Taylor & Francis
Taylor & Francis Group

Spatial Cognition & Computation

An Interdisciplinary Journal

ISSN: 1387-5868 (Print) 1542-7633 (Online) Journal homepage: http://www.tandfonline.com/loi/hscc20

Perception, cognition and reasoning about shadows

Paulo E. Santos, Roberto Casati & Patrick Cavanagh

To cite this article: Paulo E. Santos, Roberto Casati & Patrick Cavanagh (2018) Perception, cognition and reasoning about shadows, Spatial Cognition & Computation, 18:2, 78-85, DOI: 10.1080/13875868.2017.1377204

To link to this article: https://doi.org/10.1080/13875868.2017.1377204

Accepted author version posted online: 12 Sep 2017.
Published online: 12 Sep 2017.

SPATIAL COGNITION & COMPUTATION
2018, VOL. 18, NO. 2, 78–85
https://doi.org/10.1080/13875868.2017.1377204

Perception, cognition and reasoning about shadows

Paulo E. Santos [a], Roberto Casati [b], and Patrick Cavanagh [c,d]

[a]IAAA, Electrical Engineering Department, Centro Universitário da FEI, São Paulo, Brasil; [b]Institut Jean Nicod, Département d'études cognitives, ENS, EHESS, CNRS, PSL Research University Paris, France; [c]Laboratoire Psychologie de la Perception et CNRS UMR 8242, Université Paris Descartes, Sorbonne Paris Cité, Paris, France; [d]Department of Psychological and Brain Sciences, Dartmouth College Hanover, NH, USA

KEYWORDS shadows; perception; spatial reasoning; spatial perception

1. Introduction

The way we categorize shadows interfaces with the way we perceive them (Casati, 2017). Some thought experiments such as the Yale Problem (Todes & Daniels, 1975), the Intersecting Eclipses Problem (Sorensen, 2008), and the Shadow-Light Problem (Casati, 2007b) suggest that the conceptual boundaries of shadow phenomena are still not completely explored. To take just the last example, it is conceptually indeterminate whether the green spot that is cast by a green bottle intercepting light is a green light spot or a green shadow (arguments can be put forward to defend either view). Conversely, perceptual double dissociations (some physically impossible shadows are tolerated and some physically correct shadows are seen as awkward) challenge the conceptual system, as in some cases not only there is tolerance, but even preference for impossible shadows.

The conditions in which a dark patch near an object is perceived as the object's shadow have been taken into consideration in various psychophysical experiments (Kersten, Knill, Mamassian, & Bülthoff, 1996; Kersten & Mamassian, 2017) that suggest that (in some situations) cast shadows are preferred over other cues for judging motion (and location) in depth. There is also evidence that six- and seven-year-old infants use cast shadows of objects to discriminate the object's trajectories (Imura et al., 2006; Yonas & Granrud, 2006). In contrast, recent research suggests that shadows are used as coarse cues, in the sense that they are processed fast by the perceptual system and then discarded (Ehinger, Allen & Wolfe, 2016; Elder, Trithart, Pintilie & MacLean, 2004; Rensink & Cavanagh, 2004).

After a region is labelled as a shadow and the scene information is extracted, what then happens to the shadow regions? They are not objects themselves, merely accidents of light and the presence of their contours is a source of noise for perceiving and interacting with the objects in the scene. Indeed, the visual system had to evolve to recognize objects in arbitrary

CONTACT Paulo E. Santos psantos@fei.edu.br Centro Universitário da FEI, Av. Humberto de A. C. Branco, 3972, SBC, Sao Paulo, Brazil, CEP 09850-901.

lighting and so must be able to ignore cast shadows. This has led to a view that cast shadows are expendable, not especially useful after their information is extracted, and potentially distracting. Note that these demoted shadows do not vanish, they are just less salient.

A number of visual search studies have supported the argument for the discounting or demotion of shadows. In one, several shapes are presented, each casting a shadow (Figure 1(a)), and on half of the trials, one of the shadows is misaligned with the others. The task is to report whether the misaligned shadow is present. The search for the odd shadow in this case is relatively slow. In the control condition (Figure 1(b)), the cast shadows are inverted in polarity, light on the grey background rather than dark. Now the item with the odd orientation is easier to find, presumably because the lighter areas are not treated as shadows and are not discounted. The hypothesis is that the early visual system rapidly labels regions as shadows, extracts relevant scene information, and then discounts them, making their shapes and contours less of a nuisance in analysing the scene and paying attention to the objects of interest. As a result, the properties of the shadows themselves are more difficult to access. Note that the odd direction of the target shadow did not veto it as a shadow, otherwise the shadow would be easy to find.

The hypothesis of cast shadows as expendable objects has been reinforced in a recent study (Bouwmans, 2014) that, using a change blindness paradigm, investigated how well people could observe changes in the cast shadows of natural scenes when the shadow changes imply movement or disappearance of scene objects, thus, affecting the number of objects or their spatial relations in the scene. The results show that the observations of such informative shadow changes (e.g., the deletion of the shadow of an occluded object) were reported less often than the deletions of visible objects in the scene. The possibility that the shadow processing in the human perceptual system is implicit and automatic motivated the development of shadow-observation drills to be used as culturally independent tests for the identification of gifted children (Tan et al., 2012).

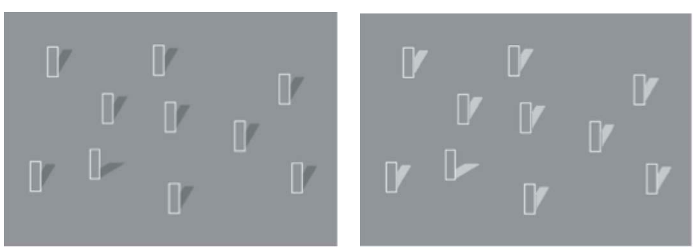

(a) (b)

Figure 1. The odd dark shadow on the left is harder to find than the geometrically identical odd lighter tab on the right. The assumption is that on the left, the dark areas are labeled as shadows and demoted as features of light rather than features of objects. Image credit: PC adapted from Rensink and Cavanagh (2004).

Notwithstanding the hypothesis of an early processing of shadows by the human perceptual system, recent research provides evidence that there is some level in the human object recognition process at which, under certain conditions, the stored representation of object's shape also contains the object's shadow (Leek, Davitt & Cristino, 2015). This seems to be related to the evidence that the shadow of a person's body contributes to the connection between personal and extrapersonal space and triggers an automatic focus of attention toward the body part casting the shadow (Pavani & Galfano, 2015). Khuu, Gordon, Balcomb and Kim (2014) present results indicating that visual awareness is necessary for the perception of illusory motion from changes in cast shadows. The authors suggest that the perception of cast shadows is guided by two processes, one implicit that is responsible for the recognition of objects, and another demanding visual awareness to infer the object's 3D location.

In spite of the importance of cast shadow perception in our understanding of the 3D world, it has been observed in art history studies that the depiction of shadows (once present in the Hellenistic and Roman paintings and mosaics) was almost completely ignored in Western art after the fall of the Roman Empire. The depiction of shadows reappear in the artistic rendering of natural scenes in an experimental phase circa 1400, in which artists try out a number of solutions to the convincing representation of shadows, as clearly shadows were felt to add great value to the realism of the depicted scene. However, in this experimental phase (that closes with the onset of mathematical algorithms for calculating the depiction of shadows, circa 1515) a great number of depictions convincingly represent shadows that are physically impossible. This may be explained by the difficulty of depicting the imprecise borders of shadows (Casati, 2007b, 2006). The fact that the human perceptual system is insensitive to shadow imperfections may indicate some intrinsic characteristics of the proprietary neural code used by visual perception (Cavanagh, 2005).

Although shadows are unavoidable parts of the visual (ecological) world, and serve as essential information for the human understanding of our three-dimensional environment (Cunningham, Beck & Mingolla, 1996; Khuu, Honson & Challinor, 2016), research in computer vision has traditionally considered shadows as noise to be suppressed from images (Dee & Santos, 2011).

Recently there has been a growing interest in the development of shadow detection algorithms to improve the performance of computer vision applications such as object tracking, identification, to infer the geometry of shadow casters, to find the light source position and in the 3D reconstruction of scenes from static images (Sasi & Govindan, 2015) (including satellite images (Elbakary & Iftekharuddin, 2014)), video (Russell, Zou & Fang, 2016; Sanin, Sanderson & Lovell, 2012) or from radar imagery (Prasath & Haddad, 2014; Raynal, Bickel & Doerry, 2014). In fact, the increasing interest in computer vision methods for detecting shadows is reflected by the number of citations that a recent review article on shadow detection (Sanin et al., 2012) received (over 200 citations in the four years from its publication and the date of writing this introduction). The article by Newey, Jones and Dee (2017) (this issue) investigates shadow modelling from video assuming a moving viewpoint.

The authors investigate two methods for shadow detection: a texture-based and an edge-based segmentation method, within a machine learning framework. The key idea of the texture-based method is to use a visual feature that is invariant under shading (the *texture*) to obtain the visible surfaces in the scene, and then use brightness to determine which of these surfaces can be classified as a shadow. In the edge-based method, the authors explore the differences in values across the set of edges detected in the images in order to train a statistical classifier to determine shadow edges. The resulting algorithms are tested on datasets containing moving shadows, moving viewpoints, and varying camera resolutions. These datasets were created by the authors and are now an important resource for the comparison of shadow-detection algorithms. The future aim of the methods investigated in (Newey et al., this issue) is to integrate the measures used and the resulting segmentations within a high-level autonomous reasoning system, which recently has been selected as one of the future computational challenges of the field (Kersten & Mamassian, 2017).

Towards the understanding of the information content in shadows within a (animal or artificial) reasoning system, the investigation of the perception-concept interface explored by linguistics and philosophy is likely to constrain some of the problems that are faced by computer scientists and psychologists. Multidisciplinary literature reviews about the investigation of shadows in philosophy, psychology and computer vision are presented in (Baxandall, 1995; Dee & Santos, 2011). In this context, Tversky (this issue) investigates the information content of shadows and dwells into the uncharted field of the *meanings* associated to the use of `shadow' as a linguistic term, and also the depiction of shadows as image metaphors.

Following a similar line, Joost Schilperoord and Lisanne van Weelden's article on "Rhetorical Shadows" (this issue) hypothesizes that there are three ways in which the depiction of cast shadows in ads and cartoons conveys meanings in situations where the commonsense shadow-caster relation is not respected. The authors suggest that resolving such an unexpected (or *incongruent*) shadow-caster relation transcends the perception of shadows, leading to the level of conceptualization. Both Tversky (this issue) and Schilperoord and van Weelden (this issue) provide an interesting contrast with Stoichita's seminal work (Stoichita, 1997) and may become a source of ideas for the emerging experimental work on the influence of shadows in showcasing products (Sharma, 2016).

Although there has been intensive progress in the investigation of the perception and cognition of shadows, there are still a number of open issues. For instance, the conditions of acceptability of incorrect shadows by the human perceptual system is still not entirely understood. Empirical results suggest that the visual system accepts incorrect (or even impossible) shadows, but the limits of incorrectness are still open for debate.

Another intriguing open question is what would be the evolutionary advantages of the selective use of the information content of shadows, because much of it is not used by our perceptual system (for instance, the number of light sources, their location and shade). The perception of shadows in movement is another issue largely neglected in the field.

The human perceptual system seems to be tuned to perceive shadows that are near the observer as darker than shadows that are further away; this seems to be linked to our attention mechanism that assigns distinct relevance measures to more important objects. However, when processing shadows, a computer vision system is incapable of making this distinction, which makes the computation of shadows in a complex scene a hard procedure. This serves as a motivation to tackle an important open issue, that involves the cross-fertilization between research in human and machine perception, namely, the investigation of models for the attention mechanism linked to shadow perception in humans that can be used as algorithms for robust computer vision systems.

Much research in shadow perception and cognition (human or machine) stumbles on one key open issue: the principles that guide the robust and consistent association of a shadow to its caster. This shadow-correspondence problem is complex for various reasons, the shadow-caster matching is naturally underconstrained and may involve non-trivial projections onto non-planar surfaces in a potential many-to-many matching procedure (Casati, 2007a; Mamassian, 2004). In addition, the definition of a satisfactory set of heuristics to guide a possible (suboptimal) shadow-object mapping procedure is still under dispute. The complex nature of this problem may be the reason why shadows are such an interesting subject for psychophysics and philosophical enquiries, but such an struggle for computer vision research.

This introduction has contextualized the articles contained in the present special issue within some of the current literature investigating the information content of cast shadows. We hope that this special issue sheds light on what is behind the cognition of shadows.

Acknowledgments

The guest editors would like to thank Schloss Dagstuhl for providing the appropriate environment for the discussions that led to the development of this special issue and to the participants of the Dagstuhl Seminar “The Message in the Shadow: Noise or Knowledge?”: Juliano Beraldo, Mehul Bhatt, Francesca Bizzarri, Marcos Cesar Danhoni Neves, Hannah

Dee, Valquiria Fenelon Pereina, Frederic Fol Leymarie, Christian Freksa, John Kennedy, John O'Dea, Ann Marie Raynal, Carl Schultz, William Sharpe, Jakob Suchan, Koichi Toyama, Barbara Tversky, and Maarten Wijntjes.

The guest editors are also indebted to the body of reviewers for their outstanding evaluation of the articles submitted for this special issue and to the coeditor-in-chief of *Spatial Cognition and Computation*, Prof. Anthony Cohn, for the strong support, from the submission of the proposal for this issue to its final edition.

Funding

This work was supported by the Fundação de Amparo à Pesquisa do Estado de São Paulo [2016/18792- 9]. Paulo E. Santos acknowledges financial support from FAPESP (Grant 2016/18792-9). Roberto Casati's research is partly funded by grants ANR-10-LABX-0087 IEC; ANR-10-IDEX-0001-02 PSL. Patrick Cavanagh received funding from the European Research Council under the European Union's Seventh Framework Program (FP7/2007-2013) ERC Grant Agreement No. AG324070 and from Dartmouth College.

ORCID

Paulo E. Santos http://orcid.org/0000-0001-8484-0354
Roberto Casati http://orcid.org/0000-0001-8215-6428
Patrick Cavanagh http://orcid.org/0000-0003-3164-2340

References

Baxandall, M. (1995). *Shadows and enlightenment*. New Haven, CT and London, UK: Yale University Press.

Bouwmans, T. (2014). Traditional and recent approaches in background modeling for foreground detection: an overview. *Computer Science Review*, *11–12*, 31–66.

Casati, R. (2006). The cognitive science of holes and cast shadows. *Trends in Cognitive Science*, *10*(2), 54–55.

Casati, R. (2007a). The copycat solution to the shadow correspondence problem. *Perception*, *37*(4), 495–503.

Casati, R. (2007b). *Shadows: Unlocking their secrets, from Plato to our time*. New York, NY: Vintage Books.

Casati, R. (2017, in press). Shadows, objects and the lexicon: On some lexicalized and non-lexicalized concepts of shadow and light. In T Crowther. & C.Mac Cumhaill (Eds.), *Perceptual ephemera*. Oxford, UK: Oxford University Press.

Cavanagh, P. (2005). The artist as neuroscientist. *Nature*, *434*, 301–307.

Cunningham, R. K., Beck, J., & Mingolla, E. (1996). Visual search for a foreground object in continuous, naturalistic displays The importance of shadows and occlusion. *Investigative Ophthalmology & Visual Science, 37*(3), S299.

Dee, H. M., & Santos, P. E. (2011). The perception and content of cast shadows: an interdisciplinary review. *Spatial Cognition & Computation*, *11*(3), 226–253. doi:10.1080/13875868.2011.565396

Ehinger, K., Allen, K., & Wolfe, J. (2016). Change blindness for cast shadows in natural scenes: even informative shadow changes are missed. *Attention, Perception & Psychophysics*, *78*(4), 978–987.

Elbakary, M. I., & Iftekharuddin, K. M. (2014, September). Shadow detection of man-made buildings in high-resolution panchromatic satellite images. *IEEE Transactions on Geoscience and Remote Sensing*, *52*(9), 5374–5386.

Elder, J., Trithart, S., Pintilie, G., & MacLean, D. (2004). Rapid processing of cast and attached shadows. *Perception*, *33*(11), 1319–1338.

Imura, T., Yamaguchi, M. K., Kanazawa, S., Shirai, N., Otsuka, Y., Tomonaga, M., & Yagi, A. (2006). Perception of motion trajectory of object from the moving cast shadow in infants. *Vision Research*, *46*(5), 652–657.

Kersten, D., Knill, D., Mamassian, P., & Bülthoff, I. (1996). Illusory motion from shadows. *Nature*, *379*, 31.

Kersten, D., & Mamassian, P. (2017). Cast shadow illusions. In A. G. Shapiro & D. Todorovic. (Eds.), *The Oxford compendium of visual illusions* (pp. 214–220). Oxford, UK: Oxford University Press.

Khuu, S. K., Gordon, J., Balcomb, K., & Kim, J. (2014). The perception of three-dimensional cast-shadow structure is dependent on visual awareness. *Journal of Vision*, *14*(3), 25. doi:10.1167/14.3.25

Khuu, S. K., Honson, V. J., & Challinor, K. L. (2016). The influence of cast shadows on the detection of three-dimensional curved contour structure. *Perception*, *45*(4), 425–442. doi:10.1177/0301006615622323 (PMID: 27107019).

Leek, E. C., Davitt, L. I., & Cristino, F. (2015). Implicit encoding of extrinsic object properties in stored representations mediating recognition: Evidence from shadow-specific repetition priming. *Vision Research*, *108*, 49–55. Retrieved from http://www.sciencedirect.com/science/article/pii/S0042698915000279

Mamassian, P. (2004). Impossible shadows and the shadow correspondence problem. *Perception*, *33*(11), 1279–1290.

Pavani, F., & Galfano, G. (2015). The multisensory body revealed through its cast shadows. *Frontiers in Psychology*, *6*, 1–11.

Prasath, V. B. S., & Haddad, O. (2014). Radar shadow detection in synthetic aperture radar images using digital elevation model and projections. *Journal of Applied Remote Sensing*, *8*(1), 083628. doi:10.1117/1.JRS.8.083628

Raynal, A. M., Bickel, D. L., & Doerry, A. W. (2014). Stationary and moving target shadow characteristics in synthetic aperture radar, *Proc. SPIE 9077, Radar Sensor Technology XVIII*, 90771B (29 May 2014).

Rensink, R. A., & Cavanagh, P. (2004). The influence of cast shadows on visual search. *Perception*, *33*(11), 1339–1358.

Russell, M., Zou, J. J., & Fang, G. (2016). An evaluation of moving shadow detection techniques. *Computational Visual Media*, *2*(3), 195–217. doi:10.1007/s41095-016-0058–0

Sanin, A., Sanderson, C., & Lovell, B. C. (2012, April). Shadow detection: a survey and comparative evaluation of recent methods. *Pattern Recognition*, *45*(4), 1684–1695. doi:10.1016/j.patcog.2011.10.001

Sasi, R. K., & Govindan, V. K. (2015). Shadow detection and removal from real images: State of art. In I. Nair (Ed.), *Proceedings of the Third International Symposium on Women in Computing and Informatics* (pp. 309–317). New York, NY: ACM. doi:10.1145/2791405.2791450

Sharma, N. (2016). Luxury implications of showcasing a product with its cast shadow. *Journal of Consumer Marketing*, *33*(7), 507–516.

Sorensen, R. (2008). *Seeing dark things: The philosophy of shadows*. Oxford, UK: Oxford University Press.

Stoichita, V. (1997). *Short history of the shadow*. Reaktion Books. Retrieved from https://books.google.com.br/books?id=1ssrLKfPY1wC

Tan, M., Mourgues, C., Aljughaiman, A., Ayoub, A., Mandelman, S. D., Zbainos, D., & Grigorenko, E. L. (in press). What the shadow knows: Assessing aspects of practical intelligence with Aurora's Toy Shadows. In H. Stoeger, A. Aljughaiman & B. Harder (Eds.), *Talent development and excellence*. Muenster, Germany: LIT.

Todes, S., & Daniels, C. (1975). Part I. beyond the doubt of a shadow. In D. Ihde & R. M. Zaner (Eds.), *Dialogues in phenomenology* (pp. 86–93). Dordrecht, The Netherlands: Springer Netherlands. doi:10.1007/978-94-010-1615-55

Yonas, A., & Granrud, C. E. (2006, January 1). Infants' perception of depth from cast shadows. *Perception & Psychophysics*, *68*(1), 154–160. doi:10.3758/BF03193665