Private-by-Design Mobile Health Applications through On-Device Processing

Sandra Servia-Rodríguez
University of Cambridge
sandra.servia-rodriguez@cl.cam.ac.uk

Cecilia Mascolo
University of Cambridge
cecilia.mascolo@cl.cam.ac.uk

The large penetration of mobile and wearable devices, equipped with a growing set of built-in sensors, have fostered the design and development of many applications aimed at assisting people with their health and wellbeing. Applications for tracking and promoting physical activity, assisting patients suffering from depression or dementia, or aiding smokers who want to quit are some example apps that have benefited from the pervasiveness and wealth of sensing capabilities of these devices.

Just-In-Time Adaptive Intervention (*JITAI*) is an emerging mobile phone intervention design based on providing the appropriate and tailored support so as to prevent negative health outcomes and promote healthy behaviours, at the right time, and only when needed. Self-reported health-related data, as well as sensing data, play a critical role on intervention triggering and delivering on *JITAI*s. Machine learning techniques can also contribute to develop sophisticated *JITAI*s that predict users' behaviour, identify the optimal time to intervene, and automatically select the content to deliver by tailoring it to the user's context.

Beyond the sensing capabilities of current wearables and smartphones, these devices come equipped with progressively more powerful CPUs, GPUs and DSPs. At the same time, machine learning libraries are being adapted and optimised to effectively work on mobile devices. Although the technology is ready to offer alternatives to cloud offloading by providing efficient on-device computation, most of the existing mobile health apps for *JITAI*s only use mobile devices to collect data from the available sensors and self-reports, and to display supporting messages, notifications and interventions. Almost all the processing of the data collected, including part or all the pre-processing, model training and message formatting, is performed on remote servers on the Internet. While this might have uses for aggregated statistics purposes and more precise user modelling, once the data leaves the device, all sorts of privacy violations become possible on this very sensitive collection.

Several issues need to be addressed to accelerate the transition from cloud-based to on-device applications. Firstly, the *resource constraints* of mobile devices compared to large clusters of servers dictates that the complexity of the machine learning models needs to be reduced in order to optimise battery consumption and storage. Some research efforts are already effectively addressing the problem of how to compress models to work on mobile devices[1,2]. Secondly, *data decentralisation* makes it difficult to train machine learning models that decide on the need of an intervention. Differentially-private distributed training[3], that works by trading accuracy for privacy, has been proposed as an effective solution. Also, for some applications, personalised models trained solely using single user data are more accurate than those trained on data from a population. Thirdly, data decentralisation also makes it difficult to monitor the proper functioning of the application, the efficacy of the intervention and the users' engagement. Techniques to aggregate and anonymise data before being sent to the server, such as the randomised response mechanism, are key to avoid data breaches. Fourthly, *battery* efficiency and *latency*, are critical to the success of any mobile phone app. Contrary to popular belief, previous research demonstrated that cloud offloading is not always the best solution to minimise devices' battery consumption nor latency[4]. This is especially true when the quality of the connectivity is poor, which would also facilitate the deployment of mobile health applications in low-income communities and countries where Internet connectivity is limited.

---

[1] Lane, N.D., Bhattacharya, S., Georgiev, P., Forlivesi, C., Jiao, L., Qendro, L., and Kawsar, F. Deepx: A software accelerator for low-power deep learning inference on mobile devices. In the Int. Conf. on Information Processing in Sensor Networks. (IPSN 2016)

[2] Wu, J., Leng, C., Wang, Y., Hu, Q., and Cheng, J. Quantized convolutional neural networks for mobile devices. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (CVPR 2016)

[3] Shokri, R. and Shmatikov, V. Privacy-preserving deep learning. In Proceedings of the 22nd ACM SIGSAC conference on computer and communications security (CCS 2015)

[4] Cuervo, E., Balasubramanian, A., Cho, D., Wolman, A., Saroiu, S., Chandra, R., and Bahl, P. MAUI: making smartphones last longer with code offload. In Proceedings of the 8th international conference on Mobile systems, applications, and services. (MobiSys 2010)