

Learning to See by Moving

Pulkit Agrawal
UC Berkeley

pulkitag@eecs.berkeley.edu

João Carreira
UC Berkeley

carreira@eecs.berkeley.edu

Jitendra Malik
UC Berkeley

malik@eecs.berkeley.edu

Abstract

The current dominant paradigm for feature learning in computer vision relies on training neural networks for the task of object recognition using millions of hand labelled images. Is it also possible to learn features for a diverse set of visual tasks using any other form of supervision? In biology, living organisms developed the ability of visual perception for the purpose of moving and acting in the world. Drawing inspiration from this observation, in this work we investigated if the awareness of egomotion (i.e. self motion) can be used as a supervisory signal for feature learning. As opposed to the knowledge of class labels, information about egomotion is freely available to mobile agents. We found that using the same number of training images, features learnt using egomotion as supervision compare favourably to features learnt using class-label as supervision on the tasks of scene recognition, object recognition, visual odometry and keypoint matching.

“We must perceive in order to move, but we must also move in order to perceive”

J.J Gibson

1. Introduction

Recent advances in computer vision have shown that visual features learnt by neural networks trained for the task of object recognition using more than a million labelled images are useful for many computer vision tasks like semantic segmentation, object detection and action classification [18, 10, 1, 33]. However, object recognition is one among many tasks for which vision is used. For example, humans use visual perception for recognizing objects, understanding spatial layouts of scenes and performing actions such as moving around in the world. Is there something special about the task of object recognition or is it the case that useful visual representations can be learnt through other modes of supervision? Clearly, biological agents perform complex visual tasks and it is unlikely that they require external supervision in form of millions of labelled examples. Unlabelled visual data is freely available and in theory this data

can be used to learn useful visual representations. However, until now unsupervised learning approaches [4, 22, 29, 32] have not yet delivered on their promise and are nowhere to be seen in current applications on complex real world imagery.

Biological agents use perceptual systems for obtaining sensory information about their environment that enables them to act and accomplish their goals [13, 9]. Both biological and robotic agents employ their motor system for executing actions in their environment. Is it also possible that these agents can use their own motor system as a source of supervision for learning useful perceptual representations? Motor theories of perception have a long history [13, 9], but there has been little work in formulating computational models of perception that make use of motor information. In this work we focus on visual perception and present a model based on egomotion (i.e. self motion) for learning useful visual representations. When we say useful visual representations [34], we mean representations that possess the following two characteristics - (1) ability to perform multiple visual tasks and (2) ability of performing new visual tasks by learning from only a few labeled examples provided by an extrinsic teacher.

Mobile agents are naturally aware of their egomotion (i.e. self-motion) through their own motor system. In other words, knowledge of egomotion is “freely” available. For example, the vestibular system provides the sense of orientation in many mammals. In humans and other animals, the brain has access to information about eye movements and the actions performed by the animal [9]. A mobile robotic agent can estimate its egomotion either from the motor commands it issues to move or from odometry sensors such as gyroscopes and accelerometers mounted on the agent itself.

We propose that useful visual representations can be learnt by performing the simple task of correlating visual stimuli with egomotion. A mobile agent can be treated like a camera moving in the world and thus the knowledge of egomotion is similar to the knowledge of camera motion. Using this insight, we pose the problem of correlating visual stimuli with egomotion as the problem of predicting the camera transformation from the consequent pairs of im-

ages that the agent receives while it moves. Intuitively, the task of predicting camera transformation between two images should force the agent to learn features that are adept at identifying visual elements that are present in both the images (i.e. visual correspondence). In the past, features such as SIFT, that were hand engineered for finding correspondences were also found to be very useful for tasks such as object recognition [23, 19]. This suggests that egomotion based learning can also result in features that are useful for such tasks.

In order to test our hypothesis of feature learning using egomotion, we trained multilayer neural networks to predict the camera transformation between pairs of images. As a proof of concept, we first demonstrate the usefulness of our approach on the MNIST dataset [21]. We show that features learnt using our method outperform previous approaches of unsupervised feature learning when class-label supervision is available only for a limited number of examples (section 3.4) Next, we evaluated the efficacy of our approach on real world imagery. For this purpose, we used image and odometry data recorded from a car moving through urban scenes, made available as part of the KITTI [12] and the San Francisco (SF) city [7] datasets. This data mimics the scenario of a robotic agent moving around in the world. The quality of features learnt from this data were evaluated on four tasks (1) Scene recognition on SUN [37] (section 5.1), (2) Visual odometry (section 5.4), (3) Keypoint matching (section 5.3) and (4) Object recognition on Imagenet [31] (section 5.2).

Our results show that for the same amount of training data, features learnt using egomotion as supervision compare favorably to features learnt using class-label as supervision. We also show that egomotion based pretraining outperforms a previous approach based on slow feature analysis for unsupervised learning from videos [36, 14, 25]. To the best of our knowledge, this work provides the first effective demonstration of learning visual representations from non-visual access to egomotion information in real world setting.

The rest of this paper is organized as following: In section 2 we discuss the related work, in section 3, 4, 5 we present the method, dataset details and we conclude with the discussion in section 6.

2. Related Work

Past work in unsupervised learning has been dominated by approaches that pose feature learning as the problem of discovering compact and rich representations of images that are also sufficient to reconstruct the images [6, 3, 22, 32, 27, 30]. Another line of work has focused on learning features that are invariant to transformations either from video [36, 14, 25] or from images [11, 29]. [24] perform feature learning by modeling spatial transformations

using boltzmann machines, but donot evaluate the quality of learnt features.

Despite a lot of work in unsupervised learning (see [4] for a review), a method that works on complex real world imagery is yet to be developed. An alternative to unsupervised learning is to learn features using intrinsic reward signals that are freely available to a system (i.e self-supervised learning). For instance, [15] used intrinsic reward signals available to a robot for learning features that predict path traversability, while [28] trained neural networks for driving vehicles directly from visual input.

In this work we propose to use non-visual access to egomotion information as a form of self-supervision for visual feature learning. Unlike any other previous work, we show that our method works on real world imagery. Closest to our method is the the work of transforming auto-encoders [16] that used egomotion to reconstruct the transformed image from an input source image. This work was purely conceptual in nature and the quality of learned features was not evaluated. In contrast, our method uses egomotion as supervision by predicting the transformation between two images using a siamese-like network model [8].

Our method can also be seen as an instance of feature learning from videos. [36, 14, 25] perform feature learning from videos by imposing the constraint that temporally close frames should have similar feature representations (i.e. slow feature analysis) without accounting for either the camera motion or the motion of objects in the scene. In many settings the camera motion dominates the motion content of the video. Our key observation is that knowledge of camera motion (i.e. egomotion) is *freely* available to mobile agents and can be used as a powerful source of self-supervision.

3. A Simple Model of Motion-based Learning

We model the visual system of the agent with a Convolutional Neural Network (CNN, [20]). The agent optimizes its visual representations (i.e. updating the weights of the CNN) by minimizing the error between the egomotion information (i.e. camera transformation) obtained from its motor system and egomotion predicted using its visual inputs only. Performing this task is equivalent to training a CNN with two streams (i.e. Siamese Style CNN or SCNN[8]) that takes two images as inputs and predicts the egomotion that the agent underwent as it moved between the two spatial locations from which the two images were obtained. In order to learn useful visual representations, the agent continuously performs this task as it moves around in its environment.

In this work we use the pretraining-finetuning paradigm for evaluating the utility of learnt features. Pretraining is the process of optimizing the weights of a randomly initialized CNN for an auxiliary task that is not the same as the target




Figure 1: Exploring the utility of egomotion as supervision for learning useful visual features. A mobile agent equipped with visual sensors receives a sequence of images as inputs while it moves in its environment. The movement of the agent is equivalent to the movement of a camera. In this work, egomotion based learning is posed as the problem of predicting camera transformation from image pairs. The top and bottom rows of the figure show some sample image pairs from the SF and KITTI datasets that were used for feature learning.

task. Finetuning is the process of modifying the weights of a pretrained CNN for the given target task. Our experiments compare the utility of features learnt using egomotion based pretraining against class-label based and slow-feature based pretraining on multiple target tasks.

3.1. Two Stream Architecture

Each stream of the CNN independently computes features for one image. Both streams share the same architecture and the same set of weights and consequently perform the same set of operations for computing features. The individual streams have been called as Base-CNN (BCNN). Features from two BCNNs are concatenated and passed downstream into another CNN called as the Top-CNN (TCNN) (see figure 2). TCNN is responsible for using the BCNN features to predict the camera transformation between the input pair of images. After pretraining, the TCNN is removed and a single BCNN is used as a standard CNN for feature computation for the target task.

3.2. Shorthand for CNN architectures

The abbreviations Ck, Fk, P, D, Op refer to a convolutional(C) layer with k filters, a fully-connected(F) layer with k filters, pooling(P), dropout(D) and the output(Op) layers respectively. We used ReLU non-linearity after every convolutional/fully-connected layer, except for the output layer. The dropout layer was always used with dropout of 0.5. The output layer was a fully connected layer with number of units equal to the number of desired outputs. As an example of our notation, C96-P-F500-D refers to a network with 96 filters in the convolution layer followed by ReLU non-linearity, a pooling layer, a fully-connected layer with

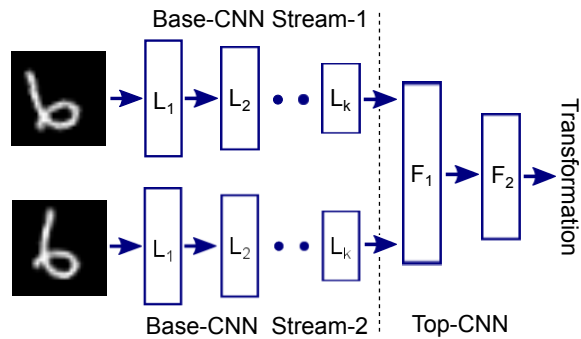


Figure 2: Description of the method for feature learning. Visual features are learnt by training a Siamese style Convolutional Neural Network (SCNN, [8]) that takes as inputs two images and predicts the transformation between the images (i.e. egomotion). Each stream of the SCNN (called as Base-CNN or BCNN) computes features for one image. The outputs of two BCNNs are concatenated and passed as inputs to a second multilayer CNN called as the Top-CNN (TCNN) (shown as layers F_1 , F_2). The two BCNNs have the same architecture and share weights. After feature learning, TCNN is discarded and a single BCNN stream is used as a standard CNN for extracting features for performing target tasks like scene recognition.

500 unit, ReLU non-linearity and a dropout layer. We used [17] for training all our models.

3.3. Slow Feature Analysis (SFA) Baseline

Slow Feature Analysis (SFA) is a method for feature learning based on the principle that useful features change slowly in time. We used the following contrastive loss formulation of SFA [8, 25],

$$L(x_{t_1}, x_{t_2}, W) = \begin{cases} D(x_{t_1}, x_{t_2}) & \text{if } |t_1 - t_2| \leq T \\ 1 - \max(0, m - D(x_{t_1}, x_{t_2})) & \text{if } |t_1 - t_2| > T \end{cases} \quad (1)$$

where, L is the loss, x_{t_1}, x_{t_2} refer to feature representations of frames observed at times t_1, t_2 respectively, W are the parameters that specify the feature extraction process, D is a measure of distance with parameter, m is a predefined margin and T is a predefined time threshold for determining whether the two frames are temporally close or not. In this work, x_t are features computed using a CNN with weights W and D was chosen to be the L2 distance. SFA pretraining was performed using two stream architectures that took pairs of images as inputs and produced outputs x_{t_1}, x_{t_2} as outputs from the two streams respectively.

3.4. Proof of Concept using MNIST

On MNIST, egomotion was emulated by generating synthetic data consisting of random transformation (translations and rotations) of digit images. From the training set of 60K images, digits were randomly sampled and then transformed using two different sets of random transformations to generate image pairs. CNNs were trained for predicting the transformations between these image pairs.

3.4.1 MNIST Data

For egomotion based pretraining, relative translation between the digits was constrained to be an integer value in the range [-3, 3] and relative rotation θ was constrained to lie within the range [-30°, 30°]. The prediction of transformation was posed as a classification task with three separate soft-max losses (one each for translation along X, Y axes and the rotation about Z-axis). SCNN was trained to minimize the sum of these three losses. Translations along X, Y were separately binned into seven uniformly spaced bins each. The rotations were binned into bins of size 3° each resulting into a total of 20 bins (or classes). For SFA based pretraining, image pairs with relative translation in the range [-1, 1] and relative rotation within [-3°, 3°] were considered to be temporally close to each other (see equation 1). A total of 5 million image pairs were used for both pretraining procedures.

3.4.2 MNIST Network Architectures

We experimented with multiple BCNN architectures and chose the optimal architecture for each pretraining method separately. For egomotion based pretraining, the two BCNN streams were concatenated using the TCNN: *F1000-D-Op*. Pretraining was performed for 40K iterations (i.e. 5M examples) using an initial learning rate of 0.01 which was reduced by a factor of 2 after every 10K iterations.

The following architecture was used for finetuning: *BCNN-F500-D-Op*. In order to evaluate the quality of BCNN features, the learning rate of all layers in the BCNN were set to 0 during finetuning for digit classification. Finetuning was performed for 4K iterations (which is equivalent to training for 50 epochs for the 10K labelled training examples) with a constant learning rate of 0.01.

3.4.3 MNIST Results

The BCNN features were evaluated by computing the error rates on the task of digit classification using 100, 300, 1K and 10K class-labelled examples for training. These sets were constructed by randomly sampling digits from the standard training set of 60K digits. For this part of the experiment, the original digit images were used (i.e. without any transformations or data augmentation). The standard test set of 10K digits was used for evaluation and error rates averaged across 3 runs are reported in table 1.

The BCNN architecture: C96-P-C256-P, was found to be optimal for egomotion and SFA based pretraining and also for training from scratch (i.e. random weight initialization). Results for other architectures are provided in the supplementary material. For SFA based pretraining, we experimented with multiple values of the margin m and found that $m = 10, 100$ led to the best performance. Our method outperforms convolutional deep belief networks [22], a previous approach based on learning features invariant to transformations [29] and SFA based pretraining.

4. Learning Visual Features From Egomotion in Natural Environments

We used two main sources of real world data for feature learning: the KITTI and SF datasets, which were collected using cameras and odometry sensors mounted on a car driving through urban scenes. Details about the data, the experimental procedure, the network architectures and the results are provided in sections 4.1, 4.2, 4.3 and 5 respectively.

4.1. KITTI Dataset

The KITTI dataset provided odometry and image data recorded during 11 short trips of variable length made by a car moving through urban landscapes. The total number of frames in the entire dataset was 23,201. Out of 11, 9

Method	# examples for finetuning			
	100	300	1000	10000
Autoencoder [17]	24.1	12.2	7.7	4.8
Ranzato et al. [29]	-	7.18	3.21	0.85
Lee et al. [22]	-	-	2.62	-
Train from Scratch	20.1	8.3	4.5	1.6
SFA (m=10)	11.2	6.4	3.5	2.1
SFA (m=100)	11.9	6.4	4.8	4.7
Egomotion (ours)	8.7	3.6	2.0	0.9

Table 1: Comparison of various pretraining methods on MNIST reveals that egomotion based pretraining outperforms many previous approaches for unsupervised learning. The performance is reported as the **error rate**.

sequences were used for training and 2 for validation. The total number of images in the training set was 20,501.

The odometry data was used to compute the camera transformation between pairs of images recorded from the car. The direction in which the camera pointed was assumed to be the Z axis and the image plane was taken to be the XY plane. X-axis and Y-axis refer to horizontal and vertical directions in the image plane. As significant camera transformations in the KITTI data were either due to translations along the Z/X axis or rotation about the Y axis, only these three dimensions were used to express the camera transformation. The task of predicting the transformation between pair of images was posed as a classification problem. The three dimensions of camera transformation were individually binned into 20 uniformly spaced bins each. The training image pairs were selected from frames that were at most ± 7 frames apart to ensure that images in any given pair would have a reasonable overlap. For SFA based pretraining, pairs of frames that were separated by at most ± 7 frames were considered to be temporally close to each other.

The SCNN was trained to predict camera transformation from pairs of 227×227 pixel sized image regions extracted from images of overall size 370×1226 pixels. For each image pair, the coordinates for cropping image regions were randomly chosen. Figure 1 illustrates typical image crops.

4.2. SF Dataset

SF dataset provides camera transformation between $\approx 136K$ pairs of images (constructed from a set of 17,357 unique images). This dataset was constructed using Google StreetView [7]. $\approx 130K$ image pairs were used for training and $\approx 6K$ pairs for validation.

Just like KITTI, the task of predicting camera transformation was posed as a classification problem. Unlike KITTI, significant camera transformation was found along all six dimensions of transformation (i.e. the 3 euler angles




Figure 3: Visualization of layer 1 filters learnt by egomotion based pretraining on (a) KITTI and (b) SF datasets. A large majority of layer-1 filters are color detectors and some of them are edge detectors. This is expected as color is a useful cue for determining correspondences between image pairs.

and the 3 translations). Since, it is unreasonable to expect that visual features can be used to infer big camera transformations, rotations between $[-30^\circ, 30^\circ]$ were binned into 10 uniformly spaced bins and two extra bins were used for rotations larger and smaller than 30° and -30° respectively. The three translations were individually binned into 10 uniformly spaced bins each. Images were resized to a size of 360×480 and image regions of size 227×227 were used for training the SCNN.

4.3. Network Architecture

BCNN closely followed the architecture of first five AlexNet layers [18]: *C96-P-C256-P-C384-C384-C256-P*. TCNN architecture was: *C256-C128-F500-D-Op*. The convolutional filters in the TCNN were of spatial size 3×3 . The networks were trained for 60K iterations with a batch size of 128. The initial learning rate was set to 0.001 and was reduced by a factor of two after every 20K iterations.

We term the networks pretrained using egomotion on KITTI and SF datasets as KITTI-Net and SF-Net respectively. The net pretrained on KITTI with SFA is called KITTI-SFA-Net. Figure 3 shows the layer-1 filters of KITTI-Net and SF-Net. A large majority of layer-1 filters are color detectors, while some of them are edge detectors. As color is a useful cue for determining correspondences between closeby frames of a video sequence, learning of color detectors as layer-1 filters is not surprising. The fraction of filters that detect edges is higher for the SF-Net. This is not surprising either, because higher fraction of images in the SF dataset contain structured objects like buildings and cars.

5. Evaluating Motion-based Learning

For evaluating the merits of the proposed approach, features learned using egomotion based supervision were compared against features learned using class-label and SFA based supervision on the challenging tasks of scene recognition, intra-class keypoint matching and visual odometry and

object recognition. The ultimate goal of feature learning is to find features that can generalize from only a few supervised examples on a new task. Therefore it makes sense to evaluate the quality of features when only a few labelled examples for the target task are provided. Consequently, the scene and object recognition experiments were performed in the setting when only 1-20 labelled examples per class were available for finetuning.

The KITTI-Net and SF-Net (examples of models trained using egomotion based supervision) were trained using only only $\approx 20K$ unique images. To make a fair comparison with class-label based supervision, a model with AlexNet architecture was trained using only 20K images taken from the training set of ILSVRC12 challenge (i.e. 20 examples per class). This model has been referred to as AlexNet-20K. In addition, some experiments presented in this work also make comparison with AlexNet models trained with 100K and 1M images that have been named as AlexNet-100K and AlexNet-1M respectively.

5.1. Evaluation on Scene Recognition

SUN dataset consisting of 397 indoor/outdoor scene categories was used for evaluating scene recognition performance. This dataset provides 10 standard splits of 5 and 20 training images per class and a standard test set of 50 images per class. Due to time limitation of running 10 runs of the experiment, we evaluated the performance using only 3 train/test splits.

For evaluating the utility of CNN features produced by different layers, separate linear (SoftMax) classifiers were trained on features produced by individual CNN layers (i.e. BCNN layers of KITTI-Net, KITTI-SFA-Net and SF-Net). Table 2 reports recognition accuracy (averaged over 3 train/test splits) for various networks considered in this study. KITTI-Net outperforms SF-Net and is comparable to AlexNet-20K. This indicates that given a fixed budget of pretraining images, egomotion based supervision learns features that are almost as good as the features using class-based supervision on the task of scene recognition. The performance of features computed by layers 1-3 (abbreviated as L1, L2, L3 in table 2) of the KITTI-SFA-Net and KITTI-Net is comparable, whereas layer 4, 5 features of KITTI-Net significantly outperform layer 4, 5 features of KITTI-SFA-Net. This indicates that egomotion based pretraining results into learning of higher-level features, while SFA based pretraining results into learning of lower-level features only.

The KITTI-Net outperforms GIST[26], which was specifically developed for scene classification, but is outperformed by Dense SIFT with spatial pyramid matching (SPM) kernel [19]. The KITTI-Net was trained using limited visual data ($\approx 20K$ frames) containing visual imagery of limited diversity. The KITTI data mainly contains images of roads, buildings, cars, few pedestrians, trees and some

Method	1	5	10	20
AlexNet-Scratch	1.1	3.1	5.9	14.1
KITTI-SFA-Net (Slowness)	1.5	3.9	6.1	14.9
KITTI-Net (Egomotion)	2.3	5.1	8.6	15.8

Table 3: Top-5 **accuracy** on the task of object recognition on the ILSVRC-12 validation set. AlexNet-Scratch refers to a net with AlexNet architecture initialized with randomly weights. The weights of KITTI-Net and KITTI-SFA-Net were learned using egomotion based and SFA based supervision on the KITTI dataset respectively. All the networks were finetuned using 1, 5, 10, 20 examples per class. The KITTI-Net clearly outperforms AlexNet-Scratch and KITTI-SFA-Net.

vegetation. It is in fact surprising that a network trained on data with such little diversity is competitive on classifying indoor and outdoor scenes with the AlexNet-20K that was trained on a much more diverse set of images. We believe that with more diverse training data for egomotion based learning, the performance of learnt features will be better than currently reported numbers.

The KITTI-Net outperformed the SF-Net except for the performance of layer 1 (L1). As it was possible to extract a larger number of image region pairs from the KITTI dataset as compared to the SF dataset (see section 4.1, 4.2), the result that KITTI-Net outperforms SF-Net is not surprising. Because KITTI-Net was found to be superior to the SF-Net in this experiment, the KITTI-Net was used for all other experiments described in this paper.

5.2. Evaluation on Object Recognition

If egomotion based pretraining learns useful features for object recognition, then a net initialized with KITTI-Net weights should outperform a net initialized with random weights on the task of object recognition. For testing this, we trained CNNs using 1, 5, 10 and 20 images per class from the ILSVRC-2012 challenge. As this dataset contains 1000 classes, the total number of training examples available for training for these networks were 1K, 5K, 10K and 20K respectively. All layers of KITTI-Net, KITTI-SFA-Net and AlexNet-Scratch (i.e. CNN with random weight initialization) were finetuned for image classification.

The results of the experiment presented in table 3 show that egomotion based supervision (KITTI-Net) clearly outperforms SFA based supervision (KITTI-SFA-Net) and AlexNet-Scratch. As expected, the improvement offered by motion-based pretraining is larger when the number of examples provided for the target task are fewer. These result show that egomotion based pretraining learns features useful for object recognition.

Method	Pretrain Supervision	#Pretrain	#Finetune	L1	L2	L3	L4	L5	L6	#Finetune	L1	L2	L3	L4	L5	L6
AlexNet-1M	Class-Label	1M	5	5.3	10.5	12.1	12.5	18.0	23.6	20	11.8	22.2	25.0	26.8	33.3	37.6
AlexNet-20K		20K	5	4.9	6.3	6.6	6.3	6.6	6.7	20	8.7	12.6	12.4	11.9	12.5	12.4
KITTI-SFA-Net	Slowness	20.5K	5	4.5	5.7	6.2	3.4	0.5	-	20	8.2	11.2	12.0	7.3	1.1	-
SF-Net	Egomotion	18K	5	4.4	5.2	4.9	5.1	4.7	-	20	8.6	11.6	10.9	10.4	9.1	-
KITTI-Net		20.5K	5	4.3	6.0	5.9	5.8	6.4	-	20	7.9	12.2	12.1	11.7	12.4	-
GIST [37]	Human	-	5	6.2					20	11.6						
SPM [37]	Human	-	5	8.4					20	16.0						

Table 2: Comparing the **accuracy** of neural networks pre-trained using motion-based and class-label based supervision for the task of scene recognition on the SUN dataset. The performance of layers 1-6 (labelled as L1-L6) of these networks was evaluated after finetuning the network using 5/20 images per class from the SUN dataset. The performance of the KITTI-Net (i.e. motion-based pretraining) fares favorably with a network pretrained on Imagenet (i.e. class-based pretraining) with the same number of pretraining images (i.e. 20K).

5.3. Evaluation on Intra-Class Keypoint Matching

Identifying the same keypoint of an object across different instances of the same object class is an important visual task. Visual features learned using egomotion, SFA and class-label based supervision were evaluated for this task using keypoint annotations on the PASCAL dataset [5].

Keypoint matching was computed in the following way: First, ground-truth object bounding boxes (GT-BBOX) from PASCAL-VOC2012 dataset were extracted and resized (while preserving the aspect ratio) to ensure that the smaller side of the boxes was of length 227 pixels. Next, feature maps from layers 2-5 of various CNNs were computed for every GT-BBOX. The keypoint matching score was computed between all pairs of GT-BBOX belonging to the same object class. For given pair of GT-BBOX, the features associated with keypoints in the first image were used to predict the location of the same keypoints in the second image. The normalized pixel distance between the actual and predicted keypoint locations was taken as the error in keypoint matching. More details about this procedure have been provided in the supp. materials.

It is natural to expect that accuracy of keypoint matching would depend on the camera transformation between the two viewpoints of the object (i.e. viewpoint distance). In order to make a holistic evaluation of the utility of features learnt by different pretraining methods on this task, matching error was computed as a function of viewpoint distance [35]. Figure 4 reports the matching error averaged across all keypoints, all pairs of GT-BBOX and all classes using features extracted from layers conv-3 and conv-4.

KITTI-Net trained only with 20K unique frames was superior to AlexNet-20K and AlexNet-100K and inferior only to AlexNet-1M. A net with AlexNet architecture initialized with random weights (AlexNet-Rand), surprisingly performed better than AlexNet-20K. One possible explanation for this observation is that with only 20K exam-

Method	Translation Acc.			Rotation Acc.		
	δX	δY	δZ	$\delta\theta_1$	$\delta\theta_2$	$\delta\theta_3$
SF-Net	40.2	58.2	38.4	45.0	44.8	40.5
KITTI-Net	43.4	57.9	40.2	48.4	44.0	41.0
AlexNet-1M	41.8	58.0	39.0	46.0	44.5	40.5

Table 4: Comparing the **accuracy** of various pretraining methods on the task of visual odometry.

ples, features learnt by AlexNet-20K only capture coarse global appearance of objects and are therefore poor at keypoint matching. SIFT has been hand engineered for finding correspondences across images and performs as well as the best AlexNet-1M features for this task (i.e. conv-4 features). KITTI-Net also significantly outperforms KITTI-SFA-Net. These results indicate that features learnt by egomotion based pretraining are superior to SFA and class-label based pretraining for the task of keypoint matching.

5.4. Evaluation on Visual Odometry

Visual odometry is the task of estimating the camera transformation between image pairs. All layers of KITTI-Net and AlexNet-1M were finetuned for 25K iterations using the training set of SF dataset on the task of visual odometry (see section 4.2 for task description). The performance of various CNNs was evaluated on the validation set of SF dataset and the results are reported in table 4.

Performance of KITTI-Net was either superior or comparable to AlexNet-1M on this task. As the evaluation was made on the SF dataset itself, it was not surprising that on some metrics SF-Net outperformed KITTI-Net. The results of this experiment indicate that egomotion based feature learning is superior to class-label based feature learning on the task of visual odometry.




Figure 4: Intra-class keypoint matching error as a function of viewpoint distance averaged over 20 PASCAL objects using features from layers conv3 (left) and conv4 (right) of various CNNs used in this work. Please see the text for more details.

6. Discussion

In this work, we have shown that egomotion is a useful source of intrinsic supervision for visual feature learning in mobile agents. In contrast to class labels, knowledge of egomotion is "freely" available. On MNIST, egomotion-based feature learning outperforms many previous unsupervised methods of feature learning. Given the same budget of pre-training images, on task of scene recognition, egomotion-based learning performs almost as well as class-label-based learning. Further, egomotion based features outperform features learnt by a CNN trained using class-label supervision on two orders of magnitude more data (AlexNet-1M) on the task of visual odometry and one order of magnitude more data on the task of intra-class keypoint matching. In addition to demonstrating the utility of egomotion based supervision, these results also suggest that features learnt by class-label based supervision are not optimal for all visual tasks. This means that future work should look at what kinds of pretraining are useful for what tasks.

One potential criticism of our work is that we have trained and evaluated high capacity deep models on relatively little data (e.g. only 20K unique images available on the KITTI dataset). In theory, we could have learnt better features by downsizing the networks. For example, in our experiments with MNIST we found that pretraining a 2-layer network instead of 3-layer results in better performance (table 1). In this work, we have made a conscious choice of using standard deep models because the main goal of this work was not to explore novel feature extraction architectures but to investigate the value of egomotion for learning visual representations on architectures known to perform well on practical applications. Future research focused on exploring architectures that are better suited for egomotion based learning can only make a stronger case for this line of work. While egomotion is freely available to mobile agents, there are currently no publicly available datasets as large as Imagenet. Consequently, we were un-

able to evaluate the utility of egomotion based supervision across the full spectrum of training set sizes. In appendix B of the supplementary materials we have provided a preliminary study of the performance of egomotion based supervision with varying amount of pretraining data.

In this work, we chose to first pretrain our models using a base task (i.e. egomotion) and then finetune these models for target tasks. An equally interesting setting is that of online learning where the agent has continuous access to intrinsic supervision (such as egomotion) and occasional explicit access to extrinsic teacher signal (such as the class labels). We believe that such a training procedure is likely to result in learning of better features. Our intuition behind this is that seeing different views of the same instance of an object (say) car, may not be sufficient to learn that different instances of the car class should be grouped together. The occasional extrinsic signal about object labels may prove useful for the agent to learn such concepts. Also, current work makes use of passively collected egomotion data and it would be interesting to investigate if it is possible to learn better visual representations if the agent can actively decide on how to explore its environment (i.e. active learning [2]).

Acknowledgements

This work was supported in part by ONR MURI-N00014-14-1-0671. Pulkit Agrawal was partially supported by Fulbright Science and Technology Fellowship. João Carreira was supported by the Portuguese Science Foundation, FCT, under grant SFRH/BPD/84194/2012. We gratefully acknowledge NVIDIA corporation for the donation of Tesla GPUs for this research.

Supplementary Materials

Supplementary Materials can be found at the first author's website.

References

- [1] P. Agrawal, R. Girshick, and J. Malik. Analyzing the performance of multilayer neural networks for object recognition. In *ECCV 2014*, pages 329–344. Springer, 2014.
- [2] R. Bajcsy. Active perception. *Proceedings of the IEEE*, 76(8):966–1005, 1988.
- [3] H. Barlow. Unsupervised learning. *Neural computation*, 1(3):295–311, 1989.
- [4] Y. Bengio, A. C. Courville, and P. Vincent. Unsupervised feature learning and deep learning: A review and new perspectives. *CoRR, abs/1206.5538*, 1, 2012.
- [5] L. Bourdev, S. Maji, T. Brox, and J. Malik. Detecting people using mutually consistent poselet activations. In *ECCV*, pages 168–181. Springer, 2010.
- [6] H. Bourlard and Y. Kamp. Auto-association by multilayer perceptrons and singular value decomposition. *Biological cybernetics*, 59(4-5):291–294, 1988.
- [7] D. M. Chen, G. Baatz, K. Koser, S. S. Tsai, R. Vedantham, T. Pylvanainen, K. Roimela, X. Chen, J. Bach, M. Pollefeys, et al. City-scale landmark identification on mobile devices. In *CVPR*, pages 737–744, 2011.
- [8] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, volume 1, pages 539–546. IEEE, 2005.
- [9] J. E. Cutting. *Perception with an eye for motion*, volume 177.
- [10] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. *arXiv preprint arXiv:1310.1531*, 2013.
- [11] P. Fischer, A. Dosovitskiy, and T. Brox. Descriptor matching with convolutional neural networks: a comparison to sift. *arXiv preprint arXiv:1405.5769*, 2014.
- [12] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun. Vision meets robotics: The kitti dataset. *IJRR*, 2013.
- [13] J. J. Gibson. *The Ecological Approach to Visual Perception*. Houghton Mifflin, 1979.
- [14] R. Goroshin, J. Bruna, J. Tompson, D. Eigen, and Y. LeCun. Unsupervised feature learning from temporal data. *arXiv preprint arXiv:1504.02518*, 2015.
- [15] R. Hadsell, P. Sermanet, J. Ben, A. Erkan, J. Han, B. Flepp, U. Muller, and Y. LeCun. Online learning for offroad robots: Using spatial label propagation to learn long-range traversability. In *RSS*, volume 11, page 32, 2007.
- [16] G. E. Hinton, A. Krizhevsky, and S. D. Wang. Transforming auto-encoders. In *ICANN*, pages 44–51. Springer, 2011.
- [17] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACM International Conference on Multimedia*, pages 675–678. ACM, 2014.
- [18] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.
- [19] S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *CVPR*, volume 2, pages 2169–2178. IEEE, 2006.
- [20] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.
- [21] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [22] H. Lee, R. Grosse, R. Ranganath, and A. Y. Ng. Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In *ICML*, pages 609–616. ACM, 2009.
- [23] D. G. Lowe. Object recognition from local scale-invariant features. In *ICCV*, volume 2, pages 1150–1157. IEEE, 1999.
- [24] R. Memisevic and G. E. Hinton. Learning to represent spatial transformations with factored higher-order boltzmann machines. *Neural Computation*, 22(6):1473–1492, 2010.
- [25] H. Mobahi, R. Collobert, and J. Weston. Deep learning from temporal coherence in video. In *ICML*, pages 737–744. ACM, 2009.
- [26] A. Oliva and A. Torralba. Building the gist of a scene: The role of global image features in recognition. *Progress in brain research*, 155:23–36, 2006.
- [27] B. A. Olshausen et al. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996.
- [28] D. A. Pomerleau. Alvin: An autonomous land vehicle in a neural network. Technical report, DTIC Document, 1989.
- [29] M. Ranzato, F. J. Huang, Y.-L. Boureau, and Y. LeCun. Unsupervised learning of invariant feature hierarchies with applications to object recognition. In *CVPR*, pages 1–8. IEEE, 2007.
- [30] M. Ranzato, A. Szlam, J. Bruna, M. Mathieu, R. Collobert, and S. Chopra. Video (language) modeling: a baseline for generative models of natural videos. *arXiv preprint arXiv:1412.6604*, 2014.
- [31] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015.
- [32] R. Salakhutdinov and G. E. Hinton. Deep boltzmann machines. In *AISTATS*, pages 448–455, 2009.
- [33] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *NIPS*, pages 568–576, 2014.
- [34] S. Soatto. Visual scene representations: Sufficiency, minimality, invariance and approximations. *arXiv preprint arXiv:1411.7676*, 2014.
- [35] S. Vicente, J. Carreira, L. Agapito, and J. Batista. Reconstructing pascal voc. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 41–48. IEEE, 2014.
- [36] L. Wiskott and T. J. Sejnowski. Slow feature analysis: Unsupervised learning of invariances. *Neural computation*, 14(4):715–770, 2002.
- [37] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR*, pages 3485–3492. IEEE, 2010.