

# Subitizing with Variational Autoencoders

Rijnder Wever, Tom F.H. Runia 

University of Amsterdam, Intelligent Sensory Information Systems

**Abstract.** Numerosity, the number of objects in a set, is a basic property of a given visual scene. Many animals develop the perceptual ability to *subitize*: the near-instantaneous identification of the numerosity in small sets of visual items. In computer vision, it has been shown that numerosity emerges as a statistical property in neural networks during unsupervised learning from simple synthetic images. In this work, we focus on more complex natural images using unsupervised hierarchical neural networks. Specifically, we show that *variational autoencoders* are able to spontaneously perform subitizing after training without supervision on a large amount of images from the Salient Object Subitizing dataset. While our method is unable to outperform supervised convolutional networks for subitizing, we observe that the networks learn to encode numerosity as a basic visual property. Moreover, we find that the learned representations are likely invariant to object area; an observation in alignment with studies on biological neural networks in cognitive neuroscience.

**Keywords:** Object counting · Numerosity · Variational autoencoders

## 1 Introduction

The ability to answer the question “How many?” is an important capability of our visual recognition system. Animals use *visual number sense* to rank, quantify and label objects in a scene [29]. There is evidence [28, 32, 2] that the human brain uses at least two distinct mechanisms for non-verbal representation of number: one for large quantity *estimation* and a *subitizing* faculty for near-instantaneous identification of a small number of objects (1-4). In this work, we propose a brain-inspired approach for learning to subitize from large image datasets.

The concepts of visual number sense and instance counting are well studied in computer vision [24, 36, 45, 27, 1, 41, 43, 13, 4]. Most recent work relies on supervised learning using deep convolutional neural networks (CNNs) [22] to assess the instance count in a given scene. While existing methods perform admirably on various counting tasks, the number of visual classes and the availability of labeled image datasets is decisive for their performance. Motivated by these observations and our visual cognition system, this paper explores hierarchical representation learning for visual numerosity in an *unsupervised* setting.

Our work is inspired by the observation of Stoianov and Zorzi [38] that visual numerosity emerges as a statistical property of images in artificial neural networks trained in an unsupervised manner. Specifically, the authors train Restricted Boltzmann Machines (RBMs) on synthetic images containing a random number of objects and show that neural response distributions correlate with number discriminability. Their observations are intriguing but the simple synthetic images



**Fig. 1.** Example images from the Salient Object Subitizing dataset [45]. Although the ability to *subitize* should allow people to identify the number of instances in each image at a glance, these scenes pose a challenge to computer vision models due to variety in appearance, saliency ambiguity, scene clutter and occlusion.

do not capture the complexity of natural visual scenes. In this work, we focus on unsupervised learning of numerosity representations from natural images containing diverse object classes (see example images in Fig. 1).

The contributions of this work are the following. We explore the emergence of visual number sense in deep networks trained in an unsupervised setting on natural images. Specifically, we propose the use of variational autoencoders with both the encoder and decoder parametrized as CNNs to effectively handle complex images and maintain spatial organization. For optimization, we include the recently proposed feature perceptual loss [14] instead of the pixel-to-pixel distance metric to aid representation learning. Finally, we present preliminary quantitative and qualitative results on unsupervised representation learning for numerosity from the Salient Object Subitizing dataset [45].

## 2 Related Work

**Numerical Cognition.** Non-verbal numerical competence is implicitly developed across humans [21, 9, 10, 7] and animal species [28, 12, 6]. These abilities possibly arise from numerosity being an integral part of the sensory world [40]. Interestingly, humans have developed the ability to *subitize* [18, 19] for near-instantaneous numerosity identification of small visual sets (1-4 items). The near-instantaneous character of subitizing and its spontaneous neural development are possibly caused by the visual system’s limited but automatic capability to process spatial configurations of salient objects [32, 17, 7]. Whereas visual number sense relates to object area and density, the neural responses of numerosity-selective cognitive systems were found to be invariant to all visual features except numerosity [28, 12]. Furthermore, studies on cognitive neuroscience have shown that the perception of number functions independently from mathematical reasoning [12, 33]. All these findings suggest that visual number sense is a perceptual property that emerges directly from the visual sensory input.

**Numerosity in Computer Vision.** Instance counting in visual scenes has received substantial interest from the deep vision community, notably in object counting [24, 4, 13, 27, 30], crowd-size estimation [43, 15], animal population estimation [1] and video repetition [25, 34]. The shared similarity between most of these recent works is their use of CNNs for supervised representation learning

from large image datasets. Of these approaches, the recent work of Zhang et al. [45] is most similar to ours as we also evaluate on the task of instance counting and use their Salient Object Subitizing dataset. While these methods are effective on specific domains, they require large amounts of labeled data and are limited to a predefined set of visual classes and numerosity range. Therefore, we here study brain-inspired unsupervised representation learning for visual number sense.

Stoianov and Zorzi [38] discovered the emergence of neural populations in artificial neural networks sensitive to numerosity while invariant to object size. Their observations align with object size invariance for visual number sense in the human brain [28]. In this work we emphasize on learning visual number representations from realistic natural images rather than the simple binary images studied in [38]. As a consequence, representation learning becomes significantly more challenging, making RBMs difficult to train. We propose variational autoencoders [20] to learn visual numerosity representations in an unsupervised setting. -

### 3 Methods

Inspired by Stoianov and Zorzi [38] we propose an unsupervised generative model to learn visual numerosity representations from natural and synthetic image datasets. Specifically, we use a variational autoencoder for encoding and reconstructing training images. The underlying principle is that numerosity is a key characteristic in the images and the network learns to encode visual numerosity in the latent representation.

#### 3.1 Variational Autoencoder

We use the original definition of the variational autoencoder (VAE) as introduced by Kingma and Welling [20]. VAEs are among the most popular approaches for unsupervised representation learning due to their generative nature and the fact that the encoder and decoder can be parameterized by neural networks trainable with stochastic gradient descent. For an excellent overview of VAEs we refer the reader to the tutorial of Doersch [8] as we here only outline the core idea.

VAEs learn to map data samples  $X$  to a posterior distribution  $Q(z | X)$  rather than a deterministic latent representation as used in conventional autoencoders. Inputs can be reconstructed by sampling latent vector  $z$  from the posterior distribution and passing it through a decoder network. To make sampling feasible, the posterior distribution is parametrized by a Gaussian distribution with its mean and variance predicted by the encoder. In addition to a reconstruction loss, the posterior  $Q(z | X)$  is regularized with its Kullback–Leibler divergence from a prior distribution  $P(z)$  which is typically also Gaussian with zero mean and unit variance such that the KL divergence can be computed in closed form [20]. Together, the VAE’s objective function is the summation of a reconstruction term (negative log-likelihood of the data) and the KL regularization:

$$\mathcal{L}_{VAE} = E[\log P(X | z)] - \mathcal{D}_{KL}[Q(z | X) || P(z)] \quad (1)$$



**Fig. 2. Top:** original images from the Salient Object Subitizing dataset [45]. **Center:** VAE reconstructions using traditional loss. **Bottom:** VAE reconstructions using feature perceptual loss. Note the improved ability to reconstruct salient objects and contour sharpness, likely beneficial for object subitizing.

We use this formulation with both encoder and decoder parametrized as convolutional neural network to learn visual representations from a large collection of images displaying scenes with a varying number of salient objects.

### 3.2 Feature Perceptual Loss

VAEs are known to produce blurry reconstructions [11]. In our preliminary experiments we observed difficulties with reconstructing multiple salient objects, negatively affecting the ability to subitize. Therefore, we employ the recent feature perceptual loss of Hou et al. [14] which uses intermediate layer representations in the objective function of the autoencoder. The authors use a VGG-19 network [37] pretrained on ImageNet [35] denoted as  $\Phi$  and define a set of layers  $l_i \in L$  for computing the perceptual loss. Specifically, during the training the mean squared error between the hidden representations of input image  $X$  and the reconstruction  $\tilde{X}$  is added to the loss for the predefined layers:

$$\mathcal{L}_{rec}^L = \sum_{l \in L} \text{MSE}(\Phi(X)^l, \Phi(\tilde{X})^l) \quad (2)$$

The intuition is that responses of layers  $l_i$  should be retained in reconstruction  $\tilde{x}$  as they represent important visual characteristics. Following their recommendations and our own findings, we use  $L = \{\text{relu1\_1}, \text{relu2\_1}, \text{relu3\_1}\}$  from the pretrained VGG-19 network to compute the loss. The feature perceptual loss and original VAE objective are weighed according to  $\mathcal{L}_{total} = \alpha \mathcal{L}_{kl} + \beta \mathcal{L}_{rec}$  in which  $\alpha$  and  $\beta$  are hyperparameters. We found this extension to our autoencoder to improve visual saliency and representation learning compared to plain pixel-by-pixel reconstruction loss (see Fig. 2 for a visual comparison).

## 4 Experiments

### 4.1 Datasets

**Salient Object Subitizing Dataset.** Proposed by Zhang et al. [45], the Salient Object Subitizing (SOS) dataset contains 14K images for the purpose of instance counting. The images originate from MS-COCO [26], ImageNet [35] and SUN [42]. Each image is annotated with an instance count label: 0, 1, 2, 3 or 4+ salient objects. The final image collection is biased towards centered dominant objects and backgrounds scenes as the authors observe [45]. In practice, the class imbalance may pose training difficulties. See Fig. 1 for some examples.

**Synthetic Data.** To counter class imbalance and increase dataset size we follow Zhang et al. [45] in pretraining our model with synthetic images and gradually adding images from the SOS dataset. The images are synthesized by cut-pasting objects from the THUS10000 dataset [5] onto backgrounds from the SUN dataset [42]. Following [45] we apply random image transforms to each object to increase diversity in appearance. For example images we refer to Fig. 7 in [45].

### 4.2 Implementation Details

**Network Architecture.** Our models are implemented in PyTorch. [31]. The VAE’s encoder and decoder are parameterized as CNNs. Denoting a convolutional layer as  $N @ K_w \times K_h - S$  with  $N$  filters of size  $K_w \times K_h$  and stride  $S$ , the encoder architecture is as follows:  $64 @ 4 \times 4 - 1 \rightarrow 64 @ 4 \times 4 - 2 \rightarrow 128 @ 4 \times 4 - 2 \rightarrow 768 @ 4 \times 4 - 2$ . Final spatial features are fed in two fully-connected layers encoding the  $\mu$  and  $\Sigma$  parameters of the posterior distribution for sampling latent vectors using the reparametrization trick [20]. The  $\Sigma$  layer uses softplus activation to ensure a positive output. The decoder uses transposed convolutions [44] to upsample latent representations and is implemented by the following architecture:  $768 @ 3 \times 3 - 1 \rightarrow 768 @ 3 \times 3 - 2 \rightarrow 256 @ 3 \times 3 - 2 \rightarrow 64 @ 3 \times 3 - 2$ . All convolutional blocks are followed by Leaky ReLU activation and batch normalization [16]. The feature perceptual loss parameters are set to  $\alpha = 1.0$  and  $\beta = 0.03$ . The size of the latent dimension is set to 180.

**Optimization.** Data augmentation of random horizontal flips, crops and color shifts is applied to all images. Preprocessed images are of  $161 \times 161$  size when fed into the network. We warm-up by pretraining on 80K synthetic images and gradually start adding natural images beyond 20 epochs. The initial learning rate is set to 0.0015 and is divided by 5 when the loss on the test set plateaus for more than 4 epochs. The VAE is trained for a total of 140 epochs. To remedy the class imbalance we follow [23] by randomly removing 10% examples from the most frequent classes (and loss weighting for the softmax classifier in Sec. 4.3).

**Table 1.** Comparison of our unsupervised approach to existing supervised approaches for *instance counting* over the Salient Object Subitizing dataset. We report the count average precision (%) over the entire test set. Results from existing methods (rows 2-4) were reported in [45]. Our method is unable to outperform the fully-supervised CNN of [45] but performs well given that our visual representations are trained unsupervised.

Count Label $\rightarrow$	0	1	2	3	4+	mean
Chance	27.5	46.5	18.6	11.7	9.7	22.8
GIST [39]	67.4	65.0	32.3	17.5	24.7	41.4
SIFT+IFV [3]	83.0	68.1	35.1	26.6	38.1	50.1
CNN_FT [45]	93.6	93.8	75.2	58.6	71.6	78.6
VAE + softmax (ours)	76.0	49.0	40.0	27.0	30.0	44.4

### 4.3 Evaluating the Visual Numerosity Representation

In this experiment we evaluate the strength of the representations learned by the VAE in unsupervised setting. As the visual reconstructions are hard to compare and do not conceal the sense visual numerosity, we perform a quantitative comparison with the state-of-the-art [45, 39, 3] by training a softmax classifier on top of the visual representations learned without supervision. The task is to predict the instance count (classification). Specifically, we fix the VAE parameters and feed the latent representations for a given image to the softmax classifier referred to as **VAE + softmax**. The softmax classifier is modeled as a two-layer perceptron with 160 units per layer and ReLU activations. It predicts the instance count from the VAE’s latent representations. We use the SOS train set *with* count labels to minimize cross-entropy loss. Note that this count classifier uses supervision on top of the unsupervised visual representations. We compare the subitizing performance with existing work: handcrafted GIST features with SVM classifier [39], a SIFT-based representation with Improved Fisher Vectors (IFV) [3] and the fully-supervised CNN specifically designed for subitizing [45].

Table 1 reports the performance of both our method and existing work on the SOS test set [45]. The subitizing performance of our softmax classifier is comparable to the performance of the GIFT and SVM classifier [39] and SIFT+IFV [3]. Our method is unable to surpass the fully-supervised CNN of [45]. This is not surprising as their network is significantly larger, pretrained on millions of images from ImageNet and uses full supervision from the SOS count labels. Our visual representations are trained in unsupervised setting without ImageNet pretraining. The quantitative results indicate that unsupervised learning of a representation of visually complex images by a VAE discovered that numerosity in the subitizing range is a key characteristic of natural scenes.

### 4.4 Size-invariant Numerosity Detectors

We investigate whether the learned visual representations for numerosity estimation are invariant to *object area*, as found in [38] for simple synthesized data and



**Fig. 3. Top Left:** original image with a single salient instance from SOS. **Remaining Images:** reconstructions of the VAE by slightly increasing the response at individual dimensions in the latent representation. A single value in the latent space can correspond to multiple image characteristics such as lighting, color, numerosity and object size.

observed in cognitive studies [28, 12]. Our methodology is similar to Stoianov and Zorzi [38]: we search for a relationship between the VAE’s latent representations and both cumulative object area and instance count in synthesized images.

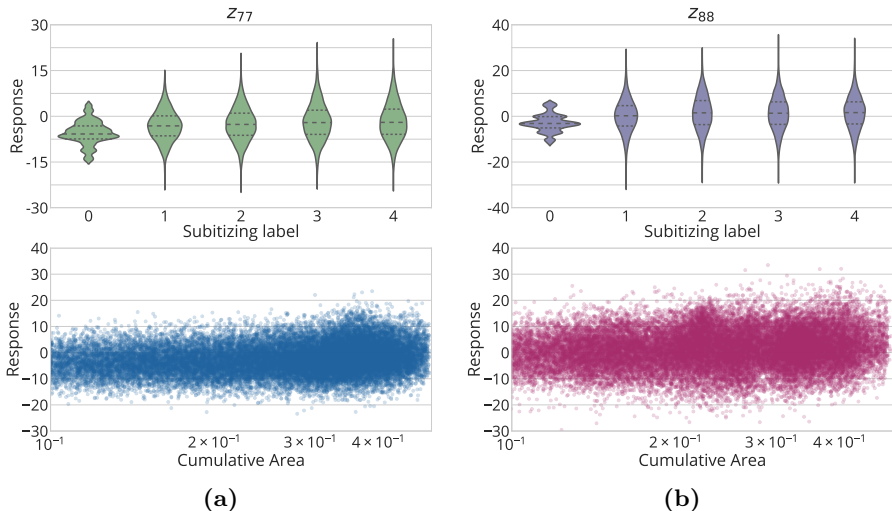
To this end, we create a dataset with synthetic images containing  $N \in [0, 4]$  copies of the same object (with  $N$  sampled uniformly at random) and their corresponding cumulative area values  $A$  (measured in the number of pixels). We fix the VAE parameters and create a set of latent vectors  $Z$  from 35K synthetic images with 15 object classes from the THUS10000 dataset [5]; the size between objects in each image varies modestly. To reduce noise in the representations created by the VAE we make sure objects are not overlapping (interestingly, overlapping objects hinder our brain’s ability to subitize [7]). Although not natural images, their appearance is more diverse than the binary images of [38].

Using all latent vectors  $z$  from  $Z$ , we search for latent dimensions  $z_i$  that serve as either numerosity or area encoders by means of a linear regression suggested in [38]. Following their approach, we fit the following relationship between  $z_i$  and the variables  $A$  and  $N$  across the entire dataset (all variables normalized):

$$z_i = \beta_1 \log(N) + \beta_2 \log(A) + \varepsilon \quad (3)$$

Stoianov and Zorzi [38] formulate two criteria for which an individual dimension  $z_i$  significantly responds to changes in object size or numerosity. First, the dimension should explain at least 10% of the variance ( $R^2 \geq 0.1$ ) in the activity, and secondly, the regression coefficient of the complementary property has an absolute value smaller than 0.1. Here, we slightly loosen the first criteria of [38] by setting the threshold to 5% variance because the visual complexity of our training images is significantly higher. More specifically, in our VAE the individual latent dimensions  $z_i$  can be responsible for encoding more than one visual characteristic. This is observed in Fig. 3, where a slight change in the latent dimension can change more than one visual characteristic in the reconstruction. Due to this fact, the responses of  $z_i$  will inherently be noisier.

From the regression over 35K synthetic images we found one or two reoccurring dimensions that responded to area  $A$  or numerosity  $N$  with regression error  $R^2 > 0.06 \pm 0.020$ . Interestingly, we also found that whenever the regression showed multiple dimensions responding to object area, the two dimensions



**Fig. 4.** Response distribution of two latent dimensions when feeding synthetic images of different subitizing label and cumulative object area. **(a)**  $z_{77}$  responds to numerosity (subitizing label) whereas being invariant to object size ( $R^2 = 0.055$  for the fit of Eq. (3)). **(b)**  $z_{88}$  shows a typical response profile of a dimension sensitive to cumulative object count ( $R^2 = 0.056$ ). The cumulative object area  $A$  is shown on logarithmic scale.

always changed in opposite direction (different sign) which is in agreement with [38]. Therefore, the latent space likely encodes object area and numerosity as independent properties of images, consistent with coding properties of numerosity-selective neurons [12]. Finally, in Fig. 4 we plot the characteristics of response profiles of the dimensions  $z_i$  that were found to encode either cumulative area or visual numerosity. For the area dimension ( $z_{88}$  shown in Fig. 4b), images with either very small or large cumulative area push the mean response distribution significantly upward or downward. On the other hand, for the numerosity encoding dimension ( $z_{77}$  shown in Fig. 4a) the response is more stable. This is evidence for dimension  $z_{77}$  encoding visual number sense while being invariant to object size.

## 5 Conclusion

We have proposed unsupervised representation learning for visual number sense on natural images. Specifically, we propose a convolutional variational autoencoder to learn the concept of number from both synthetic and natural images without supervision. In agreement with previous findings on numerosity in artificial multi-layer perceptrons [38] and biological neuronal populations [12, 28], a representation with the ability to encode numerosity within the subitizing range invariant to object area and appearance has been learned. Therefore, we present additional evidence that the concept of visual number sense emerges as a statistical property in variational autoencoders when presented a set of images displaying a varying number of salient objects.



## Acknowledgements

The authors would like to thank the Intelligent Sensory Information Systems Institute and the Informatics Institute of the University of Amsterdam for their financial contribution to the travel expenses.

## Bibliography

- [1] C. Arteta, V. Lempitsky, and A. Zisserman. Counting in the wild. In *ECCV*, 2016.
- [2] D. Burr and J. Ross. A visual sense of number. *Current biology*, 18(6): 425–428, 2008.
- [3] K. Chatfield, V. S. Lempitsky, A. Vedaldi, and A. Zisserman. The devil is in the details: an evaluation of recent feature encoding methods. In *BMVC*, 2011.
- [4] P. Chattopadhyay, R. Vedantam, R. R. Selvaraju, D. Batra, and D. Parikh. Counting everyday objects in everyday scenes. In *CVPR*, 2017.
- [5] M.-M. Cheng, N. J. Mitra, X. Huang, P. H. Torr, and S.-M. Hu. Global contrast based salient region detection. *PAMI*, 37(3):569–582, 2015.
- [6] H. Davis and R. Pérusse. Numerical competence in animals: Definitional issues, current evidence, and a new research agenda. *Behavioral and Brain Sciences*, 11(4):561–579, 1988.
- [7] S. Dehaene. *The number sense: How the mind creates mathematics*. OUP USA, 2011.
- [8] C. Doersch. Tutorial on variational autoencoders. *arXiv preprint arXiv:1606.05908*, 2016.
- [9] L. Feigenson, S. Dehaene, and E. Spelke. Core systems of number. *Trends in cognitive sciences*, 8(7):307–314, 2004.
- [10] M. C. Franka, D. L. Everettb, E. Fedorenkoa, and E. Gibsona. Number as a cognitive technology: Evidence from pirahã language and cognitionq. *Cognition*, 108:819–824, 2008.
- [11] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio. *Deep learning*, volume 1. MIT press Cambridge, 2016.
- [12] B. M. Harvey, B. P. Klein, N. Petridou, and S. O. Dumoulin. Topographic representation of numerosity in the human parietal cortex. *Science*, 341(6150):1123–1126, 2013.
- [13] S. He, J. Jiao, X. Zhang, G. Han, and R. W. Lau. Delving into salient object subitizing and detection. In *ICCV*, 2017.
- [14] X. Hou, L. Shen, K. Sun, and G. Qiu. Deep feature consistent variational autoencoder. In *WACV*, 2017.
- [15] Y. Hu, H. Chang, F. Nian, Y. Wang, and T. Li. Dense crowd counting from still images with convolutional neural networks. *Journal of Visual Communication and Image Representation*, 38:530–539, 2016.
- [16] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015.

- [17] B. R. Jansen, A. D. Hofman, M. Straatemeier, B. M. Bers, M. E. Raijmakers, and H. L. Maas. The role of pattern recognition in children’s exact enumeration of small numbers. *British Journal of Developmental Psychology*, 32(2): 178–194, 2014.
- [18] W. S. Jevons. The power of numerical discrimination, 1871.
- [19] E. L. Kaufman, M. W. Lord, T. W. Reese, and J. Volkman. The discrimination of visual number. *The American journal of psychology*, 62(4):498–525, 1949.
- [20] D. P. Kingma and M. Welling. Auto-encoding variational bayes. In *ICLR*, 2014.
- [21] G. Lakoff and R. E. Núñez. Where mathematics comes from: How the embodied mind brings mathematics into being. *AMC*, 10:12, 2000.
- [22] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *Nature*, 521(7553):436, 2015.
- [23] G. Lemaître, F. Nogueira, and C. K. Aridas. Imbalanced-learn: A python toolbox to tackle the curse of imbalanced datasets in machine learning. *JMLR*, 18(17):1–5, 2017.
- [24] V. Lempitsky and A. Zisserman. Learning to count objects in images. In *NIPS*, 2010.
- [25] O. Levy and L. Wolf. Live repetition counting. In *ICCV*, 2015.
- [26] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- [27] X. Liu, Z. Wang, J. Feng, and H. Xi. Highway vehicle counting in compressed domain. In *CVPR*, 2016.
- [28] A. Nieder. The neuronal code for number. *Nature Reviews Neuroscience*, 17(6):366–382, 2016.
- [29] A. Nieder and S. Dehaene. Representation of number in the brain. *Annual review of neuroscience*, 32:185–208, 2009.
- [30] M. Noroozi, H. Pirsiavash, and P. Favaro. Representation learning by learning to count. In *ICCV*, 2017.
- [31] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in pytorch. In *NIPS Workshops*, 2017.
- [32] M. Piazza and V. Izard. How humans count: numerosity and the parietal cortex. *The Neuroscientist*, 15(3):261–273, 2009.
- [33] M. Poncet, A. Caramazza, and V. Mazza. Individuation of objects and object parts rely on the same neuronal mechanism. *Scientific reports*, 6: 38434, 2016.
- [34] T. F. H. Runia, C. G. M. Snoek, and A. W. M. Smeulders. Real-world repetition estimation by div, grad and curl. In *CVPR*, June 2018.
- [35] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.
- [36] S. Seguí, O. Pujol, and J. Vitria. Learning to count with deep object features. In *CVPR Workshops*, 2015.

- [37] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [38] I. Stoianov and M. Zorzi. Emergence of a ‘visual number sense’ in hierarchical generative models. *Nature Neuroscience*, 15(2):194, 2012.
- [39] A. Torralba, K. P. Murphy, W. T. Freeman, M. A. Rubin, et al. Context-based vision system for place and object recognition. In *ICCV*, 2003.
- [40] P. Viswanathan and A. Nieder. Neuronal correlates of a visual “sense of number” in primate parietal and prefrontal cortices. *Proceedings of the National Academy of Sciences*, 110(27):11187–11192, 2013.
- [41] E. Walach and L. Wolf. Learning to count with cnn boosting. In *ECCV*, 2016.
- [42] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, 2010.
- [43] F. Xiong, X. Shi, and D.-Y. Yeung. Spatiotemporal modeling for crowd counting in videos. In *ICCV*, 2017.
- [44] M. D. Zeiler, D. Krishnan, G. W. Taylor, and R. Fergus. Deconvolutional networks. In *CVPR*, 2010.
- [45] J. Zhang, S. Ma, M. Sameki, S. Sclaroff, M. Betke, Z. Lin, X. Shen, B. Price, and R. Měch. Salient object subitizing. *IJCV*, 124(2):169–186, 2017.