False memories for scenes using the DRM paradigm

Filip Děchtěrenko

Charles University

Institute of Psychology, Czech Academy of Sciences

Jiří Lukavský

Charles University

Institute of Psychology, Czech Academy of Sciences

Jiří Štipl

Charles University

Author's note

Filip Děchtěrenko, Faculty of Arts, Charles University. Celetná 20, 110 00, Prague.

Institute of Psychology, Czech Academy of Sciences, Hybernská 8, 110 00, Prague, Czech

Republic. Jiří Lukavský, Faculty of Arts, Charles University. Celetná 20, 110 00, Prague.

Institute of Psychology, Czech Academy of Sciences, Hybernská 8, 110 00, Prague, Czech

Republic. Jiří Štipl, Faculty of Arts, Charles University. Celetná 20, 110 00, Prague.

Correspondence concerning this article should be addressed to Filip Děchtěrenko, Czech

Academy of Sciences, Hybernská 8, 110 00 Prague, Czech Republic. E-mail:

filip.dechterenko@gmail.com, tel: +420 774 876 252

Word count: 4642 words (excluding title, references, author affiliations, and acknowledgments)

Abstract

When we are exposed to a large number of scenes, we are remarkably good at memorizing them. The precise nature of large memory capacity is still an unresolved issue. Because visual scenes involve complex stimuli, it was difficult, until recently, to evaluate the similarity of two scenes. Researchers required a crucial manipulation check to decode this remarkable memory performance. One possible tool for helping us express how two scenes are perceptually similar are deep neural networks. In this study, we explored the extent to which we can create false visual memories using tasks inspired by the Deese-Roediger-McDermott paradigm. The similarity of scenes was evaluated by creating similarity space defined by a pre-trained deep network. The distances defined by the deep network were evaluated in a simple validation experiment. Results showed increased false alarm rates for scenes close to the averaged representation of the presented stimuli. A similar approach can be used for further studies regarding visual memory for complex scenes.

Keywords: visual scenes, deep convolutional networks, visual memory, DRM

Introduction

The world around us is full of information that we process and store. When studying visual perception, it has been observed that visual memory has massive capacity for complex stimuli (Standing, 1973; Standing, Conezio, & Haber, 1970). Recent studies have shown that we are able to store not only shallow representations of objects, but also many details (Brady, Konkle, Alvarez, & Oliva, 2008). Similar findings were observed for scenes (Konkle, Brady, Alvarez, & Oliva, 2010). Although this capacity is large, how it is organized remains an unanswered question. High recollection accuracy decreases when using yes/no tasks instead of 2AFC (Cunningham, Yassa, & Egeth, 2015). Moreover, recent results show that memory is much more limited than it would seem when it comes to details (Draschkow, Reinecke, Cunningham, & Vö, 2019).

How are the memories represented in long-term memory? In the case of objects, possible categories involve a hierarchical structure ranging from simple features to larger meaningful groups and on to whole objects (Brady et al., 2011; Ullman, 2007). The latter involves cases for scenes where the situation is more complicated: i.e. not only the objects in the scene are meant to be represented, but also their spatial arrangement and the gist of the scene (Oliva, 2005). The organization of whole scenes in memory remains an open question.

The semantic categories in our mind are formed to match the categories of the outside world. They are also meant to be economic as concerns terms for properties used to describe structure (Rosch & Loyd, 1979). Categories are formed around a prototype: a typical example that shows maximal resemblance to rest of the members in the category, while having a low resemblance to other categories (Rosch & Mervis, 1975). Prototypes are created not only for

semantic categories, but also for visual stimuli. There are two main theories around visual prototype formation. The prototype-distance model (Posner, 1970) states that a prototype is formed to minimize the distance from the presented stimuli. Meanwhile, the attribute-frequency model (Neumann, 1974) states that a prototype consists of attributes that were presented with the highest frequency. Further studies have shown higher experimental support for the prototype-distance model as concerns geometric patterns (Franks & Bransford, 1971) or faces (Sampaio, Reinke, Mathews, Swart, & Wallinger, 2018; Solso, Heck, & Mearns, 1993; Solso & Raynis, 1979). In the context of visual scenes, similar findings have not been studied. This is because it was difficult to measure what a visual prototype is.

Recent years have brought forward a possible solution to the problem of visual prototype similarity in scenes. Oliva and Torralba (2001) created a GIST descriptor that characterizes the spatial layout of a scene. It can be abstracted to the natural properties of scenes (Torralba & Oliva, 2003): something human observers can naturally use (Greene & Oliva, 2009, 2010). For this descriptor, we could compute pairwise distances, which would allow us to create a similarity space. With advances in training convolutional neural networks (CNN, LeCun, Bengio, & Hinton, 2015; Goodfellow, Bengio, & Courville, 2016), CNNs have been adopted for various task-in-vision science work such as scene categorization (Krizhevsky, Sutskever, & Hinton, 2012), object recognition (Kheradpisheh, Ghodrati, Ganjtabesh, & Masquelier, 2016) and image memorability (Khosla, Raju, Torralba, & Oliva, 2015). One particular use of CNNs is to evaluate sets of stimuli using pretrained networks, to take output from one of the inner layers, and to use it as a descriptor for the stimuli. Descriptors extracted in this way are known as *deep features*. Again, we can compute the distance between deep features using the L2 norm. Lukavský and

Děchtěrenko (2017) used this approach in a memory study that explored the properties of similarity space defined by deep features.

In this study, we explore, in two experiments, the properties of similarity space (which we denote as *image space*) defined by neural networks. First, we explored whether the distances in the image space are detectable by human observers. Do humans understand which scenes are closer to each other as defined by the image space? Second, we use image space to verify whether visual prototypes are formed for complex scenes the same as for simpler stimuli.

Experiment 1

In this experiment, we decided to validate whether the similarity space defined by convolutional neural networks can be perceived by human observers. We defined the similarity space by evaluating images using pretrained CNN (Zhou, Lapedriza, Xiao, Torralba, & Oliva, 2014) and computing pairwise L2 distances for each image. Using this structure, we repeatedly selected clusters of eight similar scenes while manipulating their distance from the remaining ninth scene. The participants' task was to detect the odd-man-out scene.

Methods

Participants Twelve subjects (mean age = 24.83, SD = 3.49, 6 males) participated in this experiment. All of them had normal or corrected-to-normal vision, and none had participated in this type of experiment before. The subjects were students of Charles University and participated as a part of their course requirements. All participants gave informed consent in advance. Sample size was based on prior power analysis done in software G*Power 3.1.92 (Faul, Erdfelder, Buchner, & Lang, 2009). We were aiming to detect large effect using one-sample t-tests while

having $\alpha = .05$ and power $1 - \beta = .8$, which did indeed occur in the above-mentioned twelve participants.

Stimuli and scene similarity Scenes were selected using a scene database (Konkle et al., 2010). From this database, we selected 64 categories: each having 32 photographs of scenes resulting in 2,048 scenes in total. Using pre-trained CNN from Zhou et al. (2014), we created a multidimensional space. Each image was evaluated by the CNN, and we took the output from the last fully inner connected layer (fc7) as the representation of the scene (a vector of 4,096 numbers). The distance metric in this structure was an L2 norm. This resulted in a 4,096 multidimensional image space. The scene categories were marked as either natural (12 categories) or man-made (52 categories).

Design For each semantic category in the database, we created experimental stimuli. To do this, we randomly sampled one scene (here denoted as *selected scene*). For this scene, we computed distances to all other scenes from the same category (31 distances). We divided the distances into 5 quintiles. As the last quintile could contain outliers with respect to typical exemplars in the category, we used only scenes from quintiles 1 to 4.

For the selected scene, we chose the 7 closest scenes (*distractors*) from quintile 1 and one scene selected from quintiles 2-4 (*target*). The selection of quintiles was randomized for each scene category (21-22 categories per quintile).

Procedure. The experiment was programmed using Psychopy1.90.2 software (Peirce et al., 2019). There were 64 trials in total. In each trial, nine scenes were shown in a rectangular grid (the size of each scene was 256×256 px). The scenes were selected as described in this paper's

Design section. The position of all scenes in the grid was randomized. The task was to detect the scene that was most different from the rest. No additional instructions, i.e. which criteria should be taken into account for oddity, were given. Participants responded by mouse selection, the time for the response was unlimited. The experiment lasted approximately 20 minutes. We created six versions of the protocols, and each protocol was supplied to two participants.

Data analysis. The data was analyzed in an R 3.6 statistical environment (R Core Team, 2018). Generalized linear mixed models were used for the analysis of differences in detection accuracy between quintiles. The division into quintiles was used as a fixed factor, while participant and scene category served as random factors. First, we used maximal models (both intercept and slopes) for the random factors, i.e. as suggested by (Barr, Levy, Scheepers, & Tily, 2013). However, as the models led to a singular fit, we reverted to other ones with random intercepts only. The significance of the fit was tested using a likelihood ratio test.

As division into quintiles loses information about distance distribution, we performed an additional analysis where we expressed the dissimilarity of the target as ratio $\frac{d_{ST}}{d_{avg}}$. d_{ST} is the distance between the selected scene and the target, and d_{avg} is the average distance of distractors from the target. Note that the sole distance between the target and the selected scene would not represent the spread of the closest neighbors. For this structure, we again used generalized linear mixed models. Moreover, we tested accuracy in detecting the odd scene from a chance level (1/9) using a sample t-test. The effect size in that case was expressed using Cohen's d .

Due to a coding error, in several conditions (23 trials out of 768), two identical distractors were used. We replicated the analysis without those trials, but the results were similar. All materials and scripts were uploaded to OSF repository (<https://osf.io/sqjbn/>).

Results

Accuracy for detecting different scenes. The accuracy of detecting different scenes increased significantly with increased distance ($\chi^2(1) = 21.7, p < .001$). As visualized in Figure 1, the accuracy for the second quintile was not significant from the chance (chance level: $1/9$, mean = 10%, SD = 9 %, $t(11) = -0.316, p = .758, d = 0.09$); however, they were significant for third quintile (mean = 21%, SD = 9 %, $t(11) = 3.9, p = .003, d = 1.12$) and fourth quintile (mean = 26%, SD = 12 %, $t(11) = 4.35, p = .001, d = 1.26$). When the task was re-parametrized using a ratio of distances, we achieved analogous results ($\chi^2(1) = 30.4, p < .001$). In Figure 2, we can see the accuracy for each category. The exploration of the image space for the top and bottom categories did not reveal any obvious pattern. In order to do that, we reduced the data using multidimensional scaling. As each pair of participants received the same set of images, we tested whether the participants made the same pattern when selecting the odd scene. We estimated agreement in selection using Cohen's κ . Average agreement was .24 (SD = .06), which is considered weak agreement (Fleiss, Levin, & Paik, 2003). Participants therefore differed in their selection of the most different scenes.

compares their differences; and Scale invariant feature transform (SIFT), (Lowe, 2004), which is also used for object detection. We selected these descriptors because they represent the scene's spatial layout (GIST), differences in color (RGB), and differences in objects (HOG, SIFT). The descriptors were computed in MATLAB 2018 using custom scripts.

Each of the descriptors was used as follows: for each of the 9 scenes, we computed pairwise L2 distances between corresponding descriptors resulting in the distance matrix. From this structure, we evaluated for each scene the distances from the rest using hierarchical clustering as suggested by Torgo (2007). This method ranked the scenes based on how distant they are from the rest. This selection was then compared with the human selection of the odd scene and with the selection of the distant scene by the image space metric.

The left part of Figure 3 shows that none of the simpler metrics could predict selection of the distant scene as defined by image space nor could they predict human performance. For the case of predicting the image space with increasing distance, prediction by GIST increased. Meanwhile, for scenes close to the distractors, metrics used for object detection delivered the best performance in explaining image space properties.

Individual metrics were not able to explain human selection of the distant scene, showing that decision strategies are complex when choosing which scene is most different.

Figure 3: Prediction of odd scene using GIST, HOG, SIFT, and RGB metrics. The left plots shows prediction of the odd scene as defined using image space structure. The right part shows prediction of human performance.

Discussion

The results showed that for distant scenes participants were able to detect them highly above chance level (Cohen’s $d_s \geq 1.12$). Moreover, the pattern showed an increasing trend. With a richer scene database, we could observe even higher accuracies. This result is surprising given the complexity of the task. For example, when discussing individual strategies for a single category (bathroom), one of the participants replied that he selected that one scene because the bathroom was more European-like, while the others were more American-like. The scale of European-like/American-like would be difficult to express using simpler decision-making criteria. Moreover, similar individual criteria could be different for each category or set of images. This was supported by our evaluation of agreement of subjects showing low agreement.

When trying to separate the decision-making criteria into simpler ones using well-known image descriptors, we observed low accuracy. This shows that participants did indeed have

various strategies for detection. Nevertheless, they were able to detect the distant scenes above chance level (with the exception of the second quintile).

In this experiment, we measured the odd scenes within semantic scene categories. We also tested three additional subjects in the experiment where the selection of distractors and target scene was not limited to one category only. Rather, it was computed for all images in the database. The task then became trivial as the distant scene was always selected from different categories. Meanwhile, the distractors were from the same category. In that case, all participants reached perfect performance.

The use of similarity space defined by CNN is widely accepted as a possible model for scene representation (Groen et al., 2018). However, the inner structure of the similarity space is difficult to understand. A promising approach for decomposing this structure is to create additional similarity spaces defined by simpler metrics and see where the structures resemble one another and where they differ. In the case of our image space, we indicated that scenes close together share similar spatial layout, while the differences for tight clusters of scenes might occur in the presence of different objects. The similarity with GIST was shown in our previous study with image space (Lukavský & Děchtěrenko, 2017)

Taken together, even with the small sample size, we showed that participants are able to estimate the distances defined by the image space. Thus, this experiment serves as a validation experiment for further use of image spaces. A similar approach could be used for other perceptual spaces: defined either by CNN e.g. (Groen et al., 2018; Lukavský & Děchtěrenko, 2017) or by multidimensional scaling (Hout, Goldinger, & Brady, 2014).

Experiment 2

In the main experiment, we focused on the creation of false memories while using the properties of image space defined by CNN. Do we form a prototypical representation from the presented stimuli and compare the presented stimuli to this prototype? If the prototype is formed, two conditions should hold. First, recognition for presented scenes closer to the prototype should be higher. Second, for the non-presented scenes close to the prototype, the false alarm rate should be higher. We focused on the latter condition in this experiment.

We tested memory performance in a paradigm inspired by the Deese-Roediger-McDermott paradigm (DRM; Deese, 1959; Roediger & McDermott, 1995). We presented a series of scenes and participants had the task of memorizing them. The similarity defined by image space served as a tool for quantifying the similarity of stimuli and describing the presented stimuli using coordinates in image space, i.e. where the center of the presented stimuli should lie. We queried our subjects with either seen or unseen scenes, while we varied whether the scenes/images were closer to the center of the presented scenes or further away. If the visual stimuli are internally represented in a structure showing resemblance to the image space, there should be more false alarms for distractors closer to the center than for distractors further away from the center. This is because the center represents the visual prototype (Posner, 1970).

Methods

Participants Twenty-seven subjects participated in the experiment (mean age = 22.96, SD = 4.65, 5 males). All of them had normal or corrected-to-normal eyesight, and none of them were colorblind. Twelve subjects participated in exchange for course credit, while the remaining 15 participated for a financial reward: 200 Czech crowns (approximately 8 euros). All participants

gave their informed consent before the experiment. Sample size was determined by prior power analysis. We needed 27 subjects for detection of a medium effect (Cohen’s $d = 0.5$) using one-sample t-test, $\alpha = .05$ and having the power of a test $1 - \beta = 0.8$. We did not perform power analysis regarding linear mixed models, as there is no scientific consensus on how to proceed in such cases. However, the sample size is typical for studies regarding visual perception.

Stimuli We selected the stimuli for the experiment from the FIGRIM database (Bylinskii et al., 2015), which contains 21 categories: each with more than 300 scenes of 700×700 px size. Before the selection of the stimuli, we inspected the photographs for artifacts (camera markers, etc.) and manually removed scenes so as to keep only quality photographs for the corresponding scene category. After the selection, 8,665 photographs remained. For those scenes, we again evaluated each of them using pre-trained CNN and created an image space from activation of the fc7 layer.

Design. The whole design is visualized in Figure 4. For each scene category, we first computed the category center. Then we divided the scenes based on their distances from the center into groups closer than median and farther than median. From the closer group, we randomly selected 15 scenes, which were used in the training block. These 15 scenes create a polygon in the image space. So, then we could compute its center representing the average target scene. From the group of closer scenes, we selected 5 more scenes closest to the center of the training scenes. We denote these scenes as *close-distractors*. From the group of farther scenes, we randomly selected 5 scenes that we denote as *far-distractors*. Finally, from the training scenes, we selected five that would be presented in the testing block. Each category was sampled three times, while unique photographs were used for each set.

Figure 4: Schema of the experimental design. The black x denotes the category center and the dotted circle the group of closer scenes. For randomly-selected targets from the closer group (green dots), we computed the center (green x) and selected close-distractors (red dots). From the farther scenes, we selected far-distractors (blue dots).

The experiment was programmed in PsychoPy. Participants were seated approximately 60 cm from the LCD display (22"). The experiment consisted of alternating parts made up of training blocks (15 training scenes) and testing blocks (15 scenes). From the 15 testing scenes, five were previously presented target scenes; five were close-distractors; and the remaining five were far-distractors. Before each training scene, a fixation cross was shown for 300 ms. Before the first scene in the experimental training block, the name of the category was shown instead. Each scene was presented for 1 s. In the testing part, participants responded using their keyboards. First, they responded using an arrow to say whether the scene was presented or not. The left arrow denoted a presented scene, while the right arrow denoted a novel scene. Then they rated their confidence on a scale of "not sure at all", "somewhat sure", and "very sure". The

selection of confidence levels was again done using arrow keys (the down arrow confirmed the selection). The whole procedure is visualized in Figure 5. Participants were instructed that the speed of answering was not important, just the accuracy. There were no breaks in the experiment, but participants were instructed to take time out when needed. The experiment lasted 50-70 minutes.

Figure 5: Experimental procedure for one training and testing block.

Data analysis. The analysis was performed in R. We were mainly focused on differences between distractor types (close-distractor and far-distractor), which we tested using generalized linear mixed models. We used participant and category type as random factors while keeping the model maximal (both as concerns intercept and slope). Similar to Experiment 1, we tested the differences against a chance level (here 50%).

We also analyzed the data using the Signal Detection Theory approach (Green & Swets, 1966; Macmillan & Creelman, 2004). Use of confidence ratings allowed us to construct empirical ROC curves with non-unit slopes for discriminating between targets and close-distractors and

targets and far-distractors. We expressed the performance using a d_a measure as an alternative to traditional d' used in the case of non-unit ROC slopes. The latter is recommended in memory studies (Loiotile & Courtney, 2015). Using this approach, the bias is not computed traditionally, but instead the perceptual distance between scene types is computed (Loiotile & Courtney, 2015; Macmillan & Creelman, 2004). Again, all materials and scripts were uploaded to OSF repository (<https://osf.io/sqjbn/>).

Results

The recognition accuracy for target scenes was 68% (SD = 12%) as visualized in Figure 6. In the case of distractors, the incorrect answer represented a false alarm. The false alarm rate for close distractors was 24% (SD = 17%), while the false alarm rate was only 11% (SD = 11%) for far-distractors.

Single sample t-tests revealed that recognition was above chance level in all three conditions ($ps < .001$, Cohen's $ds > 1.52$). As stated in the *Data Analysis* section, we used linear mixed models to test for differences in accuracy between distractor types. The difference was significant ($\chi^2(1) = 58.7, p < .001$)

Figure 6: Accuracy results and distribution for each type of scene. The vertical line shows the bootstrapped standard error of the mean; the horizontal line denotes chance level. The false alarm rate for distractors could be computed by subtracting the accuracy from 100%.

By using the confidence data taken together with old/new answers, we obtained empirical ROC curves (visualized in Figure 7) for both discrimination conditions. Red bold curve show averaged performance, while grey curves denote individual performances. Both curves had slope lesser than 1 (targets vs. close distractors: slope = 0.62, SD = 0.14; targets vs. far distractors: slope = 0.75, SD = 0.22) and the difference from unit slope was significant ($ps < .001$). Thus, we had to use index d_a for measuring sensitivity. Discriminability between target scenes and close distractors was 1.14 (SD = 0.55) and it was 1.71 (SD = 0.60) between target scenes and far distractors. The difference in sensitivity was significant (paired t-test: $t(26) = -14.58$, $p < .001$, $d = 2.81$). When visualizing the schema SDT model (Figure 8), the perceptual distance between distractors (Doshier, 1984) was small (pseudo $d_a = 0.57$).

Figure 7: ROC curves for discrimination between target and close-distractors (left) and target and far-distractors (right). The black line denotes chance level. The red line shows the average ROC curve, while the grey lines show ROC curves for individual participants.

Figure 8: SDT model for three scene types. The perceptual distance between distractors was small (pseudo $d_a = 0.57$).

The exploration of recognition accuracy per each scene category revealed that *bridge* and *bathroom* were the easiest categories, while the most difficult were *badlands* and *cockpit*. The full results are displayed in Figure 9. At the individual scene level, we explored whether memory accuracy could be explained by a photograph’s memorability (Isola, Xiao, Torralba, & Oliva, 2011). Memorability is good candidate for predicting the task, because it has been demonstrated that memorable images show good perceptual organization (Goetschalckx, Moors, Vanmarcke, & Wagemans, 2018) and are quickly categorizable (Broers, Potter, & Nieuwenstein, 2018). Therefore, we evaluated the photographs using trained CNN LaMem, which can predict memorability resembling human performance (Khosla et al., 2015). We then computed the correlation of memory accuracy with this predicted memorability. The results showed that the correlation was small and significant for targets ($r = .10, p = .003$) and for close-distractors ($r = .15, p = .007$), while insignificant for far-distractors ($r = .05, p = .335$). Therefore, the effect of memorability is only small.

Figure 9: Average accuracy (hits in the case of targets and correct rejections in the case of distractors) for individual categories.

Discussion

We explored whether we form a perceptual prototype when memorizing complex stimuli such as scenes and to what extent we can replicate DRM-style false memories for scenes close to the expected prototype. We found out that distractors closer to the average of the presented scenes show a higher false alarm rate than distractors far from the average. Similar results were observed for simpler stimuli such as geometric stimuli (Solso & Raynis, 1979) or simple faces (Solso & McCarthy, 1981): not for complex scenes.

The evidence presented here on the prototypical formation of scene averages does not offer definite proof about the use of image space as a model of memory for scenes. Recent literature, however, shows that in terms of scene classification deep features show high resemblance to functional connectivity (Groen et al., 2018). Moreover, the presented approach could be easily extended for use in comparisons between several competing models. We could select stimuli in a similar manner using different image spaces and examine which structure would show more false alarms.

Analysis of reaction times showed that scenes closer to the visual prototype required more time to respond. This is in line with findings regarding reaction time in prototype formation for numerical prototypes (Solso et al., 1993) or for scene classification using adjectives (Greene & Oliva, 2009).

One problem with image space is that representation is highly dimensional; it is very sparsely populated even for large scene databases. It limits finer experimental manipulation in order to create actual visual prototypes. One possible solution to this problem would be to use either 3D models of scenes or to take advantage of recent advances in artificial scenes created by

Generative Adversarial Networks (Goodfellow et al., 2014). The latter were recently prototyped for scene creation based on semantic description (Park, Liu, Wang, & Zhu, n.d.).

Our research complements findings from scene categorization, which show similar resemblance to scene memory. Research focused on scene categorization showed converging evidence that categorizing scenes by their function is the best model for explaining human performance. It also outperforms deep features (Greene, Baldassano, Esteva, Beck, & Fei-Fei, 2016). However, on the neural level, functional connectivity appears to be a better model: showing higher resemblance with behavioral data (Groen et al., 2018). Similar findings could be observed in the case of visual memory for the scenes.

Conclusion

In these two experiments, we explored the possible use of convolutional neural networks as a measure of perceptual or representational similarity in humans. We observed that participants were able to detect the most distant scene in the image space (Experiment 1) and when we preselected stimuli based on their distance as estimated by deep features, we observed typical effects in formation of prototypes. Taken together with our past research using deep features (Lukavský & Děchtěrenko, 2017), we believe that convolutional neural networks are a promising measure of perceptual similarity.

Acknowledgements

The work of F.D. on this research was supported by a student grant GAUK No. 604218 and by Grant No. RVO 68081740. The work of J.L. was part of a research program at the Czech

Academy of Sciences, Strategy AV21, and was supported by the Czech Science Grant Foundation No. 16-07983S.

References

- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278. doi:[10.1016/j.jml.2012.11.001](https://doi.org/10.1016/j.jml.2012.11.001)
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2009). Compression in visual working memory: using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General*, 138(4), 487–502. doi:[10.1167/8.6.199](https://doi.org/10.1167/8.6.199)
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of Vision*, 11(5), 4–4. doi:[10.1167/11.5.4](https://doi.org/10.1167/11.5.4)
- Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences*, 105(38), 14325–14329. doi:[10.1073/pnas.0803390105](https://doi.org/10.1073/pnas.0803390105)
- Broers, N., Potter, M. C., & Nieuwenstein, M. R. (2018). Enhanced recognition of memorable pictures in ultra-fast RSVP. *Psychonomic Bulletin & Review*, 25(3), 1080–1086. doi:[10.3758/s13423-017-1295-7](https://doi.org/10.3758/s13423-017-1295-7)
- Bylinskii, Z., DeGennaro, E., Rajalingham, R., Ruda, H., Zhang, J., & Tsotsos, J. (2015). Towards the quantitative evaluation of visual attention models. *Vision Research*, 116, 258–268. doi:[10.1016/j.visres.2015.04.007](https://doi.org/10.1016/j.visres.2015.04.007)

- Cunningham, C. A., Yassa, M. A., & Egeth, H. E. (2015). Massive memory revisited: Limitations on storage capacity for object details in visual long-term memory. *Learning & Memory*, 22(11), 563–566. doi:[10.1101/lm.039404.115](https://doi.org/10.1101/lm.039404.115)
- Dalal, N., & Triggs, B. (2005). Histograms of oriented gradients for human detection. In 2005 *IEEE computer society conference on computer vision and pattern recognition (cvpr'05)* (Vol. 1, pp. 886–893). IEEE. doi:[10.1109/CVPR.2005.177](https://doi.org/10.1109/CVPR.2005.177)
- Deese, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology*, 58(1), 17–22. doi:[10.1037/h0046671](https://doi.org/10.1037/h0046671)
- Doshier, B. A. (1984). Discriminating pre-experimental (semantic) from learned (episodic) associations: A speed-accuracy study. *Cognitive Psychology*, 16(4), 519–555. doi:[10.1016/0010-0285\(84\)90019-7](https://doi.org/10.1016/0010-0285(84)90019-7)
- Draschkow, D., Reinecke, S., Cunningham, C. A., & Võ, M. L. (2019). The lower bounds of massive memory: Investigating memory for object details after incidental encoding. *Quarterly Journal of Experimental Psychology*, 72(5), 1176–1182. doi:[10.1177/1747021818783722](https://doi.org/10.1177/1747021818783722)
- Faul, F., Erdfelder, E., Buchner, A., & Lang, A.-G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods*, 41(4), 1149–1160. doi:[10.3758/BRM.41.4.1149](https://doi.org/10.3758/BRM.41.4.1149)
- Fleiss, J. L., Levin, B., & Paik, M. C. (2003). *Statistical Methods for Rates and Proportions* (3rd ed., p. 800). Hoboken, NJ, USA: John Wiley & Sons, Inc. doi:[10.1002/0471445428](https://doi.org/10.1002/0471445428)

- Franks, J. J., & Bransford, J. D. (1971). Abstraction of visual patterns. *Journal of Experimental Psychology*, 90(1), 65–74. doi:[10.1037/h0031349](https://doi.org/10.1037/h0031349)
- Goetschalckx, L., Moors, P., Vanmarcke, S., & Wagemans, J. (2018). *Get the Picture? Goodness of Image Organization Contributes to Image Memorability*. doi:[10.31234/osf.io/8atj9](https://doi.org/10.31234/osf.io/8atj9)
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning* (pp. 1–800). MIT Press.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... Bengio, Y. (2014). Generative Adversarial Networks. In *Proceedings of the international conference on neural information processing systems* (pp. 2672–2680). Elsevier. Retrieved from <http://arxiv.org/abs/1406.2661>
- Green, D. M., & Swets, J. A. (1966). *Signal Detection Theory and Psychophysics* (pp. 1–455). New York: John Wiley; Sons.
- Greene, M. R., Baldassano, C., Esteva, A., Beck, D. M., & Fei-Fei, L. (2016). Visual scenes are categorized by function. *Journal of Experimental Psychology: General*, 145(1), 82–94. doi:[10.1037/xge0000129](https://doi.org/10.1037/xge0000129)
- Greene, M. R., & Oliva, A. (2009). Recognition of natural scenes from global properties: Seeing the forest without representing the trees. *Cognitive Psychology*, 58(2), 137–176. doi:[10.1016/j.cogpsych.2008.06.001](https://doi.org/10.1016/j.cogpsych.2008.06.001)
- Greene, M. R., & Oliva, A. (2010). High-level aftereffects to global scene properties. *Journal of Experimental Psychology: Human Perception and Performance*, 36(6), 1430–1442. doi:[10.1037/a0019058](https://doi.org/10.1037/a0019058)

Groen, I. I., Greene, M. R., Baldassano, C., Fei-Fei, L., Beck, D. M., & Baker, C. I. (2018).

Distinct contributions of functional and deep neural network features to representational similarity of scenes in human brain and behavior. *eLife*, 7. doi:[10.7554/eLife.32962](https://doi.org/10.7554/eLife.32962)

Hout, M. C., Goldinger, S. D., & Brady, K. J. (2014). MM-MDS: A Multidimensional Scaling Database with Similarity Ratings for 240 Object Categories from the Massive Memory Picture Database. *PLoS ONE*, 9(11), e112644. doi:[10.1371/journal.pone.0112644](https://doi.org/10.1371/journal.pone.0112644)

Isola, P., Xiao, J., Torralba, A., & Oliva, A. (2011). What makes an image memorable? In *CVPR 2011* (pp. 145–152). IEEE. doi:[10.1109/CVPR.2011.5995721](https://doi.org/10.1109/CVPR.2011.5995721)

Kheradpisheh, S. R., Ghodrati, M., Ganjtabesh, M., & Masquelier, T. (2016). Deep Networks Can Resemble Human Feed-forward Vision in Invariant Object Recognition. *Scientific Reports*, 6(1), 32672. doi:[10.1038/srep32672](https://doi.org/10.1038/srep32672)

Khosla, A., Raju, A. S., Torralba, A., & Oliva, A. (2015). Understanding and Predicting Image Memorability at a Large Scale. In *2015 IEEE international conference on computer vision (iccv)* (pp. 2390–2398). IEEE. doi:[10.1109/ICCV.2015.275](https://doi.org/10.1109/ICCV.2015.275)

Konkle, T., Brady, T. F., Alvarez, G. A., & Oliva, A. (2010). Scene Memory Is More Detailed Than You Think. *Psychological Science*, 21(11), 1551–1556. doi:[10.1177/0956797610385359](https://doi.org/10.1177/0956797610385359)

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information processing systems* (pp. 1097–1105). doi:[10.1145/3065386](https://doi.org/10.1145/3065386)

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.

doi:[10.1038/nature14539](https://doi.org/10.1038/nature14539)

Loiotile, R. E., & Courtney, S. M. (2015). A signal detection theory analysis of behavioral pattern separation paradigms. *Learning & Memory*, 22(8), 364–369.

doi:[10.1101/lm.038141.115](https://doi.org/10.1101/lm.038141.115)

Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2), 91–110. doi:[10.1023/B:VISI.0000029664.99615.94](https://doi.org/10.1023/B:VISI.0000029664.99615.94)

Lukavský, J., & Děchtěrenko, F. (2017). Visual properties and memorising scenes: Effects of image-space sparseness and uniformity. *Attention, Perception, & Psychophysics*, 79(7), 2044–2054. doi:[10.3758/s13414-017-1375-9](https://doi.org/10.3758/s13414-017-1375-9)

Macmillan NA, C. C. (2004). *Detection Theory: A Users Guide*. (Intergovernmental Panel on Climate Change, Ed.) (pp. 1–30). Cambridge: Routledge. doi:[10.4324/9781410611147](https://doi.org/10.4324/9781410611147)

Neumann, P. G. (1974). An attribute frequency model for the abstraction of prototypes. *Memory & Cognition*, 2(2), 241–248. doi:[10.3758/BF03208990](https://doi.org/10.3758/BF03208990)

Oliva, A. (2005). Gist of the Scene. In *Neurobiology of attention* (pp. 251–256). Elsevier.

doi:[10.1016/B978-012375731-9/50045-8](https://doi.org/10.1016/B978-012375731-9/50045-8)

Oliva, A., Hospital, W., & Ave, L. (2001). Modeling the Shape of the Scene : A Holistic Representation of the Spatial Envelope. *International Journal of Computer Vision*, 42(3), 145–175. doi:[10.1023/A:1011139631724](https://doi.org/10.1023/A:1011139631724)

- Oliva, A., & Torralba, A. (2001). Modeling the Shape of the Scene: A Holistic Representation of the Spatial Envelope. *International Journal of Computer Vision*, 42(3), 145–175.
doi:[10.1023/A:1011139631724](https://doi.org/10.1023/A:1011139631724)
- Park, T., Liu, M.-Y., Wang, T.-C., & Zhu, J.-Y. (n.d.). Semantic Image Synthesis with Spatially-Adaptive Normalization. In *Proceedings in Computer Vision and Pattern Recognition*.
Retrieved from <http://arxiv.org/abs/1903.07291>
- Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., ... Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, 51(1), 195–203. doi:[10.3758/s13428-018-01193-y](https://doi.org/10.3758/s13428-018-01193-y)
- Posner, M. I. (1970). Abstraction and the Process of Recognition. In *Psychology of Learning and Motivation - Advances in Research and Theory* (Vol. 3, pp. 43–100). doi:[10.1016/S0079-7421\(08\)60397-7](https://doi.org/10.1016/S0079-7421(08)60397-7)
- R Core Team. (2018). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <http://www.r-project.org/>
- Roediger, H. L., & McDermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(4), 803–814. doi:[10.1037/0278-7393.21.4.803](https://doi.org/10.1037/0278-7393.21.4.803)
- Rosch, E., & Loyd, B. B. (1979). Cognition and Categorization. *The American Journal of Psychology*, 92(3), 561. doi:[10.2307/1421578](https://doi.org/10.2307/1421578)
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7(4), 573–605. doi:[10.1016/0010-0285\(75\)90024-9](https://doi.org/10.1016/0010-0285(75)90024-9)

- Sampaio, C., Reinke, V., Mathews, J., Swart, A., & Wallinger, S. (2018). High confidence in falsely recognizing prototypical faces. *Quarterly Journal of Experimental Psychology* (2006), 71(6), 1348–1356. doi:[10.1080/17470218.2017.1329844](https://doi.org/10.1080/17470218.2017.1329844)
- Solso, R. L., Heck, M., & Mearns, C. (1993). Prototype formation in very short-term memory. *Bulletin of the Psychonomic Society*, 31(3), 185–188. doi:[10.3758/BF03337320](https://doi.org/10.3758/BF03337320)
- Solso, R. L., & McCarthy, J. E. (1981). Prototype formation of faces: A case of pseudo-memory. *British Journal of Psychology*, 72(4), 499–503. doi:[10.1111/j.2044-8295.1981.tb01779.x](https://doi.org/10.1111/j.2044-8295.1981.tb01779.x)
- Solso, R. L., & Raynis, S. A. (1979). Prototype formation from imaged, kinesthetically, and visually presented geometric figures. *Journal of Experimental Psychology: Human Perception and Performance*, 5(4), 701–712. doi:[10.1037/0096-1523.5.4.701](https://doi.org/10.1037/0096-1523.5.4.701)
- Standing, L. (1973). Learning 10,000 pictures. *Quarterly Journal of Experimental Psychology*, 25(2), 207–222. doi:[10.1080/14640747308400340](https://doi.org/10.1080/14640747308400340)
- Standing, L., Conezio, J., & Haber, R. N. (1970). Perception and memory for pictures: Single-trial learning of 2,500 visual stimuli. *Psychonomic Science*, 19(2), 73–74.
doi:[10.3758/BF03337426](https://doi.org/10.3758/BF03337426)
- Torgo, L. (2007). Resource-Bounded Fraud Detection. In *Progress in Artificial Intelligence* (pp. 449–460). Berlin, Heidelberg: Springer Berlin Heidelberg. doi:[10.1007/978-3-540-77002-2_38](https://doi.org/10.1007/978-3-540-77002-2_38)
- Torralba, A., & Oliva, A. (2003). Statistics of natural image categories. *Network: Computation in Neural Systems*, 14(3), 391–412. doi:[10.1088/0954-898X/14/3/302](https://doi.org/10.1088/0954-898X/14/3/302)

- Ullman, S. (2007). Object recognition and segmentation by a fragment-based hierarchy. *Trends in Cognitive Sciences*, 11(2), 58–64. doi:[10.1016/j.tics.2006.11.009](https://doi.org/10.1016/j.tics.2006.11.009)
- Xiao, J., Hays, J., Ehinger, K. A., Oliva, A., & Torralba, A. (2010). SUN database: Large-scale scene recognition from abbey to zoo. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (pp. 3485–3492). IEEE. doi:[10.1109/CVPR.2010.5539970](https://doi.org/10.1109/CVPR.2010.5539970)
- Zhou, B., Lapedriza, A., Xiao, J., Torralba, A., & Oliva, A. (2014). Learning Deep Features for Scene Recognition Using Places Database. In *Advances in Neural Information Processing Systems 27 (NIPS)* (pp. 487–495).