Woznowski, P. R., Tonkin, E. L., Laskowski, P., Twomey, N., Yordanova, K., & Burrows, A. (2017). Talk, text or tag? the development of a self-annotation app for activity recognition in smart environments. In *2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops 2017): Proceedings of a meeting held 13-17 March 2017, Kona, Big Island, Hawaii, USA* (pp. 123-128). [7917544] Institute of Electrical and Electronics Engineers (IEEE).
<https://doi.org/10.1109/PERCOMW.2017.7917544>

Peer reviewed version

License (if available):
Unspecified

Link to published version (if available):
[10.1109/PERCOMW.2017.7917544](https://doi.org/10.1109/PERCOMW.2017.7917544)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via IEEE at <http://ieeexplore.ieee.org/document/7917544/>. Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Talk, text or tag?

The development of a self-annotation app for activity recognition in smart environments

Przemyslaw Woznowski*, Emma Tonkin*, Pawel Laskowski*, Niall Twomey*, Kristina Yordanova^{†*}
and Alison Burrows*

*Faculty of Engineering, University of Bristol, Bristol, BS8 1UB, UK

[†]Institute of Computer Science, University of Rostock, Albert-Einstein-Str. 22, 18059 Rostock, Germany

Abstract—Pervasive computing and, specifically, the Internet of Things aspire to deliver smart services and effortless interactions for their users. Achieving this requires making sense of multiple streams of sensor data, which becomes particularly challenging when these concern people’s activities in the real world. In this paper we describe the exploration of different approaches that allow users to self-annotate their activities in near real-time, which in turn can be used as ground-truth to develop algorithms for automated and accurate activity recognition. We offer the lessons we learnt during each design iteration of a smart-phone app and detail how we arrived at our current approach to acquiring ground-truth data ‘in the wild’. In doing so, we uncovered tensions between researchers’ data annotation requirements and users’ interaction requirements, which need equal consideration if an acceptable self-annotation solution is to be achieved. We present an ongoing user study of a hybrid approach, which supports activity logging that is appropriate to different individuals and contexts.

Index Terms—Activity logging; ground-truth acquisition; NFC; self-annotation; smart-phone app; voice-logging.

I. INTRODUCTION

The assumption that human activity data generated by pervasive systems can be interpreted and acted upon is central to enabling smart environments. These smart environments are viewed as a promising means to support the prompt delivery of appropriate services in various domains, such as health and care [1]–[3]. Here, there is a concerted effort to obtain a rich picture of natural human behaviour in real-life settings. Yet automated and accurate activity recognition is a complex challenge that remains largely unsolved. One approach to this challenge seeks to train machine learning algorithms using a baseline set of training data, which has been labelled by one or more human experts. Acquiring this ground-truth can be reasonably straightforward in controlled environments such as laboratories [4]–[7]. However, these approaches are not scalable and, therefore, hold limited practical value for real world deployments.

In order to build smart environments that are capable of delivering localised and timely interventions, we must also respond to the need to train machine learning algorithms for diverse users as well as diverse contexts. One possible solution is to engage these users in self-reporting their activity

data, which introduces its own unique challenges. There is evidence to suggest that it is only feasible to expect users to self-annotate for short periods of time, to acquire coarse-grained and non-intimate activities [8]. We feel that self-reporting activities is unnatural and introduces a seemingly unnecessary cognitive load. Compliance with self-reporting can therefore be problematic due to the lack of clear and immediate benefits to the user. Herein lies an opportunity to develop usable and useful tools for self-annotating activities, which are underpinned by simple interaction models but also draw on strategies that foster compliance. It is worth noting that there is no silver bullet to this problem, though it is foreseeable that successful solutions need to be customisable, in order to reflect individual user preferences.

Our work aims to draw together researcher and user requirements in the space of ground truth acquisition, with a view to developing an effective self-annotation tool. In this paper, we present a number of user-tested design iterations, through which we derived a set of requirements for ground truth acquisition systems. Building on these experiences, we developed an app that supports various modes of logging activity and location, which we are currently evaluating with users living in a prototype smart home. We begin by exploring self-annotation requirements, in addition to available tools that provide activity, location and other relevant data on a regular basis.

II. RELATED WORK

A. Understanding requirements for self-annotation

Activity recognition has attracted a lot of research interest, yet there are many unsolved problems in this domain. This is partially because researchers themselves do not know exactly what they are after. Many developments in this space are technology- rather than requirements-driven as argued in [9]. Very few studies on activity recognition in smart environments list a comprehensive set of requirements for ground-truth annotation. [10] and [11] used a method called experience sampling to acquire user annotations in a ‘free living’ experiment. Tapia et al. [10] issued participants with a personal digital assistant (PDA) running the experience sampling method (ESM) software. Every 15 minutes, participants were notified via a

beep sound to questions about: what they were doing at the beep and for how long; and whether they were doing another activity before the beep. Their study was conducted in a single-occupancy scenario where all the sensors activations could be attributed to an individual participant. They captured activity type, time, and duration, although not very accurately. Upon interviewing the participants, they realised the weaknesses of their method: some activities were recorded by mistake; activities of short duration were difficult to capture; there were delays between the sensor firings and the labels of activities; fewer labels were collected than anticipated (low compliance); and sometimes participants specified one activity and carried out a different one [10].

The aspiration of machine learning and artificial intelligence systems is to surpass the ‘human-level’ of predictive ability on a given task. Since the requirements of any one task will define the quality of annotations that are required, there is no universally accepted set of requirements for annotation campaigns from a machine learning perspective [12]. Indeed, forcing explicit labels has been criticised as providing ‘incomplete’ descriptions of the data in classification tasks [13]. To overcome these and other issues, some researchers capture and deliver label uncertainty explicitly by averaging over multiple annotations of the same data [14] or by utilising enterprise-scale crowd-sourcing technologies such as Amazon’s Mechanical Turk [15]. Predictive models learnt on such data can be seen to model the ‘average annotator’ and will yield predictions that are less susceptible to the bias of any single annotator. When learning technologies are deployed in the wild, adaptive classification models will update their parameters in response to new annotations automatically [16]–[18]. In these scenarios, the presence of mistakenly selected annotations will significantly deteriorate the quality of predictions and such events should be avoided.

One means of delivering self-annotation tools are smart-phones or similar devices, for which there are a number of guidelines on interface design. Choi et al. in [19] ran a user study and found that for smart-phones “a simplified interface design of the task performance, information hierarchy, and visual display attributes contributes to positive satisfaction evaluations when users interact with their smartphone”. Other literature in this space, e.g. [20], advises on all aspects of user interface design, ranging from navigation, tools, and charts, to social patterns and feedback. More generally, simplicity is positively associated with perceived visual aesthetics [21] and visual aesthetics influences the perception of usability [22]. The constituent elements of ‘simplicity’ are clarity, orderliness, homogeneity, grouping, balance, and symmetry [23]–[25].

B. Alternative sources for labelling data

It is useful to note that some commonly-used applications create data that may, directly or indirectly, be used as a source of annotations. A well-known example of this type of application is the use of tools intended for personal information management (PIM), which support the creation, storage and

use of information to organise one’s roles, responsibilities and tasks [26]. Such tools may implement functionality such as notetaking, to-do lists and logging of recent activity, as well as collaborative functionality such as instant messaging or calendar sharing. Although not primarily designed for the purpose of capturing annotation data, PIM datasets are sometimes used as part of an annotation strategy (for example, [27]). Data from instant messaging [27] may also contain useful information about location and activity.

Social media services provide sites, APIs and applications that support online discourse through user-generated content [28]. Examples of services of this kind include social networks such as Facebook, microblogging services such as Twitter and Tumblr, photo sharing websites such as Instagram, and link sharing and annotation services, of which Tumblr is also an example. Data originating from other applications, such as the location-sharing service Foursquare [29], may be shared through social media services. Consequentially, social media corpora may be mined for significant amounts of information about times, places, and people [30].

Such tools and services are of interest in discussion of annotation for activity recognition. They are widely and electively used, although the usage of each platform varies by nationality, demographics [31] and personality [32]. Factors in their uptake include enjoyment and perception of usefulness [33]. Individuals are able to tailor their contributions, a presentation of self through user-selected or contributed artefacts [34].

III. EVOLUTION OF THE ANNOTATION APP

In this section, we describe how we approached the problem of ground-truth acquisition for ‘in the wild’ deployment. The agreed platform was Android or web-based apps. Due to the lack of usability and design guidelines specific to ground-truth acquisition systems, we followed general design guidelines for websites and smart-phone app development. We focused on acquiring ground-truth for activities performed at home and their time-stamp, to support the training and validation of machine learning algorithms. We aimed to meet researcher as well as user requirements, therefore we found that these requirements evolved over time as we pilot tested each version of the app with users. All versions of the app were tested on smart-phones only, with the exception of the voice logging app which was also tested on smart-watches.

A. Model-based Approach

Our first version of the smart-phone app was based on the SPHERE ADL ontology [35]. This ontology is organised hierarchically and has up to three levels of activities, ranging from broad categories in tier 1 (e.g. information interaction) through to more specific tier 2 activities (e.g. using a computer) with some including tier 3 detail (e.g. email). The app presented the user with a drop-down list of tier 1 labels and, once an item was selected, it automatically populated another drop-down list with tier 2 activities for that category.

Lessons learned: While we strived to make the app easy to use, we overlooked the fact that the ontology was researcher-

and research-driven. Users of the app tended not know which category to choose first in order to log a particular activity. It also became clear that the academic terminology used in the ontology was clunky and not in keeping with language in everyday use.

B. Voice-based Approach

Following user feedback, we developed a voice-based logging app. In addition to changing the mode of logging, we used this opportunity to experiment with alternative hardware interfaces. Therefore, the same app was implemented for Android smart-phones and Android smart-watches. We sought to keep the information displayed in the apps minimal, to reduce capture burden on the user. However, we became interested in capturing the location of logged activities so, upon terminating an activity, users were asked to specify where it had taken place. We conducted a study to evaluate the usability of voice-based logging and the two different interfaces for self-annotating activity data, which is reported in [8].

Lessons learned: Voice-based logging is a promising approach for self-annotating activity data, but the technology is not yet sufficiently mature. The speech recognition was not always accurate, especially for non-native English speakers, and the interaction is slow. Some people reported that this form of logging was impractical in noisy locations and could be annoying to use in shared spaces. Moreover, users found it burdensome to provide location in addition to activity information.

C. Location-based Approach

We found that acquiring two pieces of information, i.e. activities and their locations, can lead to an unnecessarily complicated interaction model. A user-acceptable solution to ground-truth logging ought to work quickly and efficiently without unnecessary dialogues. Some home activities are bound to particular locations; for example, people tend to prepare meals in the kitchen. Therefore, location information can be bound to activities instead of acquired from the user. Working on this assumption, the location-based app provided the option to choose from different locations in the first instance. Each location was associated with a set of activities for the user to select. Thus the location was directly associated with the activity without the need to manually log that information [36].

Lessons learned: This approach highlighted that there are activities that cannot be bound to a single location; one example of this is vacuum cleaning, which can occur across several rooms as a person cleans their home. On a small interface such as a smart-phone, there is a limit to how many activities can be displayed under each location, without requiring the user to scroll through long lists.

D. NFC-based Approach

Our previous approaches all relied on users remembering to self-report their activities, which presented a challenge

in itself. We thus became interested in exploring how the environment could prompt people to log their activities, perhaps through visual cues in locations where certain activities habitually occur. One promising approach was to leverage the NFC technology available in smart-phones, which has been shown to be usable and robust for self-logging [37]. We developed an app that automatically logged activity and location, upon scanning NFC tags that had been programmed with the relevant information. We then attached labels with the name of the activity over the NFC tags and placed them in appropriate locations in a prototype smart home. Scanning a tag with the smart-phone was used to start and stop logging an activity, but users were also able to stop logging an activity from a list of ongoing activities.

Lessons learned: Common NFC tags don't work on metal surfaces. Although there are NFC tags that are suitable for metal surfaces, we simply avoided placing the tags where we thought there might be interference. Care needs to be taken when deciding where to place the NFC tags, in order to avoid users accidentally logging activities when they put their phone down. This form of logging requires users to pair the right area of the smart-phone with the NFC tag, and some users reported that the interaction was not as immediate as they anticipated.

IV. METHOD

Based on our experiences of pilot testing the various self-annotation approaches, we developed an app that allows users to choose their preferred mode of logging from three available options. In this section we provide details of an ongoing study, in which we are testing this version of the app with people who stay in a prototype smart home.

A. App Design

In the current version of the self-annotation app we took a hybrid approach, which combines the most successful logging modes: voice-based, location-based and NFC-based (Fig. 1). We acknowledge that self-annotation can be cumbersome and that following an ontology can impose an additional cognitive load on the user, so we did not incorporate the ontology-driven approach in this hybrid version. Nevertheless, the ontology terms are still present in the location-based and the NFC-based logging yet the ontology structure is not visible to the user; the voice-based logging is unrestricted.

The main screen of the app comprises a settings cog and four buttons, which correspond to: voice-based logging (**Tell me**), location-based logging (**Choose me**), **Ongoing activities**, and **My history**. Through location-based logging, the user can choose between pre-defined locations and start activities within these locations. To log activities and location via NFC, the user holds the smart-phone in close proximity to a pre-programmed NFC tag and the app opens automatically displaying a confirmation message; repeating this process with the same NFC tag terminates the activity. NFC tags can be programmed with activity and location information through the settings cog. Fig. 2 provides an overview of the hybrid app's functions.

Fig. 1. Logging 'prepare hot drink' with the hybrid app (NFC tag in the background).

Semantic matching is performed across all logging modes, which means NFC-logged activities will show up in the location-based screen. The **Ongoing activities** button has a counter over it to indicate the number of activities being logged through any of the available modes. By clicking on this button the user can select an item from the list, edit its details, delete it or terminate it. Terminated activities are moved from **Ongoing activities** to **My history**. Alternatively, through the settings cog, the user can terminate all ongoing activities with a single button press if, for example, a user leaves the house. Users can manually edit any entry and can create additional activities under each location. With this app, we aimed to meet the following requirements:

- Allow users to log activities in a manner that is appropriate for them and their context;
- Allow users to seamlessly switch between different modes of logging (start activity via one mode and terminate using another mode);
- Allow users to log activities beyond those considered by the researchers;
- Allow users to use natural language, which will in turn help to refine the terminology used in the ontology;
- Combine activity and location information whenever possible.

B. Aim & Objectives

The aim of this study is to evaluate the self-annotation app, deployed within a smart home environment. In doing so, we hope to (a) better understand people's preferences for self-annotation with a view to maximising compliance;

(b) compare self-initiated logging (location-based and voice-based) with logging that is prompted by contextual reminders (NFC-based); (c) expand and refine the ontology to reflect language that is meaningful to end users.

C. Participants & Procedure

This study is embedded within a larger study, in which people are invited to live in a prototype smart home for previously agreed periods of between two days and two weeks. During their stay, participants are encouraged to live and behave as they do at home. Each participant is provided with a smart-phone, which has the self-annotation app installed, and asked to log activities using their preferred mode. After their stay, participants are interviewed about their experiences of living in the smart home and self-annotating using the hybrid app. Due to the characteristics of the prototype smart home, participants must be over 18 years old and able to perform usual daily activities in an unfamiliar environment, without increased risk to themselves or others.

V. EARLY FINDINGS

To date, three participants (two female) have taken part in this study. While we acknowledge that this sample is too small to draw conclusions, we present some early qualitative findings that we feel are of interest for discussion. Different participants preferred different logging approaches, with some using a single mode of logging and others using a combination. Some participants chose their mode of logging by thinking primarily about reliably capturing data rather than their own user experience, as illustrated by the following participant quote:

"I did get into the habit of using the list and once I'd gotten into the habit, it was just much easier to stick with that habit than to change modality. I learnt a method and it worked, sort of thing. [...] Although it wasn't perhaps as easy to use, in principle, I valued the reliability of using the list because I just had to do it and I knew it had been done."

Participants who used a combination of modes of logging explained that their choice depended on the context, such as the type of activity, the location of the activity, how busy they were, and if they were alone or not. While the participant sample is not sufficient to understand if particular modes of logging are better suited to certain activities or locations, we have observed that voice-based logging was the least used approach overall. Some participants mentioned that the process of self-annotating their activities was unnatural, as it required them to be aware that they intended to perform an activity before they began it. Activities such as making a cup of coffee have a relatively clear start and end time. However, as one participant mentioned, drinking that cup of coffee may span a period of time during which a person is sipping that coffee amidst a number of other activities:

Fig. 2. Flow chart of the hybrid app (user interface buttons in blue, NFC logging function in orange).

“Did I start drinking an hour ago but just had several, little periods of drinking, or did drinking start when I first brought back a coffee into my office and it hasn’t finished yet because I’ve still got a bit of cold coffee here?”

It was evident from the data that people had different interpretations of what constitutes an activity, and they also placed different value on what is worth logging. Some activities were less likely to be logged, as they were perceived as personal or intimate. We noted that participants tended to be more compliant with self-annotation in the beginning, but frequency of logging decreased over time. One participant described how she occasionally compensated for not having annotated an activity as it happened by logging it retrospectively and estimating roughly how long it had taken to complete.

Using the smart-phone for self-annotation was generally acceptable, though it could raise some challenges, particularly if the user’s hands were busy. Only a couple of participants said that they don’t habitually carry their phone with them around the house. Nevertheless, the following participant anecdote suggests that using a smart-phone may not be appropriate for all areas of the home:

“I put [the smart-phone] in my back pocket at one point and when I went to the loo, it accidentally fell out. Fortunately, it landed on the floor and not down the [toilet]. I’ve had family members who’ve lost it down the [toilet] before now.”

VI. DISCUSSION & FUTURE WORK

The hybrid approach presented in this paper evolved from taking researcher requirements as the starting point, and subsequently incorporating user feedback to produce a solution that is both useful and usable. The aim of this hybrid approach was

to allow users to self-annotate in ways that were appropriate to them and to their contexts. The NFC and location-based modes are usable and produce annotations that are in-line with an ontology, while voice-based logging is more prone to error but supports unrestricted annotations. We are currently running a study, collecting qualitative data through interviews and quantitative data logged through the annotation app, to better understand which modes of logging are most appropriate and why. We acknowledge that this study is still in the very early stages, and that the work presented in this paper focuses on self-annotation of activities in the home. Nevertheless, we anticipate that eventual learning from this study will be transferable to self-annotation tools for deployment in other environments, such as public and outdoor spaces.

While researchers may be after large quantities of high-quality annotations, it is not always realistic to expect users to provide this level of information about themselves. Given that motivation is central to achieving adequate compliance, more work needs to be done in this space. There are motivational strategies which are worth investigating, in particular given that humans already voluntarily engage in annotations by recording data in PIM systems and posting on social media. It would be worth understanding what factors motivate people to record their data using these media and how they can be leveraged for the purpose of encouraging people to provide ground truth for their data. Other topics that are beyond the scope of this work but warrant attention in future research are privacy concerns and their effect on the reliability of the annotation data. Even though the approaches reported in this paper aim to empower users by providing them with control over their data, it is foreseeable that there are instances in which they might intentionally introduce error.

ACKNOWLEDGEMENTS

This work was performed under the SPHERE IRC, funded by the UK Engineering and Physical Sciences Research Council (EPSRC), Grant EP/K031910/1. We thank our collaborators and the participants who took part in this study for their time and insights.

REFERENCES

- [1] P. N. Dawadi, D. J. Cook, M. Schmitter-Edgecombe, and C. Parsey, "Automated assessment of cognitive health using smart home technologies," *Technology and health care*, vol. 21, no. 4, pp. 323–343, 2013.
- [2] S. S. Intille, K. Larson, E. M. Tapia, J. S. Beaudin, P. Kaushik, J. Nawyn, and R. Rockinson, "Using a live-in laboratory for ubiquitous computing research," in *Pervasive Computing*. Springer, 2006, pp. 349–365.
- [3] N. Zhu, T. Diethe, M. Camplani, L. Tao, A. Burrows, N. Twomey, D. Kaleshi, M. Mirmehdi, P. Flach, and I. Craddock, "Bridging e-health and the internet of things: The sphere project," *Intelligent Systems, IEEE*, vol. 30, no. 4, pp. 39–46, 2015.
- [4] J. Pärkkä, M. Ermes, P. Korpipää, J. Mäntyjärvi, J. Peltola, and I. Korhonen, "Activity classification using realistic data from wearable sensors," *IEEE Transactions on Information Technology in Biomedicine*, 2006.
- [5] L. Atallah, B. Lo, R. Ali, R. King, and G.-Z. Yang, "Real-time activity classification using ambient and wearable sensors." *IEEE transactions on information technology in biomedicine : a publication of the IEEE Engineering in Medicine and Biology Society*, vol. 13, no. 6, pp. 1031–9, Nov 2009.
- [6] M. G. Tsipouras, A. T. Tzallas, G. Rigas, S. Tsouli, D. I. Fotiadis, and S. Konitsiotis, "An automated methodology for levodopa-induced dyskinesia: assessment based on gyroscope and accelerometer signals." *Artificial intelligence in medicine*, vol. 55, no. 2, pp. 127–35, Jun. 2012.
- [7] U. Maurer, A. Smailagic, D. Siewiorek, and M. Deisher, "Activity Recognition and Monitoring Using Multiple Sensors on Different Body Positions," in *International Workshop on Wearable and Implantable Body Sensor Networks (BSN'06)*. IEEE, 2006, pp. 113–116.
- [8] P. Woznowski, P. Laskowski, A. Burrows, E. Tonkin, and I. Craddock, "Evaluating the use of voice-enabled technologies for ground-truthing activity data," in *ARDUOUS: 1st International Workshop on Annotation of user Data for Ubiquitous Systems*. Hawaii, USA: IEEE PerCom, March 2017.
- [9] P. Woznowski, D. Kaleshi, G. Oikonomou, and I. Craddock, "Classification and suitability of sensing technologies for activity recognition," *Computer Communications*, 2016.
- [10] E. M. Tapia, S. S. Intille, and K. Larson, "Activity recognition in the home using simple and ubiquitous sensors," in *International Conference on Pervasive Computing*. Springer, 2004, pp. 158–175.
- [11] N. Kern, B. Schiele, and A. Schmidt, "Recognizing context for annotating a live life recording," *Personal and Ubiquitous Computing*, vol. 11, no. 4, pp. 251–263, 2007.
- [12] J. Whitehill, T.-f. Wu, J. Bergsma, J. R. Movellan, and P. L. Ruvolo, "Whose vote should count more: Optimal integration of labels from labelers of unknown expertise," in *Advances in neural information processing systems*, 2009, pp. 2035–2043.
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [14] N. Twomey, T. Diethe, M. Kull, H. Song, M. Camplani, S. Hannuna, X. Fafoutis, N. Zhu, P. Woznowski, P. Flach, and I. Craddock, "The sphere challenge: Activity recognition with multimodal sensor data," 2016.
- [15] G. Paolacci, J. Chandler, and P. G. Ipeirotis, "Running experiments on amazon mechanical turk," *Judgment and Decision making*, vol. 5, no. 5, pp. 411–419, 2010.
- [16] T. Diethe, N. Twomey, and P. Flach, "Active transfer learning for activity recognition," in *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*.
- [17] N. Twomey, T. Diethe, and P. Flach, "Bayesian active learning with evidence-based instance selection," in *Workshop on Learning over Multiple Contexts, European Conference on Machine Learning (ECML15)*, 2015.
- [18] T. Diethe, N. Twomey, and P. Flach, "Bayesian active transfer learning in smart homes," in *ICML Active Learning Workshop*, vol. 2015, 2015.
- [19] J. H. Choi and H.-J. Lee, "Facets of simplicity for the smartphone interface: A structural model," *International Journal of Human-Computer Studies*, vol. 70, no. 2, pp. 129–142, 2012.
- [20] T. Neil, *Mobile design pattern gallery: UI patterns for smartphone apps*. O'Reilly Media, Inc., 2014.
- [21] D. C. L. Ngo, L. S. Teo, and J. G. Byrne, "Modelling interface aesthetics," *Information Sciences*, vol. 152, pp. 25–46, 2003.
- [22] N. Tractinsky, A. S. Katz, and D. Ikar, "What is beautiful is usable," *Interacting with computers*, vol. 13, no. 2, pp. 127–145, 2000.
- [23] M. Bauerly and Y. Liu, "Effects of symmetry and number of compositional elements on interface and design aesthetics," *Intl. Journal of Human-Computer Interaction*, vol. 24, no. 3, pp. 275–287, 2008.
- [24] M. Moshagen and M. T. Thielsch, "Facets of visual aesthetics," *International Journal of Human-Computer Studies*, vol. 68, no. 10, pp. 689–709, 2010.
- [25] N. Tractinsky, A. Cokhavi, M. Kirschenbaum, and T. Sharfi, "Evaluating the consistency of immediate aesthetic perceptions of web pages," *International journal of human-computer studies*, vol. 64, no. 11, pp. 1071–1083, 2006.
- [26] W. Jones, "Personal information management," *Annual review of information science and technology*, vol. 41, no. 1, pp. 453–504, 2007.
- [27] L. Coyle, J. Ye, S. McKeever, S. Knox, M. Staelber, S. Dobson, and P. Nixon, "Gathering datasets for activity identification," 2009.
- [28] S. Asur and B. A. Huberman, "Predicting the future with social media," in *Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on*, vol. 1. IEEE, 2010, pp. 492–499.
- [29] J. Lindqvist, J. Cranshaw, J. Wiese, J. Hong, and J. Zimmerman, "I'm the mayor of my house: examining why people use foursquare—a social-driven location sharing application," in *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 2011, pp. 2409–2418.
- [30] G. D. McKenzie, "A temporal approach to defining place types based on user-contributed geosocial content," Ph.D. dissertation, UNIVERSITY OF CALIFORNIA, SANTA BARBARA, 2015.
- [31] M. Duggan, N. B. Ellison, C. Lampe, A. Lenhart, and M. Madden, "Social media update 2014," *Pew Research Center*, vol. 9, 2015.
- [32] T. Correa, A. W. Hinsley, and H. G. de Ziga, "Who interacts on the web?: The intersection of users personality and social media use," *Computers in Human Behavior*, vol. 26, no. 2, pp. 247 – 253, 2010. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0747563209001472>
- [33] K.-Y. Lin and H.-P. Lu, "Why people use social networking sites: An empirical study integrating network externalities and motivation theory," *Computers in Human Behavior*, vol. 27, no. 3, pp. 1152 – 1161, 2011, group Awareness in {CSCL} Environments. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0747563210003766>
- [34] B. Hogan, "The presentation of self in the age of social media: Distinguishing performances and exhibitions online," *Bulletin of Science, Technology & Society*, vol. 30, no. 6, pp. 377–386, 2010. [Online]. Available: <http://bst.sagepub.com/content/30/6/377.abstract>
- [35] P. Woznowski, R. King, W. Harwin, and I. Craddock, "A human activity recognition framework for healthcare applications: ontology, labelling strategies, and best practice," in *2016 International Conference on Internet of Things and Big Data (IoTBD)*. Rome, Italy: INSTICC, April 2016.
- [36] M. Schröder, K. Yordanova, S. Bader, and T. Kirste, "Tool support for the live annotation of sensor data," in *Proceedings of the 3rd International Workshop on Sensor-based Activity Recognition and Interaction*. ACM, Jun 2016.
- [37] X. Luo, P. Woznowski, A. Burrows, M. Haghighi, and I. Craddock, "Splash: Smart-phone logging app for sustaining hydration enabled by nfc," in *Proceedings of the 2016 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. ACM, 2016, pp. 1526–1532.