. The question was whether functionally distinct pupillary signatures can be mapped onto a low-dimensional space, and if so whether this space reflects the specific function mapped by the task (and experimental context) or rather unspecific physiological constraints.

Results

We started with mapping the pupillary responses of 20 healthy human participants to stimuli changing (slightly) in luminance. We asked participants to simply view a matrix of numbers superimposed on a black background: nine "zeros" were arranged in three rows and intermittently changed their shades of gray so that we could map phasic responses to 8 different luminance levels (**Figure 1A**). We found the well-known PLR, consisting in a rather quick constriction of the pupils which reaches its peak within 1/1.5 seconds and then gradually recoups to a new baseline diameter (**Figure 2A**). We performed dimensionality reduction first with Principal Components Analysis (PCA) to obtain features that maximally summarize the variability in the data. We found that one component (PC1), i.e. one score value per trial, was sufficient to represent 79% of all the data (3 scores represented 95%, **Supplementary Figure 1**). The data are also well summarized by the eigenvector of PC1 (**Figure 2B**), describing the relevance of each timepoint to the component, which clearly depicts the "shape" of the PLR. While this is not surprising, because pupil traces are strongly autocorrelated, it remains remarkable how compact such a rich dataset and physiological measure can become. Having to cope with only a handful of values per trial is more manageable for most uses. Furthermore, the odds that components start to represent noise, instead of genuine signal, increase with their number; in this sense resolving to use few, quintessential features allows one to probe the primitive shape of the process at hand, e.g. the PLR, directly, without a priori assumptions. In **figure 2C** we show that this approach is also powerful in that PC1 clearly separates all the 8 different luminance levels administered in the task ($F_{(1, 527.75)}= 144.77$, $p <.001$), thereby corroborating the notion that it could represent a latent dimension along which the PLR happens with different strength.

We then moved to mapping the pupillary responses to changing **Working Memory Load (WML)**. In this case, the presented numbers did not change their luminance. Instead, for 1.5 s, the matrix was populated with 2 to 6 numbers chosen randomly and appearing in random positions (**Figure 1B**). Participants had to scan the matrix first, and then retain the numbers in working memory for an

additional 3 s before providing their response, following the presentation of an auditory cue. We thus adapted a rather classic working memory span task in order to avoid the sequential presentation of stimuli, in consequence of which storing in memory (and pupil dilation) would presumably occur at different moments in different load conditions [22]. The adaptation was effective because: behavioral performance declined with increasing WML (**Supplementary Figure 2**); pupillary responses clearly showed robust dilation as a function of WML (**Figure 2D**, $F_{(1, 18.66)}= 137.82$, $p <.001$). As observed for the PLR mapping, this task also provided a dataset composed of very few latent dimensions. One principal component could account for 76% of the variability, 3 components about 93% (**Supplementary Figure 1**). The first component clearly distinguished different WML levels (**Figure 2F**), and its eigenvector (**Figure 2E**) suggests the primitive shape of WML is a slower, sustained dilation during the course of the trial and memory maintenance.

Next, we administered a task in which both dimensions changed in a 3x3 design, that is, both luminance and the amount of numbers to memorize were manipulated (**Figure 1C**). Both main effects were evident in pupillary recordings and resembled those described in the previous tasks (**Figure 3A**), also behaviorally (**Supplementary Figure 3**). Once again, the dataset was inherently low-dimensional, and the first component could account for 73% of the overall variability (3 components for 94%, **Supplementary Figure 1**). In this case, PC1 tracked well both features (Luminance: $F_{(1, 64.91)}= 28.62$, $p <.001$; Memory load: $F_{(1, 21.97)}= 41.18$, $p <.001$; **Figure 3C**) but did not discriminate between the two. Indeed, the eigenvector of PC1 (**Figure 3B**) fell almost perfectly midway between the PLR and WML fingerprints described above (**Figure 3D**). While, again, this remains a handy, effective way to summarize a dataset without too many arbitrary choices (e.g., a given time window) or assumptions (e.g., the shape of the pupillary function), it remains to be assessed whether this low-dimensional space maps onto distinct functional processes.

We therefore turned to non-orthogonal approaches with the aim of identifying latent processes that are potentially correlated as well as more interpretable. We focus in particular on promax-rotated

PCA (rPCA), in keeping with previous suggestions [23–25]. We choose to isolate k= 3 components following these studies and because components above three accounted for less than 3% of overall variability each in PCA. These analyses yielded three rotated components which were remarkably similar across the three different tasks, albeit accounting for a slightly different share of the overall variability (**Figure 4A**). Despite that, and despite oblique rotations are often non-optimal to account for the largest variability sources in the data, three components were sufficient to account for most of the phasic pupillary dynamics in all tasks (>87%). Because the three tasks vary substantially in terms of their requirements and pupillary signatures, these components are unlikely to reflect task-specific features. Rather, since this structure appears consistently across studies [23–25] and tasks of different nature, they are more likely to reflect underlying physiological constraints of the signal (changes in pupil size) rather than function (either pupil constriction to light or dilation to memory load). Indeed, both luminance and memory load in the combined task were equally well-tracked by the first two rotated components, which could discriminate both (Luminance: $F_{(1, 55.78)}$= 27.21 and p <.001 for RC1, $F_{(1, 3.58)}$= 23.96 and p =.011 for RC2; Memory load: $F_{(1, 26.52)}$= 43.49 and p <.001 for RC1, $F_{(1, 302.22)}$= 50.92 and p <.001 for RC2; **Figure 4B**). The exception was RC3: this component showed up consistently across tasks, with the earliest latency, and explained the smallest share of the overall variance. RC3 was sensitive to changes in luminance in the combined task ($F_{(1, 72.74)}$= 13.88, p <.001; **Figure 4C**) and in the PLR mapping (**Supplementary Figure 4**), but not cognitive load in either the combined or WML mapping tasks ($F_{(1, 22.09)}$= 0.63, p =.437; **Figure 4C** and **Supplementary Figure 5**).

Discussion

In this study we decided to map pupillary reflexes to light, which are far more robust than, for example, the modulations of the PLR that may be caused by attention [13,14]. Furthermore, these reflexes were contrasted against psychosensory modulations of pupil size induced by working memory load.

This means that we were in a rather favorable position to maximize the chances to isolate processes that are both functionally and physiologically distinct, considered that their neural generators are very different [12,26,27]. The point was not as much about showing that changes in pupil size are inherently low-dimensional – this is expected in light of the strong autocorrelation of the signal – but rather how well the latent, low-dimensional space separates distinct functional processes.

We did find one component (RC3) that appeared to map luminance changes but not pupil dilation due to cognitive load. This component accounted for a small share of the overall variance in the data and loaded especially on the earliest timepoints, which may suggest it refers to the PLR specifically. However, it is striking that this component was still recovered as such in a task only probing memory load, and even though it did not discriminate between key conditions; this points to an underlying latent structure that remains deployed, possibly silently, regardless of the task at hand (e.g., a primarily parasympathetic component). Overall, indeed, we found that the components recovered from very different tasks and pupil traces were very similar; this by itself speaks about the presence of common constraints in these traces. Moreover, the same components, beyond RC3, could track both reflexes to light and mental effort, which lead to signature changes in pupil size. The fact that these functionally different processes can be quantified very well by few components with a similar shape points again to the existence of constraints of a different nature. For example, a few authors have suggested, based on pharmacological interventions or the manipulation of environmental light levels, that pupillary dilation to mental effort may be composed of two waves [28,29]: the first, predominant in bright environments, would reflect the cortical inhibition of the parasympathetic efferent pathway; the second, occurring later in time and explaining most of the dilation, would represent instead a primary sympathetic component. One should be weary in over-interpreting oscillatory patterns in PCA [30]. However, this account fits well with our findings, and could very well outline the main constraints imposed on pupil size, beyond its autocorrelated nature. Contextual variables may tap onto the dynamic continuum between the different branches of the autonomic nervous system, thereby causing pupil size to change along low-dimensional manifolds that reflect

this balance accurately, but much more efficiently. Clearly, the blanket is too short to separate neatly distinct functional processes, as increased sympathetic activity implies, with the passing of time, decreased parasympathetic tone. Still, here we show that changes in pupil size are not different from much richer physiological signals in their potential to be mapped onto low-dimensional, interpretable spaces [31,32]. The notion of neural manifolds has enriched greatly the current debate around the neural control of movement or other cognitive processes; we surmise that the notion of pupillary manifolds could, likewise, provide fruitful ground for cognitive and computational pupillometry. The possibility to infer latent, generative processes behind the observed signal is of particular interest. These processes may sometimes be hard to see from the data, and yet constitute the source biological signal that explains most of the subsequent changes in pupil size; for most uses, this is therefore the signal of interest. Ultimately, the concept of pupillary manifold provides a tool to delve deeper into the physiological origin of phasic changes in pupil size. These changes have been clearly implicated in finely-tuning behavior toward maximal utility [33–35], so that accessing and better characterizing their fingerprint on pupil size represents a very appealing challenge, one which would contribute harmonizing the terminology between psychology and physiology. Last but not least, this feature points to the opportunity of a unified approach to analyze pupillary data via dimensionality reduction, for which we provide a toolbox (https://github.com/EBlini/Pupilla). This approach not only shields from more or less arbitrary assumptions about the latency (e.g., time windows) or the shape (e.g., pupillary response function) of pupillary dynamics (which open an entire multiverse of analytical choices, [36]): it could feed much more efficiently physiologically-informed computational models of human learning and decision-making.

Methods

All materials, raw data, and analysis scripts for this study are available through the Open Science Framework website: https://osf.io/dkpcs.

The core functions for preprocessing and analysis are available through GitHub: https://github.com/EBlini/Pupilla.

Participants

Twenty participants, all students from the University of Florence, took part in this study (17 females, M= 25.1 years, SD= 6.95 years). Inclusion criteria were normal (or corrected-to-normal) vision, and no history of neurological, psychiatric, or sensory disorders. The experimental procedures were approved by the local ethics committee (Commissione per l'Etica della Ricerca, University of Florence, July 7, 2020, n. 111). The research was carried in accordance with the Declaration of Helsinki, and informed consent was obtained from all participants prior to the experiment.

Procedures

The participants were tested in a dimly lit, quiet room, their head comfortably resting on a chinrest. They faced a remote infrared-based eye-tracker (EyeLink 1000, SR Research Ltd.), at a distance of approximately 57 cm from the screen. Each session started with a 15-points calibration of the eye-tracker, which was then set to monitor participants' pupil size continuously at a 500 Hz sampling rate. The open-source software OpenSesame [37] was used to present the stimuli, along the procedure outlined below and depicted in **Figure 1**.

PLR mapping.

We asked participants to passively view a matrix of gray numbers superimposed on a black background (0 on the grayscale, 0.8 cd/m^2, **Figure 1A**). The numbers were nine "zeros" encompassing about 3°, written in gray (30 on the grayscale). In each trial, numbers were presented as such for 1000 ms during the baseline phase. Then, for 4500 ms, the luminance of the numbers changed randomly to one of 8 grayscale values (31, 36, 43, 51, 60, 71, 84, 100). The 8 levels were chosen to be approximately linearly arranged on the logarithmic space. In absolute terms, these luminance values mapped roughly linearly in the range between 3.5 and 27 cd/m^2 (note that this reflects the color of the font, not the local luminosity of stimuli). After this phase of interest, the numbers returned to the original intensity value and were presented for 4000 to 5000 ms (with a uniform jitter), in order to allow the pupils to return to their baseline size. There were 48 trials, 6 for each luminance level.

WML mapping.

The stimuli during the baseline and ITI phases were identical to the PLR task, though their color intensity (on the grayscale) was slightly higher (51); this was because the focus was pupillary dilation to working memory load, and this minor adaptation allowed for more room to observe changes in this direction. In this task, the numbers presented in the phase of interest did not change their luminance. Instead, for 1500 ms, the matrix was populated with 2 to 6 numbers chosen randomly and appearing in random positions (**Figure 1B**). Participants had to scan the matrix first, and then retain the numbers in working memory for an additional 3000 ms before providing their response. During this time, the numbers returned identical to the baseline (nine "zeros"). Once this phase ended, an auditory cue prompted participants' responses, which was then scored manually by the researcher. This task was composed of 40 trials, 8 for each of the 5 levels of cognitive load.

Combined task.

In this task, both dimensions (luminance and cognitive load) changed in a 3x3 design (**Figure 1C**). The stimuli during the baseline and ITI phases were identical to the PLR task. Then, either 2, 4, or 6 randomly selected numbers appeared in random positions in the matrix; concurrently, the luminance of the numbers changed randomly to one of 3 brighter values (36, 51, 71 on the grayscale). After 1500 ms, the matrix returned to being composed only by zeros, though luminance manipulation remained in place for an additional 3000 ms. The participants' response was solicited by an auditory cue (as in the WML task above), and manually recorded by the experimenter. This task was composed of 90 trials, 10 for each of the 9 conditions given by the 3 x 3 design (Span by Luminance); it was administered in two blocks to avoid participants' excessive fatigue and distress.

Data processing

We only retained pupil size measured during fixations, as identified by the eyetracker (i.e., measures of pupil size during blinks and saccades were discarded). We did not further correct pupil size by regressing out the position of the gaze on the screen. However, we don't have a reason to suspect this may constitute a bias because the eyes were highly constrained to central vision (< 3°), and whenever numbers beyond zero were presented (i.e., WML and combined task) they appeared in random positions of the matrix. In addition, we further used a velocity-based criterion to identify likely artifacts; each gap in the traces was then extended by 20 ms both before and after the artifacts. Trials in which more than 40% of the data were missing were excluded, and the gaps in the remaining ones were linearly interpolated. Traces were then smoothed through cubic splines. Next, traces were down-sampled to 10 ms epochs by taking the median pupil diameter for each time bin. In order to better cope with inter-individual differences we then z-transformed pupil diameter values separately for each participant and task [15,19,38]. With normalization, a value of 0 represents the subject-specific mean

pupil diameter and, regardless of baseline values, scores represent the relative pupil size expressed as a fraction of the overall participant's variability. This is useful to account for the fact that the same amount of relative dilation or constriction (e.g., 0.1 mm) has very different meanings in participants with small vs. large baseline pupil size. Trials starting with extreme baseline values (exceeding 2 sd from the mean of the baselines of all trials) were discarded. Finally, all series were realigned by subtraction to a baseline value set to be the median of pupil size in the time window located ±250 ms from the onset of the phase of interest. Overall, data cleaning led to discard an average of: 9.5% (SD= 6.5%) trials per participant in the PLR mapping task; 11.2% (SD= 12%) trials per participant in the WML mapping task; 13.1% (SD= 14%) trials per participant in the combined task. Overall, the final pupillometric results were in keeping with classical findings regarding the pupillary light reflex [3] and psychosensory modulations (i.e., by cognitive load) [22].

Dimensionality reduction

We used R version 4.2.3 [39]. As a first step we used temporal PCA through the 'prcomp()' function, which uses single value decomposition. We did not further centered and scaled the data matrix, which included data from all participants and trials, because pupil size data were normalized and baseline-corrected beforehand, as outlined above. Temporal PCA has been used before to describe the major component(s) of pupil traces, including that of the PLR [14]. Previous studies have also attempted to link different components to different functions, for example a sustained response versus more residual responses (e.g., transient responses, cognitive variables, accommodation-related phenomena) [14,40]. However, in the case of tasks that elicit multiple, heterogeneous processes (e.g., both the PLR and psychosensory-dilation effects) this becomes less viable: PCA attempts to describe the original data in the most compact way possible, but this efficiency in data reduction sometimes hampers the interpretability of the eigenvectors as latent constructs. In the context of this study, this approach

resulted in the eigenvector of the combined task being a combination of the eigenvectors recovered in the PLR and WML tasks. This could be a convenient analytic tool to draw statistical inferences, but it appears clear that a functional interpretation of the component would be unwarranted in this case. Rather, rotations can be used to transform the best PCA solution into alternative ones which, at the cost of being less efficient in terms of data reduction (components are not "principal" anymore), yield eigenvectors that are more easily interpretable. Rotations, such as *promax* used here, can be oblique, which allows components to correlate instead of being orthogonal – often a more realistic and biologically sound scenario. Note that here we followed the '*psych*' package [41] and rescaled the eigenvectors by the square root of the eigenvalues, which makes them more comparable to factor loadings more typical of factor analysis.

Analyses

The following packages greatly eased our work: psych (v 2.4.1, Revelle, 2024); dplyr (v 1.1.4, Wickham et al., 2023); ggplot2 (v 3.4.4, Wickham, 2016); lme4 (v 1.1-35.1, Bates et al., 2015); afex (v 1.3, Singmann et al., 2023).

When analyzing pupil traces, we followed [46] by using crossvalidated Linear Mixed Effects Models (LMEM) through the lme4 package for R [44]. In this approach all trials from each participant are assigned deterministically to one of 3 folds. Two folds are circularly used as the training set; here, intercept-only LMEMs are performed for each timepoint, and the timepoint having the peak t-value (for each fixed effect or interaction) is then used in the test set to confirm the overall consistency of the target factor across folds. This approach is computationally efficient and very powerful in suggesting the presence of a consistent experimental effect *somewhere* along the time course of the trials. In order to enhance the precision in identifying a temporal cluster for any given effect, we additionally scored a consensus between folds, as timepoints in which all folds presented t-values

above |2|; timepoints in which this rather stringent, albeit arbitrary criterion was met were thus identified as a temporal cluster [38].

We then analyzed the scores obtained from dimensionality reduction techniques. The scores were obtained, for each trial, through the 'predict' method of the relevant function. For inference, we also turned to type 3 LMEMs with p-values given by the Satterthwaite approximation as implemented in afex (v 1.3, Singmann et al., 2023). In this case we could always set the matrix of random effects to be the maximal one [47,48], that is one including random slopes for all fixed effects and their interactions.

PLR mapping

There was a significant main effect of Luminance ($t_{(848.78)}$= -14.73, p<.001; peaks around 1.1 s). A clear consensus for this effect was found starting from 400 ms and lasting until the end of the trial (**Figure 2A**).

Luminance could also predict the scores of the first component (**Figure 2C**): $F_{(1, 527.75)}$= 144.77, p <.001.

WML mapping

There was a significant main effect of working memory load ($t_{(690.75)}$= 14.91, p<.001; peaks between 3 and 3.5 s. A clear consensus for this effect was found starting from 1.5 s and lasting until the end of the trial (**Figure 2D**).

Working memory load could also predict the scores of the first component (**Figure 2F**): $F_{(1, 18.66)}$= 137.82, p <.001.

Combined task

In the combined task, both main effects were significant (**Figure 3A**). There was a significant main effect of Luminance ($t_{(1541.64)}$= -9.19, p<.001; peaks around 1 s and consensus from 440 ms up to 4.2 s). There was also a main effect of memory load ($t_{(1543.15)}$= 7.69, p<.001; peaks between 2.8 s and 3 s, and consensus from 1.2 s onward). However, there was no interaction between the two ($t_{(1542.05)}$= -1.09, p=.275).

Both main effects were also significant when analyzing the first principal component scores (Luminance: $F_{(1, 64.91)}$= 28.62, p <.001; Memory load: $F_{(1, 21.97)}$= 41.18, p <.001; **Figure 3C**). In addition, however, there was a significant two-way interaction between the two ($F_{(1, 52.24)}$= 4.34, p =.042). Linear contrasts showed a trend, in the condition with the brightest luminance, for a reduced effect of working memory load, especially between 2 and 4 digits (**Figure 3C**).

The effects above were almost identical when analyzing the first fingerprint (rotated component RC1) in that both main effects (Luminance: $F_{(1, 55.78)}$= 27.21, p <.001; Memory load: $F_{(1, 26.52)}$= 43.49, p <.001) and their interaction ($F_{(1, 58.51)}$= 4.3, p =.042) were significant. The same holds for the second fingerprint, RC2 (Luminance: $F_{(1, 3.58)}$= 23.96, p =.011; Memory load: $F_{(1, 302.22)}$= 50.92, p <.001; Luminance by Load: $F_{(1, 31.48)}$= 4.74, p =.037). Results are depicted in **Figure 4B**.

However, RC3 was found to be specific for Luminance ($F_{(1, 72.74)}$= 13.88, p <.001), but not working memory load ($F_{(1, 22.09)}$= 0.63, p =.437) or their interaction ($F_{(1, 34.79)}$= 0.9, p =.348). RC3 is depicted in both **Figure 4B** and **Figure 4C**.

Acknowledgements

This publication was produced with the co-funding of the European Union - Next Generation EU, in the context of The National Recovery and Resilience Plan, Investment 1.5 Ecosystems of Innovation, Project Tuscany Health Ecosystem (THE), ECS00000017. Spoke 3. CUP: B83C22003920001.

Authors Contribution

EB: Conceptualization; Methodology; Software; Formal Analysis; Investigation; Data curation; Visualization; Writing – original draft, review and editing. **RA:** Funding acquisition; Writing – review and editing. **GA:** Writing – review and editing.

Data Availability Statement

All materials, raw data, and analysis scripts for this study are available through the Open Science Framework website: https://osf.io/dkpcs .

The core functions for preprocessing and analysis are available through GitHub: https://github.com/EBlini/Pupilla .

References

1. Campbell, F.W., and Gregory, A.H. (1960). Effect of Size of Pupil on Visual Acuity. Nature *187*, 1121–1123. https://doi.org/10.1038/1871121c0.

2. Loewenfeld, I.E. (1999). The Pupil: Anatomy, Physiology, and Clinical Applications, Vols. I and I1.

3. Mathôt, S. (2018). Pupillometry: Psychology, Physiology, and Function. J Cogn *1*, 16. https://doi.org/10.5334/joc.18.

4. Banks, M.S., Sprague, W.W., Schmoll, J., Parnell, J.A.Q., and Love, G.D. (2015). Why do animal eyes have pupils of different shapes? Science Advances *1*, e1500391. https://doi.org/10.1126/sciadv.1500391.

5. Binda, P., and Murray, S.O. (2015). Keeping a large-pupilled eye on high-level visual processing. Trends in Cognitive Sciences *19*, 1–3. https://doi.org/10.1016/j.tics.2014.11.002.

6. Einhäuser, W. (2017). The Pupil as Marker of Cognitive Processes. In Computational and Cognitive Neuroscience of Vision Cognitive Science and Technology., Q. Zhao, ed. (Springer), pp. 141–169. https://doi.org/10.1007/978-981-10-0213-7_7.

7. Laeng, B., Sirois, S., and Gredebäck, G. (2012). Pupillometry: A Window to the Preconscious? Perspect Psychol Sci *7*, 18–27. https://doi.org/10.1177/1745691611427305.

8. Mathôt, S., and Van der Stigchel, S. (2015). New Light on the Mind's Eye: The Pupillary Light Response as Active Vision. Curr Dir Psychol Sci *24*, 374–378. https://doi.org/10.1177/0963721415593725.

9. Sirois, S., and Brisson, J. (2014). Pupillometry. Wiley Interdiscip Rev Cogn Sci *5*, 679–692. https://doi.org/10.1002/wcs.1323.

10. Strauch, C., Wang, C.-A., Einhäuser, W., Van der Stigchel, S., and Naber, M. (2022). Pupillometry as an integrated readout of distinct attentional networks. Trends Neurosci *45*, 635–647. https://doi.org/10.1016/j.tins.2022.05.003.

11. Vilotijević, A., and Mathôt, S. (2023). Functional benefits of cognitively driven pupil-size changes. Wiley Interdiscip Rev Cogn Sci, e1672. https://doi.org/10.1002/wcs.1672.

12. Binda, P., and Gamlin, P.D. (2017). Renewed Attention on the Pupil Light Reflex. Trends in Neurosciences *40*, 455–457. https://doi.org/10.1016/j.tins.2017.06.007.

13. Binda, P., Pereverzeva, M., and Murray, S.O. (2013). Attention to Bright Surfaces Enhances the Pupillary Light Reflex. J. Neurosci. *33*, 2199–2204. https://doi.org/10.1523/JNEUROSCI.3440-12.2013.

14. Binda, P., and Murray, S.O. (2015). Spatial attention increases the pupillary response to light changes. Journal of Vision *15*, 1. https://doi.org/10.1167/15.2.1.

15. Blini, E., and Zorzi, M. (2023). Pupil size as a robust marker of attentional bias toward nicotine-related stimuli in smokers. Psychon Bull Rev *30*, 596–607. https://doi.org/10.3758/s13423-022-02192-z.

16. Mathôt, S., Linden, L. van der, Grainger, J., and Vitu, F. (2013). The Pupillary Light Response Reveals the Focus of Covert Visual Attention. PLOS ONE *8*, e78168. https://doi.org/10.1371/journal.pone.0078168.

17. Bradley, M.M., Miccoli, L., Escrig, M.A., and Lang, P.J. (2008). The pupil as a measure of emotional arousal and autonomic activation. Psychophysiology *45*, 602–607. https://doi.org/10.1111/j.1469-8986.2008.00654.x.

18. de Winter, J.C.F., Petermeijer, S.M., Kooijman, L., and Dodou, D. (2021). Replicating five pupillometry studies of Eckhard Hess. International Journal of Psychophysiology *165*, 145–205. https://doi.org/10.1016/j.ijpsycho.2021.03.003.

19. Dureux, A., Blini, E., Grandi, L.C., Bogdanova, O., Desoche, C., Farnè, A., and Hadj-Bouziane, F. (2021). Close facial emotions enhance physiological responses and facilitate perceptual discrimination. Cortex *138*, 40–58. https://doi.org/10.1016/j.cortex.2021.01.014.

20. Finke, J.B., Roesmann, K., Stalder, T., and Klucken, T. (2021). Pupil dilation as an index of Pavlovian conditioning. A systematic review and meta-analysis. Neuroscience & Biobehavioral Reviews *130*, 351–368. https://doi.org/10.1016/j.neubiorev.2021.09.005.

21. Hess, E.H., and Polt, J.M. (1960). Pupil Size as Related to Interest Value of Visual Stimuli. Science *132*, 349–350. https://doi.org/10.1126/science.132.3423.349.

22. Klingner, J., Tversky, B., and Hanrahan, P. (2011). Effects of visual and verbal presentation on cognitive load in vigilance, memory, and arithmetic tasks. Psychophysiology *48*, 323–332. https://doi.org/10.1111/j.1469-8986.2010.01069.x.

23. Jainta, S., and Baccino, T. (2010). Analyzing the pupil response due to increased cognitive demand: An independent component analysis study. International Journal of Psychophysiology *77*, 1–7. https://doi.org/10.1016/j.ijpsycho.2010.03.008.

24. Wetzel, N., Buttelmann, D., Schieler, A., and Widmann, A. (2016). Infant and adult pupil dilation in response to unexpected sounds. Developmental Psychobiology *58*, 382–392. https://doi.org/10.1002/dev.21377.

25. Wong, H.K., Epps, J., and Chen, S. (2020). A Comparison of Methods for Mitigating Within-Task Luminance Change for Eyewear-Based Cognitive Load Measurement. IEEE Transactions on Cognitive and Developmental Systems *12*, 681–694. https://doi.org/10.1109/TCDS.2018.2876348.

26. Wang, C.-A., and Munoz, D.P. (2015). A circuit for pupil orienting responses: implications for cognitive modulation of pupil size. Current Opinion in Neurobiology *33*, 134–140. https://doi.org/10.1016/j.conb.2015.03.018.

27. DiNuzzo, M., Mascali, D., Moraschi, M., Bussu, G., Maugeri, L., Mangini, F., Fratini, M., and Giove, F. (2019). Brain Networks Underlying Eye's Pupil Dynamics. Frontiers in Neuroscience *13*.

28. Steinhauer, S.R., Siegle, G.J., Condray, R., and Pless, M. (2004). Sympathetic and parasympathetic innervation of pupillary dilation during sustained processing. International Journal of Psychophysiology *52*, 77–86. https://doi.org/10.1016/j.ijpsycho.2003.12.005.

29. Steinhauer, S.R., and Hakerem, G. (1992). The pupillary response in cognitive psychophysiology and schizophrenia. In Psychophysiology and experimental psychopathology: A tribute to Samuel Sutton Annals of the New York Academy of Sciences. (New York Academy of Sciences), pp. 182–204.

30. Shinn, M. (2023). Phantom oscillations in principal component analysis. Proceedings of the National Academy of Sciences *120*, e2311420120. https://doi.org/10.1073/pnas.2311420120.

31. Gallego, J.A., Perich, M.G., Miller, L.E., and Solla, S.A. (2017). Neural Manifolds for the Control of Movement. Neuron *94*, 978–984. https://doi.org/10.1016/j.neuron.2017.05.025.

32. Langdon, C., Genkin, M., and Engel, T.A. (2023). A unifying perspective on neural manifolds and circuits for cognition. Nat Rev Neurosci *24*, 363–377. https://doi.org/10.1038/s41583-023-00693-x.

33. Aston-Jones, G., and Cohen, J.D. (2005). AN INTEGRATIVE THEORY OF LOCUS COERULEUS-NOREPINEPHRINE FUNCTION: Adaptive Gain and Optimal Performance. Annual Review of Neuroscience *28*, 403–450. https://doi.org/10.1146/annurev.neuro.28.061604.135709.

34. Reynaud, A.J., Froesel, M., Guedj, C., Ben Hadj Hassen, S., Cléry, J., Meunier, M., Ben Hamed, S., and Hadj-Bouziane, F. (2019). Atomoxetine improves attentional orienting in a predictive context. Neuropharmacology *150*, 59–69. https://doi.org/10.1016/j.neuropharm.2019.03.012.

35. Reynaud, A.J., Blini, E., Koun, E., Macaluso, E., Meunier, M., and Hadj-Bouziane, F. (2021). Atomoxetine modulates the contribution of low-level signals during free viewing of natural images in rhesus monkeys. Neuropharmacology *182*, 108377. https://doi.org/10.1016/j.neuropharm.2020.108377.

36. Calignano, G., Girardi, P., and Altoè, G. (2023). First steps into the pupillometry multiverse of developmental science. Behav Res. https://doi.org/10.3758/s13428-023-02172-8.

37. Mathôt, S., Schreij, D., and Theeuwes, J. (2012). OpenSesame: An open-source, graphical experiment builder for the social sciences. Behav Res *44*, 314–324. https://doi.org/10.3758/s13428-011-0168-7.

38. Blini, E., Anobile, G., and Arrighi, R. (2024). What pupil size can and cannot tell about math anxiety. Psychological Research. https://doi.org/10.1007/s00426-024-02020-0.

39. R Core Team (2023). R: A Language and Environment for Statistical Computing (R Foundation for Statistical Computing).

40. Young, R.S., and Kennish, J. (1993). Transient and sustained components of the pupil response evoked by achromatic spatial patterns. Vision Res *33*, 2239–2252. https://doi.org/10.1016/0042-6989(93)90103-4.

41. Revelle, W. (2024). psych: Procedures for Psychological, Psychometric, and Personality Research (Northwestern University).

42. Wickham, H., François, R., Henry, L., Müller, K., and Vaughan, D. (2023). dplyr: A Grammar of Data Manipulation.

43. Wickham, H. (2016). ggplot2: Elegant Graphics for Data Analysis (Springer-Verlag New York).

44. Bates, D., Mächler, M., Bolker, B., and Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. Journal of Statistical Software *67*, 1–48. https://doi.org/10.18637/jss.v067.i01.

45. Singmann, H., Bolker, B., Westfall, J., Aust, F., and Ben-Shachar, M.S. (2023). afex: Analysis of Factorial Experiments.

46. Mathôt, S., and Vilotijević, A. (2022). Methods in cognitive pupillometry: Design, preprocessing, and statistical analysis. Behav Res. https://doi.org/10.3758/s13428-022-01957-7.

47. Barr, D.J., Levy, R., Scheepers, C., and Tily, H.J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. Journal of Memory and Language *68*, 255–278. https://doi.org/10.1016/j.jml.2012.11.001.

48. Blini, E., Tilikete, C., Farnè, A., and Hadj-Bouziane, F. (2018). Probing the role of the vestibular system in motivation and reward-based attention. Cortex *103*, 82–99. https://doi.org/10.1016/j.cortex.2018.02.009.

Figure captions.

Figure 1: experimental procedures. We started with mapping the pupillary reflex to light (PLR, **panel A**). We asked participants to passively view a matrix of gray numbers superimposed on a black background. In each trial, the luminance of the numbers changed randomly to one of 8 brighter grayscale values, whereas the numbers themselves never changed. In **panel B**, on the other hand, we mapped pupil size changes to different working memory load (WML) conditions. While the luminance of the stimuli never changed, 2 to 6 numbers were selected randomly to replace the zeros, in random positions in the matrix. These target numbers remained in position for 1500 ms before being replaced by the baseline (nine "zeros") for additional 3000 ms; numbers thus had to be kept in memory throughout this duration, until a response was solicited by an auditory sound. Finally, **panel C** depicts the combined task, in which both luminance levels and cognitive load were manipulated in a 3 x 3 design. Here, either 2, 4, or 6 randomly selected numbers appeared in random positions in the matrix, for 1500 ms, and then disappeared; concurrently, the luminance of the numbers changed randomly to one of 3 brighter values.

Figure 2: mapping pupillary responses to light and cognitive load. Panels A and D depict the time course of pupil size changes in the respective tasks; the shaded areas depict the 95% confidence interval. The first principal component from each task could explain a sizeable portion of the overall variability in the data. The respective eigenvectors are depicted in **panels B and E**, and capture the primitive shapes of the processes at hand. In these panels, the color scale ranges linearly from yellow\smallest weight to red\largest weight. Both principal components' scores map efficiently a latent space along which the respective processes happen with a distinct strength. This mapping is depicted in **panels C and F** (error bars represent 95% confidence intervals).

Figure 3: mapping pupillary responses to changes of both light and cognitive load. Panel A depicts the time course of pupil size changes as a function of different working memory load

conditions (colors) and luminance levels (facets); the shaded areas depict the 95% confidence interval. This combined task was also characterized by a very low dimensionality. The eigenvector of the first component is depicted in **panel B** (the color scale ranges linearly from yellow\smallest weight to red\largest weight). The scores obtained from this component mapped very efficiently both dimensions (**panel C** depicts the mean scores, error bars represent 95% confidence interval). **Panel D** depicts the eigenvector (as in B) together with those obtained when mapping PLR and WML in isolation (as in **2B** and **2E**, though the sign was flipped for the former).

Figure 4: latent pupillary dynamics unveiled by oblique factor analysis. Regardless of the task and cognitive processes at hand, we found three remarkably similar structures behind the data (**panel A**). These three components could explain a large portion of the data and mapped efficiently both PLR and WML. **Panel B** depicts biplots with scores from pairs of the three components obtained in the combined task; plotted are mean scores (error bars are 95% confidence intervals). The first two components mapped constriction to light and dilation to cognitive load equally well. On the other hand, one component, RC3, mapped reflexes to light specifically (**panel C**).

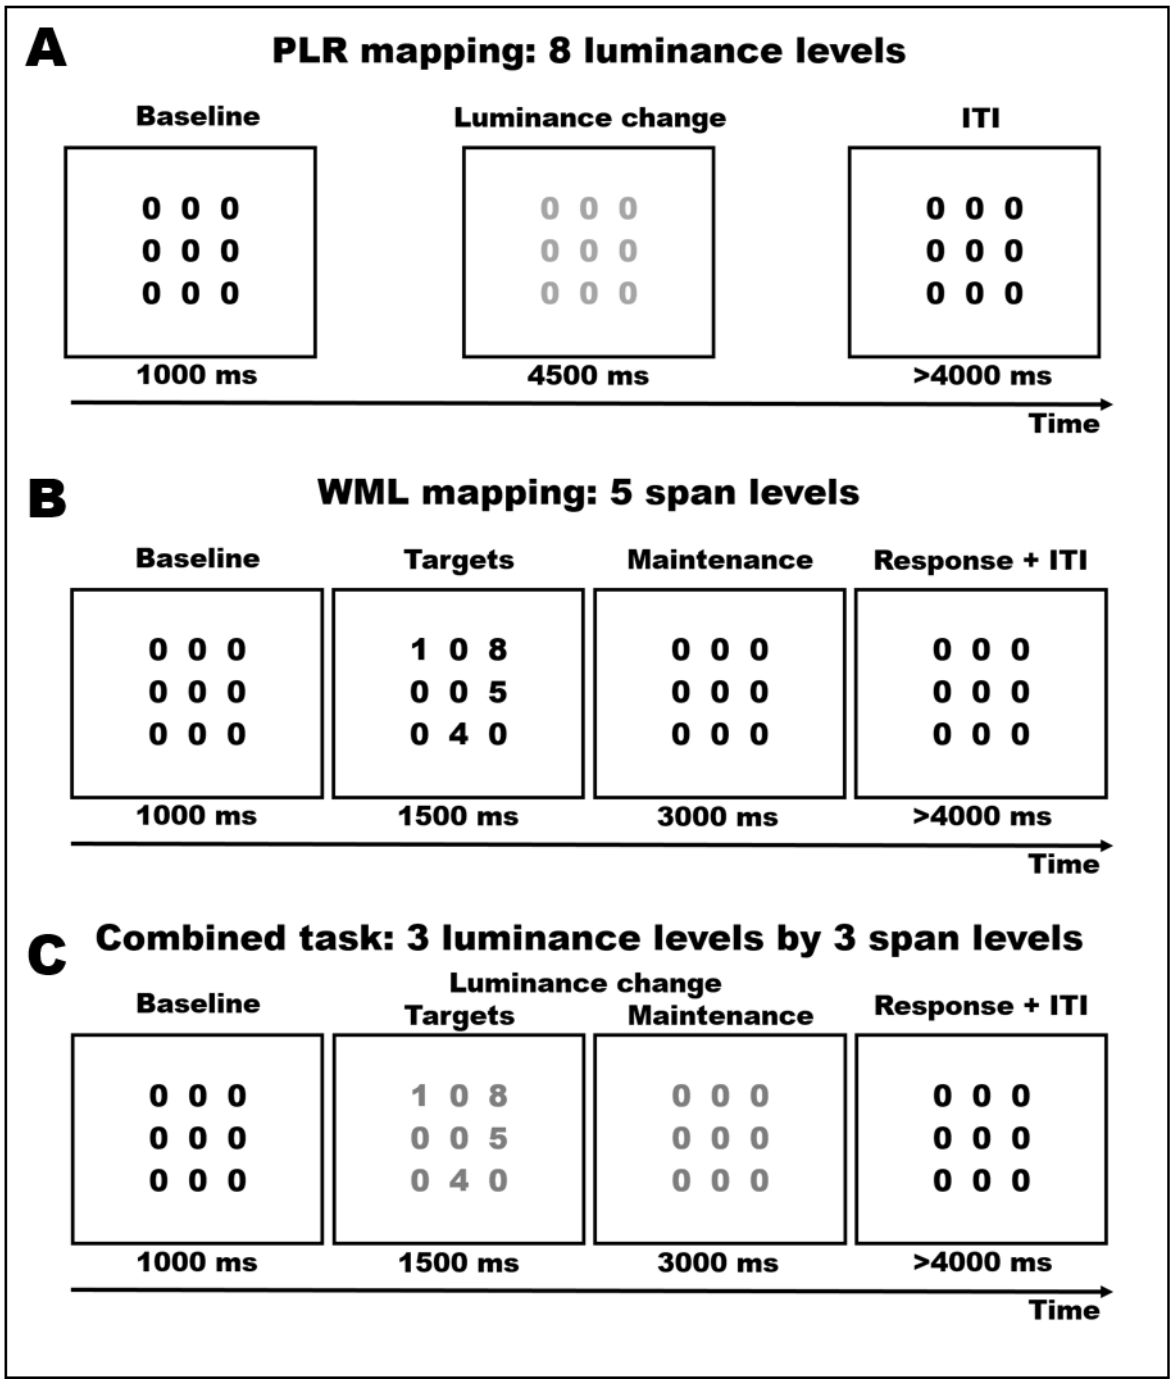
A
PLR mapping: 8 luminance levels
Baseline
Luminance change
ITI
0 0 0
0 0 0
0 0 0
1000 ms
4500 ms
>4000 ms
Time
B
WML mapping: 5 span levels
Baseline
Targets
Maintenance
Response + ITI
1 0 8
0 0 5
0 4 0
1500 ms
3000 ms
C
Combined task: 3 luminance levels by 3 span levels
Luminance change

Figure 1.

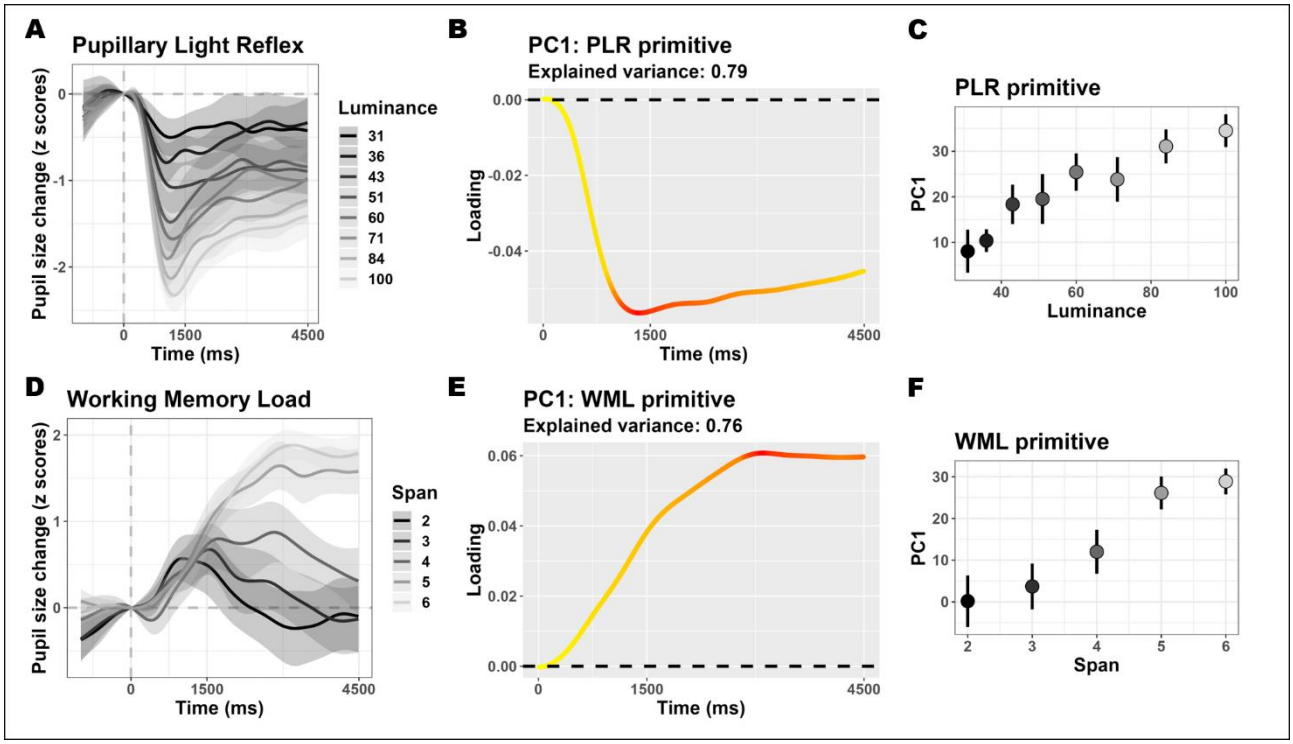
A
Pupillary Light Reflex
Pupil size change (z scores)
Time (ms)
Luminance
31
36
43
51
60
71
84
100
B
PC1: PLR primitive
Explained variance: 0.79
Loading
Time (ms)
C
PLR primitive
PC1
Luminance
D
Working Memory Load
Pupil size change (z scores)
Time (ms)
Span
2
3
4
5
6
E
PC1: WML primitive
Explained variance: 0.76
Loading
Time (ms)
F
WML primitive
PC1
Span

Figure 2.

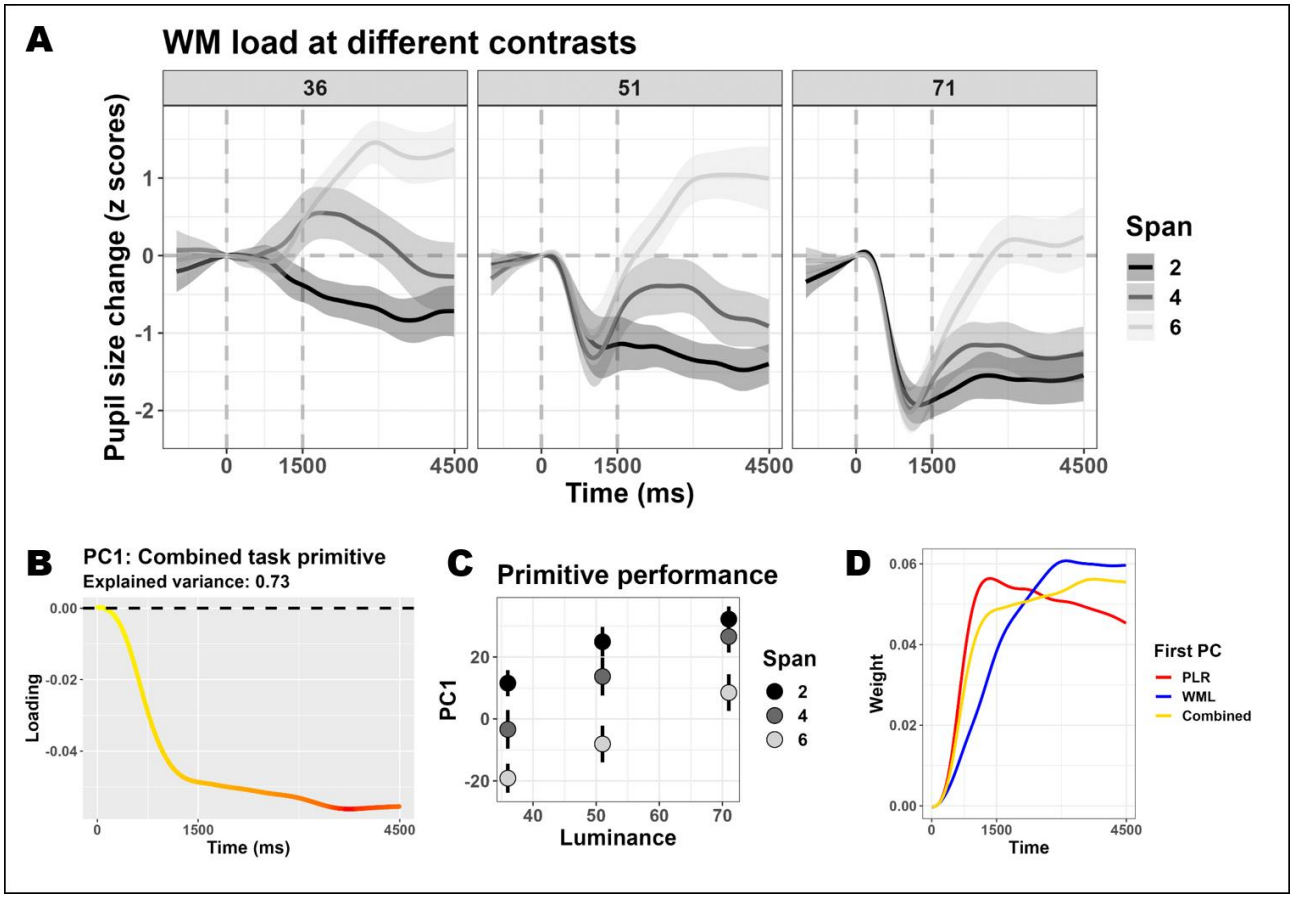
A
WM load at different contrasts
36
51
71
Pupil size change (z scores)
Time (ms)
Span
2
4
6
B
PC1: Combined task primitive
Explained variance: 0.73
Loading
Time (ms)
C
Primitive performance
PC1
Luminance
Span
2
4
6
D
Weight
Time
First PC
PLR
WML
Combined

Figure 3.

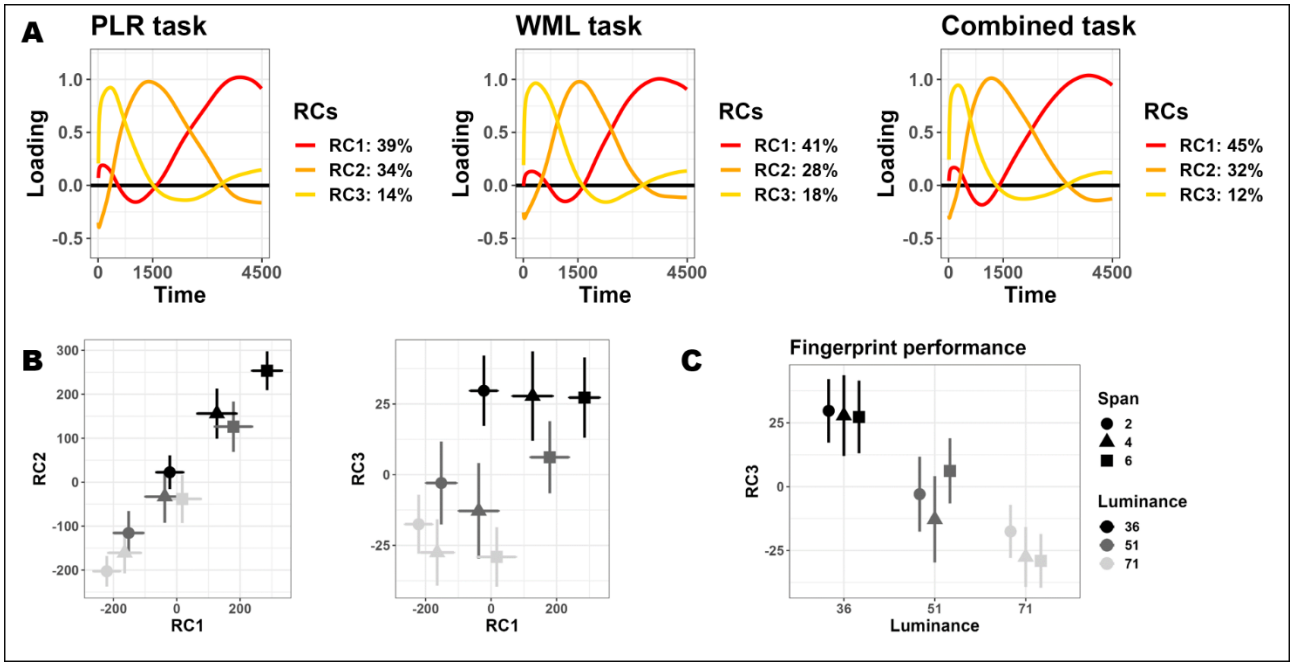
A
PLR task
WML task
Combined task
Loading
Time
RCs
RC1: 39%
RC2: 34%
RC3: 14%
RC1: 41%
RC2: 28%
RC3: 18%
RC1: 45%
RC2: 32%
RC3: 12%
B
RC1
RC2
RC3
C
Fingerprint performance
Luminance
Span
2
4
6
Luminance
36
51
71

Figure 4.

Supplementary figures captions

Figure S1: phasic pupil size changes are intrinsically low-dimensional. The figure depicts the scree plots from principal component analysis on the three tasks. Three components are generally sufficient to explain more than 88% of the overall variability in the generating data. Components above 3 account for less than 3% of the remaining variance.

Figure S2: behavioral performance in the WML mapping task. Mean (95% confidence interval) accuracy in reporting the presented numbers as a function of cognitive load. Performance was near ceiling up until 4 digits, then decreased rapidly with 5 and, especially, 6 digits.

Figure S3: behavioral performance in the combined task. Mean (95% confidence interval) accuracy in reporting the presented numbers as a function of cognitive load and luminance level. Performance was near ceiling up until 4 digits, then sensibly decreased, regardless of luminance.

Figure S4: fingerprint component of reflexes to light. One rotated component, RC3, obtained from the PLR mapping task, mapped efficiently the entire luminance space on a latent dimension. Values depict mean scores (95% confidence interval).

Figure S5: the fingerprint of reflexes to light do not map cognitive load. One rotated component, RC3, obtained from the WML mapping task, did not map distinct cognitive load conditions. Values depict mean scores (95% confidence interval).

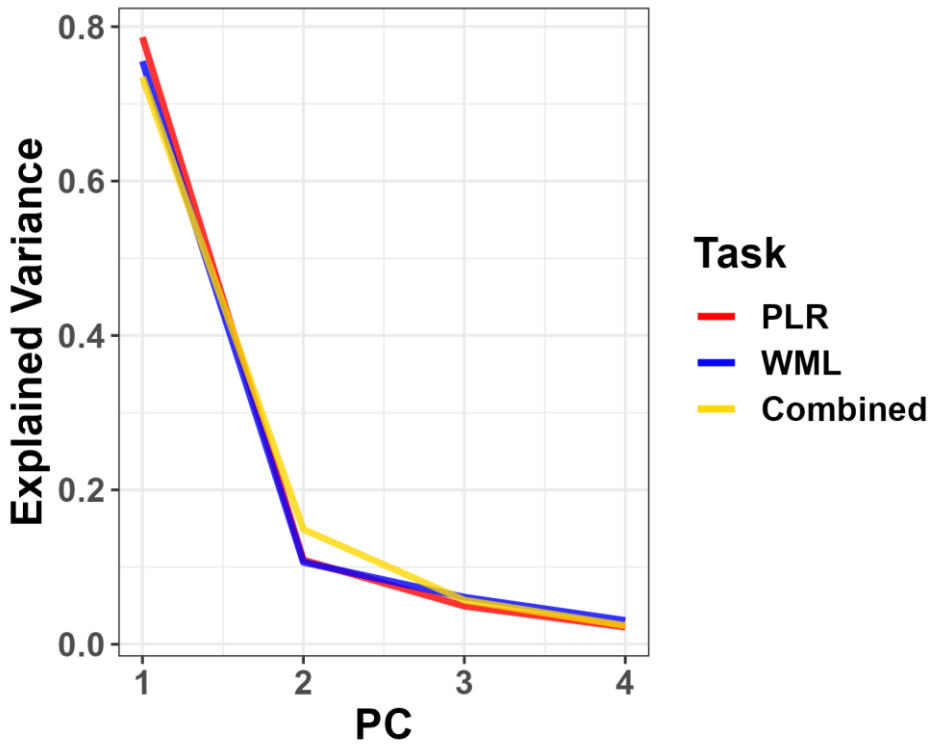
0.8
0.6
0.4
0.2
0.0
Explained Variance
1
2
3
4
PC
Task
PLR
WML
Combined

Figure S1.

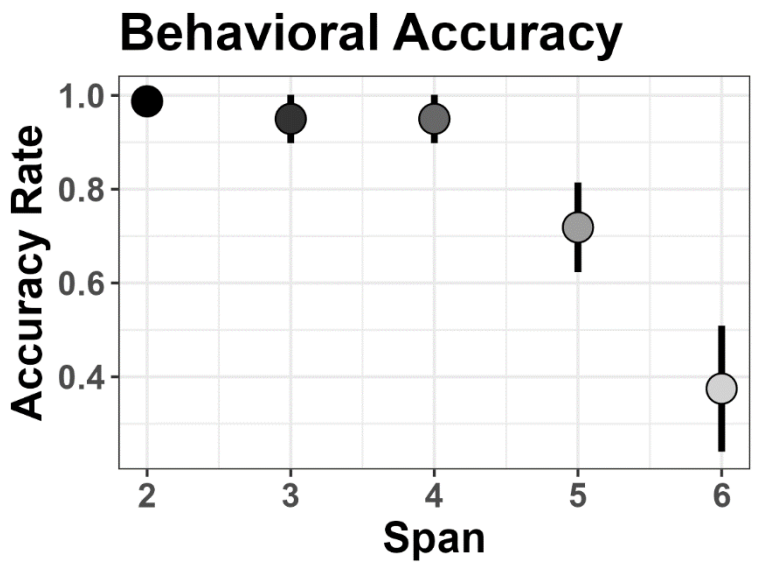

Figure S2.

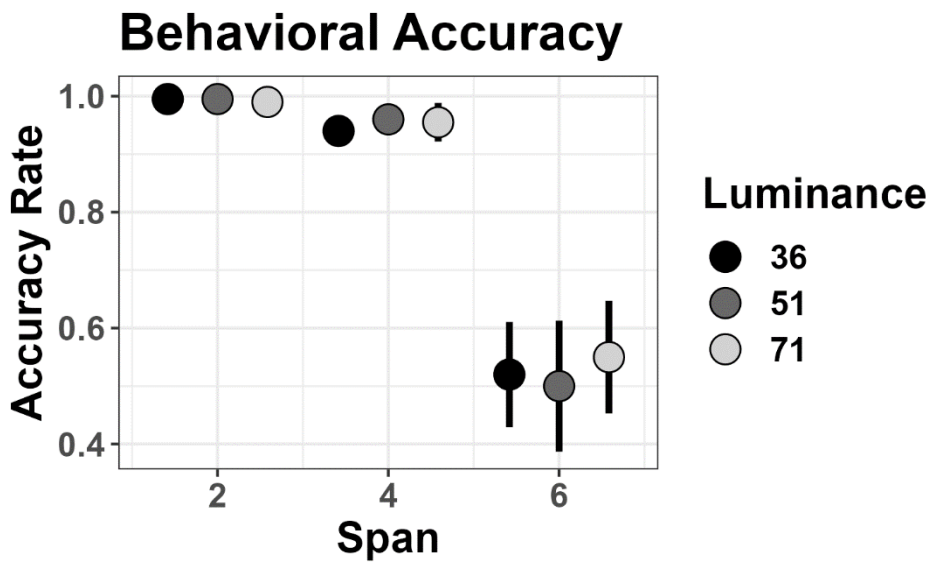
Behavioral Accuracy
Accuracy Rate
1.0
0.8
0.6
0.4
2
4
6
Span
Luminance
36
51
71

Figure S3.

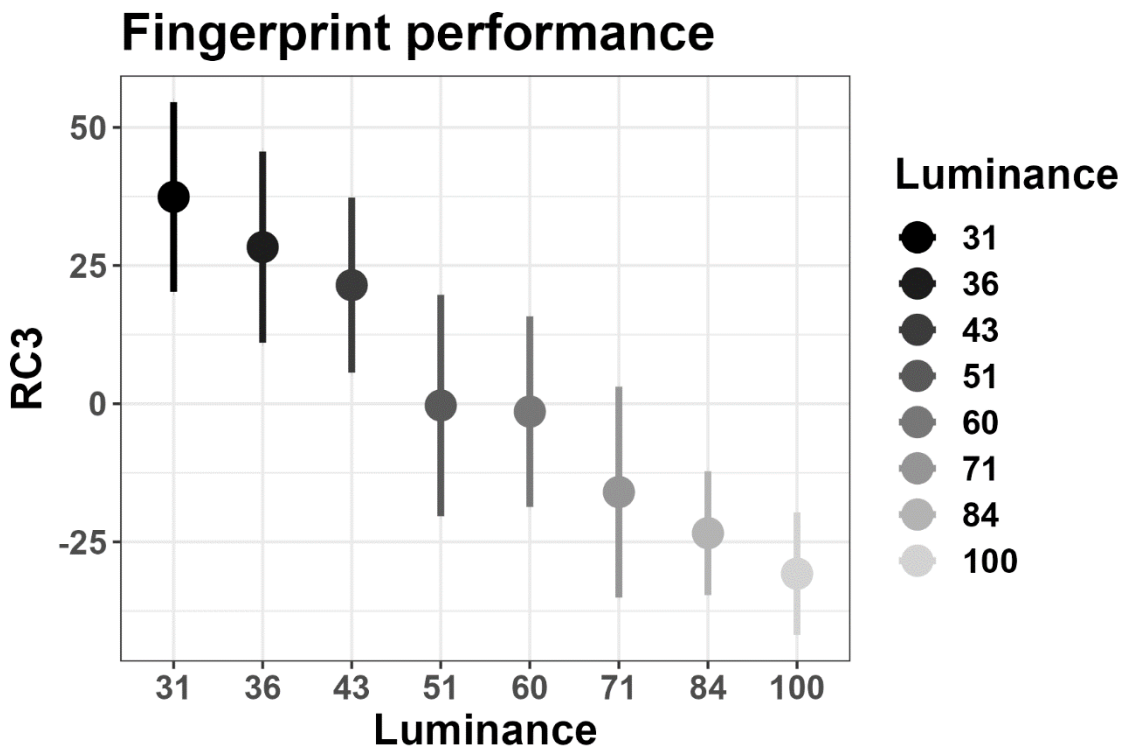
Fingerprint performance
RC3
50
25
0
-25
31
36
43
51
60
71
84
100
Luminance
Luminance
31
36
43
51
60
71
84
100

Figure S4.

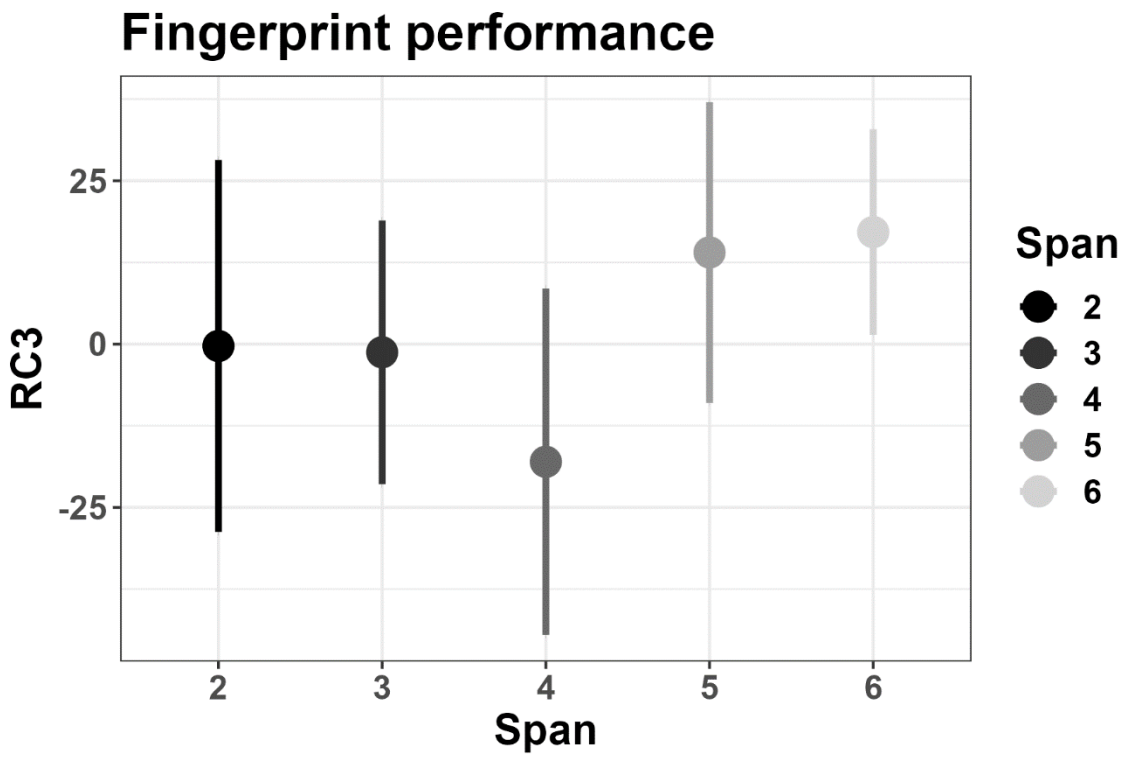

Figure S5.