



OPEN

DATA DESCRIPTOR

Human EEG recordings for 1,854 concepts presented in rapid serial visual presentation streams

Tijl Grootswagers^{1,2}✉, Ivy Zhou², Amanda K. Robinson², Martin N. Hebart³
& Thomas A. Carlson²

The neural basis of object recognition and semantic knowledge has been extensively studied but the high dimensionality of object space makes it challenging to develop overarching theories on how the brain organises object knowledge. To help understand how the brain allows us to recognise, categorise, and represent objects and object categories, there is a growing interest in using large-scale image databases for neuroimaging experiments. In the current paper, we present THINGS-EEG, a dataset containing human electroencephalography responses from 50 subjects to 1,854 object concepts and 22,248 images in the THINGS stimulus set, a manually curated and high-quality image database that was specifically designed for studying human vision. The THINGS-EEG dataset provides neuroimaging recordings to a systematic collection of objects and concepts and can therefore support a wide array of research to understand visual object processing in the human brain.

Background & Summary

Humans are able to visually recognise and meaningfully interact with a large number of different objects, despite drastic changes in retinal projection, lighting or viewing angle, and the objects being positioned in cluttered visual environments. Object recognition and semantic knowledge, our ability to make sense of the objects around us, have been the subject of a large amount of cognitive neuroscience research¹⁻³. However, previous neuroimaging research in this field has often relied on a manual selection of a small set of images^{1,4,5}. In contrast, recent developments in computer vision have produced very large image sets for training artificial intelligence, but the individual images in these sets are minimally curated and therefore make them often unsuitable for research in psychology and neuroscience. To overcome these issues, recent work has created large, curated image sets that are designed for studying the cognitive and neural basis of human vision^{4,6}. One of these is THINGS⁴, which is an image set containing 1,854 object concepts representing a comprehensive set of nameable concepts used in the English language, accompanied by 26,107 associated manually-curated high-quality image exemplars and human behavioural annotations. This rich collection of stimuli and behavioural data has already been used to study the core representational dimensions underlying human similarity judgements⁷. The next phase is to collate corresponding neural responses to stimuli in THINGS. This would contribute to an emerging landscape of large datasets of neural responses to curated image sets that accelerate research in visual, computational, and cognitive neuroscience⁸.

Collecting neurophysiological data for the THINGS dataset, with over 26,000 images, is unachievable in a traditional neuroimaging experiment: Typically, classic object vision experiments present around one image per second^{e.g. 9-11}. Collecting one trial for each image in THINGS would take more than seven hours, which is infeasible to achieve with a single session and would require a complex design involving multiple scanning sessions. However, we have recently shown that it is possible to uncover detailed information about visual stimuli presented in rapid serial visual presentation (RSVP) streams using electroencephalography (EEG)¹²⁻¹⁵. In one of these studies¹², participants viewed over 16,000 visual object presentations at 5 and 20 images per second, in a single 40-minute EEG session. Results from multivariate pattern classification and representational similarity analysis revealed detailed temporal dynamics of object processing that were similar to work that used slower

¹The MARCS Institute for Brain, Behaviour and Development, Western Sydney University, Sydney, Australia. ²School of Psychology, The University of Sydney, Sydney, Australia. ³Vision and Computational Cognition Group, Max Planck Institute for Human Cognitive and Brain Sciences, Leipzig, Germany. ✉e-mail: t.grootswagers@westernsydney.edu.au

presentation speeds (around one image per second). Therefore, fast presentation paradigms are highly suitable for collecting neural responses to the large number of visual object stimuli in the THINGS database.

Here, we present THINGS-EEG, a dataset of human ($n = 50$) electro-encephalography responses to 22,248 images from all 1,854 concepts in the THINGS object database. In the main session, each image was repeated once, and in the second part, 200 validation images were repeated 12 times each to be able to assess data quality and compare the data to future datasets acquired with other modalities. In total, 26,248 visual images were presented in a 1-hour EEG session. This was achieved using a rapid serial visual presentation paradigm. Technical validation results indicated the dataset contains detailed neural responses to images, which shows that the dataset is a high-quality resource for future investigations into the neurobiology of visual object recognition.

In this study, we presented over 25,000 trials in a 1-hour EEG experiment. While this is an exceptionally large number in a visual object perception study, the paradigm has several limitations. Firstly, by presenting the images in rapid succession at 10 Hz, new information is being presented while previous trials are still being processed. While our previous work has shown that a great amount of detail about objects can be extracted from brain responses to RSVP streams^{12–15}, the images are being forward and backward masked, and therefore the data does not capture the full brain response to each image¹³. For example, cognitive functions such as memories or emotions may not have enough time to be instantiated at such rapid presentation rates. Our design also involved one presentation per image, which makes image-specific analyses challenging, placing the focus of this work at the level of the 1,854 object concepts. The benefit of this is that the data has a built-in control for image-level confounds. For example, visual regularities that are specific to an image but vary across a concept will not generalise to the concept level. Visual statistics that reliably differ at the concept level of course may need to be accounted for, depending on the goals of the experimenter. For example, recent work has pointed out concept-level differences in mean luminance between images in the THINGS image set¹⁶, which can be controlled for in future analyses of the THINGS-EEG dataset. Another point to consider is that our recording setup did not include EOG or EMG channels, which means the dataset does not contain external recordings of eye or other muscle movements. These movement patterns are unlikely to contain informative stimulus-specific information, due to the fast presentation paradigm. However, they still cause noise artefacts in the EEG data. Future users could consider detecting and correcting for eye movements using the frontal EEG channels.

The THINGS-EEG dataset has strong potential for investigating the neurobiology of visual object recognition and semantic knowledge. The dataset presents a very large set of non-invasive neural responses to visual stimuli in human participants. We foresee many possible uses of this dataset. For example, the dataset could be used to test models of visual object representation, such as different semantic models, or deep neural networks. It could be used to test the generalisability of previous studies that were limited by small stimulus sets⁵. The rapid presentation paradigm allows to examine sequential effects in the data, such as how a specific object concept influences the encoding of the subsequent presentations. The consistent presentation frequency also lends itself to separate the data into oscillatory components. Instead of the RSA and classification analyses presented here, it is also possible to analyse the data in an encoding framework, for example by creating an encoding model from the semantic information in the THINGS-dataset. In sum, as THINGS is a high-quality stimulus set of record size, THINGS-EEG accompanies this resource with a comprehensive set of human neuroimaging recordings.

Methods

50 individuals volunteered to take part in the experiment in return for course credit. Participants were recruited from the undergraduate student population at the University of Sydney. They were 36 females and 14 males, mean age 20.44 (sd 2.72), age range 17 – 30. Participants had different language profiles, with 26 native English speakers, 24 non-native speakers, 24 monolinguals, and 25 bilinguals. Handedness was not recorded. All participants reported normal or corrected-to-normal vision and reported no neurological or psychiatric disorders. There are 4 participants marked for potential exclusion due to notably poor signal quality or equipment failure (marked in the *participants.tsv* file). These participants are included in the release for completeness. The study was approved by the University of Sydney ethics committee. Informed consent was obtained from all participants at the start of the experiment.

Stimuli were obtained from the THINGS database⁴ (Fig. 1A). For detailed information on the contents and organisation of this stimulus database, readers are referred to the accompanying publication⁴. THINGS contains 1,854 objects concepts, with 12 or more images per concept. The first 12 images for each concept were used for this experiment, resulting in 22,248 different visual images, which we divided into 72 sequences of 309 stimuli. Individual concepts were never repeated within one sequence. To be able to assess within- and between-subject variance on the same images, we presented another 200 validation images from the THINGS database 12 times to every subject after the main part of the experiment. For every subject we used the same 200 validation images (listed in the *test_images.csv* file). We repeated these images in 12 sequences of 200, where we presented them in random order. The experiment thus contained 84 sequences in total. The subjects were not explicitly made aware of the two different parts.

The experiment was programmed in Python (v3.7), using the Psychopy¹⁷ library (version 3.0.5). The sequences were presented at 10 Hz, with a 50% duty cycle (Fig. 1C). That is, each image was presented for 50 ms, followed by a 50 ms blank screen. Participants were seated about 57 cm from the screen, and the stimuli subtended approximately 10 degrees visual angle. Overlaid at the centre of each image was a bullseye (0.5 degrees visual angle) to help participants maintain fixation. To increase attention and engagement, each sequence contained 2 to 5 random target events, where the bullseye turned red for 100 ms, and the participants were instructed to press a button on a button box using their right hand. These target events are marked in the dataset, as researchers may want to consider excluding these target events, depending on the aim of their analysis. At the end of each sequence, the display showed the progress through the experiment, and participants were able to start the next sequence with a button press. Participants were asked to sit still and minimise eye movements



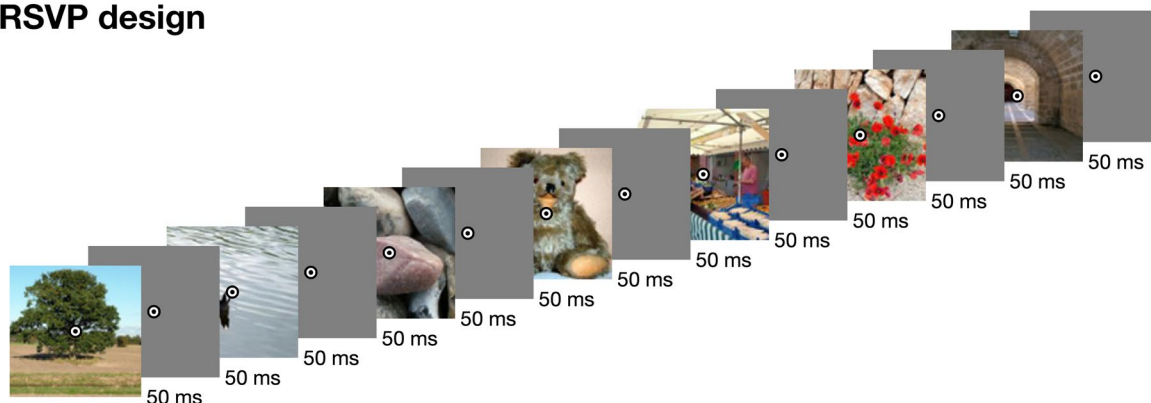
a Example images**b EEG setup****c RSVP design**

Fig. 1 Example stimuli, design, and EEG setup image. (A) Example images similar to the stimuli used in the experiment. (B) EEG experimental setup (photo credit AKR). (C) Rapid serial visual presentation design. For illustration purposes, only part of the sequence is shown. For this figure, all images were replaced by public domain images with similar appearance (obtained from PublicDomainPictures.net: Brunhilde Reing).
.....

during the sequences and to use the time between sequences as breaks and relax, and start the next sequence using a button press when they were ready. The experiment lasted around one hour.

We used a BrainVision ActiChamp system to record continuous data while participants viewed the sequences (Fig. 1B). Conductive gel was used to reduce impedance at each electrode site below 10 kOhm where possible. The median electrode impedance was under 18 kOhm in 40/50 participants and under 60 kOhm in all participants. We used 64 electrodes, arranged according to the international standard 10–10 system for electrode placement^{18,19}. The signal was digitised at a 1000-Hz sample rate with a resolution of 0.0488281 μ V. Electrodes were referenced online to Cz. An event trigger was sent over the parallel port at the start of each sequence (trigger code E3), and at every stimulus onset event (trigger code E1) and stimulus offset event (trigger code E2).

To perform basic quality checks and technical validation, for each subject, we ran a standard decoding analysis. We decoded pairwise images for the 200 validation images, and we created the full time-varying 1,854 \times 1,854 Representational Dissimilarity Matrix^{20,21} reflecting the pairwise decoding accuracies between all 1,854 object concepts. We used a minimal preprocessing pipeline derived from our previous RSVP-MVPA studies^{12–15}. Using Matlab (R2020b) and the EEGLab (v14.0.0b) toolbox²², data were filtered using a Hamming




Fig. 2 Results for the 200 validation images that were repeated 12 times at the end of the session. **(A)** mean pairwise classification accuracy over time. **(B)** Mean pairwise decoding accuracy over time, per subject, sorted by peak classification accuracy. Subjects 1 and 6 are not shown as they did not have data on the validation images. **(C)** Noise ceiling over time shows the expected correlation of the ‘true’ model with the RDMs of the validation images and reflects the between-subject variance in the RDMs.

windowed FIR filter with 0.1 Hz highpass and 100 Hz lowpass filters, re-referenced to the average reference, and downsampled to 250 Hz. Epochs were created for each individual stimulus presentation ranging from [−100 to 1000 ms] relative to stimulus onset. No further preprocessing steps were applied for the technical validation analysis presented here (as in our previous work using similar presentation paradigms^{12–15}). Researchers may want to consider popular preprocessing steps such as baseline correction or eye movement correction. The channel voltages at each time point served as input to the decoding analysis.

Decoding analyses were performed in Matlab using the CoSMoMvPA toolbox²³. We first decoded between the 200 validation images. For a given pairs of images, we used a leave-one-sequence out (total: 12 sequences) cross-validation procedure and trained a regularised ($\lambda = 0.01$) linear discriminant classifier to distinguish between the images. The mean classification accuracy on the image pair in the left-out sequences were stored in a $200 \times 200 \times 275 \times 50$ (image \times image \times time point \times subject) RDM, which is symmetrical across its first diagonal. A similar procedure was performed for the main experiment, using the 1,854 different image concepts. This resulted in an $1,854 \times 1,854 \times 275 \times 50$ (concept \times concept \times time point \times subject) RDM. For the 200 validation images, we also computed noise ceilings by comparing between subject RDMs, as described in previous work²⁴. The noise ceilings estimate the lower and upper bound of the highest achievable performance of a model that attempts to explain variance in the data.

Data Records

All data and code are publicly available. The raw EEG recordings are hosted in BIDS^{25,26} standard on OpenNeuro (<https://doi.org/10.18112/openneuro.ds003825.v1.1.0>)²⁷. The preprocessed (Matlab/EEGLAB²² format) data, and group-average RDMs are included for convenience (in the data/derivatives directory). The RDMs for individual subjects are hosted in Matlab format in a separate repository on Figshare (<https://doi.org/10.6084/m9.figshare.14721282>)²⁸. All custom code is available from the Open Science Framework (<https://doi.org/10.17605/OSF.IO/HD6ZK>)²⁹, which also contains links to the above repositories.

Technical Validation

We computed representational dissimilarity matrices for the 200 validation images, by calculating time-varying decoding accuracy between all pairs of images. Mean, subject-wise decoding accuracy (Fig. 2A) showed an initial peak around 100 ms, and a second, lower peak around 200 ms after stimulus onset. The shape of the time-varying decoding was similar to previous object decoding studies^{6,8–9,10}, and was also similar to previous results on images presented in fast succession^{12–14}, indicating data quality was similar to these studies. For most




Fig. 3 Results for the 1,854 image concepts that were repeated 12 times (using a different image each time). **(A)** Full $1,854 \times 1,854$ RDM at 200 ms, arranged by high-level category. **(B)** Zoomed in section of the full RDM. **(C)** Mean pairwise classification accuracy between concepts over time. **(D)** Correlation over time between the neural RDM and four high-level categorical models.

subjects, this shape was apparent from their individual data (Fig. 2B). The noise ceiling (Fig. 2C) indicates an average similarity (correlation) of up to 0.2 between the subject-specific dissimilarity matrices.

Next, we computed the full RDM for all 1,854 images concepts (1,717,731 pairs). The average accuracy within this RDM (Fig. 3A) was lower than for the validation images, which is likely due to the fact that accuracy reflects generalised concept-similarity across images. Figure 3B shows the full RDM at one time point (200 ms). To test if the values in the RDM contain meaningful information, we computed the correlation between the full RDM and four example categorical models (Fig. 3C). The models coded for the presence of a certain category (e.g., animal). Figure 3C shows each model reaches an above-zero correlation, with the 'natural' model reaching the highest correlation, around 200 ms.

Code availability

Code and detailed instructions to reproduce the technical validation analyses and figures presented in this manuscript are available from the Open Science Framework (<https://doi.org/10.17605/OSF.IO/HD6ZK>)²⁹, which also contains links to the data repositories.

Received: 26 July 2021; Accepted: 3 November 2021;

Published online: 10 January 2022

References

1. Wardle, S. G. & Baker, C. Recent advances in understanding object recognition in the human brain: deep neural networks, temporal dynamics, and context. *FL1000Research* **9**, 590 (2020).
2. DiCarlo, J. J. & Cox, D. D. Untangling invariant object recognition. *Trends Cogn. Sci.* **11**, 333–341 (2007).
3. Gauthier, I. & Tarr, M. J. Visual Object Recognition: Do We (Finally) Know More Now Than We Did? *Ann. Rev. Vis. Sci.* **2**, 377–396 (2016).
4. Hebart, M. N. *et al.* THINGS: A database of 1,854 object concepts and more than 26,000 naturalistic object images. *PLOS ONE* **14**, e0223792 (2019).
5. Grootswagers, T. & Robinson, A. K. Overfitting the Literature to One Set of Stimuli and Data. *Front. Hum. Neurosci.* **15**, (2021).
6. Mehrer, J., Spoerer, C. J., Jones, E. C., Kriegeskorte, N. & Kietzmann, T. C. An ecologically motivated image dataset for deep learning yields better models of human vision. *Proc. Natl. Acad. Sci.* **118**, (2021).
7. Hebart, M. N., Zheng, C. Y., Pereira, F. & Baker, C. I. Revealing the multidimensional mental representations of natural objects underlying human similarity judgements. *Nat. Hum. Behav.* **4**, 1173–1185 (2020).
8. Naselaris, T., Allen, E. & Kay, K. Extensive sampling for complete models of individual brains. *Curr. Opin. Behav. Sci.* **40**, 45–51 (2021).
9. Carlson, T. A., Tovar, D. A., Alink, A. & Kriegeskorte, N. Representational dynamics of object vision: The first 1000 ms. *J. Vis.* **13**, 1 (2013).
10. Cichy, R. M., Pantazis, D. & Oliva, A. Resolving human object recognition in space and time. *Nat. Neurosci.* **17**, 455–462 (2014).
11. Kaneshiro, B., Guimaraes, M. P., Kim, H.-S., Norcia, A. M. & Suppes, P. A Representational Similarity Analysis of the Dynamics of Object Processing Using Single-Trial EEG Classification. *PLOS ONE* **10**, e0135697 (2015).
12. Grootswagers, T., Robinson, A. K. & Carlson, T. A. The representational dynamics of visual objects in rapid serial visual processing streams. *NeuroImage* **188**, 668–679 (2019).
13. Robinson, A. K., Grootswagers, T. & Carlson, T. A. The influence of image masking on object representations during rapid serial visual presentation. *NeuroImage* **197**, 224–231 (2019).
14. Grootswagers, T., Robinson, A. K., Shatek, S. M. & Carlson, T. A. Untangling featural and conceptual object representations. *NeuroImage* **202**, 116083 (2019).
15. Grootswagers, T., Robinson, A. K., Shatek, S. M. & Carlson, T. A. The neural dynamics underlying prioritisation of task-relevant information. *Neurons Behav. Data Anal. Theory* **5**, 1–17 (2021).
16. Harrison, W. J. Luminance and contrast of images in the THINGS database. Preprint at <https://doi.org/10.1101/2021.07.08.451706> (2021).
17. Peirce, J. *et al.* PsychoPy2: Experiments in behavior made easy. *Behav. Res. Methods* **51**, 195–203 (2019).
18. Jasper, H. H. The ten-twenty electrode system of the International Federation. *Electroencephalogr Clin Neurophysiol* **10**, 371–375 (1958).
19. Oostenveld, R. & Praamstra, P. The five percent electrode system for high-resolution EEG and ERP measurements. *Clin. Neurophysiol.* **112**, 713–719 (2001).
20. Grootswagers, T., Wardle, S. G. & Carlson, T. A. Decoding Dynamic Brain Patterns from Evoked Responses: A Tutorial on Multivariate Pattern Analysis Applied to Time Series Neuroimaging Data. *J. Cogn. Neurosci.* **29**, 677–697 (2017).
21. Kriegeskorte, N., Mur, M. & Bandettini, P. A. Representational Similarity Analysis - Connecting the Branches of Systems Neuroscience. *Front. Syst. Neurosci.* **2**, 4 (2008).
22. Delorme, A. & Makeig, S. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J. Neurosci. Methods* **134**, 9–21 (2004).
23. Oosterhof, N. N., Connolly, A. C. & Haxby, J. V. CoSMoMPPA: Multi-Modal Multivariate Pattern Analysis of Neuroimaging Data in Matlab/GNU Octave. *Front. Neuroinformatics* **10**, (2016).
24. Nili, H. *et al.* A Toolbox for Representational Similarity Analysis. *PLoS Comput Biol* **10**, e1003553 (2014).
25. Gorgolewski, K. J. *et al.* The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. *Sci. Data* **3**, 160044 (2016).
26. Pernet, C. R. *et al.* EEG-BIDS, an extension to the brain imaging data structure for electroencephalography. *Sci. Data* **6**, 103 (2019).
27. Grootswagers, T., Zhou, I., Robinson, A., Hebart, M. N. & Carlson, T. A. Human electroencephalography recordings from 50 subjects for 22,248 images from 1,854 object concepts. *OpenNeuro* <https://doi.org/10.18112/openneuro.ds003825.v1.1.0> (2021).
28. Grootswagers, T., Zhou, I., Robinson, A., Hebart, M. N. & Carlson, T. A. THINGS-EEG: Human electroencephalography recordings for 22,248 images from 1,854 object concepts. *figshare* <https://doi.org/10.6084/m9.figshare.14721282> (2021).
29. Grootswagers, T., Zhou, I., Robinson, A., Hebart, M. N. & Carlson, T. A. THINGS-EEG: Human electroencephalography recordings from 50 subjects for 22,248 images from 1,854 object concepts. *Open Science Framework* <https://doi.org/10.17605/OSF.IO/HD6ZK> (2021).

Acknowledgements

This research was supported by ARC DP160101300 (TAC), ARC DP200101787 (TAC), and ARC. DE200101159 (AKR). The authors acknowledge the University of Sydney HPC service for providing High Performance Computing resources.

Author contributions

TG: Conceptualization, Methodology, Investigation, Formal analysis, Visualization, Data Curation, Writing – Original Draft, Writing – Review & Editing, Supervision, Project administration. IZ: Conceptualization, Investigation, Data Curation, Writing – Review & Editing, Project administration. AKR: Conceptualization, Methodology, Writing – Review & Editing. MNH: Conceptualization, Methodology, Writing – Review & Editing. TAC: Conceptualization, Methodology, Writing – Review & Editing, Supervision, Funding acquisition, Project administration.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to T.G.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

The Creative Commons Public Domain Dedication waiver <http://creativecommons.org/publicdomain/zero/1.0/> applies to the metadata files associated with this article.

© The Author(s) 2022