Implicit Interaction Through Machine Learning: Challenges in Design, Accountability, and Privacy

Donald McMillan(✉)

Stockholm University, Stockholm, Sweden
donald.mcmillan@dsv.su.se

Abstract. Implicit Interaction takes advantage of the rise of predictive algorithms, trained on our behaviour over weeks, months and years, and employs them to streamline our interactions with devices from smartphones to Internet connected appliances. Implicit Interaction provides users the advantage of systems that learn from their actions, while giving them the feedback and controls necessary to both understand and influence system behaviour without having to rely on an application for every connected device. This is an active area of research and as such presents challenges for interaction design due, in part, to the use of user-facing machine learning algorithms. This paper discusses the challenges posed by designing in accountability for system actions and predictions, the privacy concerns raised by both the sensing necessary to power these predictions and in how the predictions and systems actions themselves can expose behavioural patterns, and the challenges inherent in designing for the reality of machine learning techniques rather than the hype.

Keywords: Implicit interaction · Internet of things · Machine learning · Privacy · Interaction design

1 Introduction

The Internet of Things (IoT) presents several visions of the future. One, which extrapolates from current practice of interacting with nascent IoT technology, is of providing apps for everything [21] to be installed, updated, and learned alongside the tasks of installing, updating, and learning to use the IoT devices themselves. A competing vision is that of Implicit Interaction.

Implicit interactions stay in the background, thriving on analysis of speech, movement, and other contextual data, avoiding unnecessarily disturbing us or grabbing our attention. When we turn to them, depending on context and functionality, they either shift into an explicit interaction-engaging us in a classical interaction dialogue informed by the analysis of the context at hand, alternatively they continue to engage us implicitly using entirely different modalities that do not require an explicit dialogue that is through the way we move or engage in other tasks, the smart objects respond to us. For example, one form of implicit interaction is when mobile phones listen to surrounding conversation and continuously adapt to what might be a relevant starting

I. Kompatsiaris et al. (Eds.): INSCI 2017, LNCS 10673, pp. 352–358, 2017.
https://doi.org/10.1007/978-3-319-70284-1_27

point once the user decides to turn to it [17]. As the user activates the mobile, we can imagine how the search app already has search terms from the conversation inserted, the map app shows places discussed in the conversation, or if the weather was mentioned and the person with the mobile was located in their garden, the gardening app may have integrated the weather information with the sensor data from the humidity sensor in your garden to provide a relevant starting point. This is of course only possible through providing massive data sets and making continuous adaptations to what people say, their indoor and outdoor location, their movements and any smart objects in that environment – thriving off the whole ecology of artefacts, people and their practices. The implicit interaction paradigm, however, presents unique challenges when dealing with accountability, privacy, and control. This paper discusses each of these in turn, and sets up a series of challenges for researchers in this field.

2 Background

Implicit Interaction builds on the history of intelligent agents [20], behavioural inference [12] and motion sensing [2]. Implicit interactions stay in the background [6, 13, 18], thriving on data analysis of speech [17] and movements [16]. On a trajectory towards this vision, HCI researchers have used machine learning in interaction in a variety of, from turning the body into an interactive interface [10] to creating adaptive interfaces that automate and facilitate user's tasks reduce [15]. Machine learning has been used to understand the routines of users [5] including their interruptability during certain tasks [8, 9], in order to better support interaction. Better understanding the user themselves, rather than their routines, has been used to provide systems that can detect depression [7], recommend products and services [11], and produce intelligent tutors [1].

However designing with machine learning is an ongoing challenge and area of research. Dove *et al.* surveyed UX practitioners [16] noting problems in understanding the capabilities of Machine Learning, and integrating it into design practices such as prototyping. Systems based on machine learning may be unable to show understanding of users' intent, leading them to be perceived as useless and unintuitive [22], and in studies of self-driving vehicles human behaviour has been shown to be responsive to real and perceived levels of control [4].

3 Privacy and Accountability

The vision of Implicit Interaction is built on collection and training on large amounts of personal data, and making the sensing and collection of this data happen in a privacy sensitive manner is a challenge facing any number of fields concerned with human subjects looking to harness the growing power of machine learning. The recent court case in the USA surrounding data collected, possibly inadvertently, by the Amazon Echo home assistant [19] highlights one issue with collecting large amounts of data not explicitly directed at controlling the devices provided. Building on Privacy by Design recommendations [14], we can suggest that systems should not store raw sensor data, such as the unrecognized audio recordings under scrutiny in the Echo case above, but this in turn causes problems for the future of said systems. The cases where the system was unable to understand the users' intention are the cases where the raw data can provide valuable input to improve the algorithm, not training the system using these pieces of data would require significantly more effort in collecting and maintaining a training data set. Even without the raw sensor data, information about the actions of users could be retrieved by interrogating the machine learning algorithm. By comparing the output of the underlying algorithm taken in its base state and the trained state for a particular user when presented with different inputs it would be possible to provide a probability that a certain input pattern had been presented to that algorithm. While this could not provide someone with new data it holds the possibility that such techniques would allow law enforcement, or other actors, to determine with some certainty if a certain action had been taken or phrase had been spoken in a location.

More than just the data, the actions of the system can cause problems with privacy. The actions are a reflection of the dominant behavior of the user or users, and in a system where the goal is to pre-empt users to adjust the environment to them without any explicit commands this could expose behavior that the user would rather keep private. The precursor to this problem has already been seen, when Target (a large retailer in the USA) through their loyalty card system predicted a teen's pregnancy and, with targeted adverts and discounts, inadvertently informed the new grandfather before his daughter had informed him that she was pregnant[1]. This calls for research into how best to give control of the learning function of the algorithms that will watch our everyday lives back to the users being watched, and it is two fold. On one hand, users would benefit from the ability to pause the learning of systems in their home in unusual situations (such as renting their home to strangers on AirBnB) or in situations they would not want reflected in the presentation of self that the automatic actions of a system trained on their actions would become. On the other hand, research is needed in ways to allow users to interrogate algorithms and make the current actions of the system accountable to past actions of the users. Not only would this allow 'explaining away' embarrassing situations caused by such systems, but it would provide understanding of the effect their actions have on the learning algorithm necessary to consciously change their behavior to result in the system behaving in a manner they choose.

As the learning algorithms watch user activity over time, finding and reinforcing predictive correlations between behaviour and the user's interaction with the actuators under the control of the system, understanding why the system behaves as is does can require a level of reflection not necessary in short-loop interactions. The challenge here is not only providing an understandable representation of the progression of training in relation to the recorded actions of the users, it also requires research to extract and understand the current state of training of the learning algorithm which in itself is a current topic of research. Indeed, the 'black box' nature of many deep learning algorithms hinders the accountability of the algorithms significantly. In the vision of Implicit Interaction there would be cascade of learning algorithms, with those closest to

[1] https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/.

the user detecting simple individual actions (such as room occupancy, wakefulness, or low level activities like reading a book), which would in turn be used as input to those dedicated to higher level understanding of activity over time, and ideally these would both be inputs into a further stage of machine learning algorithms looking at user intent and supporting higher level longer term goals for user engagement.

In such a tiered system, it would be possible to store and display the inputs that had been detected by the lower level algorithms alongside the results of the higher-level ones, providing at least some semblance of meaning and accountability to system actions. However, before research into making the underlying algorithms accountable bears fruit this will provide only a small window through which the user must make sense of the actions of the system and how they can influence it.

4 Implicit Control

Controlling a system via a machine learning algorithm will present a number of challenges in itself. In challenging the accountability problem above, the users must be informed through design of the possibilities and limitations with machine learning algorithms, and before this can be attempted designers must develop the understanding necessary to design control mechanisms that fit within these limitations. This presents a challenge as the current narrative on machine learning is described in the Gartner Hype Cycle[2] as being at the "Peak of Inflated Expectations" – meaning that many of the assumptions about machine learning disseminated outside the community of practitioners and researchers working directly with these algorithms can be taken with a pinch of salt.

The most important 'Inflated Expectation' sold in relation to work on Implicit Interaction is that given 'enough' data a system could learn to understand rare and infrequent contexts of human behaviour. Machine learning excels in categorising complex data into common and relatively balanced categories, but when the data ventures into the realm of 'Imbalanced Domains' [3] there are a number of extra challenges in developing algorithms that can effectively categorise the data. There are two distinct challenges in imbalanced domains that should be understood by those designing for and with machine learning, but in order to explain them the default behaviour of a machine learning algorithm over a data set must also be described.

Take, for example, a data set consisting of 1 million tweets relating to a political party. A machine learning algorithm trained to determine if the tweets are positive or negative in sentiment would be shown a subset of these tweets categorised by an expert. Each one of these training tweets is fed into the algorithm and it provides its guess as to if the tweet is positive or negative. As this is during the training of the algorithm, the outcome is checked against the expert categorisation. If the machine matches the expert then the algorithm parameters (which differ from one method to the next) are reinforced, if the machine gets it wrong then the parameters are changed in the other direction. This is done over thousands and thousands of tweets with the overall

[2] http://www.gartner.com/newsroom/id/3412017.

accuracy a goal for the system. Carrying on this overly simplistic abstraction we can see that if the data set consisted of 20% of positive tweets and 80% negative tweets then the algorithm could return an 80% accuracy simply by categorising all data as negative – something that beginner's implementations of machine learning algorithms often do.

At this level there are any number of techniques, beyond simply looking in more detail at the type of error, to ensure that this doesn't happen in production systems yet the underlying problem is still the same when the number of categories increases and the size of the small categories decreases. The two problems with Imbalanced Domains, rephrased from [3], can be described as: (1) It is more important for the user to get accurate results from some categories in the data than it is from others, and (2) the cases that are more important to the user are under-represented in the data.

In the context of home automation, it may be more important to the user for the system to recognise rare occurrences accurately (a burglary or an acute illness, for example) at the expense of occasionally misrecognising the intention to open the window as closing the curtains.

Given enough information about the preferences of the user the choice of algorithm and the learning method can be adjusted to improve accuracy, but getting that information from the user is difficult. This is a major challenge for interaction design: Providing methods and metaphors to allow users to understand and influence the learning algorithms possibly by exposing the imbalances in the data in order to provide a counterbalance of user preferences.

5 Conclusion

This paper presents challenges faced in HCI by the advances in machine learning and their ongoing incorporation into user-facing systems. The Implicit Interaction project described in brief above is one of many research initiatives striving to bridge the gaps in knowledge and practice between users, machine learning experts, interaction designers, and the machine learning algorithms themselves. By increasing the awareness of the problems presented by the data that feeds the algorithms, the social and societal implications of the inferences that they make, and the disconnect between expectations of their abilities and the realities of implementing such algorithms this paper provides a starting point for more nuanced discussion of the ever increasing influence that machine learning algorithms have in our everyday lives.

References

1. Aleven, V., Mclaren, B.M., Sewall, J., Koedinger, K.R.: A new paradigm for intelligent tutoring systems: example-tracing tutors. Int. J. Artif. Intell. Ed. **19**(2), 105–154 (2009)
2. Bachynskyi, M., Palmas, G., Oulasvirta, A., Weinkauf, T.: Informing the design of novel input methods with muscle coactivation clustering. ACM Trans. Comput. Hum. Interact. (TOCHI) **21**(6), 30 (2015)

3. Branco, P., Torgo, L., Ribeiro, R.P.: A Survey of predictive modeling on imbalanced domains. ACM Comput. Surv. **49**(2), 31:1–31:50 (2016). doi:10.1145/2907070
4. Casner, S.M., Hutchins, E.L., Norman, D.: The challenges of partially automated driving. ACM Commun. **59**(5), 70–77 (2016). doi:10.1145/2830565
5. Davidoff, S., Ziebart, B.D., Zimmerman, J., Dey, A.K.: Learning patterns of pick-ups and drop-offs to support busy family coordination. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2011), pp. 1175–1184 (2011). https://doi.org/10.1145/1978942.1979119
6. Dix, A.: Beyond intention-pushing boundaries with incidental interaction. In: Proceedings of Building Bridges: Interdisciplinary Context-Sensitive Computing (2002). https://pdfs.semanticscholar.org/d0a8/28c43654b1e16527a92863732cb13189fba9.pdf. Accessed 10 June 2017
7. Doryab, A., Min, J.K., Wiese, J., Zimmerman, J., Hong, J.I.: Detection of behavior change in people with depression. In: Proceedings of AAAI Workshop on Modern Artificial Intelligence for Health Analytics (MAIHA) (2014). http://repository.cmu.edu/hcii/271
8. Fogarty, J., Hudson, S.E., Atkeson, C.G., Avrahami, D., Forlizzi, J., Kiesler, S., Lee, J.C., Yang, J.: Predicting human interruptibility with sensors. ACM Trans. Comput. Hum. Interact. **12**(1), 119–146 (2005). doi:10.1145/1057237.1057243
9. Fogarty, J., Hudson, S.E., Lai, J.: Examining the robustness of sensor-based statistical models of human interruptibility. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2004), pp. 207–214 (2004). https://doi.org/10.1145/985692.985719
10. Harrison, C., Tan, D., Morris, D.: Skinput: appropriating the body as an input surface. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2010), pp. 453–462, (2010). https://doi.org/10.1145/1753326.1753394
11. Herlocker, J.L., Konstan, J.A., Riedl, J.: Explaining collaborative filtering recommendations, pp. 241–250 (2000). https://doi.org/10.1145/358916.358995
12. Höök, K.: Steps to take before intelligent user interfaces become real. Interact. Comput. **12** (4), 409–426 (2000). doi:10.1016/S0953-5438(99)00006-5
13. Wendy, J., Leifer, L.: The design of implicit interactions: making interactive systems less obnoxious. Des. Issues **24**(3), 72–84 (2008). doi:10.1162/desi.2008.24.3.72
14. Langheinrich, M.: Privacy by design—principles of privacy-aware ubiquitous systems. In: Ubicomp 2001: Ubiquitous Computing, pp. 273–291 (2001). https://doi.org/10.1007/3-540-45427-6_23
15. Langley, P.: Machine learning for adaptive user interfaces. In: Brewka, G., Habel, C., Nebel, B. (eds.) KI 1997. LNCS, vol. 1303, pp. 53–62. Springer, Heidelberg (1997). doi:10.1007/3540634932_3
16. Lim, C.H., Vats, E., Chan, C.S.: Fuzzy human motion analysis: a review. Pattern Recogn. **48** (5), 1773–1796 (2015). doi:10.1016/j.patcog.2014.11.016
17. McMillan, D., Loriette, A., Brown, B.: Repurposing conversation: experiments with the continuous speech stream. In: Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI 2015), pp. 3953–3962 (2015). https://doi.org/10.1145/2702123.2702532
18. Schmidt, A.: Implicit human computer interaction through context. Pers. Technol. **4**(2–3), 191–199 (2000). doi:10.1007/BF01324126
19. Smith, I.: Amazon releases Echo data in murder case, dropping first amendment argument. PBS (2017). http://www.pbs.org/newshour/rundown/amazon-releases-echo-data-murder-case-dropping-first-amendment-argument/
20. Wooldridge, M., Jennings, N.R.: Intelligent agents: theory and practice. Knowl. Eng. Rev. **10**(2), 115–152 (1995). doi:10.1017/S0269888900008122

21. Xiao, R., Laput, G., Zhang, Y., Harrison, C.: Deus EM Machina: on-touch contextual functionality for smart IoT appliances. In: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI 2017), pp. 4000–4008 (2017). https://doi.org/10.1145/3025453.3025828
22. Yang, R., Newman, M.W.: Learning from a learning thermostat: lessons for intelligent systems for the home. In: Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing (UbiComp 2013), pp. 93–102 (2013). https://doi.org/10.1145/2493432.2493489