SANI, S., WIRATUNGA, N., MASSIE, S. and COOPER, K. 2018. Matching networks for personalised human activity recognition. In Bichindaritz, I., Guttmann, C., Herrero, P., Koch, F., Koster, A., Lenz, R., López Ibáñez, B., Marling, C., Martin, C., Montagna, S., Montani, S., Reichert, M., Riaño, D., Schumacher, M.I., ten Teije, A. and Wiratunga, N. (eds.) *Proceedings of the 1st Joint workshop on artificial intelligence in health, organized as part of the Federated AI meeting (FAIM 2018)*, co-located with the 17th International conference on autonomous agents and multiagent systems (AAMAS 2018), the 35th International conference on machine learning (ICML 2018), the 27th International joint conference on artificial intelligence (IJCAI 2018), and the 26th International conference on case-based reasoning (ICCBR 2018), 13-19 July 2018, Stockholm, Sweden. CEUR workshop proceedings, 2142. Aachen: CEUR-WS [online], pages 61-64. Available from: http://ceur-ws.org/Vol-2142/short4.pdf

Matching networks for personalised human activity recognition.

SANI, S., WIRATUNGA, N., MASSIE, S. and COOPER, K.

2018

Matching Networks for Personalised Human Activity Recognition

Sadiq Sani[1], Nirmalie Wiratunga[1], Stewart Massie[1], and Kay Cooper[2]

[1] School of Computing Science and Digital Media,
[2] School of Health Sciences,
Robert Gordon University,
Aberdeen AB10 7GJ, Scotland, UK
{s.sani,n.wiratunga,s.massie,k.cooper}@rgu.ac.uk

Abstract. Human Activity Recognition (HAR) has many important applications in health care which include management of chronic conditions and patient rehabilitation. An important consideration when training HAR models is whether to use training data from a general population (subject-independent), or personalised training data from the target user (subject-dependent). Previous evaluations have shown personalised training to be more accurate because of the ability of resulting models to better capture individual users' activity patterns. However, collecting sufficient training data from end users may not be feasible for real-world applications. In this paper, we introduce a novel approach to personalised HAR using a neural network architecture called a matching network. Matching networks perform nearest-neighbour classification by reusing the class label of the most similar instances in a provided support set. Evaluations show our approach to substantially out perform general subject-independent models by more than 5% macro-averaged F1 score.

1 Introduction

Automatic recognition and tracking of human activity using wearable sensors is increasingly being adopted for health care applications e.g. management of chronic low back pain in SelfBACK [1] [1]. An important consideration for HAR applications is classifier training, where training examples can either be acquired from a general population (subject-independent), or from the target user of the system (subject-dependent). Previous works have shown using subject-dependent data to result in superior performance [5, 2, 3, 6]. The relatively poorer performance of subject-independent models can be attributed to variations in activity patterns, gait or posture between different individuals [4]. However, training a classifier exclusively with user provided data is not practical in a real-world configuration as this places significant burden on the user to provide sufficient amounts of training data required to build a personalised model.

[1]The SelfBACK project is funded by European Union's H2020 research and innovation programme under grant agreement No. 689043.

In this paper, we introduce an approach to personalised HAR using matching networks. Matching Networks are a type of neural network architecture introduced for the task of one-shot learning [7] which is a scenario where an algorithm is trained to recognise a new class from just a few examples of that class. Given a (typically small) support set of labelled examples, matching networks are able to classify an unlabelled example by reusing the class labels of the most similar examples in the support set. At the same time, because classification is only conditioned on the support set, matching networks behave like non-parametric models and can reason with any set of examples that are provided at runtime, without the need for retraining the network. This makes our system potentially able to continuously adapt to changes in the user's context.

2 Personalised HAR using Matching Networks

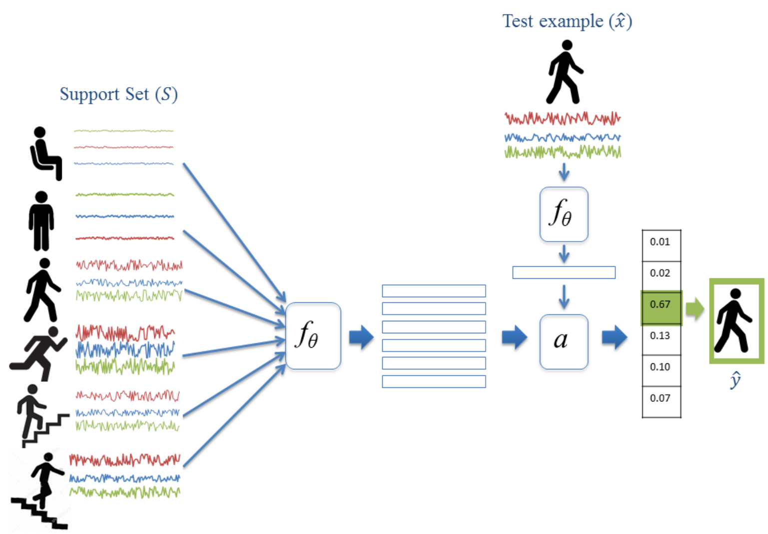

Fig. 1. Illustration of matching network for HAR.

The aim of matching networks is to learn a model that maps an unlabelled example $\hat{x}$ to a class label $\hat{y}$ using a small support set S of labelled examples. This is illustrated in Figure 1. Given a set of instances $X = \{x\}$ and a set of class labels $L = \{y\}$, an embedding function f_θ which in this case a neural network parameterised by θ, the function a is an attention mechanism that takes the embedded representation of a test instance $\hat{x}$ and a support set S and returns a probability distribution $P(y|\hat{x}, S)$ over class labels y of instances in S. To train the matching network for personalised HAR, we also define a set of users U where each user $u_j \in U$ is comprised of a set of labelled examples as follows:

$$u_j = \{(x, y)|x \in X, y \in L\} \quad (1)$$

Next we define a set of training instances T_j for each user u_j as follows:

$$T_j = \{(S_j, B_j)\}^l \quad (2)$$

i.e., T_j is made up of user-specific support and target set pairs S_j and B_j respectively, where $S_j = \{(x, y)|x \in u_i, y \in L\}$ and $B_j = \{(x, y)|x \in u_j, x \notin S_j\}$. Note that the set of labels in S_j is always equivalent to L because we are interested in learning a classifier over the entire set of activity labels. Accordingly, S_j contains m examples for each class $y \in L$ and the cardinality of S_j is $|S_j| = m \times |L|$. Both S_j and B_j are sampled at random from u_j l times to create T_j. Each B_j is used with it's respective S_j by classifying each instance in B_j using S_j and computing loss using categorical cross entropy. The network is trained using stochastic gradient descent and back propagation.

3 Evaluation

Evaluation is conducted on a dataset of 50 users with 9 activity classes (Standing, Sitting, Lying, Walking Slow, Walking Normal, Walking Fast, Up Stairs, Down Stairs) where each user performs each activity for about 3 minutes. We adopt a hold-out validation strategy where 8 out of the 50 users are randomly selected for testing. To simulate user provided samples for creating personalised support sets, we hold out the first 30 seconds of each test user's data for creating the support set. This leaves approximately 150 seconds of data per activity which are used for testing, Performance is reported using macro-averaged F1 score.

In the evaluation, we explore the performance of our matching network against a number of baseline approaches. Accordingly we compare the following algorithms:

- kNN: Nearest-neighbour classifier trained on the entire training set
- SVM: Support Vector Machines trained on the entire training set
- MLP: A Feed-forward neural network trained on the entire training
- MNet: Our personalised matching network approach

All algorithms use 5-second window sizes for signal partitioning and discrete cosine transform for feature representation.

Table 1. Results of different algorithms showing F1 scores.

Algorithm	*kNN*	*SVM*	*MLP*	*MNet*
F1 Score	0.661	0.734	0.723	**0.788**

It can be observed from Table 1 that MNet produces the best result, SVM and MLP have comparative performance while kNN comes in last. MNet out performs both SVM and MLP by more than 5% which shows the effectiveness of our matching network approach at exploiting personal data for activity recognition.

4 Conclusion

In this paper, we presented a novel approach for personalised HAR using matching networks. There are two main advantages to the approach we presented. Firstly, our approach is able to achieve high accuracy using only a small set of user provided examples (30 seconds in this work) which makes it more practical for real-world applications compared to subject-dependent training which requires the end user to provide large amounts (possible hours) of labelled training data. Secondly, our approach does not require retraining the model at runtime when new data becomes available which makes it very adaptable.

References

1. Bach, K., Szczepanski, T., Aamodt, A., Gundersen, O.E., Mork, P.J.: Case representation and similarity assessment in the selfback decision support system. In: Proceedings of 24th International Conference on Case-Based Reasoning. pp. 32–46. Springer International Publishing (2016)
2. Berchtold, M., Budde, M., Gordon, D., Schmidtke, H.R., Beigl, M.: Actiserv: Activity recognition service for mobile phones. In: Proceedings of International Symposium on Wearable Computers. pp. 1–8 (2010)
3. Jatoba, L.C., Grossmann, U., Kunze, C., Ottenbacher, J., Stork, W.: Context-aware mobile health monitoring: Evaluation of different pattern recognition methods for classification of physical activity. In: Proceedings of 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society. pp. 5250–5253 (2008)
4. Longstaff, B., Reddy, S., Estrin, D.: Improving activity classification for health applications on mobile devices using active and semi-supervised learning. In: Proceedings of 4th International Conference on Pervasive Computing Technologies for Healthcare. pp. 1–7 (2010)
5. Sani, S., Wiratunga, N., Massie, S., Cooper, K.: knn sampling for personalised human activity recognition. In: Proceedings of International Conference on Case-Based Reasoning. pp. 330–344. Springer (2017)
6. Tapia, E.M., Intille, S.S., Haskell, W., Larson, K., Wright, J., King, A., Friedman, R.: Real-time recognition of physical activities and their intensities using wireless accelerometers and a heart rate monitor. In: Proceedings of 11th IEEE International Symposium on Wearable Computers. pp. 37–40. IEEE (2007)
7. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D.: Matching networks for one shot learning. In: Proceedigns of Advances in Neural Information Processing Systems. pp. 3630–3638 (2016)