

Straight to Shapes: Real-time Detection of Encoded Shapes

Saumya Jetley * Michael Sapienza *[†] Stuart Golodetz Philip H.S. Torr
 Department of Engineering Science
 University of Oxford

{sjetley, smg, phst}@robots.ox.ac.uk, m.sapienza@samsung.com

Abstract

Current object detection approaches predict bounding boxes that provide little instance-specific information beyond location, scale and aspect ratio. In this work, we propose to regress directly to objects' shapes in addition to their bounding boxes and categories. It is crucial to find an appropriate shape representation that is compact and decodable, and in which objects can be compared for higher-order concepts such as view similarity, pose variation and occlusion. To achieve this, we use a denoising convolutional auto-encoder to learn a low-dimensional shape embedding space. We place the decoder network after a fast end-to-end deep convolutional network that is trained to regress directly to the shape vectors provided by the auto-encoder. This yields what to the best of our knowledge is the first real-time shape prediction network, running at 35 FPS on a high-end desktop. With higher-order shape reasoning well-integrated into the network pipeline, the network shows the useful practical quality of generalising to unseen categories that are similar to the ones in the training set, something that most existing approaches fail to handle.

1. Introduction

Automatically detecting [13] and delineating [7] object instances in images is a core problem in computer vision, with wide-ranging applications. For example, knowing the individual boundaries of nearby objects can allow robots to grasp them [31]. To help visually-impaired people become more independent, specially-designed glasses can highlight the boundaries of objects with which to interact [18].

For these applications, real-time frame processing is crucial, and the information provided by bounding box predictions [34] or a pixel-wise semantic segmentation of the scene [43] is not enough. A bounding box captures no more than the location, scale and aspect ratio of an object, a fairly coarse representation that provides no information

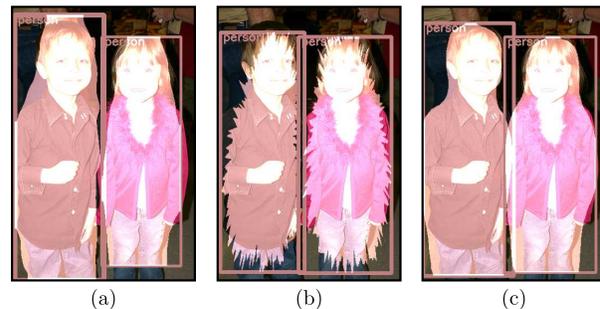


Figure 1: Testing direct regression to shape on a validation image from PASCAL VOC [11] using (a) binary shape masks, (b) radial vectors, and (c) a learned shape embedding.

about the object's boundary. On the other hand, bottom-up pixel-labelling approaches have no explicit notion of local and global object shape. To achieve adherence to local boundaries and impose spatial and appearance consistency in the semantic space, recent works [4, 43] post-process using conditional random fields (CRFs). Such post-processing quickly becomes intractable for the higher-order constructs required to capture object structure, pose and occlusions.

In response to the above-mentioned shortcomings, we implement an embedding space that incorporates notions of object shape, pose and occlusion patterns. A deep regression network is then trained to map input image patches to this embedding space to tease apart an object's category and its shape mask. Fig. 1 shows an example of our network regressing to three different shape representations. Observably, object shape and category are correlated for many classes, so training the learner with a single objective function that includes both terms is mutually-reinforcing.

We propose an embedding space that is learnt on the binary instance masks in a class-agnostic manner. By design, the embedding space is *compact*, *decodable* (supports mapping of binary instance masks in and out of the space), *continuous* (instance masks degrade gracefully around the point of interest in the space) and *interpretable* (one in which we can reason about the similarities between shapes and their corresponding categories). Our formulation is thus able to leverage shape reasoning and extend the prediction of shape

* Authors contributed equally

[†]M. Sapienza performed this research at the University of Oxford, and is currently with Samsung Research America, Mountain View CA.

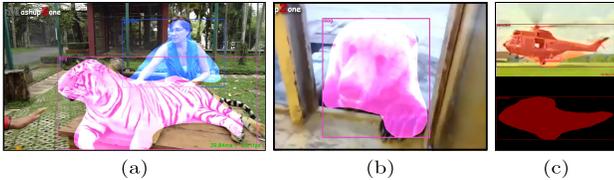


Figure 2: Detecting the shapes of previously-unseen categories in images from YouTube videos: (a) a person strokes a tiger, (b) a bear is fed through a window, and (c) a helicopter crashes into the ground (see supplementary video). The tiger and bear are both detected as class ‘dog’, and the helicopter is detected as ‘aeroplane’. Despite wrong categories, the predicted shape masks capture useful information about the types of objects that are present.

masks to object categories our network has never seen before, but which have similar characteristics to the ones seen during training, like the tiger, bear and helicopter in Fig. 2.

It is also important to note that we tackle the detection of shapes of individual instances in a top-down paradigm with a single sweep of the network. This is in contrast to recent works [8, 2] combining top-down and bottom-up paradigms with various sequential arrangements of bounding-box detection, pixel-wise segmentation and category recognition to achieve instance segmentation. Due to a sequential processing, the error in the above networks is additive. For instance, if bounding-box detection comes first and cuts through an object boundary, pixel-wise labelling inside the box can never retrieve the correct object boundary. Likewise, if detection follows pixel-wise labelling, it is difficult to estimate object boundaries for overlapping instances of the same class.

Contributions. The proposed approach overcomes the aforementioned difficulties by directly regressing to multiple object locations, shapes, and categories, as shown in Fig. 3. Crucial to the process is the learning of a compact and decodable embedding space in which shapes can be described and compared. To this end, we demonstrate the use of a denoising auto-encoder [39] in encoding real-world shape templates. Moreover, the single-sweep processing affords our implementation real-time capabilities. To the best of our knowledge, this is the first real-time shape prediction network, running at 35 FPS on a high-end desktop with an i7-4960 Processor (3.6GHz, 12-core) and a Titan X GPU. Our shape prediction network thus builds upon what is currently possible at the intersection of object detection [33], instance segmentation [17] and joint embeddings [22].

2. Related Work

Importance of shape cues. In the past, shape cues have been successfully used to guide object recognition [3] and localisation [20], by virtue of their ability to capture discriminative, category-specific details. With the advent of deep networks, suitable features best adapted to the task are assumed to be learnt by the network during training [21], by

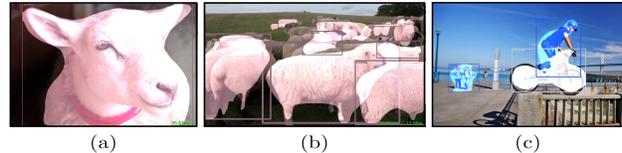


Figure 3: Regressing to single or multiple instances with large intra-category shape variation: consider e.g. the shape of (a) a sheep’s head, in contrast to (b) its body. In (c), we detect multiple overlapping shapes, in this case a person on a bike.

optimising an appropriate loss function. In existing object detection pipelines, this loss is computed on the 4D bounding box specifications [33], ignoring the use of shape to guide learning. By contrast, regressing to detailed shape vectors affords the network a more informed supervision.

Predicting object shape masks. Predicting object shape masks has previously been identified as an important practical challenge [12, 27]. However, these approaches contain independently-optimised processing blocks and have shown limited success, working only for images containing only a single object and/or with high background contrast. More recently, Pinheiro *et al.* [29, 30] have explored the practical benefits of learning to predict object proposals as pixel-wise segmentation masks. However, their networks are composed of disjoint stages for predicting segmentation masks, object locations and object classes. By contrast, and in a spirit similar to YOLO [33], the architecture that we propose for predicting shape embeddings (from which the shape masks can be reconstructed) is simpler and contains only one network, with a single objective.

Extending bounding box detectors to shape. Recent progress in bounding box object detection has been fuelled by the representational power of deep networks. R-CNN [14] made massive leaps in detection accuracy by using a deep network to predict the category of pre-generated bounding box proposals. Soon after, Faster R-CNN [34] emerged. It does away with separate stages for region proposal generation and classification, instead using a region proposal network that shares convolutional features with a detection network, achieving benefits in computational efficiency and accuracy. More recently, the YOLO [33] detection network streamlined object detection by replacing the separate proposal generation and classification stages with regression to spatially-separated bounding boxes and class probabilities. This resulted in real-time object detection rates. Inspired by YOLO, we extend the network to predict shape encodings, achieving similar speeds, but with the advantage of incorporating additional shape information.

Shape embeddings. The simplest embedding space we explore in this work consists of appropriately down-sampled binary masks. Other shape representations are also possible starting points for establishing this space; a comprehensive review of shape representations can be found in [41]. Moreover, inspired by the way in which [28] learned

word embeddings to find a feature space in which similar words are close together, we aim to learn a *shape* embedding in which to reason about the shape of objects. In contrast to [28], however, we seek a lower-dimensional encoding (compared to the dimensions of a binary shape mask) to make detection by regression feasible. We observe the compression, noise-handling and reconstruction capabilities of the learned shape encodings to be significantly superior to the hand-crafted shape representations.

Learned shape embeddings. Due to the recent successes of learned latent representations for small binary images of digits and faces [19], we learn a shape embedding by training an auto-encoder. The Shape Boltzmann Machine (SBM) of [10] uses a deep restricted Boltzmann architecture to model a shape space that is realistic and generalisable. However, the above networks are trained to model shapes and are not optimised for reconstruction. Furthermore, it is important to note that the fully-connected construction of SBM restricts it to small synthetic images with a resolution of around 50×50 . To circumvent this practical limitation, we incorporate convolutional layers, and train the model as a denoising auto-encoder [39]. Thus, in a first, we extend the full benefit of auto-encoders to real-world shape templates. Interestingly, by learning a continuous and decodable shape embedding, we are able to predict masks for categories that have not been observed previously but share close shape similarity with those in the train set (see Fig. 2). We are able to achieve this performance at a fraction of the dimensionality of hand-crafted shape representations.

Our work is also related to that of Li *et al.* [22], which handles single image-to-embedding mappings [37] and assumes that objects in the input image are prominent. By contrast, we regress to multiple shape embeddings, thereby extending the approach to object detection (see Fig. 3). Another important distinction is that the shape embeddings in [22] are hand-crafted and non-decodable, preventing the synthesis of new shapes from arbitrary coordinates in the embedding space and limiting the applications to retrieval with nearest-neighbour search.

Potential applications: instance segmentation and occlusion handling. Recent works [15, 7] on instance segmentation attempt to combine the benefits of the top-down and bottom-up paradigms by performing pixel-wise class labelling within each bounding box prediction. Such bounding box-based approaches err primarily for occluded objects and their boundaries. A bottom-up approach like [5] resorted to post-hoc object-level or category-specific reasoning to combine the separated parts of object instances. Also noteworthy is the work by Dai *et al.* [6], which extends a fully-convolutional network (FCN) [25] to predict position-sensitive instance score maps that then need to be assembled into object instances at the output. Another approach [36] makes use of a recurrent neural network (RNN)

architecture, which is trained with a single loss function in order to predict object instances sequentially. Although this has shown promise on a limited set of object categories, it has difficulty remembering all past predictions, and therefore finding those that it needs to predict next.

Ultimately, describing an object as a set of pixels with an associated category has its limitations. In particular, shapes represented as pixel sets are hard to compare: there is no obvious way of computing *meaningful* distances between high-dimensional object masks. This is unfortunate, because an ability to compare shapes is extremely useful: the shape of an object is not merely an abstract property, but can indicate something fundamental about the object’s capabilities and role. For example, whilst there is significant variety amongst the many legged animals in the world, they are clearly more similar to each other than to an inanimate object such as a lorry, and the shapes involved reflect this. As such, if we can find a representation in which we are able to compare shapes quantitatively, we can use the distance values to provide important clues about objects, even for categories on which we have not trained. For these reasons, we tackle the shape prediction problem as part of an end-to-end pipeline, using an intermediate embedding space that knows about realistic category-level poses, occlusions and shape priors, as learnt from the training data.

3. Proposed Approach

Deep regression network. An object’s shape is one of its fundamental properties, which, when coupled with its location, scale and category, gives rich information about its possible pose and coarse depth, and how one can interact with it. To take advantage of this, we therefore extend the state-of-the-art YOLO [33] object detection network that we mentioned in §2 to regress to not only object locations, confidence scores and conditional probabilities for each category, but also detailed shape encodings, as illustrated in Fig. 4. The details of how we do this are described in §4.

Decodable shape representation. We require a shape embedding that is both compact and decodable for the object shape prediction application. Moreover, we would benefit immensely from a representation that embodies higher-order understanding of object shape, realistic poses and occlusion patterns. For instance, we would like the distances between shapes in the representation to reflect our own understanding of the similarity between shapes. To this end, in §5 we investigate three decodable shape representations: i) downsampled binary masks, ii) radial vectors and iii) learned shape encodings.

In §6, we join the ideas from §4 and §5, and evaluate the merits of our real-time shape detection pipeline.

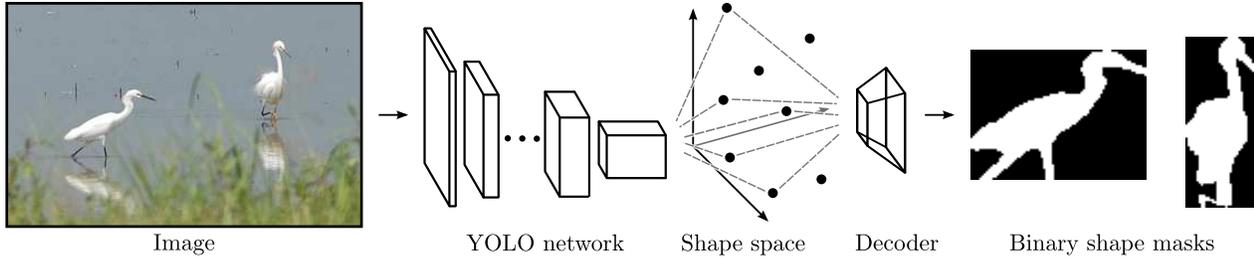


Figure 4: From left, an input image is rescaled and passed through a single network that regresses directly to multiple object hypotheses (category, location, aspect and shape). The learned shape encoding is compact, decodable, continuous and interpretable, allowing the prediction of shapes not seen in the training data, and more meaningful distance comparisons between shapes. Note that predicting encoded shapes enables higher-order reasoning about the pose of the bird (upright, facing right) and how it compares to other birds: this would be impossible from a bounding box alone. Moreover, the current shape space is learned in an unsupervised manner and any annotations superimposed on the shape space would add to the predicted information at test time.

4. Deep Regression Network

This section draws inspiration from the YOLO detection algorithm [33] and extends it for the prediction of shapes. Like YOLO, our network reasons globally about objects in the image and is trained to minimise a single objective function. However, we augment YOLO’s object representation with an encoded vector that denotes an object’s shape.

The YOLO pipeline starts by dividing the input image into an $S \times S$ grid. If the centre of an object lies within a cell, then that cell becomes responsible for detecting that object. Here, each grid cell predicts $B = 2$ sets of boxes, confidence scores and shapes with a dimensionality of N per set, as well as a probability mass function.

The box prediction corresponds to the 4-coordinates that fully specify the tightest fitting bounding box to the object. The confidence score is calculated as the intersection over union (IoU) between the predicted and target box, if a target box exists, and zero otherwise. Each shape encoding represents the shape of an object, independent of its location and aspect ratio. For instance, in order to calculate the binary shape mask that we provide as a target for learning, we take the ground truth object segmentation mask, binarise it and rescale it as part of the encoding process. Overall, if the shape is to be encoded by a 16×16 binary mask, then $N = 1 + 4 + 256$.

Each grid cell also predicts a conditional probability mass function, which, when multiplied by the object confidence score, results in a score reflecting the confidence for the category and its overlap:

$$p(c|o) * p(o) * IoU = p(c) * IoU. \quad (1)$$

As input to our network, we feed in 448×448 images, which are transformed using random rotations, translations, spatial scaling and pixel scaling. The network’s target is augmented by transforming the ground truth shapes with the same geometric transformations as above. The final size of the target vector for a single image becomes

$$D = S \times S \times (N \times B + |C|), \quad (2)$$

where $|C|$ is the number of dataset categories, N includes parameters for the confidence score, minimum bounding box and shape encoding, and $S \times S$ is the total number of cells in the image grid.

For training and inference, we use Darknet [32]. We optimise the sum of squared errors between the predicted network output and a target tensor. The four-component loss function used to train the network can be expressed as $L = L_{\text{box}} + L_{\text{conf}} + L_{\text{shape}} + L_{\text{pmf}}$, where

$$L_{\text{shape}} = \lambda_{\text{shape}} \sum_{i=0}^{S^2} \sum_{j=0}^B \sum_{k=0}^{N-5} \mathbb{1}_{ij}^{\text{obj}} (\tau_{ijk} - \hat{\tau}_{ijk})^2. \quad (3)$$

Here, τ is the target shape representation and $\hat{\tau}$ is the predicted shape representation. Definitions for the other three loss components can be found in [33].

Since a prediction only requires one forward pass through the network, it is very fast at test time (§6). We use non-maximal suppression to reduce the number of predicted overlapping shapes.

5. Decodable Shape Representation

The shape embedding space is a crucial part of the proposed formulation. We believe that the tasks of constructing the embedding space and learning a mapping from the image domain to this embedding space are intrinsically independent and should therefore be treated separately. Moreover, doing so makes the two tasks more tractable [22]. We thus experiment with two hand-crafted representations and one learned shape representation to establish the embedding space, described as follows.

5.1. Downsampled binary shape masks

A very simple shape descriptor can be obtained just by downsampling the full-size binary shape mask. Given a fixed descriptor size $d = k \times k$, we downsize the image using OpenCV’s INTER_AREA resampling approach. For reconstruction, the descriptor can be resized back up again using bicubic interpolation.

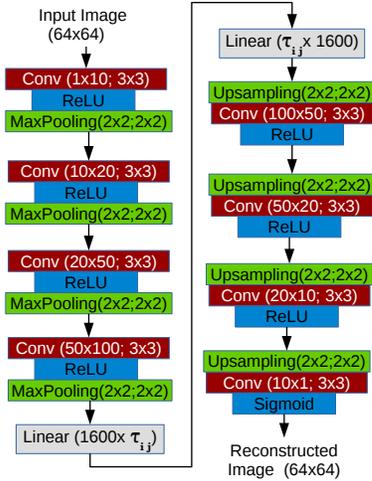


Figure 5: Block diagram of the denoising autoencoder using convolutional layers. The learnt shape representations are used to train the regression network.

5.2. Radial representation

A radial descriptor represents a shape as a series of distances between some centre point within the shape and points deterministically distributed over the shape’s boundary. We choose the boundary points by finding where rays cast outwards from the centre point at angles uniformly distributed over $[0, 2\pi)$ intersect the boundary. To improve shape reconstruction it makes sense to choose a centre point that has a direct line of sight to as much of the boundary as possible. In practice, we construct several radial descriptors for each shape and pick the one with maximal IoU. More details can be found in the supplementary material.

5.3. Learned shape encoding

The aim is to implement a network that can compress an input binary mask to a comparatively lower-dimensional space. Autoencoders have already been shown to improve over PCA and logistic PCA at a similar task for digits, curves and faces [19]. Thus, in an initial experiment, we use the fully-connected autoencoder architecture (784-1000-500-250-30) of [19] to learn a shape embedding space for the Caltech-101 silhouettes contained in the trainval set. In addition to the restricted Boltzmann machine-style pre-training of [19], we fine-tune the network weights using cross-entropy error between the predicted and target shape masks. This setup yields a reconstruction IoU of 0.85 on the test set. It becomes evident that this architecture, also employed in [10, 9], is not able to exploit the spatial redundancy in visual data due to its fully-connected nature; for the same reason, it is not scalable to large images.

It is important to note that similar shape patterns or boundary formations can be found at different locations in the input image: for example, the front and rear tyre of a bicycle or motorbike look similar. We can leverage this

repetition of visual patterns by the use of a convolutional construct. Accordingly, a block diagram of our autoencoder network is shown in Fig. 5. To prevent the network from learning an identity mapping to a higher-dimensional space in the initial layers of the network meant for data-disentangling, we introduce white noise in the input image so the network generalises to a denoiser [39]. We minimize the binary cross-entropy loss to adjust the network weights. It can be simplified as

$$L_{bce} = \sum_{i=1}^P \begin{cases} -\log(\hat{p}_i) & \text{if } p_i = 1 \\ -\log(1 - \hat{p}_i) & \text{if } p_i = 0, \end{cases} \quad (4)$$

where P is the total number of pixels in the image, \hat{p}_i is the predicted probability and p_i is the pixel value.

6. Experiments and Results

Our method allows us to unify detection and segmentation by integrating higher-order concepts of shape into the end-to-end prediction pipeline. We thus study the application of our proposed prediction network to instance segmentation (§6.2), where we found that having the concept of shapes allows it to predict shape masks even for unseen categories. Section 6.3 presents a comparative analysis of this zero-shot segmentation capability. Essential to our formulation is the need to reconstruct accurate shape masks from an interpretable and low-dimensional shape space. We thus also analyse the trade-offs between three different shape spaces in §6.1.

6.1. Comparing the Shape Representations

The choice of representation is crucial, as it dictates what aspects of a shape are made explicit and the ease with which it can be manipulated [26]. We thus compare the shape representations outlined in §5 on three different axes: i) their ability to accurately represent shapes at low dimensionalities, ii) the continuity of the associated shape space and iii) the structure of the shape space.

Representation ability: We compute the average shape reconstruction error (one minus the IoU) yielded by each of our three descriptors at different sizes on the SBD dataset. Table 1 shows how the error for each descriptor varies with the representation size. At low descriptor dimensionalities our encodings are better able to preserve shape than either of the two hand-crafted alternatives. However, the benefits fade away at larger dimensionalities. A qualitative example of the artifacts introduced by lowering the dimensionality of each shape representation can be seen in Fig. 6: both of the two hand-crafted representations experience a severe decrease in quality as the size decreases, whereas our representation still manages to capture the major topological features of the shape even at size 20.

SBD	20 (25)	50 (49)	100	200 (196)	256
Downsampled mask	0.15	0.10	0.07	0.05	0.04
Radial	0.13	0.08	0.07	0.06	0.06
Shape encoding	0.125	0.08	0.07	0.06	0.06

Table 1: Average shape reconstruction error (one minus the IoU) achieved by each of our three descriptors at different sizes on the SBD validation set [11, 16]. The unbracketed numbers indicate the sizes used for the radial and shape encoding descriptors; the bracketed numbers indicate the sizes used for the downsampled mask descriptors, which are square by design.

Shape space continuity: To study the continuity of the shape spaces, we add Gaussian noise to the encodings and visualise its effects on the reconstructed shape masks in Fig. 7. As expected, an increase in noise degrades all representations; however, a little bit of noise in the radial vector distorts the shape beyond recognition. The binary mask loses precision at the boundaries, as the wheels of the bicycle blend into the border of the image. The learned shape representation is able to preserve details such as the seat and the circular wheels better than the hand-crafted alternatives.

Shape space structure: We visualise the structure of the shape spaces associated with two representations – the learnt encodings and the radial descriptors – in Fig. 8. Even at a first glance, it is obvious that the shape space for the learnt encoding provides a better category-based clustering of the shape masks. For instance, as we move from the top left region of the space to the bottom right along a diagonal, we observe a transition from samples with large boundary projections (horses and birds) to upright shapes of humans and finally to more streamlined shapes such as bottles. There is no such clear structure in the shape space for radial descriptions: indeed, the categories seem to be scattered across various disparate clusters. We also visualise the nearest neighbour of each of a pre-selected set of 10 ‘anchor’ images, where the learnt embedding space observably shows a closer match between neighbouring shapes.

Following this qualitative analysis, we use the PASCAL3D+ [40] dataset to quantify the difference in shape space structure by gathering object pose and category statistics around neighbourhoods in the respective shape spaces. We divide the PASCAL3D+ dataset into a train set of 11578 images and a val set of 120 images. The val set is formed by randomly selecting 10 images from each of the 12 rigid object categories of PASCAL3D+ that are a subset of the 20 PASCAL VOC categories. For each val set image, we find the 50 nearest neighbours in the train set and calculate the variance in pose and category with respect to the validation image. The averaged results over all val images and categories are shown in Table 2.

We find that poses (elevation, azimuth, distance) of objects in the vicinity of a shape in the learned shape encoding space are more similar to each other than those in the

	Radial (20D)	Shape Encoding (20D)
mAV (azimuth)	1.163	1.008
mAV (elevation)	0.294	0.246
mAV (distance)	21.537	17.968
# GT class = majority class	7 out of 12	12 out of 12
mA% majority class	0.547	0.800

Table 2: Comparing the interpretability of the radial representation with that of the learned shape encoding on the PASCAL3D+ dataset. The statistics are obtained from the 50 nearest neighbours of a validation shape in each respective space, and then averaged over all 10 validation samples and 12 object categories.

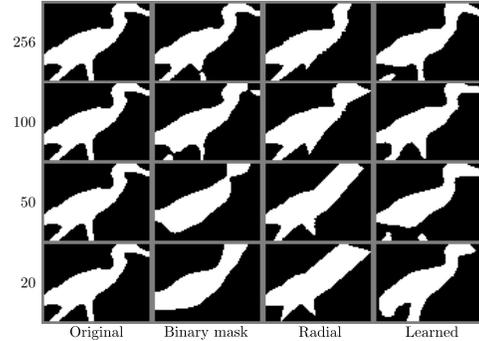


Figure 6: As the dimensionality reduces, both binary masks and radial vectors lose precious details, such as the legs and head of the bird, whilst the learned embedding is able to preserve the overall topology of the shape and degrades more gracefully.

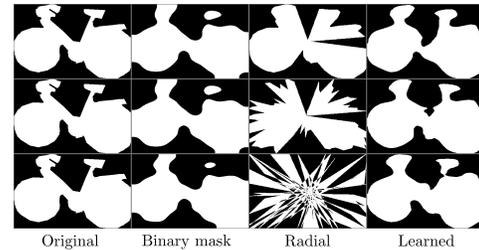


Figure 7: The effect of adding Gaussian noise of zero mean and increasing variance to each of the shape representations.

radial space (smaller mean average variance - mAV). Moreover, the majority category amongst neighbouring objects is more likely (80% of the time) to match to the category of the central shape. These numbers demonstrate that the shape space is more interpretable than the radial space.

6.2. Instance Segmentation as an Application

The hypothesis that we would like to test is that we are able to regress directly to an object-based representation of the scene. For this experiment, we train a YOLO-style network that maps directly from a single image to a collection of object masks and their locations that can be compared against those generated by state-of-the-art instance segmentation methods [17, 2, 6].

Dataset, splits and performance metrics: We train and

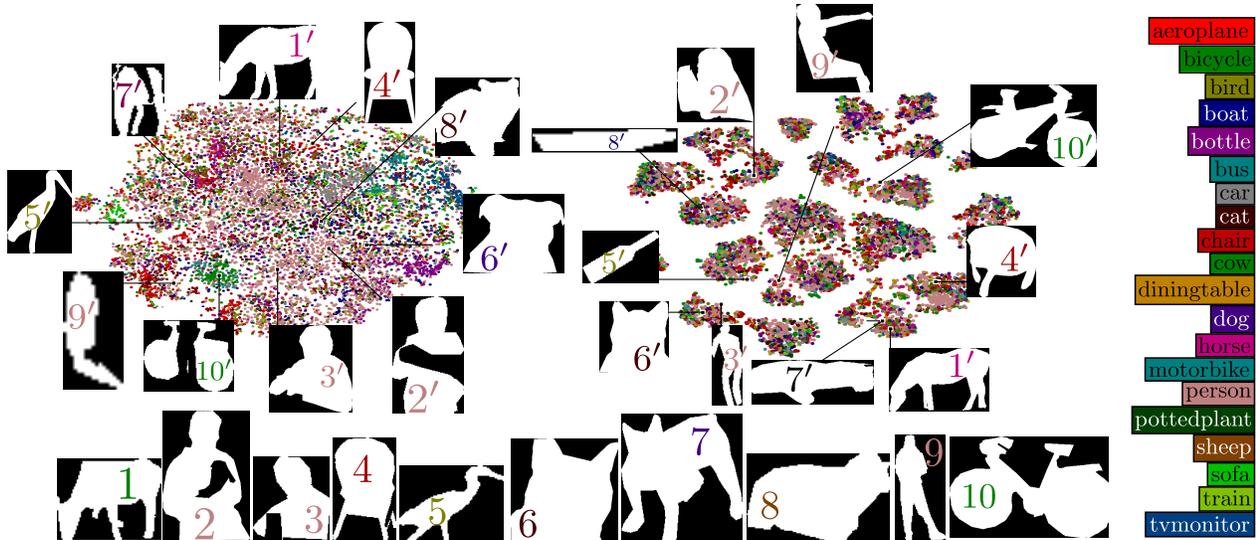


Figure 8: Visualising the 20D learned embedding space (left) next to the 256D radial vector space (right) with *t*-SNE [38]. Notice the differences between the anchors’ nearest neighbours (NNs) in the two representations. For instance, consider the bicycle anchor (no. 10), which is facing left: its NN in the learned embedding is also facing left (even though it is occluded), whilst its NN in the radial space is facing to the right.

evaluate our method on the established PASCAL VOC with SBD annotations [11, 16] dataset, with the same split sets as used by [17, 6]. We report the mean average precision at various overlap thresholds (mAP^r , mAP_{vol}^r) [17].

Results and discussion: In Table 3, we report the instance segmentation results we obtain when regressing to the following shape representations: i) downsampled binary shape masks ($d = 256$), ii) radial vectors ($d = 256$), and iii) learned embeddings ($d \in \{20, 50\}$). The dimensionalities were chosen based on the reconstruction errors obtained in Table 1, where we observed that we are able to use lower-dimensional shape codes with the learned embedding than with the hand-crafted ones (see Fig. 6). This reduction in the number of parameters per shape means that we can train neural network models with fewer parameters, and with less chance of overfitting. This intuition is consistent with our empirical findings in Table 3, where we achieve better performance as the dimensionality of the embedding is reduced from 50 to 20. In comparison to both of our embeddings, both downsampled 16×16 binary masks and 256D radial descriptors perform worse, despite having more space available in which to represent the shape. One possible reason for this in the case of radial descriptors is their higher sensitivity to noise, as can be seen in Fig. 7. Qualitative results showing the performance of our method on the SBD dataset can be seen in Fig. 9, and in the supplementary material.

It will be seen in Table 3 that our instance segmentation results lag somewhat behind current state-of-the-art methods; this parallels the lag in object detection results experienced by YOLO with respect to state-of-the-art methods in that field. We hypothesise that this is because both ap-

	SBD (2857 val images)			Time ms
	$mAP^r@.5$	$mAP^r@.7$	mAP_{vol}^r	
BinaryMask	32.3	12.0	28.6	26.3
Radial	30.0	6.5	29.0	27.1
Embedding (50)	32.6	14.8	28.9	30.5
Embedding (20)	34.6	15.0	31.5	28.0
SDS [17]	49.7	–	41.4	48k
MNC [6]	65.0	46.7	–	330

Table 3: Quantitative instance segmentation results on PASCAL SBD 2012 val. The timing results were obtained on a high-end desktop containing a Titan X.

proaches struggle to precisely localise certain objects, especially small ones (see Figure 9). However, like YOLO, our approach runs in real time, something that cannot be achieved by any of the *offline* methods against which we compare (as can be seen in Table 3, existing methods are in no way real time, taking at least an order of magnitude longer to run). We are also much better than the state-of-the-art instance segmentation method of Arnab *et al.* [2] at recognising unseen categories (see Fig. 10), as discussed in the next section.

6.3. Zero-shot segmentation

Traditionally, the task of zero-shot learning has been achieved through the use of attributes [35, 42]. A recent work on zero-shot boundary segmentation [23] follows a similar approach. By contrast, our method doesn’t need an explicit definition of attributes to extend shape segmentation to unseen categories (see Fig. 10); shape attributes are implicitly captured in the definition of the shape space. We are thus able to scale to object categories that have similar shape appearances at test time. This is very similar to how



Figure 9: (a) Correct predictions using downsampled binary masks. (b) Correct predictions using 20D learnt shape encodings. In the 3rd image, the horns of the cow are missed and the human shape mask gets elongated due to an incorrect bounding box prediction. (c) Missed detections using the 20D shape encodings. The network misses out or false fires on small objects (2nd column). The dogs in the images are falsely categorised as cats, and the sofa incorrectly includes the nearby dining table.

humans perform. For instance, someone who has never seen a tiger before would still be able to segment it properly, even though he/she might compare it to a *cat- or dog-like animal*, and this is precisely how our network approaches the task.

In order to evaluate our segmentation of unseen class objects, we take our shape prediction model (50D learned encoding) trained on PASCAL VOC train and val, and test it on the 60 object categories in MS-COCO [24] not present in PASCAL VOC. We filter images from the entire MS-COCO dataset such that no image contains an object of a seen category; yielding 8037 images in our test-bed with a total of 24,396 instances of unseen object classes. We gather all object predictions with a detection score greater than $t = 0.05$, since the unseen category scores are likely to be low, and then run them through the evaluation code provided by [24], which we modified to calculate IoU over shapes instead of boxes. The results, in Table 4, are specially reported for large objects as the PASCAL VOC dataset used for training predominantly contains objects ($> 80\%$ objects) in the large size range as defined in the MS-COCO dataset. Given that the existing methods perform very poorly for segmentation of unseen classes (Fig. 10), we report our results as a strong baseline. We hope that these results will help catalyse research in this direction.

COCO (8037 val images)			
	mAP ^r @.5 (all)	mAP ^r @.5 (large)	AR ^r @.5 (large)
Embedding (50)	3.6	7.1	23.2

Table 4: Quantitative zero-shot segmentation results obtained by running our embedding (50) prediction model on the 60 categories of MS-COCO not present in PASCAL VOC.

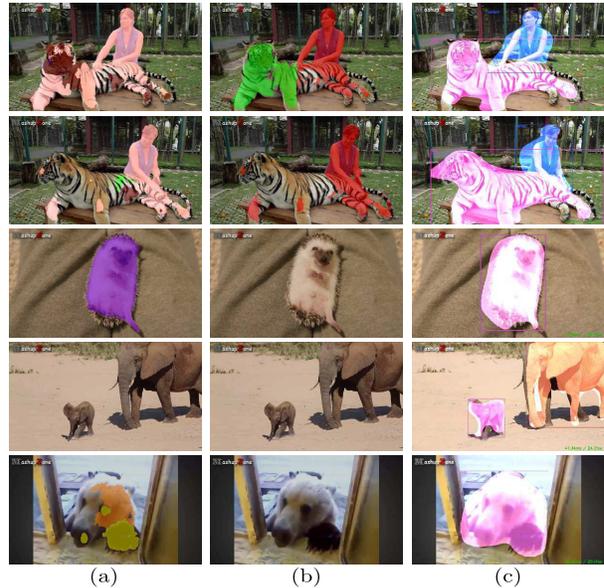


Figure 10: A comparison between the state-of-the-art (a) semantic segmentation [1], (b) instance segmentation [2], and (c) our shape detection results, on images from YouTube videos of animals that are not present in the PASCAL VOC training set. In the first two rows, instance segmentation predicts that the legs of the tiger are human. Our method is more consistent over the tiger images taken from the same video. In the lower rows, the instance segmentation approach (b) fails to predict any segments, whilst our method predicts class ‘dog’ for the tiger, hedgehog, baby elephant and bear, and class ‘horse’ for the large elephant.

7. Conclusion

In this work, we show for the first time that it is possible to regress directly and simultaneously to multiple object representations that incorporate the notion of shape. We combine ideas from object detection, instance segmentation and low-dimensional embedding spaces to create a real-time system to detect encoded shapes. One key factor that allows us to do this is the introduction of a shape embedding space that is compact, decodable, continuous and interpretable. We find that imbuing object detectors with the knowledge of shape allows us to predict plausible masks for previously-unseen categories, allowing us to detect the presence of objects in images for which current state-of-the-art instance segmentation methods perform poorly or fail.

Our next step is to investigate how our shape prediction approach can be extended to cope with a larger variety of shapes and object categories like those in MS-COCO [24].

Acknowledgements. This work was supported by the EP-SRC, ERC grant ERC-2012-AdG 321162-HELIOS, EP-SRC grant Seebibyte EP/M013774/1 and EPSRC/MURI grant EP/N019474/1.

References

- [1] A. Arnab, S. Jayasumana, S. Zheng, and P. H. Torr. Higher order conditional random fields in deep neural networks. In *European Conference on Computer Vision*, pages 524–540. Springer, 2016. 8
- [2] A. Arnab and P. H. S. Torr. Bottom-up instance segmentation using deep higher-order crfs. In *Proc. British Machine Vision Conference*, 2016. 2, 6, 7, 8
- [3] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(4):509–522, Apr. 2002. 2
- [4] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *arXiv preprint arXiv:1606.00915*, 2016. 1
- [5] Y.-T. Chen, X. Liu, and M.-H. Yang. Multi-instance object segmentation with occlusion handling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3470–3478, 2015. 3
- [6] J. Dai, K. He, Y. Li, S. Ren, and J. Sun. Instance-sensitive fully convolutional networks. *CoRR*, abs/1603.08678, 2016. 3, 6, 7
- [7] J. Dai, K. He, and J. Sun. Instance-aware semantic segmentation via multi-task network cascades. *CoRR*, abs/1512.04412, 2015. 1, 3
- [8] J. Dai, K. He, and J. Sun. Instance-aware semantic segmentation via multi-task network cascades. *arXiv preprint arXiv:1512.04412*, 2015. 2
- [9] S. Eslami and C. Williams. A generative model for parts-based object segmentation. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 100–107. Curran Associates, Inc., 2012. 5
- [10] S. A. Eslami, N. Heess, C. K. Williams, and J. Winn. The shape boltzmann machine: a strong model of object shape. *International Journal of Computer Vision*, 107(2):155–176, 2014. 3, 5
- [11] M. Everingham, L. Van Gool, C. Williams, J. Winn, and A. Zisserman. Pascal visual object classes challenge results. Available from www.pascal-network.org, 2005. 1, 6, 7
- [12] V. Ferrari, F. Jurie, and C. Schmid. From images to shape models for object detection. *Int. Journal of Computer Vision*, 2009. 2
- [13] R. Girshick. Fast R-CNN. In *Proc. Int. Conf. Computer Vision*, 2015. 1
- [14] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, 2014. 2
- [15] S. Gupta, R. B. Girshick, P. A. Arbeláez, and J. Malik. Learning rich features from RGB-D images for object detection and segmentation. In *Proc. European Conf. Computer Vision*, 2014. 3
- [16] B. Hariharan, P. Arbeláez, L. Bourdev, S. Maji, and J. Majik. Semantic contours from inverse detectors. In *Proc. Int. Conf. Computer Vision*, 2011. 6, 7
- [17] B. Hariharan, P. Arbeláez, R. Girshick, and J. Malik. Simultaneous detection and segmentation. In *Proc. European Conf. Computer Vision*, 2014. 2, 6, 7
- [18] S. L. Hicks, I. Wilson, L. Muhammed, J. Worsfold, S. M. Downes, and C. Kennard. A depth-based head-mounted visual display to aid navigation in partially sighted individuals. *PLoS ONE*, 8(7):1–8, 07 2013. 1
- [19] G. E. Hinton and R. R. Salakhutdinov. Reducing the dimensionality of data with neural networks. *Science*, 2006. 3, 5
- [20] A. K. Jain, Y. Zhong, and S. Lakshmanan. Object matching using deformable templates. *IEEE Transactions on pattern analysis and machine intelligence*, 18(3):267–278, 1996. 2
- [21] Y. LeCun and Y. Bengio. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995. 2
- [22] Y. Li, H. Su, C. R. Qi, N. Fish, D. Cohen-Or, and L. J. Guibas. Joint embeddings of shapes and images via CNN image purification. In *Siggraph Asia*, 2015. 2, 3, 4
- [23] Z. Li, E. Gavves, T. E. J. Mensink, and C. G. M. Snoek. Attributes make sense on segmented objects. In *European Conference on Computer Vision*, 2014. 7
- [24] T. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. In *Proc. European Conf. Computer Vision*, 2014. 8
- [25] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional models for semantic segmentation. In *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, 2015. 3
- [26] D. Marr. *Vision*. W. H. Freeman and Company, 1982. 5
- [27] M. Marszałek and C. Schmid. Accurate object localization with shape masks. In *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, 2007. 2
- [28] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781, 2013. 2, 3
- [29] P. O. Pinheiro, R. Collobert, and P. Dollár. Learning to segment object candidates. In *Advances in Neural Information Processing Systems*, 2015. 2
- [30] P. O. Pinheiro, T.-Y. Lin, R. Collobert, and P. Dollár. Learning to refine object segments. In *Proc. European Conf. Computer Vision*, 2016. 2
- [31] D. Rao, Q. V. Le, T. Phoka, M. Quigley, A. Sudsang, and A. Y. Ng. Grasping Novel Objects with Depth Segmentation. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, pages 2578–2585, 2010. 1
- [32] J. Redmon. Darknet: Open source neural networks in c. <http://pjreddie.com/darknet/>, 2013–2016. 4
- [33] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *IEEE Int. Conf. on Computer Vision and Pattern Recognition*, 2016. 2, 3, 4

- [34] S. Ren, K. He, R. Girshick, and J. Sun. Faster R-CNN: towards real-time object detection with region proposal networks. *Advances in Neural Information Processing Systems*, 2015. [1](#), [2](#)
- [35] B. Romera-Paredes. An embarrassingly simple approach to zero-shot learning. In *Int. Conf. on Machine Learning*, 2015. [7](#)
- [36] B. Romera-Paredes and P. H. S. Torr. Recurrent instance segmentation. *CoRR*, abs/1511.08250, 2015. [3](#)
- [37] B. Tekin, I. Katircioglu, M. Salzmann, V. Lepetit, and P. Fua. Structured prediction of 3d human pose with deep neural networks. In *Proc. British Machine Vision Conference*, 2016. [3](#)
- [38] L. van der Maaten and G. Hinton. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605, 2008. [7](#)
- [39] P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning*, pages 1096–1103. ACM, 2008. [2](#), [3](#), [5](#)
- [40] Y. Xiang, R. Mottaghi, and S. Savarese. Beyond PASCAL: A benchmark for 3D object detection in the wild. In *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2014. [6](#)
- [41] D. Zhang and G. Lu. Review of shape representation and description techniques. *Pattern Recognition*, 37, 2004. [2](#)
- [42] Z. Zhang and V. Saligrama. Zero-shot learning via joint latent similarity embedding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6034–6042, 2016. [7](#)
- [43] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. H. S. Torr. Conditional Random Fields as Recurrent Neural Networks. In *Proc. Int. Conf. Computer Vision*, pages 4321–4329, 2015. [1](#)