Unsupervised part representation by Flow Capsules

Sara Sabour [1 2] **Andrea Tagliasacchi** [1 2] **Soroosh Yazdani** [2] **Geoffrey E. Hinton** [1 2] **David J. Fleet** [1 2]

Abstract

Capsule networks are designed to parse an image into a hierarchy of objects, parts and relations. While promising, they remain limited by an inability to learn effective low level part descriptions. To address this issue we propose a novel self-supervised method for learning part descriptors of an image. During training, we exploit motion as a powerful perceptual cue for part definition, using an expressive decoder for part generation and layered image formation with occlusion. Experiments demonstrate robust part discovery in the presence of multiple objects, cluttered backgrounds, and significant occlusion. The resulting part descriptors, a.k.a. part capsules, are decoded into shape masks, filling in occluded pixels, along with relative depth on single images. We also report unsupervised object classification using our capsule parts in a stacked capsule autoencoder.

1. Introduction

Humans learn to perceive shapes in terms of their parts and their spatial relationships (Singh & Hoffman, 2001). Studies show that infants form early object perception by dividing visual inputs into units that move rigidly and separately (Spelke, 1990), and they do so in a largely unsupervised way. Inspired by this and recent work on parts discovery, we propose a self-supervised method for learning visual part descriptors for Capsule networks (Hinton et al., 2018; Sabour et al., 2017).

Capsule networks represent objects in terms of primary part descriptors, in local canonical frames, and coordinate transformations between parts and the whole. As a result of their architecture, they are robust to various challenges, including viewpoint and adversarial attacks. Stacked capsule network architectures (Kosiorek et al., 2019) have shown very promising results on a number of image datasets. But because they are trained with an image reconstruction loss,

foreground-background separation and part discovery in cluttered images remain challenging.

This paper introduces a way to learn encoders for object parts (a.k.a., primary capsules) to address these challenges. The encoder takes as input a single image. During training, however, parts discovery is accomplished through self-supervision with video (cf., (Mahendran et al., 2018)). Like the classical literature on perceptual organization and common fate in Gestalt psychology (e.g., (Spelke, 1990; Wagemans et al., 2012)), we exploit the fact that regions of the image that move together often belong together. This is a strong perceptual cue that facilitates foreground-background segmentation and parts discovery, and allows one to disentangle texture and other aspects of appearance from shape.

The proposed part encoder (see Fig. 1) captures the underlying part shapes, their relative poses, and their relative depth ordering. The introduction of depth ordering is particularly useful in order to account for occlusion, much like it has been used in classical layered motion models (Wang & Adelson, 1994). In this way, through learning, we naturally aggregate information over many training images even though a given part may rarely be visible in its entirety in any single frame. In essence, the model prefers a simple part-based description, where many variations in appearance can be explained by a coordinate transform or by occlusion, rather than by a change in shape.

We demonstrate the approach on several datasets showcasing challenges due to texture, occlusions, scale, and instance variation. We compare our model to recent related work including PSD (Xu et al., 2019) and R-NEM (Van Steenkiste et al., 2018), where part masks and dynamics are learnt using motion as input during training to generate the next frame. The Flow Capsule model provides unsupervised segmentation of shapes, even in the face of texture and background, outperforming PSD (Xu et al., 2019). Flow capsules also provide explicit depth ordering to account for occlusion, with the added benefit that part inference facilitates shape completion when observed parts are partially occluded.

We also report unsupervised classification of images using Flow Capsule part embeddings. We compare our results on several datasets with different challenges against SCAE (Kosiorek et al., 2019). Our experiments show Flow Cap-

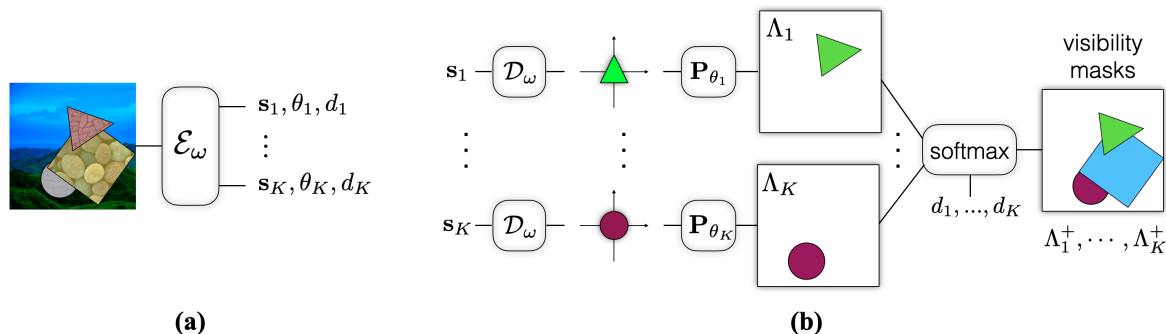[1]Department of Computer Science, University of Toronto [2]Google research. Correspondence to: Sara Sabour <sasabour@google.com>.

(a) **(b)**

Figure 1: **Segmentation Inference Architecture** – (a) The encoder $\mathcal{E}_\omega$ parses an image into part capsules, each comprising a shape vector s_k, a coordinate transform θ_k, and a scalar depth value d_k. (b) A shared shape decoder $\mathcal{D}_\omega$ represents part shapes in a canonical coordinate frame. The shape decoder is an implicit function that takes as input a shape vector, s_k, and a location in canonical coordinates and returns the probability that location is inside the shape. Shapes are then transformed into image coordinates, using θ_k, and layered according to the relative depths d_k, yielding part visibility masks.

sules consistently outperform SCAE in unsupervised object classification, especially on images with background.

2. Related Work

Given the vast literature of part-based representations for visual tasks, we focus here only on the most closely related recent work.

Transforming autoencoders (Hinton et al., 2011) introduced the idea of capsule networks. (Sabour et al., 2017) revisited the capsule concept and introduced capsule hierarchies for object classification, and subsequent work has produced improved routing algorithms (Hinton et al., 2018; Hahn et al., 2019; Ahmed & Torresani, 2019). Nevertheless, learning primary capsules from images has remained largely untouched. An analogy to text understanding would be a language that has a well defined grammar and parser, but no good definition or representation of words. Our work addresses this shortcoming with a technique for learning primary capsules.

Unsupervised capsule learning has been explored by (Kosiorek et al., 2019) and (Rawlinson et al., 2018). Such methods rely on an image reconstruction loss for part discovery and object-background discrimination is difficult with cluttered backgrounds and texture. (Srivastava et al., 2019; Zhao et al., 2019) have proposed unsupervised learning of capsule autoencoders for 3D objects from point clouds. With the exception of capsule models trained with classification labels (Hinton et al., 2018) or segmentation masks (LaLonde & Bagci, 2018; Duarte et al., 2018), previous methods have struggled with natural images. Using visual motion in video for training, flow capsules are able to classify and segment images without ground truth labels or segmentation masks.

Our current flow capsule formulation models 2D objects, producing representations reminiscent of classical layered

models. Many layered models assign pixels to layers independently, using linear mixture models (Wang & Adelson, 1994), and hence they often do not capture the coherence or compactness of object occupancy. Some methods used soft MRF priors to encourage spatial coherence (Weiss, 1997), while others enforced spatial coherence through local parametric support masks (Jepson et al., 2002). In flow capsules, by comparison, we learn the correspondence shape masks.

Flow capsules use visual motion for part discovery, as it is well-known to be a strong cue for self-supervised learning. For example, (Vijayanarasimhan et al., 2017) learns to infer depth, segmentation, and relative 3D motion from consecutive frames, and can be trained in a self-supervised manner using photometric constraints. Optical flow has also been used to learn hierarchical part-based models of shape and dynamics in terms of a layered image model (Xu et al., 2019). Given a sequence of frames, SfM-Net (Vijayanarasimhan et al., 2017) predicts depth, segmentation, camera and rigid object motions, converts those into a dense frame-to-frame motion field (optical flow), differentiably warps frames in time to match pixels and back-propagates. The model can be trained with various degrees of supervision: 1) self-supervised by the re-projection photometric error (completely unsupervised), 2) supervised by ego-motion (camera motion), or 3) supervised by depth (e.g., as provided by RGBD sensors). SfM-Net extracts meaningful depth estimates and successfully estimates frame-to-frame camera rotations and translations. It often successfully segments the moving objects in the scene, even though such supervision is never provided. These and related methods take *optical flow* or multiple frames as input. While flow capsules use video during training, the learned part detector (see Fig. 1), takes as input a single frame, as also done by (Mahendran et al., 2018).

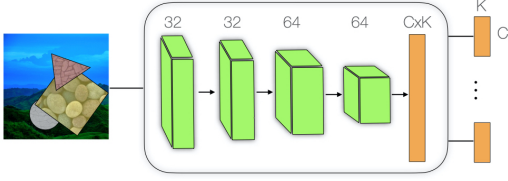Other work has focused on representation learning for shape

Figure 2: **Encoder Architecture** – Image is passed through a series of convolutional layers with ReLU activation. After each convolution layer representations are downsampled by 2x2 AveragePooling. Last Convolutional Layer is attached to a tanh fully connected layer followed by a simple fully connected layer that is grouped into K capsules of C dimensions each.

Figure 3: **Decoder Architecture** – We use an implicit function decoder to decode the shape vectors into mask values. The shape vector s is concatenated with each pixel position u and then passed through an MLP with ReLU activations. The final output is the one dimensional logit of the masks.

via generative models. (Eslami & Williams, 2012) use an RBM for learning 2D part-based models, with block Gibbs sampling for inference. More closely related to flow capsules, (Huang & Murphy, 2016) learn a generative model for shape within a layered image model, including depth ordering and occlusion. Given an image, variational inference is used to infer shape and foreground/background separation. Flow capsule encoders, by comparison, are trained as auto-encoders and are therefore somewhat easier to learn. And there are many other approaches to learning generative models of shape that disentangle shape and deformation to explain observed images (e.g., (Skafte & Hauberg, 2019)). Flow capsules also disentangle shape and transformations from canonical coordinates to image coordinates, but they decompose shapes into multiple, (approximately) rigid parts with occlusions.

3. Model

Our goal is to learn an encoder that parses images of familiar shapes into parts. To facilitate training, and subsequent downstream tasks, we also learn a decoder capable of generating segment masks for the parts in the image. (see Fig. 1). In what follows we first describe the form of the proposed capsule encoder and the mask decoder. The subsequent section then describes the objective and training procedure.

Image Encoder: Our goal is to learn a capsule encoder $\mathcal{E}_\omega$, with parameters ω, that encodes a given image a collection of K primary capsules. The architecture we propose is depicted in Fig. 2. Each capsule, $\mathbf{c}_k$, comprises a vector $\mathbf{s}_k$ that encodes the *shape* of the part, a *pose* vector $\boldsymbol{\theta}_k$ that parameterizes the mapping from part-centric coordinates into image coordinates (or scene coordinates more generally), and a scalar d_k that specifies relative inverse depth (ie larger for foreground objects):

$$\mathcal{E}_\omega(I) = \{\mathbf{c}_0, \dots, \mathbf{c}_k\}, \quad \mathbf{c}_k = (\mathbf{s}_k, \boldsymbol{\theta}_k, d_k) \,. \quad (1)$$
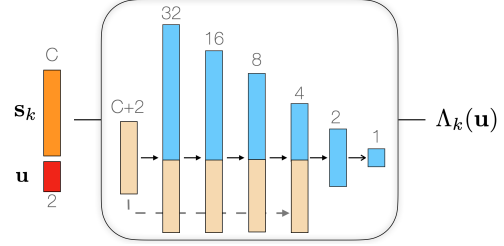
One can think of the geometry as expressing the part shape in canonical coordinates and the pose as the transformation from part coordinates to image coordinates. At present we use a four-parameter conformal mapping from part to image coordinates, i.e., uniform scaling, rotation, and translation (as depicted in Fig. 1).

Mask Decoder: A mask decoder is introduced to facilitate self-supervised learning of the encoder, as well as downstream tasks for which the support of the part in the image is useful. It allows one to visualize the part and connect it to observations in the image. A mask decoder $\mathcal{D}_\omega$ generates an object silhouette (or mask) that specifies the part shape in terms of the region occupied by it. Fig 3 visualizes our current mask decoder architecture, while Fig. 1 captures the entire process of decoding the part and mapping the part from canonical coordinates $\mathbf{v}$ into image coordinates $\mathbf{u}$, i.e., $\mathbf{u} = \mathbf{P}_{\boldsymbol{\theta}_k}\mathbf{v}$.

In more detail, the mask decoder represents the shape of the part in its canonical coordinate frame, $\mathcal{D}_\omega(\mathbf{v}; \mathbf{s}_k)$. This is then mapped into image coordinates according to the pose vector, $\boldsymbol{\theta}_k$, yielding the shape mask in the image frame, denoted Λ_k:

$$\Lambda_k(\mathbf{u}) \;=\; \mathcal{D}_\omega(\mathbf{P}_{\boldsymbol{\theta}_k}^{-1}(\mathbf{u}); \mathbf{s}_k) \,, \quad (2)$$

where $\mathbf{P}_{\boldsymbol{\theta}_k}$ is conformal mapping with parameters $\boldsymbol{\theta}_k$. Note that Λ_k is a function of spatial position.

We employ capsule codes of size $\mathbf{c}_k \in \mathbb{R}^C$, where $\boldsymbol{\theta}_k \in \mathbb{R}^4$, $d_k \in \mathbb{R}$, and therefore $\mathbf{s}_k \in \mathbb{R}^{C-5}$. As we focus on planar layered models with depth d, we extract the translational component as the first two dimensions $[\boldsymbol{\theta}_k]_{0,1}$, the rotation and scale as $[\boldsymbol{\theta}_k]_{2,3}$. More concretely,

$$\mathbf{P}_{\boldsymbol{\theta}_k} = \begin{bmatrix} \cos([\boldsymbol{\theta}_k]_2)[\boldsymbol{\theta}_k]_3 & -\sin([\boldsymbol{\theta}_k]_2)[\boldsymbol{\theta}_k]_3 & [\boldsymbol{\theta}_k]_0 \\ \sin([\boldsymbol{\theta}_k]_2)[\boldsymbol{\theta}_k]_3 & \cos([\boldsymbol{\theta}_k]_2)[\boldsymbol{\theta}_k]_3 & [\boldsymbol{\theta}_k]_1 \\ 0 & 0 & 1 \end{bmatrix}$$

$$(3)$$

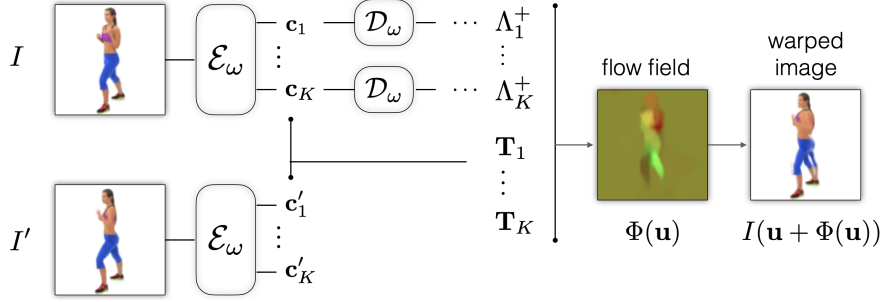**Occlusion:** With opaque objects one does not expect parts

Figure 4: **Self Supervised Training** – We seek to train a single image encoder $\mathcal{E}_\omega$ capable to decompose the scene into a collection of *primary capsules*. We perform this task via the proxy task of regressing a flow field Φ that is used to warp image $\mathbf{I}$ into $\mathbf{I}'$. The flow field is determined by the parts visibility masks Λ_k^+ in (4), and the part coordinate transforms $\mathbf{T}_k$ in (5). The rendering loss for training (7) is the residual error between the warped version of $\mathbf{I}$ and the next frame $\mathbf{I}'$.

will be visible in their entirety. To account for occlusion, part masks are layered according to their depth order, thereby providing the visible portion of each part in a given image. In order to do so in a differentiable manner, to enable gradient-based learning, we treat the scalar d_k as a logit, and apply a softmax across the logits (depths) at every pixel to generate a visibility mask for each part (see Fig. 1). The visible portion of the kth part is given by

$$\Lambda_k^+ = \text{Softmax}_k(d_k \Lambda_k) . \tag{4}$$

As the gap between the largest d_k and other values grows, the softmax approaches the argmax, which of course would be ideal for opaque layers.

A typical auto-encoder might then reconstruct the image in terms of these masks in order to formulate the image reconstruction loss. The problem with such an approach is that the encoder would also be forced to encode other properties of the images, such as texture, lighting and the background, with pixel level accuracy. To avoid this problem, here we aim only to learn an encoder for the part shapes, positions and depth layering, and then consider a form of self-supervised learning that relies on motion (optical flow) between consecutive frames in video. The use of flow provides a strong image cue for the segmentation of parts, without the need to model texture, lighting and other fine-grained properties of appearance.

4. Self-Supervised Learning

The self-supervised loss requires training data comprising pairs of images in temporal succession. From a given image pair, the encoder separately provides an ordered set of capsules detected in each image. The poses from corresponding capsules and their visibility masks determine an optical flow field for the two images. In particular, this flow field can be used to warp one frame toward the other, from which brightness constant and other common objectives in optical flow optimization are used to specify the training loss.

In more detail, let the two images of a training pair be denoted $\mathbf{I}$ and $\mathbf{I}'$. As shown in Fig. 4, the capsule encoded is used to extract an ordered set of capsules from both images. The part capsules are denoted $\mathbf{c}_k = (\mathbf{s}_k, \boldsymbol{\theta}_k, d_k)$ and $\mathbf{c}_k' = (\mathbf{s}_k', \boldsymbol{\theta}_k', d_k')$. From the corresponding part capsules we compute the predicted optical flow field Φ in terms of the capsule poses, which determine a mapping $\mathbf{T}_k$ from one image to the next, and the layered visibility masks; i.e.,

$$\mathbf{T}_k = \mathbf{P}_{\boldsymbol{\theta}_k'} \circ (\mathbf{P}_{\boldsymbol{\theta}_k})^{-1} \tag{5}$$

$$\Phi(\mathbf{u} \,|\, \mathcal{E}_\omega(\mathbf{I}), \mathcal{E}_\omega(\mathbf{I}')) = \sum_k \underbrace{\Lambda_k^+(\mathbf{u})}_{\text{visibility}} \underbrace{[\mathbf{T}_k(\mathbf{u}) - \mathbf{u}]}_{\text{flow of } k\text{-th capsule}} \tag{6}$$

where $\mathbf{u} \in [0,1]^2$ denotes 2D image coordinates, and and $\mathbf{T}_k$ is a transformation that maps points from image locations in $\mathbf{I}$, to the canonical coordinate frame of part k, and then into the other frame $\mathbf{I}'$. Note that the use of $\mathbf{T}_k(\mathbf{u}) - \mathbf{u}$ in (6) ensures that the generation of an identity flow is the easiest prediction for the network $\mathbf{T}_k(\mathbf{u})$ to achieve – i.e. a residual connection.

Given the estimated flow between the training image pair, we warp the pixels of $\mathbf{I}$ according to Φ to estimate $\mathbf{I}'$. Then we optimize an L2 brightness constancy loss between the warped first frame and the second frame:

$$\mathcal{L}_{\text{render}} = \mathbb{E}_{\mathbf{u} \sim [0,1]^2}$$
$$\| \mathbf{I}(\,\mathbf{u} + \Phi(\mathbf{u} \,|\, \mathcal{E}_\omega(\mathbf{I}), \mathcal{E}_\omega(\mathbf{I}'))\,) - \mathbf{I}'(\mathbf{u}) \|_2^2 . \tag{7}$$

To facilitate training the model without any ground truth segmentation masks or flow fields we incorporate two regularizers. The first, $\mathcal{L}_{\text{smooth}}$, is a flow smoothness term often used in unsupervised flow estimation methods.

$$\mathcal{L}_{\text{smooth}} = \left\| \frac{\partial \Phi}{\partial u_x}, \frac{\partial \Phi}{\partial u_y} \right\|_2^2 . \tag{8}$$

This loss enhances gradient propagation through larger movements.

The second regularizer, $\mathcal{L}_{\text{center}}$, encourages the model to center the part shape close to the origin of the canonical coordinate frame:

$$\mathcal{L}_{\text{center}} = \frac{1}{K} \sum_k \frac{\sum_{\mathbf{v}} \|\mathbf{v}\, \Lambda_k(\mathbf{v})\|_2^2}{\sum_{\mathbf{v}} \Lambda_k(\mathbf{v})} \qquad (9)$$

Theoretically, our calculations are equivariant to the relative position of the center of image and center of the canonical coordinate frame. But keeping both centered at $(0,0)$ significantly improves the inference of rotations. If the part is located far from the original of the canonical frame, then during training for many rotations the object will be projected out of the image grid. Keeping it near the center helps make the loss much smoother. [1]

The finial loss is a weighted some of the render loss and the two regularizers.

5. Experiments

Dataset. We evaluate our method on images with different dynamics, shapes, backgrounds and textures. In particular, we evaluate the approach on image data of scenes containing multiple occluding colored geometrical shapes (Geo), Geo with ImageNet (Deng et al., 2009) background and textured objects (Geo$^+$), and on real human motions (Exercise) (Xu et al., 2019)}. For the synthetic Geo dataset, we use the same code and setup as (Xu et al., 2019). We generate 100k images for training, 1k for validation, and 10k for testing. Images have different background colors, with geometrical shapes (circle, triangle, square) of various color, scale and position. The objects in Geo only change their translational position between two frames.

We further analyze robustness of FlowCapsules to texture and background by augmenting the Geo dataset into Geo$^+$. In Geo$^+$ the simple background is replaced with random Imagenet images. Also random textured patches from the Brodatz texture dataset (Picard et al., 1993) are added to the geometrical shapes.

The Exercise dataset contains real images of trainers performing demo exercises. This dataset contains articulated, out of plane movements including rotation. Ideally Flow-Capsules should automatically learn parts that correspond to rigid body segments between two joints. The Exercise dataset has 49356 pairs of images for training, which are extracted from 20 exercise demo videos. The test set has 30 single images, for which (Xu et al., 2019) provided ground truth segmentation masks.

Experimental setup. We train our models using the Adam optimizer (Kingma & Ba, 2014) with a fixed learning rate

of $1e - 4$ for 150 epochs. We use $C = 16$ and $K = 16$ for our models. A regularization factor of $1e - 2$ is used for $\mathcal{L}_{\text{center}}$ and $1e - 4$ for $\mathcal{L}_{\text{smooth}}$. To calculate an intersection-over-union (IoU) measure of performance on the visibility masks, we normalize and then threshold the masks at 0.1 to get a binary $(0, 1)$ mask.

5.1. Estimating Flow during training

To verify that our architecture and setup is able to train and estimate the flow on video frames in an unsupervised manner, we first inspect the quality of FlowCapsule estimated flow after training for each dataset. Fig. 6 shows the estimated flow Φ alongside Φ_{gt} for pairs of images in Geo and Geo$^+$ training set. Even in the presence of texture and background the flow Φ is extremely accurate. Comparing the last column in Fig. 6, and the warped first frame, with the second column, and the ground truth next frame, one can appreciate some of the challenges in unsupervised flow estimation. Taking into consideration that estimating $\mathbf{I}'$ using Φ does not account for pixels that become unoccluded, $\mathcal{L}_{\text{render}}$ is not expected to reach 0. Given that ground truth flow in these two datasets is composed of only translation, we verify that our estimated flow can handle extra degrees of freedom. It successfully estimates the rotation and scale as 0 and 1.

Fig. 5 shows the estimated flow for the Exercise dataset. The ground truth flow in this case consists of several articulated movements in various poses. Also, the parts here are much smaller than in Geo. The estimated flows are blurrier than the synthetic dataset, but they still capture different movements in different parts, which means that it has successfully approximated the articulated movements with overlapping conformal transformations.

5.2. Unsupervised Segmentation

The first task on which we evaluate our part representations is unsupervised Segmentation. Our model is able to decompose a single image into its *movable* parts. Given a single image, the model has not yet seen the second image. Therefore, it has to provide the part shapes and coordinate transforms for all the parts in the scene that can be moved. Typical optical flow estimation algorithms which use masks to estimate the flow, only generate masks for *moved* parts. Because their model builds the masks by comparing the two images.

The two most relevant prior works to FlowCapsules which focus on part discovery are SCAE (Kosiorek et al., 2019) and PSD (Xu et al., 2019). Fig. 7 shows SCAE's transformed templates and image reconstructions. Even in the simple case without background or texture, one can see that SCAE fails to segment the image into cohesive parts. This failure becomes markedly worse as we add object texture
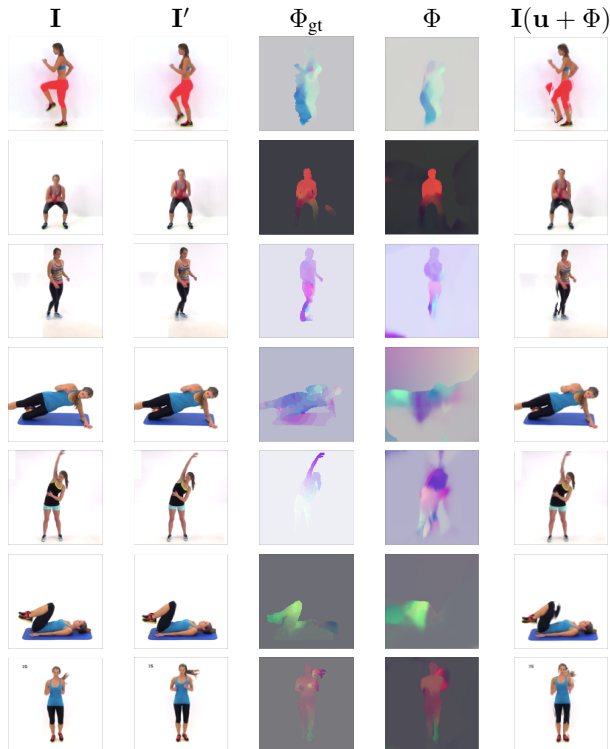
---

[1]This numerical issue can also be solved by training with a larger grid (padded around the original image).

Figure 5: Estimated flows and predicted next frames on exercise training set.

and a background in Geo$^+$.

Tab. 1 compares the intersection over union (IoU) of our masks against PSD and R-NEM (Van Steenkiste et al., 2018). Although PSD receives the ground truth optical flow during training, FlowCapsules consistently have better or equal IoUs during testing on both the Geo and Exercise datasets. One other significant difference between PSD and FlowCapsules lies in how they generate the masks. FlowCapsules generate the $[128, 128]$ masks from the capsule shape vectors of $16D$. The results in Table 1 verify that shape details are captured in the part encoding. In comparison, PSD goes straight from image to segmentation masks with convolutional layers, and there is no part encoding per se. In effect, PSD labels the pixels directly while FlowCapsules generate the shape masks.

Interestingly, FlowCapsule parts learn to explain the objects most efficiently. Masks in Fig. 8 show that FlowCapsules learn to explain a meaningful shape (e.g. a triangle or a circle) as a part. Even though we have no regularizer to force the parts to encompass one full shape (a full circle). The model can potentially learn half circles and use those to explain occluded objects. However, it opts to explain the images with relatively few parts. Hence, the capsule masks often cover all the pixels of one shape in the Geo dataset. To further investigate the quality of FlowCapsules we compare
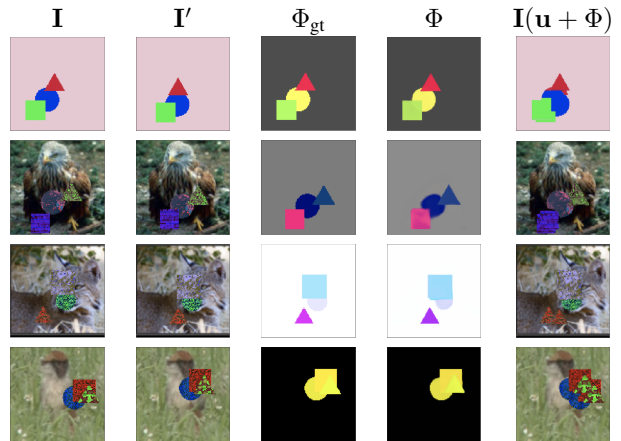

Figure 6: Estimated flows and predicted next frames on Geo and Geo$^+$ training set.

Figure 7: SCAE reconstructions after trained on Geo and Geo$^+$ on the left. The learned part shape templates are visualized on the right. Although SCAE is able to reconstruct to some degree, the part templates do not correspond to any coherent parts. The learned parts are diverse to enable reconstructing the color, texture and background in the images.

IoU of the inference masks with ground truth masks, as in (Xu et al., 2019). Table 1 shows that our generated masks have better overall quality compared to previous work.

Geo is generated synthetically, but we also have the ground truth segmentations for the full shapes. Since FlowCapsules generate the masks separately, first via Λ_k and then layered to form visibility masks, Λ_k^+, we can compare Λ_k against the full shape ground truth segmentations. In this setup Flow Capsules achieve and IoU of **0.96** on all the shapes, circle, square, and triangle equally. First this results indicates our model indeed encodes the full shape of a part rather than only the visible region. Second, comparing this results against Table 1 indicates that the FlowCapsules layering mechanism is not efficient. The model looses accuracy after layering, specially for the circle.
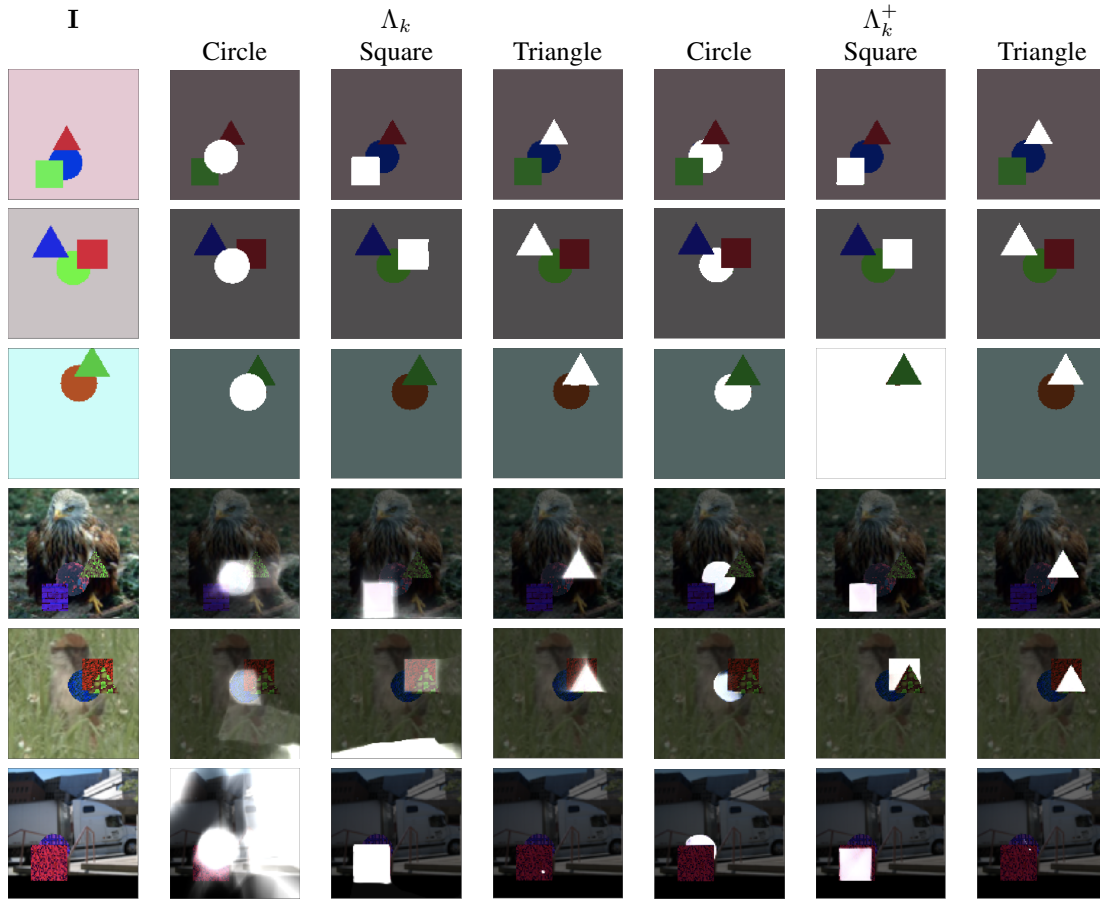
Figure 8: Flow Capsule inferred object shapes and their visible segmentations masks on Geo and Geo$^+$ test set. In the last row of each dataset where one of the objects is absent, the corresponding mask is empty.
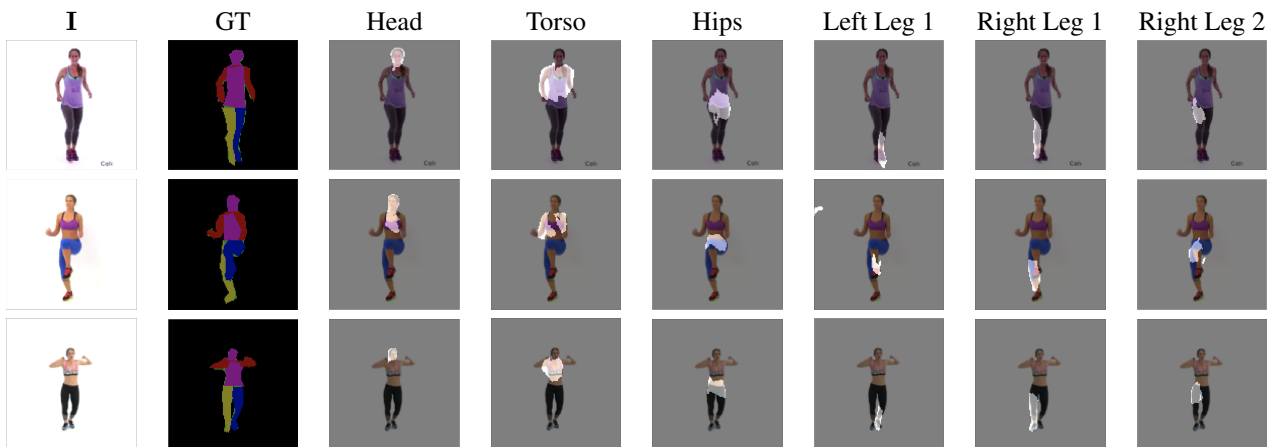

Figure 9: The ground truth segmentation masks along with sample Flow Capsule visible masks Λ_k^+ on the exercise test set.

On Geo$^+$, flow capsules achieve the IoU of **0.85** (circle), **0.93** (square), **0.90** (triangle) and overall of **0.89**. We were not able to train PSD on Geo$^+$ successfully. But the IoUs indicate added texture and background reduces our IoUs by about 10%, which still yields significantly more accurate segmentations compared to PSD on simple Geo.

On the Exercise dataset, FlowCapsule segments the body into roughly rigid parts. Fig. 9 shows decodings of some

Table 1: Intersection over union of inferred segmentation masks with ground truth masks for Geo and Exercise dataset. Flow Capsules consistently have better IoUs.

		R-NEM	PSD	Flow Capsules
Geo	Circle	0.54	0.93	**0.93**
	Square	0.56	0.82	**0.98**
	Triangle	0.58	0.90	**0.98**
	All	0.56	0.88	**0.96**
Exercise	Torso	0.32	0.57	**0.62**
	Left Leg	0.29	0.37	**0.59**
	Right Leg	0.23	0.34	**0.54**
	All	0.28	0.43	**0.58**

	Geo		Geo+	
	K=4	K=100	K=4	K=100
SCAE	0.48	0.59	0.49	0.51
FlowCapsule	**0.79**	**0.99**	**0.52**	**0.74**

Table 2: K-Means clustering accuracy with 4 or 100 clusters for Geo and Geo$^+$ datasets. Using FlowCapsule part representations consistently results in better classification accuracy as opposed to reconstruction based features of SCAE.

of the part capsules. The masks show that a single capsule consistently picks up the pixels of a part such as head or right leg regardless of the input image. So capsule identities are tied to a semantic part rather than a spatial position. Also, interestingly it breaks the object by their joints and separates the hips (lower torso) from the legs and from the upper torso. Table 1 verifies that Flow Capsule generated segments are more accurate than PSD and R-NEM. Considering that capsule segments are more fine grained than the provided ground truth masks, the quantitative IoUs are significantly better.

5.3. Unsupervised Classification

To further evaluate FlowCapsules, we replace the primary capsule autoencoder (bottom of the stack) in SCAE (Kosiorek et al., 2019) with the FlowCapsules. We call the new model FlowSCAE. Then we train the top SCAE object capsules to reconstruct the pose of FlowCapsules, just as SCAE is trained. We report unsupervised classification results using K-means clustering with various numbers of clusters on object Capsule presences in Table 2. We set the predicted label for each cluster to be the label that shows up the most in each cluster.

We compare our results against SCAE trained on reconstructing the image for the Geo, and Geo$^+$ datasets. We modified SCAE training slightly to produce coloured templates for the GEO dataset, and to produce textured templates for Geo$^+$ attached to the primary capsules (See supplementary material for the details). We compute the accuracy with numbers of clusters $K = 4$ and $K = 100$. We remark that for FlowSCAE with $K = 100$ the representations contained only 28 clusters for Geo dataset.

6. Conclusion

We introduce a method for learning capsule part representations (primary capsules) called FlowCapsules. The capsule

encoder takes as input a single frame and estimates a set of primary capsules, each comprising a shape mask in canonical coordinates, a pose transformation into image coordinates, and relative depth information. Training is done in a self-supervised manner from consecutive frames in video. To that end we use a Siamese architecture to estimate a parametric optical flow field between two frames, for which the flow is determined by the poses of corresponding part capsules in the two frames. Given the flow, one can use brightness constancy (or a related loss) to train the model. That said, the capsule encoder per se only takes one frame as input, and does not require flow at test time. Therefore, our capsule encoder learns to detect and encode the movable parts in a image. As such, this approach differs significantly from other approaches that essentially segment the flow field itself into moving parts (vs movable parts).

Our particular model makes our capsule encoder suitable for extracting a part encoding from a single image and the part encodings disentangle position from shape, which makes them applicable to both segmentation and classification tasks. We evaluate capsule part shape encodings in terms of segmentation and IoU of ground truth masks. We show that our method is not only able to learn an encoding but also able to generate more precise masks than previous related work. Furthermore, we evaluate the position encoding of our capsules by replacing the primary encoder of an SCAE (Kosiorek et al., 2019). The new, unsupervised object classification model, FlowSCAE, outperforms SCAE, especially when the complexity of images increase with additional background and texture.

Future directions include scaling this method to large scale video datasets. One way to do so would be to use off the shelf flow estimators for pre-training FlowCapsules. Given recent advances in the unsupervised flow estimation literature, e.g., incorporating different degrees of smoothness and accounting for the edges, we hope to further improve our part encodings. Another interesting direction would be to model the flow, masks and transformations in 3D using point clouds. A 3D setup would make natural image understanding easier for the model, especially with change of perspective and camera movement.

Acknowledgements

References

Ahmed, K. and Torresani, L. Star-caps: Capsule networks with straight-through attentive routing. In *Advances in Neural Information Processing Systems*, pp. 9098–9107, 2019.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.

Duarte, K., Rawat, Y., and Shah, M. Videocapsulenet: A simplified network for action detection. In *NeurIPS*, pp. 7610–7619, 2018.

Eslami, A. and Williams, C. A generative model for parts-based object segmentation. *NIPS*, 2012.

Hahn, T., Pyeon, M., and Kim, G. Self-routing capsule networks. In *Advances in Neural Information Processing Systems*, pp. 7656–7665, 2019.

Hinton, G. E., Krizhevsky, A., and Wang, S. D. Transforming auto-encoders. In *International Conference on Artifical Neural Networks*, 2011.

Hinton, G. E., Sabour, S., and Frosst, N. Matrix capsules with em routing. In *International Conference on Learning Representations*, 2018.

Huang, J. and Murphy, K. Efficient inference in occlusion-aware generative models of images. *ICLR Workshop (arXiv:1511.06362)*, 2016.

Jepson, A., Fleet, D. J., and Black, M. J. A layered motion representation with occlusion and compact spatial support. *ECCV*, pp. 692–706, 2002.

Kingma, D. and Ba, J. Adam: A method for stochastic optimization. *CoRR*, 2014.

Kosiorek, A., Sabour, S., Teh, Y. W., and Hinton, G. E. Stacked capsule autoencoders. *Advances in Neural Information Processing Systems*, pp. 15486–15496, 2019.

LaLonde, R. and Bagci, U. Capsules for object segmentation. *arXiv preprint arXiv:1804.04241*, 2018.

Mahendran, A., Thewlis, J., and Vedaldi, A. Self-supervised segmentation by grouping optical-flow. *ECCV Workshop*, 2018.

Picard, R. W., Kabir, T., and Liu, F. Real-time recognition with the entire brodatz texture database. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pp. 638–639. IEEE, 1993.

Rawlinson, D., Ahmed, A., and Kowadlo, G. Sparse unsupervised capsules generalize better. *CoRR*, 2018.

Sabour, S., Frosst, N., and Hinton, G. E. Dynamic routing between capsules. In *Advances in Neural Information Processing Systems*, 2017.

Singh, M. and Hoffman, D. Part-based representations of visual shape and implications for visual cognition. In Kellman, P. and Shipley, T. (eds.), *From fragments to objects: Segmentation and grouping in vision*, chapter 9, pp. 401–459. Elsevier Science, 2001.

Skafte, N. and Hauberg, S. Explicit disentanglement of appearance and perspective in generative models. *NeurIPS*, pp. 1018–1028, 2019.

Spelke, E. S. Principles of object perception. *Cognitive science*, 14(1):29–56, 1990.

Srivastava, N., Goh, H., and Salakhutdinov, R. Geometric capsule autoencoders for 3d point clouds. *arXiv preprint arXiv:1912.03310*, 2019.

Van Steenkiste, S., Chang, M., Greff, K., and Schmidhuber, J. Relational neural expectation maximization: Unsupervised discovery of objects and their interactions. *arXiv preprint arXiv:1802.10353*, 2018.

Vijayanarasimhan, S., Ricco, S., Schmid, C., Sukthankar, R., and Fragkiadaki, K. Sfm-net: Learning of structure and motion from video. *arXiv preprint arXiv:1704.07804*, 2017.

Wagemans, J., Elder, J. H., Kubovy, M., E., P. S., A., P. M., M., S., and von der Heydt, R. A century of gestalt psychology in visual perception: I. perceptual grouping and figure–ground organization. *Psychological Bulletin*, 138 (6):1172–1217, 2012.

Wang, J. and Adelson, E. H. Representing moving images with layers. *IEEE Transactions on Image Processing*, 3 (5):625–638, 1994.

Weiss, Y. Smoothness in layers: Motion segmentation using nonparametric mixture estima-tion. *IEEE CVPR*, pp. 520–526, 1997.

Xu, Z., Liu, Z., Sun, C., Murphy, K., Freeman, W. T., Tenenbaum, J. B., and Wu, J. Unsupervised discovery of parts, structure, and dynamics. In *International Conference on Learning Representations*, 2019.

Zhao, Y., Birdal, T., Deng, H., and Tombari, F. 3d point capsule networks. *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1009–1018, 2019.