2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 25-29, 2020, Las Vegas, NV, USA (Virtual)

Personalized Online Learning with Pseudo-Ground Truth

Viktor Losing[1], Martina Hasenjäger[1], and Taizo Yoshikawa[2]

Abstract— **Personalized online machine learning allows a very accurate modelling of individual behavior and demands. In particular, a system that dynamically adapts during runtime can initiate a continuous collaboration with its user where both alternately adjust to each other to maximize the system's utility. However, in application scenarios based on supervised learning it is often unclear how to obtain the required ground truth for such dynamic systems. In this paper, we focus on applications where a real-time classification of sequential data is crucial. Concretely, we propose to adapt an online personalized model solely based on pseudo-ground-truth information which is provided by another machine learning model. This model has the advantage to classify sequences in retrospective with a small delay and thus is able to achieve a higher performance than real-time systems. In particular, it is a pre-trained offline model, which means that no ground-truth information is necessary during runtime. We apply the proposal on the task of online action classification, for which the benefits of personalization have been recently emphasized.**

I. Introduction

Personalized machine learning has been recently analyzed in different domains such as motion classification [1]–[3], advanced driver assistant systems [4]–[6] and recommendations [7]–[9]. The general idea is that the focus on one person drastically reduces the variance within the data, enabling a better performance with a smaller amount of data. In other words, products/services trying to serve a broad range of customers have to make sacrifices in order to perform well on average. In contrast, the focus on a single user enables a very specific adaptation to individual demands. Hence, personalization is particularly crucial in the case of highly diverse customers. Additionally, the major problem of inter-person generalization is avoided, reducing the computational complexity as post-processing steps such as normalization or temporal integration can often be omitted [3].

In the literature, personalization has mainly been considered in combination with offline machine learning [1], [2], [4], [8], [9]. However, personalization is a more natural fit to online machine learning, since products are often exclusively used by a single person. Therefore, systems that are able to adapt during their application are implicitly personalized. This combination is more powerful as it allows even to handle changes in user preferences and environments. We envision an adaptive online system that collaborates with its user where both alternately adjust to each other over time and maximize the utility of the system. The advantages of personalized online learning models have recently been emphasized in [3], [5].

One crucial challenge in case of supervised online learning is the acquisition of labels on the fly, drastically limiting the application range. Getting supervised information is in general challenging, and particularly large models, e.g. Deep Neural Networks, require huge amounts of labeled data. Still nowadays, the majority of the data is manually labeled by human annotators and therefore expensive. In the case of offline learning, the labels can be collected in advance over a long time period. However, such an approach is not viable for dynamically adapting systems and a different solution is required.

Nonetheless, there are a few options to obtain the ground truth during runtime. One common approach is to obtain the feedback explicitly from humans. For instance, an individual user marks emails as spam for spam classification, but also in human-robot interactions the labels may be explicitly demanded [10], [11]. Clearly, this is rather tedious, requiring the willingness of humans to cooperate. Even though the label burden can be reduced with active-learning techniques [12], where only the most-promising examples are annotated, human input is still required. In case of prediction tasks, an automatic extraction of ground truth can often be done in retrospective. For instance, personalized online learning was recently applied in the task of driver-maneuver prediction at intersections [5]. As soon as the car has passed the corresponding intersection the executed action is automatically labeled within the recorded data and fed back into the online learning model. However, various problems such as online time series classification remain challenging even in retrospective, requiring a different solution.

Usually, offline models are more accurate and robust than online ones, therefore, online models are often neglected for real-world applications. However, in case of personalization they often provide, amongst other benefits, a significant performance advantage. In this paper, we exploit this advantage in order to obtain an accurate ground truth for the purpose of online personalization. Concretely, we propose a system that leverages the benefits of personalized online learning in the task of online action classification. The main idea is that even though the classification in retrospective may still be challenging, there are various application domains where it significantly facilitates the task. Therefore, we propose to pre-train an offline model that performs the classification in retrospective, yielding a higher performance than a model classifying without delay. Its labels are then used as pseudo-ground truth to adapt a personalized online model. In particular, our approach does not require ground truth during

[1]HONDA Research Institute Europe, Carl-Legien-Str. 30, 63073 Offenbach am Main, Germany
[2]HONDA R&D Co., Ltd. HGRX, 8-1 Honcho, Wako-shi, Saitama, 351-0188 Japan

978-1-7281-6211-9/20/$31.00 ©2020 IEEE

runtime.

The approach is related to semi-supervised learning [13] as one model, trained with supervised examples, provides the labels for the adaptation of another model. However, the novelty of our contribution is to use the classification output of a pre-trained average-user model, which classifies sequences in retrospective, as pseudo-ground truth for an online personalized model to exploit the performance benefits of online personalization. We show that the pseudo-ground truth is accurate enough to enable the personalized online model to outperform an average-user offline model, which was trained with real-ground truth.

II. Framework

The proposal is designed for applications that require a real-time classification of time-series data. Concretely, we focus on the evaluation of off- and online learning models for the task of online action classification [14]. Given an input stream of data the goal is to instantly classify the current motion a human is performing. In the following, we define the overarching problem and describe the characteristics of both learning schemes.

A. Online time-series classification

A stream $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_t\}$ of feature vectors (IMU sensor measurements in our case) arrives one after another, where $\mathbf{x}_i \in \mathbb{R}^n$. As t is the current time, the goal is to assign to $\mathbf{x}_t$ the correct class among the set of predefined C classes. In our case, a class corresponds to an action such as walking, standing, etc. Algorithms are allowed to use not only the current feature vector $\mathbf{x}_t$ but also those of the past. The classification is done for each feature vector presented in the order of the stream. A model assigns the class in the form of:

$$y_t^\star = \underset{y_t \in \{1,\dots,C\}}{\arg\max} \; P(y_t | \mathbf{x}_{t-l-1}, \dots, \mathbf{x}_t),$$

where l is the number of past feature vectors an algorithm is allowed to use. Naturally, online time-series classification is more challenging than its offline counterpart, since methods are not allowed to “peek in the future” and must instantaneously determine the class of x_t. This fact is the foundation of our approach. We delay the classification for a predefined number of m time units which enables the retrospective model to utilize also the feature vectors $\mathbf{x}_{t+1} \dots \mathbf{x}_{t+m}$, leading to a higher performance.

B. Offline Learning

In the offline learning setting, an algorithm generates a model function $h : \mathbb{R}^n \mapsto \{1, \dots, c\}$ based on a training set $D_{\text{train}} = \{(\mathbf{x}_i, y_i) \,|\, i \in \{1, \dots, j\}\}$. In the subsequent test phase, the model is applied on another set $D_{\text{test}} = \{(\mathbf{x}_i, y_i) \,|\, i \in \{j+1, \dots, k\}\}$, whose labels are kept hidden. The model provides a label $\hat{y}_i = h(\mathbf{x}_i)$ for every point $x_i \in D_{\text{test}}$ and the 0-1 loss $\mathcal{L}(\hat{y}_i, y_i) = \mathbb{1}(\hat{y}_i \neq y_i)$ is calculated. The test error

$$E(D_{\text{test}}) = \frac{1}{k-j} \sum_{i=j+1}^{k} \mathcal{L}(h(\mathbf{x}_i), y_i) \qquad (1)$$

is the commonly applied performance metric. In our case, an average-user offline model is tested on the data of one hold-out subject and trained with the data of the remaining ones.

C. Online Learning

The online learning setting is more challenging, since the data is accessed one-by-one in a predefined order and the algorithm provides a model after each datapoint. Therefore, models initially tend to deliver a lower performance compared to their offline counterparts. However, they provide the benefits of a low time and space complexity, are able to process datasets of arbitrary sizes and allow particular tuning to a special problem domain.

Formally, a potentially infinite sequence $S_t = (s_1, s_2, \dots, s_t)$ of tuples $s_i = (\mathbf{x}_i, y_i)$ arrives one after another. In contrast to the offline setting, a model function is generated after each tuple. As t represents the current time stamp, the classification $\hat{y}_t = h_{t-1}(\mathbf{x}_t)$ is done according to the previously learned model h_{t-1}. After the true label y_t is revealed, the applied learning algorithm generates a new model $h_t = \text{train}(h_{t-1}, s_t)$ on the basis of the current tuple s_t and the previous model h_{t-1}. Usually, the interleaved-test-train error is used for performance evaluation and is defined as:

$$\hat{E}(S_t) = \frac{1}{t} \sum_{i=1}^{t} \mathcal{L}(h_{i-1}(\mathbf{x}_i), y_i). \qquad (2)$$

We iteratively train the personalized online models from scratch. Precisely, each time series of one user is first classified by the model and subsequently used for training.

The fundamental difference between off- and online approaches is that the offline methods have generally a much larger set of training data available, whereas the online algorithms have the capability to adapt to the actual test data. The natural consequence is that online methods using few data are only applicable if the variation in the test condition is not too high, which is particularly the case for personalized learning. Online algorithms are even able to adapt to non-stationary environments and efficient methods have been recently published [15], [16]. However, concept drift goes beyond the scope of this contribution and is not explicitly considered here.

III. Proposal: Personalized Online Learning with Pseudo-Ground Truth

The main idea of our approach is based on the fact that even though classifying a sequence in retrospective may still be difficult, it is often substantially easier than classifying it instantly. Therefore, it is plausible to assume that a model classifying in retrospective achieves a higher performance. We leverage this performance gap to adapt the personalized model during runtime.

We propose to construct a retrospective ground-truth model that classifies the time series with delay. Given the

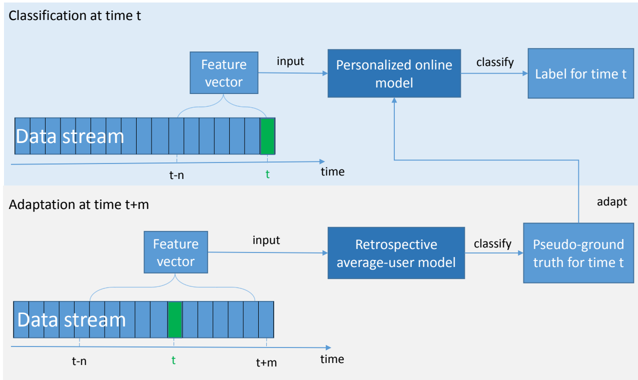

Fig. 1: Overview of the system architecture. The personalized model is performing online action classification without delay. However, the adaptation of the model is delayed due to the buffering of the retrospective model that provides the pseudo-ground.

feature vector $\mathbf{x}_t$ at time t the model classifies it with a delay of m time units at time $t+m$:

$$\hat{y}_t = \underset{y_t \in \{1,\dots,C\}}{\arg\max} \; P(y_t | \mathbf{x}_{t-l-1}, \dots, \mathbf{x}_{t+m}). \tag{3}$$

The delay allows the model to leverage features after time t which facilitates the task and yields a higher performance. In our setting, it allows to classify the current motion with information about how the motion progresses after the query time t.

The retrospective model is an average-user offline model which is trained in advance, preferably utilizing a large dataset including many different users. During runtime, this model is applied to feed its classification outputs as pseudo-ground truth to the personalized online model for adaptation. The personalized model is performing online classification for the time t, meaning it classifies the current feature $\mathbf{x}_t$ immediately without any delay. However, it is adapted at least with a delay of m time units, such that at time $t+m$ it can use the tuple $(\mathbf{x}_t, \hat{y}_t)$ for adaptation. Figure 1 gives an overview of the overall architecture. Naturally, the performance of the personalized model is bounded by the one of the retrospective model.

A. Conditions for Applicability

There are several prerequisites for the applicability of our approach. For a given task at hand the following conditions must apply:

1) Classification in retrospective is not trivial.
2) Online classification without delay is absolutely necessary.
3) Classification in retrospective yields a higher performance than online classification.
4) Personalized online models (using ground truth information) are more accurate than average-user offline models.

The first two conditions exclude alternative solutions that are preferable for certain scenarios. For instance, heuristic approaches are often sufficient to obtain the ground truth in retrospective for prediction tasks, therefore an additional model is not necessary. In the case that a classification delay is acceptable, a retrospective model could directly be used for classification. The last two conditions ensure that there is potential to improve the performance of average-user offline models with the proposal.

B. Drawbacks

Even though our approach is quite intuitive, applicable to various scenarios and can easily be realized, it has various disadvantages. For instance, it trades an adaptation delay of the personalized model to obtain the pseudo-ground truth, which in most cases is negligible. More relevant is the increase of the computational burden as well as the higher system complexity due to the usage of two models, potentially relying on different sets of features. The potentially biggest drawbacks follow from the fact that the system depends on the retrospective model, which is a static, average-user model. Hence, it needs to generalize across different users, potentially requiring post-processing steps such as normalization and temporal integration. Furthermore, the architecture cannot deal with changes in the data distribution over time, also known as concept drift.

IV. Dataset

In this paper, we utilize the dataset that we recently introduced in [3]. It was recorded using the popular Xsens bodysuit with seventeen IMUs, measuring linear and angular motions with a triad of gyroscopes and accelerometers, distributed on different body locations [17]. The data was sampled with a rate of 60 Hz.

Four different subjects performed nine movement sequences consisting of several single actions of sixteen different classes. These sequences were repeated 10-20 times. Figure 2 depicts some action sequences. Altogether, the dataset encodes 2755 actions represented by 329021 single instances covering a time period of around $\sim$ 90 minutes. The distribution of the motion durations is depicted in Figure 3.

V. Experiments

We focus on comparing an average-user offline model against a personalized online one, which is dynamically adapted based on pseudo-ground truth. The pseudo-ground truth is provided by another average-user model which classifies the data with delay. Table I provides an overview of all analyzed models. Please note that the personalized model trained with real ground truth (PERS) is not available in practice. It is merely included to analyse the performance loss caused by the pseudo-ground truth, which can be seen as the untapped potential of personalization by our method.

We evaluate our approach in two alternative scenarios which empirically provide the upper- and lower performance bounds of our method on the evaluated data. In the first, we

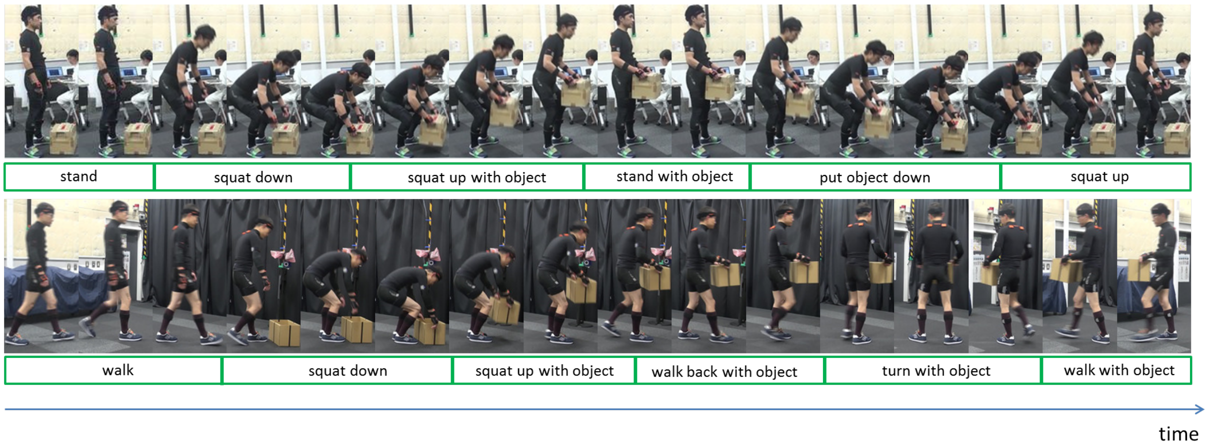

Fig. 2: Exemplary motion sequences consisting of different subsequent actions.

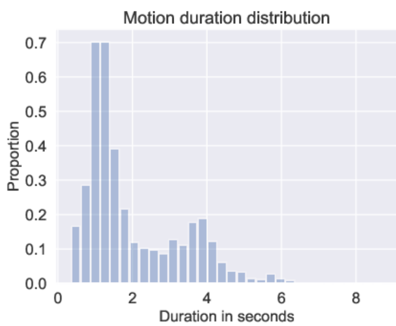

Fig. 3: The distribution of the motion duration. The mean duration of an action is $1.9\,\mathrm{s}$ seconds and the longest motion has a duration of $8.5\,\mathrm{s}$.

TABLE I

THE EVALUATED MODELS IN OUR EXPERIMENTS. IN PRACTICE, ONLINE-GROUND TRUTH IS NOT AVAILABLE SUCH THAT THE MODEL PERS CANNOT BE CONSTRUCTED, REQUIRING ALTERNATIVE APPROACHES. RF=RANDOM FOREST [18]; ORF=ONLINE RANDOM FOREST [19]

Abbr.	Algorithm	Description
AVG	RF	Average user model
$\mathrm{AVG_R}$	RF	Average-user model that classifies with delay and provides the pseudo-ground truth
PERS	ORF	Personalized model adapted with ground truth
$\mathrm{PERS_P}$	ORF	Personalized model adapted with pseudo-ground truth

assume to have perfectly segmented motions, whereas in the second we use no segmentation at all. A perfect segmentation enables the creation of feature vectors that maximize the coverage of the current motion without the risk of encoding portions of temporally adjacent motions. Naturally, this is the ideal case for our approach as the performance advantage of the retrospective model is maximized, leading to accurate pseudo-ground truth. Conversely, having no segmentation limits the leeway of the retrospective model due to the trade-off between the coverage of the current motion and the risk to include portions of temporally adjacent motions.

A. Model evaluation

Average-user models are trained in leave-one-subject-out scheme. Precisely, they are tested with the data of one specific subject, whereas data of the remaining subjects is used for training. This is done repeatedly such that each subject is used for testing once.

Personalized models are evaluated in the online learning setting, as described in Section II-C. The model classifies first the label of one sample and uses it afterwards for model adaption. This is done for all samples in the dataset. However, the order of each instance within one performed action is predefined by the recording time, and therefore, there is a high degree of label auto-correlation, since each action consists of a number of samples with the same class. In this case, the ordinary online scheme is misleading because a naive classifier, simply using the previously seen label for classification, achieves a very low error rate without learning any mapping between the input stream and the corresponding labels. Therefore, we perform an action-wise evaluation. Precisely, the model classifies all samples of one action, before the corresponding labels get revealed and the models are adapted. Thereby, model PERS uses the real-ground truth, whereas model $\mathrm{PERS_P}$ relies on the pseudo-ground truth provided by the restrospective model $\mathrm{AVG_R}$. Personalized models are trained from scratch for each subject in single pass. Consequently, they are without any form of pre-training and only access the data of one subject. Please note, that we calculate both errors (off- and online) using the same data for testing, but the online algorithms continuously adapt their model to the test subject.

We apply on- and offline variants of the popular Random Forest (RF) [18], [19] to enable a possibly fair comparison. The RF is a well known state-of-the-art learning algorithm, delivering highly competitive results [20], [21] and is easy to apply out of the box. Concretely, we use decision forests consisting of 100 trees and rely on the class entropy as impurity function [22].

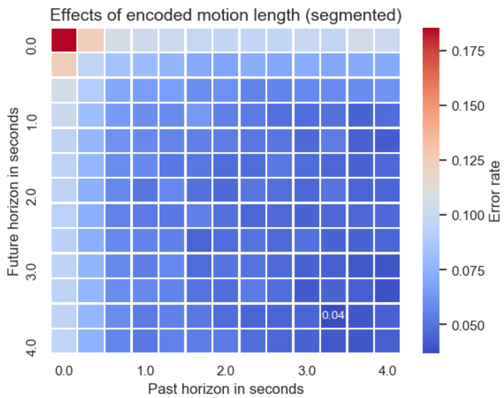

Fig. 4: Effects of varying the size of the encoded time-window in both temporal directions on the model error rate. The configuration yielding the lowest error is highlighted and achieved by encoding $3.33\,\mathrm{s}$ of the past motion data and $3.66\,\mathrm{s}$ into the future. In general, a large horizons in both directions lead to a high performance.

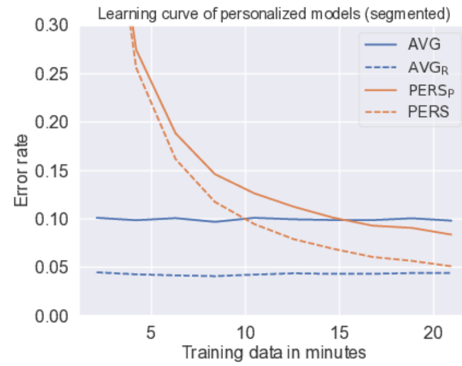

Fig. 5: The classification performance of the different models. The personalized model using pseudo-ground truth ($\mathrm{PERS_P}$) performs better than the average-user model after $\sim 15\,\mathrm{min}$ of motion data. Personalized models are trained from the scratch and continuously adapt to the specific subject, thereby reducing their error rate. Contrary, the average-user models are pre-trained and static.

B. Data Encoding

Similar to our previous work [3], we performed a wrapper-based feature selection [23] and selected the 5 sensors that yielded the highest performance for personalized and average-user models respectively. We refer the interested reader for more details about the feature selection process to [3]. Based on the selected sensors we construct the feature vector in the following way:
For each time t we stack the IMU features of the covered time duration (l instances before t for all models and additionally m instances after t for the retrospective model) and encode them via the Discrete Cosinus Transformation (DCT) [24]. Finally, the 5 largest coefficients are used as feature vectors for each IMU input signal. The resulting feature vectors of the personalized models consist of 135 dimensions, whereas 165 dimensions were used by the average-user ones.

C. Segmented Motions

Initially, we assume perfectly segmented motions and correspondingly pad the feature vectors with zeros at the beginning and end of a motion.

1) Varying the length of encoded window: In preliminary experiments, we determine the hyperparameters for the size of the encoded window (variables l, m in Equation 3). Figure 4 shows the performance of an offline classifier trained in leave-one-subject-out with varying sizes of encoded time windows. As expected, the performance increases with a larger time window for perfectly segmented motions. It can be seen, that encoding a time horizon of more than $3.66\,\mathrm{s}$ in either direction (past or “future”) does not provide further benefits for our dataset. Hence, we set the parameters accordingly, i.e. $l = 3.33\,\mathrm{s}$; $m = 3.66\,\mathrm{s}$. Therefore, the online action classification models (AVG, PERS, $\mathrm{PERS_P}$) are fed with feature vectors that encode 3.3 s of the past motion data, whereas the retrospective model processes feature vectors that encode $7\,\mathrm{s}$ of motion data ($3.33\,\mathrm{s}$ of the past data and $3.66\,\mathrm{s}$ of the “future” input). Essentially, the retrospective model is mostly able to access the whole motion for classification considering the tpyical motion durations in the dataset (see Figure 3).

2) Main Experiment: We performed our main experiment of training personalized models with pseudo-ground truth 100 times re-shuffling the order of the motions for each repetition. As the total number of data instances varies across the subjects we truncate the amount of considered data instances for each subject to the same length ($\sim$ 22 minutes) after reshuffling.
Figure 5 depicts the learning curve of the personalized models. The average-user model classifying the samples in retrospective ($\mathrm{AVG_R}$) clearly performs better than the one classifying it immediately (AVG), nearly halving the error rate. As already mentioned, these models are static and therefore their performance is more or less constant over time. In contrast, the personalized models are continuously adapting thus reducing the error rate with increasing amount of processed data. It seems likely that the error rate will be further reduced with more data. The personalized model relying on the pseudo-ground truth ($\mathrm{PERS_P}$) is able to outperform the average-user model (AVG) after $\sim 70\%$ of the data. Please note that the models $\mathrm{PERS_P}$ and AVG are the only ones that can be applied in practice to perform online action classification. The retrospective model $\mathrm{AVG_R}$ is classifying with delay and model PERS uses ground truth information that is usually not available. Their curves are only shown for analytical purposes. In particular, it can be seen that the personalized models using real-ground truth information (PERS) is able to deliver a similar performance as $\mathrm{AVG_R}$, but it classifies the actions without delay. Hence, the potential of personalization in this task seems to be sufficient to expect that $\mathrm{PERS_P}$ converges towards the performance of $\mathrm{AVG_R}$ with more data.

TABLE II

THE RESULTS OF WELCH'S T-TEST COMPARING THE ERROR RATES OF THE AVERAGE-USER MODEL AND THE PERSONALIZED ONE USING THE PSEUDO-GROUND TRUTH IN THE TASK OF ONLINE ACTION CLASSIFICATION. SIGNIFICANT DIFFERENCES ARE MARKED IN BOLD. THE EXPERIMENTS WERE REPEATED 100 TIMES WITH RESHUFFLED MOTION ORDER. THE DATA WAS SUBDIVIDED IN 10 CHRONOLOGICALLY ORDERED PARTS AND THE TEST WAS PERFORMED FOR EACH OF THOSE.

Part	PERS_P	AVG	P-Value	Better/Worse
1	0.447($\pm$0.05)	0.899($\pm$0.04)	~ 0	**worse**
2	0.725($\pm$0.04)	0.901($\pm$0.04)	3.3E-224	**worse**
3	0.812($\pm$0.04)	0.899($\pm$0.04)	6.2E-107	**worse**
4	0.854($\pm$0.04)	0.903($\pm$0.04)	2.0E-48	**worse**
5	0.874($\pm$0.04)	0.899($\pm$0.04)	2.2E-14	**worse**
6	0.887($\pm$0.04)	0.900($\pm$0.04)	7.5E-05	**worse**
7	0.899($\pm$0.03)	0.901($\pm$0.04)	4.3E-01	worse
8	0.907($\pm$0.03)	0.901($\pm$0.04)	6.7E-02	better
9	0.909($\pm$0.03)	0.899($\pm$0.03)	1.0E-03	**better**
10	0.916($\pm$0.03)	0.902($\pm$0.03)	2.5E-06	**better**

We subdivided the data in 10 portions and tested the corresponding error rates for significant differences between the models AVG and PERS_P. Concretely, we performed Welch's t-test for unequal variances [25] with $\alpha = 0.005$. The alpha value was chosen based on the Bonferroni correction [26] to keep the probability of false positives below 5%. Table II shows the results of the test for each portion of the data. The error rates with corresponding standard deviations are listed as well. The personalized model using the pseudo-ground truth is significantly better than the average-user model within the last 20% of the processed data. Hence, we can conclude that the proposed approach yields a personalized model on the basis of pseudo-ground truth that significantly outperforms an average-user offline model, pre-trained on real-ground truth.

D. Unsegmented motions

In practice, it is unrealistic to expect perfectly segmented motions. Hence, we perform the same experiments but assume to have no segmentation of the motions at all. Concretely, the feature vectors are constructed independently from the underlying duration of the motion classes, therefore, they cover sometimes multiple and different motion classes.

1) Varying the length of encoded window: The impact of varying the size of the encoded time window is illustrated by the heatmap of Figure 6. As expected, a relatively small window size of 1.3 s (in both directions) yields the best compromise which is less than half the size yielding the best performance in the segmented case. Since the classification task is more challenging with unsegmented data, the error rates are also distinctly higher.

2) Main Experiment: The model performances achieved with the unsegmented motion data are depicted in Figure 7. The personalized model using the pseudo-ground truth (PERS_P) has a flatter learning curve and achieves towards the end a similar performance as the average model. The plausible explanation is that even though the retrospective model AVG_R yields also in this case a similar performance

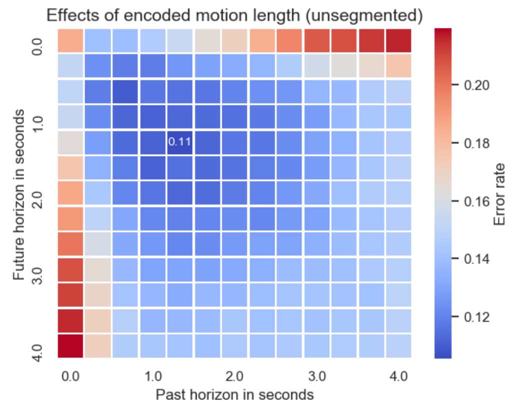

Fig. 6: Error rates of varied length of encoded time-window in the feature vector. The sweet spot of the hyperparameter setting is clearly pronounced for unsegmented motions, as too large horizons are detrimental.

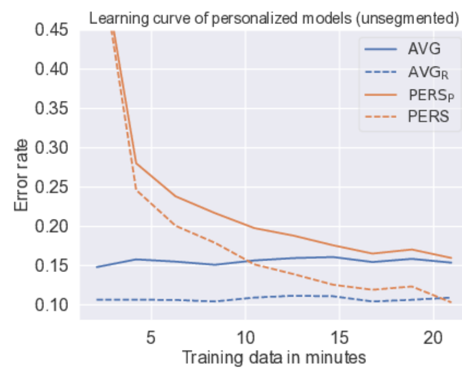

Fig. 7: The learning curve of the personalized model using the pseudo-ground truth is flatter for unsegmented data as the pseudo-ground truth is less accurate, requiring more data instances to converge. Therefore, more data is necessary to evaluate whether a benefit can be achieved with the proposal.

advantage of $\sim 5\%$ compared to AVG, it has in absolute terms a higher error rate. Consequently, the pseudo ground-truth is less accurate and the personalized model depending on it PERS_P requires more data instances to converge. However, the course of the learning curve suggests that PERS_P likely outperforms the AVG model with additional data.

The results in both scenarios underline the potential of our approach, where the effectiveness was more pronounced in the case of segmented data. As a segmentation of motion data is usually beneficial in terms of the general classification performance [27] it has even a higher value in combination with our approach.

VI. CONCLUSION AND DISCUSSION

In this paper, we tackled the challenge of obtaining ground truth information for personalized online learning. Concretely, we proposed a novel approach to adapt a personalized online model based on pseudo-ground truth information that is provided by another offline model, classifying the sequences in retrospective. This strategy was applied for the task of online action classification using segmented and unsegmented motion data. Thereby, the approach was able to outperform average-user offline models, pre-trained using real-ground truth, with significance in case of segmented motions. Without the segmentation the learning curve of the model was flatter and it achieved the same performance as the average-user model. However, the shape of the curve suggests that it probably delivers a better performance with additional motion data.

Our intuitive approach can be easily applied in practice and does not require any ground-truth information during runtime. The main conditions for the applicability of the proposal is that personalization and classification in retrospective provide a significant performance advantage over average-user real-time classification. Please note, that the approach is independent from the underlying learning algorithms. These are interchangeable as long as the required conditions are still met. For instance, end-to-end deep learning architectures such as Residual Networks (ResNet) or Long Short-term memorys (LSTM) [28], [29] could also be used as offline models. However, we utilized the established Random Forests as we aimed for similar online- and offline learning architectures to isolate the benefits of personalization. Furthermore, the dataset is rather small (altogether $\sim 90\,\mathrm{min}$ motion data) and prohibits the training of large networks due to overfitting effects.

We are interested in extending the work in two directions. First, we want to investigate whether our envisioned goal of an alternately adaptation between system and user can be observed in practice. In that regard, active learning approaches could be used for ground-truth acquisition. Furthermore, we want to explore different ways to combining average-user models with personalized ones to get the best of both worlds, in particular we aim for a higher robustness. For instance, the "cold-start" of personalized models as they have initially insufficient data could be avoided. Average-user models could be used as fall-back whenever the performance of the personalized model drops, potentially due to the intrinsically higher volatility. Another promising direction to increase the robustness is data augmentation [30].

References

[1] C. Medrano, I. Plaza, R. Igual, Á. Sánchez, and M. Castro, "The effect of personalization on smartphone-based fall detectors," *Sensors*, vol. 16, no. 1, p. 117, 2016.

[2] G. M. Weiss and J. W. Lockhart, "The impact of personalization on smartphone-based activity recognition," in *AAAI Workshop on Activity Context Representation: Techniques and Languages*, 2012, pp. 98–104.

[3] V. Losing, B. Hammer, and H. Wersing, "Personalized online learning of whole body motions using multiple inertial measurement units," in *IEEE International Conference on Robotics and Automation*, 2019.

[4] V. Butakov and P. Ioannou, "Personalized driver/vehicle lane change models for ADAS," *IEEE Transactions on Vehicular Technology*, vol. 64, no. 10, pp. 4422–4431, October 2015.

[5] V. Losing, B. Hammer, and H. Wersing, "Personalized maneuver prediction at intersections," in *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, Oct 2017, pp. 1–6.

[6] M. Hasenjäger and H. Wersing, "Personalization in advanced driver assistance systems and autonomous vehicles: A review," in *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2017, pp. 2205–2211.

[7] A. S. Das, M. Datar, A. Garg, and S. Rajaram, "Google news personalization: scalable online collaborative filtering," in *Proceedings of the 16th international conference on World Wide Web*. ACM, 2007, pp. 271–280.

[8] Y. Zhang, M. Chen, D. Huang, D. Wu, and Y. Li, "idoctor: Personalized and professionalized medical recommendations based on hybrid matrix factorization," *Future Generation Computer Systems*, vol. 66, pp. 30–35, 2017.

[9] M. F. Alhamid, M. Rawashdeh, H. Dong, M. A. Hossain, and A. El Saddik, "Exploring latent preferences for context-aware personalized recommendation systems," *IEEE Transactions on Human-Machine Systems*, vol. 46, no. 4, pp. 615–623, 2016.

[10] J. Luo, A. Pronobis, B. Caputo, and P. Jensfelt, "Incremental learning for place recognition in dynamic environments," in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Oct 2007, pp. 721–728.

[11] S. Amershi and M. Cakmak, "Power to the people: The role of humans in interactive machine learning," *AI Magazine*, December 2014.

[12] B. Settles, "Active learning literature survey," University of Wisconsin-Madison Department of Computer Sciences, Tech. Rep., 2009.

[13] J. Prakash and D. Nithya, "A survey on semi-supervised learning techniques," *International Journal of Computer Trends and Technology*, vol. 8, 02 2014.

[14] Y. Li, C. Lan, J. Xing, W. Zeng, C. Yuan, and J. Liu, "Online human action detection using joint classification-regression recurrent neural networks," in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds. Cham: Springer International Publishing, 2016, pp. 203–220.

[15] V. Losing, B. Hammer, and H. Wersing, "KNN Classifier with Self Adjusting Memory for Heterogeneous Concept Drift," in *16th International Conference on Data Mining (ICDM)*. IEEE, 2016.

[16] H. M. Gomes, A. Bifet, J. Read, J. P. Barddal, F. Enembreck, B. Pfharinger, G. Holmes, and T. Abdessalem, "Adaptive random forests for evolving data stream classification," *Machine Learning*, vol. 106, no. 9, pp. 1469–1495, Oct 2017.

[17] D. Roetenberg, H. Luinge, and P. Slycke, "Xsens mvn: Full 6dof human motion tracking using miniature inertial sensors," vol. 3, 01 2009.

[18] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.

[19] A. Saffari, C. Leistner, J. Santner, M. Godec, and H. Bischof, "On-line random forests," in *IEEE 12th International Conference on Computer Vision Workshops (ICCV Workshops)*. IEEE, 2009, pp. 1393–1400.

[20] M. Fernández-Delgado, E. Cernadas, S. Barro, and D. Amorim, "Do we need hundreds of classifiers to solve real world classification problems," *J. Mach. Learn. Res*, vol. 15, no. 1, pp. 3133–3181, 2014.

[21] V. Losing, B. Hammer, and H. Wersing, "Choosing the Best Algorithm for an Incremental Learning Task," in *European Symposium on Artificial Neural Networks (ESANN)*, 2016.

[22] U. M. Fayyad and K. B. Irani, "The attribute selection problem in decision tree generation," in *AAAI*, 1992, pp. 104–110.

[23] J. Tang, S. Alelyani, and H. Liu, "Feature selection for classification: A review," *Data classification: Algorithms and applications*, p. 37, 2014.

[24] N. Ahmed, T. Natarajan, and K. R. Rao, "Discrete cosine transform," *IEEE Transactions on Computers*, vol. C-23, pp. 90–93, Jan 1974.

[25] G. D. Ruxton, "The unequal variance t-test is an underused alternative to Student's t-test and the MannWhitney U test," *Behavioral Ecology*, vol. 17, no. 4, pp. 688–690, 05 2006.

[26] C. E. Bonferroni, *Teoria statistica delle classi e calcolo delle probabilita*. Libreria internazionale Seeber, 1936.

[27] E. Keogh, S. Chu, D. Hart, and M. Pazzani, "Segmenting time series: A survey and novel approach," in *Data mining in time series databases*. World Scientific, 2004, pp. 1–21.

[28] Z. Wang, W. Yan, and T. Oates, "Time series classification from scratch with deep neural networks: A strong baseline," in *International Joint Conference on Neural Networks (IJCNN)*, May 2017, pp. 1578–1585.

[29] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[30] X. Ren, J. Luo, E. Solowjow, J. A. Ojea, A. Gupta, A. Tamar, and P. Abbeel, "Domain randomization for active pose estimation," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 7228–7234.