ViSpa (Vision Spaces): A computer-vision-based representation system for individual images and concept prototypes, with large-scale evaluation

Fritz Günther¹, Marco Marelli^{2,3}, Sam Tureski⁴, and Marco Alessandro Petilli²

¹Humboldt-Universität zu Berlin, Berlin, Germany

²University of Milano–Bicocca, Milan, Italy

³NeuroMI, Milan Center for Neuroscience, Milan, Italy

⁴University of Tübingen, Tübingen, Germany

Author Note

Note: This is the author’s preprint version of the article (date: 14. 09. 2022). The final article is accepted for publication in *Psychological Review*.

This work was supported by a Research Grant (project number 452322374) from the German Research Foundation (DFG), awarded to Fritz Günther. We thank Dermot Lynott and two anonymous reviewers for their invaluable comments on earlier versions of this manuscript, and we thank Martin Hebart for many helpful pointers to the literature. All authors declare that they have no conflict of interest.

All materials, datasets, and analysis scripts for the studies reported here are available at <https://osf.io/qvw9c>. All studies received ethical approval by the Ethics Committee for Psychological Research at the University of Tübingen (Az Günther_2020_0603_190). This study was not preregistered. All images in this paper are in the public domain and visually similar to the actual image material used in the studies; see the Supplementary Material for sources.

Abstract

Quantitative, data-driven models for mental representations have long enjoyed popularity and success in psychology (for example, distributional semantic models in the language domain), but have largely been missing for the visual domain. To overcome this, we present *ViSpa* (Vision Spaces), high-dimensional vector spaces that include vision-based representation for naturalistic images as well as concept prototypes. These vectors are derived directly from visual stimuli through a deep convolutional neural network (DCNN) trained to classify images, and allow us to compute vision-based similarity scores between any pair of images and/or concept prototypes. We successfully evaluate these similarities against human behavioral data in a series of large-scale studies, including off-line judgments – visual similarity judgments for the referents of word pairs (Study 1) and for image pairs (Study 2), and typicality judgments for images given a label (Study 3) – as well as on-line processing times and error rates in a discrimination (Study 4) and priming task (Study 5) with naturalistic image material. *ViSpa* similarities predict behavioral data across all tasks, which renders *ViSpa* a theoretically appealing model for vision-based representations and a valuable research tool for data analysis and the construction of experimental material: *ViSpa* allows for precise control over experimental material consisting of images and/or words denoting imageable concepts, and introduces a specifically vision-based similarity for word pairs. To make *ViSpa* available to a wide audience, this article a) includes (video) tutorials on how to use *ViSpa* in R, and b) presents a user-friendly web interface at <http://vispa.fritzguenther.de>.

Keywords: Vision models; Convolutional neural networks; Mental Representations; Perceptual Representations; Visual similarity

ViSpa (Vision Spaces): A computer-vision-based representation system for individual images and concept prototypes, with large-scale evaluation

As theoretical constructs – explanatory variables that cannot be observed directly – it is essential that the scientific concepts postulated in psychology and cognitive science can be quantitatively characterized and thus operationalized in a precise, non-subjective manner. This includes mental representations (Pitt, 2020; E. R. Smith, 1998), which can be defined as “an encoding of some information, which an individual can construct, retain in memory, access, and use in various ways” (E. R. Smith, 1998, p. 391).

Oftentimes, such quantitative characterizations of mental representations are induced from behavioral data, and thus from the *outcome level*. Typical examples are studies in which participants are asked to generate semantic features for given words in order to extrapolate semantic representations for these words, which are then used to predict behavioral performance in other tasks (McRae, De Sa, & Seidenberg, 1997; Vigliocco, Vinson, Lewis, & Garrett, 2004; see the discussion in Andrews, Vigliocco, & Vinson, 2009). Similarly, other studies use data from free association tasks (where participants are presented with words and are asked to generate the first words that come to their mind, De Deyne, Kenett, Anaki, Faust, & Navarro, 2016; Kenett, Levi, Anaki, & Faust, 2017; Kenett, Kenett, Ben-Jacob, & Faust, 2011) or similarity judgments (Benedek et al., 2017; Hebart, Zheng, Pereira, & Baker, 2020) to estimate the semantic representations that produced this behavior (in the form of, for example, semantic networks or spatial models). The distances between these representations can then be used to predict behavior in other tasks.

However, such approaches move from the same level of description of the phenomenon to be investigated – behavioral data in one task is ultimately always explained by (condensed) behavioral data in some other task, inevitably confusing the *explanans* with the *explanandum* (see Westbury, 2016, for an in-depth discussion). This is analogous to asking people how similar given colors are in order to set up a psychological model of color

perception – a good *description* of the system that can however not *explain* why some colors are perceived as more similar than others. A more sophisticated account would instead move from the actual stimuli experienced by humans (light of a certain wavelength) and build on how they are processed by the cognitive system in order to explain behavior. In fact, it should ultimately be the aim of our scientific endeavor to model mental representations starting from the *input level* – the arrangements of stimuli that humans experience, and from which they form their representations (Petilli, Günther, Vergallito, Ciapparelli, & Marelli, 2021; Westbury, 2016). One highly influential class of psychological models adopting this principle are the connectionist models developed in the 1980s (McClelland, Rumelhart, & The PDP Research Group, 1986; Rumelhart, McClelland, & The PDP Research Group, 1986), which adapted and further developed the neural network models that nowadays dominate computer science and artificial intelligence research (especially since the arrival of deep learning models; LeCun, Bengio, & Hinton, 2015).

In the field of semantic memory research, the class of distributional semantic models has risen to considerable prominence (Günther, Rinaldi, & Marelli, 2019; Jones, Willits, & Dennis, 2015; Landauer & Dumais, 1997; Mandera, Keuleers, & Brysbaert, 2017). Focusing on the input level as well, these models approximate the accumulated language experience of humans via large text corpora, and derive semantic representations for words as high-dimensional numerical vectors that are based on the word distributional patterns over linguistic contexts. While earlier models were built on word-by-context counts (Landauer & Dumais, 1997; Lund & Burgess, 1996), current state-of-the-art models are based on neural-network architectures (Baroni, Dinu, & Kruszewski, 2014; Mandera et al., 2017; see however also Lenci, Sahlgren, Jeuniaux, Gyllensten, & Milianni, 2021; Wingfield & Connell, 2022). The cognitive adequacy of these models – their theoretical plausibility as well as empirical performance in predicting human behavioral data – has been established in a wide range of studies, corroborating the view that the semantic representations they induce can indeed serve as quantitative operationalizations for human mental representations (see

Günther et al., 2019; Jones et al., 2015, for recent reviews).

The pioneering efforts in the distributional modelling tradition had a heavy focus on language information in their characterization of semantic representations. This focus was mostly dictated by practical convenience rather than theoretical considerations (Günther et al., 2019): The required linguistic data approximating human (language) experience has long been available in the form of large digitalized text corpora, which are relatively easy to handle and process computationally. However, from the very start, the resulting language-centric approach has been heavily criticized as only reflecting a fraction of human experience (Andrews et al., 2009; Borghesani & Piazza, 2017; Glenberg & Robertson, 2000; Rogers et al., 2004). Any model of mental representations – even when one only considers representations in semantic memory – that only considers linguistic experience is therefore deemed to be necessarily incomplete (Andrews et al., 2009). This view that linguistic or language-based representations are only one part of the picture is firmly established in many prominent models of human memory: The working memory model by Baddeley includes a visuospatial alongside a phonological system (Baddeley, 1992; Baddeley & Hitch, 1974), which are assumed to interact with a “visual semantics” as well as a “language” system in long-term memory (Baddeley, 2000). Similarly, the dual coding theory by Paivio centrally assumes two different processing channels, verbal as well as visual, that create separate types of mental representations (Paivio, 1986, 1991). Sensory representations, similar or equal to actual perceptual experience, are also at the theoretical core of the embodied or grounded cognition approach (Barsalou, 1999; Glenberg & Robertson, 2000; Zwaan & Madden, 2005), which assumes that all human cognition is inherently linked to and grounded in perception and action.

In the cognitive sciences, this gap between language – where data-driven, quantitatively characterized models of mental representations have a successful history – and perception – where such models have long been missing – has been identified as an unsatisfactory imbalance (Andrews et al., 2009; Pezzulo et al., 2011, 2013). However, in

the visual domain – the by far most prominent perceptual dimension for humans (Lynott, Connell, Brysbaert, Brand, & Carney, 2020) and the one featured in well-established theoretical models (Baddeley, 2000; Paivio, 1986) – the advent of powerful computer-vision models (typically again employing (deep) neural network architectures; Chatfield, Simonyan, Vedaldi, & Zisserman, 2014; Krizhevsky, Sutskever, & Hinton, 2012) has provided us with promising candidates to fill up this gap (see Andrews, Frank, & Vigliocco, 2014). While such models have been used in psychological studies before (see the section *Deep convolutional neural network models in cognitive science*), a thorough large-scale evaluation across different tasks and paradigms is still missing.

In the present article, we employ such a model that can generate vision-based representations for individual images, and extend it to also produce prototypical representations capturing vision-based conceptual representations. We then present a large-scale evaluation of the model-obtained estimates in a series of five large-scale studies, each employing hundreds of participants and thousands of items, in order to empirically establish its validity for psychological research.

Vision-based representations in deep convolutional neural network models

In order to model vision-based mental representations, we employ a deep convolutional neural network model (DCNN) that originated in the field of computer vision (Krizhevsky et al., 2012). These models are trained as image classifiers, aimed at predicting an image label from the RGB-codes for each pixel of an input image, with high levels of performance in this task (Chatfield et al., 2014, report top-5 error rates of around .13, meaning that the correct label was included in the top five model predictions in 87% of cases).

Specifically, we examine the VGG-F model (Chatfield et al., 2014), an eight-layer variant of the VGG architecture (of which there are also deeper variants with 16 or 19 layers, Simonyan & Zisserman, 2014). We opted for a VGG model because these architectures are extremely prominent in computer vision, and have also performed well in

cognitively-oriented studies (Jozwik, Kriegeskorte, Storrs, & Mur, 2017; Kubilius, Bracci, & Op de Beeck, 2016; Peterson, Abbott, & Griffiths, 2016, 2018; Petilli et al., 2021; Roads & Love, 2021, see the section *Deep convolutional neural network models in cognitive science*). Although deeper models perform better for the image classification task they are trained on (Simonyan & Zisserman, 2014), we selected the relatively simpler VGG-F variant since previous studies have shown that these smaller models perform at least as well for explaining human behavioral data while being substantially less complex (Petilli et al., 2021; see also Jozwik et al., 2017).

A schematic representation of the network architecture is presented in Figure 1; for a more detailed mathematical description, see Chatfield et al. (2014) and Vedaldi and Lenc (2015). The first five layers (conv1–conv5), after the input layer 0, constitute the convolutional layers of the network (i.e., their neurons are arranged in a three-dimensional structure, see Figure 1). This is most intuitive for the input layer: the images fed to the network are standardized to a resolution of 224*224 pixels, which determine the height and width of layer 0, and the depth of 3 is determined by the RGB code of these pixels. Differently from a regular neural network, neurons in a convolutional layer do not receive input from all neurons in the previous layer, but only from a local region (as indicated by the pyramid-shaped connections in Figure 1). The final three layers (fc6–fc8) are fully-connected layers where, as suggested by the name, each neuron receives input from all neurons in the previous layer – that is, the type of layer used in regular neural networks.

Here, we employ a pre-trained version of the network available in the *MatConvNet* library for MATLAB (Vedaldi & Lenc, 2015). This model was trained on around 1.2 million individual images from 1,000 different ImageNet categories (Deng et al., 2009), a large collection of labelled images adopting the hierarchical category structure of WordNet (Fellbaum, 1998). As can be seen in Figure 1, any given input information is passed on from one layer to the next until the model arrives at an output. For any given input, the activation values in layer n determine the activation values in layer $n + 1$. *How* exactly this

Figure 1. A graphical illustration of the VGG-F model architecture, with a tensor encoding the RGB values for each pixel of an image serving as input (the images are standardized to 224*224 pixels), and a one-hot vector for image labels as output.

information is passed on (i.e., the weights between the elements in these layers) is optimized during model training.

DCNNs primarily rely on various visual features of the images to perform their task (Geirhos, Jacobsen, et al., 2020), most importantly (combinations of) color and shape (Kriegeskorte, 2015; LeCun et al., 2015; van Dyck & Gruber, 2020; these features include line orientations in various degrees of complexity, including local shapes, Baker, Lu, Erlikhman, & Kellman, 2018; as well as textures, Geirhos, Rubisch, et al., 2018). Previous studies have argued that the different layers of these models encode different types of information: While the first layers encode low-level perceptual information like line orientations and color blobs, deeper layers start to capture basic textures and object parts, and the final layers already represent object classes and complex gestalt representations that ultimately form the basis for label classifications (LeCun et al., 2015; Mahendran & Vedaldi, 2015; Zeiler & Fergus, 2014; see also

<https://donglaiw.github.io/proj/mneuron> for a demonstration). Similar to network-based models in the language domain such as *word2vec* (Mandera et al., 2017; Mikolov, Chen, Corrado, & Dean, 2013), the hidden layers of the VGG-F network can be taken as candidates for mental representations of the input – with their psychological plausibility to be established in empirical investigations.

In the present work, we aim for a large-scale empirical evaluation of these representations from a cognitive perspective, as quantitative characterizations of vision-based representations in humans. We are thus focusing on the *outputs* produced by the model (i.e., the activation values in the hidden layers) for a given input, and in particular on the similarities between these representations and to what extent they correspond to how similar human representations of these stimuli are. Since we cannot observe these human representations directly, our evaluation relies on large collections of behavioral data that we expect to be affected by these similarities. In this context, we assume the representations produced by the model to be cognitively plausible as long as the similarities derived from them are able to successfully describe and capture these empirical data (Lindsay, 2021).

This does not necessarily imply that *how* the model arrives at these representations (the model architecture itself and the mechanisms it builds on) has to resemble the human visual processing system. However, divergences at this level do not by themselves necessarily affect the adequacy of the resulting representations: In a simplified analogy, assume a case where one system’s representation of a two-dimensional image is just a point reflection of the image. Another system that performs a 180° rotation around this point, or a specific sequence of two line reflections will obtain the same representations for the same input, even though they all apply different transformations to it (and arrive at different intermediate products along the way).

This being said, recent research in neuroscience has pointed out striking similarities between the model architectures and mechanisms at play in DCNNs on the one hand and

the human visual processing system on the other, while dissociations and divergences between the two have also been identified and acknowledged (for extensive reviews on this topic, see for example Cichy & Kaiser, 2019; Jacobs & Bates, 2019; Kriegeskorte, 2015; Lindsay, 2021; Majaj & Pelli, 2018; Serre, 2019).

Deep convolutional neural network models in cognitive science. The VGG models and similar DCNNs are extremely popular in the fields of computer vision/computer science, and have been successfully applied in a vast amount of studies there (Rawat & Wang, 2017; in July 2021, the seminal paper by Krizhevsky et al., 2012, had over 84,000 citations on Google Scholar). However, such studies typically pursue an engineering objective, investigating model performance in pre-defined computational tasks (such as classifying images, the original purpose of these models).

In cognitive science, vision-based representations obtained from computer-vision DCNNs have received considerably less attention as compared to computer science; however, some studies have started to investigate to what extent these models can predict human behavior and performance (see Battleday, Peterson, & Griffiths, 2021). Several studies found solid correlations between image similarities estimated by the models with corresponding human judgments of image similarity (Jozwik et al., 2017; Peterson et al., 2016, 2018; Roads & Love, 2021; Zhang, Isola, Efros, Shechtman, & Wang, 2018, see Study 2 for more detailed descriptions), especially for the deeper layers of the networks (Jozwik et al., 2017). Findings by Kubilius et al. (2016) corroborate the adequacy of DCNNs as models for the human visual system in their ability to predict human judgments on the shape of objects, even when shape and category membership of the employed stimuli was dissociated.

Other studies found that DCNNs mirror how humans organize images into different categories (Battleday, Peterson, & Griffiths, 2020; Singh, Peterson, Battleday, & Griffiths, 2020), especially when their original representations are further modified to reflect prevailing computational accounts of human category learning. Lake, Zaremba, Fergus,

and Gureckis (2015) found that the probabilities a DCNN assigned to the possible output labels for given input images predicted category typicality ratings for the same images.

Another line of studies has aimed at enriching more traditional representations (such as text-based distributional vectors) using vision-based representations from DCNNs, and found significant improvements of model performance on a series of different benchmarks such as semantic similarity judgments or categorization (De Deyne, Navarro, Collell, & Perfors, 2021; Lazaridou, Pham, & Baroni, 2015; Lazaridou, Marelli, & Baroni, 2017).

Going beyond these more explicit judgments, an extensive series of studies by Geirhos and colleagues (Geirhos, Jacobsen, et al., 2020; Geirhos, Meding, & Wichmann, 2020; Geirhos et al., 2021; Geirhos, Rubisch, et al., 2018; Geirhos, Temme, et al., 2018) has compared DCNN to human performance in more psychophysics-oriented paradigms. In these, they observed notable discrepancies in behavior between the two systems, such as a different emphasis on texture (DCNNs) versus shape (humans) (Geirhos, Rubisch, et al., 2018), or a much higher robustness against image distortions for humans (Geirhos, Jacobsen, et al., 2020; Geirhos, Temme, et al., 2018); however, more recent work has found that these gaps are progressively closing (Geirhos et al., 2021).

While the studies discussed so far investigated vision-based representations for individual images, more recent research also considered vision-based prototypes. In an analysis of rating data (De Groot, Koelewijn, Huettig, & Olivers, 2016), Petilli et al. (2021) found that only prototypical VGG similarities (obtained by averaging over different images labelled with the same word; see below for more detailed information) but not text-based distributional similarities predicted participants' *visual* similarity ratings for word-picture pairs, while the opposite was true for *linguistic* similarity ratings. In the same study, Petilli et al. (2021) found that these prototypical VGG similarities predict response times in a lexical priming task (Hutchison et al., 2013) even when controlling for the text-based distributional similarity between the prime and the target word.

Finally, some studies have employed vision-based representations derived from

DCNNs in order to investigate theoretical psychological questions, implicitly assuming the cognitive plausibility of these representations. Günther, Petilli, and Marelli (2020) used vision-based representations to demonstrate perception-based conceptual combination effects during the processing of compound words such as *swordfish*. In another study, Günther, Petilli, Vergallito, and Marelli (2020) found that vision-based representations can, with some success, be predicted from text-based distributional vectors, offering a mechanism via which non-experienced objects can be grounded in visual experience. Similarly, Lazaridou et al. (2017) demonstrated that visual intuitions for novel words learned from text alone can be predicted from their model combining textual with visual representations.

ViSpa: Vision Spaces

In the present article, we present the *ViSpa* (*Vision Spaces*) system, an adaptation of the pre-trained VGG-F model by Vedaldi and Lenc (2015) that, in addition to representations for individual images, also produces prototypical vision-based representations for concepts (see also Anderson, Bruni, Lopopolo, Poesio, & Baroni, 2015). Precursors of *ViSpa* were already employed in some of the studies outlined in the previous section (Günther, Petilli, & Marelli, 2020; Günther, Petilli, Vergallito, & Marelli, 2020; Petilli et al., 2021). In this system, an individual image representation serves to model the human mental representation when processing a specific visual stimulus. We assume that this representation influences behavior involving a given stimulus. The prototype representations are concept-level representations formed by averaging over many individual image representations from the same category. We further assume that these individual representations can be stored in memory (in line with the definition of mental representations by E. R. Smith, 1998), as exemplars of the category. The system thus resembles the classical exemplar-prototype distinction in concept research (Murphy, 2002; E. E. Smith & Medin, 1981; see also Singh et al., 2020 for a different approach to

implement this distinction in DCNNs). We assume that these concept-level representations not only play a role in visual processing, but also in other scenarios involving conceptual processing, such as reading written words or text (Anderson et al., 2015). In the studies presented in this article, these assumptions will be investigated empirically.

Representations for individual images: The IMG spaces

The pre-trained VGG-F model described earlier (Chatfield et al., 2014), which is available in the *MatConvNet* library for MATLAB (Vedaldi & Lenc, 2015), can be applied to any input image – importantly, also images outside its training set – in order to induce the corresponding distributional representation (see Figure 2). These are defined as the activation levels of the units of the network at a given layer for the input image. For a given layer, all these representations are high-dimensional numerical vectors sharing the same dimensions, and thus all populate the same vector space. Since the network has eight layers plus the input layer, we obtain a total of nine such image spaces, which will be abbreviated as *IMG spaces* from here on out.

For the present study, we obtained representations for images from 7,801 different image categories from the ImageNet database (Deng et al., 2009) (specifically, the ImageNet fall 2011 release), using the following procedure: First, we excluded all image categories with less than 100 images in ImageNet. Second, since ImageNet employs the categorical structure established in WordNet (Fellbaum, 1998; Miller, 1995), many image categories are annotated with more than one label, and the same label is sometimes assigned to different categories. Different methods to arrive at categories of images denoted by a unique word label have been proposed (De Deyne et al., 2021; Petilli et al., 2021; see also Beyer, Hénaff, Kolesnikov, Zhai, & Oord, 2020, who collected new labels for the ImageNet images altogether). Here, we used the method introduced in Günther, Petilli, and Marelli (2020): We set up a distinct category for each word label by selecting all images annotated with the respective label in the whole Imagenet database. When two or

more category labels included the exact same set of images, we always selected the category label with the highest word frequency in the 2.8-billion word corpus by Baroni et al. (2014) and excluded the others. Finally, we considered only image categories whose label was amongst the 300,000 most common words in the corpus by Baroni et al. (2014).

From each of these categories, we randomly sampled up to 200 different images. For each image, we extracted the activations from all layers (including layer 0 and layer 8) of the pre-trained VGG-F network included in the *MatConvNet* library for MATLAB (Vedaldi & Lenc, 2015). The activations from the early layers of the network (i.e., the convolutional layers conv1–conv5), which are originally represented as three-dimensional tensors, were vectorized. Due to this, all our vision-based representations are high-dimensional vectors, which will be referred to as layer 0 – layer 8 throughout this article. For a given layer, the activation values for any pair of images can then be compared using standard similarity metrics (see Figure 2) in order to obtain a model-based image similarity score (henceforth *IMG similarity*). Representations from two different layers cannot be compared to one another, since they populate different vector spaces. Following previous work (Petilli et al., 2021; Zhang et al., 2018), we employed the cosine similarity

$$\cos(a, b) = \frac{\sum A_i B_i}{\sqrt{\sum A_i^2} \sqrt{\sum B_i^2}}$$

as the metric for all similarities.¹

Vision-based prototypes: The PRO spaces

In the previous section, we have described how representations for individual images can be induced from the model. This section presents a method of obtaining concept-level representations that satisfy two criteria: (1) each concept is represented via a single representation (i.e., high-dimensional vector), in concordance with other concept-level

¹ The results reported in this paper change vary minimally when considering Euclidean distance or exponential cosine distance as alternative metrics, with a consistent slight advantage for the cosine similarity used here.

Figure 2. The IMG and PRO spaces, and the similarity measures derived from them.

representations such as language-based distributional vectors (Landauer & Dumais, 1997; Mandera et al., 2017); and (2) the concept representations populate the same vector space as the individual image representations, allowing for direct comparisons between them. The most straightforward of those are prototype representations, which play a central role in concept research (Murphy, 2002; Rosch & Mervis, 1975; Rosch, 1978). Such prototypes can straightforwardly be obtained as the centroid of a category, by averaging individual instances (i.e., images) of a category (see Anderson et al., 2015, who used such prototype representations to predict brain responses for written text; see also Battleday et al., 2020,

who refer to these averaged representations as their Classic Prototype model, and De Deyne et al., 2021.). These averaged representations are a straightforward implementation of the prototype-distance conceptualization by Posner (1970), where the prototype minimizes the distances (in case of the centroid, the squared Euclidean distance) between all category members and itself (see Battleday et al., 2020 for other options, for example based on the Mahalanobis distance). As pointed out by Děchtěrenko, Lukavský, and Štipl (2021), previous empirical research on visual prototypes suggests that prototype-distance representations perform better than other conceptualizations such as the attribute-frequency model (Neumann, 1974). While it has been proposed that exemplar models or hybrid models offer a better account for human concept representations than prototype models (Love, Medin, & Gureckis, 2004; McKinley & Nosofsky, 1995), this advantage has not been confirmed for high-dimensional models of visual representation as obtained from DCNNs (Battleday et al., 2020), which corroborates the usefulness of the relatively simpler prototype representations.

As described in the previous section, all individual images in the IMG space are associated to category labels. This allows us to calculate prototypical representations for categories as the centroid of the 100 to 200 individual vision-based vectors extracted for each category (Günther, Petilli, & Marelli, 2020; Günther, Petilli, Vergallito, & Marelli, 2020; Petilli et al., 2021; see also Anderson et al., 2015). This was done for all layers of the network, thus resulting in nine different prototype (PRO) spaces. Our PRO space thus contains vector representations for 7,801 different *words* serving as category labels – due to the averaging process, the representations in this model are no longer associated to individual images². These averaged vectors represent aggregated visual experience with the word referents, abstracted from the idiosyncrasies of individual exemplars, and are thus analogous to the prototype representations for concepts that feature prominently in the

² An alternative setup where the prototype was modelled as the closest available image representation to this aggregate representation hardly had any effect on the reported model fits

psychological literature (Rosch & Mervis, 1975; Rosch, 1978; E. E. Smith & Medin, 1981). Since these prototype representations are constructed from a specific set of labeled images, our PRO representations depend fully on the label applied to a given image and the composition of the consequent image sets. This will have different implications for different PRO representations: For example, the representation for a superordinate concept such as “pet” can be expected to be constructed from less visually consistent images (and thus more diverse IMG representations) than for the basic-level concept “dog” or the subordinate-level concept “poodle”. Potential implications of this will be discussed in the General Discussion.

As for images, we can compute the similarity between prototypes (*PRO similarity*) as the cosine proximity between their vector representations (Petilli et al., 2021; see Figure 2). In addition, since individual images and vision-based prototypes at a given layer of the network all populate the same vector space, we can compute the cosine similarity for any image-prototype pair (or image-label pair). In the specific case of a similarity between image vectors and the prototype vectors for their own label, we refer to this similarity as the model-based *typicality* of the image (see Figure 2). In the research on concepts, the notion of typicality plays a central role: As stated by Murphy (2002, p. 24): “as a general observation, one can say that whenever a task requires someone to relate an item to a category, the item’s typicality influences performance. This kind of result is extremely robust”. For example, one key finding is that typicality determines the speed and accuracy of category membership judgments (Rips, Shoben, & Smith, 1973; for a review of other results, see Murphy, 2002). Therefore, typicality is seen as one main organizing principle of categorization, which forms the basis for concepts (Rosch & Mervis, 1975; Rosch, 1978).

Outline for a large-scale evaluation

Even though the pioneering results reviewed earlier speak for DCNNs as promising candidates for modeling many aspects of human cognition, a thorough large-scale

evaluation in this respect is still pending. Such an evaluation is however urgently needed if one aims to employ these models to investigate theoretical questions (see Günther, Petilli, Vergallito, & Marelli, 2020; Günther, Petilli, & Marelli, 2020; Lazaridou et al., 2015, 2017). Rating studies have focused on pairs of individual images, and were based on relatively limited item sets (Jozwik et al., 2017; Peterson et al., 2016, 2018; Zhang et al., 2018; but see Roads & Love, 2021). The analysis of word-image similarity ratings conducted by Petilli et al. (2021) was also based on a rather small item set (De Groot et al., 2016), and substituted the actual image-word similarity with prototypical similarities between a word and the image label instead of testing a genuine prototype-image similarity (see Figure 2). A proper evaluation of vision-based prototype similarities against participants ratings is still missing, as is an evaluation of image-to-prototype similarities such as the typicality score (Figure 2). Conversely, while vision-based prototype representations have been found to predict response times in processing tasks such as lexical priming (Petilli et al., 2021), we do not currently have such evaluations for processing times of individual images.

In this article, we significantly extend upon this previous work by subjecting the VGG-F representations to a wide-breadth cognitive evaluation, employing a range of different large-scale studies to examine their capability to predict human behavior and performance. In Study 1, the similarities between prototype vectors (PRO similarities, Figure 2) will be evaluated against human ratings of visual similarity as well as semantic similarity for word pairs (that is, the word referents). Analogously, in Study 2, the similarities between image vectors (IMG similarities, Figure 2) will be evaluated against human ratings of image similarities. Study 3 ties together prototype and image vectors, by evaluating model-based typicality scores (Figure 2) against human judgments of the typicality of images given a category word. Finally, Study 4 and Study 5 go beyond ratings by evaluating image similarities against response times in a discrimination task (Study 4) and an image priming paradigm (Study 5). All studies received ethical approval by the Ethics Committee for Psychological Research at the University of Tübingen.

In all experiments, implemented using the *jsPsych* toolkit (de Leeuw, 2015; de Leeuw & Motz, 2016) and run in web browsers, we collect large datasets for thousands of items from hundreds of participants using online crowdsourcing (Keuleers & Balota, 2015), with at least 30 observations per item in every study. This large-scale approach comes with several advantages for our evaluation studies: a) the studies will be founded on a relevant portion of all the vision-based representations in our system, b) the large-scale approach will ensure an adequate power for our statistical analyses (Brysbaert & Stevens, 2018; Loken & Gelman, 2017), and c) it will mitigate increased noise in our data that might result from our online-testing format.

In all studies employing these similarities as continuous predictors for the respective dependent variables (i.e., all studies except for Study 3, see below), the data will be analyzed using Generalized Additive Models (GAMs Wood, 2006; Wood, Pya, & Säfken, 2016) to account for potential non-linear effects³. To exclude the impact of overly influential outliers and ensure the robustness of these GAMs, data points will be removed on the basis of a threshold of 3 SD standardized residual model errors, and the GAMs will then be re-fitted on the respective truncated datasets (Baayen, 2008). For each study, we run the same analysis for all eight layers of the VGG-F model to ensure the general robustness of the results (since the model produces a number of different representations, it is desirable that not only a subset of them produces reliable estimates. Furthermore, we can identify the best-performing representations produced by the model (compare Jozwik et al., 2017; Petilli et al., 2021). Besides providing a systematic evaluation of all representations produced by the model, this can also offer theoretical insights: As described earlier, deeper layers tend to capture more abstract gestalt-level representations (LeCun et al., 2015; Mahendran & Vedaldi, 2015; Zeiler & Fergus, 2014), and DCNNs

³ Throughout all studies reported in this article, the variance explained by these models was consistently higher than for linear mixed-effect models (Bates, Mächler, Bolker, & Walker, 2015), which however showed the same patterns of results with respect to the significance of predictors.

appears to resemble the human visual processing stream (e.g. Kriegeskorte, 2015). A comparison between the layers thus allows us to examine which of these representations best model the human vision-based representations involved in different behavioral tasks.

In addition, we considered as baselines the similarities between two types of non-learned, physical-level image representations (and prototypes derived from them) (Battleday et al., 2020): One is the input layer (*layer 0*) to the VGG-F network, that is, the RGB values for each pixel of the standardized images. The other is a Histogram of Oriented Gradients (HOG) representation, which describes the relative orientations of local features in the images (Lowe, 2004). Following Battleday et al. (2020), we used the default parameters of the OpenCV library (Bradski, 2000) to construct these HOG representations from downsampled 32×32 -pixel versions of the original images.

Where applicable, we also considered the rating values collected in the present study as additional predictors (see Table 1 and the individual descriptions of the studies below). As an additional predictor besides our main variables of interest (IMG similarity and PRO similarity), we consider language-based similarities, derived from the best-performing model in a large-scale evaluation by Baroni et al. (2014), which was constructed using the *cbow* variant of the *word2vec* algorithm (Mikolov, Chen, et al., 2013; Mikolov, Sutskever, Chen, Corrado, & Dean, 2013) applied on a ~ 2 billion word corpus of natural language (a concatenation of an English Wikipedia dump, the ukWaC corpus, Baroni, Bernardini, Ferraresi, & Zanchetta, 2009; and the British National Corpus, BNC Consortium, 2007). The *cbow* model is trained to predict a word given its immediately surrounding words, and thus tends to assign similar representations to mutually interchangeable words rather than systematically co-occurring words. Therefore, as other distributional models relying on these local word-by-word contexts, it tends to produce high similarities for words denoting entities that *are* similar (i.e., semantically similar words such as *spider* – *tarantula*) rather than entities that often occur together (i.e., associatively similar words such as *spider* – *web*) (see Jones, Kintsch, & Mewhort, 2006; Sahlgren, 2008, for more details). Controlling

for such language-based similarities is especially relevant as previous studies clearly demonstrate that language data already encode substantial amounts of perceptual information (Johns & Jones, 2012; Louwerse, 2011; also consider the *spider – tarantula* example): According to the Symbol Interdependency Hypothesis (Louwerse, 2011), a core purpose of language is to communicate our experiences with the world to others, and therefore the relations between words tend to be structurally similar to the relations between their referents. This is perhaps best illustrated by congenitally blind persons’ knowledge about the visual appearance of animals (Kim, Elli, & Bedny, 2019) or the colors of objects (Connolly, Gleitman, & Thompson-Schill, 2007), without any direct access to sensory information (see Lewis, Zettersten, & Lupyan, 2019). Therefore, we employ language-based models as a baseline in our studies, to investigate to which extent vision-based models add explanatory power over and above what can be induced from language.

Table 1 provides an overview over the results of all these studies, including split-half reliabilities which provide an upper bound for the variance that can be explained by the stimulus characteristics, the overall performance of the different layers in terms of variance explained (R^2) by the respective GAMs (with a graphical indication of the .95-confidence interval for the best performing layers’ R^2), as well as the significant predictors in these GAMs. The detailed results will be described in the individual study sections. The reported values are relatively specific for the GAM analysis presented here; we also provide bivariate Spearman rank correlations between the different similarity measures on the one hand and the different behavioral measures on the other hand in Table 2 to ensure comparability with other studies. Note that Table 2 is a very rough approximation, as it does not take into account collinearities between the different similarity measures; for example, correlations between PRO and IMG similarities for the same image pair range between $r_S = .15$ (layer 0) and $r_S = .92$ (layer 8). This can result in apparent discrepancies between the bivariate correlations in Table 2 and the F -values in Table 1, which provide a

clearer picture of the actual relative influence of the different similarity measures when they are considered simultaneously.

The studies reported in this article were not preregistered. All materials, datasets, and analysis scripts for the studies reported here are available at <https://osf.io/qvw9c> (Günther, Marelli, Tureski, & Petilli, 2022).

Study 1: Visual similarity ratings for word pairs

Previous studies employing the vision-based system presented here (Figure 2) mainly focused on word material and vision-based prototype representations (Günther, Petilli, & Marelli, 2020; Petilli et al., 2021), which we thus take as the starting point for our evaluation. These previous studies examined the effects of PRO similarities on processing tasks such as lexical decision, naming and timed sensibility judgments. While interesting and informative from a psychological perspective, this can be argued to take the second step before the obvious first one: Traditionally, model-based similarity scores are primarily evaluated against human similarity judgments (Baroni et al., 2014; Bruni, Tran, & Baroni, 2014; Jozwik et al., 2017; Zhang et al., 2018).

With this first study, we provide such an evaluation for PRO similarities, testing the extent to which they predict human ratings of visual similarity for the referents of word pairs (i.e, how similar do *knife* and *sword* look like?). In a follow-up analysis, we then examine whether PRO similarities indeed encode specifically visual similarities by comparing the results for this task with those for a *semantic* relatedness judgment task (the MEN dataset by Bruni et al., 2014).

Method

Participants. We originally recruited 601 participants via the *Prolific* crowdsourcing platform (Palan & Schitter, 2018; Peer, Brandimarte, Samat, & Acquisti, 2017), with the study only visible to participants who indicated that they were native

Table 1

Summary of the results from all studies including a layer comparison (Studies 1, 2, 4, and 5). The upper part of the table indicates the split-half reliabilities ρ as well as the (marginal) adjusted R^2 of the GAM (i.e., the variance explained by all fixed effect smooth terms in the model) predicting the respective dependent variables of the study for all eight layers of the VGG-F model, the RGB values in layer 0, and the HOG representations. R^2 -values outside of the .95-CI for the best-performing layer's R^2 are displayed in grey. The lower part indicates the F values for significant smooth terms in the GAM for the best-performing layer. A dash (–) indicates a non-significant smooth term that was not included in the final GAM; an empty cell indicates a smooth term that cannot sensibly be included in the respective analysis.

		Study 1		Study 2	Study 4		Study 5	
		word sim.		image sim.	discrim.		priming	
		(visual)	(semantic)	(visual)	RT	PC	RT	PC
split-half reliability ρ		.925	–	.941	.481	.378	.366	.253
baselines	RGB (VGG layer 0)	.600	.691	.433	.110	.072	.001	.002
(R^2 values)	HOG	.607	.692	.525	.106	.070	.005	.001
VGG layers	layer 1	.594	.695	.499	.112	.085	.002	.001
(R^2 values)	layer 2	.614	.697	.559	.113	.075	.006	.000
	layer 3	.630	.698	.620	.118	.080	.009	.003
	layer 4	.635	.695	.638	.119	.085	.006	.002
	layer 5	.641	.701	.616	.121	.076	.005	.001
	layer 6	.638	.699	.706	.125	.097	.013	.003
	layer 7	.639	.705	.690	.116	.081	.008	.003
	layer 8 (labels)	.626	.707	.690	.119	.081	.005	.003
smooth	cbow	236.3	216.4	64.3	–	–	–	–
terms	rating value				9.9	19.1	–	–
(F -values)	IMG sim.			139.4	15.8	15.0	13.0	20.7
	PRO sim.	118.3	53.4	26.6	2.5	–	5.6	–

Note. RT: response time; PC: percent correct

Table 2

Bivariate Spearman rank correlations between the different similarity measures and behavioral dependent variables in Studies 1, 2, 4, and 5.

		Study 1		Study 2		Study 4		Study 5	
		word sim.		image sim.		discrim.		priming	
		(visual)	(semantic)	(visual)		RT	PC	RT	PC
cbow		.72	.80	.57		.14	-.11	-.03	.03
rating value (Study 2)		.59	.57			.23	-.17	-.08	.05
IMG	HOG			.36		.10	-.09	-.05	.03
	layer 0 (RGB)			.05		.05	-.05	(.01)	(-.01)
	layer 1			.26		.12	-.06	-.03	(.02)
	layer 2			.37		.12	-.11	-.04	.03
	layer 3			.56		.18	-.13	-.09	.06
	layer 4			.58		.18	-.14	-.07	.04
	layer 5			.33		.14	-.10	-.03	(.02)
	layer 6			.73		.23	-.15	-.10	.06
	layer 7			.72		.19	-.14	-.08	.05
	layer 8			.74		.21	-.15	-.06	.06
PRO	HOG	.47	.39	.58		.14	-.11	-.06	.03
	layer 0 (RGB)	.41	.31	.45		.17	-.11	(-.01)	.03
	layer 1	.46	.41	.56		.17	-.12	-.03	.03
	layer 2	.52	.44	.63		.20	-.13	-.05	.06
	layer 3	.59	.52	.68		.20	-.13	-.06	.06
	layer 4	.62	.54	.70		.20	-.14	-.06	.05
	layer 5	.63	.56	.70		.20	-.13	-.05	.05
	layer 6	.63	.56	.71		.22	-.14	-.06	.06
	layer 7	.64	.59	.72		.20	-.14	-.06	.05
	layer 8	.62	.58	.71		.20	-.13	-.05	.06

Note. RT: response time; PC: percent correct. Non-significant correlations ($p \geq .05$) are displayed in brackets. Consider the intercorrelations between the similarity measures when interpreting these results.

English speakers with normal or corrected-to-normal vision. Of these, 121 were excluded because they were classified as not following the instructions (114 participants, see below) or because they reported a native language other than English (5 additional participants) in the actual study. Of the remaining 480 participants that entered the data analysis

($M_{Age} = 36.21$ years, $SD_{Age} = 13.25$ years), 292 identified as female, 181 as male, and 7 indicated another gender identity.

Material. The complete item material consisted of 3,000 word pairs including 2,228 unique nouns. The basis of the item material was collected from two prominent sources: (1) The MEN dataset (Bruni et al., 2014), which contains semantic relatedness ratings for 3,000 word pairs, and (2) the Semantic Priming Project (SPP; Hutchison et al., 2013), which contains 6,644 word pairs classified as semantically related or unrelated. From these resources, we selected all 2,164 pairs for which both words appeared in our prototype-level VIS space (1,167 from the MEN dataset, 997 from the SPP), and calculated their vision-based similarity as the cosine similarity of their respective prototype vectors. As would be expected given that vision-based representations were available for the 2,228 unique nouns in our item material (which implies that they were used as image labels), most of these words are very concrete: The subset of 1,822 words that were also available in the rating study by (Brysbaert, Warriner, & Kuperman, 2014) had a median concreteness rating of 4.59 on a scale from 1 (abstract) to 5 (concrete) (min = 1.45, $Q_1 = 4.07$, $Q_3 = 4.85$, max = 5.00), as compared to a median of 2.88 for their whole set of 40,000 widely-known English words (min = 1.04, $Q_1 = 2.12$, $Q_3 = 3.04$, max = 5.00).

For creating the item material, we aimed for a relatively uniform distribution of vision-based similarities, in order to evaluate the whole range of similarity measures produced by our model⁴. For this purpose, we focused on the similarities based on layer 7 of the VGG-F network, which was used in most previous studies (Günther, Petilli, Vergallito, & Marelli, 2020; Günther, Petilli, & Marelli, 2020; Lazaridou et al., 2017; Petilli et al., 2021). However, in the current material of 2,164 pairs sampled from the two different sources, the distribution of vision-based similarities was heavily left-skewed. To counteract this, we filled up the lower-density parts of the distribution (mostly on the low end of the

⁴ By randomly sampling 500,000 word pairs, we determined this range to fall almost exclusively within the [0.4, 1.0]-interval.

Figure 3. Distribution of vision-based PRO similarity (layer 7) for the word pairs in Study 1 by source (*MEN*: Bruni et al., 2014; *SPP*: Hutchison et al., 2013; *new*: new items added for the present study).

vision-based similarity scale, but also at the very high end; see Figure 3 by sampling 836 new word pairs from the entries of the PRO space. We only selected nouns with a frequency of at least 1,000 in a 2.8-billion word corpus (Baroni et al., 2014), created from concatenating the ukWaC web corpus (Baroni et al., 2009), a Wikipedia dump, and the British National Corpus (BNC Consortium, 2007). For the final list of word pairs, we calculated PRO similarities for all layers of the VGG-F network, as well as language-based similarities from a distributional semantic model built from the same corpus (Baroni et al., 2014; henceforth *cbow similarities*). A correlation matrix for these different similarity measures is provided in Figure 4.

In order to obtain human rating scores for the visual similarity of the word referents, we employed the best-worst scaling technique (Hollis, 2018; Hollis & Westbury, 2018). In

Figure 4. Correlations between the PRO similarities for the word pairs in Study 1, including vision-based similarities from all layers of the VGG-F network, the HOG similarities, as well as language-based similarities from the cbow model. As can be seen, starting from layers 3 and 4 the similarities are highly correlated. Also, the correlation with cbow similarities increases for the later layers of the network.

best-worst scaling, participants are presented with trials including tuples of N different items, and are instructed to select the item that scores the highest on a pre-defined dimension (best) and the item that scores the lowest (worst). Since each item thus generates information for the rank order of a number of pairwise comparisons (between the best item and all other items, and between the worst item and all other items), best-worst scaling is an efficient data-collection method. In addition, simulations and empirical studies have demonstrated that best-worst scaling provides better estimates of the latent

underlying variable than other rating techniques (Hollis, 2018; Hollis & Westbury, 2018).

In our present study, each trial included 4 word pairs for which participants had to indicate the pair whose referents were the most and least visually similar. Although recent studies have shown that tuples of 6 items are generally preferable (Hollis, 2020), we opted for the standard number of 4 items, since each item itself included a comparison between the two words in the pair, and the advantage for 6 items would be offset by the high cognitive load required to produce a judgment (see the discussion in Hollis, 2020). We collected 32 observations per item, which is sufficient for near-asymptote elimination of measurement errors (Hollis, 2018). With 4 word pairs per trial, our complete experimental material thus consisted of $3,000 * 32/4 = 24,000$ trials. We used software provided by Hollis (2018) to optimally assign the word pairs to these trials.

Each participant was presented with an individualized set of 50 trials, plus 5 catch trials which were the same for all participants and served as a compliance check for which we pre-determined the best and the worst item (for example [*baker* – *pretzel*, *tulip* – *grass*, *knife* – *sword*, *couch* – *bathtub*], where *knife* – *sword* was the most similar and *baker* – *pretzel* the least similar item; or [*nail* – *screw*, *gun* – *axe*, *submarine* – *football*, *willow* – *rainbow*], where *nail* – *screw* was the most similar and *willow* – *rainbow* the least similar item). It was taken care that, across all catch trials, the normative best and worst items were determined by their visual similarity, rather than other semantic/taxonomic or associative/thematic relations between the items. In addition, the experimental material included two practice trials which were the same for all participants, and where participants were provided with detailed feedback if their response deviated from the pre-determined correct answer (for example [*spear* – *arm*, *apple* – *peach*, *salt* – *stallion*, *pebble* – *diamond*], where *apple* – *peach* was the most similar and *salt* – *stallion* the least similar item).

Procedure. Participants were recruited with the study title “Which object pair looks the most and least similar?”. After giving informed consent and providing their demographic data, they were instructed that their task was to select out of 4 word pairs

presented to them the one for which the described object pair looks the most similar and the one for which it looks the least similar. Using the example [*bottle – pyramid, dog – wolf, lamp – soldier, car – street*], participants were explicitly instructed that the task was about judging *visual* similarity – *car – street* was presented as the least similar pair, even though *car* and *street* often occur together (i.e., are highly associated). Naturally, participants were not able to select the same item as the best and worst in any given trials.

Participants were instructed that they would be presented with two practice trials (for which they would receive feedback) before the actual study with 55 trials started. The catch trials were inserted at random positions between the actual experimental trials. Participants received £1.88 for their participation in the study, which was estimated to take 15 minutes to complete (the actual median completion time was 16.7 minutes).

Results

We classified as compliant only participants who selected the pre-determined correct response in the catch trials in at least (a) 60% for both the best option and the worst option, (b) 40% for the best option and 80% for the worst option, or (c) 80% for the best option and 40% for the worst option. Following these criteria, 114 participants were excluded from the analysis as non-compliant, and their individualized item sets were re-tested with another participant.

Using the software provided by Hollis (2018), the participant answers were then scaled via Value learning. Although other scaling techniques are available, Value learning is usually the most robust method (Hollis, 2018). We estimated the reliability of these rating values by scoring two independent random half-splits of our dataset; the Spearman-Brown split-half reliability is $\rho = .925$. The pattern of results reported here does not change for any other scoring algorithm provided by the software. Examples for the ten highest-scoring items and the ten lowest-scoring items are provided in Table 3.

As displayed in Table 1, the PRO similarities obtained from layer 5 of the VGG-F

Table 3

The items with the highest and lowest scores (i.e., for which the object pairs looked the most and least similar according to participant judgments) in Study 1 (word similarity ratings).

highest scores			lowest scores	
	item	value		item value
	feline – kitty	.821		inn – jellyfish .204
	coke – pepsi	.819		salt – teacher .206
	cream – milk	.817		uphill – gravy .212
	tangerine – orange	.816		giraffe – jelly .212
	chimpanzee – ape	.814		flamingo – office .212
	aircraft – airplane	.813		airplane – communist .214
	dog – puppy	.811		socket – cliff .214
	taxi – cab	.810		sandbank – gorilla .214
	crocodile – alligator	.810		coyote – spectrometer .215
	dog – husky	.809		chairlift – pastry .216

model best predicted the data (with layers 6 and 7 showing only slightly lower performance). In a GAM (estimated using the *mgcv* R package; Wood, 2015) predicting the rating values from a PRO similarity smooth term, this smooth term is significant ($F = 297.2$, $edf = 7.2$, $rdof = 8.2$, $p < .001$, with $R^2_{adj.} = .450$). This GAM neither contains random intercepts for word pairs, since this would require the GAM to estimate more parameters than data points available; nor it does contain random effects for the first and the second word of a pair, since this prevents a reliable parameter estimation for some of the GAMs reported here.

Since participants in our study provided ratings for word pairs, a major control variable of interest is the language-based similarity of these latter (Petilli et al., 2021), that is, the similarity of the word meanings in language usage. In line with previous studies, we

employed a distributional semantic model to measure these language-based similarities (Günther, Dudschig, & Kaup, 2016; Mandera et al., 2017; Petilli et al., 2021); as described above, we employed the best-performing model by (Baroni et al., 2014), which was constructed using *word2vec*’s cbow algorithm (Mikolov, Chen, et al., 2013).

Within our set of 3,000 word pairs, the rank correlation indicates a moderate relationship between these cbow similarities and our PRO similarities ($\rho = .57, p < .001$). Due to this correlation, we examined whether vision-based similarities accounted for variance in the rating values over and above the variance already explained by the language-based cbow similarities. To this end, we estimated a GAM predicting rating values with smooth terms of (layer 5) PRO similarity as well as cbow similarity, and compared it to a GAM only including a smooth term for cbow similarity using a likelihood-ratio test. The additional effect of PRO similarity significantly improved the GAM predictions ($\Delta Deviance = 6.25, \Delta df = 5.0, p < .001$), with an improvement in adjusted R^2 from $R^2_{adj.} = .564$ to $R^2_{adj.} = .641$ ($\Delta R^2_{adj.} = .077$). This also holds the other way around: A GAM that includes both smooth terms performed significantly better than one only including a PRO similarity smooth term ($\Delta Deviance = 16.33, \Delta df = 4.2, p < .001$), with an improvement in adjusted R^2 from $R^2_{adj.} = .440$ to $R^2_{adj.} = .641$ ($\Delta R^2_{adj.} = .201$).

We calculated .95-confidence intervals for the final model’s R^2 with a bootstrapping technique (Efron & Tibshirani, 1994): For this, we sampled (with replacement) 1,000 bootstrapping datasets of the same size as the original dataset from this original set. In these bootstrapping datasets, some items will appear multiple times while others will be removed. We estimate a GAM with the same parameter structure as the final model for each of these bootstrapping datasets, extracted their R^2 -values, and calculated the .025-quantile and .975-quantile of these 1,000 values. For the final model reported here, we obtained $CI_{.95}^{R^2} = [.623, .662]$.

In the resulting GAM, both smooth terms indicate a positive relation between the predictors and the rating values, as displayed in Figure 5 (created using the *mgcViz* R

Study 1: Visual similarity ratings (words)

Figure 5. The partial effects of PRO similarity (left panel) and cbow similarity (right panel) on visual similarity rating values for word pairs in Study 1, with .95-confidence intervals.

package; Fasiolo, Nedellec, Goude, & Wood, 2018).

Follow-up analysis: Semantic relatedness judgments. To put the findings of the previous section into perspective, we additionally analyzed the semantic relatedness scores in the MEN dataset (obtained from expert annotators; Bruni et al., 2014) for the 1,167 word pairs shared between our dataset and the MEN dataset. Despite these clear task differences, the rank correlation between these two ratings is very high at $\rho = .79$ ($p < .001$).

In this analysis, layer 8 of the VGG-F model performed best (closely followed by layer 7 and layer 5); see Table 1. We performed the same model comparison as reported in the previous section for the MEN ratings; also here, adding PRO similarities to the GAM already containing cbow similarities significantly improves model predictions, as indicated by a likelihood-ratio test ($\Delta Deviance = 2.13$, $\Delta df = 2.0$, $p < .001$). Analogous results are obtained by adding cbow similarities to a GAM already containing PRO similarities ($\Delta Deviance = 25.87$, $\Delta df = 2.7$, $p < .001$). In the resulting model, both effects indicate a positive relation between the predictors and MEN ratings, as displayed in Figure 6. Across

Figure 6. The partial effects of PRO similarity (left panel) and cbow similarity (right panel) on MEN semantic relatedness ratings, with .95-confidence intervals.

tasks, the vision-based model is thus able to explain portions of the variance in the ratings that the language-based model does not explain. However, for semantic relatedness judgments (Bruni et al., 2014), the improvement in explained variance by adding visual similarity (from $R_{adj.}^2 = .681$ to $R_{adj.}^2 = .709$; $\Delta R_{adj.}^2 = .028$) is not as large as for ratings when considering the same subset of 1,167 word pairs in our dataset ($\Delta R_{adj.}^2 = .069$). For the final model, we obtain $CI_{.95}^{R^2} = [.683, .736]$.

Discussion

In Study 1, we observed that PRO similarities predicted human similarity judgments for word pairs, with substantial levels of variance explained (see Table 1). This was the case for both visual (Figure 5) as well as semantic similarity judgments (Figure 6) – which are highly associated, but still distinguishable. Importantly, in both cases the vision-based PRO similarities predicted variance in these judgments over and above language-based cbow similarities, even though the tasks were performed solely on word material. This suggests that, when evaluating linguistic materials, perceptual information concerning the referent is taken into account, information that cannot be reduced to aspects redundantly

encoded in language usage. In line with previous studies (Günther, Petilli, & Marelli, 2020; Petilli et al., 2021), these results speak for the impact of visual information on language processing. However, in interpreting this pattern we also need to consider that the vision-based and language-based model differ not only in the modality of the training data, but also in their training objective and modelling architecture (categorizing images based on their physical features versus predicting words from their context words) – for further considerations in these respect see the section “Vision-based and language-based similarities” in the General Discussion.

At first glance, the high correlation between visual and semantic similarity judgments can appear suspicious – did participants in our visual similarity rating task actually follow the instructions? Notwithstanding the association between visual and semantic judgments, it seems likely that they did: with very explicit instructions and practice trials, as well as the filtering of data via catch trials, our study was clearly optimized to capture this specifically visual type of similarity. Indeed, for concrete nouns serving as labels for images, a high correlation between visual and semantic similarity can be expected: On the one hand, visual similarity is one main factor determining semantic similarity (since it adds to the number of shared features between two objects), and on the other hand, visual similarity usually is a good indicator for functional and conceptual similarity (that is, similar concepts such as *coke* and *pepsi* tend to be more visually similar, Deselaers & Ferrari, 2011; see also the other examples in Table 3).

However, more direct evidence for a dissociation is provided by the data itself: In our analyses, we find that – in terms of variance explained – the contribution of vision-based PRO similarities to predicting visual similarity ratings is considerably larger than for semantic similarity ratings. If participants were performing the exact same type of judgment in both tasks, such a difference in the relative contributions of the predictors would be difficult to explain. Rather, this difference indicates that different processes and mental operations are involved for the two types of ratings. The specific pattern of this

systematic difference thus suggests that PRO similarities encode specifically visual information that is not captured by language-based models.

Finally, anticipating the other studies described below, we also need to note that the proportion of participants excluded because they did not pass our compliance check was way higher in Study 1 (about 20%) than in the other rating studies (about 8% in Study 2 and 4% in Study 3). This may be due to the task at hand, because participants had to make a visual similarity rating for words and were thus required to engage in an additional mental operation – visualizing objects – rather than just performing judgments on the very stimuli presented to them. In addition, the literature shows that word pair ratings are susceptible to be influenced by many factors other than the intended target dimension (e.g. Honke & Kurtz, 2019), which may have lead some participants to perform judgments other than the intended visual similarity ratings.

Study 2: Visual similarity ratings for image pairs

Having examined the extent to which vision-based PRO similarities predict visual similarity ratings for word pairs, we now turn to visual similarity ratings for pairs of individual images. Such ratings have been examined against similarities between DCNN representations before, using among others (variations of) the VGG model employed here: In a study by Zhang et al. (2018), participants had to choose which of two distorted images was more similar to the original. In the present study, we instead use pairs of two different naturalistic images. Similar data has been collected in several other studies (Jozwik et al., 2017; Peterson et al., 2016, 2018), but we extend these by testing a much larger set of stimuli without restricting the domains from which these images are collected. Furthermore, we select images from the very same categories for which the words in Study 1 were used as labels. This allows us to compare the two types of ratings, and to examine the effect of PRO similarities on image ratings in addition to that of IMG similarities.

Another large dataset has been examined by Roads and Love (2021), who collected similarity judgments for a stimulus set of 50,000 ImageNet images and compared them against DCNN predictions. In each trial, they presented one target image to participants, and asked them to choose the two most similar images out of nine candidates. In our Study 2, we employ a conceptually similar technique to Roads and Love, namely best-worst scaling, as a state-of-the-art technique with desirable psychometric properties to collect rating values (Hollis, 2018; Hollis & Westbury, 2018). This technique allows us to estimate a reliable similarity score for each image pair, based on multiple observations per item.

Method

Participants. In this study, we originally recruited 519 participants via *Prolific*, again only presenting the study to participants who indicated that they were native English speakers with normal or corrected-to-normal vision, and who did not participate in Study 1. Of these, 39 were excluded because of incomplete data (one participant), because they were classified as not following the instructions (19 participants, applying the same criteria based on the catch trials as in Study 1) or because they reported a native language other than English (19 additional participants) in the actual study. Of the remaining 480 participants that entered the data analysis ($M_{Age} = 35.57$ years, $SD_{Age} = 13.22$ years), 288 identified as female, 186 as male, and 6 indicated another gender identity.

Material. We aimed at establishing a reasonable comparability between this study on image pairs and Study 1 on word pairs. To this end, our image pairs were modeled after the word pairs in Study 1, by substituting each word with an individual image for which this word served as an ImageNet label, thus keeping the pairs intact (as an example, since *tiger* – *lion* is an item in Study 1, one item in Study 2 is a pair of an image labeled TIGER and image labeled LION).

In the image selection procedure, we only considered images with a width-to-height ratio of 4:3 (the most common ratio – about 25% – in our set of images), and an original

resolution of at least 299×224 pixels. All image candidates with a higher resolution were then resized to a resolution of 299×224 pixels.⁵ This was done to standardize the images and prevent that “technical” differences between images such as size, width-to-height-ratio, or image resolution affect similarity judgments. In addition, this standardization is especially relevant for the studies on on-line processing reported later in this manuscript, where parts of the same item material will be employed.

In order to further increase the comparability of this study to Study 1 where we employed the prototype vectors of the PRO model, we then selected for each label the image that was closest (in terms of cosine similarity) to its respective prototype vector. The large majority of selected images were very close to the prototype vector in comparison to all 100 – 200 images used to create these prototypes (median position = 3 [$Q_1 = 2$; $Q_3 = 9$]; mean cosine similarity = .86, $SD = .05$). For the final list of image pairs, we calculated IMG similarities for all layers of the VGG-F network; a correlation matrix for these different similarity measures is provided in Figure 7.

We employed the same best-worst scaling technique (Hollis, 2018; Hollis & Westbury, 2018) as described in Study 1 in order to obtain rating scores for the items, except that the words were replaced with images. Again, each participant was presented with an individualized subset of 50 out of the 24,000 different trial constellations. We also employed a new set of the two practice trials and a new set of five catch trials to make the instructions and compliance control as clear as possible. The images of a pair were always presented side-by-side, with the different pairs being arranged below one another (see Figure 8).

⁵ Note that as a preprocessing step of the VGG-F algorithm all the images are scaled to the fixed size of 224×224 in order to standardize the input layer (Chatfield et al., 2014); we are thus not destroying information in the images that could be picked up by the VGG-F algorithm and affect image similarity scores

Figure 7. Correlations between the similarities for the image pairs in Study 1, including vision-based IMG similarities from all layers of the VGG-F network. As can be seen, there are high correlations between layer 3 and 4, as well as between the latest layers (6-8) of the network.

Procedure. Other than a change in the instruction (participants were instructed to judge which image pair looked the most and least similar), the procedure of Study 2 was identical to Study 1. Participants were paid £1.88 for the participation. The median completion time for all 57 trials (2 practice trials + 50 experimental trials + 5 catch trials) in this study was 19.8 minutes.

Which look the most similar?

Figure 8. Example for an experimental trial in Study 2 (similarity ratings for image pairs).

Results

We again scored the best-worst ratings using Value learning (Hollis, 2018), although again very little difference can be observed between the different scaling techniques (both in terms of their pairwise correlations as well as the pattern of effects reported here). The Spearman-Brown split-half reliability (see Study 1) of these values is $\rho = .941$. Examples

Figure 9. Examples for the highest- and lowest-scoring items in Study 2 (similarity ratings for image pairs).

for the two highest- and lowest-scoring items are provided in Figure 9.

In Study 2, layer 6 of the VGG-F network was identified as the best-performing

layer (see Table 1). In a GAM predicting the rating values from a smooth term for IMG similarity, this smooth term was significant ($F = 704.1$, $edf = 6.1$, $rdof = 7.3$, $p < .001$, with $R^2_{adj.} = .633$).

Although participants in Study 2 were presented only with images, each of these images is associated to a word label in our data (these labels were in fact used to create the material for Study 2 from the material of Study 1). This allows us to calculate two additional similarity scores for each image pair, beyond the IMG similarities between the images themselves: (a) The similarity between the visual prototypes for their image labels (PRO similarity), and (b) the language-based similarity between their image labels (cbow similarity). Moreover, these image labels allows us to compare the image ratings to the word ratings in Study 1 – where we observe a moderate rank correlation ($\rho = .59$, $p < .001$) – and to the MEN semantic similarity ratings (Bruni et al., 2014) – again with a moderate correlation in the shared set of 1,167 items ($\rho = .57$, $p < .001$).

In layer 6 of the VGG-F model, the correlation between the IMG similarities (for the actual image pairs) and the PRO similarities (for their respective visual prototypes) is very high at $\rho = .76$ ($p < .001$), as expected due to the fact that the individual images in our study were selected to be as close to the prototype as possible. The correlation between the IMG similarities and the language-based similarities for their labels is moderate ($\rho = .46$, $p < .001$).

In a next step, we examined the predictive power of all three types of similarity (IMG, PRO, cbow) on the visual similarity rating values (for this model containing all three parameters, we obtain $R^2_{adj.} = .706$, $CI^2_{.95} = [.689, .727]$). To this purpose, we estimated a GAM predicting ratings from all three types of similarity, and systematically tested whether the removal of parameters for these similarities significantly reduced model fit, as indicated by likelihood-ratio tests. This was the case for all three similarities:

- IMG similarity: $\Delta Deviance = 8.98$, $\Delta df = 6.1$, $p < .001$, with a change in variance explained from $R^2_{adj.} = .706$ to $R^2_{adj.} = .609$; $\Delta R^2_{adj.} = .097$

- PRO similarity: $\Delta Deviance = 1.29$, $\Delta df = 4.4$, $p < .001$, with a change in variance explained from $R_{adj.}^2 = .706$ to $R_{adj.}^2 = .692$; $\Delta R_{adj.}^2 = .014$
- language-based similarity: $\Delta Deviance = 3.14$, $\Delta df = 4.1$, $p < .001$, with a change in variance explained from $R_{adj.}^2 = .706$ to $R_{adj.}^2 = .672$; $\Delta R_{adj.}^2 = .034$

Thus, all types of similarity explain significant parts of the variance in the human ratings on their own, with IMG similarity contributing the most to variance explained, followed by cbow similarity and PRO similarity. In the model containing all three types of similarity, the relations between all three predictors and the rating values are positive, as displayed in Figure 10; Fasiolo et al., 2018 (the falloff for very high values of language-based similarity is based only on very few items, and should therefore not be over-interpreted, as is also indicated by the wide confidence interval band in that area).

Discussion

Among the variables tested in Study 2, IMG similarities were shown to be the strongest predictor for image-pair ratings, explaining a high amount of variance in these ratings. This corroborates previous studies which also identified image similarity scores derived from DCNNs as a very effective predictor of human similarity judgments (Jozwik et al., 2017; Zhang et al., 2018). Notably, the DCNN-based similarity measures clearly outperform the physical-level baselines (similarities between RGB- and HOG-representations; see Table 1), suggesting that human ratings are not based on purely physical similarities between the presented images.

Over and above these effects of IMG similarity, we also identified effects of PRO similarities and cbow similarities on visual similarity ratings for image pairs. This indicates that these similarity judgments do not involve a purely perceptual comparison of the actual visual stimuli (the images), but also a certain degree of higher-level analyses, including a conceptual classification of the images and a comparison of the respective semantic representations formed through aggregated experience. In other words, when comparing

Study 2: Visual similarity ratings (images)

Figure 10. The partial effects of vision-based similarity (upper panels: IMG similarity on the left, PRO similarity on the right) and language-based cbow similarity (lower panel) on visual similarity rating values for image pairs in Study 2, with .95-confidence intervals.

specific images of a tiger and a lion, human judges are also, to a certain degree, biased by how similar tigers and lions look and are in general. Importantly, using the IMG, PRO, and cbow representations, we are able to capture and quantify this diverse influences, evidencing the scientific usefulness of these quantitative models.

Notably, we observe these effects using only image stimuli, without providing any additional information about their categories or corresponding labels. These categories and labels are to some degree arbitrary and ambiguous: The words *breakfast*, *bread*, *eggs*,

melon, *food*, etc., are all sensible labels for the image in the bottom left of Figure 8. Importantly, participants in Study 2 never *see* the ImageNet label (Deng et al., 2009) on which our model depends, only the image itself. The fact that we still observe effects of variables tied to these ImageNet labels (PRO similarities and cbow similarities) for similarity judgments on these images indicates that the participants have to follow the original image annotators in their categorization and labelling process at least to some degree. This highlights that, although the ImageNet categories and their labels – and by consequence the models employed here – are somewhat arbitrary and ambiguous, the system nevertheless captures relevant categorical information, lending further to the psychological plausibility and adequacy of the model representations presented in Figure 2.

Study 3: Typicality ratings for image-word pairs

In the previous experiments, we have investigated the effects of IMG similarities (between two individual images) and PRO similarities (between two vision-based prototypes) on human similarity judgments. While the results of these studies (and especially the PRO similarity effects on image similarity ratings in Study 2) have begun to shed some light on the interplay between the IMG space and the PRO space, we have not yet evaluated the connection between the IMG space and PRO space on the modeling side.

More specifically, a central property of the system presented in Figure 1 is that, as a result of the averaging process, the vision-based prototype representations populate the very same vector space as the individual image representations. This allows us to not only to compute similarities between two prototypes denoted by word labels (PRO similarities) or two images (IMG similarities), but also between a prototype and an individual image. In principle, these metrics capture the degree of visual typicality of a given image with respect to a certain category; to our knowledge, these similarities have not yet been subject to any empirical investigation. While Lake et al. (2015) also investigated category typicality ratings, they predicted these from the probabilities the DCNN assigns to the

different output labels when fed with a given input image. These are naturally limited to the (1,000) different image labels on which the network is initially trained.

In Study 3, we aim to fill this gap in a further large-scale evaluation. For this purpose, we focus on typicality scores as an important subset of image-word similarities (with typicality playing a central role in the research on concept representations, as discussed earlier Rosch & Mervis, 1975; Rosch, 1978), and investigate the extent to which these model-based scores predict human typicality judgments.

Method

Participants. We originally recruited 940 participants, again via *Prolific*, who indicated that they were native English speakers with normal or corrected-to-normal vision, and who did not participate in any of the previous studies. Of these, 38 were excluded because of incomplete data (seven participants), because they were classified as not following the instructions (31 participants, applying the same criteria based on the catch trials as in Study 1 and 2). Of the remaining 902 participants that entered the data analysis ($M_{Age} = 33.67$ years, $SD_{Age} = 12.60$ years), 544 identified as female, 340 as male, and 18 indicated another gender identity.

Material and Procedure. Our item material for Study 3 consisted of the most frequent 1,500 nouns out of the 2,228 different nouns included in Study 2 (with frequencies based on the corpus by (Baroni et al., 2014)). For each of these words, we selected five different associated images, out of the set of the 100-to-200 images that were used to estimate the visual prototype vector. These images were selected based on the similarity between their respective vectors and the prototype vector: For each word, we selected the images that were the most similar (100-percentile) and the least similar (0-percentile) to the prototype, as well images at the remaining quartiles (25-percentile,

50-percentile/median, and 75-percentile).⁶ To ensure that the same percentiles could be used for each word, we did not apply any restrictions on image size, ratio, or resolution.

Similar to Study 1 and 2, we used a best-worst-scaling technique to obtain typicality ratings for the images. Participants were presented with the word on top of the page, and the five images listed below it in random order (see Figure 11). Participants were asked to “select the image that shows the least typical and the most typical exemplar of the word (i.e., the worst and the best representation of the word)”. We provided further clarification by stating that participants could think of the task as choosing “for which image would the word presented to you be the worst and the best description”, or “if you hear or read the word, which image is closest to the first image that comes to your mind?”.

The material was arranged into 30 unique lists, each consisting of 50 word-image sets that served as experimental trials. Each of these lists was presented to 30 different participants, resulting in 30 data points per item⁷. Similar to Study 1 and 2, each participant was additionally presented with two practice trials (the only trials for which they received feedback), and five catch trials with reasonably unambiguous answers. The 50 actual experimental trials and the five catch trials were presented in randomized order after the practice trials. Participants received £1.88 for their participation in the study, which was estimated to take 15 minutes to complete (the actual median completion time was 15.8 minutes).

Results

The best-worst ratings for each of the 1,500 word-image sets were scored individually via Value learning (Hollis, 2018), resulting in 7,500 typicality rating values. The Spearman-Brown split-half reliability (see Study 1) of these values is $\rho = .946$. Examples

⁶ These similarities were computed based on layer 7 of the VGG-F model, but the rank order of images was the same for layer 6 which generally performs better in the experiments presented here.

⁷ Two lists were accidentally presented to 31 participants

mosaic

least typical		most typical
☒		☐
☐		☐
☐		☒
☐		☐
☐		☐

Figure 11. Example for an experimental trial in Study 3 (typicality ratings for word-image pairs).

Figure 12. Examples for extremely high-scoring and very low-scoring items in Study 3 (typicality ratings for word-image pairs). The numbers over the images indicate rating values (the bold-face numbers representing extremely high or low values in the data).

for the two words that included the highest- and lowest-scoring items (i.e., the most typical and the least typical image) are provided in Figure 12.

In Study 3, all images were selected as specific percentiles (000, 025, 050, 075, 100) with respect to their similarity to the prototype vector. In our analysis, we therefore operationalized model-based typicality (the similarity between an image and the respective prototype) as a categorical predictor. More specifically, we estimated a Linear

Mixed-Effects Model (LMEM Bates et al., 2015; Kuznetsova, Brockhoff, & Christensen, 2017) predicting participants' rating values from this model-based typicality (with the 050-percentile set as the reference level for this categorical predictor) as well as random intercepts for the words.

The rating values for the different model-based typicality percentiles are displayed in Figure 13. As can be seen, higher model-based typicalities come with higher participant ratings ($F(4, 7495) = 333.9$, $p < .001$ for the typicality predictor). This positive relation holds across all conditions: Ratings values for the 000-percentile ($b = -0.10$, $t = -19.65$, $p < .001$) and the 025-percentile ($b = -0.02$, $t = -4.32$, $p < .001$) are smaller than the ones for the 050-percentile, while ratings values for the 075-percentile ($b = 0.03$, $t = 5.51$, $p < .001$) and the 100-percentile ($b = 0.07$, $t = 15.31$, $p < .001$) are larger. By changing the reference level accordingly, we furthermore find differences between the 025-percentile and the 000-percentile ($b = -0.07$, $t = -15.34$, $p < .001$), as well as between the 075-percentile and the 100-percentile ($b = 0.05$, $t = 9.80$, $p < .001$).

Discussion

In Study 3 we find that, when presented with several images belonging to the same category, human judges prefer to choose images with high model-based typicality scores as typical visual representations for a given word (and images with low scores as non-typical representations). Indeed, we find graded differences between all the levels of model-based typicality scores employed in this study: The higher this score is, the higher the typicality ratings are.

These results play a central part in establishing the validity and cognitive plausibility of the presented vision-based representations (Figure 1) as a whole, by directly tying together the IMG space and the PRO space derived from it into a cohesive, internally consistent system. This is especially important since, unlike the IMG space that is created by feeding images to a pre-trained model (Vedaldi & Lenc, 2015), the creation of the PRO

Figure 13. Violin plots for the distributions of typicality rating values in Study 3 by model-based typicality scores; red lines represent mean values, the (very narrow) red-shaded areas around those mean values .95-confidence intervals.

space involves some and modelling choices and free parameters and is thus to some degree arbitrary: The PRO space is created from 100 to 200 images per category, which are necessarily a non-exhaustive subset of all possible images. These categories can sometimes comprise several ImageNet categories sharing the same label, which also introduces some modelling choices, as described in the section on the creation of the ViSpa system above. Furthermore, these categories are instantiated via the image labels, which are themselves severely under-determined (for example, the uppermost image in Figure 11 could also be labeled as *woman*, *Greek*, *Roman*, or *art*).

Still, as demonstrated in Study 3, the typicality scores derived from these under-specified prototype representations reliably predict human typicality judgments, even though the participants were neither aware of these representations nor did they know

about the 100-to-200 individual images from which they were created. Alongside the initial results from Study 1 and Study 2 (where we observe PRO similarity effects on word-word and image-image similarity ratings), these results present a strong piece of evidence for the validity and plausibility of the PRO space, in particular, and the whole vision-based system presented here (Figure 2), in general.

Study 4: Discrimination task with image pairs

In the previous experiments, we evaluated the model-based similarity scores against explicit human judgments. This is the most common approach to model evaluation, especially in computational fields where human judgments are usually treated as gold standards (Baroni et al., 2014; Bruni et al., 2014; Hill, Reichart, & Korhonen, 2015; Jozwik et al., 2017; Zhang et al., 2018).

However, for an evaluation from a cognitive perspective, we should not solely rely on these off-line judgments, but also investigate the predictive power of the vision-based model for on-line processing measures. This approach has played a vital role in establishing the cognitive plausibility of language-based distributional semantic models (Günther et al., 2016; Jones et al., 2006; Mandera et al., 2017). Previous studies have also taken this approach to investigate vision-based models (Günther, Petilli, & Marelli, 2020; Petilli et al., 2021), but so far only considered the effects of PRO similarities in word-based tasks such as lexical decision.

In Study 4, we therefore aim at investigating the effect of vision-based similarities on processing times in an image-based task. More specifically, we employed a well-established perceptual paradigm: the discrimination task, in which participants are presented with two stimuli and have to decide whether the two stimuli are identical (Ehrenstein & Ehrenstein, 1999; Pelli & Farell, 1995). In this paradigm, correct “different” responses are typically fewer and slower the more similar the stimuli are, since it is more difficult to discriminate between them, and hence to appreciate them as different from one

another (Ehrenstein & Ehrenstein, 1999). If the vision-based similarity scores derived from our model (and especially IMG similarities) predict response times and accuracy rates in a discrimination task employing images as the stimulus material, this will provide further evidence for the cognitive plausibility of the model-induced representations.

Method

Participants. The data for this study was collected from 750 participants, recruited via *Prolific*, with the study only visible to native, monolingual speakers of English with normal or corrected-to-normal vision. Of these ($M_{Age} = 35.59$ years, $SD_{Age} = 13.31$ years), 497 self-identified as female, 244 as male, and nine indicated another gender identity. Participants were explicitly and repeatedly advised not to participate in this study if the rapid succession of visual stimuli might have negative consequences for their well-being or health (for example, if they had a history of epilepsy).

Data from 26 additional participants were excluded due to high error rates (see Results section), and a further 20 participants were excluded due to technical errors during data collection (incomplete data files).

Material. The employed experimental material was the exact same 3,000 image pairs (including 2,228 unique images) already used in Study 2. These were arranged into 25 different experimental lists of 120 image pairs each, so that repetitions of individual images within each list were kept at a minimum (43 cases where an image appeared three times, 381 cases where an image appeared 2 times, and 5109 cases where an image appeared only once per list). In addition to these experimental items, each list contained 120 identical pairs in which the same image appeared twice. These were sampled from the (about) 2,108 images that were not already included in a list. Within each list, the experimental items and identical pairs were randomly arranged into four blocks, so that each block contained 30 experimental items and 30 identical pairs. Each list additionally included 32 practice trials, which were randomly sampled with the restriction that none of their images already

appeared in the experimental items or identical pairs. All images were presented in their original size (299×224 pixels) in the center of the screen, in front of a light-gray background, and the experiment was presented in full-screen mode.

We additionally created four sets of visual masks with different spatial resolutions, with four different masks per set (following the study by Aly, Wansard, Segovia, Yonelinas, & Bastin, 2014, who also employed a discrimination task with naturalistic images). These masks are included to ensure that responses to the target were not influenced by the after-image of the first stimulus, with dynamic masks being especially effective for this purpose. The size of all masks was set to the same size as the images (299×224 pixels). The four masks of each set were created by arranging randomly-colored squares with a size of either 1×1 pixels, 2×2 pixels, 4×4 pixels, or 8×8 pixels (see Figure 16).

Procedure. After providing informed consent, participants were instructed that they would be presented with pairs of images in short succession, and that their task was to decide whether the two images were identical or not by pressing the *F* or *J* key on their keyboard. The key-to-response assignment was randomized for each participant.

All presented images were pre-loaded at the start of an experimental session, in order to avoid lags and timing discrepancies that loading them during the experiment would cause. The experiment then started with the 32 practice trials in random order. The trial procedure is illustrated in Figure 14. Each trial started with a fixation cross for 500 ms. Then, the first image appeared on the screen for 100 ms, followed by four masks of 30 ms each (this constitutes a dynamic noise mask, consisting of one randomly-sampled mask from each of the four visual-resolution sets in random order, Aly et al., 2014). A visual mask – especially a dynamic mask – decreases the potential for interference or facilitation caused by visual after-images, thus providing more experimental control over the actual duration in which the visual information is available to participants. In addition, the reduced availability of visual information to participants can be expected to make the (otherwise quite easy) task more difficult. The specific timings for the first image and mask

Figure 14. The trial procedure in Study 5.

duration were the most effective in producing the standard similarity effects in discrimination tasks, as identified in a pilot study employing the same procedure as the study described here with unambiguously very dissimilar, medium similar, or highly similar image pairs (according to both the ratings from Study 2 and the IMG model). After the mask, the second image (the target) was presented until participants made their response. If this response was incorrect, the feedback “ERROR!” appeared on the screen for 1500 ms; no feedback appeared for correct answers. The next trial was then presented after a pause of 500 ms.

After the practice trials, the four experimental blocks were presented to the participants. The order of blocks, as well as the order of items within each block, was randomized for each participant. The trial procedure of the experimental blocks was identical to that of the practice trials. Before each experimental block, participants had the opportunity to make a pause.

Participants received £1.88 for their participation in the study, which was estimated to take 15 minutes to complete (the actual median completion time was 11.9 minutes).

Results

As described in the Method section, participants were excluded from analyses if they gave correct answers in fewer than 85 % of experimental trials (that is, the practice trials were not considered). We further excluded trials with response times lower than 150 ms or higher than 1,500 ms (1.9 % of trials) before analysis.

Response time analysis. Response times were best predicted by layer 6 of the VGG-F model (see Table 1). For our analysis, response times were log-transformed (Baayen & Milin, 2010), and we only analyzed experimental trials (i.e., trials with two different images) for which correct responses were made. In order to keep the results of this study (and especially the R^2 values) comparable to the rating studies, response times were aggregated at the item level (by averaging over participants) before being subjected to analysis.⁸ For this purpose, the two orderings of a pair (CHEETAH – PHONE and PHONE – CHEETAH) were considered as two different items, since the actual responses were made to two different target images. Our aggregated dataset therefore consisted of 6,000 image pairs, with an average response time of 643.9 ms ($SD = 57.3$ ms) across items for correct responses. The Spearman-Brown split-half reliability for these average response times was estimated at $\rho = .481$.

We first estimated a GAM (Wood, 2006, 2015) containing a smooth term for IMG similarities. Random intercepts for the target images could not be included because the GAMs could not be reliably estimated. In this model, the smooth term of IMG similarities was significant ($F = 129.7$, $edf = 4.6$, $rdof = 5.7$, $p < .001$), with the variance explained by the GAM at $R^2_{adj.} = .110$. This effect corresponds to a difference of $\Delta(\overline{RT}) = 96.9$ ms

⁸ This does not affect the pattern of results, which was identical in a trial-level analysis with the GAMs including random intercepts for participants.

between the GAM estimate for the lowest IMG similarity ($\overline{\text{RT}}(.033) = 634.2$ ms) and the highest IMG similarity ($\overline{\text{RT}}(.892) = 731.1$ ms) in our dataset. Interestingly, this level of performance was even slightly superior to a GAM that predicted response times only from a smooth term for the participant ratings of image similarity collected in Study 2 (with this smooth term being significant at $F = 77.5$, $edf = 8.4$, $rdof = 8.9$, $p < .001$), where we observed $R^2_{adj.} = .103$.

We then examined whether these vision-based IMG similarities add *additional* predictive value for response times over and above the participant ratings. This was indeed the case: Adding a smooth term for IMG similarity to a GAM that already contained a smooth term for participant ratings of image similarity (as well as random intercepts for participants) significantly improved this model, as indicated by a likelihood-ratio test ($\Delta Deviance = 0.86$, $\Delta df = 4.7$, $p < .001$; with an increase in variance explained from $R^2_{adj.} = .105$ to $R^2_{adj.} = .123$; $\Delta R^2_{adj.} = .018$).

As in Study 2, we then additionally tested for potential effect of vision-based prototype similarities (PRO similarities) and language-based similarities (cbow similarities) for the labels used to annotate the images of each pair (note again that these labels were never presented to participants, see Figure 14). Adding a smooth term for PRO similarities further improved the model prediction, ($\Delta Deviance = 0.08$, $\Delta df = 4.4$, $p = .020$), but with only a very minor increase in variance explained to $R^2_{adj.} = .125$, $\Delta R^2_{adj.} = .002$). For this model, we obtained $CI^{R^2}_{.95} = [.109, .147]$. Adding a smooth term for cbow similarities did not improve model predictions ($\Delta Deviance = 0.01$, $\Delta df = 1.1$, $p = .172$).

In the resulting model (see Figure 15), higher IMG similarity, PRO similarity, and rating values all predicted slower responses. As can clearly be seen, these relations are almost entirely restricted to images with very high similarities – there are hardly any differences in processing times between low- and medium-similarity pairs, and the curves only start trending upwards towards the right end of the plots.⁹ These plots also highlight

⁹ Note that unlike for the cbow effect in Study 2, this shift for highly-similar items is based on a large

that, although significant in the GAM, the actual variance explained by the PRO similarity term is almost negligible.

Percent correct analysis. We repeated the analyses presented in the previous section for percent-correct data (overall, 97.29% of all responses were correct, $SD = .05$ across items, with a Spearman-Brown reliability for the average percentage of correct answers estimated at $\rho = .378$). To this end, we employed GAMs with a model structure identical to the one described in the section on response times, except for the change in the dependent variable. As for the response time analysis, layer 6 best predicted the data.

In a GAM predicting the percentage of correct answers from a smooth term for IMG similarities, this smooth term was significant ($F = 51.94$, $edf = 8.2$, $rdof = 8.8$, $p < .001$; $R_{adj.}^2 = .071$), as was the case for a GAM predicting correct answers only from a smooth term for rating values ($F = 56.12$, $edf = 8.7$, $rdof = 8.98$, $p < .001$; with a relatively similar amount of variance explained at $R_{adj.}^2 = .078$). As for response times, a GAM that included a smooth term IMG similarities in addition to a smooth term for rating values significantly outperformed a GAM that only included rating values ($\Delta Deviance = 0.21$, $\Delta df = 8.9$, $p < .001$ in a likelihood-ratio test, with $R_{adj.}^2 = .097$ and $\Delta R_{adj.}^2 = .019$ for the resulting model). For this model, we obtained $CI_{95}^{R^2} = [.074, .129]$. Additionally including a smooth term for PRO similarity did not significantly improve the GAM ($\Delta Deviance = 0.01$, $\Delta df = 2.2$, $p = .149$), and neither did including a smooth term for cbow similarity ($\Delta Deviance = 0.00$, $\Delta df = 0.91$, $p = .548$).

In the resulting GAM, higher values of IMG similarities predicted fewer correct responses, as did higher rating values (see Figure 15). As for response times, this effect was mainly due to a falloff in correct responses for very similar items, with a much smaller effect across the rest of the spectrum.

number of observations and is thus far more reliable.

Study 4: Discrimination task (images)

Figure 15. The effects of rating values (upper row), IMG similarity (middle row), and PRO similarity (lower row) on log-transformed response times (left column) and error rates (right column) in the discrimination task (Study 5), with .95-confidence intervals.

Discussion

In Study 4, we find the expected similarity effects on participants' performance in a discrimination task: For pairs of different but very similar images, participants were slower

and made more errors in their judgments, than when evaluating two less similar images were different. Remarkably, performance was predicted by the image-similarity (IMG) estimates obtained from the proposed models, indicating that the system captures dimensions that are relevant in the processing of visual stimuli by human participants.

As can be seen in Figure 15, this effect mainly emerged for very high similarities (with essentially no differences between low- and medium-similarity pairs), which underlines the importance of considering non-linear effects in order to get a precise picture of the relations between variables. In other words, once two images are sufficiently dissimilar, it does not really matter *how* dissimilar they are; participants just have a hard time discriminating between two very similar images. The fact that we find a very similar pattern of effects for rating values (which might be considered the “gold standard” for similarity scores) and our model-derived IMG similarities indicates that this non-linear pattern is a genuine psychological effect (in discrimination tasks, humans specifically struggle for very similar stimuli) rather than a systematic measurement error of our model (such as a scenario in which the model is only able to discriminate very similar image pairs from pairs that are not very similar, but fails to give sensible similarity scores for the rest of the similarity scale).

Intriguingly, this effect of IMG similarities even emerges when controlling for human rating values for the same items, and thus accounts for variance in human implicit responses that is not captured by human explicit judgments. This may partly be explained by both predictors imperfectly and complementarily tapping into the same latent variable that influences responses in the behavioral task. However, another possible explanation refers to a genuine difference between what these similarity measures capture. For instance, human ratings tend to be more heavily influenced by higher-level conceptual, semantic, and associative relations, scene properties and even cultural factors than DCNN similarities (which, if at all, only have very indirect access to such information provided via the labels during model training); DCNN similarities on the other hand lean more towards encoding

lower-level perceptual features such as color, shape, and texture (Jozwik et al., 2017; Peterson et al., 2016; Rosenfeld, Solbach, & Tsotsos, 2018; see also the language-based similarity effect in Study 2). This is especially relevant here since the discrimination task calls for a purely perceptual comparison, and accordingly performance in this task is primarily influenced by perceptual similarities between the stimuli (Ehrenstein & Ehrenstein, 1999; Pelli & Farell, 1995); in this context, human ratings might be too “diluted” by higher-level conceptual factors to fully capture the behavioral phenomenon.

In addition to the IMG similarity effects, our statistical analysis also yields a PRO similarity effect on response times, but not on response accuracy. However, a critical examination of this effect raises some doubts: The effect is extremely small (only accounting for 0.1% of variance explained) and its regression line is essentially flat (see Figure 14). In our view, this does not form a sufficiently solid empirical basis for the quite profound theoretical implications of such an effect (an activation of vision-based conceptual representations during a discrimination task). We therefore refrain from interpreting this effect.

Taken together, results of Study 4 thus demonstrate that measures derived from our vision-based models predict human behavior in on-line processing tasks. Importantly, the predictive power of the model does not fall behind human judgments that are traditionally taken as gold standard measures: When considered on their own, the variance explained by IMG similarities is higher than that explained by rating values, and (as indicated by a likelihood-ratio test), IMG similarities explain variance in the behavioral data that rating values are not able to account for. Considering the fact that IMG similarities are also a strong predictor for rating values (which themselves are behavioral data that a psychological model needs to explain rather than taking them as a given; Westbury, 2016), these results demonstrate the usefulness and predictive power of our vision-based model representations.

Notably, the split-half reliability for the behavioral data in Study 4 is considerably

lower than the one observed for the ratings in the previous studies (compare Table 1). This is not simply due to the nature of these variables (explicit rating vs. response time): Studies with lexical material report split-half reliabilities for response-time data between .70 and .80 (Ferrand et al., 2010; Keuleers, Diependaele, & Brysbaert, 2010; Keuleers, Lacey, Rastle, & Brysbaert, 2012). It is also not due to the web-based data collection: Mandera, Keuleers, and Brysbaert (2020) report a convergent validity of .73 with the lab-based response times for the same items from Keuleers et al. (2012). However, the task and type of item material plays an important role: For (web-collected) lexical decision and sensibility judgment times on compound words in the datasets by Günther and Marelli (2020), we observe $\rho = .63$ and $\rho = .71$, respectively, for the “yes”-responses for existing words, but only $\rho = .44$ and $\rho = .46$, respectively, for the “no”-responses for novel words. Since images are much more complex than words, and since the images we used are not highly familiar, over-learned stimuli, more variability in the present data is to be expected.

Study 5: Priming paradigm (“visual decision task”) with image pairs

In Study 4, we employed a discrimination task to investigate the effects of vision-based similarity on participants’ on-line processing performance, as measures by response times and accuracy. Of course, as any experimental paradigm, the discrimination task comes with specific task demands that may influence performance. For instance, the discrimination task necessarily involves a direct comparison between the two stimuli presented.

To increase the generalizability of our results, we evaluated our vision-based similarities against human behavior in another on-line processing paradigm: Priming (Lashley, 1951; Meyer & Schvaneveldt, 1971; Tulving & Schacter, 1990). In a standard priming paradigm, a prime stimulus is presented before a target stimulus, and participants have to provide a response concerning the target. Processing times for the target are typically faster and more accurate the more similar the prime and the target are, with a common explanation being that the prime pre-activates mental representations which then

facilitate the access to related targets (Hutchison et al., 2013).

Priming effects due to the visual similarity between prime and target images (as determined by human raters) have been observed in previous studies, typically employing a picture naming task (Carr, McCauley, Sperber, & Parmelee, 1982; Sperber, McCauley, Ragain, & Weil, 1979). However, picture naming introduces several additional higher-level task demands other than image processing, such as categorization and labeling (and thus, an activation of the language system). To avoid this, we developed the “visual decision task”, in analogy with the prominent lexical decision task from lexical priming studies (Hutchison et al., 2013; Lucas, 2000). In the lexical priming task, participants are presented with two words – the prime, which typically is a real word, and the target, which may be a real word or a nonword. Participants have to decide whether the target is a real word or a nonword. When it comes to real-word target, a decision is usually faster and less error-prone the more similar prime and target are (a priming effect; Hutchison et al., 2013; Meyer & Schvaneveldt, 1971). In our visual decision task, we included as targets either real images (in analogy to words) or distorted images (in analogy to nonwords), asking participants to decide whether the target is a real or distorted image. Targets were preceded by real images as primes, irrelevant for the required task. In fact, unlike the discrimination experiment, the visual decision task does *not* require a comparison between images – in principle, participants could always give the correct response even if they ignored the prime.

Based on the results from the visual priming literature (Carr et al., 1982; Sperber et al., 1979) as well as the lexical decision results from the lexical priming literature (Hutchison et al., 2013), we expect the vision-based similarity measures derived from our model to predict priming effects – faster responses and fewer errors towards real-image targets the more these are similar to the prime image – in the visual decision task.

Method

Participants. The data for this study was collected from 753 participants, recruited via *Prolific*, with the same participation restrictions as Study 4 (native, monolingual speakers of English with normal or corrected-to-normal vision). Of these ($M_{Age} = 35.16$ years, $SD_{Age} = 12.50$ years), 490 self-identified as female, 257 as male, and five indicated another gender identity. As in Study 4, participants were explicitly and repeatedly advised not to participate in this study if the rapid succession of visual stimuli might have negative consequences for their well-being or health.

Data from 16 additional participants were excluded due to high error rates (see Results section), and a further 14 participants were excluded due to technical errors during data collection (incomplete data files).

Material. We employed as experimental material the same 3,000 image pairs (including 2,228 unique images) which were already used in Study 2 and study 4. These were again arranged into the 25 different experimental lists of 120 image pairs created for Study 4. In each list, we included an additional 120 image pairs to be used in the trials requiring a “no”-response in the visual decision (see below). Each of these pairs consisted of two different images that were not yet included in the experimental pairs. Within each list, the experimental items and “no”-response items were randomly arranged into four blocks, so that each block contained 30 experimental items and 30 “no”-response items. Each list additionally included 32 practice trials, which were randomly sampled with the restriction that none of their images already appeared in the experimental or “no”-response items. Finally, we constructed a second version for each list by reversing the order of images within each pair, thus producing a total of 50 different experimental versions (25 lists in two versions each).

After constructing these lists, we transformed the targets of all “no”-response pairs. As described before, the setup for this priming study was inspired by the lexical decision task, in which participants have to decide whether a target word is a real word (TURTLE)

or a nonword (TURPLE). As in this example, the no-response targets are often pronounceable nonwords that are derived by altering existing words (Keuleers & Brysbaert, 2010), maintaining the word surface properties and respect for the language orthotactics while removing its arbitrary association to a given meaning. We adopted this rationale to generate “no”-response targets for the visual priming task: These targets were created by applying a diffeomorphic scrambling transformation (Stojanoski & Cusack, 2014), which “preserves the basic perceptual properties of the image while removing meaning”. We set the amount of scrambling to 20 and the number of scrambling steps to 30 (Stojanoski & Cusack, 2014). These parameters were identified as producing the strongest priming effects in a pilot study, using the same procedure as the study described here with unambiguously very dissimilar, medium similar, or highly similar image pairs (according to both the ratings from Study 2 and the IMG model). Finally, since this transformation also scrambles the borders of images (producing an extremely obvious cue for whether an image is real or scrambled), we cropped the borders and slightly zoomed into the images (using the *ImageMagick* software; The ImageMagick Development Team, 2021) to create images of size 299×224 pixels with straight borders. Examples for the experiment targets are presented in Figure 16.

As in Study 4, all images were presented with a size of 299×224 pixels in the center of the screen (for the real images, this was their original size; note again that the scrambled images were cropped to this size), in front of a light-gray background. The experiment was presented in full-screen mode, to minimize the interference of non-relevant visual stimuli. In addition to the images, we employed the same set of visual masks used in Study 4 (Aly et al., 2014). In the pilot study, we found that priming effects were more pronounced when a mask was presented after the prime, as compared to a procedure in which the prime was followed by a blank page for the same duration.

Procedure. After informed consent and a questionnaire for demographic data, participants were instructed that they would be presented with pairs of images in short

succession, and that their task was to decide whether the second image was real or distorted by pressing the *F* or *J* key on their keyboard. Note that the response in this “visual decision task” only depends on the properties of the target stimulus, and the prime (or its relation to the target) is irrelevant for the purpose of producing a response – as it is the case for the lexical decision task in lexical priming. The key-to-response assignment was randomized for each participant.

As in Study 4, all presented images were pre-loaded at the start of an experimental session. The experiment started with the 32 practice trials in random order. The trial procedure is illustrated in Figure 16. Each trial started with a fixation cross for 750 ms. Then, the first image (the prime) appeared on the screen for 250 ms, followed by four masks of 30 ms each¹⁰ – a dynamic noise mask (Aly et al., 2014), consisting of one randomly-sampled mask from each of the four visual-resolution sets in random order. The total SOA thus amounts to 370 ms, with image priming effects having been observed at SOAs of 300 ms and 400 ms (Carr et al., 1982). After the mask, the second image (the target) was presented until participants made their response. If this response was incorrect, the feedback “ERROR!” appeared on the screen for 1500 ms; no feedback appeared for correct answers. The next trial was then presented after a pause (blank page) of 750 ms.

After the practice trials, the four experimental blocks were presented to the participants. The order of blocks, as well as the order of items within each block, was randomized for each participant. The trial procedure of the experimental blocks was identical to that of the practice trials. Before each experimental block, participants had the opportunity to make a pause.

Participants received £1.88 for their participation in the study, which was estimated to take 15 minutes to complete (the actual median completion time was 15.0 minutes).

¹⁰ Due to factors such as different internet connections and end devices, the actual presentation times (derived from timestamps in the data) were slightly above the times indicated here (although this only ranged between 0 and 10 percent for the overwhelming majority of cases).

Figure 16. (Upper panel) Examples for scrambled images serving as “no”-response targets in Study 4, alongside the original images from which they were created. (Lower panel) The trial procedure in Study 4.

Results

We followed the same analysis procedure described for Study 4. “No”-response items trials – in Study 5, trials with scrambled target images – were excluded from analysis. For the

analysis, the two orderings of each pair were considered as two separate items (since responses were made for two different targets). Besides datasets from the 16 participants with high error rates (less than 85 % correct), trials with response times lower than 150 ms or higher than 1,500 ms (3.6 % of trials) were excluded from the analysis.

Response time analysis. For the response time analysis, we only considered trials in which participants gave correct answers. In these data, the average response time was 671.3 ms ($SD = 61.1$ ms across items), with a Spearman-Brown split-half reliability for these average response times estimated at $\rho = .366$. Log-transformed response times in the priming task were best predicted by layer 6 of the VGG-F model (see Table 1).

In the GAM (Wood, 2006, 2015) containing only a smooth term for IMG similarities, this effect was small but significant ($F = 16.3$, $edf = 3.8$, $rdof = 4.8$, $p < .001$), with the model's adjusted R^2 -value at $R^2_{adj.} = .013$. This effect corresponds to a difference of $\Delta(\overline{RT}) = 38.6$ ms between the GAM estimate for the lowest IMG similarity ($\overline{RT}(.033) = 694.4$ ms) and the highest IMG similarity ($\overline{RT}(.892) = 655.9$ ms) in our dataset (note again that, differently from the discrimination task, *faster* responses are expected for more similar items in a priming paradigm). Even though the effect of IMG similarities is relatively small in terms of R^2 , the predictions of a GAM only containing this effect are superior to a GAM predicting response times only from a smooth term for the participant ratings of image similarity collected in Study 2 ($F = 42.1$, $edf = 1.0$, $rdof = 1.0$, $p < .001$), for which we observed $R^2_{adj.} = .007$. In fact, adding IMG similarities to a GAM that already contained rating values significantly improved the model fit ($\Delta Deviance = 0.28$, $\Delta df = 4.8$, $p < .001$, with an increase in variance explained from $R^2_{adj.} = .007$ to $R^2_{adj.} = .013$; $\Delta R^2_{adj.} = .006$) (for this model, we obtained $CI^{R^2}_{.95} = [.009, .022]$). However, we did *not* observe such an improvement when adding rating values to a model that already contained IMG similarities ($\Delta Deviance = 0.01$, $\Delta df = 1.1$, $p = .323$).

On the other hand, adding a smooth term for PRO similarities to a GAM that

already contained IMG similarities significantly improved the predictions ($\Delta Deviance = 0.04$, $\Delta df = 0.9$, $p = .017$, with a very small increase in variance explained to $R^2_{adj.} = .013$; $\Delta R^2_{adj.} = .001$). We did not observe such an improvement when adding smooth terms for cbow similarities to the model ($\Delta Deviance = 0.00$, $\Delta df = 1.0$, $p = .485$). In the resulting model, higher values of IMG similarities predicted faster responses, with a relatively monotonous effect across the whole spectrum of values (see Figure 17). The effect of PRO similarities appears to run in the opposite direction, but this may be the result of statistical suppression resulting from collinearity (Friedman & Wall, 2005; the correlation between IMG and PRO similarities in layer 6 is $\rho = .75$): In terms of zero-order rank correlations without considering additional predictors, both the effects of IMG similarity ($r_s = -.10$) and PRO similarity ($r_s = -.06$) would be considered facilitatory (see Table 2).¹¹ In any case, note again that the PRO parameter explains hardly any additional variance in the data.

Percent correct analysis. Overall, 97.87% of all responses were correct ($SD = .05$ across items, with a Spearman-Brown reliability for the average percentage of correct answers estimated at $\rho = .253$). In the GAM analysis, layer 8 of the VGG-F network best predicted the data (see Table 1), closely followed by layer 6 and layer 7. In a GAM predicting the percent-correct data only from a smooth term for IMG similarities, this effect was small but significant ($F = 20.7$, $edf = 1.0$, $rdof = 1.0$, $p < .001$; $R^2_{adj.} = .003$), with considerably more variance explained than a GAM predicting correct answers only from a smooth term for rating values ($F = 7.0$, $edf = 1.0$, $rdof = 1.0$, $p = .008$; with $R^2_{adj.} = .001$). A GAM including both rating values and IMG similarities performed significantly better than a GAM only including ratings values ($\Delta Deviance = 0.02$, $\Delta df = 1.0$, $p < .001$ with $R^2_{adj.} = .003$ and $\Delta R^2_{adj.} = .002$ for the resulting GAM), but it did not perform better than a GAM only including IMG similarities ($\Delta Deviance = 0.00$, $\Delta df = 1.0$, $p = .177$). Rating values were therefore dropped as predictors. Similarly,

¹¹ We do not find these changes in sign that can result from statistical suppression in the other Studies.

Figure 17. The effects of IMG similarity (upper row), PRO similarity (lower row, left panel) and rating values (lower row, right panel) on log-transformed response times (left column) and error rates (right column) in the priming paradigm (Study 5), with .95-confidence intervals.

additionally including PRO similarities ($\Delta Deviance = 0.00$, $\Delta df = 1.8$, $p = .558$) or cbow similarities ($\Delta Deviance = 0.00$, $\Delta df = 1.0$, $p = .364$) did not further improve the GAM already containing IMG similarities.

In the resulting model (for which we obtain $R_{adj.}^2 = .003$, $CI_{.95}^{R^2} = [.001, .009]$), higher values of IMG similarities predicted a higher propability of correct responses, in a relatively linear manner across the spectrum of values (see Figure 17).

Discussion

The observed IMG similarity effect – faster responses and fewer errors for more similar prime-target images – in the priming task using the visual decision paradigm again demonstrates that the proposed visual system is able to capture behavioral effects in on-line processing tasks, thus generalizing the results of Study 4. However, Study 5 also shows several interesting differences with respect to Study 4: First, differently from the discrimination task, the visual decision task does not require a comparison between the images, indicating an implicit appreciation of prime-target similarity by part of the participants (in line with the lexical priming literature; Hutchison et al., 2013; Petilli et al., 2021). This offers a plausible explanation for the smaller amount of variance explained by IMG similarity in Study 5 (where evaluating the similarity between images is not required to perform the task) as compared to Study 4 (where an explicit comparison between the two images need to be carried out) (see Table 1). Second, while the effect of IMG similarities in the discrimination task appears to play a role only for very similar item (Figure 15), its effect in the visual decision task can be observed across the whole range of values (Figure 17). This shows that IMG similarities don't only differentiate very similar from less similar images, but also capture subtler differences between less similar images, and demonstrates the importance of evaluating model-based measures across multiple different tasks. Third, we found already in Study 4 that IMG similarities are a better predictor for behavioral data than human ratings of image similarity; this finding is further emphasized in Study 5, where IMG similarities completely supplant human ratings as predictor for both processing times and error rates. These findings highlight the potential dissociation between the factors influencing explicit human judgments on the one hand and their on-line processing behavior on the other hand (see again the Discussion of Study 4 for possible explanations). Thus, in addition to theoretical reservations against taking human judgments as gold standard measures of psychological variables (Westbury, 2016), models building on the actual properties of stimulus (i.e. the input level) can surpass human-based

measures also in terms of empirical performance.

In addition to the effect of IMG similarities, our statistical analysis also yields an effect of PRO similarities on response times (but not on error rates). However, as it was already the case for Study 4, this effect is extremely small, more or less flat, and explains hardly any variance. As argued in Study 4, we doubt that this forms a reliable enough empirical foundation, considering the theoretical implications of such an effect, and thus refrain from further interpreting it.

General Discussion

In the present paper, we examined the high-dimensional vision-based representations produced by the *ViSpa* system, an adaptation of a pre-trained VGG-F model (Chatfield et al., 2014; Vedaldi & Lenc, 2015), from a cognitive perspective. To this end, we evaluated the similarities between these representations against behavioral data in a series of large-scale experiments, thus providing a thorough examination and evaluation of the model.

Summary of results

The results of this evaluation are summarized in Table 1. At a general level, we observe that PRO similarities predict visual as well as semantic similarity ratings for word pairs (Study 1), IMG similarities predict visual similarity ratings for image pairs (Study 2), and typicality scores (i.e., similarities between an image representation and its prototype representation) predict typicality ratings (Study 3). Thus, all types of scores produced by the model predict some form of participant ratings, which are usually taken as standard measure for model evaluations (Baroni et al., 2014; Bruni et al., 2014; Jozwik et al., 2017). Going beyond ratings, we also demonstrate that IMG similarities predict processing times in an image discrimination task (Study 4) and a visual decision priming task (Study 5). Notably, the similarities derived from the model representations consistently predict the empirical data better than those derived from physical-level representations (RGB and

HOG, see Table 1), indicating that the observed behavior is influenced by higher-level vision-based representations on the path to object recognition and classification.

More specifically, the results of our studies can be summarized as follows (see again Table 1): First, for all experiments employing image pairs (Studies 2, 4, and 5) we find IMG similarity effects. For processing times (Studies 4 and 5), the predictive power of IMG similarities actually exceeds that for participants' rating values for the same image pairs, to the point where IMG similarities even supersede rating values as predictors for processing times in the priming paradigm (Study 5). These results provide a strong case for the cognitive plausibility of the IMG representations. Since the individual images in the IMG spaces form the core foundation of the *ViSpa* system – the PRO representations are derived from the IMG representations by averaging – this attested validity of the IMG representations is a necessary prerequisite for the plausibility of the entire representational system.

Of course, the averaging itself does not guarantee that the PRO representations of the *ViSpa* system are actually plausible prototype representations. Evidence for this plausibility is however provided in Study 3, where we find that the model indeed captures human intuitions on typicality – the central theoretical construct in prototype theories (Rosch & Mervis, 1975; Rosch, 1978). Thus, Study 3 provides vital evidence tying together the IMG and PRO representations as a unitary system.

Vision-based and language-based similarities. The PRO representations themselves (or, more specifically, the similarities between them) also predict behavioral data across different tasks. Most importantly, these PRO measures predict visual as well as semantic similarity ratings in Study 1 – where participants gave judgments on *word* pairs – over and above the semantic similarity of the same word pairs, as measured by a distributional semantic model (*cbow*, Baroni et al., 2014; Mikolov, Chen, et al., 2013). This mirrors results for on-line processing by Petilli et al. (2021), who observed PRO similarity effects over and above language-based similarity effects in a lexical priming task. Taken

together, these findings offer evidence that both visual experience on the one hand and language experience on the other hand provide distinguishable and relevant contributions to building our conceptual representations (Andrews et al., 2009, 2014). However, as it would be expected, their relative influence depends on task demands as well as the item material (see Table 1 and Table 2): For semantic similarity judgments of word pairs, language-based similarity has the strongest relative effect as compared to visual similarity; the relative effect of vision-based similarity becomes already more pronounced for *visual* similarity judgments of word pairs, and stronger than the language-based effect for tasks involving images. Thus, a higher focus on visual information – in task demands or material – comes with stronger effects of vision-based representations.

However, the impact of either source of information might vary between different items. De Deyne et al. (2021), in a study comparing the relative contribution of vision-based and language-based representations for predicting human similarity ratings for words, observed results for the MEN dataset consistent with those reported in Table 1 and Table 2; however, they also showed that the relative influence of linguistic and visual information depends on the level of description for these words: In their study, visual information added to the predictions for basic-level concept pairs of the same category (such as *dog-cat* or *Dalmatian-poodle*), but hardly for higher-level or non-taxonomic pairs (*dog-bone*). Similarly, although visual and linguistic data can in principle be integrated for any concept, irrespective of its degree of concreteness, the more a word is perceived as abstract by the speakers of a language the more a model will have a hard time aligning a sensible visual representations with its language-based one (Günther, Petilli, Vergallito, & Marelli, 2020). As a result the importance of visual information in similarity estimates will depend, up to a certain extent, on the very concepts that are being compared, over and above task requirements.

In the interpretation of these patterns, it is also important to consider that the modality of data considered – images versus text – is not the only difference between the

vision-based model (VGG-F) and language-based model (cbow) used here. The VGG-F model is trained to categorize images starting from their basic form-level (pixels of the image), and in its layers learns to exploit more and more complex visual feature combinations to make these judgments. Except for a background that is part of some images, it has no information about which other visual objects appear in similar contexts. The cbow model on the other hand is trained to predict words from the other words surrounding them, and has no explicit access to independent information related to the word referent (such as its features) – any such information picked up by the model is entirely inferred from the contexts in which the word is used (Günther, Petilli, Vergallito, & Marelli, 2020). It might therefore be possible that the VGG-F model tends to pick up semantic/taxonomic information (i.e., obtains more similar representations for objects that share more features), while the cbow model tends to pick up more associative/thematic information (i.e., obtains more similar representations for entities occurring in similar contexts). This could affect the empirical patterns described in this section, in addition to modality-specific effects: For example, the finding that language-based similarities have a stronger effect on word ratings might partially be explained by the fact that speakers in such tasks have a hard time ignoring thematic relations when making similarity judgments, even when they are asked to focus on taxonomic relations only (Honke & Kurtz, 2019). In order to more thoroughly disentangle the effects of modality and this potential taxonomic/thematic-tendencies, one direction for future research could be to employ vision-based models that consider co-occurrence patterns between visual objects to build their representations (e.g. Sadeghi, McClelland, & Hoffman, 2015), thus applying the mechanism underlying the cbow algorithm to vision.

Prototype and image similarities. We further observe effects of PRO similarity in the tasks employing image material (Studies 2, 4, and 5). In Study 2, these can be interpreted as participants also considering more general visual experience with the entities depicted on the images, rather than only the specific images themselves, in order to form

their similarity judgments. In principle, one could arrive at a similar interpretation on the basis of the statistically significant PRO similarity effects in Studies 4 and 5 – participant could be assumed to routinely activate such generalized, conceptual information about a depicted entity when processing an image. However, as pointed in the Discussions to these studies, these effects are extremely small, hence their reliability is arguably very limited.

Notably, the comparison between Study 1, 2 4 and 5 shows a clear dissociation between IMG similarities and PRO similarities in terms of their predictive power (see Table 1). This predictive power is not only determined by the type of stimuli used (words or images), but also by the specific task at hand (for example, off-line judgments versus on-line processing). These results strongly indicate that, even though they are correlated, both measures are useful and add unique, non-redundant explanatory power.

In the *ViSpa* system, all prototype representations are operationalized as the centroid of the individual image representations in a category. However, there are differences in the internal structure of these categories: For some categories (such as *spaghetti* or *daisy*), all images are relatively similar to one another, and consequently, the prototype representation is relatively similar to all images, while for other categories (such as *angler* or *fish*) all these representations can be relatively far from one another. There is a possibility that these differences are systematic: For example, one could expect that the individual exemplars will become more different for higher-level categories (e.g., *animal*), as compared to subordinate categories (e.g., *fish*, or even *salmon*). Determining to what extent this is the case, and if it affects the quality and predictive power of the prototype representations at the different levels of description, requires future studies with careful annotations for this level of description.

However, we want to note that the reduced internal consistency of prototype representations as a function of the taxonomic level is not unique to PRO representations in the *ViSpa* system; the concrete mechanism of averaging over multiple individual representations that can be more or less diverse just makes it more evident. For example,

also in language usage words at a higher level of description tend to be used in more diverse contexts (Lenci & Benotto, 2012; Kotlerman, Dagan, Szpektor, & Zhitomirsky-Geffet, 2010): A speaker can use the word “dog” in more or less the same contexts in which s/he uses “poodle”, and also in additional contexts (including, but not only, the contexts in which she can use other subordinates of “dog”, such as “labrador” or “terrier”). This has a clear impact on language-based models that estimate semantic representations from word distributions over contexts (distributional models such as the cbow model employed here): in our example, the representation for “dog” (or “pet”, or “animal”) will be aggregated over more different contexts than the representation for “poodle”, and hence be less internally consistent. The same consideration even applies to annotation-based quantitative representations: For instance, in the Lancaster Sensorimotor Norms (Lynott et al., 2020), word meanings/concepts are represented via their ratings on several perceptual and action-related dimensions. However, when judging the amount of haptic experience people have made with “pets”, these judgments will also aggregate (most likely in a biased manner) over different pets for which such experience is not consistent (e.g., dogs and cats, but also goldfishes). Alternatively, in semantic feature norms such as by Buchanan, Valentine, and Maxwell (2019), whether a concept such as “animal” is annotated with features such as “domesticated” or “furry” (both of which were selected as features for the concept) becomes less and less clear and consistent for higher-level concepts (these two features are only true for a subset of animals). Thus, reduced consistency for higher-level concepts is a feature of prototype representations in general, rather than of *ViSpa* PRO representations specifically.

Task differences. A comparison between the various studies also reveals substantial differences in the variance explained by the different similarity measures (see Table 1). Here, the main factors at play are the nature of the task and the type of response requested, with variance explained being much higher for off-line ratings (Study 1 and 2) than for on-line response times and accuracy rates (Study 4 and 5). In this context, it is

important to consider the large differences in the split-half reliability between these dependent variables (Table 1; see the Discussion of Study 4), which set an upper bound to the variance that can be explained by stimulus characteristics. Still, in the off-line rating studies, a larger proportion of the variance that can be explained is actually explained by the predictor variables (see Table 1). This closely reflects the literature on language-based distributional models, where the same divergence has been observed (Mandera et al., 2017). Overall, the on-line responses appear to be much more subject to systematic influences of factors other than similarities (for lexical stimuli, this can include variables such as frequency and length; for image stimuli, this can for example include visual familiarity or complexity). Furthermore, participants in the rating tasks (Study 1 and 2) were explicitly instructed to make judgments about the variable of interest, the similarity between stimuli. This is not the case for the on-line tasks: The discrimination task (Study 4) only required same-different judgments, which can at best be seen as a very coarse binary similarity judgment (between perfect similarity and anything else). And as mentioned in the Discussion of Study 5, the priming task does not require any comparison between the stimuli. Thus, similarity effects in these tasks result automatically from the way in which the cognitive system operates while performing them, rather than from task demands themselves.

The layers of the network. We compared the performance of the different layers in an exploratory manner, to subject *all* model representations (and not just an arbitrary selection) to empirical examination. Across all studies, the best predictions (in numerical terms) are derived from the deeper layers (5–8) of the VGG-F network (see Jozwik et al., 2017), which are all relatively similar to one another (Figure 4 and 7). These layers further down the processing stream (Jozwik et al., 2017; Kriegeskorte, 2015) have been shown to capture higher-level, more abstracted gestalt information rather than simple visual features (LeCun et al., 2015; Mahendran & Vedaldi, 2015; Zeiler & Fergus, 2014). In fact, the deeper layers of the network show higher levels of performance than the purely

physical image representations given by the RGB codes in layer 0 (as well as the HOG representations, Lowe, 2004), and typically outperform them statistically (see Table 1), thus justifying the additional complexity added by the network architecture. Given the results of our studies (Table 1), we recommend to use the representations in layer 6 for most general purposes – layer 6 often shows the best performance (Studies 2, 4, and the response times in Study 5), and even in the other cases (Study 1 and error rates in Study 5) it is not far behind the best-performing layer and among the top candidates. However, this recommendation should be taken as a general-level rule-of-thumb – for specific tasks and theoretical questions, it is possible that another layer that better captures lower- or higher-level visual information, or even another model architecture altogether, might be more adequate or produce better results on an empirical level (compare Baroni et al., 2014; Lenci et al., 2021; Wingfield & Connell, 2022, in the language domain).

Interim conclusion. Taken together, the present studies return a very positive evaluation of the *ViSpa* system as a model for vision-based mental representations. The item sets considered in our evaluation covered thousands of prototypes and images which, even though they are only a subset of the thousands of prototype representations and especially the tens of thousands of image representations currently available in *ViSpa*, still provide a solid empirical basis. Thus, *ViSpa* offers a valuable resource for researchers in many different fields, from those using naturalistic images as experimental material to those studying the role of sensory information in conceptual representation and language processing. For those researchers, *ViSpa* is openly available as a resource with convenient interfaces – see the section *Availability of the resources* at the end of this paper.

Training data and the relation between vision-based representations, the conceptual system, and language

As pointed out in the previous section, the vision-based representation of the network layers of the *ViSpa* system are not simply representations of the physical visual properties

of images – this would be the representations in layer 0, which were clearly outperformed by the deeper layers. Rather, the representations of the network layers capture the visual information used to classify and categorize images, according to the labels used to describe them (Chatfield et al., 2014; Krizhevsky et al., 2012). Therefore, all representations except for layer 0 are inherently influenced by a (human) conceptual system – the network is trained to categorize images, and these categories are an interpretation of humans segmenting physical stimuli into discrete classes, rather than an a-priori property of the physical world. In fact, the network is not only influenced by human categories, but even by *language* (see the discussion in Günther, Petilli, & Marelli, 2020): categorization by part of the model is characterized as selecting the correct label for an image (see Figure 1). Here, these labels were provided by a small sample of English-speaking expert annotators and thus depend on the linguistic classifications and intuitions of these speakers; however, different speakers or speaker communities can of course apply different linguistic conceptualizations for the same stimuli or even the same category structure (for example, using the same label or two different labels for two different categories, Wolff & Holmes, 2011). As a direct consequence of this, the instance of the *ViSpa* system presented here is influenced by how English speakers label visual stimuli (note again that the DCNN used here was originally trained to predict the English word labels for pictures Vedaldi & Lenc, 2015). This of course introduces some degrees of freedom to the model: Decisions had to be made as to which images are classified as belonging and not belonging to the same category, and which label is or is not used for this category. Thus, through the mechanisms of model training and the network weights learned in this process, already the IMG representations end up being influenced by a human-centered conceptual perspective. At the same time, however, once model training is complete, the network can generate representations for any input image purely from its physical properties (the RGB values), without requiring any labels for newly-provided images.

The same does not hold for the PRO representations of the *ViSpa* system: These are always constructed in a top-down manner, as they are directly defined by human-produced category labels - and, as a direct consequence, specific to a given language (here, English). In addition, these externally-selected category labels directly affect the set of images whose representations are averaged to form the prototype representation (with an additional degree of freedom introduced by our sampling of images for this purpose). Thus, the influence of human categorizations and language is especially relevant for the PRO representations (see also Günther, Petilli, & Marelli, 2020). In addition, the selected images can never be an exhaustive set of all possible images for that category, and thus the PRO representations – which are taken as an abstracted concept-level representations – are necessarily based only on the visual experience with a subset of all possible exemplars. As a consequence, the PRO representations also depend on the composition and representativeness of the images included in a given category - for example, if these images are selected to cover as many exemplars as possible (also unusual and atypical ones) without necessarily considering these exemplars’ saliency or actual occurrence in the environment, then this might skew the prototype representations. While the results of our studies (and especially Study 3 on typicality ratings) attest a general validity of our prototype representations, there is potential room of improvement by relying on exemplars that are carefully selected to be representative of a typical person’s actual experience.

More generally, vision-based models centrally rely on a specific training set – for the VGG-F model investigated here, the ImageNet database (Deng et al., 2009; Vedaldi & Lenc, 2015). This database is not necessarily a representative sample of human visual experience – for instance, dog breeds and bird species are heavily over-represented (Elgendy, 2020). At the same time, DCNNs are known to typically perform quite poorly for out-of-distribution inputs, stimuli that deviate substantially from the training data which the models tend to mis-classify while still being confident about their estimate (Goodfellow, Shlens, & Szegedy, 2015; Hendrycks & Gimpel, 2016). If a model were to

only encounter images of cars and airplanes during training, it would end up inadequately capturing features that are not needed to discriminate between these categories (but which might be very relevant for other categories, such as animals). This is known as *shortcut learning*, the tendency to develop decision rules that provide good outcomes for a given learning set but fail to generalize – a characteristic of both artificial and biological learning systems (Geirhos, Jacobsen, et al., 2020). This results in poor performances of DCNNs (or, more precisely, drastic divergences from human intuitions) for various kinds of image distortions and alterations (Baker et al., 2018; Golan, Raju, & Kriegeskorte, 2020; van Dyck & Gruber, 2020; see Geirhos, Jacobsen, et al., 2020 for an overview). Since the VGG model used here is trained on (a portion of) ImageNet, its representations will likely only be accurate for items that are structurally analogous to the naturalistic images/photos included in ImageNet. While this is not a major issue for the instance of the *ViSpa* system presented here – which is entirely based on ImageNet categories and naturalistic images – its generalizability for substantially different stimuli will most likely be limited. However, recent studies suggest that reasonable representations for other types of images (such as line drawings and even simplistic sketches) can be extracted from DCNNs trained on naturalistic images (Singer, Seeliger, Kietzmann, & Hebart, 2021; as long as these depict some internal structure in addition to just the outer contours of objects, Baker et al., 2018) allowing for some cautious optimism in this matter.

Moreover, as most datasets, ImageNet is known to be subject to several other biases than over- or under-representations of certain classes (Torralba & Efros, 2011) – for example, it often does not accurately distinguish brides and grooms from non-American countries (Shankar et al., 2017). Biases such as the ones discussed here will directly influence the resulting model representations. However, it is important to point out that exactly the same arguments apply for more traditional language-based distributional semantic models, where concept-level representations are derived from word distributions in large corpora of natural language. Also here, some speaker had the freedom to choose

which word to use or not to use to refer to a concept, and for which concepts they use the same or different words. Furthermore, the language corpora forming basis for distributional models are also only a sample of all the language experience that is possibly available to a speaker, and also subject to biases (Caliskan, Bryson, & Narayanan, 2017). Still, distributional models have been firmly established as cognitively plausible models for language-based semantic representations (Günther et al., 2019; Jones et al., 2015). As shown in the present study, the same holds for vision-based representations in the *ViSpa* system, which successfully predict the behavior of participants despite these under-specifications. We can only expect the model performance to increase even further when the *ViSpa* system or a similar model is estimated on a more representative collection of visual experience, which we leave to future studies.

Alternative models and fine-tuning

In the present study, we implemented the *ViSpa* system on a specific model architecture (Chatfield et al., 2014), which is extremely prominent in the computer vision literature and has been found to perform well in cognitive studies (Jozwik et al., 2017; Kubiľius et al., 2016; Peterson et al., 2016, 2018; Petilli et al., 2021; Roads & Love, 2021). Our objective was to subject this instance of the system (which has been employed in previous studies, Günther, Petilli, & Marelli, 2020; Günther, Petilli, Vergallito, & Marelli, 2020; Petilli et al., 2021) to a wide array of empirical evaluations against various large-scale behavioral datasets.

At the same time, the data collected for the evaluation studies presented here constitute invaluable resources in their own right, and are at least as large (in most cases much larger) than many prominent datasets typically used in the evaluation of computational models (Bruni et al., 2014; Finkelstein et al., 2001; Hill et al., 2015; Rubenstein & Goodenough, 1965). Other recent studies have also begun to provide larger datasets: For example, the work by Roads and Love (2021) described in Study 2 provides

similarity ratings for 50,000 ImageNet images. However, while this data is much “wider” in that it covers more images and image pairs, our Study 2 is much “deeper” in that it contains more individual responses per image pair, and is ultimately based on a much larger sample of individual participant responses (which allows us to estimate robust similarity scores; Hollis, 2018; Hollis & Westbury, 2018). We thus believe that our data – openly available at <https://osf.io/qvw9c> (Günther et al., 2022) – can serve as an ideal gold standard for future evaluations and comparisons of different computer-vision models.

There are several directions that can be explored from here. One is to compare different conceptually similar models (i.e., other supervised image classifier models) that differ in their specific architecture (for example, other variants of the VGG model, Simonyan & Zisserman, 2014; CaffeNet/AlexNet, Krizhevsky et al., 2012; GoogLeNet, Szegedy et al., 2015; or ResNet; He, Zhang, Ren, & Sun, 2016).

Another direction would be to explore conceptually different models. One interesting class would be semi-supervised (e.g. Assran et al., 2021) and self-supervised (e.g. Donahue, Krähenbühl, & Darrell, 2016; Dumoulin et al., 2016; Ericsson, Gouk, Loy, & Hospedales, 2021) vision models, which move around some of the issues discussed in the previous section on training data. While fully-supervised models such as the VGG model employed are trained to predict “gold standard” labels (which ultimately need to be provided by human annotators) as accurately as possible, self-supervised models are trained on tasks that require no such outside information, but leverage the training data itself (see Chen et al., 2021): To only list some examples, some models slice up images and learn to recover their original relative spatial positions (Doersch, Gupta, & Efros, 2015) or remove parts of images and train the model to fill them up again based on the surrounding context information (Pathak, Krähenbühl, Donahue, Darrell, & Efros, 2016), while others generate “pseudo-categories” by applying various transformations to original images and learn to classify those (Dosovitskiy, Springenberg, Riedmiller, & Brox, 2014) or are trained to find the most similar image to a given target within a set of candidates, with the correct

solution being an altered (“augmented”) version of the same image (Chen, Kornblith, Norouzi, & Hinton, 2020). Semi-supervised on the other hand make use of a small fraction of the available image labels, which they leverage to learn about the unlabelled images (while often also partly relying on self-supervised techniques, Assran et al., 2021; Chen, Kornblith, Swersky, Norouzi, & Hinton, 2020; Zhai, Oliver, Kolesnikov, & Beyer, 2019). As might be expected given the same general goal on the one hand and the vastly different model architectures and training objectives on the other hand, the representations learned by self-, semi-, and fully-supervised models differ in some aspects and show resemblance in others (Battleday et al., 2020; Geirhos, Narayanappa, et al., 2020).

Since human visual experience does not always include feedback about the correct classification of a visual stimulus, and humans do not require vast amounts of labelled data to obtain object representations, taking into account such semi- and self-supervised techniques is a promising venue for further models and studies in the field of cognitive science (see, for example, Luo, Sexton, & Love, 2021). Initial explorations in this direction suggests that these models might capture human behavioral data not better (Geirhos, Narayanappa, et al., 2020) or even considerably worse (Battleday et al., 2020; Peterson et al., 2018) than supervised models. However, as stated by Geirhos, Narayanappa, et al. (2020, p. 1), “we are clearly just at the beginning of what could be called a self-supervised revolution of machine vision”, and future research might very well produce more promising results. On the other hand, the current lower performance of these approaches may also reflect more fundamental issues: Humans don’t only passively perceive their environment, but actively organize their experiences, categorize them, and, differently from other animals, label it, thus arranging this experience by means of a linguistically-informed semantic system. This labelling, which is conceptually in line with supervised learning techniques, has been found to have profound consequences on the representation and organization of the perceptual experience (Luo et al., 2021; Lupyan, 2012; Lupyan, Abdel Rahman, Boroditsky, & Clark, 2020).

Finally, another intriguing approach is the psychologically-motivated fine-tuning of vision-based representations. For example, Peterson et al. (2016, 2018) applied a weighting function on representations obtained from a supervised DCNN. This weighting was trained to optimize human similarity ratings for a training set of images, and then applied to other images outside this training set. These transformed representations outperformed the original representations in predicting human similarity judgments. Note however that this approach adds an entirely outcome-based transformation, inferred from additional behavioral data, to pre-trained DCNN representations.

Other approaches instead augment DCNN representations (or DCNNs themselves) directly by taking into account psychological models of categorization (Battleday et al., 2020; Singh et al., 2020) – which is after all the task the models were trained on. Battleday et al. (2020) rely from the Luce-Shepard choice rule, where the probability of assigning an item to a specific category is a function of the similarities between an item and all different possible categories. They train their model to implicitly re-weight the dimensional values of their vision-based representations when computing these similarities, so that this categorization is optimized. In their study, Battleday et al. (2020) trained and tested this model on a set of ten different categories of animals and vehicles (such as dogs, birds, trucks, or ships), so this model would need to be scaled up to be tested on larger datasets. In a conceptually similar approach, Singh et al. (2020) directly modify a DCNN architecture by channelling the representations in the last layer through such a decision rule to assign probabilities to each category. This model can be trained end-to-end and learn image as well as category representations from raw image inputs, instead of applying a transformation to fixed, pre-trained representations. Even though this model has so far only been tested on the task it was trained to perform (i.e., image categorization), it appears to be a very promising candidate for vision-based representations informed by psychological theory.

Ultimately, when modelling a representational system, one has the choice between general-level representations or representations that are fine-tuned for specific tasks (see Wingfield & Connell, 2022). In our view, there surely are modelling objectives where task-specific fine-tuning and maximizing the overall performance of a model are desirable, for example when creating a software application that performs a given task with maximum accuracy (i.e., an engineering-oriented approach). However, in the scientific enterprise of creating a model of the human representational system, we advocate for a unified general-level representation model instead of multiple representations for the same entity, each of which is fine-tuned for a specific task (see also Günther et al., 2019). We argue that the astounding flexibility of the cognitive system should instead be modelled in the task-specific cognitive processes that retrieve and use these general-level representations (see also Jorge-Botana, Olmos, & Luzón, 2020). Of course, in this framework, this general-level model can and should be parameterized and adjusted to be as high-performing as possible *across tasks* (compare Baroni et al., 2014; Pereira, Gershman, Ritter, & Botvinick, 2016).

Conclusion

We presented *ViSpa* as a data-driven model for vision-based mental representations of individual images and concept prototypes, which are the product of the stimuli experienced by the model. The psychological adequacy of these representations has been successfully validated with a set of large-scale studies employing different experimental paradigms and tasks. We are optimistic that, in addition to their theoretical appeal, the *ViSpa* system and the measures derived from it can serve as valuable tools for the design of experimental material and data analysis across a wide variety of psychological research fields.

Availability of the resources

The set of studies described in this article presents a large-scale evaluation of a resource – vision-based prototype and image representations – that can be useful to researchers from many different fields of psychology and cognitive science, studying mental representations and image processing, or employing word or image material in their experimental studies. In studies employing word material, controlling for linguistic properties such as semantic neighborhood density and semantic similarities with distributional semantic models has become a firmly established practice (e.g., Danguecan & Buchanan, 2016; Mandera et al., 2017; Reilly & Desai, 2017; see also Günther et al., 2019; Jones et al., 2015), while achieving this level of control for image stimuli or for the visual properties of the object denoted by a word has remained notoriously difficult. Relying on ratings to operationalize these variables is problematic, both from an empirical (see for example the results of Study 5) as well as a theoretical perspective (Westbury, 2016).

To make the *ViSpa* system presented here easily accessible to as wide an audience as possible, we release two publicly available tools to conveniently and intuitively work with these models. On the one hand, the models themselves can be downloaded from www.fritzguenther.de/software-resources/vispa. This allows them to be seamlessly integrated with existing toolkits such as the R package *LSAfun* (Günther, Dudschig, & Kaup, 2015). On the other hand, we set up an intuitive web interface at <http://vispa.fritzguenther.de> which includes many of the same functionalities without requiring any programming experience.

Integration in R using *LSAfun*

The PRO as well as IMG spaces are provided in plain text format and .rda format for R (R Core Team, 2021) in an online repository at www.fritzguenther.de/software-resources/vispa, from where they can be downloaded. Users can then interact with these spaces using their preferred software. One

ideal option for this is the R package *LSAfun* (Günther et al., 2015). This package was originally developed for interactions with distributional semantic spaces, but can be used for any vector space containing named vectors (in the PRO space, the vectors names are the image category labels; in the IMG space, the vector names are the file names of the individual images). After loading such a vector space into the R workspace, the *LSAfun* package then provides functionalities such as various forms of similarity computations, neighborhood explorations, plots for areas of the vector spaces, and various other functions. A detailed instruction manual for the package is provided in Günther et al. (2015); video tutorials for the *LSAfun* package in general and how to use it with *ViSpa* in particular are provided at www.fritzguenther.de/software-resources/video_tutorials.

Web interface

As stated by Mandera et al. (2017), although many specialized toolkits for interactions with large-scale vector spaces are available, the required technical skill set is not necessarily part of the curriculum for many potentially interested researchers in various fields. Due to this, user-friendly web interfaces such as the LSA homepage (Dennis, 2007) or *snaut* (Mandera et al., 2017) have been widely popular and well-received. We therefore set up our own web interface modeled after *snaut*, using the computational infrastructure provided at <https://github.com/pmandera/snaut>, which can be accessed at <http://vispa.fritzguenther.de> (see Figure 18 for a screenshot). This interface also allows for similarity computations and neighborhood explorations, and comes with a *Help* page explaining how to use these tools in detail. In addition, the website hosts a searchable database of the ImageNet images (Deng et al., 2009) included in the IMG space (see Figure 18), so that users can see which exact images the model is representing with its vectors.

vispa - vision spaces website

ApplicationHelp

Using vector space: all images and prototypes

★ Server ready

More info

★ NeighboursMatrixPairwise

Note: To compute typicality scores, just calculate the similarity between an image and its own prototype vector (e.g., between LION_001 and LION_PRO).

Calculate

cosine similarity

DANDELION_012:DANDELION_013

Load from a file

Check availability

Usage info

Picture picker

Use this side of the screen to search for images, find out their names, and copy & paste these names into the text box on the left-hand side for comparison.

For more information on the ImageNet categories included under a given label [click here](#).

DANDELIONSearch

DANDELION

DANDELION_012Copy name

DANDELION_013Copy name

Figure 18. Screenshot of the *ViSpa* website at <http://vispa.fritzguenther.de>

References

- Aly, M., Wansard, M., Segovia, F., Yonelinas, A. P., & Bastin, C. (2014). Cortical and subcortical contributions to state-and strength-based perceptual judgments. *Neuropsychologia*, *64*, 145–156.
- Anderson, A. J., Bruni, E., Lopopolo, A., Poesio, M., & Baroni, M. (2015). Reading visually embodied meaning from the brain: Visually grounded computational models decode visual-object mental imagery induced by written text. *NeuroImage*, *120*, 309–322.
- Andrews, M., Frank, S., & Vigliocco, G. (2014). Reconciling embodied and distributional accounts of meaning in language. *Topics in Cognitive Science*, *6*, 359–370.
- Andrews, M., Vigliocco, G., & Vinson, D. (2009). Integrating experiential and distributional data to learn semantic representations. *Psychological Review*, *116*, 463–498.
- Assran, M., Caron, M., Misra, I., Bojanowski, P., Joulin, A., Ballas, N., & Rabbat, M. (2021). Semi-supervised learning of visual features by non-parametrically predicting view assignments with support samples. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (p. 8443-8452).
- Baayen, R. H. (2008). *Analyzing linguistic data: A practical introduction to statistics using R*. Cambridge, UK: Cambridge University Press.
- Baayen, R. H., & Milin, P. (2010). Analyzing reaction times. *International Journal of Psychological Research*, *3*(2), 12–28.
- Baddeley, A. (1992). Working memory. *Science*, *255*(5044), 556–559.
- Baddeley, A. (2000). The episodic buffer: a new component of working memory? *Trends in Cognitive Sciences*, *4*, 417–423.
- Baddeley, A., & Hitch, G. (1974). Working memory. In *Psychology of learning and motivation* (Vol. 8, pp. 47–89). New York, NY: Academic Press.
- Baker, N., Lu, H., Erlikhman, G., & Kellman, P. J. (2018). Deep convolutional networks

- do not classify based on global object shape. *PLoS Computational Biology*, *14*(12), e1006613.
- Baroni, M., Bernardini, S., Ferraresi, A., & Zanchetta, E. (2009). The WaCky wide web: a collection of very large linguistically processed web-crawled corpora. *Language resources and evaluation*, *43*, 209–226.
- Baroni, M., Dinu, G., & Kruszewski, G. (2014). Don’t count, predict! A systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings of ACL 2014* (pp. 238–247). East Stroudsburg, PA: ACL.
- Barsalou, L. W. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences*, *22*, 637–660.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 1–48.
- Battleday, R. M., Peterson, J. C., & Griffiths, T. L. (2020). Capturing human categorization of natural images by combining deep networks and cognitive models. *Nature Communications*, *11*, 5418.
- Battleday, R. M., Peterson, J. C., & Griffiths, T. L. (2021). From convolutional neural networks to models of higher-level cognition (and back again). *Annals of the New York Academy of Sciences*, *1505*, 55–78.
- Benedek, M., Kenett, Y. N., Umdasch, K., Anaki, D., Faust, M., & Neubauer, A. C. (2017). How semantic memory structure and intelligence contribute to creative thought: a network science approach. *Thinking & Reasoning*, *23*, 158–183.
- Beyer, L., Hénaff, O. J., Kolesnikov, A., Zhai, X., & Oord, A. v. d. (2020). Are we done with imagenet? *arXiv preprint arXiv:2006.07159*.
- BNC Consortium. (2007). *The British National Corpus, version 3*. Oxford, UK: Bodleian Libraries. Retrieved from <http://www.natcorp.ox.ac.uk/>
- Borghesani, V., & Piazza, M. (2017). The neuro-cognitive representations of symbols: the case of concrete words. *Neuropsychologia*, *105*, 4–17.

- Bradski, G. (2000). The opencv library. *Dr. Dobb's Journal: Software Tools for the Professional Programmer*, 25(11), 120–123.
- Bruni, E., Tran, N.-K., & Baroni, M. (2014). Multimodal distributional semantics. *Journal of Artificial Intelligence Research*, 49, 1-47.
- Brysbaert, M., & Stevens, M. (2018). Power analysis and effect size in mixed effects models: A tutorial. *Journal of Cognition*, 1(1).
- Brysbaert, M., Warriner, A. B., & Kuperman, V. (2014). Concreteness ratings for 40 thousand generally known English word lemmas. *Behavior Research Methods*, 46, 904–911.
- Buchanan, E. M., Valentine, K. D., & Maxwell, N. P. (2019). English semantic feature production norms: An extended database of 4436 concepts. *Behavior Research Methods*, 51, 1849–1863.
- Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334), 183–186.
- Carr, T. H., McCauley, C., Sperber, R. D., & Parmelee, C. (1982). Words, pictures, and priming: on semantic activation, conscious identification, and the automaticity of information processing. *Journal of Experimental Psychology: Human Perception and Performance*, 8, 757.
- Chatfield, K., Simonyan, K., Vedaldi, A., & Zisserman, A. (2014). Return of the devil in the details: Delving deep into convolutional nets. *arXiv preprint; arXiv:1405.3531*.
- Chen, T., Frankle, J., Chang, S., Liu, S., Zhang, Y., Carbin, M., & Wang, Z. (2021). The lottery tickets hypothesis for supervised and self-supervised pre-training in computer vision models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 16306–16316).
- Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning* (pp. 1597–1607).

- Chen, T., Kornblith, S., Swersky, K., Norouzi, M., & Hinton, G. E. (2020). Big self-supervised models are strong semi-supervised learners. In *Advances in Neural Information Processing Systems* (pp. 22243–22255).
- Cichy, R. M., & Kaiser, D. (2019). Deep neural networks as scientific models. *Trends in Cognitive Sciences*, 23, 305-317.
- Connolly, A. C., Gleitman, L. R., & Thompson-Schill, S. L. (2007). Effect of congenital blindness on the semantic representation of some everyday concepts. *Proceedings of the National Academy of Sciences*, 104, 8241–8246.
- Danguécan, A. N., & Buchanan, L. (2016). Semantic neighborhood effects for abstract versus concrete words. *Frontiers in psychology*, 7, 1034.
- Děchtěrenko, F., Lukavský, J., & Štipl, J. (2021). False memories for scenes using the drm paradigm. *Vision Research*, 178, 48–59.
- De Deyne, S., Kenett, Y. N., Anaki, D., Faust, M., & Navarro, D. (2016). Large-scale network representations of semantics in the mental lexicon. In M. N. Jones (Ed.), *Big Data in Cognitive Science: From Methods to Insights* (pp. 174–202). New York, NY: Routledge.
- De Deyne, S., Navarro, D. J., Collell, G., & Perfors, A. (2021). Visual and affective multimodal models of word meaning in language and mind. *Cognitive Science*, 45, e12922.
- De Groot, F., Koelewijn, T., Huettig, F., & Olivers, C. N. (2016). A stimulus set of words and pictures matched for visual and semantic similarity. *Journal of Cognitive Psychology*, 28, 1–15.
- de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, 47, 1–12.
- de Leeuw, J. R., & Motz, B. A. (2016). Psychophysics in a web browser? comparing response times collected with javascript and psychophysics toolbox in a visual search task. *Behavior Research Methods*, 48, 1–12.

- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). ImageNet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition* (pp. 248–255).
- Dennis, S. (2007). How to Use the LSA Website. In T. K. Landauer, D. S. McNamara, S. Dennis, & W. Kintsch (Eds.), *Handbook of Latent Semantic Analysis* (pp. 57–70). Mahwah, NJ: Erlbaum.
- Deselaers, T., & Ferrari, V. (2011). Visual and semantic similarity in imagenet. In *Ieee computer vision and pattern recognition (cvpr) 2011* (pp. 1777–1784).
- Doersch, C., Gupta, A., & Efros, A. A. (2015). Unsupervised visual representation learning by context prediction. In *Proceedings of the IEEE International Conference on Computer Vision* (pp. 1422–1430).
- Donahue, J., Krähenbühl, P., & Darrell, T. (2016). Adversarial feature learning. *arXiv preprint, arXiv:1605.09782*.
- Dosovitskiy, A., Springenberg, J. T., Riedmiller, M., & Brox, T. (2014). Discriminative unsupervised feature learning with convolutional neural networks. *Advances in Neural Information Processing Systems*, 27, 766–774.
- Dumoulin, V., Belghazi, I., Poole, B., Mastropietro, O., Lamb, A., Arjovsky, M., & Courville, A. (2016). Adversarially learned inference. *arXiv preprint, arXiv:1606.00704*.
- Efron, B., & Tibshirani, R. J. (1994). *An introduction to the bootstrap*. Chapman & Hall/CRC.
- Ehrenstein, W. H., & Ehrenstein, A. (1999). Psychophysical methods. In *Modern techniques in neuroscience research* (pp. 1211–1241). Heidelberg: Springer.
- Elgendy, M. (2020). *Deep learning for vision systems*. Shelter Island, NY: Manning.
- Ericsson, L., Gouk, H., Loy, C. C., & Hospedales, T. M. (2021). Self-supervised representation learning: Introduction, advances and challenges. *arXiv preprint, arXiv:2110.09327*.

- Fasiolo, M., Nedellec, R., Goude, Y., & Wood, S. N. (2018). Scalable visualisation methods for modern Generalized Additive Models. *Arxiv preprint*. Retrieved from <https://arxiv.org/abs/1707.03307>
- Fellbaum, C. (1998). *WordNet: An Electronic Lexical Database*. Cambridge, MA: MIT Press.
- Ferrand, L., New, B., Brysbaert, M., Keuleers, E., Bonin, P., Méot, A., . . . Pallier, C. (2010). The french lexicon project: Lexical decision data for 38,840 french words and 38,840 pseudowords. *Behavior Research Methods*, 42, 488–496.
- Finkelstein, L., Gabrilovich, E., Matias, Y., Rivlin, E., Solan, Z., Wolfman, G., & Ruppin, E. (2001). Placing search in context: The concept revisited. *ACM Transactions on Information Systems*, 20, 116–131.
- Friedman, L., & Wall, M. (2005). Graphical views of suppression and multicollinearity in multiple linear regression. *The American Statistician*, 59, 127–136.
- Geirhos, R., Jacobsen, J.-H., Michaelis, C., Zemel, R., Brendel, W., Bethge, M., & Wichmann, F. A. (2020). Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2, 665–673.
- Geirhos, R., Meding, K., & Wichmann, F. A. (2020). Beyond accuracy: Quantifying trial-by-trial behaviour of CNNs and humans by measuring error consistency. *arXiv preprint, arXiv:2006.16736*.
- Geirhos, R., Narayanappa, K., Mitzkus, B., Bethge, M., Wichmann, F. A., & Brendel, W. (2020). On the surprising similarities between supervised and self-supervised models. *arXiv preprint, arXiv:2010.08377*.
- Geirhos, R., Narayanappa, K., Mitzkus, B., Thieringer, T., Bethge, M., Wichmann, F. A., & Brendel, W. (2021). Partial success in closing the gap between human and machine vision. *arXiv preprint, arXiv:2106.07411*.
- Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F. A., & Brendel, W. (2018). ImageNet-trained CNNs are biased towards texture; increasing shape bias

- improves accuracy and robustness. *arXiv preprint, arXiv:1811.12231*.
- Geirhos, R., Temme, C. R. M., Rauber, J., Schütt, H. H., Bethge, M., & Wichmann, F. A. (2018). Generalisation in humans and deep neural networks. *arXiv preprint, arXiv:1808.08750*.
- Glenberg, A. M., & Robertson, D. A. (2000). Symbol Grounding and Meaning: A Comparison of High-Dimensional and Embodied Theories of Meaning. *Journal of Memory and Language, 43*, 379–401.
- Golan, T., Raju, P. C., & Kriegeskorte, N. (2020). Controversial stimuli: Pitting neural networks against each other as models of human cognition. *Proceedings of the National Academy of Sciences, 117*, 29330–29337.
- Goodfellow, I. J., Shlens, J., & Szegedy, C. (2015). Explaining and harnessing adversarial examples. In *International Conference on Learning Representations (ICLR)*.
- Günther, F., Dudschig, C., & Kaup, B. (2015). LSAfun—An R package for computations based on Latent Semantic Analysis. *Behavior Research Methods, 47*, 930–944.
- Günther, F., Dudschig, C., & Kaup, B. (2016). Latent semantic analysis cosines as a cognitive similarity measure: Evidence from priming studies. *The Quarterly Journal of Experimental Psychology, 69*, 626–653.
- Günther, F., & Marelli, M. (2020). Trying to make it work: Compositional effects in the processing of compound “nonwords”. *Quarterly Journal of Experimental Psychology, 73*, 1082–1091.
- Günther, F., Marelli, M., Tureski, S., & Petilli, M. A. (2022, April). *ViSpa (Vision Spaces): A computer-vision-based representation system for individual images and concept prototypes, with large-scale evaluation*. The Open Science Framework. Retrieved from <https://osf.io/qvw9c> doi: 10.17605/OSF.IO/QVW9C
- Günther, F., Petilli, M. A., & Marelli, M. (2020). Semantic transparency is not invisibility: A computational model of perceptually-grounded conceptual combination in word processing. *Journal of Memory and Language, 112*, 104104.

- Günther, F., Petilli, M. A., Vergallito, A., & Marelli, M. (2020). Images of the unseen: Extrapolating visual representations for abstract and concrete words in a data-driven computational model. *Psychological Research*, Advance online publication. doi: 10.1007/s00426-020-01429-7
- Günther, F., Rinaldi, L., & Marelli, M. (2019). Vector-space models of semantic representation from a cognitive perspective: A discussion of common misconceptions. *Perspectives on Psychological Science*, 14, 1006–1033.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Hebart, M. N., Zheng, C. Y., Pereira, F., & Baker, C. I. (2020). Revealing the multidimensional mental representations of natural objects underlying human similarity judgements. *Nature Human Behaviour*, 4, 1173–1185.
- Hendrycks, D., & Gimpel, K. (2016). A baseline for detecting misclassified and out-of-distribution examples in neural networks. *arXiv preprint, arXiv:1610.02136*.
- Hill, F., Reichart, R., & Korhonen, A. (2015). Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41, 665–695.
- Hollis, G. (2018). Scoring best-worst data in unbalanced many-item designs, with applications to crowdsourcing semantic judgments. *Behavior Research Methods*, 50, 711–729.
- Hollis, G. (2020). The role of number of items per trial in best–worst scaling experiments. *Behavior Research Methods*, 52, 694–722.
- Hollis, G., & Westbury, C. (2018). When is best-worst best? A comparison of best-worst scaling, numeric estimation, and rating scales for collection of semantic norms. *Behavior Research Methods*, 50, 115–133.
- Honke, G., & Kurtz, K. J. (2019). Similarity is as similarity does? a critical inquiry into the effect of thematic association on similarity. *Cognition*, 186, 115–138.
- Hutchison, K. A., Balota, D. A., Neely, J. H., Cortese, M. J., Cohen-Shikora, E. R., Tse,

- C.-S., ... Buchanan, E. (2013). The semantic priming project. *Behavior Research Methods*, 45, 1099–1114.
- Jacobs, R. A., & Bates, C. J. (2019). Comparing the visual representations and performance of humans and deep neural networks. *Current Directions in Psychological Science*, 28(1), 34–39.
- Johns, B. T., & Jones, M. N. (2012). Perceptual Inference Through Global Lexical Similarity. *Topics in Cognitive Science*, 4, 103–120.
- Jones, M. N., Kintsch, W., & Mewhort, D. J. K. (2006). High-dimensional semantic space accounts of priming. *Journal of Memory and Language*, 55, 534–552.
- Jones, M. N., Willits, J., & Dennis, S. (2015). Models of semantic memory. In J. Busemeyer, Z. Wang, J. Townsend, & A. Eidels (Eds.), *Oxford Handbook of Mathematical and Computational Psychology* (pp. 232–254). New York, NY: Oxford University Press.
- Jorge-Botana, G., Olmos, R., & Luzón, J. M. (2020). Bridging the theoretical gap between semantic representation models without the pressure of a ranking: some lessons learnt from LSA. *Cognitive Processing*, 21, 1–21.
- Jozwik, K. M., Kriegeskorte, N., Storrs, K. R., & Mur, M. (2017). Deep convolutional neural networks outperform feature-based but not categorical models in explaining object similarity judgments. *Frontiers in psychology*, 8, 1726.
- Kenett, Y. N., Kenett, D. Y., Ben-Jacob, E., & Faust, M. (2011). Global and local features of semantic networks: Evidence from the hebrew mental lexicon. *PloS one*, 6(8), e23912.
- Kenett, Y. N., Levi, E., Anaki, D., & Faust, M. (2017). The semantic distance task: Quantifying semantic distance with semantic network path length. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43, 1470–1489.
- Keuleers, E., & Balota, D. A. (2015). Megastudies, crowdsourcing, and large datasets in psycholinguistics: An overview of recent developments. *The Quarterly Journal of*

- Experimental Psychology*, 68, 1457–1468.
- Keuleers, E., & Brysbaert, M. (2010). Wuggy: a multilingual pseudoword generator. *Behavior Research Methods*, 42, 627–633.
- Keuleers, E., Diependaele, K., & Brysbaert, M. (2010). Practice effects in large-scale visual word recognition studies: A lexical decision study on 14,000 dutch mono-and disyllabic words and nonwords. *Frontiers in psychology*, 1, 174.
- Keuleers, E., Lacey, P., Rastle, K., & Brysbaert, M. (2012). The British Lexicon Project: Lexical decision data for 28,730 monosyllabic and disyllabic English words. *Behavior Research Methods*, 44, 287–304.
- Kim, J. S., Elli, G., & Bedny, M. (2019). Furry hippos and scaly sharks: Knowledge of animal appearance among sighted and blind adults. *PsyArXiv preprint*. Retrieved from <https://doi.org/10.31234/osf.io/hw5pm>
- Kotlerman, L., Dagan, I., Szpektor, I., & Zhitomirsky-Geffet, M. (2010). Directional distributional similarity for lexical inference. *Natural Language Engineering*, 16, 359–389.
- Kriegeskorte, N. (2015). Deep neural networks: a new framework for modeling biological vision and brain information processing. *Annual Review of Vision Science*, 1, 417–446.
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet Classification with Deep Convolutional Neural Networks. In F. Pereira, C. J. C. Burges, L. Bottou, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems 25* (pp. 1097–1105).
- Kubilius, J., Bracci, S., & Op de Beeck, H. P. (2016). Deep neural networks as a computational model for human shape sensitivity. *PLoS Computational Biology*, 12(4), e1004896.
- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2017). lmerTest package: Tests in linear mixed effects models. *Journal of Statistical Software*, 82(13), 1–26.

- Lake, B. M., Zaremba, W., Fergus, R., & Gureckis, T. M. (2015). Deep neural networks predict category typicality ratings for images. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 1243–1248).
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato’s problem: The Latent Semantic Analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*, 211–240.
- Lashley, K. S. (1951). The problem of serial order in behavior. In L. A. Jeffress (Ed.), *Cerebral mechanisms in behavior* (pp. 112–131). New York, NY: Wiley.
- Lazaridou, A., Marelli, M., & Baroni, M. (2017). Multimodal word meaning induction from minimal exposure to natural text. *Cognitive Science*, *41*, 677–705.
- Lazaridou, A., Pham, N. T., & Baroni, M. (2015). Combining language and vision with a multimodal skip-gram model. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies* (pp. 153–163). East Stroudsburg, PA.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*, 436–444.
- Lenci, A., & Benotto, G. (2012). Identifying hypernyms in distributional semantic spaces. In *Proceedings of *SEM* (pp. 75–79).
- Lenci, A., Sahlgren, M., Jeuniaux, P., Gyllensten, A. C., & Miliani, M. (2021). A comprehensive comparative evaluation and analysis of distributional semantic models. *arXiv preprint; arXiv:2105.09825*.
- Lewis, M., Zettersten, M., & Lupyan, G. (2019). Distributional semantics as a source of visual knowledge. *Proceedings of the National Academy of Sciences*, *116*, 19237–19238.
- Lindsay, G. W. (2021). Convolutional neural networks as a model of the visual system: Past, present, and future. *Journal of Cognitive Neuroscience*, *33*, 2017–2031.
- Loken, E., & Gelman, A. (2017). Measurement error and the replication crisis. *Science*, *355*(6325), 584–585.

- Louwerse, M. M. (2011). Symbol interdependency in symbolic and embodied cognition. *Topics in Cognitive Science*, 3, 273–302.
- Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). Sustain: a network model of category learning. *Psychological Review*, 111(2), 309–332.
- Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60, 91–110.
- Lucas, M. (2000). Semantic priming without association. *Psychonomic Bulletin & Review*, 7, 618–630.
- Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instrumentation, and Computers*, 28, 201–208.
- Luo, X., Sexton, N. J., & Love, B. C. (2021). A deep learning account of how language affects thought. *Language, Cognition and Neuroscience*, Advance online publication. doi: 10.1080/23273798.2021.2001023
- Lupyan, G. (2012). Linguistically modulated perception and cognition: The label-feedback hypothesis. *Frontiers in psychology*, 3, 54.
- Lupyan, G., Abdel Rahman, R., Boroditsky, L., & Clark, A. (2020). Effects of language on visual perception. *Trends in Cognitive Sciences*, 24, 930–944.
- Lynott, D., Connell, L., Brysbaert, M., Brand, J., & Carney, J. (2020). The lancaster sensorimotor norms: Multidimensional measures of perceptual and action strength for 40,000 english words. *Behavior Research Methods*, 52, 1271–1291.
- Mahendran, A., & Vedaldi, A. (2015). Understanding deep image representations by inverting them. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (pp. 5188–5196).
- Majaj, N. J., & Pelli, D. G. (2018). Deep learning—using machine learning to study biological vision. *Journal of Vision*, 18(13), 2.
- Mandera, P., Keuleers, E., & Brysbaert, M. (2017). Explaining human performance in

- psycholinguistic tasks with models of semantic similarity based on prediction and counting: A review and empirical validation. *Journal of Memory and Language*, *92*, 57–78.
- Mandera, P., Keuleers, E., & Brysbaert, M. (2020). Recognition times for 62 thousand english words: Data from the english crowdsourcing project. *Behavior Research Methods*, *52*, 741–760.
- McClelland, J. L., Rumelhart, D. E., & The PDP Research Group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2: Psychological and Biological Models). Cambridge, MA: MIT Press.
- McKinley, S. C., & Nosofsky, R. M. (1995). Investigations of exemplar and decision bound models in large, ill-defined category structures. *Journal of Experimental Psychology: Human Perception and Performance*, *21*, 128–148.
- McRae, K., De Sa, V. R., & Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General*, *126*, 99–130.
- Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, *90*, 227–234.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv:1301.3781v3*.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems (NIPS) 2013* (pp. 3136–3144). Red Hook, NY: Curran Associates.
- Miller, G. A. (1995). WordNet: a lexical database for English. *Communications of the ACM*, *38*(11), 39–41.

- Murphy, G. L. (2002). *The Big Book of Concepts*. Cambridge, MA: MIT Press.
- Neumann, P. G. (1974). An attribute frequency model for the abstraction of prototypes. *Memory & Cognition*, 2, 241–248.
- Paivio, A. (1986). *Mental representations: A dual coding approach*. Oxford, UK: Oxford University Press.
- Paivio, A. (1991). Dual coding theory: Retrospect and current status. *Canadian Journal of Psychology/Revue canadienne de psychologie*, 45, 255–287.
- Palan, S., & Schitter, C. (2018). Prolific.ac – a subject pool for online experiments. *Journal of Behavioral and Experimental Finance*, 17, 22–27.
- Pathak, D., Krähenbühl, P., Donahue, J., Darrell, T., & Efros, A. A. (2016). Context encoders: Feature learning by inpainting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 2536–2544).
- Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of Experimental Social Psychology*, 70, 153–163.
- Pelli, D. G., & Farell, B. (1995). Psychophysical methods. In M. Bass, E. W. V. Stryland, D. R. Williams, & W. L. Wolfe (Eds.), *Handbook of optics* (2nd ed., pp. 29.21–29.13). New York: McGraw-Hill.
- Pereira, F., Gershman, S., Ritter, S., & Botvinick, M. (2016). A comparative evaluation of off-the-shelf distributed semantic representations for modelling behavioural data. *Cognitive Neuropsychology*, 33, 175–190.
- Peterson, J. C., Abbott, J. T., & Griffiths, T. L. (2016). Adapting deep network features to capture psychological representations. *arXiv preprint, arXiv:1608.02164*.
- Peterson, J. C., Abbott, J. T., & Griffiths, T. L. (2018). Evaluating (and improving) the correspondence between deep neural networks and human representations. *Cognitive Science*, 42, 2648–2669.
- Petilli, M. A., Günther, F., Vergallito, A., Ciapparelli, M., & Marelli, M. (2021).

- Data-driven computational models reveal perceptual simulation in word comprehension. *Journal of Memory and Language*, 117, 104194.
- Pezzulo, G., Barsalou, L. W., Cangelosi, A., Fischer, M. H., McRae, K., & Spivey, M. (2013). Computational grounded cognition: a new alliance between grounded cognition and computational modeling. *Frontiers in psychology*, 3, 612.
- Pezzulo, G., Barsalou, L. W., Cangelosi, A., Fischer, M. H., Spivey, M., & McRae, K. (2011). The mechanics of embodiment: A dialog on embodiment and computational modeling. *Frontiers in psychology*, 2, 5.
- Pitt, D. (2020). Mental Representation. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy* (Spring 2020 ed.). Metaphysics Research Lab, Stanford University.
<https://plato.stanford.edu/archives/spr2020/entries/mental-representation/>.
- Posner, M. I. (1970). Abstraction and the process of recognition. In *Psychology of learning and motivation* (Vol. 3, pp. 43–100). Elsevier.
- R Core Team. (2021). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from <https://www.R-project.org/>
- Rawat, W., & Wang, Z. (2017). Deep convolutional neural networks for image classification: A comprehensive review. *Neural Computation*, 29, 2352–2449.
- Reilly, M., & Desai, R. H. (2017). Effects of semantic neighborhood density in abstract and concrete words. *Cognition*, 169, 46–53.
- Rips, L. J., Shoben, E. J., & Smith, E. E. (1973). Semantic distance and the verification of semantic relations. *Journal of Verbal Learning and Verbal Behavior*, 12, 1–20.
- Roads, B. D., & Love, B. C. (2021). Enriching imagenet with human similarity judgments and psychological embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 3547–3557).
- Rogers, T. T., Lambon Ralph, M. A., Garrard, P., Bozeat, S., McClelland, J. L., Hodges, J. R., & Patterson, K. (2004). Structure and deterioration of semantic memory: a

- neuropsychological and computational investigation. *Psychological Review*, 111, 205–235.
- Rosch, E. (1978). Principles of Categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and Categorization* (pp. 27–46). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573–605.
- Rosenfeld, A., Solbach, M. D., & Tsotsos, J. K. (2018). Totally looks like-how humans compare, compared to machines. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops* (pp. 1961–1964).
- Rubenstein, H., & Goodenough, J. (1965). Contextual correlates of synonymy. *Communications of the ACM*, 8, 627–633.
- Rumelhart, D. E., McClelland, J. L., & The PDP Research Group. (1986). *Parallel distributed processing: Explorations in the Microstructure of Cognition* (Vol. 1: Foundations). Cambridge, MA: MIT Press.
- Sadeghi, Z., McClelland, J. L., & Hoffman, P. (2015). You shall know an object by the company it keeps: An investigation of semantic representations derived from object co-occurrence in visual scenes. *Neuropsychologia*, 76, 52–61.
- Sahlgren, M. (2008). The Distributional Hypothesis. *Rivista di Linguista*, 20, 33–53.
- Serre, T. (2019). Deep learning: the good, the bad, and the ugly. *Annual Review of Vision Science*, 5, 399–426.
- Shankar, S., Halpern, Y., Breck, E., Atwood, J., Wilson, J., & Sculley, D. (2017). No classification without representation: Assessing geodiversity issues in open data sets for the developing world. *arXiv preprint, arXiv:1711.08536*.
- Simonyan, K., & Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint, arXiv:1409.1556*.
- Singer, J., Seeliger, K., Kietzmann, T. C., & Hebart, M. N. (2021). From photos to sketches-how humans and deep neural networks process objects across different levels

- of visual abstraction. *PsyArXiv preprint*. Retrieved from <https://psyarxiv.com/xg2uy/>
- Singh, P., Peterson, J. C., Battleday, R. M., & Griffiths, T. L. (2020). End-to-end deep prototype and exemplar models for predicting human behavior. *arXiv preprint, arXiv:2007.08723*.
- Smith, E. E., & Medin, D. L. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- Smith, E. R. (1998). Mental representation and memory. In *The Handbook of Social Psychology* (Vol. 2, pp. 391–445). New York, NY: McGraw-Hill.
- Sperber, R. D., McCauley, C., Ragain, R. D., & Weil, C. M. (1979). Semantic priming effects on picture and word processing. *Memory & Cognition*, 7, 339–345.
- Stojanoski, B., & Cusack, R. (2014). Time to wave good-bye to phase scrambling: creating controlled scrambled images using diffeomorphic transformations. *Journal of Vision*, 14(12), 6.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., ... Rabinovich, A. (2015). Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- The ImageMagick Development Team. (2021). *ImageMagick*. Retrieved from <https://imagemagick.org>
- Torralba, A., & Efros, A. A. (2011). Unbiased look at dataset bias. In *24th IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2011)* (pp. 1521–1528).
- Tulving, E., & Schacter, D. L. (1990). Priming and human memory systems. *Science*, 247(4940), 301–306.
- van Dyck, L. E., & Gruber, W. R. (2020). Seeing eye-to-eye? a comparison of object recognition performance in humans and deep convolutional neural networks under image manipulation. *arXiv preprint, arXiv:2007.06294*.

- Vedaldi, A., & Lenc, K. (2015). Matconvnet: Convolutional neural networks for Matlab. In *Proceedings of the 23rd ACM international conference on Multimedia* (pp. 689–692).
- Vigliocco, G., Vinson, D. P., Lewis, W., & Garrett, M. F. (2004). Representing the meanings of object and action words: The featural and unitary semantic space hypothesis. *Cognitive Psychology*, 48, 422–488.
- Westbury, C. (2016). Pay no attention to that man behind the curtain. *The Mental Lexicon*, 11, 350–374.
- Wingfield, C., & Connell, L. (2022). Understanding the role of linguistic distributional knowledge in cognition. *Language, Cognition and Neuroscience*, Advance online publication. doi: 10.1080/23273798.2022.2069278
- Wolff, P., & Holmes, K. J. (2011). Linguistic relativity. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2, 253–265.
- Wood, S. N. (2006). *Generalized Additive Models: An Introduction with R*. Boca Raton, FL: CRC Press.
- Wood, S. N. (2015). mgcv: Mixed GAM Computation Vehicle with GCV/AIC/REML Smoothness Estimation [Computer software manual]. Retrieved from <http://CRAN.R-project.org/package=mgcv> (R package version 1.8-7)
- Wood, S. N., Pya, N., & Säfken, B. (2016). Smoothing parameter and model selection for general smooth models. *Journal of the American Statistical Association*, 111, 1548–1563.
- Zeiler, M. D., & Fergus, R. (2014). Visualizing and understanding convolutional networks. In *European conference on computer vision* (pp. 818–833).
- Zhai, X., Oliver, A., Kolesnikov, A., & Beyer, L. (2019). S4l: Self-supervised semi-supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 1476–1485).
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., & Wang, O. (2018). The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE*

Conference on Computer Vision and Pattern Recognition (pp. 586–595).

Zwaan, R. A., & Madden, C. J. (2005). Embodied sentence comprehension. In D. Pecher & R. A. Zwaan (Eds.), *Grounding cognition: The role of action and perception in memory, language, and thinking* (pp. 224–245). Cambridge, UK: Cambridge University Press.

Supplementary Material: Image sources

- Figure 2** <https://pixabay.com/photos/handgun-pistol-firearm-gun-weapon-3149414/>
<https://pixabay.com/vectors/gun-weapon-pistol-handgun-army-161223/>
<https://pixabay.com/photos/handgun-firearms-pistol-gun-weapon-231696/>
<https://pixabay.com/photos/glock-gun-pistol-handgun-weapon-2424292/>
<https://pixabay.com/photos/smith-and-wesson-gun-handgun-smith-938834/>
<https://pixabay.com/illustrations/gun-pistol-weapon-3d-war-warrior-5517424/>
<https://pixabay.com/photos/cordless-screwdriver-screwdriver-2239493/>
<https://pixabay.com/photos/battery-construction-cordless-drill-19983/>
<https://pixabay.com/vectors/power-drill-drill-hammer-drill-154903/>
<https://pixabay.com/photos/bit-construction-cordless-diy-5025182/>
<https://pixabay.com/photos/drill-tool-screwdriver-battery-3650994/>
<https://pixabay.com/vectors/drill-electric-screwdriver-tool-1295561/>
<https://pixabay.com/vectors/screwdriver-tool-spiky-blue-159687/>
<https://pixabay.com/vectors/screwdriver-tool-construction-work-310406/>
<https://pixabay.com/vectors/screwdriver-tools-driving-removing-35126/>
<https://pixabay.com/vectors/screwdriver-tool-maintenance-handy-145111/>
<https://pixabay.com/vectors/head-tool-screwdriver-hardware-295175/>
<https://pixabay.com/photos/tester-screwdriver-check-over-280806/>
- Figure 8** <https://pixabay.com/photos/brown-bear-bear-animal-predator-6993292/>
<https://pixabay.com/photos/makeup-make-up-colors-brushes-1289325/>
<https://pixabay.com/photos/tools-box-workshop-1595429/>
<https://pixabay.com/photos/currants-berry-currant-bush-fruit-3539264/>
<https://pixabay.com/photos/monks-praying-prayer-bangkok-asia-1822569/>
<https://pixabay.com/photos/grain-women-tear-work-food-woman-4443453/>
<https://pixabay.com/photos/food-meal-lunch-dish-tasty-3443871/>
<https://pixabay.com/photos/room-tables-chairs-glasses-meal-186641/>
- Figure 9** <https://pixabay.com/photos/raspberries-fruits-berries-food-3798021/>
<https://pixabay.com/photos/raspberries-red-fruits-berries-1239435/>
<https://pixabay.com/photos/surgery-hospital-doctor-care-1822458/>
<https://pixabay.com/photos/medical-operation-surgery-doctor-5051144/>
<https://pixabay.com/photos/cheetah-feline-wild-cat-fast-5529806/>

- <https://pixabay.com/photos/audio-background-black-business-4844443/>
<https://pixabay.com/photos/firearm-pistol-weapon-violence-3050301/>
<https://pixabay.com/photos/monks-praying-prayer-bangkok-asia-1822569/>
- Figure 11** <https://pixabay.com/photos/roman-mosaic-mosaic-british-museum-252298/>
<https://pixabay.com/photos/strip-blue-fountain-water-bird-197885/>
<https://pixabay.com/photos/ancient-architecture-art-background-21586/>
<https://pixabay.com/photos/mosaic-blue-green-background-wall-1344858/>
<https://pixabay.com/photos/jordan-nebo-mountain-cunt-moses-4158437/>
- Figure 12** <https://www.flickr.com/photos/61831540@N00/2661735143> (user Phil)
<https://pixabay.com/photos/dolphins-florida-ocean-sea-nature-2597580/>
<https://pixabay.com/photos/sea-dolphin-blue-underwater-diving-1677639/>
<https://pixabay.com/photos/dolphin-fish-ocean-sea-animal-603163/>
<https://pixabay.com/photos/dolphin-trees-zoo-nuernberg-water-4509296/>
<https://pixabay.com/photos/elder-blossoms-syrup-bottles-vial-2381649/>
<https://pixabay.com/photos/salad-melon-paprika-spring-onion-3470732/>
<https://pixabay.com/photos/tomatosoup-tomato-soup-sauce-482403/>
<https://pixabay.com/photos/soup-bowl-china-oak-table-dinner-2268278/>
<https://pixabay.com/photos/lemon-lemon-tree-fruit-3496250/>
<https://pixabay.com/photos/glass-water-reflection-956263/>
<https://pixabay.com/photos/lemon-tree-lemon-tree-fruit-citrus-2111048/>
- Figure 14** <https://pixabay.com/photos/raspberries-fruits-berries-food-3798021/>
<https://pixabay.com/photos/raspberries-red-fruits-berries-1239435/>
- Figure 16** <https://pixabay.com/photos/bird-parrot-exotic-colorful-5124424/>
<https://pixabay.com/photos/shark-underwater-sea-fish-ocean-4676799/>
<https://pixabay.com/photos/apartment-architecture-brick-1866634/>
<https://pixabay.com/photos/parrot-bird-colorful-green-red-620345/>
- Figure 18** <https://pixabay.com/photos/dandelion-taraxacum-composites-4119846/>
<https://pixabay.com/photos/dandelion-flower-nature-seeds-5161504/>