

## Context-Aware Wearables

The last thing we need is a pandemic of stray cats

KIM SNOOKS

ImaginationLancaster, Lancaster University, United Kingdom

JOSEPH LINDLEY

ImaginationLancaster, Lancaster University, United Kingdom

DANIEL RICHARDS

ImaginationLancaster, Lancaster University, United Kingdom

ROGER WHITHAM

ImaginationLancaster, Lancaster University, United Kingdom

We present Connected Companion (CoCo), a health tracking wearable that provides users with timely, context-relevant notifications aimed at improving wellness. Traditionally, self-tracking wearables report basic health data such as resting heart rate; these data are visualised and positive behaviours (e.g. exercising often) are encouraged with rudimentary gamification (e.g. award badges) and notification systems. CoCo is the first wearable to combine caffeine, alcohol and cortisol sensors, a context network (which predicts user context), and a wellness model (which establishes per-user wellness measures). Working in tandem these provide users with notifications that encourage discrete behaviours intended to optimise user-wellness per very specific biological and social contexts. The paper describes the (sometimes unexpected) results of a user-study intended to evaluate CoCo's efficacy and we conclude with a discussion about the power and responsibility that comes with attempts to build context-aware computing systems.

**CCS CONCEPTS** • **Human-centered computing** → *User interface design*; **Information visualization**; • Applied computing → *Health informatics*; • **Computer systems organization** → Sensors and actuators; • **General and reference** → Design.

**Additional Keywords and Phrases:** design fiction, imaginary paper, sensors, health trackers

### ACM Reference Format:

First Author's Name, Initials, and Last Name, Second Author's Name, Initials, and Last Name, and Third Author's Name, Initials, and Last Name. 2018. The Title of the Paper: ACM Conference Proceedings Manuscript Submission Template: This is the subtitle of the paper, this document both explains and embodies the submission format for authors using Word. In Woodstock '18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY. ACM, New York, NY, USA, 10 pages. NOTE: This block will be automatically generated when manuscripts are processed after acceptance.

## 1 LOTUS, IVY, AND BLOOM: DATA DRIVEN WELLNESS AND CONTEXT MODELLING

CoCo is the first wearable to combine caffeine, alcohol and cortisol sensors, external data streams, and machine learning to implement context-aware sensing for providing timely and relevant notifications for users intended to help them optimise their wellness<sup>1</sup>. CoCo's software is built around three interlinked components a context-network (IVY), a wellness model (LOTUS), and a management layer (BLOOM). In terms of hardware, the CoCo wearable uniquely combines multiple biosensors into a single device providing the software components with the data necessary to create and maintain bespoke wellbeing and context models. These include a cortisol (potentiostatic circuit and chronoamperometry [43]) sensor; caffeine (electro-chemical differential pulse voltammetry sensing [46]); and alcohol sensor (using proprietary ION sensor cartridges [24]). Alongside these specialised biosensors, data pertaining to heartrate, blood oxygen, movement, ambient sound, and temperature are also captured. CoCo users are required to give permission for CoCo to access additional 3<sup>rd</sup> party data in order to provide the IVY context-network with sufficient data to reach reasonably confident context assessments. These data include location, content and meta-data of calendar entries and messaging apps (compatible with various Email Clients, WhatsApp, Facebook Messenger, iOS/Android SMS apps). To support the functions of the LOTUS wellness model, CoCo prompts users to generate wellness labels at regular intervals (Figure 1c 'Wellness Check').

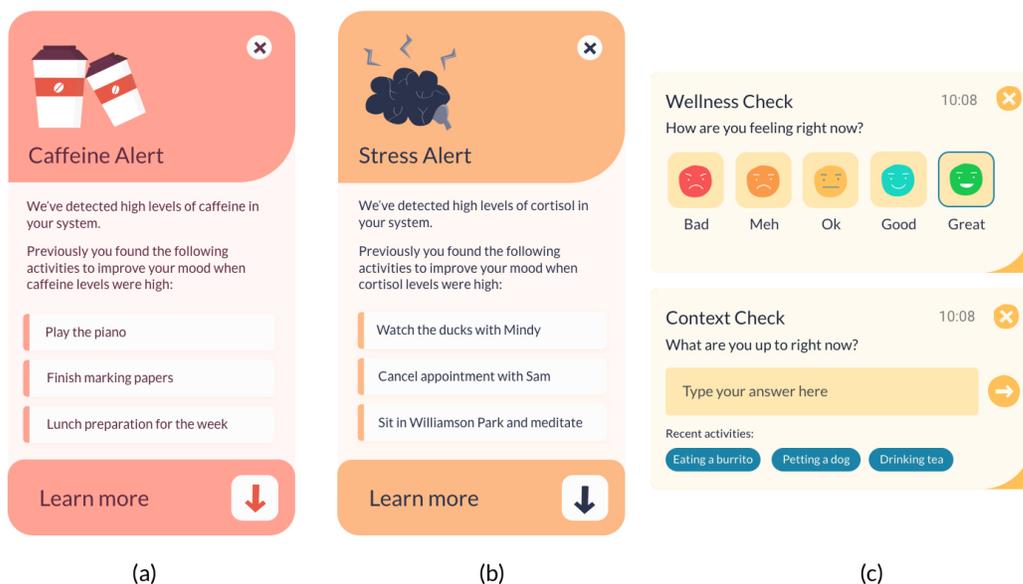


Figure 1: CoCo Notifications sent to users during the user study

The IVY context network (Figure 2) correlates live sensor data and 3<sup>rd</sup> party data (e.g. calendar, location) in order to create robust labels which can describe context to a high level of accuracy. For, example calendar and messenger entries may suggest a 'coffee shop meeting' which will be reinforced by relevant sensor data such as location, sound signature (e.g. the sound signature of a coffee shop), and increased caffeine. Using supervised

<sup>1</sup> A formal definition of wellness is beyond the scope of this paper and incorporates aspects of mental and physical health alongside subjective accounts of happiness and emotional well-being [8].

learning, the model increases confidence via manual interventions with the user (Figure 1c, 'Context Check'). The result is a bespoke model which can take partial datapoints to infer context to a high level of certainty on a per-user basis.

### Context-Network Model (IVY)

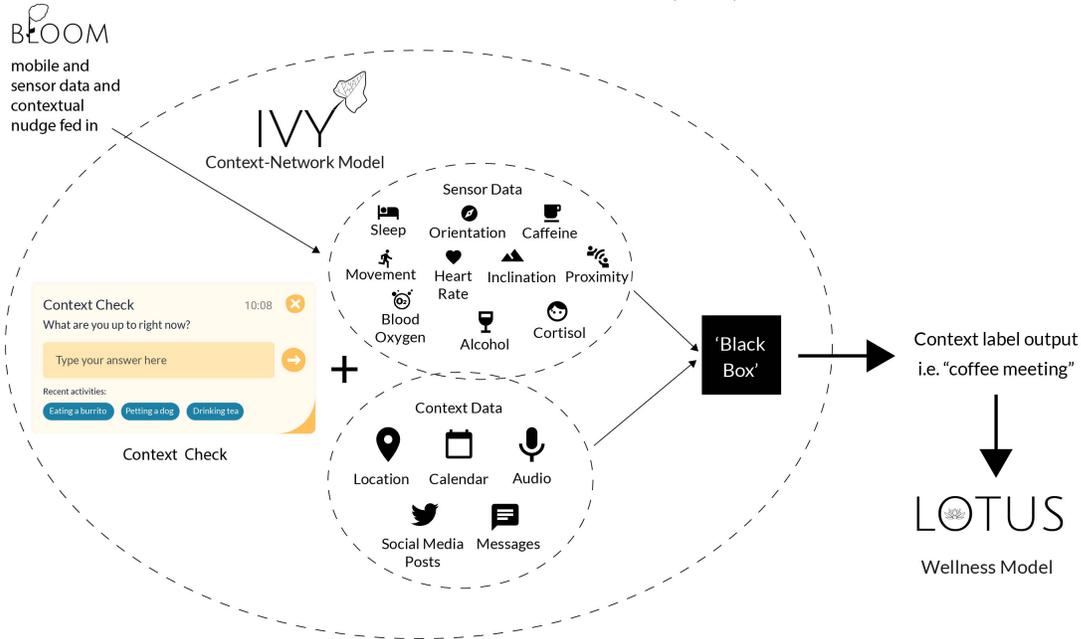


Figure 2: Context-Network Model (IVY)

The wellness model LOTUS (Figure 3) correlates labels from the IVY context network (e.g. 'having a coffee meeting with a colleague') with user-generated wellness tags and relevant sensor data<sup>2</sup> (Figure 3). In contrast to traditional health-related wearables that assume an average or standard interpretation of wellness for all users, through LOTUS, CoCo *learns* what wellness means for each individual user. The result is an architecture which has the capability to adapt to nuances of both context and perceived wellness, and to do so for each user uniquely.

<sup>2</sup> Notably, sensor data (e.g. high caffeine level) is processed by the LOTUS model *independently* of IVY, this means that biosensor data exert influence on CoCo's understanding context and wellness independently of each other.

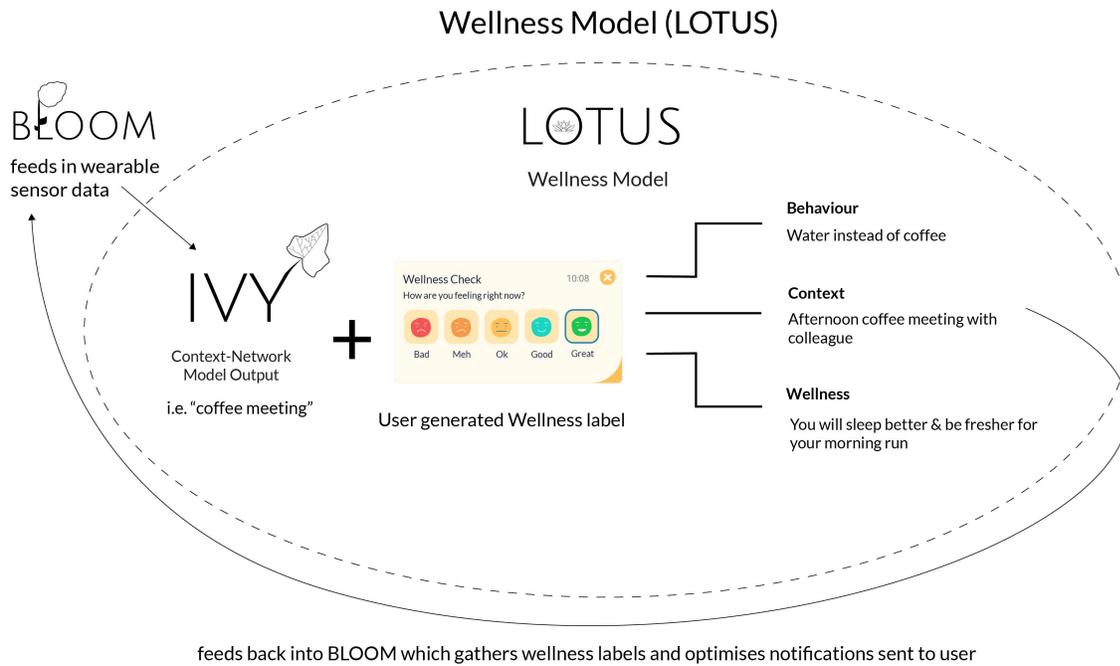


Figure 3: Wellness Model (LOTUS)

BLOOM is CoCo’s management layer (Figure 4). BLOOM can query both IVY and LOTUS models in order to highlight and promote correlations between established context and wellness baselines. Until the model meets a pre-determined confidence threshold BLOOM runs in a training mode, allowing IVY and LOTUS models to learn based on each user’s data. Once baselines are in place BLOOM provides notifications (these are referred to as *Welltexts*—a portmanteau of wellness and context) to users. *Welltexts* are notifications designed to encourage users to adopt specific behaviours (e.g. reduce caffeine intake or get more sleep) which, at particular times or depending on context, may increase their predicted wellness. In addition to managing *Welltexts*, BLOOM is also responsible for enhancing the confidence that each model has by prompting users to generate further training data to label what they are doing and how they feel about it at key inflection points (see Figure 1). Figure 4 shows the system architecture: BLOOM manages notifications; sensor data from the wearable feed both IVY context network and LOTUS wellness model independently; 3<sup>rd</sup> party data (location, messaging, audio, etc) are fed into IVY; IVY’s outputs (i.e. computed contexts) are utilised by LOTUS; both models are continually reinforced by additional user labelling, managed by BLOOM. The components are configured to encourage users towards behaviours which will increase their reported wellness as much as possible.

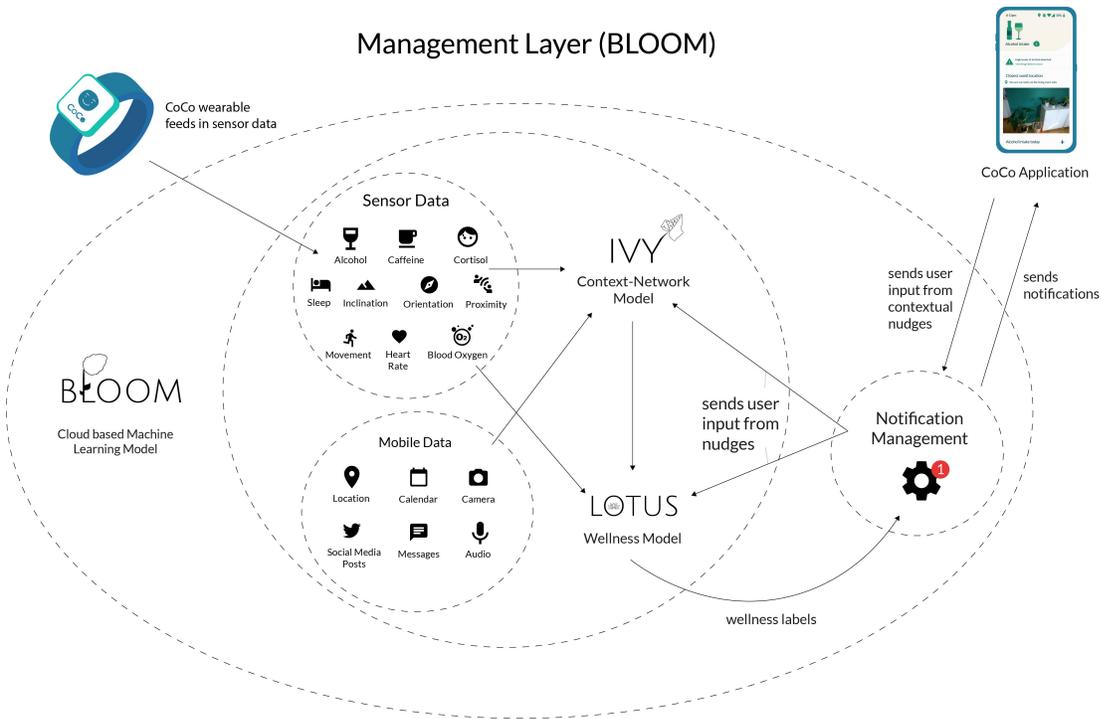


Figure 4: Diagram showing how CoCo’s software and hardware components function together.

## 2 BACKGROUND AND RELATED WORK

Historically, health trackers embed sensors such as ambient light sensors and accelerometers to infer basic facts about a user’s life, e.g. step count, sleep cycles, heart rate. Miniaturisation has allowed the current generation to combine biosensors which can detect not only heart rate but, alcohol, caffeine and cortisol levels into a single device. Moreover, while self-tracking devices provide users with a plethora of visual means to interpret data, users struggle to derive meaning from the information [17]. This is, in part, due to a separation of data and context [31]. The upshot is that users find it very difficult to make meaningful decisions based on their tracking data [10,23].

Context-aware computing [14,15,47] focuses on detecting movements, routines and actions to provide relevant contextual information to a user. Whilst contextual computing has been a long-term aspiration of HCI and Ubiquitous Computing it has struggled with transitions between states [38], differentiating activities [20], human perception [38] and supporting specific goals [6]. The crossover with affective computing, to recognise and interpret human emotion, further highlights the complexity and challenges in creating context-aware systems [37]. However, improvements in machine learning [39], increased availability of relevant data [9], an enhanced battery and network performance mean that efficacious context models are increasingly practicable [11]. Although research has attempted to encourage behaviour change through recommendations [18,20,21,38,42,49], machine learning models [28,41] and timely interventions [30,32,39], CoCo is the first system, that we know of, which combines alcohol, caffeine and cortisol sensors with a functional context model in order to encourage specific user behaviours.

In the remainder of the paper, we describe preliminary findings of a user study. Specifically, we draw upon data from user interviews to evaluate user experience. Initial results show that CoCo elicits a high level of engagement with users and can encourage behaviour change. However, due to identifying several unintended consequences of the system, we suggest that further work is needed in order to define systems which can reliably support a positive user experience and minimise unanticipated negative outcomes.

### 3 USER STUDY

We deployed CoCo wearables to 15 participants who were asked to use the system for 6 months. The participant cohort comprised 3 co-habiting childless couples, 2 who lived alone, 1 family with children and 1 household of shared occupancy; in total there were 8 male and 7 were female participants; the median age was 28, mean age was 34, youngest 9 months and eldest 80 years old. Participants were given a CoCo wearable and instructions on how to download and use the mobile application. Participants have been pseudonymised throughout.

Table 1: List of participants interviewed in the user study

Household	Name	Age	Gender	Profession
1	Philip T	50	Male	Doctor
1	Kat T	47	Female	Swimming Instructor
2	Euan C	30	Male	Veterinarian
2	Gina C	28	Female	Barista (part time)
2	Piper C	9 months	Female	n/a
2	Darlene C	5	Female	n/a
3	Ron G	80	Male	Retired
3	Doris G	76	Female	Retired
4	Kelsie B	55	Female	Botanist
5	Cecilia L	22	Female	Biology Student
5	Mohammed A	19	Male	Computing Student
5	Russell F	18	Male	Psychology Student
6	Gabriel A	33	Male	Author
7	Mark G	26	Male	Software Engineer
7	Jason G	27	Male	Teacher

We interviewed participants in their households at 3 and 6 months into the trial using an unstructured ethnographical interview approach<sup>3</sup>. The purpose of this data gathering exercise was to understand the user experience of the CoCo system. In particular we wished to identify different perspectives, motivations, attitudes of the participants, as well as highlighting any problematic aspects of the system or social tensions that arose due to participants' adoption of the technology. The data, we suggest, is relevant to this particular implementation of a data-driven context-awareness system but may offer other researcher insights into generalisable challenges associated with encoding context. In the following we present 7 themes which have emerged from our preliminary engagement with the data.

<sup>3</sup> Please note that the trial took place in the latter half of 2019 before the Covid-19 pandemic and social distancing restrictions were in place.

### 3.1 Exercising

CoCo routinely recommended exercise to participants who usually found those *Welltexts* to be useful: Kelsie said "I upped my training regime" because "it decreased my stress levels" whilst Philip noted that CoCo "would often remind me that I might relax more if I went for a run". Due to the unique sensor implementation Kelsie was also able to use CoCo with tattoos, which was previously a problem with other wearables [40]. Euan also noted that the personalisation of *Welltexts* helped motivate him to exercise "because of running I'm destressing significantly". However, on numerous occasions the app's tendency to propose exercise was also problematic. For example, Philip's stress was often highest while he was engaged at work seeing patients (he is a Doctor) and those were the occasions that CoCo suggested he exercise (it is likely the IVY context network couldn't determine context because Philip's work diary was private). Conversely, for some users, the correlation between wellness and exercise was lost due to per-user training. For example, Russel—a first year psychology student—experienced that "rather than telling me to exercise more and eat better, encouraged me to spend more time socialising (Figure 5a)". This short-term gain (which enhanced his reported wellness) became problematic "in the end I had to stop using it. The time spent making friends was great, but it also made the end of term very stressful".

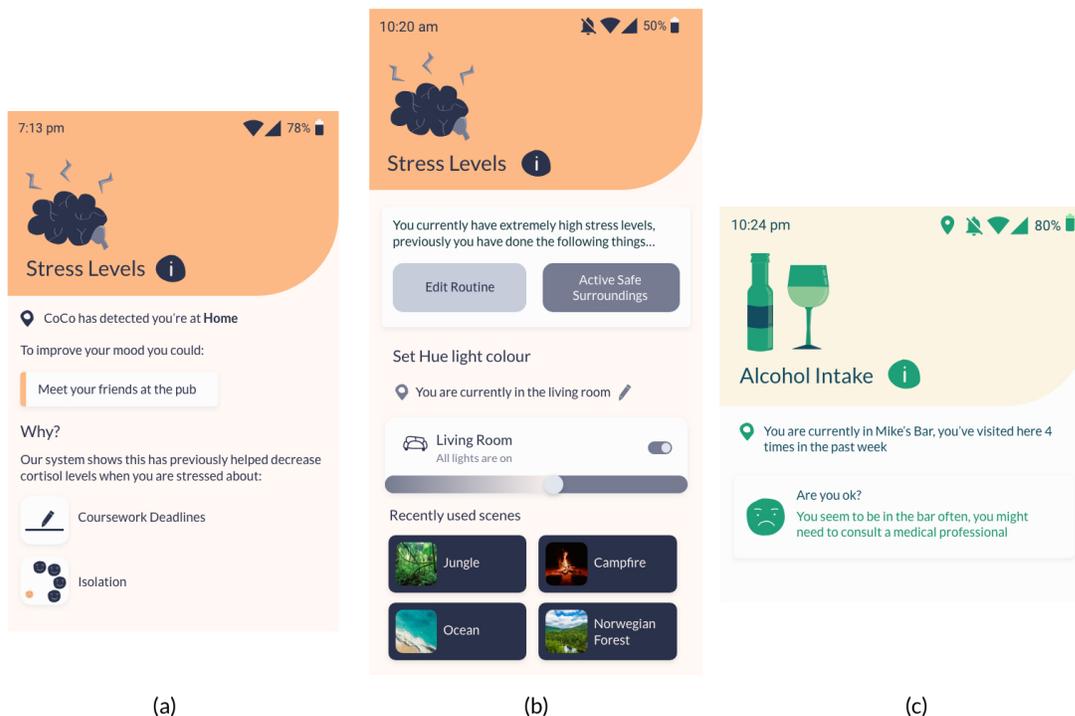


Figure 5: Screenshots of users visualised health data in the CoCo app

### 3.2 Relationships

CoCo did help people become more social, Cecilia felt happier after using CoCo, catching up with "all the mates I hadn't seen in forever" and Doris went to knitting club more often developing new skills like "finally learning how to cable knit". However, whilst CoCo was learning, it suggested Doris—a 76-year-old with a knee problem—go

mountain climbing. This suggestion was problematic in part because Doris wasn't able to do it, but moreover by reminding her that she wasn't able to CoCo inadvertently lowered her mood. These suggestions improved as users continued to interact with the system, but Philip suggested that maybe "technology is not the solution". His friend only improved his health through "a combination of his wife nagging him, and me going around his house on my daily run to pick him up". Others also found issues with the attention CoCo required, Cecilia finding that CoCo's frequent notifications and requests for additional context information to be "annoying" and "needy" when she was trying to catch up with friends. Frustrated that she was unable to use her phone to read and respond to CoCo's notifications during her work shift (she is a swimming pool lifeguard and phones in the pool room are banned), Kat eventually stopped taking part in the trial due to the "irritating" notification system. Ron found the constant "beeping and buzzing, asking what I'm doing, how I'm feeling" quite overwhelming, mentioning that the phone interfered with his day-to-day life; "I'm supposed to be retired, but this made looking after my phone a full-time job". Gina thought the app's recommendations were things she would enjoy, but that it repeatedly suggested things that were not possible at the given time "I can't choose mood lighting or listen to whale noises when the baby is crying" (Figure 5b). Euan (Gina's partner) also indicated that CoCo interfered with major life decisions; "seeing all the disruption that Darlene [their child] causes in a graph was quite startling [...] it made us think twice about a second child". Jason (a teacher) also mentioned CoCo was not able to distinguish between professional and social situations. After a successful parent's evening meeting which CoCo interpreted as a social occasion, the app later suggested that he go for a drink with the student and their parent; "It definitely needs some kind of filter so I can say do not under any circumstances suggest this again".

### 3.3 Alcohol

CoCo made people more aware of their alcohol consumption. Gina became more stressed after seeing the graph stating that "it looks like we spend half our lives pissed while looking after the kids" but later says that "I think we're just normal"; similarly, Cecilia said "I don't think I drink more than a typical student". In both cases CoCo had learned that increased alcohol intake had a short-term positive impact on self-reported wellness, and hence suggested alcohol use more frequently. Perhaps unsurprisingly this was not a successful strategy (ultimately Cecilia needed to retake exams after following the app's advice, which she attributed to the increased alcohol use). Encouraged by CoCo, Gina and Euan began to drink to relax after putting their kids to bed, data which when it was presented back to the young parents caused a mild social anxiety. CoCo flagged Kelsie as a problem user of alcohol (Figure 5c) as it confused her working at a bar with social alcohol use, this decreased her confidence in the system; "I'm not sure I trust it now, to be honest". Algorithmic bias was also revealed as an issue after CoCo mistakenly suggested Mark cut down on his drinking; "I looked at the word cloud and noticed the slurred words [...] that's when it hit me, the system thinks I'm drunk because of my speech impediment! It gets worse when I'm tired, and I've been working late recently" (Figure 6a). While Mark reported that "we had a good laugh about it" he went on to note that "others could be more sensitive, you should be careful about that kind of thing... it could really affect someone's confidence".

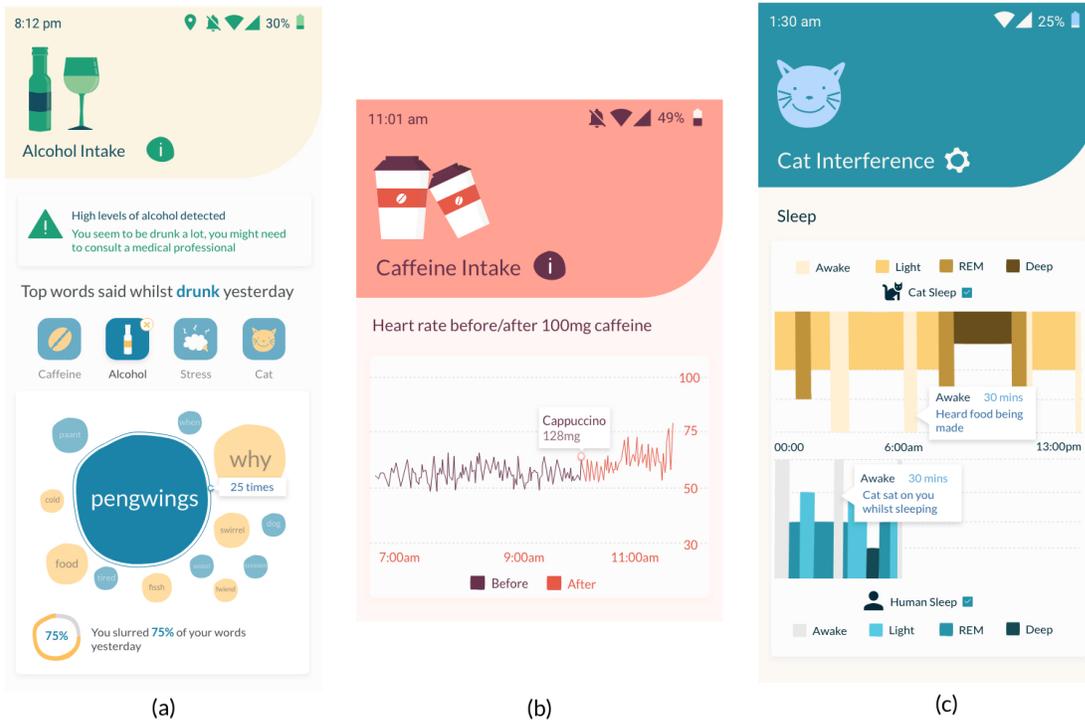


Figure 6: Screenshots of users visualised health data in the CoCo app

### 3.4 Stress

CoCo was able to correctly identify activities that were causing people stress. For Jason a regular “all staff” meeting was one such cause that, based on calendar data, CoCo suggested he avoid. Adopting CoCo’s suggestion meant that Jason ultimately had to explain this to the headteacher, which “increased my stress levels a lot more than attending the staff meeting”. Moreover, because CoCo learns according to previous data it is unable to determine why a certain activity, might not be a possibility, Gina said it suggested “doing Darlene’s medication *after* chatting with friends over a glass of wine”. The medication which Darlene (Gina’s child) is taking is time-sensitive, hence CoCo’s suggestion—despite being based on relevant data—is an impossibility. Russel’s account highlights “Hard work is necessary sometimes [...] I need to feel some pressure in order to take on the challenge of University”. However, because CoCo’s wellness model is primarily driven by short-term self-reported wellness and the data gathered from the biosensors, it was more likely to encourage Russel to engage in social activities rather than study. Toward the end of the 6-month trial this resulted in a significant net increase in Russel’s stress.

### 3.5 Caffeine

CoCo helped Gabriel to see how stressed his caffeine intake was making him (Figure 6b) but this also caused more problems; “the heating bills increased because CoCo kept telling me to turn the heating up [...] when I saw the bill my heart sank [...] Plus, I wasn’t talking to anyone—I’m a social person, so that stressed me out [...] And the bloody caffeine withdrawal...I had awful headaches”. Whilst CoCo correctly correlated an increased caffeine level

to some stress markers, the consequences of Gabriel following its suggestion decision led to new issues which were ultimately more significant. Kat noted that apps encouraging young people to decrease their caffeine intake could cause “a big downturn in footfall for local businesses, some cafés might struggle to survive”. Based on this view, Kat ignored the suggestions provided by CoCo to decrease her coffee intake. Philip liked that the app made him think about his coffee consumption, but believes, like the app, that his caffeine consumption is making him stressed. During our interview it transpired that Philip was “trying to teach myself latte art, it’s a bit of a hobby. I’m not very good, so, that winds me up as well”. Whilst Philip’s stress was associated with caffeine, a portion of the stress around the coffee intake was due to trying to learn latte art. However, as it has no data trace, the latte art activity was entirely absent from CoCo’s reasoning.

### **3.6 Cats**

Both Gabriel and Ron point out learning more about previously hidden behaviours of those around us can have an impact on our lives. Ron liked the sleep charts in CoCo (see Figure 6c) as it supported his own theories to the causes of his health problems, claiming that the cat (Muffin) was waking him up, “Muffin had to go, ten times she’s woken me up in the past week”. On the other hand, his wife Doris had a different theory “it was actually the TV waking him up. Every single day he’d refuse to go to bed saying he’s not tired and then he falls asleep. Muffin likes to sit on him when he’s sleeping in the chair”. Based on Ron’s (likely incorrect) conclusion about the cat, Muffin was rehomed. Gabriel also put his cat (Pickle) up for adoption after seeing that interactions with the cat increased his heart rate. With a recent hypertension diagnosis, Gabriel concluded that Pickle’s presence was too much of a risk. The long-term data showed that Pickle’s departure represented a significant decrease in Gabriel’s wellness, and with the cat out of the equation, CoCo then began to cite coffee intake and house temperature as potential causes. “I really miss Pickle, that’s my biggest regret about using the app [...] Imagine, if everyone used it, there’d be a pandemic of stray cats!”.

### **3.7 Confidentiality**

In order to even have a chance of working CoCo needs access to lots of data, often that data implicates 3<sup>rd</sup> parties in an unanticipated or hard to predict way. Mark had to stop using the app as “it suddenly started referring to a project that’s currently under an NDA [non-disclosure agreement] [...] I can only imagine it picked up on that via the audio? Not cool, so I stopped using it immediately”. These kinds of findings bring up questions to how the app warns users of the implications of using the app and whether the responsibility to solving these issues is on the user, the company, or the user’s friends/client. Mohammed expressed similar concerns with the amount of data gathered, he felt that the app was trying to “control all my decisions” and that he didn’t “feel comfortable giving it the power to change and dictate my life”. The increasing ubiquity of sensing devices also presents causes for concern; Kat mentioned that “the pool is supposed to be a private space, even though it’s shared physically phones are banned because people were worried about cameras and recording and stuff [...] I’m not sure what the rule would be with biosensors though, can they detect other people nearby?”.

## **4 CONCLUSION**

This paper discusses preliminary findings of a novel system which combines state of the art biosensors with machine learning to provide users with timely, context-relevant notifications intended to increase wellness. Whilst we can report that our system certainly has the potential to encourage positive behaviours, we are duty bound to

report on the more problematic aspects of the system. CoCo was extremely effective at identifying causal relationships between specific activities and relative wellness, this did not represent a meaningful handling of context. The user study highlights the difficulty of distinguishing hedonistic (and enjoyable) behaviour from long-term positive behaviours. The attempt to implement context-awareness repeatedly failed to predict that changes in the understood part of the system (e.g. proposing a reduced caffeine intake) could have an impact on the *non-understood* part of the system (e.g. Gabriel's increased heating bill); this results in a form of context-awareness which is very naïve. Whilst every attempt was made to reduce algorithmic bias it arose in even the most unexpected places (e.g. interpreting Mark's speech impediment and late nights at work as an alcohol problem). Although CoCo encouraged exercise, this applied mostly to users who already exercised frequently and in the case of Doris, by suggesting a mountain climb, CoCo actually introduced a new barrier to wellness. In most cases, users followed their intuition and ignored CoCo's suggestions when they were nonsensical or dangerous (e.g. Gina ignored the suggestion to not give her baby its medication at the right time). However, in the case of Pickle and Muffin, the owners were convinced that the feline presence was causing them harm and the cats were rehomed. Whilst about cats and not humans, this supports similar findings to Tolmie et al [48] who point out the 'invasive step' of data with relationships.

While our context model utilises an unprecedented amount of data and biosensors and builds per-user models based on those data, clearly the approach is limited. While sophisticated the model is only ever aware of the datapoints and attributes which it is aware of. The reality of a human sense of context was significantly more sophisticated than our design assumptions. The aspiration of CoCo is to improve the wellness of users using data and wearable technology. Our study suggests that encouraging behaviour change based on data-driven models is possible, but that determining whether the behaviour change is positive reliably is an unsolved problem, which, in order to solve, we must involve users throughout the design process.

## 5 DISCUSSION

Up until this point, and notwithstanding its viability, CoCo is a speculation. This paper is an example of Design Fiction as World Building [13] presented in the form of a fictional paper [3,33,34]. The aspiration for the paper is to provide enough detail for the speculation to appear plausible enough to engender a 'suspension of disbelief'. Whilst the logic of fictional papers is discussed elsewhere [35], it is important to discuss the significance of this *specific* paper, what its original contributions are, and how it represents a valid—but intentionally experimental—contribution to the (alt) HCI discourse.

The paper aims to contribute to a series of contemporary HCI concerns<sup>4</sup>. These include the potential applications of machine learning [50]; guidelines for implementing responsible, ethical and transparent AI systems [26,29]; and the emergence of Human-Data Interaction as a sub-discipline of HCI. Alongside the explosion of applications of AI (which, more often than not refers to variants of Machine Learning), recent years have seen a proliferation of frameworks, guidelines, and manifestos intended to support, encourage, or underpin ethical, transparent, and responsible system development [1,2,5,25]. While such initiatives are worthy and valuable endeavours until applied to specific applications or contexts, they remain abstract and are of limited use. Conversely, once systems are implemented, it is often too late to substantively change their design—an issue which is particularly salient in

---

<sup>4</sup> Notably, although they are overlapping concerns, we deliberately do *not* attempt to engage with the privacy and trust risks associated with such a system, this is a deliberate decision in order to maintain the central focus of the paper.

industry contexts where a minimum viable product may underpin a company's financial viability. To this end, we advocate for design-inspired research—in this case a Design Fiction/fictional paper—as a viable means to test proposed systems designs against guidelines and frameworks. For example, researchers may use fictional papers such as this one in order to check whether a proposed implementation would meet proposed design guidelines [4].

As it sits within a relatively small niche, we would like to draw attention to the paper's methodological contribution. Originally proposed as part of Design Fiction's maturation within the HCI community, fictional papers (a relative of Imaginary Abstracts [7]) write up research which never actually happened. To date there are a handful of examples of fictional papers in the HCI canon, some of these are situated far into the future and are more irreverent [27], whilst others are more 'realistic' (so much so that they may even appear to be deceptive [12]). By providing an additional example to the body of published realistic fictional papers, we hope that this work will help contribute to the further maturation of the approach.

Lastly, looking back at first and second generation health sensors such as Fitbit Sense [44] and Withings ScanWatch [19] we can see how much the integrated context-aware capabilities of CoCo enable intelligent and timely interventions to improve wellness. The findings of this (fictional) study also show the need for greater context-awareness in systems seeking to shape behaviours relating to health but also show that this must be combined with facilities to maximise user autonomy and to support users in making informed decisions based on the transparent processing of their data. Whilst in computing, modelling context often comes down to measurable features such as location, time, activity [15] etc., there are many uncertainties in human behaviour [16] which make it difficult to predetermine many situations in a computational system. Sociologists have considered how our actions [45] and knowledge [22] changes given a situation to help us understand something. We put to readers what is context? Perhaps the unintended consequences experienced by users in this study could have been addressed with better modelling of contextual factors and the intentions behind the actions users took. Conversely, perhaps the aspiration to fully model context is ill-advised, unattainable "Heffalump Trap" [36]. We suspect that the reality lies somewhere in between these extremes, and that attempts to model context may yield many benefits and come hand-in-hand with limitations. As such, we advocate for the use of future-focused and speculative research methods to concretise and explore the realities of these HCI challenges *before* such systems are implemented.

## ACKNOWLEDGMENTS

This research has been funded by Design Research Works UKRI Grant reference [MR/T019220/1](#) and by Research England Expanding, Excellence in England (E3).

## REFERENCES

- [1] ACM. ACM Code of Ethics and Professional Conduct. *ACM*. Retrieved December 11, 2020 from <https://ethics.acm.org/>
- [2] ACM. 2017. *Statement on Algorithmic Transparency and Accountability*. Retrieved December 11, 2020 from [http://www.acm.org/binaries/content/assets/public-policy/2017\\_joint\\_statement\\_algorithms.pdf](http://www.acm.org/binaries/content/assets/public-policy/2017_joint_statement_algorithms.pdf)
- [3] Haider Akmal and Paul Coulton. 2020. The divination of things by things. In *Conference on Human Factors in Computing Systems - Proceedings*, Association for Computing Machinery, New York, NY, USA, 1–12. DOI:<https://doi.org/10.1145/3334480.3381823>
- [4] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N. Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, and Eric Horvitz. 2019. Guidelines for human-AI interaction. In *Conference on Human Factors in Computing Systems - Proceedings*,

- Association for Computing Machinery, New York, New York, USA, 1–13.  
DOI:<https://doi.org/10.1145/3290605.3300233>
- [5] Oliver Bates, Kathy New, Matthew Louis Mauriello, Roy Bendor, Simran Chopra, Adrian K. Clear, Samantha Mitchell-Finnigan, Christian Remy, Samuel Mann, and Chris Preist. 2019. Towards a responsible innovation agenda for HCI. In *Conference on Human Factors in Computing Systems - Proceedings*, Association for Computing Machinery, New York, NY, USA, 1–8.  
DOI:<https://doi.org/10.1145/3290607.3299017>
- [6] Frank Bentley, Konrad Tollmar, Peter Stephenson, Laura Levy, Brian Jones, Scott Robertson, Ed Price, Richard Catrambone, and Jeff Wilson. 2013. Health mashups: Presenting statistical patterns between wellbeing data and context in natural language to promote behavior change. *ACM Trans. Comput. Interact.* 20, 5 (November 2013), 1–27. DOI:<https://doi.org/10.1145/2503823>
- [7] Mark Blythe. 2014. Research through design fiction: Narrative in real and imaginary abstracts. In *Conference on Human Factors in Computing Systems - Proceedings*.  
DOI:<https://doi.org/10.1145/2556288.2557098>
- [8] Ash Buchanan. 2017. The nature of wellbeing. *Medium*. Retrieved from <https://medium.com/benefit-mindset/the-nature-of-wellbeing-53e7bcd592ac>
- [9] Mattia Giovanni Campana, Franca Delmastro, Dimitris Chatzopoulos, and Pan Hui. 2018. Lightweight modeling of user context combining physical and virtual sensor data. In *UbiComp/ISWC 2018 - Adjunct Proceedings of the 2018 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2018 ACM International Symposium on Wearable Computers*, Association for Computing Machinery, Inc, New York, New York, USA, 1309–1320. DOI:<https://doi.org/10.1145/3267305.3274178>
- [10] Eun Kyoung Choe, Nicole B. Lee, Bongshin Lee, Wanda Pratt, and Julie A. Kientz. 2014. Understanding quantified-selfers' practices in collecting and exploring personal data. In *Conference on Human Factors in Computing Systems - Proceedings*, Association for Computing Machinery, New York, USA, 1143–1152.  
DOI:<https://doi.org/10.1145/2556288.2557372>
- [11] Gabriele Civitarese, Riccardo Presotto, and Claudio Bettini. 2019. Demo: Hybrid data-driven and context-aware activity recognition with mobile devices. In *UbiComp/ISWC 2019 - 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*, Association for Computing Machinery, Inc, New York, New York, USA, 266–267. DOI:<https://doi.org/10.1145/3341162.3343844>
- [12] Paul Coulton, Joseph Lindley, and Haider Ali Akmal. 2016. Design Fiction: Does the search for plausibility lead to deception? In *DRS2016: Future-Focused Thinking*, Design Research Society.  
DOI:<https://doi.org/10.21606/drs.2016.148>
- [13] Paul Coulton, Joseph Lindley, Miriam Sturdee, and Mike Stead. 2017. Design Fiction as World Building. In *Research through Design (RTD)*, Edingburgh, United Kingdom.
- [14] Anind K. Dey. 2001. Understanding and using context. *Pers. Ubiquitous Comput.* 5, 1 (2001), 4–7.  
DOI:<https://doi.org/10.1007/s007790170019>
- [15] Paul Dourish. 2004. What we talk about when we talk about context. *Pers. Ubiquitous Comput.* 8, 1 (February 2004), 19–30. DOI:<https://doi.org/10.1007/s00779-003-0253-8>
- [16] Kevin Duhamel. 2014. Finding the sweet spot. *Occup. Heal. Saf.* July (2014).  
DOI:<https://doi.org/10.1145/3191761>
- [17] Paul Dulaud, Ines Di Loreto, and Denis Mottet. 2020. Self-Quantification Systems to Support Physical Activity: From Theory to Implementation Principles. *Int. J. Environ. Res. Public Health* 17, 24 (December 2020), 9350. DOI:<https://doi.org/10.3390/ijerph17249350>
- [18] Asma Ghandeharioun, Asaph Azaria, Sara Taylor, and Rosalind W. Picard. 2016. “Kind and Grateful”: A Context-Sensitive Smartphone App Utilizing Inspirational Content to Promote Gratitude. *Psychol. Well.*

- Being*, 6, 1 (December 2016), 1–21. DOI:<https://doi.org/10.1186/s13612-016-0046-2>
- [19] Samuel Gibbs. 2020. Withings ScanWatch review: health-tracking watch with 30-day battery. *The Guardian*. Retrieved from <https://www.theguardian.com/technology/2020/sep/14/withings-scanwatch-review-health-tracking-watch-with-30-day-battery>
- [20] Li Guo. *Quantified-Self 2.0: Using Context-Aware Services for Promoting Gradual Behaviour Change*.
- [21] Amelie Gyrard and Amit Sheth. 2020. IAMHAPPY: Towards an IoT knowledge-based cross-domain well-being recommendation system for everyday happiness. *Smart Heal.* 15, (March 2020). DOI:<https://doi.org/10.1016/j.smhl.2019.100083>
- [22] Donna Haraway. 1988. Situated Knowledges: The Science Question in Feminism and the Privilege of Partial Perspective. *Fem. Stud.* 14, 3 (1988), 575. DOI:<https://doi.org/10.2307/3178066>
- [23] Nils B. Heyen. 2020. From self-tracking to self-expertise: The production of self-related knowledge by doing personal science. *Public Underst. Sci.* 29, 2 (February 2020), 124–138. DOI:<https://doi.org/10.1177/0963662519888757>
- [24] ION. ION - Wearable Wellness. Retrieved October 30, 2020 from <https://www.ionwearable.com/ionwearable>
- [25] Marina Jirotko, Barbara Grimpe, Bernd Stahl, Grace Eden, and Mark Hartswood. 2017. Responsible research and innovation in the digital age. *Commun. ACM* 60, 5 (May 2017), 62–68. DOI:<https://doi.org/10.1145/3064940>
- [26] Anna Jobin, Marcello Ienca, and Effy Vayena. 2019. The global landscape of AI ethics guidelines. *Nat. Mach. Intell.* (2019). DOI:<https://doi.org/10.1038/s42256-019-0088-2>
- [27] Ben Kirman, Conor Linehan, Dan O’Hara, and Shaun Lawson. 2013. CHI and the Future Robot Enslavement of Humankind; A Retrospective. In *Conference on Human Factors in Computing Systems - Proceedings*, Association for Computing Machinery, New York, New York, USA, 2199–2208. DOI:<https://doi.org/10.1145/2468356.2468740>
- [28] Santosh Kumar. 2020. Sensitivity, Specificity, Generalizability, and Reusability Aspirations for Machine Learning (ML) Models in mHealth. Association for Computing Machinery (ACM), 1–1. DOI:<https://doi.org/10.1145/3396868.3402495>
- [29] James Larus, Chris Hankin, Siri Granum Carson, Markus Christen, Silvia Crafa, Oliver Grau, Claude Kirchner, Bran Knowles, Andrew McGettrick, Damian Andrew Tamburri, Sandra Wachter, and Hannes Werthner. 2018. *When Computers Decide: European Recommendations on Machine-Learned Automated Decision Making*. DOI:<https://doi.org/10.1145/3185595>
- [30] Jisoo Lee, Erin Walker, Winslow Burleson, Matthew Kay, Matthew Buman, and Eric B. Hekler. 2017. Self-experimentation for behavior change: Design and formative evaluation of two approaches. In *Conference on Human Factors in Computing Systems - Proceedings*. DOI:<https://doi.org/10.1145/3025453.3026038>
- [31] Ian Li, Anind K Dey, and Jodi Forlizzi. 2011. Understanding My Data, Myself: Supporting Self-Reflection with Ubicomp Technologies. In *UbiComp ’11: Proceedings of the 13th international conference on Ubiquitous computing*, ACM Press, New York, NY, USA, 405–414.
- [32] Peng Liao, Kristjan Greenewald, Susan Murphy, and Predrag Klasnja. 2020. Personalized HeartSteps: A Reinforcement Learning Algorithm for Optimizing Physical Activity. 18, (2020), 22. DOI:<https://doi.org/10.1145/3381007>
- [33] Joseph Lindley, Haider Ali Akmal, Franziska Pilling, and Paul Coulton. 2020. Researching AI Legibility through Design. Association for Computing Machinery (ACM), 1–13. DOI:<https://doi.org/10.1145/3313831.3376792>
- [34] Joseph Lindley and Paul Coulton. 2015. Game of drones. In *CHI PLAY 2015 - Proceedings of the 2015 Annual Symposium on Computer-Human Interaction in Play*, Association for Computing Machinery, Inc, New York, New York, USA, 613–618. DOI:<https://doi.org/10.1145/2793107.2810300>

- [35] Joseph Lindley and Paul Coulton. 2016. Pushing the limits of design fiction: The case for fictional research papers. In *Conference on Human Factors in Computing Systems - Proceedings*, Association for Computing Machinery, 4032–4043. DOI:<https://doi.org/10.1145/2858036.2858446>
- [36] Joseph Lindley, Paul Coulton, and Rachel Cooper. 2019. The IoT and unpacking the Heffalump’s trunk. In *Advances in Intelligent Systems and Computing*, Springer Verlag, 134–151. DOI:[https://doi.org/10.1007/978-3-030-02686-8\\_11](https://doi.org/10.1007/978-3-030-02686-8_11)
- [37] Philip Moore. 2017. Do We Understand the Relationship between Affective Computing, Emotion and Context-Awareness? *Machines* 5, 3 (July 2017), 16. DOI:<https://doi.org/10.3390/machines5030016>
- [38] George Wamamu Musumba and Henry O. Nyongesa. 2013. Context awareness in mobile computing: A review. *Int. J. Mach. Learn. Appl.* 2, 1 (May 2013). DOI:<https://doi.org/10.4102/ijmla.v2i1.5>
- [39] Nathalia Nascimento, Paulo Alencar, David R Cheriton, Carlos Lucena, and Donald Cowan. 2018. A Context-Aware Machine Learning-based Approach. DOI:<https://doi.org/10.5555/3291291.3291297>
- [40] Benjamin W. Nelson, Carissa A. Low, Nicholas Jacobson, Patricia Areán, John Torous, and Nicholas B. Allen. 2020. Guidelines for wrist-worn consumer wearable assessment of heart rate in biobehavioral research. *npj Digit. Med.* 3, 1 (December 2020), 1–9. DOI:<https://doi.org/10.1038/s41746-020-0297-4>
- [41] Kurt Plarre, Andrew Raji, Syed Monowar Hossain, Amin Ahsan Ali, Motohiro Nakajima, Mustafa Al’Absi, Emre Ertin, Thomas Kamarck, Santosh Kumar, Marcia Scott, Daniel Siewiorek, Asim Smailagic, and Lorentz E. Wittmers. 2011. Continuous inference of psychological stress from sensory measurements collected in the natural environment. In *Proceedings of the 10th ACM/IEEE International Conference on Information Processing in Sensor Networks, IPSN’11*.
- [42] Mashfiqui Rabbi, Angela Pfammatter, I. Zhang, Bonnie Spring, and Tanzeem Choudhury. 2015. Automated Personalized Feedback for Physical Activity and Dietary Behavior Change with Mobile Phones: A Randomized Controlled Trial on Adults. *JMIR mHealth uHealth* 3, 2 (June 2015). DOI:<https://doi.org/10.2196/mhealth.4160>
- [43] Paul Rice, Sayali Upasham, Badrinath Jagannath, Roshan Manuel, Madhavi Pali, and Shalini Prasad. 2019. CortiWatch: watch-based cortisol tracker. *Futur. Sci. OA* 5, 9 (October 2019), FSO416. DOI:<https://doi.org/10.2144/fsoa-2019-0061>
- [44] Adrienne So. 2020. Review: Fitbit Sense. *Wired*. Retrieved from <https://www.wired.com/review/fitbit-sense-watch/>
- [45] Lucy A Suchman. 1987. *Plans and situated actions: The problem of human-machine communication*. Cambridge university press.
- [46] Li-Chia Tai, Wei Gao, Minghan Chao, Mallika Bariya, Quynh P Ngo, Ziba Shahpar, Hnin Y Y Nyein, Hyejin Park, Junfeng Sun, Younsu Jung, Eric Wu, Hossain M Fahad, Der-Hsien Lien, Hiroki Ota, Gyoujin Cho, and Ali Javey. 2018. Methylxanthine Drug Monitoring with Wearable Sweat Sensors. *Adv. Mater.* 30, 23 (2018), 1707442. DOI:<https://doi.org/10.1002/adma.201707442>
- [47] Marvin M. Theimer and Bill N. Schilit. 1994. Disseminating Active Map Information to Mobile Hosts. *IEEE Netw.* 8, 5 (1994), 22–32. DOI:<https://doi.org/10.1109/65.313011>
- [48] Peter Tolmie, Andy Crabtree, Tom Rodden, James Colley, and Ewa Luger. 2016. This has to be the cats - Personal data legibility in networked sensing systems. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work, CSCW*. DOI:<https://doi.org/10.1145/2818048.2819992>
- [49] Longqi Yang, Jenny Chen, Nicola Dell, Michael Sobolev, Drew Dunne, Mor Naaman, Yu Wang, Christina Tsangouri, and Deborah Estrin. 2019. How intention informed recommendations modulate choices: A field study of spoken word content. In *The Web Conference 2019 - Proceedings of the World Wide Web Conference, WWW 2019*, Association for Computing Machinery, Inc, New York, New York, USA, 2169–2180. DOI:<https://doi.org/10.1145/3308558.3313540>
- [50] Qian Yang, Nikola Banovic, and John Zimmerman. 2018. Mapping machine learning advances from HCI

research to reveal starting places for design innovation. In *Conference on Human Factors in Computing Systems - Proceedings*, Association for Computing Machinery.

DOI:<https://doi.org/10.1145/3173574.3173704>