CNNs Reveal the Computational Implausibility of the Expertise Hypothesis

Pranjul Gupta[1,2], Nancy Kanwisher[3,4], & Katharina Dobs[1,2,4]

[1]Dept. Psychology, Justus-Liebig University Giessen, [2]CMBB, JLU,
[3]Dept. Brain and Cognitive Sciences, MIT, [4]CBMM, MIT

Introduction

- Research on face perception provides strong evidence for domain specificity in the ventral visual pathway (e.g. Kanwisher, 2010).
- An alternative hypothesis suggests that neural mechanisms for face perception are not specific to that task, but instead for the fine-grained discrimination of exemplars of any visual categories (e.g. Gauthier et al., 2000; Postle, 2020).
- Here, we used task-optimized convolutional neural networks (CNNs) to test the computational plausibility of this "expertise hypothesis".

Face-specific or expertise areas?

Methods

Task-specific and domain-general CNNs

We trained a task-specific CNN (VGG-16) on face identity discrimination (Face CNN) and a generic CNN on object categorization (Object CNN).

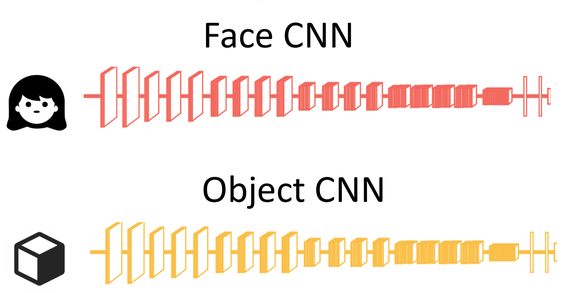

Optimizing both CNNs to fine-grained car discrimination

We choose a fine-grained car dataset with 1109 car model/make categories due to their wide usage in previous work and because they do not contain or resemble faces. We fine-tuned both systems to the car task by updating all except early layers (up to pool3).

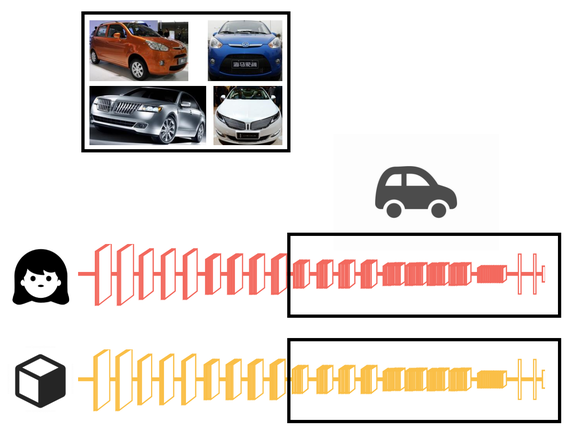

Generic features more useful for fine-grained car discrimination

Decoding car categories from task-optimized CNNs

We tested whether generic or task-specific features are more useful to readily discriminate fine-grained categories by training an SVM on (fixed) activation patterns of the penultimate layer of each CNN for 100 car model/make categories (chance-level: 1%).

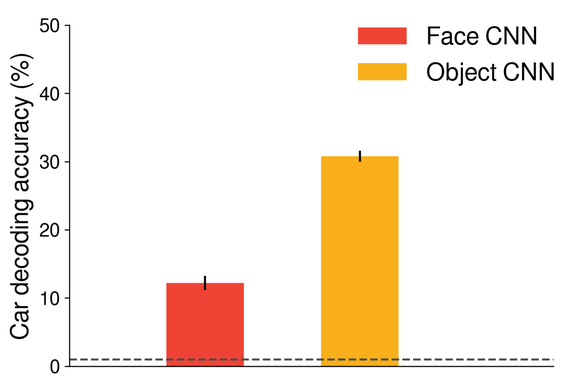

Fine-tuning CNNs to fine-grained car discrimination

We tested whether generic or task-specific features would provide a better foundation for learning a new fine-grained task by fine-tuning each CNN to 1009 car model/make classes.

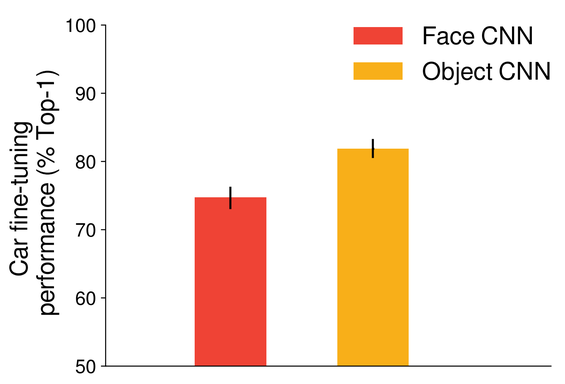

Car task recycles generic instead of task-specific features

Fine-tuning a dual-task CNN optimized for generic and face-specific tasks to discriminate cars

We tested whether a fine-grained discrimination task would recycle generic or face-specific features by fine-tuning a dual-task CNN (Dobs et al., 2022) to the car task and by performing lesion experiments.

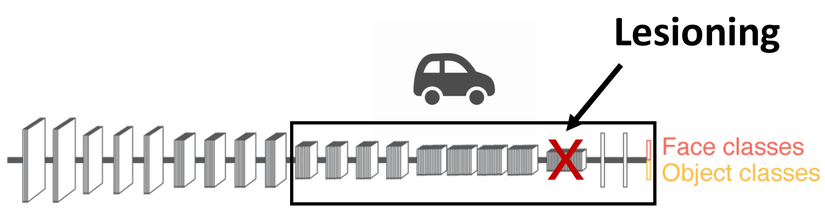

Lesioning the top-20% face-specific features harmed the car performance less than lesioning the top-20% object-specific features.

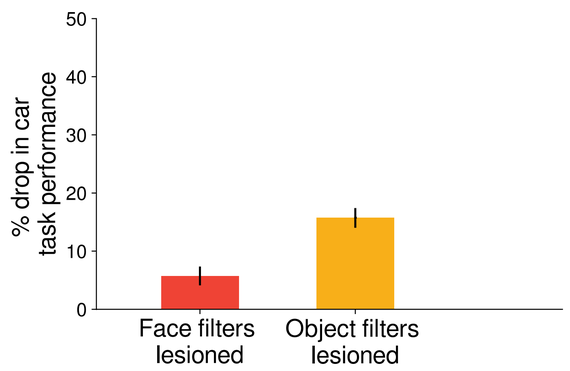

The top-20% car filters recycled more object- than face-specific filters, and additionally relied on redundant filters.

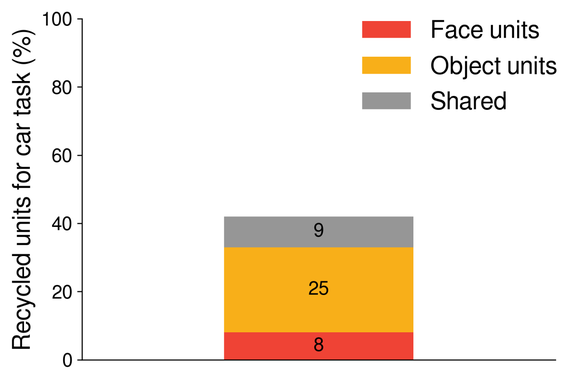

Conclusion

Findings indicate that systems optimized more broadly for object recognition serve as a better foundation than systems optimized for face recognition for subsequent acquisition of car discrimination expertise.

→Our results reveal that the expertise hypothesis does not make sense computationally.

Acknowledgements: This work was supported by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) Project No. 222641018–SFB/TRR 135 TP C9 and S.

@ Pranjul.Gupta@psychol.uni-giessen.de
@_pranjul_gupta

CNNs Reveal the Computational Implausibility of the Expertise Hypothesis

Pranjul Gupta[1,2], Nancy Kanwisher[3,4], & Katharina Dobs[1,2,4]

[1]Dept. Psychology, Justus-Liebig University Giessen, [2]CMBB, JLU,
[3]Dept. Brain and Cognitive Sciences, MIT, [4]CBMM, MIT

Introduction

- Research on face perception provides strong evidence for domain specificity in the ventral visual pathway (e.g. Kanwisher, 2010).
- An alternative hypothesis suggests that neural mechanisms for face perception are not specific to that task, but instead for the fine-grained discrimination of exemplars of any visual categories (e.g. Gauthier et al., 2000; Postle, 2020).
- Here, we used task-optimized convolutional neural networks (CNNs) to test the computational plausibility of this "expertise hypothesis".

Face-specific or expertise areas?

Methods

Task-specific and domain-general CNNs

We trained a task-specific CNN (VGG-16) on face identity discrimination (Face CNN) and a generic CNN on object categorization (Object CNN).

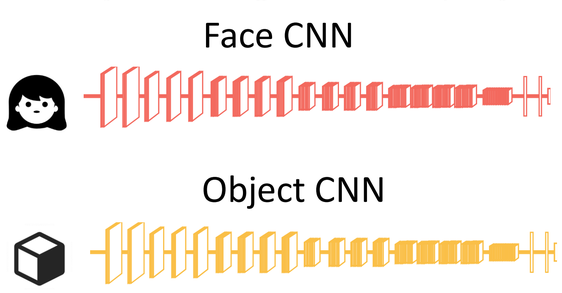

Optimizing both CNNs to fine-grained car discrimination

We choose a fine-grained car dataset with 1109 car model/make categories due to their wide usage in previous work and because they do not contain or resemble faces. We fine-tuned both systems to the car task by updating all except early layers (up to pool3).

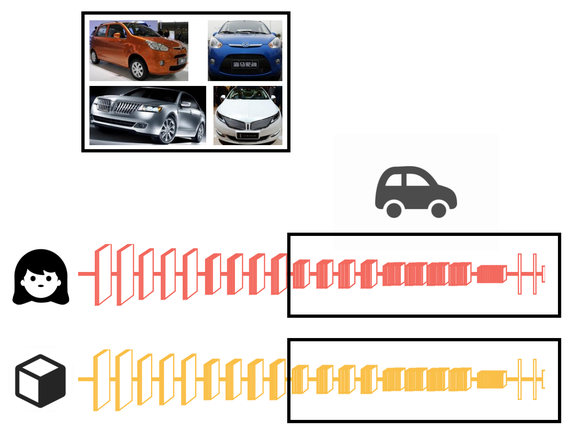

Generic features more useful for fine-grained car discrimination

Decoding car categories from task-optimized CNNs

We tested whether generic or task-specific features are more useful to readily discriminate fine-grained categories by training an SVM on (fixed) activation patterns of the penultimate layer of each CNN for 100 car model/make categories (chance-level: 1%).

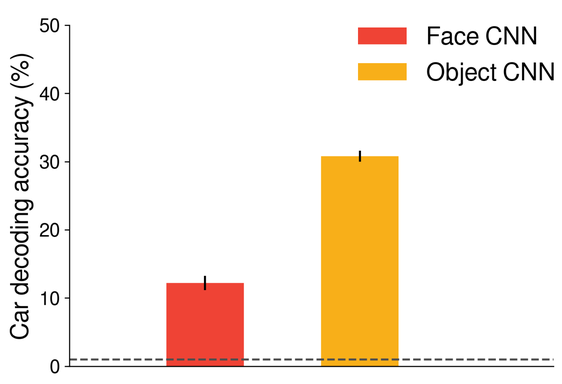

Fine-tuning CNNs to fine-grained car discrimination

We tested whether generic or task-specific features would provide a better foundation for learning a new fine-grained task by fine-tuning each CNN to 1009 car model/make classes.

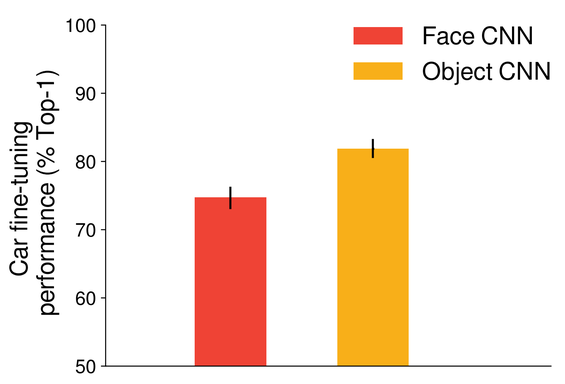

Car task recycles generic instead of task-specific features

Fine-tuning a dual-task CNN optimized for generic and face-specific tasks to discriminate cars

We tested whether a fine-grained discrimination task would recycle generic or face-specific features by fine-tuning a dual-task CNN (Dobs et al., 2022) to the car task and by performing lesion experiments.

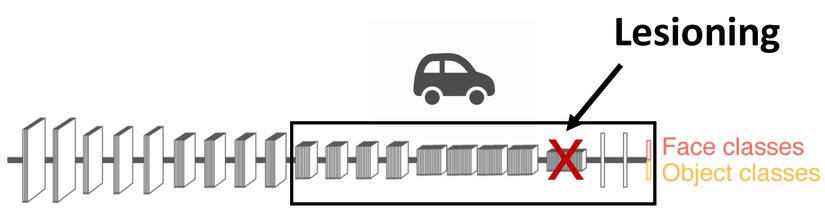

Lesioning the top-20% face-specific features harmed the car performance less than lesioning the top-20% object-specific features.

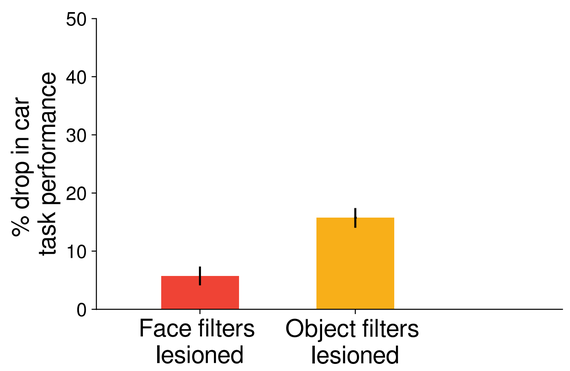

The top-20% car filters recycled more object- than face-specific filters, and additionally relied on redundant filters.

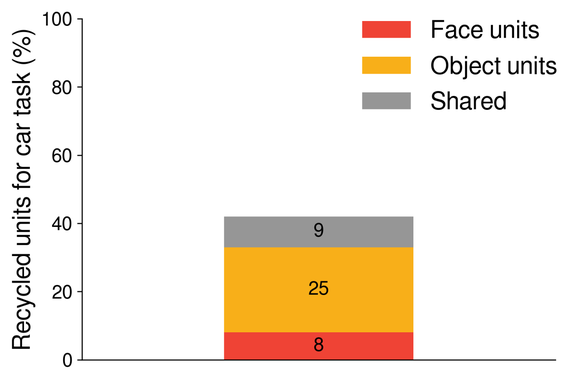

Conclusion

Findings indicate that systems optimized more broadly for object recognition serve as a better foundation than systems optimized for face recognition for subsequent acquisition of car discrimination expertise.

→Our results reveal that the expertise hypothesis does not make sense computationally.

Acknowledgements: This work was supported by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) Project No. 222641018–SFB/TRR 135 TP C9 and S.

@ Pranjul.Gupta@psychol.uni-giessen.de
@_pranjul_gupta