Is Querying Users Acceptable for Human Activity Recognition Based on Active Learning?

1st Riccardo Presotto
Dept. of Computer Science
University of Milan
Milan, Italy
riccardo.presotto@unimi.it

2nd Gabriele Civitarese
Dept. of Computer Science
University of Milan
Milan, Italy
gabriele.civitarese@unimi.it

3rd Claudio Bettini
Dept. of Computer Science
University of Milan
Milan, Italy
claudio.bettini@unimi.it

***Abstract*—Labeled data scarcity is one of the major open problems of sensor-based Human Activity Recognition (HAR). To mitigate this issue, several research groups proposed solutions based on Active Learning (AL), where the user is explicitly asked to provide feedback about the performed activity when the classifier's confidence is low. Despite existing methods trigger a limited number of queries that decreases over time, they do not take into account user's context. Indeed, a subject may not wish to be interrupted when receiving the query's notification. Delaying the query may be critical, since it is known that the feedback's quality decreases with the length of the delay. In this work, we aim at answering the following question: is real-time AL-based HAR practical in real-world scenarios? By using a novel evaluation methodology based on the Wizard of Oz approach, we performed a user-based study with 30 subjects performing physical activities and using earphones and a smartwatch as vocal and touch interfaces. Our results evaluate the impact of the AL strategy, the user's context, and the interaction modality on the effectiveness of the AL approach.**

***Index Terms*—human activity recognition, active learning, context-awareness, interruptibility**

I. Introduction

The continuous evolution of sensing devices and AI methods makes it possible to efficiently deploy sensor-based Human Activity Recognition (HAR) models to recognize in real-time human activities [1]. The most accurate HAR models are trained in a fully supervised fashion [2]. However, despite their effectiveness, supervised methods require a significant amount of labeled data. Data annotation is often prohibitive due to its costs and its obtrusiveness [3].

To mitigate this problem, semi-supervised learning methods for HAR have been proposed [4]. These approaches require a small labeled dataset to initialize the model, which is then improved by taking advantage of the large unlabeled data stream. Among the proposed semi-supervised approaches for HAR, Active Learning (AL) is particularly effective [5]–[8]. Specifically, when the classifier is uncertain about the activity currently performed by the user, the system explicitly asks her to provide a feedback. The annotated data samples obtained thanks to AL are used to improve the activity model.

Even though AL-based HAR methods achieve high recognition rates with a low number of queries, existing methods have been evaluated on public datasets assuming that users are constantly available to provide a feedback. However, in real-world scenarios, the user's context significantly impacts users' predisposition to answering AL queries. Notifications prompted at inappropriate times may lead to stress and frustration, further increasing the feedback error rate and disaffection with the system [9]. As an example, consider the *sitting* activity. If the user is sitting alone in a train station, it is very likely that an AL notification (e.g., *"are you sitting or standing?"*) would not significantly bother them. However, if the same notification is received while the user is sitting on the train and talking with other people, they may not be predisposed to provide a feedback.

Moreover, since short-term memory is more reliable than long-term memory, the quality of the feedback decreases with the time elapsed since the activity was actually performed [10]. Hence, AL notifications for HAR should be prompted as soon as possible to obtain reliable feedback.

In this paper, we aim at answering the following research question: *is real-time AL-based HAR practical in real-world scenarios?* To answer this question, we propose a novel evaluation methodology to estimate the acceptability of AL queries in different context scenarios. Our methodology is based on the Wizard of Oz (WOz) approach: while the user believes to perform activities wearing an actual AL-based HAR system, a hidden human operator (i.e., the Wizard) actually triggers AL queries at specific instants to evaluate interruptibility in different context conditions. We implemented a prototype of the system and performed a preliminary user-based study with 30 volunteers. In our experiments, we considered a well-known HAR setting in the literature, based on physical activities and mobile/wearable devices. Our results suggest that prompting AL queries at appropriate times may be acceptable, and that context information is crucial to obtain reliable feedback.

The contributions of this paper are the following:

- We propose a novel user-based evaluation methodology to study the acceptability of AL-based HAR.
- We implemented a prototype of the system for HAR based on mobile/wearable devices.
- We performed a preliminary study with 30 volunteers, showing that AL may be actually practical in real-world scenarios by taking context into account.

II. Related work

A. Interruptibility when receiving notifications

Several research works studied the general interruptibility of users when receiving notifications on mobile/wearable devices. The existing studies show that environmental and context conditions are fundamental aspects that heavily influence the user's interruptibility level [11]. Indeed, context data can be used to train a classifier in charge of deciding whether to prompt a notification or postpone it to a more acceptable context scenario [12], [13]. Other works also suggest that breakpoints between activities are another important factor that should be considered to prompt notifications [14], [15].

However, existing studies did not consider the acceptability evaluation for AL notifications in HAR, which have specific characteristics (e.g., they can not be postponed too long due to the challenge of remembering the activity being performed). Moreover, this work also considers higher-level context aspects, like the cognitive commitment of the users in other tasks and their involvement in social interactions.

B. Wizard of Oz

The Wizard of Oz (WOz) [16] is a widely adopted evaluation method where subjects are told that they are interacting with a working user interface, even though they are not. Indeed, the interaction from the interface to the user is actually mediated by a hidden human operator, called Wizard. Thanks to WOz, it is possible to systematically evaluate the actual users' experience when interacting with a prototype, before investing time and effort in developing the complete system.

WOz-based strategies have been frequently used in the pervasive computing domain. For instance, several works proposed to use WOz to evaluate the acceptability of smart-home voice interfaces for elderly subjects [17]. As another example, WOz has been adopted to design smart interfaces for blind people [18], [19]. In the context of HAR, WOz was proposed to study the impact of the automatic recognition of physical activities on productivity, by proposing interfaces prompting notifications when users have some spare time during their working hours [20].

To the best of our knowledge, there is no work in the literature that proposed WOz to study acceptability aspects of active learning in HAR.

III. Our Evaluation Framework

In this work, we are interested in evaluating the acceptability of Active Learning (AL) in HAR systems considering different context conditions. Even though several AL-based methods have been proposed in the literature, using them to automatically evaluate user experience is challenging, since it is not possible to know in advance when queries are actually triggered and in which contexts. Moreover, we are also interested in context conditions that can not be easily captured by currently available sensing solutions (e.g., involvement in social interactions). For these reasons, we propose a framework based on Wizard of Oz (WOz). In the following, we describe our approach and we present the software implementation of a prototype of the envisioned system.

A. The proposed WOz-based approach

In our evaluation methodology, the subject is told that they have to perform a set of activities while wearing mobile/wearable devices (a smartwatch and a smartphone in our experiments) that automatically recognize her activities and that periodically prompt AL notifications. The subject does not know that there is no an actual HAR recognition model running on the devices, and AL notifications are triggered by a hidden human operator (i.e., the Wizard). The wizard secretly observes the subject while executing activities and sends AL notifications to the subject's devices in specific context conditions. Figure 3a shows an example of an AL query prompted to the subject's smartwatch. Our system asks the subject to choose between two activities. If the subject is not actually performing any of the two activities (this simulates a wrong prediction of the HAR system), they can select the option "other".

Our method quantifies interruptibility considering different metrics:

- **Annoyance**: besides explicitly asking the subject to select the performed activity, we also ask her to quantify how much answering that query bothered her (see Figure 3b).
- **Reaction time**: it is computed as the time elapsed since the AL notification is prompted on the device and the actual answer of the subject. Intuitively, the longer the reaction time and the less the subject was interruptible.
- **Ignored notifications**: AL notifications automatically disappear from subjects' devices after a Time-To-Live (TTL) decided by the Wizard (e.g., 30 seconds).

Our study aims at correlating interruptibility with several subject's context variables: the activity being performed when the AL notification is received, the semantic location, the weather condition, the time of day (e.g., morning, afternoon), the presence of other persons, possible social interactions, and whether the subject is cognitively involved in a task (e.g., working at PC). Moreover, we experiment with two different AL strategies:

- **Instant-AL**: the Wizard sends AL notifications during the activity execution. This strategy simulates an AL approach that prompts a notification as soon as it is required by the HAR model.
- **TransitionAware-AL**: the Wizard sends AL notifications in the transition between two activities. This approach was suggested in the literature to mitigate annoyance [14].

Finally, we also consider two alternative user interfaces: touch and vocal. Indeed, user experience is strictly connected with the specific interface [21].

B. Software components

In this section, we introduce the main software components that we implemented for our WoZ-based system. As Figure 1 shows, our architecture consists of three modules: 1) the

WizardAPP: the application used by the Wizard to generate AL queries, transmit them to the subjects and obtain the results, 2) the *AL-HAR* app installed on the mobile/wearable devices of the subjects, which provides to the user a UI to visualize and answer queries, and 3) the *WOz-MW* middleware that enables the communication between the *WizardAPP* and the *AL-HAR* applications. In the following, we illustrate the implementation details of those components and how they interact with each other.

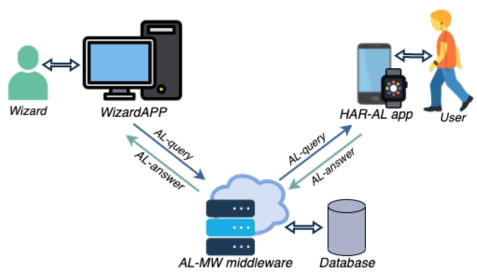

Fig. 1: Implemented WoZ-based architecture

1) WizardAPP: We implemented the *WizardAPP* as a web application. The wizard secretly monitors the subject and uses the *WizardAPP* to periodically generate and transmit AL queries to the *HAR-AL* app [1]. Besides setting the parameters for the AL query (i.e., the two activities appearing in the query and the ground truth), the Wizard also has to annotate the subject's current high-level context. Figure 2 shows a screenshot of the *WizardAPP*.

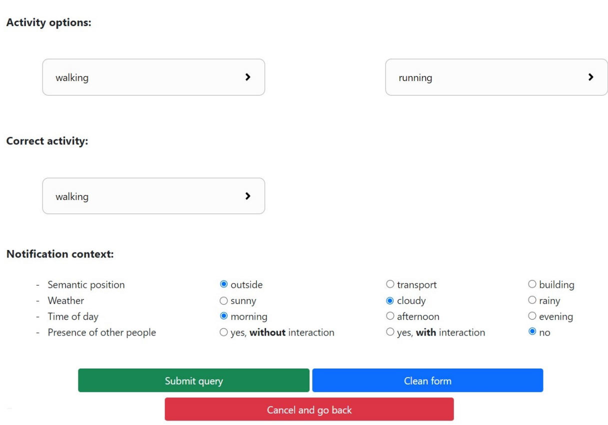

Fig. 2: A screenshot of WizardAPP: AL query generation

Note that, before starting an experimental session, the Wizard also has to select the active learning strategy and the TTL.

2) AL-HAR: Our prototype assumes that the subject carries a smartphone in the front pants pocket and a smartwatch on the wrist. We propose two types of interfaces to provide feedback to AL queries: a touch interface on the smartwatch and a vocal interface thanks to earphones connected through the smartphone. The *AL-HAR* application is subdivided into two modules: an Android application installed on the smartphone and a WearOS application installed on the smartwatch. The smartphone module is in charge of handling the network communication with the *AL-MW* middleware. Moreover, it implements the AL vocal interface. This interface uses a voice synthesizer to ask AL queries, while the feedback is processed by a *text2speech* module and then semantically interpreted to provide the feedback. The smartwatch is constantly communicating with the smartphone, and it only implements the touch interface (see screenshots in Figure 3). The smartwatch notifications are associated with a short vibration to alert the subject.

[1]In Section IV-A we describe how secret monitoring is performed in our experimental setup.

(a) Activity query interface

(b) Annoyance query interface

Fig. 3: Touch user interfaces

3) AL-MW: The *AL-MW* middleware simply connects the *WizardAPP* and the *AL-HAR* applications by using the MQTT protocol. Moreover, the middleware is also in charge of storing in a database the collected experimental data.

IV. Experimental Evaluation

We performed preliminary user-based studies with the framework proposed in Section III to study how contexts impact users' willingness to provide AL feedback. In the following, we introduce our experimental setup and the major insights obtained from our experiments.

A. Experimental setup

We recruited 30 volunteers aged between 18 and 28 years old (33% of women and 67% of men). The *AL-HAR* application presented in the previous section runs on mobile/wearable devices worn by the volunteers. In particular, we provided to each volunteer a *Nexus 5x* smartphone, a *LG G-watch R* smartwatch, and the Xiaomi Mi Piston earphones.

The Wizard (i.e., a member of our research group) remotely and secretly monitored the volunteers from the research lab, running the *WizardAPP* on a laptop. In order to enable remote monitoring, each volunteer was also provided with an action camera on the chest (we experimented both a *Insta360 go 2* and a *GoPro Hero 8*). Figure 4 summarizes the devices worn by each volunteer during the experiments.

The video stream from the action camera was remotely observed by the Wizard in real-time to generate AL queries at appropriate times. Figure 5 shows a screenshot of the video stream observed by the Wizard.

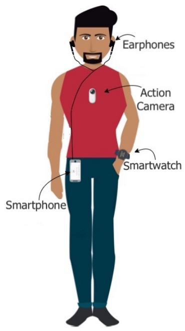

Fig. 4: The devices worn by volunteers during the experiments

Fig. 5: A screenshot of the video stream observed by the *Wizard*

B. Experimental Methodology

We designed four different scenarios for our experiments. Each scenario includes ≈ 20 minutes of outdoor and indoor activities nearby our University campus. In order to ensure that volunteers carried out activities in a realistic fashion, we did not instruct them on each specific activity to carry out, but we presented them the scenario as a set of objectives to accomplish (e.g., "reach the park near the university by taking a bus, relax for a while, and then take a bicycle to go back to the starting point"). Our scenarios sometimes included the presence of some operators (i.e., members of our lab) that interacted with the volunteers. We introduced this aspect to obtain contexts with realistic social interactions. Moreover, in order to include cognitive involvement as a context of interest, some scenarios also invited the subject to engage in some logic-based games (e.g., Sudoku). Overall, our scenarios include the following activities: *Running, Sitting, Standing, Going upstairs, Going downstairs, Cycling, Lying, Sitting/Standing on transport, and Stretching*. These activities are representative of well-known public datasets of sensor-based HAR with mobile/wearable devices [2]. During the experiments, the volunteers were free to answer AL queries by using either the touch or the vocal interface. Each volunteer performed only a single scenario.

Before starting an experiment, we instructed each volunteer on AL-based systems for HAR and their potential advantages in their daily life. Then, we presented them our AL user interfaces and we tricked them into thinking they were wearing a fully functional HAR system. In order to collect meaningful data, we also informed the volunteers that the action camera has the only objective of monitoring the experiment and that it is not actually part of the HAR system. Hence, we suggested them not to consider it as a discomfort when answering queries regarding annoyance. Moreover, we also informed them that real-world AL systems would actually trigger a low number of queries, while the system in the experiment was set to prompt a large number of queries for data analytic purposes.

C. Results

1) Impact of the AL strategy: Figures 6 show the impact of the AL strategy on the volunteers' annoyance. We observe that, consistently with the literature, *TransitionAware-AL* is associated with a reduced annoyance compared to *Instant-AL*.

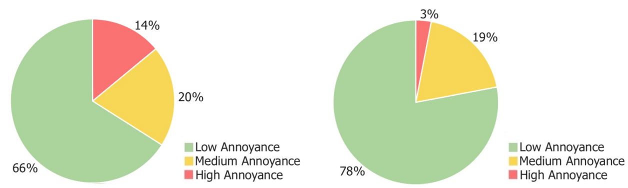

(a) Instant-AL (b) TransitionAware-AL

Fig. 6: Impact of the AL strategy on notifications' annoyance

Hence, our volunteers were more inclined in providing AL feedback during the transitions between activities, rather than randomly in the middle of activity execution. The same data analysis is reported at the activity granularity in Figure 7.

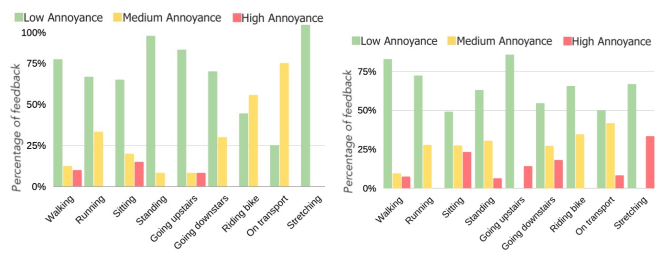

(a) TransitionAware-AL (b) Instant-AL

Fig. 7: Impact of the AL strategy on the annoyance for each activity

As expected, we observed that most of the activities requiring intense physical involvement (e.g., *running*, *taking the stairs*, and *stretching*) are related to a reduced interruptibility considering the *Instant-AL* strategy. Nonetheless, considering *cycling*, we noted a higher inclination to provide feedback with *Instant-AL*. This is due to the fact that, for safety reasons, we required the volunteers to use the vocal interface when cycling, with a positive effect on the perceived annoyance.

Surprisingly, we also observe that Instant-AL is associated with high annoyance even considering static activities (e.g., *sitting* and *standing*). This is possibly due to the fact that the volunteers could be standing and talking with another person, or sitting while performing cognitive tasks.

Our results also show that the *Transition-AL* strategy may be sometimes perceived negatively. For instance, when receiving AL queries at the end of the *cycling* activity, volunteers were less inclined in providing feedback since during the transition they were involved in carefully parking the bicycle. A similar phenomenon was observed with ”*sitting/standing on public transport*”. Indeed, the transition may involve the user getting off a crowded bus, which is a not suitable moment for providing feedback to the system. In general, the ability to detect context during transitions would be extremely useful for AL systems to determine the level of interruptibility.

The *Transition-AL* strategy also has other disadvantages. Figure 8 shows the difference in reaction time between the different AL strategies.

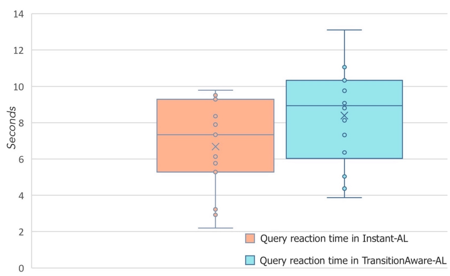

Fig. 8: Reaction time: Instant-AL vs Transition-AL

We observed that *Transition-AL* is associated with a higher reaction time. This is probably due to the fact that the volunteers had to reason more about the correct activity to choose (i.e., the one performed immediately before the query). For the same reason, Figure 9 shows that *Transition-AL* is associated with a higher number of mistakes when providing the feedback.

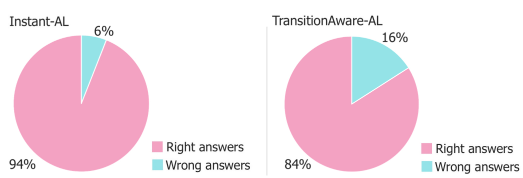

Fig. 9: Distribution of correct and wrong AL feedback

On the other hand, the reduced reaction time of *Instant-AL* may be due to the fact that our volunteers wanted to continue the current activity as soon as possible.

2) Impact of high-level context: In the following, we show the impact of high-level context on the perceived annoyance. As a representative example, Figure 10 compares the interruptibility level of the volunteers performing *walking* with or without the interaction with other persons. We observe that AL queries received during this activity are positively perceived by the volunteers, independently of the considered social context. This is probably due to the fact that this activity requires limited physical efforts, and social interactions did not discourage the volunteers from interacting with the system. However, this may also change depending on the specific type of interaction. Figure 11 shows the same result for the *taking the stairs* activity. In this case, the social context seems to significantly affect interruptibility. Possibly because this activity is associated with higher physical intensity, the presence of other persons significantly reduced the volunteers’

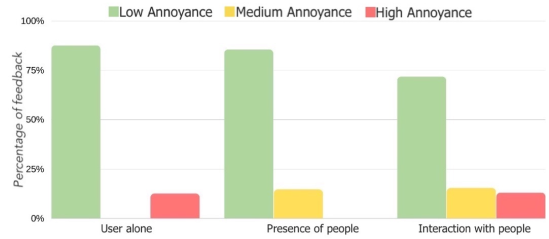

Fig. 10: *Walking*: the impact of the social context on notifications’ annoyance

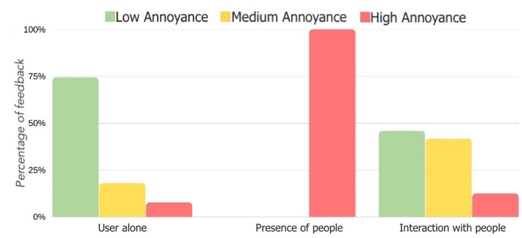

Fig. 11: *Taking stairs*: the impact of the social context on notifications’ annoyance

inclination to answer AL queries. Note that the slight reduction of annoyance associated with social interactions is due to the fact that, in our experiments, the volunteers performed this activity by interacting with acquaintances or friends. Hence, we suppose that this type of interaction slightly reduced the negative annoyance perceived by the volunteers. Finally, Figure 12 shows how the cognitive involvement of the volunteers (e.g., solving the Sudoku) affected their availability to answer.

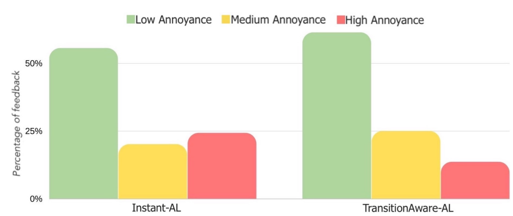

Fig. 12: Cognitive involvement and different AL strategies

As expected, we observed that the volunteers were less inclined to provide feedback during the cognitive task (*Instant-AL*) rather than at the end of it (*TransitionAware-AL*).

3) Touch vs Vocal Interface: Figure 13 shows, for each activity, the distribution between the interfaces used by the volunteers. For each activity, we observe a slight preference towards the touch interface. This is likely due to the fact that, in the considered outdoor scenarios, the vocal interface may lead to social embarrassment. This is particularly evident in *sitting/standing on public transport*, where the volunteers only used the touch interface. Nonetheless, the choice of the interface is also heavily influenced by the high-level user’s context. For instance, Figure 14 shows that, when cognitively involved in other tasks, volunteers preferred to use the vocal interface since it was perceived as less distracting.

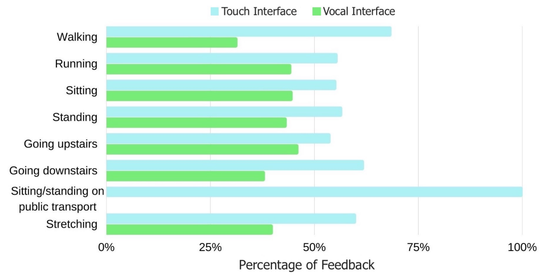

Fig. 13: Touch vs Vocal Interface

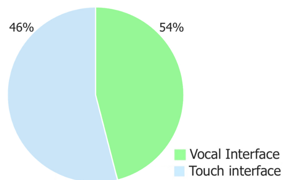

Fig. 14: Cognitive involvement and different user interfaces

V. Conclusion and Future Work

In this work, we proposed an evaluation framework based on the WOz approach to evaluate the acceptability of AL-based methods for HAR in real-world scenarios. Our results show that a limited number of queries conveyed by an AL system with appropriate interfaces at appropriate times may indeed be acceptable. However, context information should be carefully considered in order to maximize the likelihood of receiving accurate feedback.

Despite the promising results, this work has several limitations that we would like to address in future work. First, our experimental evaluation is limited in terms of considered activities and contexts. While WOz makes it possible to precisely control when to prompt queries, it significantly limits the evaluation in more realistic scenarios. Since in our study the users were observed for short periods in a fun study setup, and the annoyance level may be underestimated with respect to the one that might be observed during daily living. Moreover, the automatic recognition of high-level contexts (e.g., social interaction) is still an open challenge and part of our plans for future work.

Besides acceptability in terms of context conditions, it is also important to consider that the adoption of AL should be incentivised with rewards for the users. Besides making users aware of the benefits of AL in terms of model personalization, economic rewards and/or gamification strategies may be considered [22].

References

[1] S. Rani, H. Babbar, S. Coleman, A. Singh, and H. M. Aljahdali, “An efficient and lightweight deep learning model for human activity recognition using smartphones,” *Sensors*, vol. 21, no. 11, p. 3845, 2021.

[2] K. Chen, D. Zhang, L. Yao, B. Guo, Z. Yu, and Y. Liu, “Deep learning for sensor-based human activity recognition: Overview, challenges, and opportunities,” *ACM Comput. Surv.*, vol. 54, no. 4, pp. 1–40, 2021.

[3] D. J. Cook, K. D. Feuz, and N. C. Krishnan, “Transfer learning for activity recognition: A survey,” *Knowl. Inf. Sys.*, vol. 36, no. 3, pp. 537–556, 2013.

[4] Z. S. Abdallah, M. M. Gaber, B. Srinivasan, and S. Krishnaswamy, “Activity recognition with evolving data streams: A review,” *ACM Computing Surveys (CSUR)*, vol. 51, no. 4, p. 71, 2018.

[5] R. Smith and M. Dragone, “A dialogue-based interface for active learning of activities of daily living,” in *27th International Conference on Intelligent User Interfaces*, 2022, pp. 820–831.

[6] H. S. Hossain, M. A. A. H. Khan, and N. Roy, “Active learning enabled activity recognition,” *Perv. Mob. Comput.*, vol. 38, pp. 312–330, 2017.

[7] C. Bettini, G. Civitarese, and R. Presotto, “Caviar: Context-driven active and incremental activity recognition,” *Knowledge-Based Systems*, vol. 196, p. 105816, 2020.

[8] R. Presotto, G. Civitarese, and C. Bettini, “Semi-supervised and personalized federated activity recognition based on active learning and label propagation,” *Personal and Ubiquitous Computing*, vol. 26, no. 5, pp. 1281–1298, 2022.

[9] K. Kushlev and E. W. Dunn, “Checking email less frequently reduces stress,” *Computers in Human Behavior*, vol. 43, pp. 220–228, 2015.

[10] T. Miu, P. Missier, and T. Plötz, “Bootstrapping personalised human activity recognition models using online active learning,” in *2015 IEEE International Conference on Computer and Information Technology; Ubiquitous Computing and Communications; Dependable, Autonomic and Secure Computing; Pervasive Intelligence and Computing*. IEEE, 2015, pp. 1138–1147.

[11] C. Anderson, I. Hübener, A.-K. Seipp, S. Ohly, K. David, and V. Pejovic, “A survey of attention management systems in ubiquitous computing environments,” *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 2, no. 2, pp. 1–27, 2018.

[12] A. Mehrotra, J. Vermeulen, V. Pejovic, and M. Musolesi, “Ask, but don’t interrupt: the case for interruptibility-aware mobile experience sampling,” in *Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers*, 2015, pp. 723–732.

[13] C. Forlivesi, U. G. Acer, M. Van Den Broeck, and F. Kawsar, “Mindful interruptions: a lightweight system for managing interruptibility on wearables,” in *Proceedings of the 4th ACM Workshop on Wearable Systems and Applications*, 2018, pp. 27–32.

[14] T. Okoshi, J. Ramos, H. Nozaki, J. Nakazawa, A. K. Dey, and H. Tokuda, “Attelia: Reducing user’s cognitive load due to interruptive notifications on smart phones,” in *2015 IEEE International Conference on Pervasive Computing and Communications (PerCom)*. IEEE, 2015, pp. 96–104.

[15] ——, “Reducing users’ perceived mental effort due to interruptive notifications in multi-device mobile environments,” in *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, 2015, pp. 475–486.

[16] N. Dahlbäck, A. Jönsson, and L. Ahrenberg, “Wizard of oz studies—why and how,” *Knowledge-based systems*, vol. 6, no. 4, pp. 258–266, 1993.

[17] F. Portet, M. Vacher, C. Golanski, C. Roux, and B. Meillon, “Design and evaluation of a smart home voice interface for the elderly: acceptability and objection aspects,” *Personal and Ubiquitous Computing*, vol. 17, no. 1, pp. 127–144, 2013.

[18] Q. T. Tran, G. Calcaterra, and E. D. Mynatt, “Cook’s collage,” in *International Conference on Home-Oriented Informatics and Telematics*. Springer, 2005, pp. 15–32.

[19] S. M. Billah, V. Ashok, and I. Ramakrishnan, “Write-it-yourself with the aid of smartwatches: A wizard-of-oz experiment with blind people,” in *23rd IUI conference*, 2018, pp. 427–431.

[20] B. Altakouri, G. Kortuem, A. Grünerbl, K. Kunze, and P. Lukowicz, “The benefit of activity recognition for mobile phone based nursing documentation: A wizard-of-oz study,” in *International Symposium on Wearable Computers (ISWC) 2010*. IEEE, 2010, pp. 1–4.

[21] M. Luria, G. Hoffman, and O. Zuckerman, “Comparing social robot, screen and voice interfaces for smart-home control,” in *CHI conference*, 2017, pp. 580–628.

[22] N. Mairittha, T. Mairittha, and S. Inoue, “Optimizing activity data collection with gamification points using uncertainty based active learning,” in *Adjunct Proceedings of the 2019 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2019 ACM International Symposium on Wearable Computers*, 2019, pp. 761–767.