



Context-Driven Visual Object Recognition Based on Knowledge Graphs

Sebastian Monka^{1,2(✉)}, Lavdim Halilaj¹, and Achim Rettinger²

¹ Bosch Research, Renningen, Germany
{sebastian.monka,lavdim.halilaj}@de.bosch.com

² Trier University, Trier, Germany
rettinger@uni-trier.de

Abstract. Current deep learning methods for object recognition are purely data-driven and require a large number of training samples to achieve good results. Due to their sole dependence on image data, these methods tend to fail when confronted with new environments where even small deviations occur. Human perception, however, has proven to be significantly more robust to such distribution shifts. It is assumed that their ability to deal with unknown scenarios is based on extensive incorporation of contextual knowledge. Context can be based either on object co-occurrences in a scene or on memory of experience. In accordance with the human visual cortex which uses context to form different object representations for a seen image, we propose an approach that enhances deep learning methods by using external contextual knowledge encoded in a knowledge graph. Therefore, we extract different contextual views from a generic knowledge graph, transform the views into vector space and infuse it into a DNN. We conduct a series of experiments to investigate the impact of different contextual views on the learned object representations for the same image dataset. The experimental results provide evidence that the contextual views influence the image representations in the DNN differently and therefore lead to different predictions for the same images. We also show that context helps to strengthen the robustness of object recognition models for out-of-distribution images, usually occurring in transfer learning tasks or real-world scenarios.

Keywords: Neuro-symbolic · Knowledge graph · Contextual learning

1 Introduction

How humans perceive the real world is strongly dependent on the context [1, 30]. Especially, in situations with poor quality of visual input, for instance caused by large distances, or short capturing times, context appears to play a major role in improving the reliability of recognition [43]. Perception is not only influenced by co-occurring objects or visual features in the same image, but also by experience and memory [39]. There is evidence that humans perceive similar images




Fig. 1. The mental representation for ambiguous images can change based on the context, although the perceived image is still the same.

differently considering the given context [11]. A famous example are ambiguous figures as shown in Fig. 1.

Depending on the context, i.e. if it is Easter or Christmas [10], Fig. 1a can be either a duck or a rabbit. Likewise, influenced by own-age social biases [37], Fig. 1b can be either a young lady or an old woman. Humans categorize images based on various types of context. Known categories are based on visual features or semantic concepts [6], but may also be based on other information such as attributes describing their function. Accordingly, neuroscience has shown that the human brain encodes visual input into individual contextual object representations [16, 18, 49], namely visual, taxonomical, and functional [33]. Concretely, in a visual context, images of a drum and a barrel have a high similarity, as they share similar visual features. In a taxonomical context, a drum would be similar to a violin, as they both are musical instruments. And in a functional context, the drum would be similar to a hammer, since the same action of hitting can be performed with both objects [8].

Whereas there is much evidence that intelligent machines should also represent information in contextualized embeddings, deep neural networks (DNNs) form their object representations based only on the feature distribution of the image dataset [9, 56]. Therefore, they fail if the objects are placed in an incongruent context that was not present in previous seen images [5].

For the scope of this work we investigate the following research questions:

- **RQ1** - Can context provided in form of a KG influence learning image representations of a DNN, the final accuracy, and the image predictions?
- **RQ2** - Can context help to avoid critical errors in domain changing scenarios where DNNs fail?

To enable standard DNNs to build contextual object representations, we provide the context using a knowledge graph (KG) and its corresponding knowledge graph embedding (\mathbf{h}_{KG}). Similar to the process in the human brain, we conduct experiments with three different types of contexts, namely visual context, taxonomical context, and functional context 3. We provide two versions of knowledge infusion into a DNN and compare the induction of different contextual models in depth by quantitatively investigating their learned contextual embedding spaces using class-related cosine similarities. In addition we evaluate

our approach quantitatively by comparing their final accuracy on object recognition tasks on source and target domains and provide insights and challenges. The structure of this paper is organized as follows: Sect. 6 outlines related work. In Sect. 3.1 we introduce the three different types of context and an option to model these views in a contextual knowledge graph. Section 3 shows two ways of infusing context into a visual DNN. In Sect. 4 we conduct experiments on seven image datasets in two transfer learning scenarios. In Sect. 5 we answer the research questions and summarize the main insights of our approach.

2 Preliminaries

Contextual Image Representations in the Brain. Cognitive and neuroscience research has recently begun to investigate the relationship between viewed objects and the corresponding fMRI scan activities of the human brain. It is assumed that the primate visual system is organized into two separate processing pathways in the visual cortex, namely, the *dorsal pathway* and the *ventral pathway*. While the dorsal pathway is responsible for the spatial recognition of objects as well as actions and manipulations such as grasping, the ventral pathway is responsible for recognizing the type of object based on its form or motion [52]. Bonner et al. [7] recently showed that the sensory coding of objects in the ventral cortex of the human brain is related to statistical embeddings of object or word co-occurrences. Moreover, these object representations potentially reflect a number of different properties, which together are considered to form an object concept [33]. It can be learned based on the context in which the object is seen. For example, an object concept may include the visual features, its taxonomy, or the function of the object [18, 49].

Image Representations in the DNN. Recent work has shown that while the performance of humans, monkeys, and DNNs is quite similar for object-level confusions, the image-level performance does not match between different domains [49]. In contrast to visual object representations in the brain, which also include high level contextual knowledge of concepts and their functions, image representations of DNNs only depend on the statistical co-occurrence of visual features and a specific task. We consider the context extracted from the dataset as dataset bias. Even in balanced datasets, i.e., datasets containing the same number of images for each class, there still exists imbalance due to overlap of features between different classes. For instance, it must be taken into account that a cat and a dog have similar visual features and that in composite datasets certain classes can have different meta-information for the images, such as illumination, perspective or sensor resolution. This dataset bias leads to predefined neighborhoods in the visual embedding space, as well as predefined similarities between distinct classes. In a DNN, an *encoder network* $E(\cdot)$ maps images \mathbf{x} to a visual embedding $\mathbf{h}_v = E(\mathbf{x}) \in \mathbb{R}^{d_E}$, where the activations of the final pooling layer and thus the representation layer have a dimensionality d_E , where d_E depends on the encoder network itself.




Fig. 2. Our approach to learn contextual image representations consists of two main parts: 1) the *contextual view extraction*; and 2) the *contextual view infusion*.

Contextual Representations in the KG. A knowledge graph is a graph of data aiming to accumulate and convey real-world knowledge, where entities are represented by nodes and relationships between entities are represented by edges [21]. We define a *generic knowledge graph (GKG)* as a graph of data that relates different classes of a dataset based on defined contextual properties. These contextual properties can be both learned and manually curated. They bring in prior knowledge about classes, even those that may not necessarily be present in the image dataset, and thus place them in contextual relationships with each other. A KG comprises a set of triples $G = H, R, T$, where H represents entities, $T \subseteq E \times L$ denotes entities or literal values and R , is a set of relationships connecting H and T .

3 Learning Contextual Image Representations

The framework, as shown in Fig. 2 consists of two main parts: 1) the *contextual view extraction*, where task relevant knowledge is extracted from a generic knowledge graph; and 2) the *contextual view infusion*, where the contextual view is infused into the DNN.

3.1 Contextual View Extraction

A knowledge graph can represent prior knowledge encoded with rich semantics in a graph structure. A *GKG* encapsulating n contextual views:

$$GKG \supseteq \{GKG^1, GKG^2, \dots, GKG^n\}$$

is a collection of heterogeneous knowledge sources, where each contextual view defines specific relationships between encoded classes. However, for a particular




Fig. 3. Context can occur in various ways. Aligned to insights of how humans perceive the world, we present three contextual views of a generic knowledge graph, namely the visual, taxonomical, and functional view.

task only a specific part of a GKG can be relevant. Thus, a subgraph containing a single contextual view:

$$GKG^{view} = query(GKG; view)$$

or a combination of views is extracted from a GKG . Since object recognition models are deployed in the real world that differs from their training domain, it is necessary to encode prior knowledge that is not present in the dataset.

Based on image representations in our brain and on how humans tend to classify objects, we introduce three distinct types of contextual views as shown in Fig. 3. The first contextual view is based on visual, the second view is based on taxonomical, and the third view is based on functional properties.

Visual Context. The visual view (GKG^v) describes high-level visual properties of the classes, for instance properties describing color, shape, or texture. These properties may or may not be present in the image data set. For example if all horses in the dataset are white, we want to encode that horses can also occur in different colors.

Taxonomical Context. The taxonomical view (GKG^t) describes class relationships based on hierarchical schemes. A taxonomy is built by experts and can contain categories based on concepts from biology, living place, feeding method, etc. For instance, a biological taxonomy separate animals from vehicles and divides them into further subcategories.

Functional Context. The functional view (GKG^f) contains properties describing the function of a class. It is known that tools are categorized in the human brain based on their function [33]. In that sense properties as hit, rub, or drill would determine the category of a given tool. However, to broaden the scope, additional functional properties such as noise, transport, or smell can be introduced.

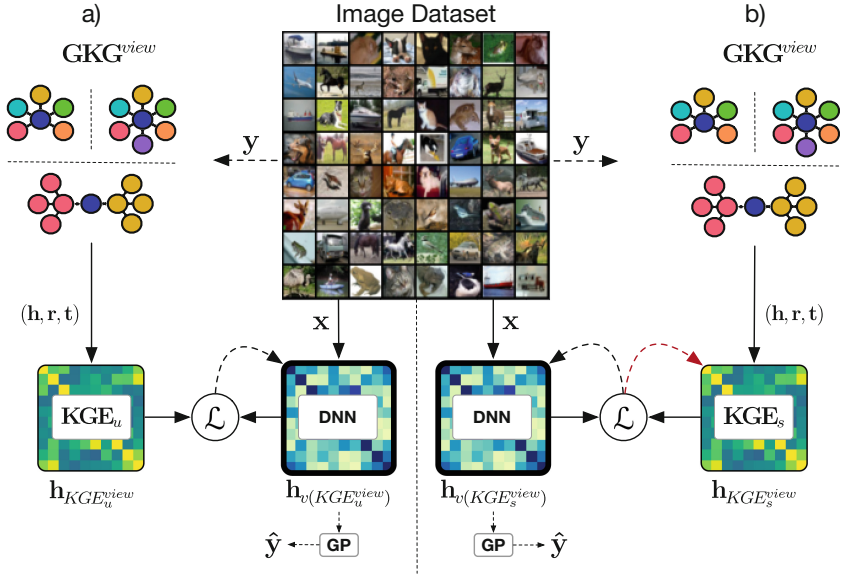


Fig. 4. Contextual view infusion. The contextual object recognition model (DNN) is trained in two different ways: a) using the KG as a trainer, where KGE_u uses no supervision of the image data; or b) using the KG as a peer, where KGE_s uses supervision of the image data. Images \mathbf{x} are fed into the DNN, producing $\mathbf{h}_{v(KGE^{view})}$ which is compared with $\mathbf{h}_{KGE^{view}}$ using the KG-based contrastive loss. In a second step, a gaussian process (GP) or linear layer is trained to predict the class labels \mathbf{y} of \mathbf{x} based on the trained $\mathbf{h}_{v(KGE^{view})}$.

3.2 Contextual View Infusion

When transferring the knowledge from the GKG^{view} using a knowledge graph embedding method (KGE) into a knowledge graph embedding:

$$\mathbf{h}_{KGE^{view}} = KGE(GKG^{view})$$

graph based relationships are transferred into spatial relationships. Intuitively, a different context leads to a different representation in the vector space, where $\mathbf{h}_{KGE^{view}}$ reflects all relationships that are modelled in GKG^{view} .

As illustrated in Fig. 4, we present two different ways of learning a visual context embedding $\mathbf{h}_{v(GKG^{view})}$ following Monka et al. [34]. The first one is $DNN_{KGE_u^{view}}$, which uses the knowledge graph as a trainer [35] and thus learns $\mathbf{h}_{KGE_u^{view}}$ without any supervision of image data. The second version is $DNN_{KGE_s^{view}}$, which uses the knowledge graph as a peer and thus learns $\mathbf{h}_{v(KGE_s^{view})}$ and $\mathbf{h}_{KGE_s^{view}}$ jointly with additional supervision of image data.

Both versions use the contrastive loss to align the image embedding $\mathbf{h}_{v(KGE^{view})}$ of the images \mathbf{x} and the DNN with the knowledge graph embedding $\mathbf{h}_{KGE^{view}}$ of the label information. A batch consists of N augmented training samples. The KG-based contrastive loss is constructed using the individual

anchor losses as given by:

$$\mathcal{L}_{KGE^{view}} = \sum_{i=1}^N \mathcal{L}_{KGE^{view},i}.$$

Within a batch, an anchor image $i \in \{1 \dots 2N\}$ is selected that corresponds to a specific class label \mathbf{y}_i , where \mathbf{y}_i points to its knowledge graph embedding $\mathbf{h}_{KGE^{view},i}$. Positive images j are all images of the batch that correspond to the same class label as the anchor i . The numerator in the loss function computes a similarity score between $\mathbf{h}_{KGE^{view},i}$ and the image embeddings $\mathbf{h}_{v(KGE^{view}),j}$. The denominator computes the similarity score between $\mathbf{h}_{KGE^{view},i}$ and the image embeddings $\mathbf{h}_{v(KGE),k}$ of all images of the other classes in the batch. As a similarity score, we choose the cosine similarity, which however can be replaced by others. $\mathbb{1}_{k \neq i} \in \{0, 1\}$ is an indicator function that returns 1 iff $k \neq i$ evaluates as true, and $\tau > 0$ is a predefined scalar temperature parameter.

$$\mathcal{L}_{KGE^{view},i} = \frac{-1}{2N_{\mathbf{y}_i} - 1} \sum_{j=1}^{2N} \mathbb{1}_{i \neq j} \cdot \mathbb{1}_{\mathbf{y}_i = \mathbf{y}_j} \cdot \log \frac{\exp(\mathbf{h}_{KGE^{view},i} \cdot \mathbf{h}_{v(KGE^{view}),j} / \tau)}{\sum_{k=1}^{2N} \mathbb{1}_{i \neq k} \exp(\mathbf{h}_{KGE^{view},i} \cdot \mathbf{h}_{v(KGE^{view}),k} / \tau)}$$

Prediction. To predict the class labels of unknown images it is common to train a linear layer (*LL*) or to use a gaussian process (*GP*) on top of $\mathbf{h}_{v(KGE^{view})}$. For *GP*, we run the whole training dataset through the trained DNN and calculate the mean and covariance matrices for all the classes in $\mathbf{h}_{v(KGE^{view})}$. *GP* and *LL*, both calculate decision boundaries in $\mathbf{h}_{v(KGE^{view})}$ for all the classes of the dataset. At inference, where the goal is to predict the class label of an unknown image, *GP* or *LL* assign probabilities if an image belongs to a specific class. The maximal probability is chosen to be the final prediction.

4 Experiments

The goal of our empirical investigations is to provide an answer to **RQ1** and **RQ2**. Therefore we conduct experiments with seven datasets in the two specific domain generalization settings, Cifar10 and Mini-ImageNet. For both experiments, we build separate GKGs that include three different contextual views, the visual (GKG^v), the taxonomical (GKG^t), and the functional (GKG^f) view, respectively. Based on the framework in Sect. 3, we use GKG^{view} to learn a contextual DNN in combination with image data. We evaluate and compare both versions of our approach, $DNN_{KGE_u^{view}}$ and $DNN_{KGE_s^{view}}$.

4.1 Implementation Details

For both experiments, we use a similar implementation of our approach. From the *GKG*, we extract various GKG^{view} s using respective SPARQL queries. A

ResNet-18 architecture is used as a DNN-backend, with a 128-dimensional MLP as the head. We train all configurations using an ADAM optimizer, a learning rate of 0.001, no weight decay, and a cosine annealing scheduler with a learning decay rate of 0.1. The images are augmented via random cropping, random horizontal flipping, color jittering, random grayscaling, and resizing to 32×32 pixels. All models are trained for 500 epochs. For a) $DNN_{KGE_u^{view}}$ we transform GKG^{view} into vector space using a graph auto encoder (GAE) [28], which we denote as the $DNN_{GAE^{view}}$ model. Our GAE comprises two convolutional layers, with a hidden layer dimension of 128. We train the GAE using an ADAM optimizer with a learning rate of 0.01 for 500 epochs. For b) $DNN_{KGE_g^{view}}$, a graph attention network (GAT) [45] is trained in combination with the image data, denoted as the $DNN_{GAT^{view}}$ model. The GAT consists of two GAT-layers with 256 hidden dimensions, 8 heads, and an output dimension of 128. Training is performed via the same KG-based contrastive loss from the images in addition to the GKG^{view} input. We optimize the GAT using an ADAM optimizer with a learning rate of 0.001 and no weight decay.

4.2 Experiments on Cifar10

Dataset settings. The source domain Cifar10 [29] consists of 6000 32×32 color images for each of the 10 classes, namely airplane, bird, automobile, cat, deer, dog, horse, frog, ship, and truck. The target domain Stl10 [14] includes 500 96×96 color images for each of the 10 classes, namely airplane, bird, automobile, cat, deer, dog, horse, monkey, ship, and truck.

Knowledge graph construction. We build a GKG that includes the previously discussed three types of context, as shown in Fig. 3. GKG^v contains visual properties like: *hasBackground*: air, forest, water; *hasColor*: black, blue, brown; *hasPart*: eyes, legs, wings; *hasShape*: rectangular, ellipsoid, cross; *hasSize*: large, medium, small; or *hasTexture*: dotted, striped, uniform. GKG^t contains a taxonomy of the classes using the type-relation. For example, the class *Horse* is-a *Mammal* and is-an *Animal* or the class *Ship* is-a *Water-vehicle* and is-a *Vehicle*. GKG^f defines the function of the class, e.g. properties like: *hasMovement*: drive, fly, swim; *hasSound*: bark, meow, vroom; *hasSpeed*: fast, medium, slow; *hasWeight*: heavy, light, middle. Our GKG contains in total 34 classes, 16 object properties, and 65 individuals. Please note that our GKG is only an example and we are aware that there are unlimited possibilities of how and what type of knowledge can be modeled in a knowledge graph.

Evaluation. To evaluate our approach we first investigate the learned embeddings, if and how semantic relationships from GKG^{view} are reflected in $\mathbf{h}_{GKG^{view}}$. Second, we compare the individual class accuracies to see how these relationships influence the final object recognition. Figure 5 shows an analysis: a) the visual view; b) the taxonomical view; and c) the functional view. For every cell in $\mathbf{h}_{GAE^{view}}$ we calculate the cosine similarity between the corresponding nodes, i.e. the classes of the image dataset, and for $\mathbf{h}_{v(GAE^{view})}$ we calculate

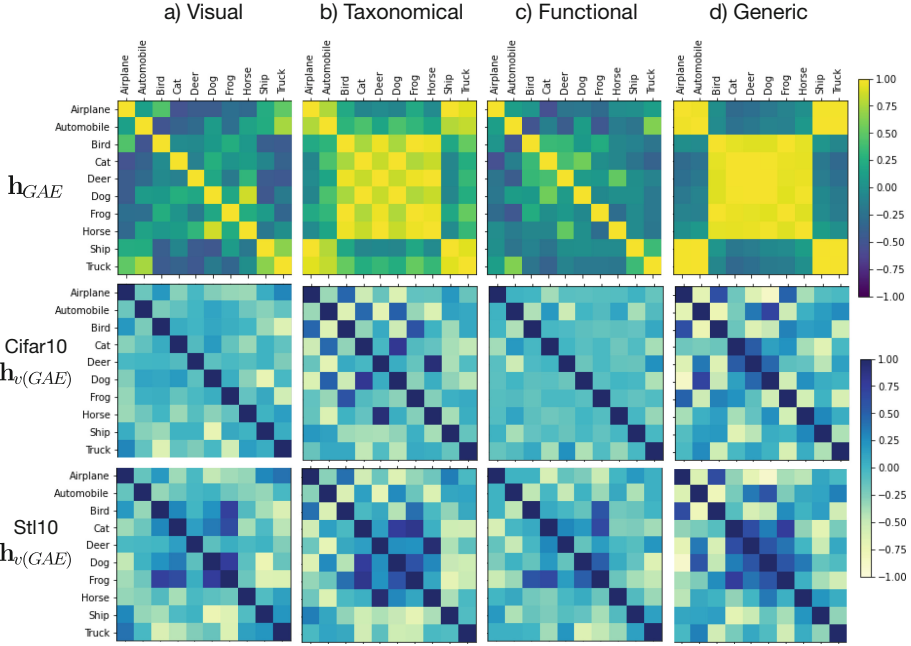


Fig. 5. We compare $h_{GAE^{view}}$ and $h_{v(GAE^{view})}$ based on: a) the visual view; b) the taxonomical view; c) the functional view; and d) the full generic KG. To investigate how the semantic relationships are reflected in the embeddings, we illustrate the individual cosine similarities between the classes of the Cifar10 and the Stl10 dataset.

the class-means of the image representations. Since the goal is to learn contextual image classifiers, we investigate if context is transferred to $h_{GKG^{view}}$ and $h_{v(GKG^{view})}$, respectively. It can be seen that semantic relationships provided by the GKG^{view} are reflected in $h_{GAE^{view}}$. In h_{GAE^v} , the airplane has the highest similarity to the truck and the bird, in h_{GAE^t} , the airplane has the highest similarity to the ship, in h_{GAE^f} , the airplane has the highest similarity to the automobile, and h_{GAE} the airplane has a high similarity to all vehicles. Further, one notices that taxonomical and generic h_{GAE} have two main distinctive groups in the embedding space. In h_{GAE^t} and h_{GAE} vehicles and animals have a high inter-cluster, but a small intra-cluster variance. For $h_{v(GAE^{view})}$, we observe that similarities in the GKG^{view} and $h_{GAE^{view}}$ are only partially reflected. All $h_{v(GAE^{view})}$ seem to have a similar underlying pattern of the class distribution, with minor differences. We think that implicit relations between class features interfere with the similarities given by h_{GAE} and the GKG . Further we retrieve different distributions for either Cifar10 or Stl10. This behaviour can be explained by the distribution shift between source and target domain. While the network attempts to separate classes in the training domain Cifar10, this separation is less successful in the testing domain Stl10.

Table 1. Comparison of the individual class accuracies for the Cifar10 dataset as training domain and the Stl10 dataset as testing domain. We compare the contextual view trained DNNs against their baseline SupSSL.

(a) Results on Cifar10											
Cifar10	Airplane	Auto	Bird	Cat	Deer	Dog	Frog	Horse	Ship	Truck	All
SupSSL	95.1	97.0	91.8	83.9	92.9	85.7	96.0	93.5	96.8	95.9	92.9
ΔDNN_{GAE^v}	-1.2	0.5	-2.6	-0.2	2.3	-0.8	-0.2	-1.1	-0.5	-0.9	-0.5
ΔDNN_{GAE^t}	-0.9	-0.6	-1.2	-30.8	-29.8	-0.2	-2.2	-1.6	-1.5	-1.3	-7.0
ΔDNN_{GAE^f}	1.0	0.2	-1.1	1.9	0.1	0.6	0.7	1.2	-0.1	-0.4	0.4
ΔDNN_{GAE}	-0.7	0.0	-2.3	0.4	0.6	-0.6	-1.1	0.0	0.3	-1.8	-0.5
ΔDNN_{GAT^v}	-0.6	-0.3	0.2	0.3	0.1	-1.0	0.3	0.9	0.7	-0.8	-0.0
ΔDNN_{GAT^t}	-0.9	0.0	-1.7	1.8	0.1	1.0	0.4	0.5	-0.1	0.3	0.1
ΔDNN_{GAT^f}	-0.4	0.5	-3.0	1.7	1.5	-0.4	0.4	0.8	0.4	-0.1	0.1
ΔDNN_{GAT}	-1.0	0.3	-1.8	1.2	-0.3	2.0	-0.5	1.7	0.0	0.7	0.2
(b) Results on Stl10											
Stl10	Airplane	Auto	Bird	Cat	Deer	Dog	Frog	Horse	Ship	Truck	All
SupSSL	85.4	86.9	82.4	56.6	91.5	60.5	—	76.5	84.5	74.1	77.6
ΔDNN_{GAE^v}	1.0	0.2	-2.6	3.4	1.2	-4.5	—	-4.8	-0.6	3.9	-0.3
ΔDNN_{GAE^t}	2.4	-1.0	-1.5	-10.1	-32.9	0.2	—	-0.5	-1.6	-1.6	-5.2
ΔDNN_{GAE^f}	1.9	-0.8	-1.3	1.4	-2.4	-0.5	—	3.4	0.9	3.3	0.7
ΔDNN_{GAE}	0.4	0.5	-1.9	1.8	-1.5	2.6	—	-1.4	-0.6	2.1	0.2
ΔDNN_{GAT^v}	0.5	-0.9	2.6	-0.6	-0.1	0.5	—	0.5	1.0	0.0	0.4
ΔDNN_{GAT^t}	1.0	-2.1	-0.5	1.9	-0.4	0.8	—	0.5	1.6	3.0	0.6
ΔDNN_{GAT^f}	2.7	-0.3	-1.5	-0.7	-1.0	-2.6	—	0.0	0.4	1.8	-0.1
ΔDNN_{GAT}	-1.6	-1.0	-2.6	-2.2	-1.2	2.8	—	3.1	1.2	4.3	0.3

In Table 1 we compare the final object recognition accuracy of the contextual DNNs, compared to their baseline SupSSL. SupSSL is the same model trained with the supervised contrastive loss [27] and without auxiliary context. We observe that for different contextual infusions the overall accuracy is not significantly impacted. For Cifar10 ΔDNN_{GAE^t} with -7.0 is the worst performing model, whereas ΔDNN_{GAE^f} with 0.4 is the best performing model. We marked the best performing model for every class in bold. It can be seen that for every class a different contextual model is outperforming the others. It also shows that context influences the focus a DNN puts on predicting a specific class. Table 1b shows the relative accuracies of the contextual models on the Stl10 dataset. Note that the models are only trained on Cifar10 data. The goal of that domain generalization scenario is to test the robustness of the models. When evaluated on the target domain, it can be observed that almost in every contextual model the relative accuracy is increased compared to the baseline with no contextual knowledge. In scenarios where the domain changes, we observe strange phenomena occurring such that the model with the second worst performance DNN_{GAE}^t for the class Aircraft of the Cifar10 dataset is the model with the second best performance for Aircraft on Stl10. However, for most of the classes, we see a

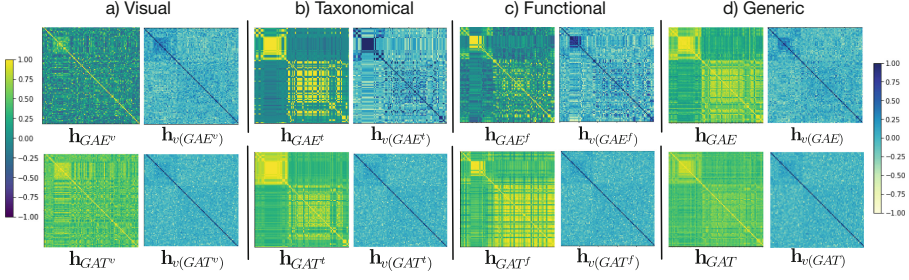


Fig. 6. We compare $\mathbf{h}_{GAE^{view}}$ and $\mathbf{h}_{v(GAE^{view})}$, as well as $\mathbf{h}_{GAT^{view}}$ and $\mathbf{h}_{v(GAT^{view})}$ based on, a) the visual view, b) the taxonomical view, c) the functional view, and d) the full generic KG. To investigate how the semantic relationships are reflected in the embeddings, we illustrate the individual cosine similarities between the classes of the Mini-ImageNet dataset.

trend that the best performing model for a class in Cifar10 tends to perform also better on the target domain.

4.3 Experiments on Mini-ImageNet

Dataset settings. We use Mini-ImageNet, a subset of the ImageNet dataset, as our training domain. It contains 100 classes, each having 600 images of size 84×84 . As testing domain we use ImageNetV2 [40] comprising 10 new test images per class, ImageNet-Sketch [46] with 50 images per class, ImageNet-R [19], which has 150 images in the style of art, cartoons, deviantart, and ImageNet-A [20] with 7.500 unmodified real-world examples.

Knowledge Graph Construction. Our GKG is build using the three contextual views as depicted in Fig. 3. GKG^v contains visual properties, e.g. *hasColor*: black, blue, brown; *hasTexture*: dotted, striped, uniform; *hasSize*: large, medium-large, small; and *hasShape*: ellipsoid, quadratic, rectangular. GKG^t contains a taxonomy of the classes using the type-relation. Following DBpedia [3], the class *Malamute* is-a *Dog*, is-a *Mammal*, is-an *Animal*, is-an *Eukaryote*, and is-a *Species*. GKG^f defines the function of a class with properties like: *hasSpeed*: fast, static, slow; *hasWeight*: heavy, light, middle; or *hasTransportation*: goods, none, people. Our GKG contains in total 166 classes, 14 object properties, and 183 individuals.

Evaluation. Due to the difficulty of deeply investigating 100×100 class similarities, we provide a qualitative overview of the embedding spaces. Figure 6 shows a qualitative comparison of $\mathbf{h}_{KGE^{view}}$ and $\mathbf{h}_{v(KGE^{view})}$ of a) the visual view; b) the taxonomical view; c) the functional view; and d) the generic knowledge graph. Complementing the experiment in Sect. 4.2, we illustrate the class similarities of \mathbf{h}_{GAT} and $\mathbf{h}_{v(GAT)}$ learned using image data as supervision. Interestingly, it can be observed that the similarities in \mathbf{h}_{GAT} and \mathbf{h}_{GAE} follow a similar pattern, but \mathbf{h}_{GAE} seems to have a stronger contrast. However, when investigating

Table 2. Comparison of the contextual view models and their SupSSL baseline on the Mini-ImageNet and its derivatives, Mini-ImageNet (Mini), ImageNetV2 (V2), ImageNet-Sketch (Sketch), ImageNet-R (R), and ImageNet-A (A).

ImageNet	Mini	V2	Sketch	R	A
SupSSL	58.6	43.0	20.3	4.3	1.2
ΔDNN_{GAE^v}	-0.3	0.0	-0.6	0.2	-0.2
ΔDNN_{GAE^t}	-19.6	-13.7	-8.8	-2.8	0.0
ΔDNN_{GAE^f}	-5.2	-3.3	-2.3	-0.7	0.3
ΔDNN_{GAE}	0.8	1.6	-0.6	-0.1	-0.1
ΔDNN_{GAT^v}	0.9	2.3	0.2	0.0	0.3
ΔDNN_{GAT^t}	1.3	0.6	0.1	0.1	0.0
ΔDNN_{GAT^f}	0.4	0.4	0.0	-0.1	-0.1
ΔDNN_{GAT}	0.5	0.6	0.1	0.0	0.0



Fig. 7. Contextual Predictions of DNN_{GAE} (GAE) and DNN_{GAT} (GAT) and their contextual view on Mini-ImageNet. The contextual view influences the image representation and therefore the final prediction for the same input image.

the learned image representations in $\mathbf{h}_{v(GAT)}$ it is hard to spot the differences between the individual contextual models.

As depicted in Table 2 DNN_{GAE}^t and DNN_{GAE}^f are outperformed by the baseline $SupSSL$ and the other models with different contextual views by a large margin. In contrast to the Cifar10 experiment where the least performing model is only 8% worse than the baseline, in Mini-ImageNet the worst is around 34%. Further, we see that DNN_{GAT}^t does not suffer from constraints given by GKG^t . This finding confirms our assumption that a joint training can soften the constraints of the GKG .

Similar to the example of ambiguous figures in Fig. 1, our approach enables DNNs to interpret the same image in various ways using contextual views given by a knowledge graph. The results in Fig. 7 show that for out of distribution images the contextual views play a major role for giving reasonable predictions. The idea is that some class confusions are not that critical as others. In that sense, for some tasks it is uncritical to confuse a goose with a house finch as they

are both part of the bird family, however confusing a music instrument (oboe), with a dog (malamute) could lead to problems. We also see that DNN_{GAE} (GAE) and DNN_{GAT} (GAT) do not necessarily predict the same image based on the given context. We believe that further research is needed w.r.t. investigating how to best incorporate context in combination with image data.

5 Discussion and Insights

With our work, we provided a method to infuse context in form of GKG^{view} into DNNs for visual object recognition. However, knowledge infusion is not straightforward, as problems of machine learning, such as hyper-parameter selection, weight initialization, or dataset dependence, strongly influence the learned representations. Regarding **RQ1** - Can context provided in form of a KG influence learning image representations of a DNN, the final accuracy, and the image predictions? - we list the insights obtained from our investigations:

- **GKG^{view} defines class-relationships.** We showed that various contextual views can be extracted from a GKG and that different views lead to different relationships between classes of the dataset.
- **$h_{KGE^{view}}$ needs to reflect GKG^{view} .** The embedding method itself also influences the $h_{KGE^{view}}$ and the performance of the final prediction model. Context can get lost when transferring GKG^{view} into $h_{KGE^{view}}$. Hard constraints either in GKG^{view} or produced by the KGE-method, e.g. to represent dissimilar classes in $h_{KGE^{view}}$ together, can drastically reduce the prediction accuracy.
- **$h_{GAE^{view}}$ is only partially reflected in $h_{v(GAE^{view})}$.** Since data-driven approaches have a strong dependence on the dataset distribution, $h_{GAE^{view}}$ only influences $h_{v(GAE^{view})}$ to form a hybrid representation. We see that data augmentation weakens the dataset bias and helps to align $h_{v(GAE^{view})}$ with $h_{GAE^{view}}$.
- **Joint training reduces the impact of GKG .** Both the learned $h_{v(KGE_s)}$ and the achieved accuracy values are only slightly affected by the induced GKG . Neither the qualitative evaluation of $h_{v(KGE_s)}$ nor the quantitative evaluation based on accuracy show any significant contextual changes.
- **Context shifts the focus on learning specific classes.** We assume that the context constrains the DNN and its hypothesis space. It is known that DNNs tend to memorize spurious correlations that can lead to catastrophic errors in the real world. We think that the task of our contextual models is to prevent exactly these errors. In our experiments, we showed that specific contextual models performed better on specific classes. We assume that context can shift the overall interest of a DNN to predict a certain class.

- **Context rather influences individual image predictions.** Similar to the proposed motivation of how humans interpret ambiguous figures we see context influencing the prediction of difficult or undefinable images in the dataset. Regarding **RQ2** - Can context help to avoid critical errors in domain changing scenarios where DNNs fail?
- **Context makes more robust against domain changes.**

It can be seen that almost every contextual model increases its relative accuracy compared to the baseline when evaluated on the target domain. Moreover, contextual models that performed better on the source dataset tend to perform better if domain change occurs. We argue that GKG^{view} regularizes the strong dependency on the source domain and thus increases the performance on the target domain.

6 Related Work

Contextual information has always been of great interest for improving computer vision systems. We structure related work into implicit-contextual visual models, explicit-contextual visual models, and contextual knowledge graph embeddings.

Implicit-Contextual Visual Models. Contextualize relationships between visual features that occur in the image itself. They are used for object priming, where the context defines a prior on the detection parameters [43] or for object detection and segmentation, where boosting is used to relate objects in an image [44]. Wu et al. [51] improved object recognition by processing object regions and context regions in parallel. To overcome the drawback of small receptive fields from standard CNNs, extensions that incorporate visual features from far image regions [24, 25] or alternative architectures, such as vision transformers (ViTs) [53] have been established recently. Moreover, Gao et al. [17] proposed that all modern DNNs are part of the implicit-contextual models since they aggregate contextual information over image regions.

Explicit-Contextual Visual Models. Use higher level information like object co-occurrences or semantic concept relationships. They induce additional contextual information that is either not in the dataset or cannot be automatically extracted by the DNN [22]. To create explicit context based on object relations, most methods use scene graphs which describe a scene based on symbolic representations of entities and their spatial and semantic relations. Scene graphs have been applied to the task of collective or group activity recognition [13, 15], object recognition [55, 56], object detection [12, 32] and visual question answering [42]. Label graphs [23] apply fine-grained labels to an image and are used to improve object recognition and reasoning over object relationships [4]. Semantic scene graphs extend scene graphs by textual descriptions and fine-grained labels of a scene [31]. Context-aware zero-shot learning for object recognition [54] or compositional zero-shot learning methods [36] add observed visual primitive states

(e.g. old, cute) to objects (e.g. car, dog) to build an embedding space based on visual context. However, scene correlations need to be addressed very carefully, as implicit-contextual models can heavily depend on learned contextual relationships that are only valid for a specific dataset configuration. Therefore, work was already done to decorrelate objects and their visual features to improve model generalization [41].

Contextual Knowledge Graph Embeddings. Whereas our approach extracts the contextual views in a previous step before the actual knowledge graph embedding, there exist works that create contextualized KG embeddings based on the full KG. Werner et al. [50] introduced a KG embedding over temporal contextualized KG facts. Their recurrent transformer enables to transform global KGEs into contextual embeddings, given the situation-specific factors of the relation and the subjective history of the entity. Ning et al. [38] proposed a lightweight framework for the usage of context within standard embedding methods. Wang et al. [47] presented a deep contextualized knowledge graph embedding method that learns representations of entities and relations from constructed contextual entity-relation chains. Wang et al. [48] introduced the contextualized KG embedding method (CoKE). They propose to take the contextual nature of KGs into account, by learning dynamic, flexible, and fully contextualized entity and relation embeddings.

7 Conclusion and Future Work

In this work, we proposed a framework for context-driven visual object recognition based on knowledge graphs. We qualitatively and quantitatively investigated how different contextual views, as well as their embedding and their infusion method, influence the learned DNN. Further, we have seen that contextual models tend to have a minor impact on the final accuracy, but a major impact on how individual classes or images are represented and predicted. In particular, for out of distribution data, where data-driven approaches suffer from less knowledge, contextual image representations help to constrain the hypothesis space, leading to more reasonable predictions. However, there are still challenges to be faced. We conducted intensive research about a possible context infusion approach and emerging challenges. On the one hand, we have the implementation of the infusion method, which itself heavily depends on modeling choices, weight initialization, as well as network and hyper-parameter selection. On the other hand, there is a strong dependence on the image data, which originally comes with an initial dataset bias. This dataset bias limits the ability to influence image data representations and thus predictions influenced by prior knowledge. However, our work showed that with deeper investigations of all the influencing parameters knowledge-infused learning is a promising approach to build context-driven and future intelligent systems.

Acknowledgement. This publication was created as part of the research project "KI Delta Learning" (project number: 19A19013D) funded by the Federal Ministry for Economic Affairs and Energy (BMWi) on the basis of a decision by the German Bundestag.

References

1. The role of context in object recognition. *Trends in Cognitive Sciences* (2007)
2. Attneave, F.: Multistability in perception. *Sci. Am.* **225**(6), 63–71 (1971)
3. Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z.G.: DBpedia: a nucleus for a web of open data. In: *The Semantic Web - 6th International Semantic Web Conference ISWC* (2007)
4. Battaglia, P.W., Pascanu, R., Lai, M., Rezende, D.J., Kavukcuoglu, K.: Interaction networks for learning about objects, relations and physics. In: *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems* (2016)
5. Beery, S., Van Horn, G., Perona, P.: Recognition in terra incognita. In: Ferrari, V., Hebert, M., Sminchisescu, C., Weiss, Y. (eds.) *ECCV 2018. LNCS*, vol. 11220, pp. 472–489. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01270-0_28
6. Biederman, I.: Recognition-by-components: a theory of human image understanding. *Psychol. Rev.* **94**(2), 115–147 (1987)
7. Bonner, M., Epstein, R.: Object representations in the human brain reflect the co-occurrence statistics of vision and language. *Nat. Commun.* **12**, 4081 (2021)
8. Bracci, S., Daniels, N., Op de Beeck, H.: Task Context Overrides Object- and Category-Related Representational Content in the Human Parietal Cortex. *Cereb. Cortex* **27**(1), 310–321 (2017)
9. Brendel, W., Bethge, M.: Approximating CNNs with bag-of-local-features models works surprisingly well on imageNet. In: *7th International Conference on Learning Representations, ICLR* (2019)
10. Brugger, P., Brugger, S.: The Easter bunny in October: is it disguised as a duck? *Percept. Mot. Skills* **76**, 2 (1993)
11. Chambers, D., Reisberg, D.: Can mental images be ambiguous? *J. Exp. Psychol. Human Perception Perform.* **11**(3), 317–328 (1985)
12. Chen, X., Li, L., Fei-Fei, L., Gupta, A.: Iterative visual reasoning beyond convolutions. In: *Conference on Computer Vision and Pattern Recognition, CVPR* (2018)
13. Choi, W., Savarese, S.: A unified framework for multi-target tracking and collective activity recognition. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) *ECCV 2012. LNCS*, vol. 7575, pp. 215–230. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33765-9_16
14. Coates, A., Ng, A., Lee, H.: An analysis of single-layer networks in unsupervised feature learning. In: *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics* (2011)
15. Deng, Z., Vahdat, A., Hu, H., Mori, G.: Structure inference machines: recurrent neural networks for analyzing relations in group activity recognition. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR* (2016)
16. DiCarlo, J.J., Cox, D.D.: Untangling invariant object recognition. *Trends in Cogn. Sci.* **11**(8), 333–341 (2007)

17. Gao, P., Lu, J., Li, H., Mottaghi, R., Kembhavi, A.: Container: context aggregation networks. In: *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems* (2021)
18. Greene, M.R., Hansen, B.C.: Disentangling the independent contributions of visual and conceptual features to the spatiotemporal dynamics of scene categorization. *bioRxiv* (2020)
19. Hendrycks, D., et al.: The many faces of robustness: a critical analysis of out-of-distribution generalization. *arXiv preprint [arXiv:2006.16241](https://arxiv.org/abs/2006.16241)* (2020)
20. Hendrycks, D., Zhao, K., Basart, S., Steinhardt, J., Song, D.: Natural adversarial examples. *arXiv preprint [arXiv:1907.07174](https://arxiv.org/abs/1907.07174)* (2019)
21. Hogan, A., et al.: Knowledge graphs. *Synthesis Lectures on Data, Semantics, and Knowledge* (2021)
22. Hoiem, D., Efros, A.A., Hebert, M.: Geometric context from a single image. In: *International Conference on Computer Vision ICCV*, Computer Society (2005)
23. Hu, H., Zhou, G., Deng, Z., Liao, Z., Mori, G.: Learning structured inference neural networks with label relations. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, Computer Society (2016)
24. Hu, J., Shen, L., Albanie, S., Sun, G., Vedaldi, A.: Gather-excite: exploiting feature context in convolutional neural networks. In: *Advances in Neural Information Processing Systems: Annual Conference on Neural Information Processing Systems* (2018)
25. Hu, J., Shen, L., Albanie, S., Sun, G., Wu, E.: Squeeze-and-excitation networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **42**, 1–34 (2020)
26. Jastrow, J.: *Fact and fable in psychology*. D Appleton & Company, New York (1900)
27. Khosla, P., et al.: Supervised contrastive learning. In: *Advances in Neural Information Processing Systems: Annual Conference on Neural Information Processing Systems* (2020)
28. Kipf, T.N., Welling, M.: Variational graph auto-encoders. *arXiv preprint [arXiv:1611.07308](https://arxiv.org/abs/1611.07308)* (2016)
29. Krizhevsky, A.: Learning multiple layers of features from tiny images. *Technology Reports* (2009)
30. Lauer, T., Schmidt, F., Vö, M.: The role of contextual materials in object recognition. *Sci. Rep.* **11**, 21988 (2021)
31. Li, Y., Zhang, D., Mu, Y.: Visual-semantic matching by exploring high-order attention and distraction. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR* (2020)
32. Liu, Y., Wang, R., Shan, S., Chen, X.: Structure inference net: Object detection using scene-level context and instance-level relationships. In: *IEEE Conference on Computer Vision and Pattern Recognition, CVPR* (2018)
33. Martin, A.: GRAPES—Grounding representations in action, perception, and emotion systems: how object properties and categories are represented in the human brain. *Psychon. Bull. Rev.* **23**(4), 979–990 (2016). <https://doi.org/10.3758/s13423-015-0842-3>
34. Monka, S., Halilaj, L., Rettinger, A.: A survey on visual transfer learning using knowledge graphs. *Semantic Web* **13**(3), 477–510 (2022)
35. Monka, S., Halilaj, L., Schmid, S., Rettinger, A.: Learning visual models using a knowledge graph as a trainer. In: *The Semantic Web - 20th International Semantic Web Conference, ISWC* (2021)

36. Naeem, M.F., Xian, Y., Tombari, F., Akata, Z.: Learning graph embeddings for compositional zero-shot learning. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2021)
37. Nicholls, M.E.R., Churches, O., Loetscher, T.: Perception of an ambiguous figure is affected by own-age social biases. *Sci. Rep.* **8**, 12661 (2018)
38. Ning, Z., Qiao, Z., Dong, H., Du, Y., Zhou, Y.: LightCAKE: a lightweight framework for context-aware knowledge graph embedding. In: Karlapalem, K., et al. (eds.) PAKDD 2021. LNCS (LNAI), vol. 12714, pp. 181–193. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-75768-7_15
39. Rafetseder, E., et al.: Children struggle beyond preschool-age in a continuous version of the ambiguous figures task. *Psychol. Res.* **85**(2), 828–841 (2019). <https://doi.org/10.1007/s00426-019-01278-z>
40. Recht, B., Roelofs, R., Schmidt, L., Shankar, V.: Do imageNet classifiers generalize to imageNet? In: ICML (2019)
41. Singh, K.K., Mahajan, D., Grauman, K., Lee, Y.J., Feiszli, M., Ghadiyaram, D.: Don't judge an object by its context: Learning to overcome contextual bias. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR. Computer Vision Foundation (2020)
42. Teney, D., Liu, L., van den Hengel, A.: Graph-structured representations for visual question answering. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR (2017)
43. Torralba, A.: Contextual priming for object detection. *Int. J. Comput. Vis.* **53**, 169–191 (2003). <https://doi.org/10.1023/A:1023052124951>
44. Torralba, A., Murphy, K.P., Freeman, W.T.: Contextual models for object detection using boosted random fields. In: Neural Information Processing Systems NIPS (2004)
45. Velickovic, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph attention networks. In: International Conference on Learning Representations ICLR (2018)
46. Wang, H., Ge, S., Lipton, Z., Xing, E.P.: Learning robust global representations by penalizing local predictive power. In: NeurIPS (2019)
47. Wang, H., Kulkarni, V., Wang, W.Y.: Dolores: deep contextualized knowledge graph embeddings. In: Conference on Automated Knowledge Base Construction, AKBC 2020, Virtual, 22–24 June 2020 (2020)
48. Wang, Q., et al.: Coke: contextualized knowledge graph embedding (2019)
49. Wardle, S.G., Baker, C.I.: Recent advances in understanding object recognition in the human brain: deep neural networks, temporal dynamics, and context. *F1000Res.* **9**, 590 (2020)
50. Werner, S., Rettinger, A., Halilaj, L., Lüttin, J.: RETRA: recurrent transformers for learning temporally contextualized knowledge graph embeddings. In: Verborgh, R., et al. (eds.) ESWC 2021. LNCS, vol. 12731, pp. 425–440. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-77385-4_25
51. Wu, K., Wu, E., Kreiman, G.: Learning scene gist with convolutional neural networks to improve object recognition. In: 52nd Annual Conference on Information Sciences and Systems CISS (2018)
52. Yang, X., Yan, J., Wang, W., Li, S., Hu, B., Lin, J.: Brain-inspired models for visual object recognition: an overview. *Artificial Intelligence Review* (2022)
53. Yu, Q., Xia, Y., Bai, Y., Lu, Y., Yuille, A., Shen, W.: Glance-and-gaze vision transformer (NeurIPS) (2021)

54. Zablocki, E., Bordes, P., Soulier, L., Piwowarski, B., Gallinari, P.: Context-aware zero-shot learning for object recognition. In: Proceedings of the 36th International Conference on Machine Learning ICML (2019)
55. Zhang, M., Feng, J., Montejo, K., Kwon, J., Lim, J.H., Kreiman, G.: Lift-the-flap: Context reasoning using object-centered graphs. arXiv preprint [arXiv:1902.00163](https://arxiv.org/abs/1902.00163) (2019)
56. Zhang, M., Tseng, C., Kreiman, G.: Putting visual object recognition in context. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition CVPR (2020)