Behavioural markers of recurrent processing in the brain

Tim Maniquet (tim.maniquet@kuleuven.be)

Hans Op de Beeck (hans.opdebeeck@kuleuven.be
Brain & Cognition Research Unit
KULeuven, Leuven (Belgium)

Abstract

The human brain processes visual input quickly and reliably, even in the face of large perturbations in visual input. How cortical computations solve object identity under complex visual changes is not well explained by current models. Recent work using deep neural networks (DNNs) points to recurrent processing, with models shown to be more resilient to visual complexity when equipped with recurrent connections. Here, we explore recurrent processing, and ask whether we can pinpoint the substrates of these computations in the brain. We design a stimulus set including visual manipulations known to trigger recurrent processing, from which we establish a behavioural benchmark of performance markers of recurrent processing. We then record fMRI on participants tasked with classifying the images from this stimulus set, and compare our benchmark markers with cortical representations. Preliminary results using this approach indicate that regions in the prefrontal cortex as potential candidates for hubs of recurrent connectivity in the context of challenging object recognition.

Keywords: recurrent processing; visual complexity; fMRI

Introduction

Object recognition is a fast, reliable process (DiCarlo, Zoccolan, & Rust, 2012). While DNN models of object recognition mimic human behaviour in some tasks, they are also limited in explaining it in the presence of complexity. The computations that solve object identity under perturbation are hence not clear. Recent research has studied this by making object recognition difficult through visual transformations such as occlusion or clutter. While models typically fail on such tasks, humans are still able to recognise objects under vast amounts of visual transformations. In these circumstances, however, recognition takes more time behaviourally (longer RTs), and accurate decoding emerges later on in the brain (Rajaei, Mohsenzadeh, Ebrahimpour, & Khaligh-Razavi, 2019).

This delay when images are noisy, occluded, or otherwise degraded, points to extra-feedforward computations needed to solve object identity. Evidence suggests that the mechanisms responsible for denoising, discarding distractors, filling in missing parts etc. are supported by recurrent processing (Kar, Kubilius, Schmidt, Issa, & DiCarlo, 2019). This is reinforced by the observation that DNNs are better at complex recognition when equipped with recurrence (Kietzmann et al., 2019). However, the specific computations implemented through recurrent processing are not well known yet.

Here, we and ask whether we can pinpoint the neural substrates of recurrent processing when it solves different kinds of visual challenges. We design a stimulus set to trigger recurrence, and use behavioural markers of the need for recurrent processing across its conditions to investigate where these computations take place in the brain.

Methods

Stimulus set

We built a stimulus set comprising 10 exemplar images from 8 different semantic categories covering a large range of animacy and real world sizes (person, bird, cat, banana, tree, fire hydrant, bus, building). We then created 4 challenging versions of each image, by applying manipulations known in the field to trigger recurrent processing (total stimulus set: 400 images from 40 conditions - 8 categories * 5 manipulations including a control condition, see Fig. 1A). We applied **deletion** by removing parts of the object, **occlusion** by adding grey disks on top of objects (Tang et al., 2018), **clutter** by adding cluttered backgrounds behind objects (Seijdel et al., 2021) and **low-pass phase scrambling** by replacing the phase of high spatial frequencies with noise (Goddard, Carlson, Dermody, & Woolgar, 2016).

To confirm the effects of our manipulations on recurrent processing, we collected data online from 85 participants performing a categorisation task on the 8 categories of the stimulus set. Images were presented for 50ms, followed by a mask on half of the trials. We established a benchmark of recurrent processing effects by extracting backward masking effect on accuracy and average reaction times (RT) per condition.

fMRI experiment

We recorded 3T fMRI data from a preliminary set of 11 participants (25.5 ± 4.8 years old, 5 females). The design of the task during the recording followed that of the behavioural experiment, with presentation times of 50ms, only without masking and with response times limited to 2.95s. Each image was presented twice during the session.

Results

Visual manipulations trigger recurrent processing

Results from our behavioural benchmark show a significantly larger backward masking effect on accuracy (all $p < 0.05$), and significantly longer RTs (all $p < 0.001$) for each of the

4 challenging manipulations as compared to the control condition. Additionally, we found a significant correlation between average RT and masking effect across our manipulations (Pearson's correlation, $r = 0.9787$, $p = 0.0037$, see Fig. 1B), showing that our indicators of recurrent processing converge.

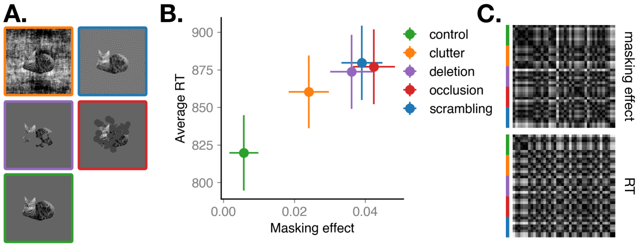

Figure 1: **A.** Example image across 4 manipulations and without manipulation (control). **B.** Scatter plot of average RT per average masking effect across manipulations. **C.** RDMs of absolute RT and masking effect differences per all 40 conditions, organised by the 5 manipulation conditions.

From these indicators, we built matrices using the absolute difference in values across all pairs of our 40 conditions (see Fig. 1C). The resulting representational dissimilarity matrices (RDMs) were used as indices of recurrent processing for our stimulus set, and compared with brain RDMs in the following analyses.

Discarding of manipulations in brain representations

To explore the content of representations across the brain, we used a support vector machine to run pairwise decoding across our 40 conditions. We looked at results from a set of ROIs defined through 44 parcels (Glasser et al., 2016). We then built RDMs from the pairwise decoding accuracies (see Fig. 2A) and used them to perform representational similarity analysis (RSA) with several target models (see Fig. 2B & C).

We found limited correlations between ROIs and the manipulation model, with early visual cortex showing the highest correlation (mean 0.07 ± 0.01 SEM). This is surprising considering the importance of manipulations in the visual aspects of our stimulus set (see GIST model (Oliva & Torralba, 2001), Fig. 2C). Conversely, we found evidence for category representation in visual areas, especially in ventral visual stream which showed correlations to the category model (mean 0.11 ± 0.01 SEM) and to the animacy model (mean 0.17 ± 0.02 SEM).

A potential explanation is that the brain can successfully discard the noise that our manipulations represent during the categorisation task. Supporting this view is the correlation of superior parietal cortex with the category model (mean 0.09 ± 0.01 SEM) but less so with the animacy model (mean 0.03 ± 0.01 SEM), which could indicate a task-driven discarding of visual manipulations, in favour of categories.

Correlates of recurrent effects in the brain

We next investigated whether we could find correlates of recurrent processing in the brain, as benchmarked for our stimulus set. To do so we used the RT and masking effect RDMs in our RSA, and looked for ROIs that showed a correlation with these markers of recurrent processing.

We did not find correlations with the RT model, except for somatosensory and motor cortex (mean 0.05 ± 0.02 SEM). The latter also shows correlation with the category model (mean 0.23 ± 0.01 SD), which seems to indicate that any correlation with the RT model should be response-driven.

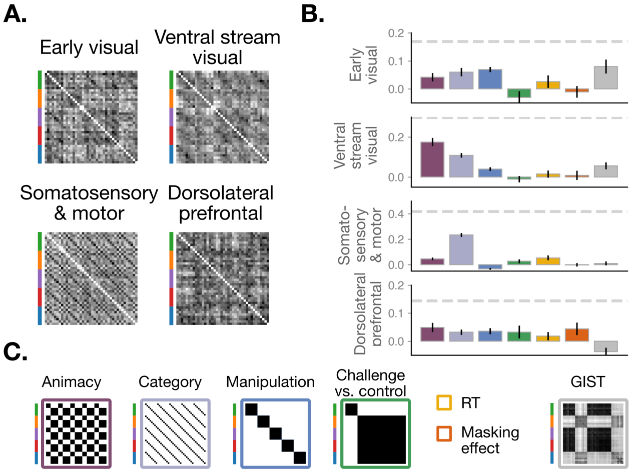

Figure 2: **A.** Pairwise decoding RDMs. **B.** RSA results for the ROIs shown in *A.*. Dotted gray lines indicate split-half noise ceilings for each ROI over the 11 participants. **C.** Theoretical models used in RSA (for RT and masking effect, see Fig. 1).

We found the strongest correlations with the masking model for dorsolateral prefrontal cortex (mean 0.04 ± 0.01 SEM, see Fig. 2) and anterior cingulate and medial prefrontal cortex (mean 0.04 ± 0.01 SEM). Both ROIs did not appear to be visually driven (mean correlation with GIST: -0.04 ± 0.01 and -0.02 ± 0.01 for the latter), which could point them as potential recurrence-driven regions for our task.

Conclusion

Our preliminary results demonstrate the use of our behavioural benchmark for studying recurrent processing in the brain. Using a stimulus set with clearly defined category and manipulation dimensions, we were able to find representations for the former more than for the latter, which points to the successful discarding of task-irrelevant visual complexity during visual categorisation. Moreover, we could find some correlations with one of our recurrent processing indicators, masking effect, in two regions in the prefrontal cortex. This shows the potential of this approach in finding the potential sources of recurrent connectivity during challenging object recognition.

References

DiCarlo, J. J., Zoccolan, D., & Rust, N. C. (2012, February). How Does the Brain Solve Visual Object Recognition? *Neuron*, *73*(3), 415–434. doi: 10.1016/j.neuron.2012.01.010

Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., … Van Essen, D. C. (2016, August). A multi-modal parcellation of human cerebral cortex. *Nature*, *536*(7615), 171–178. doi: 10.1038/nature18933

Goddard, E., Carlson, T. A., Dermody, N., & Woolgar, A. (2016, March). Representational dynamics of object recognition: Feedforward and feedback information flows. *NeuroImage*, *128*, 385–397. doi: 10.1016/j.neuroimage.2016.01.006

Kar, K., Kubilius, J., Schmidt, K., Issa, E. B., & DiCarlo, J. J. (2019, June). Evidence that recurrent circuits are critical to the ventral stream's execution of core object recognition behavior. *Nature Neuroscience*, *22*(6), 974–983. doi: 10.1038/s41593-019-0392-5

Kietzmann, T. C., Spoerer, C. J., Sörensen, L. K. A., Cichy, R. M., Hauk, O., & Kriegeskorte, N. (2019, October). Recurrence is required to capture the representational dynamics of the human visual system. *Proceedings of the National Academy of Sciences*, *116*(43), 21854–21863. doi: 10.1073/pnas.1905544116

Oliva, A., & Torralba, A. (2001, May). 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

Rajaei, K., Mohsenzadeh, Y., Ebrahimpour, R., & Khaligh-Razavi, S.-M. (2019, May). Beyond core object recognition: Recurrent processes account for object recognition under occlusion. *PLOS Computational Biology*, *15*(5), e1007001. doi: 10.1371/journal.pcbi.1007001

Seijdel, N., Loke, J., van de Klundert, R., van der Meer, M., Quispel, E., van Gaal, S., … Scholte, H. S. (2021, July). On the Necessity of Recurrent Processing during Object Recognition: It Depends on the Need for Scene Segmentation. *Journal of Neuroscience*, *41*(29), 6281–6289. doi: 10.1523/JNEUROSCI.2851-20.2021

Tang, H., Schrimpf, M., Lotter, W., Moerman, C., Paredes, A., Ortega Caro, J., … Kreiman, G. (2018, August). Recurrent computations for visual pattern completion. *Proceedings of the National Academy of Sciences*, *115*(35), 8835–8840. doi: 10.1073/pnas.1719397115