natural environment. We show that the precision with which human observers estimate object weight varies as a function of both weight and volume, and this relation is consistent with the estimated joint distribution of these properties among everyday objects. Recent theoretical work has shown that optimal inference about features with a gradient of discriminability leads to perceptual biases that may counteract and even overpower the usual Bayesian attraction to the prior. We show that participants' observed "anti-Bayesian" biases (the SWI) are consistent with Bayesian inference when taking into account the gradients of precision induced by efficient coding. Related phenomena such as the material-weight illusion can be accounted for by the same principles. These results resolve a century-old paradox of multisensory perception.

Painting, Architecture, Space – Eye and body movements in VR disambiguate illusions.

Johannes Zanker[1], Harriet O'Neill[2], Doga Gulhan[1]

[1]RHUL, United Kingdom;[2]National Gallery, United Kingdom

The mutual interest between arts and visual science has a tradition, to which ECVP made large contributions. This cross-disciplinary interaction can reach a new level using new technology, like VR studies with eye tracking enabled headsets, that can shed new light on difficult and exciting questions about the experience of artwork in the context of active exploration. We previously constructed VR environments that can bring artworks into the lab to study observers' exploration behaviour, including eye movements, that can be compared to behaviour in the real museum or gallery (e.g. Gulhan, Durant & Zanker, 2021 Scientific Reports 11:18913). Here we report eye movements recorded in a VR reconstruction of a historical monument - the Carafa Chapel (Santa Maria sopra Minerva in Rome, Filippina Lippi, 1493) - which was first presented at ECVP 2022.

This renaissance 'Gesamptkunstwerk' is somewhat unique by mixing frescos, painted panels, frames and sculptures, and architectural columns, which create a dazzling environment of illusory and real objects in space. The interaction between real and illusory elements incorporated in a large, but confined space will be demonstrated for. In the lab, issues with access, illumination, and interferences by other observers are minimised, and participants can focus on the particular task – in the present case whether a column or frame element is real or illusory. The eye movement patterns recorded from the VR environment show a focus on vertical and horizontal edges that appear to indicate spatial structures that can be convex or concave – drawing in the attention of our initial participants, whilst the physical contours between the orthogonal walls of the chapel attract less interest.

Measuring the "rocking line" illusion.

Ian M. Thornton[1], Dejan Todorović[2]

[1]Department of Cognitive Science, Faculty of Media & Knowledge Sciences, University of Malta, Msida, Malta;[2]University of BelgradeLaboratory of Experimental Psychology, Department of Psychology, University of Belgrade, Belgrade, Serbia

Recently, we introduced a new motion effect called the "rocking line" illusion (RLI; Thornton & Todorović, 2023). When a single target rectangle moves horizontally through the midline of a static checkerboard pattern, its perceived path is strongly perturbed so that it appears to rock systematically around its own center. This dynamic variant of well-known polarity-dependent orientation illusions (Todorović, 2021), is probably closely related to the footstep/inchworm illusion (Anstis, 2001) and possibly the slalom illusion (Cesaro and Agostini, 1998). The RLI is robust, in the sense that all observers experience the effect, but has a strong dependence on overall display scale as well as the checkerboard frequency and contrast. The current work has two goals. First, to introduce a freely-available online resource where the illusion can be experienced and relevant parameters manipulated. Second, to describe a method for objectively quantifying the magnitude of the orientation shifts experienced during motion. To measure the RLI, we created a variant in which the target object remains fixed at the display center, while the checkerboard scrolls across the screen. The induced motion causes the physically horizontal target object to rock in place, and its fixed location makes it easy to compare to an identical adjustment probe. The probe is positioned directly below the target object, and is controlled by the observer, using a single parameter: angle of deflection. This parameter is used to drive the apparent motion of the probe, with the frequency of angular update locked to the periodicity of the scrolling background. A phenomenal match between the rocking target and apparent motion probe is easy to achieve and is surprisingly compelling. In summary, we believe the RLI, with its strong dependence on scale, is a useful tool for exploring the spatial and temporal mechanisms that underly the perception of orientation in dynamic contexts.

Sound frequencies modulate afterimage saturation

Tiziano Agostini[1], Giulia Parovel[2], Marco Prenassi[1], Serena Cattaruzza[1], Walter Coppola[1]

[1]University of Trieste, Italy;[2]University of Siena, Italy

Many studies showed that there are strong analogies between the nature of synaesthetic experiences and those reported by non-synaesthetes in imagery, matching tasks or crossmodal interference paradigms suggesting the existence of common processes between synaesthetic perception and audio-visual processing of non-synaesthetes. In the thirties, Zietz (1931) and Werner (1934) were the first to show that a sound presented simultaneously with an afterimage can affect afterimage appearance in non-synaesthetes. By asking subjects to verbally describe the temporal evolution of afterimages, they reported that the colours of afterimages “disintegrate” with low frequency sound and “concentrate” with a high frequency sound. By using a new instrumental method, we systematically investigated this phenomenon. After being exposed to the stimulus colour (blue, red, green, yellow), subjects were asked to adjust a cursor to temporally modulate afterimages saturation while listening a sound having different frequencies (low or high). The afterimage induced by the yellow stimulus is the most affected by sounds. Subjects perceive a faster colour temporal desaturation while exposed to low frequency sounds and a slower colour temporal desaturation while exposed to high frequency sounds. In contrast, the recordings concerning the stimulus with blue colour show a sudden drop, i.e. an extremely fast temporal desaturation of the colour regardless of the frequency used. These data are coherent with the crossmodal correspondences between both pitch and loudness in audition and lightness and brightness in vision reported in the literature.

The honeycomb illusion: Peripheral vision of contours, unlike that of objects, is refractory to predictions, extrapolation and memory effects.

Marco Bertamini[1, 2], Giulio Contemori[1], Luca Battaglini[1], Carolina Maria Oletto[3]

[1]University of Padova, Italy;[2]University of Liverpool, United Kingdom;[3]University of Padua, Italy

Visual acuity decreases with eccentricity, but the differences between central and peripheral vision are almost absent in our experience of the visual world. In particular, we experience uniform extended textures as uniform and objects retain shape and colour at any location. Perception of complex scenes has therefore been described as a “grand Illusion”. Various mechanisms play a role, including the fact that the system can rely on predictions based on prior assumptions (for instance that textures are uniform) or experience (from previous fixations). On the other hand, there are known examples where priors do not affect amodal completion (Kanizsa, 1980). We studied extended textures created with a square grid and some additional lines (Bertamini et al., 2016). These lines are visible/invisible depending on whether they are located at the corners of the grid, or separated from the grid (control condition). In the first case the texture appears non uniform (Honeycomb illusion, HI). Participants judged the extend of the texture with lines, and we compared cases in which the central region was informative (same) or non-informative (different). We also tested sensitivity to shape information in the periphery in a forced-choice task. The drop in sensitivity (quantified as d') matched the size of the region in which lines were perceived in the HI, even in the control case where lines were seen over the whole texture. We conclude that mechanisms that control perception of contours operate differently in the periphery, and override prior expectations, including that of uniformity. Conversely, when elements are detected in the periphery we assign to them properties based on central vision, and are unaware that these shapes cannot be identified correctly when the task requires such discrimination.

Deep reconciliation of categorical colour perception

Arash Akbarinia[1]

[1]Justus-Liebig-Universität Gießen, Germany

We perceive colours categorically. Our perceptual system separates a continuous space into distinct categories. The most prominent example is the rainbow, there are no discontinuities in its colour spectrum, yet we see discrete bands. The underlying reason, particularly the role of language, has spawned a heated debate between universalists and relativists. We reconcile these two explanations by studying vision-language and pure-vision deep neural networks (DNN). The results of our odd-one-out experiments show that pure-vision models (e.g., ImageNet object recognition networks) explain 85% of human data. In turn, suggesting a large part of our categorical colour perception is purely driven by visual signals. The remaining 15% is explained with vision-language models (e.g., CLIP text-image matching networks) even when tested without their language module. In turn, suggesting colour categories is a free-from-language representation, yet linguistic colour terms have influenced its development. We investigated whether colour categories emerge in all pure-vision models by studying Taskonomy networks trained on 24 visual tasks. Human-like colour categories appear only in less than half of those models, namely, networks trained on semantic (e.g., image segmentation, object recognition, and scene classification) or 3D tasks (e.g., shade from shaping, surface normal prediction, and depth estimation). Our results show low-level tasks (e.g., autoencoding and denoising) never obtain human-like colour categories. It also matters whether a network is trained on 2- or 3-dimensional outputs for the same perceptual task. Networks trained on 3D tasks of edge and keypoint detection obtain human-like colour categories but not their corresponding 2D networks. Overall, our findings provide evidence for the utility of categorical colour representation in several visual tasks but also indicating a portion of categorical colour perception can only be explained by the language component, reconciling both universal and relative theories.