

Carnegie Mellon University

From the Selected Works of Ole J Mengshoel

December, 2017

Semi-Supervised Convolutional Neural Networks for Human Activity Recognition

Ming Zeng
Tong Yu
Xiao Wang
Le T. Nguyen
Ole J. Mengshoel, et al.



Available at: https://works.bepress.com/ole_mengshoel/

Semi-Supervised Convolutional Neural Networks for Human Activity Recognition

Ming Zeng, Tong Yu, Xiao Wang, Le T. Nguyen, Ole J. Mengshoel, Ian Lane
Carnegie Mellon University, Moffett Field, CA 94043
{ming.zeng, tong.yu, xiao.wang, le.nguyen, ole.mengshoel, ian.lane}@sv.cmu.edu

Abstract—Labeled data used for training activity recognition classifiers are usually limited in terms of size and diversity. Thus, the learned model may not generalize well when used in real-world use cases. Semi-supervised learning augments labeled examples with unlabeled examples, often resulting in improved performance. However, the semi-supervised methods studied in the activity recognition literatures assume that feature engineering is already done. In this paper, we lift this assumption and present two semi-supervised methods based on convolutional neural networks (CNNs) to learn discriminative hidden features. Our semi-supervised CNNs learn from both labeled and unlabeled data while also performing feature learning on raw sensor data. In experiments on three real world datasets, we show that our CNNs outperform supervised methods and traditional semi-supervised learning methods by up to 18% in mean F1-score (F_m).

Keywords-Human Activity Recognition; Deep Neural Networks; Semi-Supervised Learning; Convolutional Neural Networks

I. INTRODUCTION

Human activity recognition (HAR) is an important application area for mobile, on-body, and worn mobile technologies. Supervised learning for human activity recognition has shown great promise. Among supervised methods, deep neural networks (DNNs) have emerged as a method with much potential, in that they are less dependent on clever feature engineering and has strong generalization ability [1] compared to other supervised methods [2], [3].

Unfortunately, the problem of data labeling remains. Compared to many other machine learning applications, the problem of data labeling for HAR is substantial, since human activity data sets typically (i) have few labeled samples and (ii) are highly personal and varying.

(i) Activity data sets typically have very few labeled examples for some activities. Thus, they may not characterize well the distribution of test data collected in different situations than the training data. For example, the labeled training data may only cover walking at certain speeds. In reality, humans walk at a range of speeds. They can walk slowly when being relaxed and can walk very fast when in a hurry. The problem of limited labels is even more severe for models with high parameter complexity, such as deep neural networks.

(ii) Activity data sets are highly personal and varying, because people may perform the same activity in very different ways. For example, what one person considers jogging

may be very similar to what another person considers fast walking. With a model trained only on data where a human walks at normal speed, it is very difficult to correctly predict the behavior of a human walking in a hurry. Walking in a hurry can easily be confused with running, especially when little data of walking in a hurry is collected for training.

To address challenges (i) and (ii), many semi-supervised learning methods have been proposed to leverage the abundance of unlabeled data and provide higher generalizability. Although the labeled data of walking in a hurry may be limited, there are large amounts of *unlabeled* data recording the behavior of walking in a hurry. Semi-supervised learning from both labeled and unlabeled data can thus potentially provide better predictions for human walking in a hurry, compared to supervised learning using only labeled data.

When labeled data is limited, we may potentially improve HAR performance via adjustments to labeled data's feature representations with unlabeled data, so-called feature learning. In contrast, previous semi-supervised HAR approaches usually rely on handcrafted features [4], [5], [6]. With handcrafted features, the benefit of the unlabeled data is limited, since there is no opportunity for feature learning with the unlabeled data.

In this paper, we study how to train accurate and generalizable DNNs with limited labeled data and large scale unlabeled data for HAR. Specifically, we present two semi-supervised deep convolutional neural network methods, the convolutional encoder-decoder (CNN-Encoder-Decoder) and the convolutional ladder network (CNN-Ladder). The contributions of our work are the following.

- To our best knowledge, this is the first paper to leverage unlabeled data in CNNs in HAR applications. We utilize unlabeled data in both feature learning and model learning using CNN-Encoder-Decoder and CNN-Ladder architectures for semi-supervised HAR.
- The presented methods can achieve up to 18% F1-score improvement compared to baseline methods, on three real-world activity recognition datasets.
- To understand why our methods improve F1-score, we show the importance of adjusting low level features based on unlabeled data in semi-supervised HAR. Besides, we visualize the features in the last layers of CNN-Ladder and CNN to demonstrate that better high-level features can be learned with unlabeled data added.

II. RELATED WORKS

In this section, we discuss related work on (i) machine learning in HAR and (ii) semi-supervised learning in HAR.

A. Machine Learning for Activity Recognition

In early studies of HAR [7], machine learning models using handcrafted features shows good performance. Raw sensor data is collected from various sensors on mobile devices. From this collected data, handcrafted features are designed using domain knowledge. With the handcrafted features, machine learning models, such as random forest, naive Bayes, or SVMs, are trained and used in HAR.

Designing handcrafted features requires domain knowledge [8]. Therefore, it is desirable to develop a systematic feature learning approach to model the time series signals in HAR [9]. Deep neural networks (DNNs) are emerging feature extraction approaches to HAR, and they have made great advances in many domains [10]. They are also applied to HAR (e.g., [11], [12], [13]). The first HAR deep learning approach [11] explores unsupervised feature extraction. It outperforms principal component analysis (PCA) and statistical features. After that, convolutional neural networks (CNNs) became popular due to their locality preservation and translation invariance. A 1D CNN is used to model sensor modality [12] while a 2D CNN regards the set of signals as an image and handles multichannel sensor readings [9]. In order to capture the temporal dependencies of the sensor data, deep recurrent networks, especially long short-term memory cells (LSTMs), have achieved promising performance in HAR [3], [14]. However, due to the complexity of LSTM, they require much labeled data to avoid overfitting.

B. Semi-Supervised Learning

In semi-supervised learning, the model is trained on both labeled and unlabeled data [15]. Utilizing unlabeled data may improve a model's generalization ability.

Semi-supervised learning has been applied to HAR. An on-line adaptation method is proposed for semi-supervised learning for HAR [16]. The self-learning based approaches [4], [17] iteratively annotate the unlabeled data and selectively add them to the training dataset. The graph-based approach [5] connects labeled and unlabeled data and builds multiple graphs to propagate the labels based on similarity between features. However, these approaches treat the label propagation and classification as two separate processes. Thus, correlations between labeled data and unlabeled data may be ignored in the model.

A recent semi-supervised method, ladder networks [18], can simultaneously train a deep auto encoder on an unlabeled dataset and a neural network on a labeled dataset. The ladder network shows superior performance in semi-supervised image classification for the MNIST and CIFAR-10 dataset.

III. SEMI-SUPERVISED CNN BASED MODELS

We adopt the CNN since it provides stable latent representations at each network level, which preserved locality. It also has great potential to identify the various salient patterns of activity signals [9]. We use the multi-sensor based CNN structure [9] for both our supervised and semi-supervised learning approaches.

A. CNN for Supervised Learning

Consider a dataset with N labeled sliding windows $(x_1, t_1), (x_2, t_2), \dots, (x_N, t_N)$, where x_i is a sliding window input with length T and t_i is the activity label. A CNN maps the input $x_i = [x_{i1}, x_{i2}, \dots, x_{iT}]$ to hidden values $z_i^l = [z_{i1}^l, z_{i2}^l, \dots, z_{id}^l]$ by convolutional kernels (to be learned in the training phase), where l denotes the l -th layer (the input x_i is also regarded as 0-th layer, z^0). The CNN structure can be represented as:

$$z_i^{(1)}, \dots, z_i^{(L)}, y_i = \text{CNN}(x_i), \quad (1)$$

where $\text{CNN}(\cdot)$ contains at least one temporal convolutional layer, one pooling layer, and at least one fully connected layer prior a top-level softmax classifier. Then the supervised CNN cost function is of the form:

$$C_s = -\frac{1}{N} \sum_{i=1}^N \log P(y_i = t_i | x_i). \quad (2)$$

It requires a lot of labeled data to train a good CNN model.

B. CNN Encoder-Decoder for Unsupervised Learning

Assume that we also have M unlabeled examples $x_{N+1}, x_{N+2}, \dots, x_{N+M}$. The CNN-Encoder-Decoder consists of an encoder mapping f and a decoder mapping g . The encoder adopts the CNN feed-forward process while the decoder contains upsampling and convolution operations. Our encoder-decoder structure is similar to a denoising autoencoder (DAE) [19]. In the training, noise is injected into each layer in the network (including the input layer). The CNN-Encoder-Decoder minimizes the difference between the clean input x_i and the reconstructed decoder output \hat{x}_i . Therefore, we have the cost function:

$$C_r^{(0)} = \frac{\lambda}{M} \sum_{i=N+1}^{N+M} \|\hat{x}_i - x_i\|_2^2, \quad (3)$$

where \hat{x}_i is the reconstructed input. The decoder in the CNN-Encoder-Decoder [20] contains upsampling for max-pooling decoding and another convolutional operation for deconvolution. The upsampling uses the memorized max-pooling indices from the corresponding encoder feature map(s) to produce sparse feature maps(s) as an input of the convolutional layer in the CNN-Encoder-Decoder [20]. Then the sparse features are convolved with a trainable decoder filter bank to produce dense features.

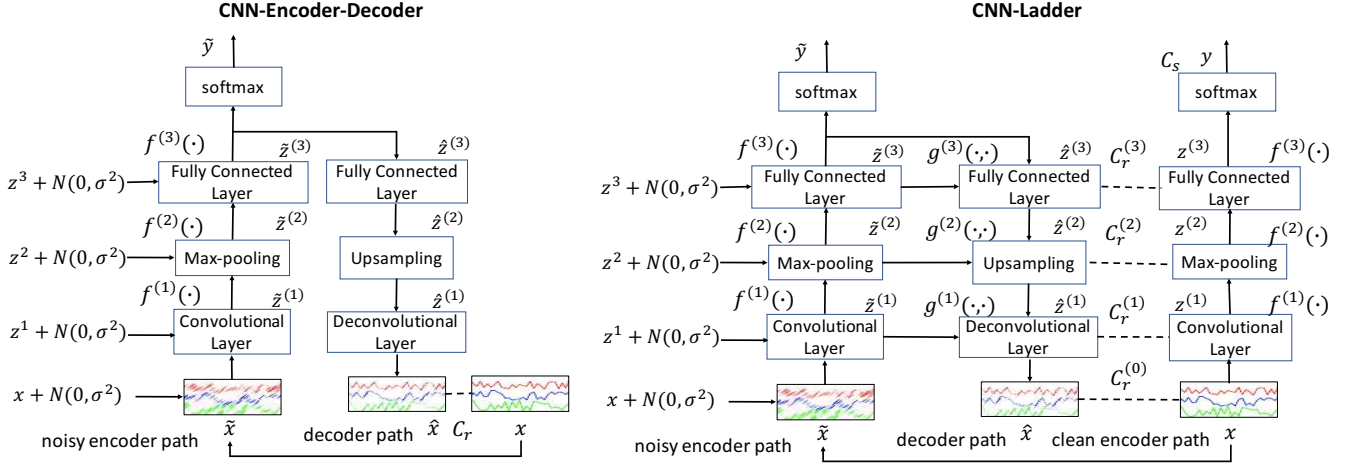


Figure 1: Structure of the CNN-Encoder-Decoder (left) and CNN-Ladder (right) applied to HAR. CNN-Ladder has two kinds of connections: *lateral connections include $g^{(l)}(\cdot, \cdot)$ and reconstructed cost function $C_r^{(l)}$* . Vertical connections contain clean encoder path ($x \rightarrow z^{(l)} \rightarrow y$), noisy encoder ($\tilde{x} \rightarrow \tilde{z}^{(l)} \rightarrow \tilde{y}$) path and decoder path ($\tilde{z}^{(3)} \rightarrow \tilde{z}^{(2)} \rightarrow \hat{x}$). The noisy encoder and clean encoder share the same mapping function f . The function g is the denoising function, which is for reconstructing the clean input from high-level representation, $\tilde{z}^{(3)}$. When we only consider the vertical connections and the lateral cost in the bottom $C_r^{(0)}$, the CNN-Ladder reduces to the CNN-Encoder-Decoder model (left).

C. Semi-Supervised CNN-Encoder-Decoder for HAR

We combine the supervised CNN and CNN-Encoder-Decoder to perform semi-supervised learning for HAR. Besides a set of labeled pairs $\{(x_i, t_i) \mid 1 \leq i \leq N\}$, semi-supervised learning [15] uses unlabeled data $\{x_i \mid N+1 \leq i \leq N+M\}$ to help in training a classifier.

In the case of a semi-supervised CNN-Encoder-Decoder, there are three paths for the labeled and unlabeled data: The clean encoding, noisy encoding, and the decoding:

$$z_i^{(1)}, \dots, z_i^{(L)} = \text{Encoder}_{\text{clean}}(x_i) \quad (4)$$

$$\tilde{x}_i, \tilde{z}_i^{(1)}, \dots, \tilde{z}_i^{(L)} = \text{Encoder}_{\text{noisy}}(x_i) \quad (5)$$

$$\hat{x}_i = \text{Decoder}(\tilde{z}_i^{(L)}). \quad (6)$$

Both labeled and unlabeled clean data pass through the clean encoder path to compute hidden variables in the middle layers, z_i^l . For the noisy encoder path, both labeled and unlabeled data are corrupted by Gaussian noise and then transformed to a more abstract representation, \tilde{z}_i^l , by the noisy encoder. For labeled data ($\tilde{x}_i, 1 \leq i \leq N$), we carry out the prediction for labeled data on the top-level softmax classifier based on cross entropy cost. The predicted label is denoted by \tilde{y}_i . For the noisy unlabeled data ($\tilde{x}_i, N+1 \leq i \leq N+M$), the decoder tries to reconstruct it (\hat{x}_i) to be the same as the corresponding clean input (x_i). We use square error to evaluate this reconstruction error. The clean and noisy encoder paths share the same parameters, only the inputs are different in Fig 1. (When we only consider the vertical connections and the lateral cost, CNN-Ladder in Fig 1 reduces to CNN-Encoder-Decoder.)

The CNN-Encoder-Decoder the cost function involves the supervised cross entropy cost from labeled data in the supervised CNN and the unsupervised denoising square error cost between the clean input and its noisy reconstruction output. Thus the cost function is

$$C_e = C_s + \lambda C_r^{(0)} \\ = -\frac{1}{N} \sum_{i=1}^N \log P(\tilde{y}_i = t_i | x_i) + \frac{\lambda}{M} \sum_{i=N+1}^{N+M} \|\hat{x}_i - x_i\|_2^2, \quad (7)$$

where the supervised cost C_s is the averaged cross entropy of the noisy output \tilde{y}_i matching the target t_i given the input x_i . The unsupervised cost C_r is the averaged square error between the reconstruction output \hat{x}_i and the clean input x_i . By using a semi-supervised CNN-Encoder-Decoder, we can potentially learn the network and features simultaneously from the data.

D. Semi-Supervised CNN-Ladder for HAR

The semi-supervised Convolutional Ladder Network (CNN-Ladder) contains two kinds of connections: the vertical connections and the lateral connections (Fig 1). The vertical connections have clean and noisy encoders (Eq 4, Eq 5) and a decoder. The reconstruction $\hat{z}_i^{(l)}$ in the decoder is not only inferred from the upper layer $\hat{z}_i^{(l+1)}$, but also estimated from its corresponding layer in the noisy encoder. The estimation is a linear function $\hat{z}^{(l)} = g(\tilde{z}^{(l)}, \hat{z}^{(l+1)})$, where $\tilde{z}^{(l)}$ is the lateral noisy signal in the encoder and $\hat{z}^{(l+1)}$ is the reconstruction of its upper layer by batch normalization [18]. These vertical skip-connections enable us to find better

middle-level representations compared to regular encoder-decoder structures.

To improve the middle-level features reconstruction in the CNN-Encoder-Decoder, we also force the intermediate layers in the decoder to be similar to the corresponding layers in the encoder. In other words, the cost function of CNN-Ladder is

$$C_l = C_s + \sum_{l=0}^L \lambda_l C_r^{(l)} = -\frac{1}{N} \sum_{i=1}^N \log P(\tilde{y}_i = t_i | x_i) + \frac{1}{M} \sum_{i=N+1}^{N+M} \sum_{l=0}^L \lambda_l \|z_i^{(l)} - z_i^{(l)}\|_2^2 \quad (8)$$

If we train neural networks on limited unlabeled data, learned hidden features may have high variance and can be unstable. With the constraints from the lateral connection, the CNN-Ladder makes every layer, $C_r^{(l)}$, contribute to the cost function. As a result, more stable hidden features can be learned from large amount of unlabeled data. Stable hidden features can generate accurate representation of the middle level features, and lead to precise recognition of complicated activities. For example, jumping jack activity prediction relies on stable and accurate representation of sub-components (spreading hands and legs, and clapping hands).

IV. EXPERIMENTS

We validate our HAR approaches on three public datasets. First, we compare our methods to other neural network methods for HAR in a supervised learning setting. Second, we compare our methods to traditional semi-supervised learning methods for HAR. Third, we conducted experiments with varying amounts of labeled and unlabeled data, to understand the usability of our methods. Fourth, we discuss why our methods perform better than traditional semi-supervised learning methods in utilizing the unlabeled data for semi-supervised HAR. Deep learning (CNN, Pretrained CNN, Pseudo-label CNN, CNN-Encoder-Decoder, CNN-Ladder) is performed on a server equipped with a Tesla K20c GPU and 64G memory. The traditional learning algorithms (LR, Self-training) are run on the same server with an Intel Xeon E5 CPU. The implementation of CNN-Ladder is based on the Ladder Networks.¹

A. Datasets

The raw sensor data is segmented by a common sliding window technique. The window size is 2 seconds with 50% overlap. Data within each window is denoted as an *example*. All the results are averaged using leave-one-subject-out cross validation. To ensure that labeled training data includes all the activity classes, we construct balanced labeled training datasets. The datasets used are as follows.

The **ActiTracker** [21] dataset contains 6 daily activities collected in a controlled laboratory environment. The activities are “jogging,” “walking,” “ascending stairs,” “descending stairs,” “sitting” and “standing.” The data are recorded from 36 users, with a 20Hz sampling rate resulting in 1,098,207 examples. After segmentation, there are around 110,000 examples (sliding windows). The number of examples for testing varies from 1,000 to 5,000.

The **PAMAP2** dataset [22] consists of 12 lifestyle activities (“walking,” “lying down,” “knees bending,” etc.) by 9 participants. Accelerometer, gyroscope, magnetometer, temperature, and heart rate data are recorded from inertial measurement units located on the hand, chest and ankle over 10 hours, resulting in 52 dimensions. The number of examples is 3,850,505. To have a temporal resolution comparable to the ActiTracker dataset, we downsampled the data to 33.3Hz, resulting in around 33,000 examples. The number of examples of test data in each experiment is around 4,500.

The **mHealth** dataset [23] contains recordings from 10 participants while performing 12 physical activities, including daily life activities (“standing,” “lying down,” etc.) and exercise activities (“cycling,” “jogging,” etc). Accelerometers, gyroscopes, magnetometer and ECG data are recorded from inertial measurement units placed on a participant’s chest, right wrist and left ankle. The data has 43,744 examples with 23 dimensions. In our experiment, we downsampled the data to 20Hz, resulting in around 8,000 examples. The number of examples used for testing is around 1,000.

B. Experimental Setup

We consider these supervised learning baselines:

- **Logistic Regression (LR)** [7]: We using traditional logistic regression for supervised learning in combination with statistical features (mean, standard deviation, correlation, max, min).
- **Supervised Convolutional neural network (CNN)** [9]: The structure of the supervised CNN is the same as the clean path in our CNN-Ladder.³

We also study traditional semi-supervised learning baselines.

- **Unsupervised Pretrained CNN** [24]: The pretrained CNN uses the unlabeled data to initialize the network parameters. We use an unsupervised pretraining method similar to multi-layer perceptron (MLP) pretraining. In the first step, the pretrained CNN uses the unlabeled data to perform encoding and decoding with the CNN structure to initialize the parameters of the network.
- **Self-training method with logistic regression (Self-training)** [4]: In self-training, an LR classifier is first trained using (a small amount of) labeled data. Then

¹<https://github.com/CuriousAI/ladder>

³Network structure: convv:40:5:1:1-maxpool:2:2-convv:50:3:1:1-maxpool:2:2-convv:20:3:1:1-convv:50:1:1:1-fc. [18]

Data	Previous Papers ^{1 2 3}		This Paper	
	LR	CNN	LR	CNN
ActiTracker (Accuracy)	78.10 ¹	90.88 ²	89.27	93.84
PAMAP2 (F_m)	-	93.70 ³	86.86	92.24

Table I: We reproduce the results of LR and CNN on ActiTracker and PAMAP2, reported in [Kwapisz et al. 2011, Zeng et al. 2014, Hammerla et al. 2016]. We show the results under the previous papers’ settings. On ActiTracker, the results are in accuracy. On PAMAP2, the results are in F_m .

we use the trained LR model to predict the labels of unlabeled data. In each iteration, predictions with high confidence are added to the labeled training set, where these predictions are now considered as the labels. In our experiments, the confidence threshold is 0.95.

- **Pseudo-label** [25]: The pseudo-label approach is essentially a self-training method. The predicted labels of the unlabeled data are used in a fine-tuning phase to improve the recognition performance.

A result is averaged across all leave-one-subject-out cross validation experiments. Thus, in each experiment, we use one user for test and the rest of the users for training. We evaluate the results using mean F1-score because the activity datasets are highly biased. The F1-score is a harmonic mean of precision and recall. The **mean F1-score**, F_m , is the mean F1-score across all the classes:

$$F_m = \frac{2 \cdot \text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}} \quad (9)$$

where for a given class

$$\text{precision} = \frac{TP}{TP + FP}, \quad \text{recall} = \frac{TP}{TP + FN}.$$

Here, FP and FN are counts of False Positives and False Negatives, respectively.

Table I shows that our baseline of supervised CNN is comparable to the results in previous papers. We also evaluate our CNN baseline on all users, instead of using the setting in the previous works. The mean F1 scores are 79.54, 75.38 and 92.83 on ActiTracker, PAMAP2 and mHealth, respectively.

C. Comparing with Supervised Methods

We compare our methods to several supervised methods, to study how our methods utilize unlabeled data in HAR. The baseline methods LR and CNNs do not use unlabeled data.

¹We only carry out 10-fold cross validation for [Kwapisz et al. 2011].

²We only carry out 10-fold cross validation for [Zeng et al. 2014].

³User 6 is for the test set, user 5 is for the validation set and the rest of the users are used for the training set [Hammerla et al. 2016].

The results are shown in Table II. On all the three datasets, CNN-Encoder-Decoder and CNN-Ladder perform consistently better than LR and CNN. In particular, CNN-Ladder achieves 17.64%, 3.59%, 9.65% improvements in mean F1-score on the three datasets, compared to the best of LR and CNN. Second, CNN-Ladder has higher F_m score than CNN-Encoder-Decoder on the three datasets.

Those results suggest that CNN-Encoder-Decoder and CNN-Ladder can effectively make use of the unlabeled data, to significantly improve accuracy. CNN-Ladder performs better than CNN-Encoder-Decoder, perhaps because better hidden features are trained. In CNN-Ladder, the loss function considers the difference between each layer of CNN and its decoder, while CNN-Encoder-Decoder only considers the difference between the final reconstructed output and the original input.

D. Comparing with Traditional Semi-supervised Methods

We compare our methods to traditional semi-supervised methods, to study how our methods can utilize the same unlabeled data to achieve more accurate predictions.

The comparisons between Pretrained CNN, Self-Training, Pseudo-Label, CNN-Encoder-Decoder, and CNN-Ladder are shown in Table II. It can be observed that CNN-Encoder-Decoder and CNN-Ladder perform better than Pretrained CNN, Self-Training, and Pseudo-Label. Specifically, CNN-Ladder achieves about 16.46%, 4.11%, 8.5% improvements in mean F1-scores on the three datasets, compared to the best of Pretrained CNN, Self-Training, and Pseudo-Label.

One disadvantage of Self-Training and Pseudo-Label that we observed is that these iterative methods need careful selection of the confidence threshold. If the confidence threshold is not appropriately selected, some unlabeled data will be assigned wrong labels and the errors will propagate in later iterations. However, in semi-supervised CNNs, no confidence threshold is needed and all available unlabeled data are input together with labeled data to train the models. Without using confidence thresholds, training neural network requires less domain knowledge and is much easier compared to training Self-Training and Pseudo-Label models.

E. Varying Amount of Labeled Data

In this section, we study the performance of our models trained with varying amounts of *labeled* data. We evaluate the F_m score of supervised CNN, CNN-Encoder-Decoder and CNN-Ladder trained on 50, 100, 200, 500, and 1,000 labeled examples. The rest of the samples in the training set are regarded as unlabeled.

Figure 2 shows the F_m trend when we vary the number of labeled examples. There are three observations. First, the F_m scores of supervised CNN, CNN-Encoder-Decoder and CNN-Ladder generally improve when we have more labeled examples. Second, with the same number of labeled examples, CNN-Encoder-Decoder, and CNN-Ladder usually

	Supervised		Semi-Supervised			Our Semi-Supervised		Improvement	
	LR	CNN	Pretrained CNN	Self-Training	Pseudo-Label	CNN-Encoder-Decoder	CNN-Ladder	Δ Supervised	Δ Semi-Supervised
ActiTracker	39.34	48.68	49.86	41.52	46.00	63.58	66.32	17.64	16.46
PAMAP2	51.31	50.22	48.54	47.86	50.79	52.68	54.90	3.59	4.11
mHealth	57.73	59.73	60.88	59.43	60.31	66.61	69.38	9.65	8.50

Table II: The F_m score of supervised methods (LR and CNN), traditional semi-supervised methods (Self-training and Pseudo-label) and our presented methods (CNN-Encoder-Decoder and CNN-Ladder). We circle the best F_m scores from supervised, semi-supervised and our semi-supervised approaches, respectively. Both of our methods (CNN-Encoder-Decoder, CNN-Ladder) are significantly better compared to the CNN and the Pretrained CNN with p -value < 0.05 .




Figure 2: The F_m scores of CNN, CNN-Encoder-Decoder, and CNN-Ladder, with varying number of labeled examples. The F_m scores of supervised CNN on all labeled training examples are also shown as red lines.

achieves higher F_m scores than CNN. Third, when CNN-Ladder is learned from 1,000 examples, its mean F_m score is already very competitive with supervised CNN learned from more than 100,000, 10,000, and 8,000 labeled examples from ActiTracker, PAMAP2 and mHealth, respectively. These results indicate that compared to CNN, CNN-Ladder can achieve similar accuracy but with much smaller number of labeled examples.

F. Varying Amount of Unlabeled Data

We now study the performance of our models trained with varying amounts of *unlabeled* data. We evaluate the F_m score of supervised CNN, CNN-Encoder-Decoder, and CNN-Ladder trained on 50 labeled examples and varying amounts of unlabeled examples. On ActiTracker, the number of unlabeled examples varies from 100 to 50,000. On PAMAP2 and mHealth, the number varies from 100 to 10,000, as these two datasets are relatively small.

Figure 3 shows the experimental results. With an increasing amount of unlabeled data, the F_m score typically improves for both CNN-Encoder-Decoder and CNN-Ladder. This suggests that better latent features in the auto-encoder can be trained with more unlabeled examples and help adjust the latent CNN features, thereby improving accuracy.

G. The Impact of Adjusting Features in Different Layers

We now study the importance of adjusting different layers in CNN-Ladder with unlabeled data. Specifically, we adjust

the λ_l of CNN-Ladder in Equation 8 to observe the impact of making the latent features between CNN and autoencoder more or less similar, in different layers. We run a set of experiments for different layers l , where $l \in \{0, 1, \dots, L\}$. In each experiment, we emphasize layer l by setting $\lambda_l = 1$ and $\lambda_k = 0.1$, where $k \in \{0, 1, \dots, l-1\} \cup \{l+1, \dots, L\}$. In our CNN-Ladder, $L = 9$.

The resulting F_m scores when varying the weights of different layers of CNN-Ladder are shown in Figure 5. A high F_m score can typically be achieved by setting a large λ_l for the layers representing low-level features. This indicates that low-level features of the neural networks can be much improved by using the unlabeled data.

In contrast, utilizing the unlabeled data for low-level features is missed in traditional semi-supervised learning methods for HAR, such as Self-Training. Self-Training uses the unlabeled data only after feature engineering is already done. That is, the handcrafted features are independent from whether unlabeled data is available or not. In a similar way, low-level features of traditional neural network methods, such as CNN, may be not as good as the low-level features of CNN-Ladder in the semi-supervised HAR.

H. How Does CNN-Ladder Achieve Higher F_m ?

As discussed in Section IV-G, CNN-Ladder can adjust low-level features with unlabeled data, while traditional semi-supervised methods' low-level features are independent from unlabeled data. This section seeks to better understand

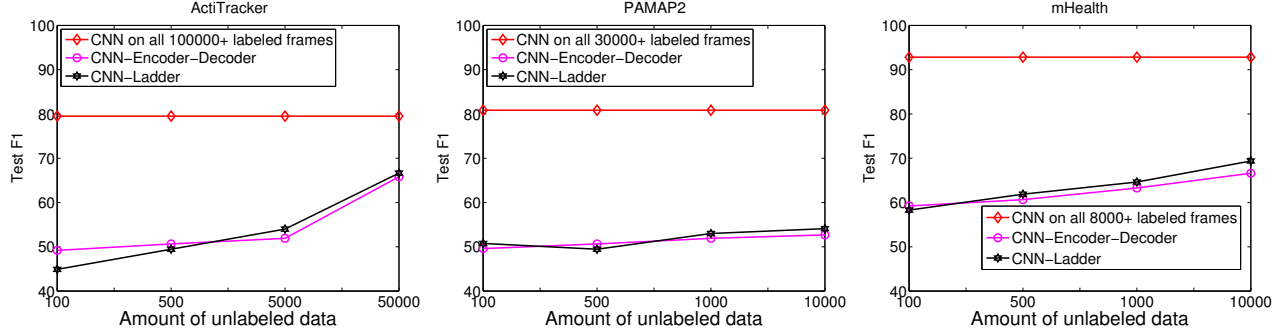


Figure 3: The F_m scores of CNN-Ladder and CNN-Encoder-Decoder, with 50 labeled examples and varying amount of unlabeled examples. The F_m scores of supervised CNN on a large number of labeled examples are also shown as red lines. The result of CNN-Ladder is significantly better than the CNN approach with p -value < 0.05 .




Figure 4: The visualizations of the low-level features of traditional CNN and CNN-Ladder using PCA. The black dots are labeled data of jogging activity from users in the training set. The red dots are unlabeled data of the same activity from a different user not in the training set. Although the red and black dots belong to the same class, they are badly scattered in the traditional CNN. In CNN-Ladder, the red dots are more concentrated around the black dots.

how CNN-Ladder’s low-level features help achieve high F_m scores in HAR.

We visualize the features in the last layer of (i) CNN-Ladder with unlabeled data versus (ii) CNN without unlabeled data. PCA is used to reduce the dimensionality of the data, and only the eigen-vectors with the largest two eigen-values are selected as axes in Figure 4. To understand how CNN-Ladder benefits from varying low-level features, we show two cases where CNN-Ladder achieves high F_m score while CNN does not.

In the prediction of jogging activity for User 11, the features in the last layer of CNN-Ladder with unlabeled data and CNN without unlabeled data are shown in Figure 4(a) and 4(b). In this case, CNN fails to predict the jogging activity of different users as the same activity. This is caused by the varying behaviors of different users, especially when the labeled examples are limited as shown in Figure 4(a). Interestingly, in the two-dimensional visualization of features in CNN-Ladder in Figure 4(b), the test examples concentrate in the region where the labeled data locate. Using the low-level feature representations trained with additional unlabeled data, the jogging activities of different users become similar, even with differences in jogging behaviors between different users. Figure 4(c) and 4(d) show another similar case for User 30.




Figure 5: The impact of making different layers’ latent features of CNN and autoencoder very similar in CNN-Ladder. Utilizing the unlabeled data starting from very low-level features is very important in semi-supervised HAR, but beyond the scope of traditional semi-supervised learning methods.

The visualization results indicate that with unlabeled data, CNN-Ladder can learn discriminative high-level features even when labeled training data is very limited. Consequently, it is easier for CNN-Ladder to achieve higher F_m .

V. CONCLUSION

We study the CNN-Encoder-Decoder and CNN-Ladder architectures for semi-supervised human activity recognition. The experimental results demonstrate that our proposed methods can achieve significant F_m improvements, compared to supervised learning methods and traditional semi-supervised learning methods. We carefully study how CNN-Ladder achieves higher F_m in human activity recognition. The empirical results show that it is very helpful to use unlabeled data to better learn low-level features in CNNs human activity recognition.

VI. ACKNOWLEDGEMENT

This research is supported in part by the National Science Foundation under the award 1346066: “SCH: INT: Collaborative Research: FITTLE+: Theory and Models for Smartphone Ecological Momentary Intervention”.

REFERENCES

- [1] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, “Understanding deep learning requires rethinking generalization,” *ICLR*, 2017.
- [2] N. D. Lane and P. Georgiev, “Can deep learning revolutionize mobile sensing?” in *HotMobile*, 2015, pp. 117–122.
- [3] F. J. Ordóñez and D. Roggen, “Deep convolutional and lstm recurrent neural networks for multimodal wearable activity recognition,” *Sensors*, vol. 16, no. 1, p. 115, 2016.
- [4] M. Stikic, K. Van Laerhoven, and B. Schiele, “Exploring semi-supervised and active learning for activity recognition,” in *ISWC*, 2008, pp. 81–88.
- [5] M. Stikic, D. Larlus, and B. Schiele, “Multi-graph based semi-supervised learning for activity recognition,” in *ISWC*, 2009, pp. 85–92.
- [6] L. Yao, F. Nie, Q. Z. Sheng, T. Gu, X. Li, and S. Wang, “Learning from less for better: semi-supervised activity recognition via shared structure discovery,” in *UbiComp*, 2016, pp. 13–24.
- [7] L. Bao and S. S. Intille, “Activity recognition from user-annotated acceleration data,” in *PerCom*, 2004, pp. 1–17.
- [8] A. Bulling, U. Blanke, and B. Schiele, “A tutorial on human activity recognition using body-worn inertial sensors,” *ACM Computing Surveys (CSUR)*, vol. 46, no. 3, 2014.
- [9] J. B. Yang, M. N. Nguyen, P. P. San, X. L. Li, and S. Krishnaswamy, “Deep convolutional neural networks on multichannel time series for human activity recognition,” in *IJCAI*, 2015, pp. 3995–4001.
- [10] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [11] T. Plötz, N. Y. Hammerla, and P. Olivier, “Feature learning for activity recognition in ubiquitous computing,” in *IJCAI*, vol. 22, no. 1, 2011, pp. 1729–1734.
- [12] M. Zeng, L. T. Nguyen, B. Yu, O. J. Mengshoel, J. Zhu, P. Wu, and J. Zhang, “Convolutional neural networks for human activity recognition using mobile sensors,” in *MobiCASE*, 2014, pp. 197–205.
- [13] N. Y. Hammerla, S. Halloran, and T. Ploetz, “Deep, convolutional, and recurrent models for human activity recognition using wearables,” *IJCAI*, pp. 1533–1540.
- [14] Y. Guan and T. Ploetz, “Ensembles of deep lstm learners for activity recognition using wearables,” *arXiv preprint arXiv:1703.09370*, 2017.
- [15] O. Chapelle, B. Schlkopf, and A. Zien, “Semi-supervised learning,” 2010.
- [16] B. Cvetkovic, B. Kaluza, M. Luštrek, and M. Gams, “Semi-supervised learning for adaptation of human activity recognition classifier to the user,” in *IJCAI*, 2011, pp. 24–29.
- [17] A. Lopes, J. Mendes-Moreira, and J. Gama, “Semi-supervised learning: predicting activities in android environment,” in *Workshop on Ubiquitous Data Mining*, 2012, p. 38.
- [18] A. Rasmus, M. Berglund, M. Honkala, H. Valpola, and T. Raiko, “Semi-supervised learning with ladder networks,” in *NIPS*, 2015, pp. 3546–3554.
- [19] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol, “Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion,” *JMLR*, vol. 11, pp. 3371–3408, 2010.
- [20] V. Badrinarayanan, A. Kendall, and R. Cipolla, “Segnet: A deep convolutional encoder-decoder architecture for image segmentation,” *arXiv preprint arXiv:1511.00561*, 2015.
- [21] J. R. Kwapisz, G. M. Weiss, and S. A. Moore, “Activity recognition using cell phone accelerometers,” *ACM SigKDD Explorations Newsletter*, vol. 12, no. 2, pp. 74–82, 2011.
- [22] A. Reiss and D. Stricker, “Introducing a new benchmarked dataset for activity monitoring,” in *ISWC*, 2012, pp. 108–109.
- [23] O. Banos, R. Garcia, J. A. Holgado-Terriza, M. Damas, H. Pomares, I. Rojas, A. Saez, and C. Villalonga, “mhealthdroid: a novel framework for agile development of mobile health applications,” in *International Workshop on Ambient Assisted Living*, 2014, pp. 91–98.
- [24] D. Erhan, P.-A. Manzagol, Y. Bengio, S. Bengio, and P. Vincent, “The difficulty of training deep architectures and the effect of unsupervised pre-training,” in *AISTATS*, vol. 5, 2009, pp. 153–160.
- [25] D.-H. Lee, “Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks,” in *Workshop on Challenges in Representation Learning, ICML*, vol. 3, 2013, pp. 2–7.