

Contrastive Learning with Auxiliary User Detection for Identifying Activities

Wen Ge[§], Guanyi Mou[§], Emmanuel O. Agu, Kyumin Lee

Computer Science Department, Worcester Polytechnic Institute Worcester, MA, USA

{wge, gmou, emmanuel, kmlee}@wpi.edu

Abstract—Human Activity Recognition (HAR) is an important task in ubiquitous computing, with impactful real-world applications. While recent state-of-the-art HAR research has demonstrated impressive performance, some key aspects remain under-explored. First, we believe that for optimal performance, HAR models should be both Context-Aware (CA) and personalized. However, prior work has predominantly focused on being Context-Aware (CA), largely ignoring being User-Aware (UA). We argue that learning user-specific differences in performing various activities is as critical as considering user context while performing HAR tasks. Secondly, we believe that the predictions of HAR models should be unified, reliably recognizing the same activity even when performed by different users. As such, the representations utilized by CA and UA models should explicitly place different users performing the same activity closer together. Moreover, identifying the user performing an activity is useful in applications such as thwarting cheating by having another person perform medically-prescribed activities.

To bridge this gap, we introduce Contrastive Learning with Auxiliary User Detection for Identifying Activities (CLAUDIA), a novel framework designed to address these issues. Specifically, we expand the contextual scope of the CA-HAR task by integrating User Identification (UI) within the CA-HAR framework, jointly predicting both CA-HAR and UI in a new task called User and Context-Aware HAR (UCA-HAR). This approach enriches personalized and contextual understanding by jointly learning user-invariant and user-specific patterns. Inspired by state-of-the-art designs in the visual domain, we introduce a supervised contrastive loss objective on instance-instance pairs to enhance model efficacy and improve learned feature quality. Through theoretical exposition, empirical analysis of real-world datasets, and rigorous experimentation, we demonstrate the significance of each component of CLAUDIA and discuss its relationship with existing methodologies. Evaluation across three real-world CA-HAR datasets reveals substantial performance enhancements, with average improvements ranging from 11.7% to 14.2% in Matthew’s Correlation Coefficient (MCC) and 5.4% to 7.3% in Macro F1 score. To encourage further research, we share code with additional supplement material in repository <https://github.com/GMouYes/CLAUDIA>.

Index Terms—Ubiquitous and mobile computing, supervised learning, human activity recognition

I. INTRODUCTION

Human Activity Recognition (HAR), which identifies activities such as walking or running from sensor data, plays a crucial role in health and medical applications [1]–[4]. HAR is challenging due to co-occurring activities (e.g., talking

while walking), making it a multi-label classification task that requires modeling these relationships.

Human activities are highly *contextualized*; people perform different activities in different contexts/situations. Consequently, recent studies have increasingly focused on Context-Aware Human Activity Recognition (CA-HAR) [5]–[7] as a promising approach. The fundamental premise of CA-HAR lies in acknowledging and capturing the significant impact of context on sensor signals by integrating context classification as an auxiliary task [8]–[10]. However, another crucial perspective is often overlooked: human activities can also be highly *personalized*. The variety and heterogeneity in human behavior inevitably lead to variances in action-performing patterns [11], [12], even for the same activity. HAR frameworks focusing solely on external contextual settings may easily overlook the innate differences across users performing the activities, potentially leading to erroneous results. Inspired by this observation, our work aligns with the field of CA-HAR research while also extending beyond its scope by devising a novel framework that accounts for both external contexts and internal user identities. Specifically, we incorporate User Identification (UI) as an auxiliary subtask, effectively extending traditional context-aware (CA) HAR design into user-context-aware (UCA) HAR. Identifying the user performing an activity is useful in applications such as thwarting cheating by having another person perform medically-prescribed activities.

Adding UI as a subtask enables the model to acknowledge user identity but does not necessarily guarantee overcoming inter-user differences. To address the challenge of user-specific variability, we introduce Contrastive Learning with Auxiliary User Detection for Identifying Activities (CLAUDIA) with a supervised contrastive loss that regularizes (different-user, same-activity) representations, ensuring closer vector similarity and capturing both user-specific and user-invariant information effectively.

In summary, we propose a novel framework **Contrastive Learning with Auxiliary User Detection for Identifying Activities** (CLAUDIA) that incorporates User Identification (UI) as a subtask alongside the primary CA-HAR task, explicitly regularizing the modeling process using an additional loss objective. This User-Context-Aware HAR (UCA-HAR) framework achieves a balance between capturing user-invariant patterns and user-specific attributes in a single model. Drawing inspiration from inter-instance relationship modeling in computer vision and natural language processing domains,

DARPA grant HR00111780032-WASH-FP-031 sponsored the research.

[§]Equal contribution

we introduce a novel supervised contrastive loss on instance-instance pairs. By leveraging inter-instance relationships, the CLAUDIA framework bridges the gaps observed in current CA-HAR approaches, resulting in refined instance representations. This research offers the following contributions:

- We extend the context definition in CA-HAR by incorporating a UI subtask to capture both user-invariant and user-specific patterns, improving HAR performance.
- We propose the CLAUDIA framework, which integrates multi-modality sensor data and a novel supervised contrastive loss to enhance HAR.
- Extensive experiments demonstrate that CLAUDIA outperforms state-of-the-art baselines across several datasets, showcasing its adaptability and effectiveness.

II. RELATED WORK

This section provides an overview of related work, including prior CA-HAR model designs, recent advancements in the UI task, and research related to contrastive learning.

A. CA-HAR Models

Prior CA-HAR work primarily focused on two aspects: 1) Directly modeling instance-label relationships and 2) Enhancing (1) by auxiliary modeling of label-label relationships.

1) *Directly Modeling Instance-Label Relationships*: Early CA-HAR models used handcrafted sensor features with traditional or deep learning methods [8], [13]. However, manual feature extraction is labor-intensive and lacks generalizability. Recently, Deep Neural Networks (DNNs) have shown success by auto-extracting features, using architectures such as CNNs [14], RNNs [15], and Graph Neural Networks (GNNs) [16]. Though DNNs offer improved generalization, they often lack interpretability, which is addressed in recent works combining learned and handcrafted features [17], [18].

2) *Enhanced Auxiliary Modeling of Inter-Label Relationships*: To capture label relationships, some studies transformed the task into graph learning, such as HAR-GCNN [19], which uses graph CNNs, and HHGNN [9], encoding user, context, and activity labels as nodes and co-occurrences as hyperedges. These approaches enhanced HAR performance by capturing label dependencies. Despite these advances, inter-instance relationships remain under-explored in CA-HAR research. Our work introduces supervised contrastive learning on instance pairs to address this, improving feature quality (Section III).

B. User Identification and Applications

User identification has gained importance with the rise of smartphones. Zou *et al.* [20] used CNN and LSTM to extract gait features for user identification, while Mekruksavanich and Jitpattanakul [21] developed an ensemble model combining activity and user classifiers. However, these methods struggle with real-world complexity. Our unified UCA-HAR model integrates UI with CA-HAR, showing superior performance across a wide and diverse range of activity scenarios, demonstrating its versatility and potential.

C. Contrastive Loss for HAR

We apply supervised contrastive learning to regularize inter-instance relationships and enhance feature quality. Unlike self-supervised contrastive methods relying on augmentation or external data [22], [23], our approach utilizes labeled data to achieve improved performance [24].

III. PROPOSED CLAUDIA FRAMEWORK

We describe our novel framework in the following subsections. We first illustrate the task formulation in Section III-A. We then demonstrate a high-level picture of our design in Sec. III-B, followed by each sub-module. Lastly, we introduce the loss objectives in Sec. III-C, including a UI sub-task loss.

A. Task Formulation

Given input raw data $X_{raw} \in R^{N \times S \times T}$, where N represents the number of samples, S denotes the number of sensors, and T indicates the snapshots per sample, we compute handcrafted features f , such that:

$$X_{hc} = f(X_{raw}), f : R^{S \times T} \rightarrow R^D \quad (1)$$

X_{hc} provides a more compact representation, where $D \ll S \times T$. Traditional HAR methods M_θ map input to labels:

$$M_\theta : X_{raw/hc} \rightarrow Y, \quad (2)$$

Our innovative performance of the UI sub-task, along with our design with supervised contrastive loss objective on inter-instance relationship modeling, are our unique contributions.

B. Model Overview

CLAUDIA M contains three components (Fig. 1):

1) *Data encoder M_{θ_t}* : Extracts features from (X_{raw}, X_{hc}) , producing X_r (Eq. 3).

$$X_r = M_{\theta_t}(X_{raw}) \oplus X_{hc}, X_r \in R^{N \times d_t} \quad (3)$$

where \oplus is feature concatenation and d_t is the new feature dimension. We used a two-layer LSTM for sequence learning.

2) *Label encoder*: Encodes activities, user identities, and context settings into binary labels Y . While there are other methods beyond 0/1 label encoding [9], [25]–[27], we test with the straightforward method based on its simplicity and note these methods are supplemental to our design.

3) *Alignment module $M_{\theta_{cls}}$* : Aligns encoded data with label representations using separate linear projection followed by activation for each label type, resulting in predictions \hat{Y} , including activity \hat{Y}_A , phone placement \hat{Y}_{PP} , and UI \hat{Y}_U .

C. Objective Function

Our overall objective minimizes the total loss L , which is composed of the classification loss L_{cls} and the contrastive loss L_d , weighted by α :

$$\hat{\theta} = \arg \min_{\theta} L(M_\theta, X_{raw}, Y) \quad (4)$$

$$L = L_{cls} + \alpha L_d$$

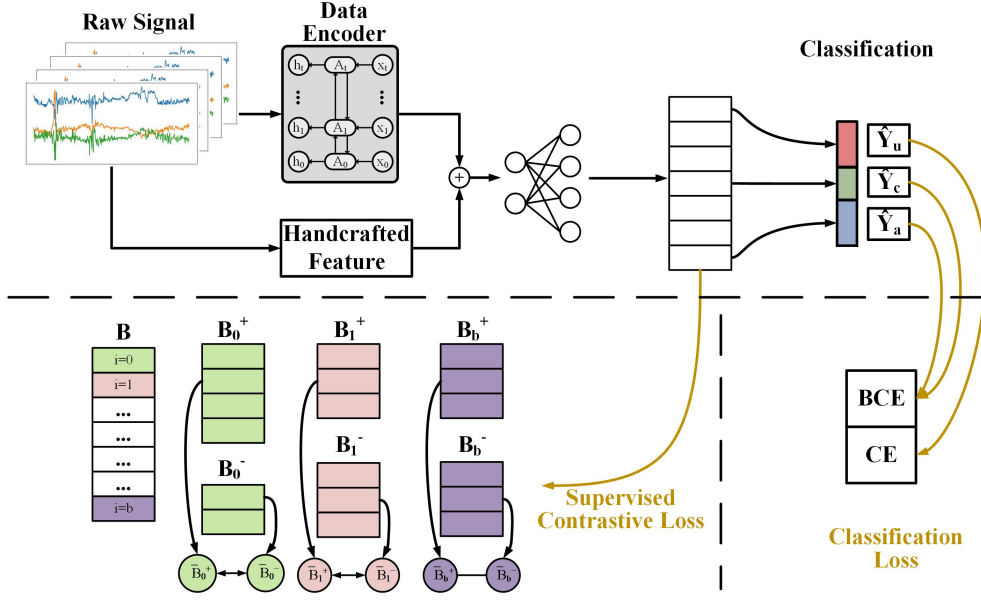


Fig. 1. Framework of CLAUDIA. The upper part illustrates the network design, comprising of data encoding and classification. The lower part depicts the objective function, consisting of two types of losses. Our technical innovation lies in two key aspects: 1) the addition of a UI sub-task, thereby broadening the scope of Context-Aware (CA) HAR to effectively include User-Aware (UA) HAR, and 2) a novel approach to improve feature quality by regularizing inter-instance relationships using supervised contrastive loss.

Classification Loss: L_{cls} includes activity L_A , phone placement L_{PP} , and UI L_U tasks, weighted by γ_1 and γ_2 :

$$L_{cls} = L_A + \gamma_1 L_{PP} + \gamma_2 L_U \quad (5)$$

As activities and general contexts may co-occur, we leverage Binary Cross Entropy (BCE) for L_A, L_{PP} (Eq. 6). In this case, class indices range from 0 to $C - 1$, where C represents the total number of classes. $\hat{y}_{n,c}$ and $y_{n,c}$ denotes the predicted and ground truth label, respectively, while $\omega_{n,c}$ represents the class weight, which is calculated based on the inverse frequency ratio to address the class imbalance issue. On the other hand, as user identities are usually mutually exclusive, we adopted general Cross Entropy (CE) for L_U (Eq. 7.)

$$L_{bce} = -\omega_{n,c} [y_{n,c} \cdot \log \sigma(\hat{y}_{n,c}) + (1 - y_{n,c}) \cdot \log(1 - \sigma(\hat{y}_{n,c}))] \quad (6)$$

$$L_{ce} = -\sum_{c=1}^C \log \frac{\exp(\hat{y}_{n,c})}{\sum_{i=1}^C \exp(\hat{y}_{n,i})} y_{n,c} \quad (7)$$

Instance-Pair Supervised Contrastive Loss: We introduce our novel design for the sample-based supervised contrastive loss in the data encoder. Given any chosen anchor x_a, y_a in any batched samples B with X_b, Y_b in a training step, we first separate the batch into positive pairs and negative pairs:

$$\begin{aligned} B &= \{(x, y) \mid x = X_b[i], y = Y_b[i], i = 0, 1, \dots\} \\ B_a^+ &= \{(x, y) \mid y \cdot y_a^T > 0, (x, y) \in B\} \\ B_a^- &= \{(x, y) \mid y \cdot y_a^T = 0, (x, y) \in B\} \\ B &= B_a^+ \cup B_a^-, \emptyset = B_a^+ \cap B_a^- \end{aligned} \quad (8)$$

Essentially, the positive pairs contain instances that share at least one label (i.e., non-zeros) with the anchor instance, while the negative pairs share no label with the anchor. Note that the size of both B_a^+, B_a^- can be volatile even under a fixed

batch size. The changing size directly impedes applying a contrastive loss similar to Contrastive Language-Image Pre-training (CLIP) [28]. To accommodate the issue, we calculated the average of both positive and negative pairs:

$$\begin{aligned} x_a^+ &= \frac{1}{|B_a^+|} \sum \{x \text{ for } (x, y) \in B_a^+\} \\ x_a^- &= \frac{1}{|B_a^-|} \sum \{x \text{ for } (x, y) \in B_a^-\} \end{aligned} \quad (9)$$

When the negative pair set is empty, we utilize a zero vector to account for its average. Another option is to sample fixed numbers k^+, k^- of the positive and negative sets. However, such a method would necessitate an undersampling and oversampling design, and would inevitably involve an extra hyperparameter search for k^+, k^- . Taking these factors into consideration, we opted for a computationally more straightforward choice as mentioned. We leave the exploration of alternate sampling methods as future work for interested readers.

Using the supervised contrastive loss function, we pull the anchor x_a towards the positive sample x_a^+ and enforce the distances between anchor x_a and the negative sample x_a^- even further away. Such an optimization direction can be reached via a cross-entropy loss:

$$\begin{aligned} p_{x_a} &= [\text{sim}(x_a, x_a^+), \text{sim}(x_a, x_a^-)] \\ q_{x_a} &= [1, 0] \\ L_d &= \frac{1}{|B|} \sum_{x \in B} CE(p_{x_a}, q_{x_a}) \end{aligned} \quad (10)$$

We re-emphasize that while prior work in the HAR domain mainly focused on classification loss (instance-label mapping regulation), we propose explicitly regularizing inter-instance relationships using a supervised contrastive loss.

TABLE I

COMPARISON OF 3 BENCHMARK CA-HAR DATASETS. THE ACCELEROMETER, GYROSCOPE, MAGNETOMETER, AND GRAVITY ARE TRI-AXIAL SENSORS. SAMPLING RATE IS ONCE PER MINUTE OR WHEN ITS VALUE CHANGED MORE THAN A THRESHOLD VALUE IF NO SAMPLING RATE WAS SPECIFIED.

Dataset	<i>WASH scripted</i>	<i>WASH unscripted</i>	<i>Extrasensory</i>
Category	<i>Scripted</i>	<i>Unscripted</i>	<i>Unscripted</i>
Instances	294,512	7,773,479	6,355,350
Features	144	139	170
Participants	107	108	60
Contexts	5	5	4
Activities	12	12	12
Common Sensors	Accelerometer(40Hz), Gyroscope(40Hz), Magnetometer(40Hz), Location, Environment Measure, Phone State		
Unique Sensors	Gravity(40Hz), Audio(46Hz)		
Context Labels	In Pocket, In Hand, In Bag		
Unique Contexts	On Table	On Table	On Table-Face Down/Up
Common Activities	Lying Down, Sitting, Walking, Sleeping, Standing, Running, Stairs-Going Down/Up		
Unique Activities	Talking on Phone, Bathroom, Jogging, Typing	Talking, Bath-Shower, Toilet, Exercising	

IV. EXPERIMENTS

We introduce real-world CA-HAR datasets in Section IV-A and describe state-of-the-art baselines in Section IV-B. We then describe our experimental setup (Sec. IV-C), evaluation metrics (Sec. IV-D), and experimental results (Sec. IV-E).

A. Datasets

CLAUDIA was evaluated on three CA-HAR datasets, encompassing scripted and unscripted data collection study designs: *WASH*¹ (Scripted and Unscripted) and *Extrasensory* [29] (Unscripted) with detailed statistics in Table I. In scripted datasets, subjects perform activities as instructed, visiting specific pre-defined target labels in a constrained order. In contrast, unscripted datasets allow participants live their lives unconstrained, then annotate their data periodically. Unlike prior studies focusing on either solely on scripted or unscripted data, we evaluated both to gain comprehensive insights. We predicted 12 activity labels common to all three datasets, ensuring a fair comparison with previous work [30]–[32].

Pre-processing: Instances with conflicting labels (e.g., sleeping while running) were removed (13% of *WASH unscripted* and 34% of *Extrasensory* data). Sensor readings were segmented with a 3-second window and 1.5-second step.

Feature extraction and selection: We extracted handcrafted features (f) identified as highly predictive in prior work [8], [9], [17], [18], [25], yielding 144, 139, and 170 features for the three datasets, respectively. Sensor signals were resampled into fixed-length (50 samples) sequences using the Fourier method [33], resulting in $x_{raw} \in \mathcal{N}^{12 \times 50}$. Features were normalized using training set statistics (i.e., $x = (f - \mu)/s$) with missing values filled with zeros.

B. Baseline HAR Models

We compared CLAUDIA to a diverse set of state-of-the-art CA-HAR models, covering graph-based (GCN [26], HGCN [27], HHGNN [9]), non-graph-based (ExtraMLP [8], LightGBM [34], CRUFT [17], GaitAuth [20], GaitIden [20]), deep learning (ExtraMLP, CRUFT, GaitAuth, GaitIden), and

machine learning (LightGBM) techniques. Due to space limits, we report the detailed descriptions of each model in our supplement materials (available in our repository).

C. Experimental Setup

We split data with training (60%), validation (20%), and testing (20%) for all users. Validation set was used to determine optimal model hyperparameters via grid search. For CLAUDIA, we used the RAdam [35] optimizer, and set batch size = 1024. We share hyperparameters in our repository.

D. Evaluation Metrics

We used Matthews Correlation Coefficient (MCC) (Eq. 11) [36] and Macro-F1 scores (Eq. 12) for imbalanced CA-HAR datasets. MCC reflects model performance across all labels, accounting for data imbalance. The Macro-F1 score, calculated as the average F1 across all labels, combines precision ($Precision = \frac{TP}{TP+FP}$) and recall ($Recall = \frac{TP}{TP+FN}$), providing a comprehensive evaluation.

$$MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}} \quad (11)$$

$$F1 = 2 * (Precision_{c_i} * Recall_{c_i}) / (Precision_{c_i} + Recall_{c_i}) \quad (12)$$

E. Experiment Results

1) *Main HAR Results:* We present comprehensive results of evaluation on HAR performance in Fig 2.

Overall Performance: CLAUDIA outperforms all baseline models across all datasets. Specifically, CLAUDIA outperforms the best baselines by an average of 11.7%, 14.2%, and 13.8% in MCC on the *WASH scripted*, *WASH unscripted*, and *Extrasensory* dataset, respectively; When measured with Macro-F1 scores, it also improved 5.4%, 7.3%, and 7.0% against the best baselines.

Noteworthy Models and Labels:

We present the detailed results for *WASH unscripted* in Table II, with other dataset results available in our repository and supplementary materials (which follow similar trends). Notably, CLAUDIA showed significant improvement in recognizing activities like *Walking* and *Bathroom/Toilet*, which

¹<https://tinyurl.com/darpaWash>

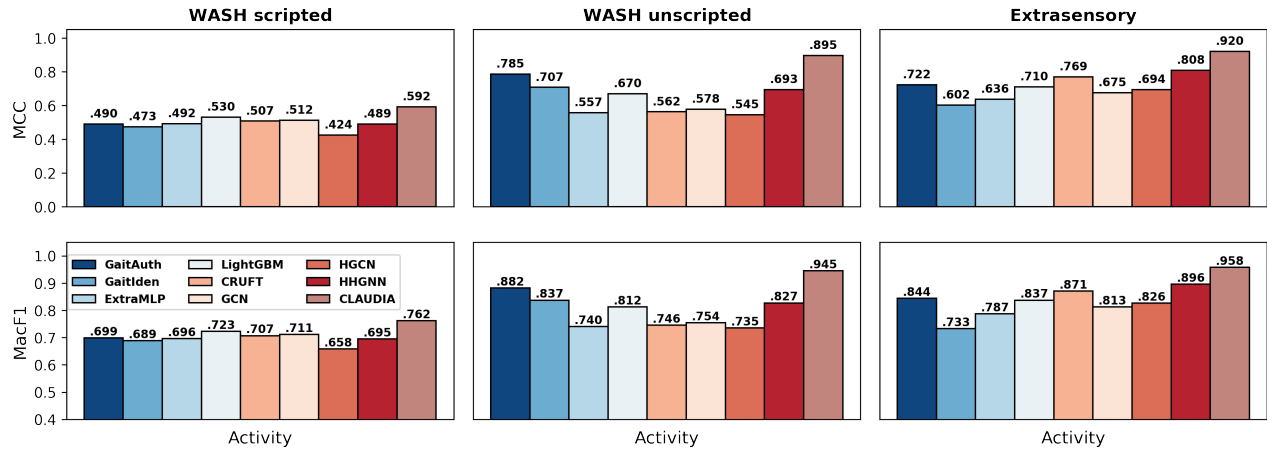


Fig. 2. Average activity recognition performance of all models across all datasets.

TABLE II
MODEL PERFORMANCE ON THE *WASH unscripted* DATASET, EACH CELL CONTAINS MCC/MACRO-F1 SCORES.

Category	UI/UA		non-GNN			GNN			Our Model	
Model	GaitAuth	GaitIden	ExtraMLP	LightGBM	CRUFT	GCN	HGNN	HHGNN	CLAUDIA	Improv.(%)
Lying Down	.842/.916	.810/.899	.831/.912	.783/.887	.820/.907	.848/.921	.729/.859	.893/.945	.975/.987	9.2/4.4
Sitting	.886/.942	.820/.908	.741/.867	.735/.864	.758/.876	.743/.868	.684/.838	.816/.907	.961/.980	8.5/4.1
Walking	.808/.899	.698/.838	.523/.731	.536/.741	.510/.721	.519/.732	.498/.716	.623/.794	.891/.945	10.3/5.1
Sleeping	.938/.969	.916/.958	.921/.960	.922/.961	.915/.957	.935/.967	.927/.963	.954/.977	.989/.994	3.6/1.8
Talking On Phone	.767/.872	.620/.779	.436/.656	.545/.730	.417/.645	.465/.675	.489/.695	.657/.806	.875/.935	14.1/7.2
Bathroom	.592/.761	.552/.740	.427/.659	.499/.704	.425/.660	.406/.641	.481/.694	.615/.780	.817/.902	32.7/15.6
Standing	.751/.864	.636/.795	.463/.686	.486/.703	.457/.681	.489/.705	.450/.677	.605/.777	.889/.943	18.3/9.1
Jogging	.836/.912	.749/.862	.520/.712	.964/.982	.551/.737	.599/.765	.409/.641	.712/.837	.922/.960	-4.3/-2.2
Running	.712/.839	.677/.821	.420/.647	.946/.973	.479/.692	.489/.695	.344/.604	.604/.769	.881/.938	-6.9/-3.6
Stairs-Going Down	.700/.833	.806/.780	.376/.626	.488/.697	.382/.636	.374/.624	.516/.721	.529/.726	.806/.897	15.1/7.7
Stairs-Going Up	.683/.822	.571/.752	.388/.634	.469/.686	.397/.645	.399/.640	.545/.740	.537/.731	.791/.888	15.9/8.0
Typing	.900/.948	.832/.912	.636/.794	.666/.814	.636/.793	.672/.817	.467/.670	.770/.876	.949/.974	5.5/2.7

involve periodic patterns and varied user movements, often with weaker labeling due to privacy concerns—aligning well with CLAUDIA’s strengths. Additionally, our model was also impressive on *Stairs-Going Up/Down*, which are challenging to distinguish [37].

Although LightGBM had an edge on *Jogging* and *Running* (inherently similar in nature and could confound each other), CLAUDIA remains the best among unified models.

Comparative Analysis across Datasets: Generally speaking, most models exhibited higher performance on the *WASH Unscripted* and *Extrasensory* datasets compared to the *WASH Scripted* dataset. This could be attributed to the *WASH Scripted* having the least number of instances and a lower instance-to-user ratio, as reported in Table I.

TABLE III
ABLATION STUDY COMPARES OUR CLAUDIA MODEL WITH ITS VARIANTS FOR THE ACTIVITY RECOGNITION TASK.

Model	<i>WASH scripted</i>	<i>WASH unscripted</i>	<i>Extrasensory</i>
CLAUDIA	.592/.762	.895/.945	.920/.958
w/o UI	.487/.696	.871/.932	.909/.952
w/o CL	.574/.751	.873/.933	.899/.947
w/o TS	.478/.688	.764/.870	.843/.915

2) **Ablation on CLAUDIA:** We conducted an ablation study on CLAUDIA to assess the contributions of each modules. Table III presents the corresponding MCC and Macro-F1 scores for the full model and its variants: 1) w/o UI:

CLAUDIA without the user identification module ($\gamma_2 = 0$); 2) w/o Contrastive Loss (CL): CLAUDIA without instance-level contrastive regularization ($\alpha = 0$); 3) w/o Time-Series (TS) encoding: Omits the sequence learning model, using only handcrafted features. The results demonstrate that the performance of CLAUDIA deteriorates when any of these key components is removed, underscoring the contribution of each module to its performance. Notably, the time-series module has the most significant impact especially on unscripted datasets. Integrating them enables CLAUDIA to achieve robust results across both scripted and unscripted datasets.

V. ANALYSIS

A. Context Identification

CLAUDIA achieves MCC/Macro-F1 scores 0.813/0.900, 0.946/0.973, and 0.974/0.987 on the *WASH scripted*, *WASH unscripted*, and *Extrasensory* datasets respectively. Detailed confusion matrices are provided in the supplementary material.

B. User Identification

CLAUDIA demonstrated near-perfect UI performance, achieving MCC and Macro-F1 scores of 0.995, 0.999, and 1.000 on the three datasets. It significantly outperformed *GaitIden* (with .959/.959, .962/.982, and .965/.963 MCC/Macro-F1 scores), especially on *WASH scripted*, indicating CLAUDIA’s effectiveness even when user distinctions were less clear.

VI. CONCLUSION AND FUTURE WORK

In this work, we improve Context-Aware Human Activity Recognition (CA-HAR) task performance by 1) expanding the contextual scope to include user identification and in learning user-specific patterns, and 2) introducing a novel inter-instance relationship objective that supplements existing instance-label modeling targets by enforcing user-invariant activity embeddings. We propose CLAUDIA, a framework that leverages both handcrafted feature and raw time-series data, innovatively performs user identification as an auxiliary sub-task to the CA-HAR framework. The resulting framework, CLAUDIA, significantly improves HAR performance across multiple real-world CA-HAR datasets with average enhancements ranging from 11.7% to 14.2% in Matthew's Correlation Coefficient (MCC) and 5.4% to 7.3% in Macro F1 score. Furthermore, we provide extensive qualitative and quantitative analysis results. Since we only experimented with recurrent-based networks for time series modeling, we leave the exploration of other architectures, such as Transformers, for future work.

REFERENCES

- [1] X. Zhou, W. Liang, I. Kevin, K. Wang, H. Wang, L. T. Yang, and Q. Jin, "Deep-learning-enhanced human activity rec. for internet of healthcare things," *IEEE Internet of Things J.*, vol. 7, no. 7, pp. 6429–6438, 2020.
- [2] S. Wan, L. Qi, X. Xu, C. Tong, and Z. Gu, "Deep learning models for real-time human activity recognition with smartphones," *Mobile Networks and Applications*, vol. 25, pp. 743–755, 2020.
- [3] K.-H. Chen, Y.-W. Hsu, J.-J. Yang, and F.-S. Jaw, "Evaluating the specifications of built-in accelerometers in smartphones on fall detection performance," *Inst. Sci. & Tech.*, vol. 46, no. 2, pp. 194–206, 2018.
- [4] J. Lindqvist and J. Hong, "Undistracted driving: A mobile phone that doesn't distract," in *Proc. Hotmobile*. ACM, 2011, pp. 70–75.
- [5] O. D. Lara and M. A. Labrador, "A survey on human activity recognition using wearable sensors," *IEEE communications surveys & tutorials*, vol. 15, no. 3, pp. 1192–1209, 2012.
- [6] C. A. Ronao and S.-B. Cho, "Human activity recognition with smartphone sensors using deep learning neural networks," *Expert systems with applications*, vol. 59, pp. 235–244, 2016.
- [7] M. M. Hassan, M. Z. Uddin, A. Mohamed, and A. Almogren, "A robust human activity recognition system using smartphone sensors and deep learning," *Future Gen. Computer Systems*, vol. 81, pp. 307–313, 2018.
- [8] Y. Vaizman, N. Weibel, and G. Lanckriet, "Context recognition in-the-wild: Unified model for multi-modal sensors and multi-label classification," *ACM IMWUT*, vol. 1, no. 4, pp. 1–22, 2018.
- [9] W. Ge, G. Mou, E. O. Agu, and K. Lee, "Heterogeneous hyper-graph neural networks for context-aware human activity recognition," in *PerCom*, IEEE. IEEE, 2023, pp. 350–354.
- [10] L. Cao, Y. Wang, B. Zhang, Q. Jin, and A. V. Vasilakos, "Gchar: An efficient group-based context-aware human activity recognition on smartphone," *J. Parallel and Distr. Comp.*, vol. 118, pp. 67–80, 2018.
- [11] N. D. Lane, Y. Xu, H. Lu, S. Hu, T. Choudhury, A. T. Campbell, and F. Zhao, "Enabling large-scale human activity inference on smartphones using community similarity networks (csn)," in *UbiComp*. ACM, 2011, pp. 355–364.
- [12] F. C. Bull, S. S. Al-Ansari, S. Biddle, K. Borodulin, M. P. Buman, G. Cardon, C. Carty, J.-P. Chaput, S. Chastin, R. Chou *et al.*, "World health organization 2020 guidelines on physical activity and sedentary behaviour," *British J. sports med.*, vol. 54, no. 24, pp. 1451–1462, 2020.
- [13] X. Gao, H. Luo, Q. Wang, F. Zhao, L. Ye, and Y. Zhang, "A human activity recognition algorithm based on stacking autoencoder and lightgbm," *Sensors*, vol. 19, no. 4, p. 947, 2019.
- [14] S. Münzner, P. Schmidt, A. Reiss, M. Hanselmann, R. Stiefelhagen, and R. Dürichen, "Cnn-based sensor fusion techniques for multimodal human activity recognition," in *Proceedings of the 2017 ACM international symposium on wearable computers*. ACM, 2017, pp. 158–165.
- [15] S. Mekruksavanich and A. Jitpattanakul, "Lstm networks using smartphone data for sensor-based human activity recognition in smart homes," *Sensors*, vol. 21, no. 5, p. 1636, 2021.
- [16] L. Shi, Y. Zhang, J. Cheng, and H. Lu, "Skeleton-based action recognition with multi-stream adaptive graph convolutional networks," *IEEE Transactions on Image Processing*, vol. 29, pp. 9532–9545, 2020.
- [17] W. Ge and E. Agu, "Cruft: Context rec. under uncertainty using fusion and temporal learning," in *Proc. ICMLA*. IEEE, 2020, pp. 747–752.
- [18] W. Ge and E. O. Agu, "Qcruft: Quaternion context recog. under uncertainty using fusion & temporal learning," in *Proc. ICSC*, IEEE. IEEE, 2022, pp. 41–50.
- [19] R. Mondal, D. Mukherjee, P. K. Singh, V. Bhateja, and R. Sarkar, "A new framework for smartphone sensor-based human activity recognition using graph neural network," *IEEE Sensors Journal*, vol. 21, no. 10, pp. 11 461–11 468, 2020.
- [20] Q. Zou, Y. Wang, Q. Wang, Y. Zhao, and Q. Li, "Deep learning-based gait recognition using smartphones in the wild," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 3197–3212, 2020.
- [21] S. Mekruksavanich and A. Jitpattanakul, "Biometric user identification based on human activity rec. using wearable sensors: An experiment using deep learning models," *Electronics*, vol. 10, no. 3, p. 308, 2021.
- [22] C. I. Tang, I. Perez-Pozuelo, D. Spathis, and C. Mascolo, "Exploring contrastive learning in human activity recognition for healthcare," *arXiv preprint arXiv:2011.11542*, 2020.
- [23] J. Wang, T. Zhu, J. Gan, L. L. Chen, H. Ning, and Y. Wan, "Sensor data augmentation by resampling in contrastive learning for human activity rec." *IEEE Sensors Journal*, vol. 22, no. 23, pp. 22 994–23 008, 2022.
- [24] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, "Supervised contrastive learning," *NeurIPS*, vol. 33, pp. 18 661–18 673, 2020.
- [25] W. Ge, G. Mou, E. O. Agu, and K. Lee, "Deep heterogeneous contrastive hyper-graph learning for in-the-wild context-aware human activity recognition," *Proc. ACM IMWUT*, vol. 7, no. 4, pp. 1–23, 2024.
- [26] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *Proc ICLR 2017*. OpenReview.net, 2017. [Online]. Available: <https://openreview.net/forum?id=SJU4ayYgl>
- [27] S. Bai, F. Zhang, and P. H. Torr, "Hypergraph convolution and hyper-graph attention," *Pattern Recognition*, vol. 110, p. 107637, 2021.
- [28] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, "Learning transferable visual models from natural language supervision," in *International conference on machine learning*, PMLR. PMLR, 2021, pp. 8748–8763.
- [29] Y. Vaizman, K. Ellis, and G. Lanckriet, "Recognizing detailed human context in the wild from smartphones and smartwatches," *IEEE Pervasive Computing*, vol. 16, no. 4, pp. 62–74, 2017.
- [30] Y. Asim, M. A. Azam, M. Ehatisham-ul Haq, U. Naeem, and A. Khalid, "Context-aware human activity recognition (cahar) in-the-wild using smartphone accelerometer," *IEEE Sensors Journal*, vol. 20, no. 8, pp. 4361–4371, 2020.
- [31] T. Gong, J. Jeong, T. Kim, Y. Kim, J. Shin, and S.-J. Lee, "Note: Robust continual test-time adaptation against temporal correlation," *Adv. Neural Information Processing Systems*, vol. 35, pp. 27 253–27 266, 2022.
- [32] F. Cruciani, A. Vafeiadis, C. Nugent, I. Cleland, P. McCullagh, K. Votis, D. Giakoumis, D. Tzovaras, L. Chen, and R. Hamzaoui, "Feature learning for human activity recognition using convolutional neural networks: A case study for inertial measurement unit and audio data," *CCF Trans. Pervasive Computing and Interaction*, vol. 2, no. 1, pp. 18–32, 2020.
- [33] J. W. Cooley and J. W. Tukey, "An algorithm for machine calc. of comp. fourier series," *Math. of Comp.*, vol. 19, no. 90, pp. 297–301, 1965.
- [34] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, "Lightgbm: A highly efficient gradient boosting decision tree," *Adv. NIPS*, vol. 30, 2017.
- [35] L. Liu, H. Jiang, P. He, W. Chen, X. Liu, J. Gao, and J. Han, "On the variance of the adaptive learning rate and beyond," *arXiv preprint arXiv:1908.03265*, 2019.
- [36] D. Chicco and G. Jurman, "The advantages of the matthews correlation coefficient (mcc) over f1 score and accuracy in binary classification evaluation," *BMC genomics*, vol. 21, pp. 1–13, 2020.
- [37] N. Ravi, N. Dandekar, P. Mysore, and M. L. Littman, "Activity recognition from accelerometer data," in *AAAI*, vol. 5, 2005, pp. 1541–1546.