

# Domain Adaptation Under Behavioral and Temporal Shifts for Natural Time Series Mobile Activity Recognition

Garrett Wilson  
garrett.wilson@wsu.edu  
Washington State University  
Pullman, WA, USA

Janardhan Rao Doppa  
jana.doppa@wsu.edu  
Washington State University  
Pullman, WA, USA

Diane J. Cook  
djcook@wsu.edu  
Washington State University  
Pullman, WA, USA

## ABSTRACT

Increasingly, human behavior is captured on mobile devices, leading to an increased interest in automated human activity recognition. However, existing datasets typically consist of scripted movements. Our long-term goal is to perform mobile activity recognition in natural settings. We collect a dataset to support this goal with activity categories that are relevant for downstream tasks such as health monitoring and intervention. Because of the large variations present in human behavior, we collect data from many participants across two different age groups. Because human behavior can change over time, we also collect data from participants over a month’s time to capture the temporal drift. We hypothesize that mobile activity recognition can benefit from unsupervised domain adaptation algorithms. To address this need and test this hypothesis, we analyze the performance of domain adaptation across people and across time. We then enhance unsupervised domain adaptation with contrastive learning and with weak supervision when label proportions are available.

## CCS CONCEPTS

• **Computing methodologies** → **Machine learning; Transfer learning; Unsupervised learning; Adversarial learning; Neural networks**; • **Mathematics of computing** → **Time series analysis**.

## KEYWORDS

datasets, time series, domain adaptation, human activity recognition

### ACM Reference Format:

Garrett Wilson, Janardhan Rao Doppa, and Diane J. Cook. . Domain Adaptation Under Behavioral and Temporal Shifts for Natural Time Series Mobile Activity Recognition. In *Proceedings of MILETS: SIGKDD International Workshop on Mining and Learning from Time Series – Deep Forecasting: Models, Interpretability, and Applications (MILETS: SIGKDD International Workshop on Mining and Learning from Time Series)*. ACM, New York, NY, USA, 9 pages.

## 1 INTRODUCTION

Humans hunger to recognize and understand behavior patterns. An individual’s activities affect their family, the workplace, society, and the environment. Human behavior is also a primary influence of a person’s own health and wellness [15, 17]. Because behavior is reflected in mobile sensor data and such data are being amassed from commercial devices, researchers are increasingly tackling the problem of automating human activity recognition (HAR) from time series sensor data. However, existing algorithmic methods and datasets are typically focused on modeling repetitive, scripted

movements that are performed in controlled settings [2, 5]. Now that these HAR systems have demonstrated success in constrained conditions, the next step is to design algorithmic methods and datasets for in-the-wild activity recognition.

Modeling and recognizing human activities in natural settings introduces numerous challenges not found in controlled settings. First, ground-truth labels are not as readily available. While the activity category is specified for scripted activities, labeling activities in the wild is burdensome, sometimes causes interruptions in the activities being modeled, and is frequently subject to error. Corresponding recognition techniques therefore need to handle cases where labels are not available for all individuals.

Second, modeling behavior that is performed in natural settings must reflect the tremendous variability that exists in human behavior. Many machine learning algorithms are designed based on the closed-world assumption, where the test data distribution is assumed to be similar to the training data distribution. In human behavior data, this assumption does not hold. In fact, variability is considered by some to be the most conspicuous characteristic of human behavior [7]. The change in distribution may lead to a drop in the performance of activity recognition algorithms. Hence, activity recognition models need to capture or adapt to these individual differences. Such differences occur not only between individuals but within a single person’s behavior at distinct time points. Behavioral facets fluctuate with changes in physiology, the external environment, and life course [16, 32]. The nonstationarity of sensor-based behavior time series data further complicates the ability to model activities, particularly in the absence of regularly-timed ground-truth labels.

We hypothesize that the performance of mobile activity recognition (MAR) in natural settings can be boosted by unsupervised domain adaptation. Adversarial learning, which utilizes a two-player game to produce a domain-invariant feature representation, has become a commonplace tool in domain adaptation [8]. We further postulate that by augmenting adversarial learning with contrastive learning, which is a technique pulling similar examples together in the representation space and pushing dissimilar examples apart [11], we can improve domain adaptation performance in the face of the increased adaptation difficulty due to the tremendous variability between people and nonstationarity in behavior over time. Finally, we hypothesize that in such naturalistic and highly variable settings, further gains may be achieved through leveraging additional weak supervision information in the form of label proportions (if this information is available), which may be easier to collect from participants than additional labeled data. For example, such proportions can be gathered through self-reported answers to questions such as “How many hours did you sleep last night?” [30].

While we do not collect this data in our datasets, we measure the performance gains achievable by utilizing simulated self-reported data, computed from the labeled data we collect.

In this paper, we present an unscripted mobile human activity recognition dataset and measure the performance gains achievable by utilizing time series domain adaptation algorithms. To achieve our goal of obtaining a model that performs well in natural settings, we collect a variety of sensor data from smartwatches worn by participants, all of whom are occasionally prompted to label what activity they are performing throughout everyday life. To cover the wide variability between people, we collect data from two different age groups, each of which consists of data from a large number of participants: one dataset component with 15 younger adult participants and another dataset component with 30 older adult participants. Additionally, to enable addressing nonstationarity in behavior over time, we collect each of these datasets over a period of time: the younger adult data were collected over two weeks separated by a month, and the older adult data were collected over four consecutive weeks.

On each of these datasets, we evaluate two time series domain adaptation methods. First, we analyze domain adaptation results in the context of cross-person adaptation, where we train a model on data from one set of participants and adapt the model to work well on a different participant’s data. Cross-person adaptation is the most commonly-studied case of adaptation for human activity recognition [30, 31]. Second, because in the creation of our datasets we collected data over a period of time, we are able to measure performance degradation of the models as the training and testing data are separated by greater lengths in time. This concept shift over time that occurs in these datasets collected in real-world settings, indicative of the nonstationarity of human behavioral data, motivates the need for cross-time domain adaptation. We repurpose time series domain adaptation methods for this cross-time domain adaptation problem and analyze the results. While domain adaptation methods yield an improvement over no adaptation, they still struggle to yield satisfactory performance on these more challenging datasets we present, motivating the need for future time series domain adaptation work.

In this paper, we describe and analyze the performance of two time series domain adaptation techniques for cross-person adaptation. We evaluate the methods on a new mobile activity recognition dataset<sup>1</sup>. Data collection spans multiple weeks, thus we determine the extent of concept drift in the data and analyze the performance of cross-time adaptation as well. For both cross-person and cross-time adaptation problems, we simulate weak supervision in the form of label proportions and measure the time series domain adaptation performance gains achievable by utilizing this additional form of information.

## 2 RELATED WORK

### 2.1 Activity Recognition

Modeling and recognizing activities from time series sensor data has become a popular research topic. Diverse classical machine

learning methods have been explored, including decision tree, nearest neighbor, support vector machine, random forest, clustering, and Bayesian classifiers [6]. However, such methods often rely on manual definition and extraction of discriminating features from the raw sensor readings. More recently, deep network architectures have been designed to model activities from the raw sensor data. Data are extracted from sensors that include triaxial accelerometers and gyroscopes, magnetometers, and heart rate. Deep learners have achieved strong recognition performance for movement-driven activity categories. Examples of these categories include body postures and repetitive movements such as standing, walking, sitting, running, lying, climbing, eating, writing, dribbling, and falling, as well as grinding, drilling, and sanding [10].

While most recent HAR research restricts activity categories to simple movements, some researchers have investigated approaches to modeling composite activities, representing combinations of simple movements. For example, Vepakomma et al. [28] model combinations of movements and location types (e.g., walk indoor versus run outdoor) as well as semantically complex activities (e.g., clean, cook). They ease the learning task by introducing handcrafted features into the model. Alternatively, Peng et al. [20] use a method which is effective when complex activities can be defined as a sequence of simple movements. This approach defines a hierarchical structure to first model simple activities, then learn combinations of the basic categories.

To evaluate these existing mobile human activity recognition approaches, researchers have created and shared HAR datasets. These popular resources include UCI HAR [1], UCI HHAR [27], WISDM AR [12], and WISDM AT [14]. In keeping with common methodologies, these datasets typically contain sensor data for activities that are scripted, performed uniformly by multiple participants, and, with the exception of WISDM AT, collected in a laboratory setting. We are interested in further extending the diversity of modeled activities. These additional categories reflect activities found in common questionnaires such as the American Time Use Survey [18] (e.g., eat, sleep, work, relax, household chores, exercise). We also include categories found in the list of basic and instrumental activities of daily living [9] because of their relevance for clinical applications (e.g., hygiene, travel, cook, errands, chores). These activity categories may be more relevant for downstream tasks such as behavior pattern and trend analysis, health monitoring and intervention, and activity-aware service provisioning.

### 2.2 Domain Adaptation

Numerous domain adaptation techniques have been developed [29]. However, only limited progress has been made toward designing robust domain adaptation approaches for time series data. Among these existing approaches are R-DANN and VRADA [22], variational domain-adversarial methods that utilize a variational recurrent neural network (RNN) as the feature extractor. While the RNN did yield promising adaptation results, other prior work [30] indicated that a 1D convolutional neural network outperformed RNNs on a variety of time-series adaptation problems. Another recent approach, AdvSKM [13], minimizes a maximum mean discrepancy metric embedded in a spectral kernel network. ContrasGAN [23] also attacks the HAR domain adaptation problem. In this case, a

<sup>1</sup><https://github.com/WSU-CASAS/smartwatch-data>

bi-directional generator adversarial network performs heterogeneous feature transfer and contrastive learning. Among these, we select CoDATS [30] as a baseline method for comparison because of its ability to perform robust unsupervised domain adaptation and incorporate weak supervision.

### 3 CASAS MAR DATASET

To evaluate our algorithms and support further development of mobile activity recognition, we create a new mobile activity recognition dataset, called CASAS MAR. This dataset consists of smart-watch data collected while participants perform their non-scripted daily routine in real-world settings. In contrast to the publicly-available human activity recognition datasets that are used in prior time series domain adaptation work [1, 12, 14, 27], this new dataset reflects the variability of movements and activity strategies that are found in natural settings.

We designed a watchOS app to collect sensor data from the Apple watch accelerometer, gyroscope, and location services. Participants are queried at periodic intervals to identify their current activity. The dataset contains two components. CASAS MAR/YA represents data collected from younger adults and CASAS MAR/OA represents data collected from older adults. These two components allow us to examine differences in labeling strategies and activity recognition performance between the two groups.

#### 3.1 CASAS MAR/YA

In the first dataset component, data were collected from 15 younger adult (college student) participants (mean age 21.0, standard deviation=2.2) over two separate weeks. The second week of data collection occurred one month after the first week started. The Apple Watch (Series 2) app continuously collected accelerometer and gyroscope data at 50 Hz when not on the charger and collected location data every 5 minutes. Every 30 minutes, the app would prompt the participants to label their current activity from among six categories: *cook*, *eat*, *hygiene*, *work*, *exercise*, and *travel*. Participants had a choice to ignore the prompt when performing an activity such as driving that inhibited responding to the prompt.

In contrast to the UCI HHAR dataset, where Stisen et al. [27] only sample either accelerometer or gyroscope sensors, our dataset includes accelerometer, gyroscope, and location readings. This expanded sensor set enables us to analyze values for 3-axis user acceleration, 3-axis rotation rate, yaw, pitch, roll, latitude, longitude, altitude, speed, and the corresponding timestamp. We stored the sensor data on the watches using Google’s Protocol Buffers (Protobuf) [21], which enables defining both an efficient storage format (each watch has space for over 2 weeks of data in this format) and provides extensibility and backwards compatibility.

The sampling frequency of 50 Hz was chosen because past work used this frequency for activity recognition [1, 4, 24–26]. After data collection, we downsampled to 10 Hz for consistency between dataset components while maintaining classification performance [26]. In the case of location data, we sampled once every 5 minutes. This represented the maximum sampling frequency that ensured the app could continuously collect data for up to 18 hours on a single charge. This constraint allowed participants to wear the

device from when they woke up to when they went to bed each day.

The domain adaptation networks are not designed to fuse time series data arriving at non-uniform rates. To provide uniform data, we maintain the last-recorded location values between readings. When a participant provides an activity label at time  $t$ , the label is applied to all readings that were collected from time  $t$  minus 5 minutes to time  $t$ . We then segment the time series into non-overlapping windows of 128 time steps. This size is selected to be consistent with other publicly-available datasets. We do not encode raw latitude/longitude/altitude location values because 1) they do not generalize effectively to new subjects and 2) they introduce a privacy risk. Instead, we perform a reverse geocode lookup to identify the location type that is associated with the location coordinate using the OpenStreetMap Nominatim database [3]. The values for location type are other, house, road, service, work, and attraction. We one-hot encode the resulting location categories for our feature vector. Since activities of daily living may be affected by time and date, we additionally include the day of week, hour of the day, and minutes past midnight as features.

#### 3.2 CASAS MAR/OA

In the second dataset component, data were collected from 30 older adult participants (mean age 70.7, standard deviation=7.8) for one month. To accommodate the preferences of this group, each participant was prompted to label their activity at most 10 times per day. Prompts were given during waking hours (e.g., between the hours of 9am and 9pm). The reduced query frequency lightened the burden of the study on participants and ensured participant responsiveness for the entire month. The prompt hours are selected to increase instance diversity over the waking hours, thus a prompt is timed for an hour where the number of labeled instances is  $\leq$  the median number of labels across all waking hours.

In contrast with the first dataset component, we collected data using Apple Watch (Series 5), which provides sufficient storage space for the increased study duration. We additionally collected data directly at 10 Hz rather than downsampling after data collection. Finally, we also expanded the label set for this study to include a total of nine activity categories: *errands*, *exercise*, *hobby*, *hygiene*, *mealtime*, *relax*, *sleep*, *travel*, and *work*.

## 4 METHODOLOGY

Our long-term goal is to perform robust activity recognition from mobile sensors, even amidst changes in the data distribution. We study two techniques to this end.

First, we study the problem of unsupervised domain adaptation. In domain adaptation, we define two domains: one or more source domains and a target domain. For unsupervised domain adaptation, the source domains consist of labeled data, and the target domain only requires unlabeled data. For example, we may collect labeled data from a number of “source” participants 1-5 with the goal of producing a model that performs well for “target” person 6. We can reduce the burden of creating this model for person 6 by not requiring this target participant to label data, but instead only wear the watch for a period of time collecting unlabeled data. We can measure the performance of this technique by selecting some

portion of a given MAR dataset to be “sources” and another portion to be “target”, e.g., for cross-person adaptation, selecting based on participant identifier, or, for cross-time adaptation, selecting based on timestamp.

Second, we enhance domain adaptation with weak supervision [30], a technique that leverages additional information about the target domain, if available. Leveraging such weak supervision is an alternative method of obtaining improved domain adaptation performance in a manner that may reduce the data collection burden on participants. Rather than labeling data instances, the target participant instead reports what proportion of the time they were performing different activities, e.g., how many hours a day they exercised or slept. While we did not collect this data in our study, we can simulate this technique by measuring the label proportions on the target domain (while taking care not to make use of the labels directly, which would violate the assumptions of unsupervised domain adaptation).

#### 4.1 Problem Setup

Following the formulation in CoDATS [30], we define several distributions in terms of the input data  $X$  and label space  $L$  consisting of labels  $Y = \{1, 2, \dots, L\}$ :  $n$  source domain distributions  $\mathcal{D}_{S_i}$  and a target domain distribution  $\mathcal{D}_T$ . During training, we draw samples i.i.d. from each labeled source domain distribution and also i.i.d. from the unlabeled marginal target distribution of  $\mathcal{D}_T$  over  $X$ , yielding  $s_i$  source samples and  $t$  target samples  $T$ :

$$S_i = \{(\mathbf{x}_j, y_j)\}_{j=1}^{s_i} \sim \mathcal{D}_{S_i}, \quad \forall i \in \{1, 2, \dots, n\} \quad (1)$$

$$T = \{(\mathbf{x}_j)\}_{j=1}^t \sim \mathcal{D}_T^X \quad (2)$$

During testing, we evaluate the model on samples drawn i.i.d. from the target distribution  $\mathcal{D}_T$ , yielding  $t$  target samples not seen during training:

$$T_{test} = \{(\mathbf{x}_j, y_j)\}_{j=1}^t \sim \mathcal{D}_T \quad (3)$$

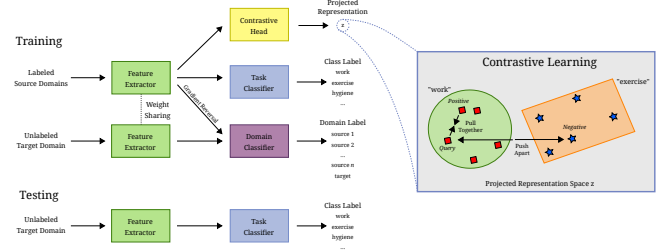
For the problem of domain adaptation with weak supervision, we are additionally provided with the target-domain label proportions, a discrete distribution:

$$P_{true}(y) = P(Y = y) = p_y, \quad \forall y \in \{1, 2, \dots, L\} \quad (4)$$

Because we focus here on time series mobile activity recognition data, each  $X$  represents a multivariate time series. A univariate time series can be represented by a list of  $H$  ordered real values  $X = [x_1, x_2, \dots, x_H]$ . Then, a multivariate time series can be represented by a list of  $K$  univariate time series  $X = [X^1, X^2, \dots, X^K]$ . Our datasets consist of sensor data, e.g., given accelerometer  $x$ ,  $y$ , and  $z$  values, we obtain a series of these readings for each univariate axis.

#### 4.2 Time Series Domain Adaptation Methods

We select two time series domain adaptation methods to boost mobile activity recognition performance. The first is CoDATS [30], and the second is called CALDA. Both methods perform unsupervised domain adaptation and can incorporate weak supervision. These two methods illustrated in Figure 1. CoDATS uses a domain adversarial neural network [8]. Building on the CoDATS architecture, CALDA additionally incorporates contrastive learning [11]. Because both of these methods support leveraging weak supervision



**Figure 1: Time Series Domain Adaptation Methods: (1) CoDATS consists of the feature extractor, task classifier, and domain classifier. (2) CALDA additionally adds a contrastive head to which a supervised contrastive loss is applied. This contrastive strategy pulls together same-labeled examples and pushes apart different-labeled examples.**

information, such as target domain label proportions that participants may be able to self-report easier than labeling additional data instances, we include the two weak supervision variations of these methods, CoDATS-WS and CALDA-WS, as well.

**4.2.1 CoDATS.** The CoDATS time series domain adaptation method consists of a time series compatible model architecture with three components: a feature extractor, a domain classifier, and a task classifier [30]. The feature extractor is based on 1D convolutions, converting the raw time series data into a learned representation space. The resulting feature representation is then input into the domain and task classifiers. The multi-class domain classifier acts as an adversary, learning to classify which domain the data originated from. By negating this gradient, CoDATS can learn a domain-invariant representation that enables improved transfer from one or more source domains to the target domain.

Formally, given a feature extractor  $F$  with parameters  $\theta_f$ , a domain classifier  $D$  with parameters  $\theta_d$ , a task classifier  $C$  with parameters  $\theta_c$ , a gradient reversal layer  $\mathcal{R}$  [8], cross-entropy losses  $\mathcal{L}_d$  for  $D$  and  $\mathcal{L}_y$  for  $C$ , and domain labels  $d_{S_i}$  for the source domains and  $d_T$  for the target domain, we obtain the following CoDATS training objective [30]:

$$\operatorname{argmin}_{\theta_f, \theta_c, \theta_d} \sum_{i=1}^n \left[ \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{D}_{S_i}} [\mathcal{L}_y(C(F(\mathbf{x})), y) + \mathcal{L}_d(D(\mathcal{R}(F(\mathbf{x}))), d_{S_i}) + \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_T^X} [\mathcal{L}_d(D(\mathcal{R}(F(\mathbf{x}))), d_T)] \right] \quad (5)$$

**4.2.2 CALDA.** Building upon CoDATS, CALDA additionally incorporates contrastive learning [31]. This domain adaptation algorithm adds an additional contrastive head component to the network architecture, to which an InfoNCE supervised contrastive loss is applied [11, 19]. In the representation space of this new contrastive head, the supervised contrastive loss pulls together source domain examples that have the same label and pushes apart source domain examples that have different labels. This process is illustrated in Figure 1 (right). Such feature-space modification is performed via pairs of examples in the training data. First, a query example is selected. Then, a positive set is constructed consisting of examples with the same label as the query, and a negative set is constructed

consisting of examples with different labels than the query. Finally, the contrastive loss in Equation 6 is minimized for each of the query examples, which pulls the positive set of examples closer to the query and pushes the negative set of examples away. Through this process, CALDA produces a feature representation that generalizes better among the multiple domains, and thereby improves target domain performance.

Formally, given a projected representation  $z$  of the query, the positive set  $P$ , the negative set  $N$ , a temperature  $\tau$ , and cosine similarity  $\text{sim}(z_1, z_2) = \frac{z_1^T z_2}{\|z_1\| \|z_2\|}$ , we can define an InfoNCE supervised contrastive loss [11, 19]:

$$\mathcal{L}_c(z, P, N) = \frac{1}{|P|} \sum_{z_p \in P} \left[ -\log \left( \frac{\exp\left(\frac{\text{sim}(z, z_p)}{\tau}\right)}{\sum_{z_k \in N \cup \{z_p\}} \exp\left(\frac{\text{sim}(z, z_k)}{\tau}\right)} \right) \right] \quad (6)$$

Next, we define the query, positive, and negative sets using set-builder notation. Given an auxiliary set  $K$  containing the input  $x$ , class label  $y$ , and domain label  $d$ , the set of labeled examples  $S_i \sim \mathcal{D}_{S_i}$  from each of the  $n$  source domains, and the projected representation from the contrastive head  $Z$  with parameters  $\theta_z$ , we define the query set  $Q_{S_i}$  for each source domain, positive set  $P_{d_q, y_q}$  for a query with domain  $d_q$  and label  $y_q$ , and negative set  $N_{d_q, y_q}$  similarly. Note that the difference between the positive and negative sets is whether the positive/negative label matches the query’s label or not.

$$K = \{(x, y, d) \mid (x, y) \in S_i, i \in \{1, 2, \dots, n\}, d = d_{S_i}\} \quad (7)$$

$$Q_{S_i} = \{(z, y) \mid (x, y, d) \in K, z = Z(F(x)), d = d_{S_i}\} \quad (8)$$

$$P_{d_q, y_q} = \{z \mid (x, y, d) \in K, z = Z(F(x)), y = y_q\} \quad (9)$$

$$N_{d_q, y_q} = \{z \mid (x, y, d) \in K, z = Z(F(x)), y \neq y_q\} \quad (10)$$

Then, given a feature extractor  $F$  with parameters  $\theta_f$ , a contrastive head  $Z$  with parameters  $\theta_z$ , the query set from each source domain  $Q_{S_i}$ , and positive and negative sets  $P_{d, y}$  and  $N_{d, y}$  for each query from domain  $d$  and with label  $y$ , we obtain the contrastive training objective:

$$\text{argmin}_{\theta_f, \theta_z} \sum_{i=1}^n \left[ \frac{1}{|Q_{S_i}|} \sum_{(z_q, y_q) \in Q_{S_i}} \mathcal{L}_c(z_q, P_{d_{S_i}, y_q}, N_{d_{S_i}, y_q}) \right] \quad (11)$$

**4.2.3 Weak Supervision.** Both CoDATS and CALDA support leveraging weak supervision information if available. By incorporating weak supervision methods into domain adaptation, we can determine the level of performance improvement from utilizing the weak supervision information. To evaluate the benefit of weak supervision, in CoDATS-WS and CALDA-WS we add a KL divergence regularization term to the training objective that aligns the predicted task classifier’s label distribution for the target domain’s unlabeled data with the known target domain label proportion distribution [30].

Formally, given a feature extractor  $F$  with parameters  $\theta_f$ , a task classifier  $C$  with parameters  $\theta_c$ , and the known target-domain label proportion distribution  $Y_{true}$ , we add the following KL divergence

**Table 1: Per-Class Variation in HAR Datasets for Average Euclidean Distance (ED) and Average KL Divergence (KL)**

Dataset	Dim	Domains	Examples	Class	ED	KL
MAR/	20	15	2122	Cook	66.6 ± 23.8	4551 ± 2333
YA	20	15	2608	Eat	64.1 ± 28.5	4301 ± 2276
	20	15	1798	Exercise	83.9 ± 28.3	5165 ± 2444
	20	15	1529	Hygiene	70.5 ± 25.5	4722 ± 2399
	20	15	2802	Travel	89.0 ± 31.7	5114 ± 2206
	20	15	13955	Work	49.2 ± 18.4	3279 ± 2117
				<b>Average</b>	<b>70.6 ± 26.4</b>	<b>4522 ± 2299</b>
MAR/	20	25	7153	Errands	83.1 ± 29.4	4651 ± 2152
OA	20	25	5707	Exercise	81.0 ± 25.9	4684 ± 2300
	20	25	7005	Hobby	58.0 ± 23.0	3833 ± 2185
	20	25	5222	Hygiene	64.2 ± 20.6	3792 ± 2110
	20	25	13723	Mealtime	58.3 ± 17.6	3786 ± 2167
	20	25	14474	Relax	51.9 ± 18.2	3426 ± 2244
	20	25	1589	Sleep	47.6 ± 18.0	3126 ± 2327
	20	25	4940	Travel	66.4 ± 26.2	4104 ± 2038
	20	25	10382	Work	56.0 ± 16.8	4016 ± 2139
				<b>Average</b>	<b>62.9 ± 22.2</b>	<b>3935 ± 2186</b>
UCI	9	30	1237	Laying	7.4 ± 5.4	350 ± 592
HAR	9	30	1133	Sitting	4.4 ± 2.5	227 ± 350
	9	30	1213	Standing	3.2 ± 1.5	107 ± 134
	9	30	1100	Walking	14.8 ± 2.7	798 ± 220
	9	30	899	Downstairs	18.3 ± 3.1	965 ± 273
	9	30	984	Upstairs	16.5 ± 3.0	899 ± 307
				<b>Average</b>	<b>10.8 ± 3.3</b>	<b>558 ± 344</b>
UCI	3	9	10068	Bike	21.1 ± 5.6	805 ± 545
HHAR	3	9	10854	Sit	14.8 ± 8.2	455 ± 596
	3	9	8808	Stairsdown	32.8 ± 7.0	840 ± 420
	3	9	9713	Stairsup	29.3 ± 6.7	768 ± 442
	3	9	10093	Stand	8.9 ± 5.8	381 ± 609
	3	9	11953	Walk	29.6 ± 5.0	863 ± 391
				<b>Average</b>	<b>22.8 ± 6.5</b>	<b>685 ± 508</b>
WISDM	3	33	492	Downstairs	21.3 ± 4.5	815 ± 358
AR	3	33	1665	Jogging	35.5 ± 3.8	1078 ± 361
	3	33	293	Sitting	16.0 ± 8.2	548 ± 819
	3	33	239	Standing	10.9 ± 5.6	393 ± 605
	3	33	605	Upstairs	20.8 ± 4.7	768 ± 369
	3	33	1972	Walking	22.6 ± 3.3	866 ± 295
				<b>Average</b>	<b>21.2 ± 5.3</b>	<b>745 ± 503</b>
WISDM	3	51	2029	Jogging	32.8 ± 7.5	1023 ± 554
AT	3	51	1194	Lying Down	12.9 ± 12.2	374 ± 815
	3	51	2718	Sitting	20.3 ± 15.9	788 ± 986
	3	51	173	Stairs	28.4 ± 9.9	849 ± 549
	3	51	1116	Standing	19.7 ± 15.9	551 ± 697
	3	51	3330	Walking	26.3 ± 7.4	919 ± 517
				<b>Average</b>	<b>23.4 ± 12.0</b>	<b>751 ± 707</b>

regularization term to the training objective:

$$\text{argmin}_{\theta_f, \theta_c} \left[ D_{KL} \left( Y_{true} \parallel \mathbb{E}_{x \sim \mathcal{D}_T^x} [C(F(x))] \right) \right] \quad (12)$$

## 5 EXPERIMENTAL RESULTS

We experimentally evaluate the impact of unsupervised domain adaptation on the CASAS MAR mobile activity recognition problem. The CASAS MAR dataset provides multiple opportunities for domain adaptation. First, we collect data from multiple participants.

The methods and movements each person uses to perform activities is highly variable, so we can view each person’s data as a separate domain for adaptation. If successful, the domain adaptation will be effective at recognizing activities for new persons who cannot provide ground-truth activity labels. Second, we view data collected at distinct time points as separate domains. Because behavior changes over time, a model that is trained from ground-truth data at one timepoint may not generalize well enough to adequately recognize the same activity classes performed at a different time. Thus, we can employ unsupervised domain adaptation to adapt models to new seasons, years, or other timeframes.

In these experiments, we evaluate the CoDATS and CALDA methods, with and without weak supervision. Following the CoDATS training methodology, for the cross-person experiments, we select 3 random sets of source domains for each of 10 different target domains for each value of  $n$  (the number of source domains). The error bars represent the average of the standard deviation across the set of 3 random sets of source domains. For the cross-time experiments, we perform adaptation from one week to another week, averaged over each of the different participants. The error bars represent standard deviation.

We additionally include a “No Adaptation” lower bound baseline, which enables us to measure the performance improvement achieved by utilizing domain adaptation. Finally, we include an upper bound “Train on Target”. This model is trained using the labels from the training set of the target domain. While this method “cheats” by accessing labeled target training data, it provides an upper bound estimate of how well a model could perform if ample target domain labels were available. We report performance in each based based on AUC values for instances in the target domain.

## 5.1 Variability of Human Behavior

This work is motivated by the need to perform activity recognition in naturalistic, unscripted settings. We expect that data collected while participants perform normal daily activities will exhibit much greater variation than is the case for scripted activities. To experimentally compare these datasets, in Table 1, we measure the variation among examples of each class for the CASAS MAR dataset as well as other publicly-available HAR datasets. We compute the distances between the raw data of all training set examples in each class using two measures: Euclidean distance and KL divergence. We observe differences between labels within CASAS MAR/YA and CASAS MAR/OA and also in the shared labels between these two dataset components (denoted by italics). The differences are indicative of differences in behavior between these groups. We note that the observed distances and standard deviations are much larger in CASAS MAR than in the prior human activity recognition datasets shown in the table, despite all of these datasets being processed in a similar manner. The increased variability is an indication that MAR will be challenging and may benefit from domain adaptation. Given the more challenging nature of CASAS MAR due to the increased variability in behavior, we next proceed to measure the benefit of utilizing time series domain adaptation methods to improve model performance under the individual differences and temporal differences present in natural activity recognition data.

**Table 2: Cross-Person Younger Adults (AUC)**

$n$	No Adaptation	CoDATS	CoDATS-WS	CALDA	CALDA-WS
2	0.686 ± 0.039	0.707 ± 0.035	0.722 ± 0.028	0.750 ± 0.030	<b>0.777 ± 0.026</b>
5	0.707 ± 0.033	0.742 ± 0.025	0.749 ± 0.019	0.790 ± 0.015	<b>0.818 ± 0.012</b>
8	0.717 ± 0.028	0.759 ± 0.018	0.763 ± 0.015	0.801 ± 0.017	<b>0.828 ± 0.013</b>
11	0.724 ± 0.017	0.779 ± 0.013	0.777 ± 0.012	0.811 ± 0.009	<b>0.837 ± 0.008</b>
14	0.735 ± 0.000	0.790 ± 0.000	0.775 ± 0.000	0.813 ± 0.000	<b>0.835 ± 0.000</b>
Avg	0.714 ± 0.024	0.755 ± 0.018	0.757 ± 0.015	0.793 ± 0.014	<b>0.819 ± 0.012</b>

**Table 3: Cross-Person Older Adults (AUC)**

$n$	No Adaptation	CoDATS	CoDATS-WS	CALDA	CALDA-WS
2	0.637 ± 0.043	0.585 ± 0.027	0.617 ± 0.021	0.700 ± 0.043	<b>0.777 ± 0.019</b>
7	0.637 ± 0.035	0.604 ± 0.022	0.625 ± 0.018	0.716 ± 0.031	<b>0.787 ± 0.022</b>
12	0.650 ± 0.037	0.614 ± 0.025	0.646 ± 0.026	0.733 ± 0.027	<b>0.796 ± 0.020</b>
17	0.663 ± 0.024	0.624 ± 0.016	0.655 ± 0.020	0.742 ± 0.020	<b>0.806 ± 0.011</b>
22	0.670 ± 0.016	0.648 ± 0.019	0.677 ± 0.022	0.738 ± 0.010	<b>0.806 ± 0.009</b>
Avg	0.651 ± 0.031	0.615 ± 0.021	0.644 ± 0.022	0.726 ± 0.026	<b>0.794 ± 0.016</b>

**Table 4: Cross-Time Younger Adults (AUC)**

No Adaptation	CoDATS	CoDATS-WS	CALDA	CALDA-WS	Train/Target
week 0 → 4					
0.819 ± 0.034	0.810 ± 0.040	0.834 ± 0.053	<b>0.872 ± 0.042</b>	0.870 ± 0.055	0.996 ± 0.002
week 4 → 0					
0.792 ± 0.044	0.780 ± 0.055	0.791 ± 0.052	<b>0.838 ± 0.044</b>	0.832 ± 0.042	0.992 ± 0.004
Avg					
0.806 ± 0.039	0.795 ± 0.047	0.812 ± 0.053	<b>0.855 ± 0.043</b>	0.851 ± 0.048	0.994 ± 0.003

**Table 5: Cross-Time Older Adults (AUC)**

No Adaptation	CoDATS	CoDATS-WS	CALDA	CALDA-WS	Train/Target
week 0 → 1					
0.636 ± 0.073	0.600 ± 0.070	0.656 ± 0.066	0.706 ± 0.072	<b>0.761 ± 0.076</b>	0.988 ± 0.009
week 0 → 2					
0.630 ± 0.081	0.604 ± 0.073	0.643 ± 0.087	0.701 ± 0.084	<b>0.747 ± 0.090</b>	0.987 ± 0.007
week 0 → 3					
0.622 ± 0.108	0.591 ± 0.090	0.648 ± 0.089	0.670 ± 0.127	<b>0.729 ± 0.096</b>	0.990 ± 0.007
week 1 → 0					
0.633 ± 0.090	0.610 ± 0.067	0.661 ± 0.082	0.673 ± 0.088	<b>0.736 ± 0.084</b>	0.984 ± 0.010
week 2 → 0					
0.630 ± 0.098	0.608 ± 0.084	0.655 ± 0.108	0.663 ± 0.098	<b>0.723 ± 0.089</b>	0.984 ± 0.010
week 3 → 0					
0.615 ± 0.098	0.601 ± 0.078	0.630 ± 0.087	0.659 ± 0.096	<b>0.717 ± 0.087</b>	0.984 ± 0.010
Avg					
0.627 ± 0.091	0.602 ± 0.077	0.649 ± 0.086	0.679 ± 0.094	<b>0.735 ± 0.087</b>	0.986 ± 0.009

## 5.2 Cross-Person Adaptation

The cross-person adaptation results are shown in Tables 2 and 3. Training on the target yields an AUC value of  $0.992 \pm 0.001$  for younger adults and  $0.979 \pm 0.003$  for older adults. With younger adults, we find the domain adaptation approaches improve over

No Adaptation in all cases. These results agree with prior cross-person time series domain adaptation work studying human activity recognition [30, 31]. With older adults, we find that CALDA always improves over No Adaptation, though CoDATS performs slightly worse. The decreased performance for CoDATS indicates the limitations of this baseline method when applied to increasingly challenging datasets. However, the successful adaptation by CALDA points to the benefit of leveraging contrastive learning in MAR domain adaptation.

### 5.3 Cross-Time Adaptation

The cross-time adaptation results are shown in Tables 4 and 5. We observe a benefit from domain adaptation with CoDATS-WS, CALDA, and CALDA-WS for both younger and older adults. Due to the limited data available for the younger adults, we cannot perform adaptation over gradually increasing time gaps. However, in the single cross-time adaptation problem from the first week of data to the week a month later shown in Table 4, we do observe concept drift as illustrated by the No Adaptation performance. For the older adults, we collected data throughout an entire month, enabling measuring the concept drift with increasing time gap. As shown in Table 5, performance in almost all cases degrades with increasing time gaps (with the exception of CoDATS and CoDATS-WS in the forward direction). This decreasing performance over time demonstrates the concept drift over time.

These results motivate the need for cross-time domain adaptation. Collected labeled data for a period of time (e.g., one week or one month) and then subsequently using that model for an extended period of time will result in sub-par performance. As shown in these results, with both younger adults and older adults, this performance is far below that of the upper bound Train on Target. However, some of this performance drop can be regained through performing unsupervised time series domain adaptation, using a combination of the original period of labeled data and additionally some unlabeled data from the more recent time period to produce an adapted model for the more recent time period.

### 5.4 Younger vs. Older Adults

In each of the cross-person and cross-time tables, if we compare the younger adult results with the older adult results, we observe that the older adult model performance is worse than the younger adult model performance. We believe there are two causes for this: increased variability in human behavior and decreased labeling density. In the Appendix, we validate that the difference in class labels is not the cause.

First, the younger adults were a more homogeneous group, largely consisting of students at one university. This means that there is limited variation for each activity label, e.g., each person’s “work” label is likely related to homework and each person’s “travel” likely consists of walking to classes. In contrast, the older adults may have vastly different “work” or “travel” labels depending on their occupation and where they live.

Second, we collected a higher density of labeled data for the younger adults by prompting each person to provide a label once every 30 minutes, due to the shorter duration of this study and the corresponding limited-time disruption to participant routines.

On the other hand, with older adults, the study duration was over twice as long, so we limited the number of daily queries for activity labels.

## 6 CONCLUSION

While human activity recognition is a popular area of research, little effort has focused on recognizing complex behaviors in the wild. Creating robust recognition methods in naturalistic settings with limited ground truth labels benefits from leveraging unsupervised domain adaptation techniques. These methods are further improved by incorporating contrastive learning and weak supervision. To validate our methods and provide a resource for mobile activity recognition, we have collected two MAR dataset components consisting of data collected from both younger and older adults performing activities in natural settings. We demonstrated the increased variability present in these datasets when compared with existing publicly-available datasets, indicating the challenging nature of our datasets. Our experimental results, both for cross-person and cross-time adaptation, verify that MAR achieves performance gains by leveraging contrastive learning in addition to adversarial learning. Future work includes developing more sophisticated time series domain adaptation methods that can better address the large domain gaps present in such challenging realistic datasets.

## ACKNOWLEDGMENTS

The authors would like to thank Justin Frow and Maureen Schmitter-Edgcombe for their help with the data collection. Research reported in this publication was supported by the National Institutes of Health under award R25AG046114 and by the National Science Foundation under award 1954372.

## REFERENCES

- [1] Davide Anguita, Alessandro Ghio, Luca Oneto, Xavier Parra, and Jorge Luis Reyes-Ortiz. 2013. A public domain dataset for human activity recognition using smartphones. In *European Symposium on Artificial Neural Networks*. 437–442.
- [2] Kaixuan Chen, Dalin Zhang, Lina Yao, Bin Guo, Zhiwen Yu, and Yunhao Liu. 2022. Deep learning for sensor-based human activity recognition: Overview, challenges, and opportunities. *Comput. Surveys* 54 (2022), 1–40. Issue 4.
- [3] OpenStreetMap contributors. 2021. OpenStreetMap Nominatim. (December 2021). Retrieved December 13, 2021 from <https://nominatim.openstreetmap.org>.
- [4] Eric Cope. 2009. *Estimating human movement using a three axis accelerometer*. Ph.D. Dissertation. Arizona State University.
- [5] L Minh Dang, Kyungbok Min, Hanxiang Wang, Md. Jalil Piran, Cheol Hee Lee, and Hyeonjoon Moon. 2020. Sensor-based and vision-based human activity recognition: A comprehensive survey. *Pattern Recognition* 108 (2020), 107561.
- [6] Florenc Demrozi, Graziano Pravadelli, Azra Bihorac, and Parisa Rashidi. 2020. Human activity recognition using inertial, physiological, and environmental sensors: A comprehensive survey. *IEEE Access* 8 (2020), 210816–210836.
- [7] Raymond Dodge. 1932. Conditions and Consequences of Human Variability. *Nature* 130 (1932), Issue 331.
- [8] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *Journal of Machine Learning Research* 17, 59 (2016), 1–35. <http://jmlr.org/papers/v17/15-239.html>
- [9] Hui Jun Guo and Amit Sapra. 2020. Instrumental activity of daily living. (2020). Retrieved December 15, 2021, from <https://www.ncbi.nlm.nih.gov/books/NBK553126/>.
- [10] Mohammed Mehedi Hassan, Md. Zia Uddin, Amr Mohamed, and Ahmad Almgren. 2018. A robust human activity recognition system using smartphone sensors and deep learning. *Future Generation Computer Systems* 81 (2018), 307–313.
- [11] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschiot, Ce Liu, and Dilip Krishnan. 2020. Supervised contrastive learning. *arXiv preprint arXiv:2004.11362* (2020).
- [12] Jennifer R. Kwapisz, Gary M. Weiss, and Samuel A. Moore. 2011. Activity recognition Using cell phone accelerometers. *SIGKDD Explor. Newsl.* 12, 2 (March 2011),

- 74–82. <https://doi.org/10.1145/1964897.1964918>
- [13] Qiao Liu and Hui Xue. 2021. Adversarial spectral kernel matching for unsupervised time series domain adaptation. In *Thirtieth International Joint Conference on Artificial Intelligence*.
- [14] Jeffrey W. Lockhart, Gary M. Weiss, Jack C. Xue, Shaun T. Gallagher, Andrew B. Grosner, and Tony T. Pulickal. 2011. Design considerations for the WISDM smart phone-based sensor mining architecture. In *Proceedings of the Fifth International Workshop on Knowledge Discovery from Sensor Data* (San Diego, California) (*SensorkDD '11*). ACM, New York, NY, USA, 25–33. <https://doi.org/10.1145/2003653.2003656>
- [15] T M Marteau, G J Hollands, and P C Fletcher. 2012. Changing human behavior to prevent disease: The importance of targeting automatic processes. *Science* 337 (2012), 1492–1495.
- [16] Lisbeth Nielsen, Melissa Riddle, and Jonathan W King. 2018. The NIH science of behavior change program: Transforming the science through a focus on mechanisms of change. *Behaviour Research and Therapy* 101 (2018), 3–11.
- [17] U.S. Department of Health and Human Services. 2021. Healthy people 2030. (2021). Retrieved December 15, 2021, from <https://www.healthypeople.gov/2020/>.
- [18] U.S. Bureau of Labor Statistics. 2021. American Time Use Survey. (2021). Retrieved December 15, 2021, from <https://www.bls.gov/tus/>.
- [19] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748* (2018).
- [20] Liangying Peng, Ling Chen, Zhenan Ye, and Yi Zhang. 2018. AROMA: A deep multi-task learning based simple and complex human activity recognition method using wearable sensors. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 2 (2018), 1–16. Issue 2.
- [21] Srdan Popić, Dražen Pezer, Bojan Mrazovac, and Nikola Teslić. 2016. Performance evaluation of using Protocol buffers in the Internet of Things communication. In *2016 International Conference on Smart Systems and Technologies (SST)*. IEEE, 261–265.
- [22] Sanjay Purushotham, Wilka Carvalho, Tanachat Nilanon, and Yan Liu. 2017. Variational adversarial deep domain adaptation for health care time series analysis. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=rk9eAFcxg>
- [23] Andrea Rosales Sanabria, Franco Zambonelli, Simon Dobson, and Juan Ye. 2021. ContrasGAN: Unsupervised domain adaptation in Human Activity Recognition via adversarial and contrastive learning. *Pervasive and Mobile Computing* (2021), 101477.
- [24] Muhammad Shoaib, Stephan Bosch, Ozlem Incel, Hans Scholten, and Paul Havinga. 2016. Complex human activity recognition using smartphone and wrist-worn motion sensors. *Sensors* 16, 4 (2016), 426.
- [25] Muhammad Shoaib, Stephan Bosch, Ozlem Durmaz Incel, Hans Scholten, and Paul JM Havinga. 2015. A survey of online activity recognition using mobile phones. *Sensors* 15, 1 (2015), 2059–2085.
- [26] Muhammad Shoaib, Stephan Bosch, Hans Scholten, Paul JM Havinga, and Ozlem Durmaz Incel. 2015. Towards detection of bad habits by fusing smartphone and smartwatch sensors. In *2015 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*. IEEE, 591–596.
- [27] Allan Stisen, Henrik Blunck, Sourav Bhattacharya, Thor Siiger Prentow, Mikkel Baun Kjærgaard, Anind Dey, Tobias Sonne, and Mads Møller Jensen. 2015. Smart devices Are different: Assessing and mitigating mobile sensing heterogeneities for activity recognition. In *PACM Conference on Embedded Networked Sensor Systems* (Seoul, South Korea) (*SenSys '15*). ACM, New York, NY, USA, 127–140. <https://doi.org/10.1145/2809695.2809718>
- [28] Praneeth Vepakomma, Debraj De, Sajal K Das, and Shekhar Bhansali. 2015. A-Wristocracy: Deep learning on wrist-worn sensing for recognition of user complex activities. In *IEEE International Conference on Wearable and Implantable Body Sensor Networks*.
- [29] Garrett Wilson and Diane J. Cook. 2020. A Survey of Unsupervised Deep Domain Adaptation. *ACM Trans. Intell. Syst. Technol.* 11, 5, Article 51 (July 2020), 46 pages. <https://doi.org/10.1145/3400066>
- [30] Garrett Wilson, Janardhan Rao Doppla, and Diane J. Cook. 2020. Multi-source deep domain adaptation with weak supervision for time-series sensor data. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (Virtual Event, CA, USA) (*KDD '20*). Association for Computing Machinery, New York, NY, USA, 1768–1778. <https://doi.org/10.1145/3394486.3403228>
- [31] Garrett Wilson, Janardhan Rao Doppla, and Diane J. Cook. 2021. CALDA: Improving Multi-Source Time Series Domain Adaptation with Contrastive Adversarial Learning. *CoRR* abs/2109.14778 (2021). [arXiv:2109.14778](https://arxiv.org/abs/2109.14778) <https://arxiv.org/abs/2109.14778>
- [32] Zhan Zhao, Haris N Koutsopoulos, and Jinhua Zhao. 2018. Detecting pattern changes in individual travel behavior: A Bayesian approach. *Transportation Research Part B: Methodological* 112 (2018), 73–88.

**Table 6: Cross-Person Younger Adults (Accuracy)**

$n$	No Adaptation	CoDATS	CoDATS-WS	CALDA	CALDA-WS	Train/Target
2	42.1 ± 6.4	44.0 ± 6.1	46.9 ± 3.6	47.6 ± 5.1	<b>50.8 ± 3.7</b>	95.1 ± 0.3
5	45.0 ± 5.3	46.8 ± 4.1	49.8 ± 2.5	52.3 ± 2.1	<b>56.1 ± 2.4</b>	95.1 ± 0.3
8	45.1 ± 3.7	48.2 ± 2.8	50.6 ± 2.1	53.3 ± 2.1	<b>56.4 ± 1.8</b>	95.1 ± 0.3
11	45.1 ± 3.2	50.9 ± 2.1	51.3 ± 2.2	55.0 ± 1.4	<b>57.3 ± 1.5</b>	95.1 ± 0.3
14	46.3 ± 0.0	51.9 ± 0.0	52.3 ± 0.0	54.4 ± 0.0	<b>57.1 ± 0.0</b>	95.1 ± 0.3
Avg	44.7 ± 3.7	48.3 ± 3.0	50.2 ± 2.1	52.5 ± 2.1	<b>55.5 ± 1.9</b>	95.1 ± 0.3

**Table 7: Cross-Person Older Adults (Accuracy)**

$n$	No Adaptation	CoDATS	CoDATS-WS	CALDA	CALDA-WS	Train/Target
2	22.1 ± 5.3	21.9 ± 3.8	27.6 ± 2.8	26.2 ± 6.0	<b>35.8 ± 3.2</b>	91.7 ± 0.6
7	20.8 ± 4.5	21.8 ± 3.6	26.4 ± 2.7	27.1 ± 5.5	<b>35.9 ± 5.3</b>	91.7 ± 0.6
12	23.4 ± 4.5	20.6 ± 2.9	26.8 ± 3.5	28.5 ± 5.4	<b>37.3 ± 4.9</b>	91.7 ± 0.6
17	23.3 ± 2.7	20.4 ± 2.1	27.4 ± 2.6	29.3 ± 3.0	<b>38.4 ± 2.7</b>	91.7 ± 0.6
22	22.9 ± 1.7	21.7 ± 2.2	28.5 ± 2.7	28.4 ± 1.9	<b>38.5 ± 3.2</b>	91.7 ± 0.6
Avg	22.5 ± 3.7	21.3 ± 2.9	27.3 ± 2.9	27.9 ± 4.4	<b>37.2 ± 3.9</b>	91.7 ± 0.6

**Table 8: Cross-Time Younger Adults (Accuracy)**

No Adaptation	CoDATS	CoDATS-WS	CALDA	CALDA-WS	Train/Target
<i>week 0 → 4</i>					
61.7 ± 5.7	61.3 ± 7.7	65.1 ± 9.7	63.1 ± 9.2	<b>67.3 ± 10.9</b>	97.1 ± 0.9
<i>week 4 → 0</i>					
60.1 ± 7.8	56.7 ± 8.9	56.8 ± 7.9	<b>61.9 ± 10.4</b>	60.1 ± 8.6	95.2 ± 1.3
Avg					
60.9 ± 6.7	59.0 ± 8.3	61.0 ± 8.8	62.5 ± 9.8	<b>63.7 ± 9.7</b>	96.2 ± 1.1

**Table 9: Cross-Time Older Adults (Accuracy)**

No Adaptation	CoDATS	CoDATS-WS	CALDA	CALDA-WS	Train/Target
<i>week 0 → 1</i>					
27.3 ± 11.4	27.3 ± 12.7	35.5 ± 11.1	34.2 ± 11.4	<b>40.9 ± 15.6</b>	96.2 ± 2.2
<i>week 0 → 2</i>					
25.6 ± 13.3	28.6 ± 12.6	33.4 ± 13.7	32.5 ± 12.0	<b>37.8 ± 16.7</b>	96.1 ± 2.0
<i>week 0 → 3</i>					
27.0 ± 16.2	25.9 ± 15.4	34.2 ± 15.1	28.1 ± 16.1	<b>34.4 ± 18.6</b>	96.1 ± 2.0
<i>week 1 → 0</i>					
28.8 ± 14.8	29.2 ± 11.5	<b>36.8 ± 13.7</b>	31.8 ± 13.2	36.1 ± 16.0	94.6 ± 3.0
<i>week 2 → 0</i>					
27.4 ± 17.4	28.8 ± 14.1	36.3 ± 18.8	29.2 ± 12.6	<b>36.5 ± 16.9</b>	94.6 ± 3.0
<i>week 3 → 0</i>					
25.2 ± 15.7	27.9 ± 13.3	32.0 ± 14.9	27.7 ± 14.0	<b>34.4 ± 15.4</b>	94.6 ± 3.0
Avg					
26.9 ± 14.8	27.9 ± 13.3	34.7 ± 14.5	30.6 ± 13.2	<b>36.7 ± 16.5</b>	95.4 ± 2.5

## A ADDITIONAL RESULTS

In the main paper we include results evaluated in terms of AUC due to the imbalanced nature of the datasets. However, prior works often use accuracy as the evaluation metric. Here we additionally include the results measured in terms of accuracy, shown in Tables 6 through 9.



**Table 10: Cross-Person Younger Adults (AUC) - Mapped**

$n$	None	CoDATS	CoDATS-WS	CALDA	CALDA-WS
2	0.686 ± 0.053	0.706 ± 0.033	0.729 ± 0.027	0.750 ± 0.038	<b>0.785 ± 0.028</b>
5	0.707 ± 0.037	0.742 ± 0.030	0.751 ± 0.025	0.795 ± 0.019	<b>0.829 ± 0.010</b>
8	0.713 ± 0.037	0.762 ± 0.017	0.764 ± 0.019	0.805 ± 0.017	<b>0.836 ± 0.010</b>
11	0.731 ± 0.022	0.781 ± 0.018	0.782 ± 0.012	0.813 ± 0.011	<b>0.841 ± 0.008</b>
14	0.730 ± 0.000	0.792 ± 0.000	0.781 ± 0.000	0.820 ± 0.000	<b>0.842 ± 0.000</b>
Avg	0.713 ± 0.030	0.757 ± 0.020	0.761 ± 0.017	0.797 ± 0.017	<b>0.827 ± 0.011</b>

**Table 11: Cross-Person Older Adults (AUC) - Mapped**

$n$	None	CoDATS	CoDATS-WS	CALDA	CALDA-WS
2	0.643 ± 0.051	0.602 ± 0.054	0.677 ± 0.037	0.695 ± 0.060	<b>0.777 ± 0.041</b>
7	0.632 ± 0.051	0.611 ± 0.041	0.677 ± 0.034	0.706 ± 0.038	<b>0.794 ± 0.022</b>
12	0.632 ± 0.059	0.620 ± 0.036	0.684 ± 0.028	0.730 ± 0.039	<b>0.799 ± 0.023</b>
17	0.637 ± 0.039	0.630 ± 0.025	0.699 ± 0.025	0.724 ± 0.022	<b>0.803 ± 0.014</b>
22	0.659 ± 0.032	0.649 ± 0.028	0.707 ± 0.020	0.734 ± 0.018	<b>0.804 ± 0.013</b>
Avg	0.640 ± 0.046	0.622 ± 0.037	0.689 ± 0.029	0.718 ± 0.035	<b>0.795 ± 0.022</b>

**Table 12: Cross-Time Younger Adults (AUC) - Mapped**

	None	CoDATS	CoDATS-WS	CALDA	CALDA-WS	Train/Target
<i>week 0 → 4</i>						
0.824 ± 0.041	0.792 ± 0.035	0.838 ± 0.046	0.859 ± 0.065	<b>0.883 ± 0.046</b>	0.996 ± 0.002	
<i>week 4 → 0</i>						
0.801 ± 0.044	0.783 ± 0.083	0.802 ± 0.060	0.838 ± 0.049	<b>0.842 ± 0.053</b>	0.991 ± 0.003	
<i>Avg</i>						
0.813 ± 0.042	0.787 ± 0.059	0.820 ± 0.053	0.849 ± 0.057	<b>0.862 ± 0.050</b>	0.993 ± 0.002	

**Table 13: Cross-Time Older Adults (AUC) - Mapped**

	None	CoDATS	CoDATS-WS	CALDA	CALDA-WS	Train/Target
<i>week 0 → 1</i>						
0.632 ± 0.142	0.645 ± 0.132	0.734 ± 0.141	0.703 ± 0.114	<b>0.769 ± 0.143</b>	0.993 ± 0.009	
<i>week 0 → 2</i>						
0.719 ± 0.136	0.660 ± 0.137	0.727 ± 0.137	0.722 ± 0.125	<b>0.802 ± 0.129</b>	0.993 ± 0.007	
<i>week 0 → 3</i>						
0.656 ± 0.156	0.590 ± 0.169	0.692 ± 0.144	0.670 ± 0.199	<b>0.766 ± 0.129</b>	0.990 ± 0.010	
<i>week 1 → 0</i>						
0.643 ± 0.121	0.625 ± 0.122	0.699 ± 0.120	0.661 ± 0.123	<b>0.756 ± 0.124</b>	0.985 ± 0.014	
<i>week 2 → 0</i>						
0.665 ± 0.134	0.648 ± 0.141	0.704 ± 0.138	0.678 ± 0.141	<b>0.709 ± 0.116</b>	0.985 ± 0.014	
<i>week 3 → 0</i>						
0.656 ± 0.156	0.629 ± 0.144	0.658 ± 0.132	0.607 ± 0.211	<b>0.741 ± 0.117</b>	0.985 ± 0.014	
<i>Avg</i>						
0.662 ± 0.141	0.633 ± 0.141	0.702 ± 0.136	0.673 ± 0.152	<b>0.757 ± 0.127</b>	0.989 ± 0.011	

Because CASAS MAR/YA and CASAS MAR/OA have different label sets, we performed experiments to validate that this is not the cause of the increased difficulty of the CASAS MAR/OA dataset. We mapped the *cook* and *eat* labels for CASAS MAR/YA to *mealtime* and removed the *errands*, *hobby*, *relax*, and *sleep* labels from CASAS MAR/OA. This yields both datasets consisting of the same 5 labels: *work*, *travel*, *exercise*, *hygiene*, and *mealtime*. The results are shown

in Tables 10 through 13. The train on target performance for cross-person younger adults is  $0.992 \pm 0.001$  and for older adults is  $0.981 \pm 0.002$ . In Table 13, we also exclude data from one participant due to there not being sufficient data for train/validation/test sets after mapping. As expected, model performance on the older adult dataset is lower than on the younger adult dataset, agreeing with the results in the main paper.