Modeling the Human Brain: Investigating Stimulus-Response Transformations in Complex, Time-Continuous Environments Using Deep Neural Networks

Sabine Haberland (sabine.haberland@tu-dresden.de)
Department of Psychology, TUD Dresden University of Technology
Zellescher Weg 17, 01069 Dresden, Germany

Hannes Ruge (hannes.ruge@tu-dresden.de)
Department of Psychology, TUD Dresden University of Technology
Zellescher Weg 17, 01069 Dresden, Germany

Holger Frimmel (holger.frimmel@tu-dresden.de)
Department of Psychology, TUD Dresden University of Technology
Zellescher Weg 17, 01069 Dresden, Germany

Abstract

Studying how humans transform complex, high-dimensional stimuli into appropriate behavior within time-continuous environments has been challenging due to the prevalent use of trial-based study designs. Deep Q-Networks (DQNs) have emerged as valuable tools for modeling stimulus-response (S-R) transformations in such environments. Here, we showed that a DQN-based encoding model approach can be used to predict neural activations and human behavior using features generated by a DQN. Therefore, we collected motor responses and fMRI data from human subjects (N=23) while playing arcade games. We hypothesized that advancements in machine learning can be leveraged to improve prediction accuracy. We compared the prediction accuracy of features generated by two recently developed DQNs and a third baseline DQN, each differing in network architecture and training procedures. We present preliminary evidence that all three DQNs predict behavior and fMRI activations significantly above chance at a fine-grained temporal scale. Features generated by the most advanced model achieved the best results. We found a hierarchical correspondence between the layers of DQNs and stages of human visuo-motor processing. These findings suggest that improved DQNs serve as suitable tools for modeling S-R transformations in a time-continuous manner.

Keywords: Encoding Model; Arcade Games; Deep Neural Network; Neuroimaging

Introduction

The human brain filters complex environmental information and transforms them into appropriate motor responses. While the conventional experimental approach, characterized by its trial structure and use of low-dimensional stimuli, has offered important insights into the brain's functionality, investigating how the brain processes high-dimensional stimuli in a complex and time-continuous environment has remained challenging. Neural networks (NNs) coupled with an encoding model have emerged as a promising approach to address this issue. NNs generate features for the encoding model, enabling, for instance, the characterization of voxel responses underlying the visual processing of stimuli. Moreover, this approach has revealed a gradient between the layers of an NN and visual processing stages in the human brain (Cichy, Khosla, Pantazis, Torralba, & Oliva, 2016; Eickenberg, Gramfort, Varoquaux, & Thirion, 2017). Deep Q-networks (DQNs), a combination of reinforcement learning (RL) and deep learning, have extended the approach to more complex and time-continuous stimuli. It has even been demonstrated that human behavior (Mohr, Cichy, & Ruge, 2019) and neural activity during gameplay of arcade games can be modeled (Cross, Cockburn, Yue, & O'Doherty, 2021). DQNs were trained to estimate, for each state and action, the expected, discounted future reward, the Q-value, without any human data. The ongoing rapid improvements in the field of RL have resulted in DQNs capable of solving complex tasks at human-level and beyond (Espeholt, Marinier, Stanczyk, Wang, & Michalski, 2019). Based on the previous results, we aimed to address the question of whether these advancements can be translated into better modeling of neural activations underlying stimulus-response (S-R) transformations. We collected fMRI data and motor responses of subjects playing arcade games. To assess progress in RL, we considered two recently developed DQNs, Ape-X (Horgan et al., 2018) and SEED (Espeholt et al., 2019), and a baseline DQN (Mnih et al., 2015) as feature-generating mappings within an encoding model. We compared the prediction accuracy of the three DQN-based encoding models in predicting human brain activity and behavior. We hypothesized that SEED, as the most complex model, exhibits the highest accuracy in modeling behavior and neural activations within task-related brain regions. We present initial evidence that advanced DQN-based encoding models have the potential to model S-R transformations in visuo-motor tasks.

Methods

Task

Data Acquisition Subjects ($N = 23$, age: 24.3 years, age range: 19-35 years) were tasked with playing three Atari games: Breakout, Space Invaders, and Enduro. The task was sampled at 45 Hz. For each game, subjects practiced the game. After training, they conducted the main experiment consisting of 5 sessions, each session lasting 7 minutes, inside an MRI scanner. Motor responses, video screens, and fMRI data were recorded. A four-button response pad was used as an input device.

Principles of the Games In Breakout, the player must control a paddle to smash a wall of bricks with a ball. In Space Invaders, the player must fight aliens with a spaceship before they reach Earth. In Enduro, the player aims to overtake as many cars as possible. Response options included 'no action', 'left', 'right', 'hit brake', 'fire', and all combinations.

Analysis

Encoding Model To predict human behavior and voxel activations we used a DQN-based encoding model approach (Figure 1). First, a DQN was used as a feature-generating mapping. We processed each frame of the human-generated videos by the DQN, resulting in a time series of activation values for each neuron of the DQN. The activations of the top layer neurons can be interpreted as Q-values associated with each possible action for a given input screen. For each DQN, game, and subject we separately employed a general linear model (GLM) to fit the generated stimulus features to human data. To model human behavior, we used the time series of Q-values as predictors and the motor responses of subjects as the dependent variable. To explore an initial validation of a hierarchical relationship between DQN layers and visuo-motor processing steps, we estimated voxel-specific GLMs with layer-specific predictors. For each voxel, we compared the prediction accuracy of two GLMs. In the first GLM, which required regularization, we tested the activations of neurons from the first layer of a DQN as predictors, while in the other GLM, we tested the activations of neurons from the output layer as predictors. The voxel's time series served as the dependent variable. Before using the time series in a GLM, they had to be preprocessed. This involved smoothing the DQN and human time series using a Gaussian kernel with a full width at half maximum of 5.3 seconds. To assess prediction accuracy, a 5-fold cross-validation procedure was used. The GLMs were fitted to four out of five sessions, and the Pearson correlation coefficient (PCC) between the predicted human time series and the actual one was calculated on the left-out session. This was repeated until each session was used once as a test session.

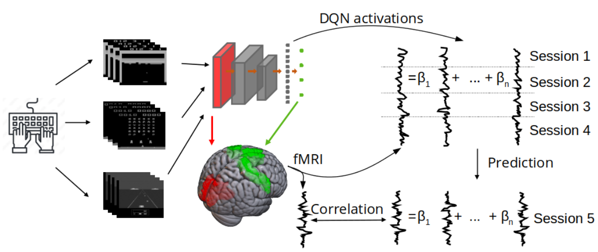

Figure 1: DQN-based encoding model: The human-generated video was processed by the DQN. The time series of neuron activations were used as predictors in a GLM to predict human data. The red-marked region of the brain represents the early visual region of interest (V1), and the green area corresponds to motor areas (MA).

DQNs We compared three encoding models based on a baseline DQN (Mnih et al., 2015), Ape-X (Horgan et al., 2018; Shirokuma, 2020)) and SEED (Espeholt et al., 2019). The baseline DQN is a "vanilla" feed-forward NN consisting of six layers, three of them convolutional. Ape-X and SEED also incorporate three convolutional layers while implementing advanced learning techniques and integrating additional components into their architecture, such as a dueling architecture that separates action evaluation and state evaluation. SEED, the most advanced model, also introduces a long short-term memory, allowing the incorporation of past experiences into decision-making processes. These advancements have led to improved model performances when playing arcade games.

Results

Prediction of Human Behavior

The accuracy of predicting human behavior is depicted in Figure 2. The three encoding models predicted human motor responses significantly above chance level (Fisher-z transformed, one-sample t-test, $p < .001$). SEED exhibited a significantly higher correlation compared to the baseline DQN and Ape-X across all three games (Fisher-z transformed, paired t-test, $p < .01$, Bonferroni correction), confirming our hypothesis at a behavioral level.

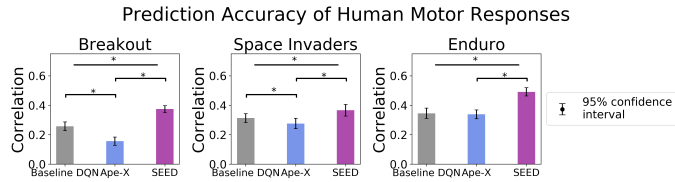

Figure 2: Predicting human behavior: Mean value of the PCC across all test blocks and subjects. Significant differences (Bonferroni correction) are denoted by '*'.

Prediction of Human Brain Activations

In Figure 3, we provided initial evidence, using the example of Enduro, that a DQN-based encoding model could achieve high accuracy in predicting activations of voxels within V1 and within MA (see Figure 1). Preliminary findings suggested a hierarchical correspondence between the layers of a DQN and the stages of visuo-motor processing. For all three DQNs, the difference (gradient) of prediction accuracy from V1 to MA, represented by $r(MA) - r(V1)$, where $r(MA)$ and $r(V1)$ denote the PCC of MA and V1 respectively, increased as the layer index increased (Fisher-z transformed, one-tailed paired t-test, $p < .001$). In SEED, the increase in gradient was most clearly observed, with V1 best predicted by the first layer and MA best predicted by the last layer (Fisher-z transformed, one-tailed paired t-test, $p < .01$).

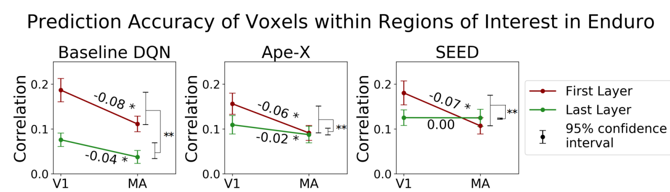

Figure 3: Predicting brain activations: Mean value of the PCC for voxels within V1 (Calcarine sulcus, Cuneus) and MA (precentral gyrus, supplementary motor area) using features of the first and last layers of the DQNs in a layer-specific encoding model. The lines are labeled with the gradient. Significant non-zero gradients are denoted by '*' (one-sample t-test, $p < .005$). A significant increase in gradients is denoted by '**' (one-tailed paired t-test, $p < .001$).

Discussion

Our findings provide initial evidence that DQN-based encoding models can predict behavior and activation patterns related to S-R transformations within time-continuous environments. They confirm our hypothesis that advancements in RL can be leveraged to enhance modeling capabilities. These predictions can be made at a fine-grained temporal scale, complementing the trial-based experimental study designs. This invites an expansion of the approach to hidden layers of the DQNs, investigating whether early layers in the DQN align with early visual areas, while higher DQN layers progress along increasingly higher-order regions of the dorsal stream.

Acknowledgment

The authors acknowledge support by the Deutsche Forschungsgesellschaft (DFG) grant 445383113.

References

Cichy, R. M., & Kaiser, D. (2019). Deep neural networks as scientific models. *Trends in Cognitive Sciences*, *23*(4), 305–317. doi: 10.1016/j.tics.2019.01.009

Cichy, R. M., Khosla, A., Pantazis, D., Torralba, A., & Oliva, A. (2016). Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Scientific Reports*, *6*(1), 27755. doi: 10.1038/srep27755

Cross, L., Cockburn, J., Yue, Y., & O'Doherty, J. P. (2021). Using deep reinforcement learning to reveal how the brain encodes abstract state-space representations in high-dimensional environments. *Neuron*, *109*(4), 724–738. doi: 10.1016/j.neuron.2020.11.021

Eickenberg, M., Gramfort, A., Varoquaux, G., & Thirion, B. (2017). Seeing it all: Convolutional network layers map the function of the human visual system. *NeuroImage*, *152*, 184–194. doi: 10.1016/j.neuroimage.2016.10.001

Espeholt, L., Marinier, R., Stanczyk, P., Wang, K., & Michalski, M. (2019). Seed rl: Scalable and efficient deep-rl with accelerated central inference. In *8th international conference on learning representations.*

Güçlü, U., & van Gerven, M. A. J. (2015). Deep neural networks reveal a gradient in the complexity of neural representations across the brain's ventral visual pathway. *Journal of Neuroscience*, *35*(27), 10005–10014. doi: 10.1523/JNEUROSCI.5023-14.2015

Horgan, D., Quan, J., Budden, D., Barth-Maron, G., Hessel, M., van Hasselt, H., & Silver, D. (2018). Distributed prioritized experience replay. In *International conference on learning representations.*

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... Demis, H. (2015). Human-level control through deep reinforcement learning. *Nature*, *518*(7540), 529–533. doi: 10.1038/nature14236

Mohr, H., Cichy, R. M., & Ruge, H. (2019). Deep neural networks can predict human behavior in arcade games. In *Proceedings of the 2019 conference on cognitive computational neuroscience.* Berlin, Germany. doi: 10.32470/CCN.2019.1043-0

Shirokuma. (2020). *distributedrl.* Retrieved commit: 049c335, from `https://github.com/neka-nat/distributed_rl`

Tomov, M. S., Tsividis, P. A., Pouncy, T., Tenenbaum, J. B., & Gershman, S. J. (2023). The neural architecture of theory-based reinforcement learning. *Neuron*, *111*(8), 1331–1344.e8. doi: 10.1016/j.neuron.2023.01.023