
APPLIED VISION
ASSOCIATION
17TH APRIL 2023
UNIVERSITY OF PLYMOUTH

ABSTRACTS

The role of bias, sensitivity, and similarity on the measurement of letter acuity

Hatem Barhoom[1, 3], Mark A. Georgeson[1, 2], Mahesh R. Joshi[1], Paul H. Artes[1], Gunnar Schmidtmann[1]
[1] Eye & Vision Research Group, School of Health Professions, University of Plymouth, UK
[2] School of Life & Health Sciences, Aston University, UK
[3] Islamic University of Gaza, Gaza, Palestine
email: hatem.barhoom@plymouth.ac.uk

We have recently introduced the noisy template model to investigate the effect of bias and sensitivity in a letter acuity task performed by 10 observers in central and peripheral vision (±3˚ eccentricity, vertical meridian) (Barhoom et al., 2021, Perception (Vol. 50, No. 1_ Suppl, pp. 83-83). Here we extend this model to identify the effect of similarity between letters on letter acuity in addition to biases and sensitivity. Results show that at all three test locations, the best model was one that combined the effects of bias, sensitivity, and similarity. The model with bias and similarity was found to be more favoured than the model with bias and sensitivity. The contribution of bias was greater than both sensitivity and similarity. There was no statistically significant effect of bias, sensitivity, or similarity on the estimated letter acuity. However, in the models that included similarity, we observed a substantial increase in the spread of the psychometric function towards the upper asymptote (mainly in the peripheral test locations). In clinical vision tests, most letter stimuli are presented at supra-threshold sizes, so it is plausible to attribute differences in performance, especially in peripheral vision, to similarity alone.

Characteristics of red and blue chromatic adaptation over 4 hours

Erin Warden-English, Heidi Baseler, Antony Morland
Department of Psychology, University of York, YO10 5DD, UK
email: ee728@york.ac.uk

Long-term chromatic adaptation may involve global renormalisation of cone signals, resulting in an enduring shift of unique hues by around 5nm (Neitz et al. 2002, Neuron, 35(4), 783-792). We sought to determine what length of time was sufficient to induce a long-term shift in unique yellow settings, and whether adaptation strength was dependent on the colour of adaptation.
Participants made unique yellow wavelength settings after adaptation to red or blue filtered light. Settings were made using the method of adjustment of a small (.67° x 1.33°) rectangle on a Wright colorimeter. Participants were adapted for 15 minutes, 1 hour, or 4 hours, and their unique yellow settings were recorded before adaptation, and twice after adaptation. Our independent variables were therefore: Colour of filter (red or blue), duration of adaptation (15, 60, or 240 minutes), and measurement time (pre-adaptation 5-minutes post-, and 60-minutes post- adaptation). A three-way mixed ANOVA detected a significant three-way interaction. This was driven by an increasingly strong effect of adaptation duration on 5-minute post-adaptation settings for the red filter condition, up to a change of 4nm, but not for the blue filter condition. Irrespective of colour or duration of adaptation, unique yellow settings returned to the pre-adaptation baseline when measured at 60-minutes post-adaptation. This suggested that 4 hours adaptation was not sufficient to induce a long-term change in unique yellow settings, but the short-term effect of adaptation on unique yellow settings did follow a different pattern at different durations of adaptation when adapting to red versus blue.

The power of a smile: Idiosyncratic information revealed through a smile improves unfamiliar face matching performance

Mila Mileva[1], Jessica Hadfield[1], Mike Burton[2].
[1] School of Psychology, University of Plymouth, Drake Circus, Plymouth, PL4 8AA
[2] Department of Psychology, University of York, Heslington, YO10 5DD
email: mila.mileva@plymouth.ac.uk

Unfamiliar face matching is a surprisingly difficult task, yet we often rely on people's matching decisions in applied settings (e.g. border control). Most attempts to improve accuracy (including training and image manipulation) have had very limited success. In a series of studies, we demonstrate that using smiling rather than neutral pairs of images brings about significant improvements in face matching accuracy, even under challenging circumstances. This is true for both match and mismatch trials, implying that the information provided through a smile helps us to identify images of the same identity as well as to distinguish between images of different identities. Study 1 compares matching performance when images in the face pair display either an open mouth smile or a neutral expression. In Study 2 we add an intermediate level, closed mouth smile, in order to identify the effect of teeth being exposed and Study 3 explores the smile advantage in a more challenging task, where an age gap of 20 or 40 years has been introduced between the two images in each face pair. Results demonstrate that an open mouth smile changes the face in an idiosyncratic way which aids face matching decisions, even when one image shows someone in their 20s and the other image shows the same person (or someone else) in their 60s. Such findings have practical implications for matching in the applied context where we typically use neutral images to represent ourselves in official documents.

The interplay of perception and cognition during novel colour category acquisition

Jasna Martinovic & Yawen Liu
School of Philosophy, Psychology and Language Sciences, The University of Edinburgh, UK
email: J.Martinovic@ed.ac.uk

Colour categories are acquired through learning, but the nature of this process is not fully understood. To simulate category acquisition in adults, participants can be trained to classify exemplars from a narrow region of colour space that corresponds to the same label (e.g. GREEN; green prototype serving as the novel hue boundary). Exemplars differ in hue and lightness but participants are trained to acquire categories along one dimension only, with their subsequent discrimination performance scrutinised for the presence of 'boundary-effects' – i.e. facilitated speed or accuracy in discriminating between-boundary as opposed to within-boundary pairs. Here, instead of looking at subsequent discrimination, we focus on error patterns during learning as well as participants' verbal reports on their performance. We report on data collected in two experiments: one in which hue and lightness learners acquire boundaries without prior extended exposure to colours, and one in which they acquire boundaries after performing hue and lightness discriminations on the same stimuli. Irrespective of prior exposure, error patterns follow a similar distribution and align with the Bezold-Brücke effect, which accounts for darker colours at the green-blue boundary appearing greener and lighter colours appearing bluer. During debrief, most learners report on using labelling strategies ('yellow'-'green' vs. 'blue'; 'light' vs. 'dark') and are able to accurately describe the most difficult exemplars. Put together, our results show that simulated category acquisition is closer to relational learning than to perceptual learning: guided by both existing categories and by the non-linear change in stimulus appearance with lightness around the green prototype.

Depth estimation in real-world scenes

Michaela Trescakova, Wendy J. Adams, Matthew D. Anderson, James H. Elder, Erich W. Graf
School of Psychology, University of Southampton, UK, Center for Vision Research, York University, Toronto, ON, Canada
email: m.trescakova@soton.ac.uk

Human depth perception is typically investigated using simple shapes defined by a small number of depth cues. Using the SYNS dataset of natural scenes, we examined the time-course of depth estimation, and the relative contribution of visual features including elevation, binocular disparity and colour to ordinal and ratio depth judgments. Participants viewed briefly presented images (17 – 267 msecs) drawn from 19 outdoor categories of SYNS images under monocular (Experiment 1) or both monocular and binocular (Experiment 2) conditions. Superimposed on the images were two crosshairs that identified a pair of locations in the scene. Participants determined which of these two locations was further away, and then used a slider to report the depth of the near location as a percentage of the depth to the far location. The depth difference and mean depth of the probed locations varied from trial to trial. In Experiments 3 and 4, we also manipulated the colour of the images (natural, unnatural and greyscale), transformed using CIELuv colour space. In all experiments, when elevation cues were informative, they dominated ordinal and ratio depth responses. Participants were able to estimate local depth even with extremely brief presentation durations and both elevation and binocular disparity informed these early depth estimates. Colour effects interacted with other depth cues such as binocular disparity, presentation time and elevation. Humans outperformed multiple computational models based on these cues, including KNN and semantic segmentation models, especially when the elevation cue was absent, and for longer presentation times.

Using the natural scenes dataset to identify brain regions responsive to the colour statistics of objects in natural scenes

Ian Pennock, Jenny Bosten, Anna Franklin.
School of Psychology, University of Sussex, Pevensey 1 Building, Falmer BN1 9QH
email: ianml.pennock@gmail.com

It has been proposed that the colour properties of objects and backgrounds differ, and that the probability that a given colour is 'from' an object rather than the background is a useful cue for object vision and recognition. In support of this, macaque inferior temporal cortex is more responsive to colours which are more frequently associated with objects than backgrounds (Rosenthal et al., 2018 JOV, 18 1). However, the colour statistics of objects and backgrounds have only been quantified using one image database which might be biased. Which human brain regions are responsive to the colour statistics of objects when viewed in natural scenes is also unclear. Here we address both needs by analysing the Natural Scenes Dataset (Allen et al., 2022 Nat. Neuro. 25 116-126), a 7T fMRI dataset in which eight participants viewed up to 10,000 unique natural scenes over multiple scans. First, we analysed the chromaticities of pixels from backgrounds and 80 segmented objects. We found that object pixels were warmer, redder (L/(L+M)), more saturated and darker (L+M) than the background. Second, pixels were categorised into 240 colour 'bins', and the probability that pixels were from objects was calculated. An average object-colour probability for each image was calculated based on the colour of the object pixels (high score indicates image has typically coloured objects). We then correlated the average object-colour probability of each image with the BOLD response. We tentatively suggest that the visual ventral pathway is responsive to the colour statistics of objects in natural scenes.

Substantial individual differences in facial expression perception

Isabelle Mareschal
Queen Mary University of London
email: i.mareschal@qmul.ac.uk

Facial expressions can convey critical information about our emotional states. Surprisingly however, there are profound individual differences in how we interpret another's facial expressions that current methods are unable to quantify, because facial expressions exist in a space that is too large to explore parametrically. To examine this, we recently developed new tools, based on a genetic algorithm, that allows participants (n>300) to quickly and efficiently create computer generated facial expressions associated with 4 emotional states (sad, happy, fear and anger). Faces were defined by 150 blendshapes, based on the Facial Action Units (FACS), and that controlled different facial features. We found substantial individual differences in the facial expressions created across emotion categories. Using Gaussian Mixture Modelling (GMM), we identified 4 normally distributed clusters and found that happy expressions were the most reliably classifiable (97% correct), but found that anger (82%), and particularly fear (63%) and sad (58%), were less reliably classifiable. 'Sad' was confused with 'fear' in 31% of instances, and 'fear' was confused as 'angry' in 24% of instances. This supports the idea that while happy and angry expressions are distinct, sad and fear are less so, and have important repercussions for the interpretation of standard emotion recognition tasks.

Higher internal noise is associated with lower sensitivity to slower speed

Mahesh R Joshi[1], Charlotte Barsdell[1], Anita J Simmers[2], Seong T Jeon[2]
[1] Eye and Vision Research Group, School of Health Professions, University of Plymouth, Plymouth, UK
[2] Vision Sciences, Department of Life Sciences, Glasgow Caledonian University, Glasgow, UK

Humans have better sensitivity to faster speeds compared to relatively slower speeds. The difference in sensitivity suggests the presence of independent processing channels for slow and fast speeds or a single channel tuned to differential sensitivity at varying speeds. In this study, we employ motion coherence and equivalent noise paradigms to investigate if differences in sensitivity to slow and fast speeds are due to varying internal noise or sampling efficiency.
A total of 8 visually normal participants discriminated the direction of motion of RDK with random noise (motion coherence) and in the presence of four noise levels generated by assigning direction of dot elements from a Gaussian distribution with different standard deviations (equivalent noise paradigm) at three dot speeds (2.5°/s, 5°/s, and 10°/s). The mean motion coherence thresholds were lowest for the dot speed of 10°/s and highest for 2.5°/s. For sensitivity at variable noise levels, the thresholds for dot speed of 10°/s were lowest, followed by 5°/s and 2.5°/s at no noise condition. The thresholds for all dots speeds however converged at the highest noise level. The nested models to differentiate the performance in terms of internal noise and sampling efficiency parameters showed that the model with difference in internal noise best described the threshold data. Our findings suggest that the sensitivity to fast and slow speeds is dependent upon the internal noise within the processing channels.

Altered high-level chromatic perception in Parkinson's disease

Ben J. Jennings[1], Salma Nimuchwala[1], Fredrick Kingdom[2], Andrew Parton[1]
[1] Brunel University London
[2] McGill University, Montreal, Canada
email: ben.jennings@brunel.ac.uk

Parkinson's Disease is a neurodegenerative disease caused by a loss of nerve cells within the brain resulting in depleted dopamine. The disease is characterised by symptoms including slowness of movement and tremors. Additionally, measurable changes in visual performance have been reported in Parkinson's patients. Chromatic sensitivity has been assessed with low-level stimuli, with studies indicating a sensitivity loss. In the current study we employed a high-level stimulus - images of natural scenes - with manipulated chromatic saturation as the dependant variable. Participants made judgments on each trial as to whether the chromatic content was over- or under-saturated. Results indicate that the Parkinson's group required an increased saturation in-order to judge the colour content of the scenes to be the 'correct' level, compared to unaltered images that were judged as under-saturated. Both an age and depression matched control group and a young control group judged the unaltered scenes as containing the correct saturation level. Additionally, a significant negative correlation existed between the time since diagnosis and the psychometric function slopes indicating an increase in response uncertainty as the disease progresses.

Can I detect cataract induced colour sensitivity changes with colour screening tests?

Zane Jansone-Langina & Maris Ozolinsh
University of Latvia, dept. of Optometry and Vision science, Riga, LV-1004, Latvia
email: jansonezane1993@gmail.com

Introduction: In the current literature, there have not been attempts to determine colour vision sensitivity shift due to cataract with classical screening tests like unsaturated Farnsworth D15 test and Hardy Rand and Rittler. Method This study analysed 54 patients with nuclear cataract (mean age 62.5 ± 0.7 years). The surgery was performed for both eyes with a two day delay between surgery of each eye. As a colour vision test we used unsaturated Farnsworth D15 panel test and Hardy Rand and Rittler (HRR) pseudoisochromatic plate tests. It comprises a 6 plate screening section plus 14 detailed diagnostic plates to determine the type and extent of colour deficiency. Results: Before the surgery 6 patients showed colour sensitivity shift to the tritan confusion axis side, did not show any colour sensitivity abnormalities after one week. Four eyes which showed normal cap arrangement sequence before the surgery, after showed nonspecific colour sensitivity changed closer to protan confusion axis. Before and after cataract surgery HRR test showed no colour vision deficiency problems. Conclusion: Before and 2 weeks after cataract surgery patient's colour vision sensitivity changes. Using unsaturated Farnsworth D15 it is possible to detect colour sensitivity shift, which cannot be done with Hardy Rand and Rittler test.

The time course of stimulus-specific perceptual learning

Patrick J. Bennett[1], Ali Hashemi[1], Jordan W. Lass[1], Allison B. Sekuler[2], Zahra Hussain[3]
[1] *Department of Psychology, McMaster University, 1280 Main Street West, Hamilton, Ontario, L8S 4K1, Canada*
[2] *Rotman Research Institute, 3560 Bathurst St, Toronto, Ontario M6A 2E1, Canada,*
[3] *School of Psychology, University of Plymouth, Portland Square, Plymouth, PL4 8AA, United Kingdom*
email: zahra.hussain@plymouth.ac.uk

Perceptual learning refers to the improvement in performance in perceptual tasks that occurs with practice. Perceptual learning usually is specific to the trained stimuli, but some proportion of improvement does generalize to novel stimuli. The time courses of general and stimulus-specific components of perceptual learning (PL) are still debated. Some researchers argue that general learning occurs first in an initial rapid phase of learning and that stimulus-specificity emerges slowly over the course of hundreds or thousands of trials, but others have found stimulus-specific learning with relatively few trials. We examined if rapid stimulus-specific perceptual learning occurs in a texture identification task, and if such learning persists over intervals of one day and one week. Participants completed 21, 63, 105, or 840 practice trials on Day 1 and then, one day or one week later, completed 840 test trials on Day 2. On Day 2, participants either saw the same stimuli presented on Day 1 or a novel set of stimuli. We found that 1) Accuracy on Day 2 in the Same stimuli condition, but not the Novel stimuli condition, increased as a linear function of the number of practice trials, 2) individual differences in initial accuracy on Day 1 were strongly associated with accuracy on Day 2, and 3) the effects of individual differences, practice, and stimulus novelty differed slightly when testing sessions were separated by 1 day and 1 week. These results support the idea that long-lasting, stimulus-specific perceptual learning emerges more-or-less continuously during practice.

The impact of notch filters on colour perception in anomalous trichromacy

Jenny Bosten & Lucy Somers
Sussex Vision Lab, School of Psychology, University of Sussex, Brighton, UK
email: j.bosten@sussex.ac.uk

Anecdotal reports imply that people with anomalous trichromacy can experience radically enhanced colour vision when using notch filter aids. However, existing empirical research has largely focussed on their effects on performance on diagnostic tests for colour vision deficiency, and has not found any substantial effects. We created a model to predict the effects of EnChroma notch filters on anomalous trichromats' colour perception for different combinations of observer, surface reflectance, illumination, filter, and natural scenes. Our model predicted that EnChroma filters do usually confer colour gamut expansions, but their precise effects depend on the particular light spectra. We tested the model predictions on 10 anomalous trichromats in three experiments: (i) colour discrimination, (ii) asymmetric colour matching between test and control filter conditions, and (iii) colour appearance from pairwise colour dissimilarity ratings. To investigate potential effects of long-term adaptation or perceptual learning, participants completed the experiments on first exposure to the filter, and after a week of regular use. We found a significant effect of the filter on colour matches in the direction predicted by the model, implying that they do have the effect of enhancing the anomalous trichromatic colour gamut. However, we found minimal effect of the filter on colour discrimination. We found an effect of the filter on enhancing the red-green axis in subjective colour space, but only at the first time point, and not after a week of regular use. Our model and empirical results provide evidence that notch filters can enhance colour perception for anomalous trichromats.

Information decoding reveals non-uniformities in neurometric hue space

Ana Rozman[1] & Jasna Martinovic[2]
[1] School of Psychology, University of Aberdeen, King's College, Aberdeen AB24 3FX
email: a.rozman@abdn.ac.uk
[2] School of Philosophy, Psychology and Language Science, The University of Edinburgh, 7 George Square, Edinburgh EH8 9JZ
email: jmartino@ed.ac.uk

The powerful multivariate analysis technique of information decoding has previously been applied to characterise the neural representation of hue (e.g. Chauhan et al., 2023 Neuroimage 268). As these studies only included a limited number of hues, the uniformity of neurometric hue space has not been fully examined. We expanded on this by characterising the neurometric hue space for 8 equidistant hues taken from the perceptually uniform CIELAB space. Electroencephalograms (EEGs) were recorded for 15 observers who viewed four cardinal (red, green, blue, yellow) and four intermediate (orange, lime, turquoise, purple) isoluminant hues. If neural hue representations correspond to perceptual distinctions as captured by CIELAB, decoding should be uniform across hues and signals of neighbouring hues should be more similar to each other and more dissimilar to more distant hues. Decoding was performed within several decoding manifolds: 1) quadruples of neighbouring hues and 2) quadruples of opposite hues. This enabled comparison of representation between hues in terms of their local and global context. Decoding accuracy peaked ~100-300ms, coinciding with the chromatic visual evoked potential (cVEP) and exhibited local asymmetries, including best decoding performance for purple and worst for green, as well as higher likelihood for hues to be mislabelled as their opposite, with red and orange being an exception. This indicates neural representation of hue reflected by EEG signals is non-uniform and at least partly driven by opponency. Future analyses will examine the degree to which decoding is driven by lower and higher-level attributes such as cone-opponency and unique hue content.

Measuring symptoms in research: Have we got it wrong?

Bruce JW Evans
Optometry & Visual Sciences, City, University of London; Northampton Square, London, EC1V 0HB
email: Bruce.Evans.1@city.ac.uk

For clinicians, the presenting symptom or "reason for visit" is paramount as this is often how the patient judges the clinician's success. In contrast, researchers typically assess symptoms using questionnaires. For example, the Convergence Insufficiency (CI) Symptom Survey (CISS) asks 15 questions about symptoms suggestive of CI. All 15 symptoms are non-specific (could result from other conditions) and some are more likely to result from other conditions than CI (e.g., "Do you feel sleepy when reading?"; "Do you have trouble remembering what you have read?"; "Do you feel like you read slowly?"). Questionnaires with non-specific questions, which might not have been reported by a given patient, will increase "noise" in symptom scores. For CI, this may explain the low correlation between CISS scores and signs of CI (Bade et al., 2013 Optom.Vis.Sci 90 988-995) and the finding that vision therapy improves signs of CI but not CISS scores (Scheiman et al., 2020 Cochrane Database of Systematic Reviews 12). A questionnaire approach is removed from the patient's main concern; whether their presenting symptom is alleviated. The literature will be reviewed for alternative ways of scoring symptoms. It will be argued that a more relevant approach is to consider, for each participant the extent to which their presenting symptom(s) has been addressed, using a simple scale of 1-10. CI is particularly problematic, because its symptoms

may manifest idiosyncratically as a variety of asthenopic symptoms and/or blur/diplopia. The extent to which these issues apply to other conditions (e.g., refractive error) will be considered.

Adapting to retinal neurodegeneration: glaucoma patients make the best of their remaining binocular function

Guido Maiello[1] & MiYoung Kwon[2]
[1] School of Psychology, University of Southampton, B44 University Rd, Southampton SO17 1PS
[2] Department of Psychology, Northeastern University, 107 Forsyth St, Boston, MA 02115
email: g.maiello@soton.ac.uk

In healthy vision, small changes in disparity-defined depth are detected foveally, peripheral vision captures the coarser structure of the environment, and cortical mechanisms combine information into a unified depth percept. Glaucoma, a blinding neurodegenerative disease, is putatively associated with peripheral vision loss, yet accumulating evidence suggests that even early glaucomatous injury may involve the macula. Here, we examine how binocularly asymmetric glaucomatous damage affects binocular disparity processing by comparing disparity sensitivity in glaucoma patients and controls (N=47). After confirming that glaucoma patients exhibited binocularly asymmetric visual field damage ($p<.001$), we employed a validated psychophysical procedure (Maiello et al., 2017 PLoS Comput Biol 16(4) e1007699) to measure how disparity sensitivity varied across visual field sectors, i.e. with full-field stimuli spanning the central 21°, and with stimuli restricted to annular regions spanning 0°-3°, 3°-9° or 9°-21°. We employed measurements with annular stimuli to model different possible scenarios regarding how disparity information is combined across visual field sectors. We adjudicated between potential mechanisms by comparing model predictions to the patterns observed with full-field stimuli. Contrary to the conventional view that glaucoma spares central vision, we found that glaucomatous damage caused widespread sensitivity loss across both foveal and peripheral vision ($p<.01$). However, similar to controls, glaucoma patients exhibited near-optimal disparity integration ($p<.001$), revealing that neurodegeneration had spared cortical processing. Our results thus suggest that glaucoma patients make the best possible use of their remaining binocular function, and pave the way for potential treatment opportunities for preserving or restoring binocular function in early disease stages.

Visual guidance of instinctive behaviours in mouse

Samuel G. Solomon, Hadrien Janbon, Adam Bimson, Thomas Wheatcroft
Institute of Behavioural Neuroscience and Department of Experimental Psychology, University College London, London, WC1H 0AP, United Kingdom
email: s.solomon@ucl.ac.uk

Visual stimuli can elicit instinctive approach and avoidance behaviours. In mouse, vision is known to be important for both instinctive avoidance of an overhead threat, and approach toward a potential terrestrial prey. Here we asked how mice responded to visual stimuli presented either above the animal or near the ground plane. We found that a looming black disc – emulating the silhouette of a rapidly approaching potential threat – induced escape behaviour when presented overhead or to the side of the animal. However, the escapes produced by side-looms were less vigorous and were often preceded by freezing behaviour. Similarly, small moving black discs – emulating a surveiling potential threat or even a potential prey – induced freezing behaviour when presented overhead or to the side of the animal. However, side-sweeps also elicited approach behaviours, such that mice explored the area of the arena near where the stimulus had been presented. Our observations therefore show that

mice combine cues to the location and features of visual stimuli when selecting among potential instinctive behaviours.

Shape-related size biases in visual area judgements

Veronica Pisu[1], Sina Mehraeen[2], Erich W. Graf[1], Marc O. Ernst[2], Wendy J. Adams[1]
[1] Centre for Perception and Cognition, Psychology, University of Southampton, Southampton, UK
[2] Department of Applied Cognitive Psychology, Ulm University, Ulm, Germany
email: V.Pisu@soton.ac.uk

Every day, we effortlessly interact with a wide range of objects of different size and shape, yet shape-related size biases have long been reported. For example, in area judgements, observers consistently judge triangles as larger than same-area squares and disks. Several authors have proposed that the biases are explained by the magnitude of one geometric features of the object, or the use of heuristic strategies (e.g., combining the object's height and width), but the picture is complicated by inconsistent conclusions from studies focusing on specific subsets of shapes. Here we explored biases in area perception, across different shapes, allowing us to test the influence of a wide range of geometric features (height, compactness, elongation, etc.). In four online experiments, observers made 2AFC judgements ("which stimulus has larger area?") for pairs of objects that differed in shape and / or area and / or orientation. We found clear shape-related biases: for example, triangular shapes were systematically perceived to have larger area than square shapes; disks were perceived to have the smallest area. When orientation was manipulated, the same shape was perceived to be larger or smaller depending on its orientation. Overall, no single geometric feature – or simple combination of two features – provides a good correlate of the biases across all shapes, and a more complete description incorporating additional geometric features is needed to explain shape-related biases in area perception. We provide a multi-predictor model that quantitatively predicts biases in perceived area across 76 shape / size / orientation combinations.