# Embracing Shifting Trends and Reviving Smartphone Sensing

Veljko Pejović, *Lancaster University, UK, University of Ljubljana and Jožef Stefan Institute, Slovenia*

*Abstract—Probably the most popular research platform originally never intended to become one, the smartphone has been instrumental in ushering us into the new era of studying human behavior. Even so, however, our analysis of research papers published in the last decade shows that the smartphone is steadily losing its prominence in mobile sensing research. The loss of a platform that is carried by six billion users at almost all times could stifle research in areas varying from human-computer interaction, over healthcare, to demography. Therefore, in this article I investigate the potential reasons behind smartphone sensing falling out of research fashion, and propose solutions for some key issues identified.*

Fifteen years have passed since the first Android and iOS smartphones hit the shelves and marked the culmination of long-standing efforts to realise a century-old prophecy of visionaries, such as Nikola Tesla[1], Isaac Asimov[2], and Arthur Clarke[3].

There is little doubt that smartphones have indeed restructured our society – it is enough to compare the way we navigate in space, make payments, watch the news, take photos and share them with social contacts today and fifteen years ago to observe the scale of the impact. What's more, there are no indications that the impact is waning. The initial opportunities enabled by sight-sound communication and ubiqutious information delivery are as of lately supplemented by the promises of artificial intelligence (AI), converting our pocket devices into a true "huge brain" that Tesla envisaged.

Researchers from different domains were quick to embrace the smartphone. The device was hailed for its potential to provide objective information about the surroundings thanks to multimodal sensors, to maintain a dialogue with its owner by being at an arm's reach at all times, and to efficiently store, process, and transfer data – all with essentially free rein that developers had when building their mobile applications. Interwoven with everyday life, the smartphone even served as a foundation for novel research fields, such as the computational social science [1].

A decade and a half later, smartphone hardware is even more powerful, the array of embedded sensors keeps growing, and the device is as pervasive as it ever was, partly due to wireless connectivity expansion. Application-wise, there is practically no aspect of human behavior left untouched. New avenues for exploration are only unfolding, primarily thanks to on-device AI.

Considering the above, it seems certain that the smartphone's role as the most promising platform for mobile computing research is unlikely to be threatened anytime soon. Yet, observing the research trends uncovered in the succeeding section of this article, it


Digital Object Identifier 10.1109/XXX.0000.0000000

[1]In a 1926 interview with Collier's magazine Tesla proclaimed that “When wireless is perfectly applied the whole earth will be converted into a huge brain, which in fact it is, all things being particles of a real and rhythmic whole. We shall be able to communicate with one another instantly, irrespective of distance. Not only this, but through television and telephony we shall see and hear one another as perfectly as though we were face to face, despite intervening distances of thousands of miles; and the instruments through which we shall be able to do his will be amazingly simple compared with our present telephone. A man will be able to carry one in his vest pocket.”

[2]In *Visit to the World's Fair of 2014*, published in 1964, Asimov ponders: “Communications will become sight-sound and you will see as well as hear the person you telephone. The screen can be used not only to see the people you call but also for studying documents and photographs and reading passages from books. The appliances of 2014 will have no electric cords, of course, for they will be powered by long-lived batteries"

[3]At a 1976 conference at MIT Clarke claimed that “The wristwatch telephone will be technologically feasible very soon. It will be completely mobile. And this would again restructure society. You'll tell the machine I'm interested in such-and-such item, sports, politics and so forth, and the machine will hunt the main central library and bring all this to you selectively.”

appears that the smartphone might slowly wane as a research tool. Why would that be?

## TRENDS

To investigate the role of the smartphone in mobile computing research, in this section I chart the prevalence of mobile platforms as a researchers' tool of choice in publications related to mobile computing. While numerous journals and conferences host such publications, journals typically have long publication cycles, preventing timely observation of recent trends, which is why I will here focus on conferences. Research utilising mobile platforms can be found in multitrack conferences, such as ACM UbiComp and ACM CHI, yet trends in platform usage can be difficult to disentangle within the broader set of topics covered by these conferences. Thus, I turn our attention to single track conferences, and among them focus on Association of Computing Machinery International Conference on Mobile Systems, Applications, and Services (ACM MobiSys) and Conference on Embedded Networked Sensor Systems (ACM SenSys). These conferences represents top-ranked[4] conferences on mobile and sensing systems research, and ones that have embraced the smartphone early on. Furthermore, according to its call for papers, MobiSys "values technical contributions with working implementations and practical evaluations", making any use of the smartphone for research clearly identifiable.

In this analysis, I examine 376 full-length research papers published in MobiSys and 321 such papers published in SenSys proceedings from 2013 to 2023. 2013 was the year when the smartphone was reported to have already reached a substantial chunk of the world's general population, with more than two billion active devices worldwide[5]. iPhone released model 5s, the first smartphone to feature a 64-bit processor and a specialized sensor data co-processor, while the Android market was dominated by similarly powerful, yet cheaper Samsung Galaxy line of devices. In addition, by 2013, researchers had fully acknowledged the smartphone's research potential, and seminal work, including the Nokia's Mobile Data Challenge had already appeared [2] , thus I assume that the selected time span will include the peak popularity of the smartphone as a research platform. From the corpus of papers, I separately count those that focus on Android, iPhone, and Windows phone-based experimentation. While the coding was in most cases straightforward, I exclude papers that use the smartphone merely as a result reporting platform (e.g. an Android app is present, but its only purpose is to show information to a user and the research itself is not inherently smartphone-specific), yet, include research that is prototyped as an Android mobile sensing app, but is deployed on tablets.

Figure 1 depicts the percentage of MobiSys and SenSys papers whose research is focused on smartphones. Android's clear dominance over other mobile operating systems (OSs) can be attributed to its open-source codebase, large developers' community, and compatibility with various hardware. However, we see

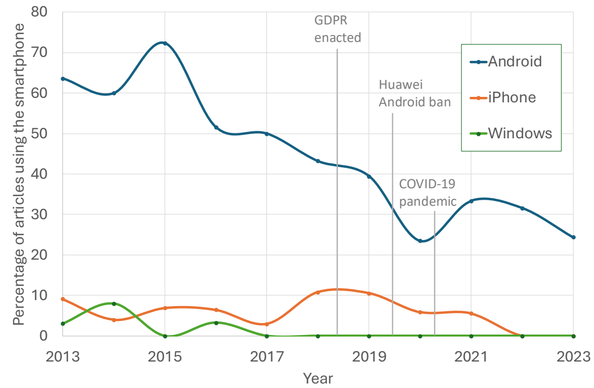


(a) Percentage of ACM MobiSys papers using a particular smartphone platform.

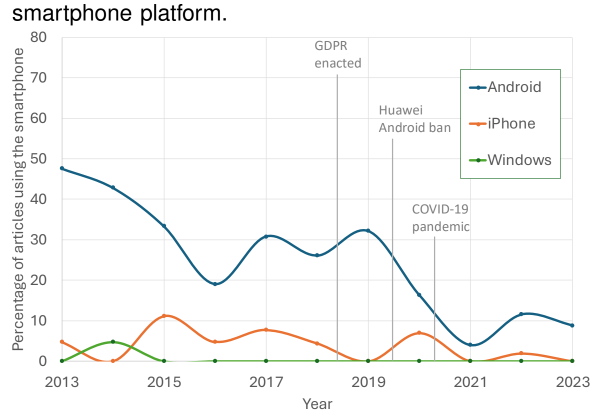


(b) Percentage of ACM SenSys papers using a particular smartphone platform.

FIGURE 1: Smartphone as a research tool of choice in ACM MobiSys and SenSys proceedings. Android used to figure in almost three quarters of the MobiSys and almost a half of SenSys research papers. A decade later, this has changed to less than a quarter and less than a tenth of papers, respectively. Meanwhile, Windows Mobile and iPhone have failed to gain much traction.

[4] CORE ranking *A* and *A**, respectively.

[5] https://www.statista.com/forecasts/1143723/smartphone-users-in-the-world

that over time, the popularity of both Android and the smartphone in general has fallen. By 2023 Android was the tool of choice in only 24% of the papers published at MobiSys, a two-third reduction over its peak in 2015, when Android was used by 72% of the MobiSys papers. Similar trends can be observed in SenSys, where the percentage of papers relying on Android sensing has fallen to 8% in 2023, c.f. 48% in 2013.

The figure also shows external events that may have impacted Android's popularity within the research community, such as its restricted use on Huawei devices, the enactment of GDPR, and the COVID-19 pandemics that made smartphone-based experimentation somewhat more difficult. Nevertheless, the popularity appears to be dwindling more persistently and long before/after these events.

The arrival of new platforms, such as earables, GPU-equipped boards, and unmanned aerial vehicles, can only partly explain the shifting research grounds, as each year's proceedings feature only a few papers using these new platforms[6]. Instead, research has also spread over platforms whose mainstream use predates the smartphone dominance, including FPGAs, microcontrollers, and software-defined radio. Therefore, I seek to answer whether changes in the smartphone itself made it a less attractive research tool over time?

## AFFORDANCES

To pinpoint the factors that have over time made the smartphone less attractive as a research tool, I will first examine the affordances that made it so desirable to begin with.

A highly-cited paper from 2012 hails the arrival of a platform that is compact and, for a piece of technology, unusually intimate and personalized [1]. The 2012 smartphone comes embedded with a range of sensors, from cameras that could be "continually streaming users' visual experiences into their smartphone memories, in case they want to share a clip with friends, family, Facebook, insurers, or police", to Bluetooth chipsets that, in combination with call logs, could reveal presence of a user's social contacts. The 2012 smartphone also allows apps to track other apps and the content of the communications, such as call logs, voice calls, and text messages, a user's interaction with the device, as well as the Web browsing history and online social network activity. For the near future, year 2025, the author envisions an ecosystem of external sensors that, together with the smartphone will enable monitoring various physiological signals, such as blood pressure, hormone levels, even ovulation. From the processing side, the 2012 smartphone already boasts multicore CPUs and gigabytes of RAM, and the author assures us that market competition among mobile operating systems will streamline the development of flexible apps that "can reach deeper into the guts of the smartphone's hardware".

Two technical affordances unique to the smartphone can be unpacked from the above description. First, the smartphone possesses immense sensing capabilities. Second, sensor sampling and data processing can be executed continuously and allow behavior tracking, thanks to the smartphone's background processing abilities. Have these affordances changed throughout the years?

## SENSING

The history of the smartphone reveals that the original intention behind embedded sensors was to facilitate interaction with the device, not provide a platform for mobile sensing. For instance, the purpose of a built-in inertial measurement unit (IMU) is to trigger screen transition to the landscape mode when a device is rotated, not to detect Parkinson's disease symptoms in the phone's user, as has been done later by Arora et al [3]. Since ingenious uses of sensors were borne out of later efforts by researchers and app developers, most sensors were not optimized for frequent use in the first place, and their energy consumption was, at least initially, rather high [4].

Substantial energy burden associated with sensor sampling was tackled by OS designers, who, over time, restricted and optimized access to sensor hardware, in order to expand the battery life of a device. The evolution of such restrictions is best observed in the case of location sensing in Android.

The initial means of accessing location in Android was through `LocationManager` class from `android.location` package. Methods of this class allow an application developer to select between GPS and network-based location determination. The latter is obtained by triangulation from signals of nearby WiFi and cellular base stations, and is often less accurate than GPS-based location determination. The GPS chip, on the other hand, remains one of the most energy hungry sensors in a smartphone [4]. While,

[6] I present a detailed breakdown of platforms used in MobiSys research from 2013 to 2023 available at https://gitlab.fri.uni-lj.si/lrk/smartphone-research-trends. The highest point of popularity of a new platform is being seen in 2023 when five MobiSys and five SenSys papers (12% and 21% of the total, respectively) presented research on GPU-equipped boards.

from the battery consumption point of view, occasional GPS sampling by a single app may go unnoticed, over time the average number of apps installed on a single user's phone grew to about 80 [5], making it highly probable that multiple applications frequently turn the GPS chipset on to obtain accurate location.

In Android smartphones, severe battery drain caused by location sensing was addressed with the introduction of `FusedLocationClient`, a class that abstracts the actual sensor used and only allows the developer to issue a query specifying whether the emphasis should be on location precision or energy needed for obtaining the location. Managed through Google Play Services, `FusedLocationClient` would then, whenever possible, "recycle" location information obtained and cached after a previous query. Consequently, an app requesting location information soon after another app has obtained the location, might be served "stale" location with virtually no battery cost incurred. While this improves the battery life, it abstracts the location access from the developer, potentially affecting mobile sensing apps.

The energy consumption overhead of location sampling was not the only issue with smartphone sensing whose prominence grew with time. Early mobile sensing apps collected sensitive data, such as raw sound and video clips and location traces, in a manner that would be completely unacceptable today due to privacy concerns. Once privacy awareness increased, mobile OSs responded with a suite of privacy control mechanisms. These mechanisms are, however, often in conflict with the needs of mobile sensing apps used for research. In early versions of Android, permissions to sample certain sensors, for instance a microphone or a camera, had to be granted at the app install time and were not retractable. Starting from Android 6.0 Marshmallow released in 2015, permissions had to be acquired at runtime, before the first time the permission-protected property was requested, making it possible for users to install the app, yet prevent the app from obtaining certain data, such as a user's location.

Considering location, further privacy-related restrictions were introduced in 2016 with Android 7.0 Nougat, where users were given an option to permit location sampling only when the app was actively used. This was additionally tightened in Android 10 Q, where background location sensing required an additional Manifest-level permission `ACCESS_BACKGROUND_LOCATION` and an explicit approval from the user. Finally, starting from Android 11, released in 2020, background location sensing can only be enabled from outside the app, through device-level settings that a lay user might not be familiar with.

Location sensing has witnessed the most radical permission overhaul, yet sensing other modalities faced similar curbing. Physical activity recognition, for instance, starting from Android 10 requires a separate permission. Camera and microphone access have been subject to permissions since the early days of Android, yet, starting from Android 9 P, even if given the permissions, apps can no longer obtain microphone and camera recordings from the background. Besides the sensing of physical properties, tighter privacy controls have also targeted the sensing of on-device events. Android 7.0 Nougat restricts sensing of certain events, such as connectivity changes. This is further tightened in Android 8.0 Oreo, where only a limited number of exempt event broadcasts can be registered. Cross-app privacy has also been enhanced, thus, it is not any more possible for an application to read others' notifications without a user explicitly granting this permission through the Settings. Similarly, obtaining information about CPU usage (via /proc/stat) is not possible since Android O.

In the era of surveillance capitalism [6], increased privacy-awareness and the resulting privacy-securing tools certainly represent a positive move, especially when it comes to a widely used personal device, such as the smartphone. However, increased restrictions on data collection severely hamper the smartphone's attractiveness as a research platform. A number of seminal smartphone-based studies from the last 15 years that have relied on unrestricted access to a user's location [7], physical activity [8], microphone recordings [9], and communication patterns [10] would be virtually impossible to conduct on today's devices.

## BACKGROUND PROCESSING

Predicting a user's depression from their mobility traces [7], or a student's grades from the places they visit and the conversations they have [11], would not be possible without the smartphone's ability to sense and process data even when the device is in one's pocket. Computation that is executed even when the app is not actively used, i.e. is in the background, was initially supported by Android's `Service` class. This is essentially a UI-less component that runs on the main computational thread. Later, the class evolved to `IntentService`, which automatically transfers the processing to a separate thread, and an additional class `AsyncTask`, that allow seamless transition between background and foreground processing. Depending on the interaction paradigm, the developers could use long-running `Services`, within them spawn a separate thread and task it with periodically collecting a user's sensor data in the background, or use an `IntentService` or an `AsyncTask` to collect data in the background, while the user is interacting with the UI.

One issue stemming from such a programming model is that a background task would keep a device's RAM occupied for extended periods of time, impacting the phone's performance. Consequently, the termination of long-running background Services, in case the system got low on resources, was implemented in Android. An alternative method to perform frequent data collection and processing, despite this limitation, was to periodically spawn a new `IntentService`, complete the data collection and processing, and then the background thread would terminate automatically. The component tasked with periodic initiation of the `IntentService` was `AlarmManager`. This class allowed one-off and periodic actions, even from within an app that is not currently active and/or interacted with.

Another issue with background processing is that it prevents the phone from transitioning to a power-saving mode, leading to battery depletion. Just as in the case of location sensing, while a periodic wake up of a single app initiated by `AlarmManager` is unlikely to have substantial impact on the battery use, frequent wake ups by a myriad of apps installed on the same phone will likely significantly reduce the battery charge and hamper the phone's usability. Project Volta, first introduced with Android 5.0 Lollipop in 2014, introduced an overhaul in the Android's energy management. The initial focus was on giving the developers a tool to optimize their apps. For instance, the Battery Historian tool would assist with identifying hardware features and processes consuming excessive energy, while `JobScheduler` would help schedule background tasks for times when a device is charging, so as to minimize the impact on the battery. Users, on the other hand, were provided with a switch for a Battery Saving mode, which dims the screen, pauses animations, and reduces background processing.

Except for the Battery Saving mode, however, the decision to save the energy remained in the developers' hands. The resulting energy savings were apparently insufficient, and thus, additional energy saving capabilities were introduced in 2015 with Android 6.0 Marshmellow. The most important in this case was the *Doze* mode, depicted in Figure 2: when in this mode, the phone would prevent network accesses, it would prevent any but high-priority notifications from arriving, and most critically for mobile sensing, it would prevent background processes from running and would

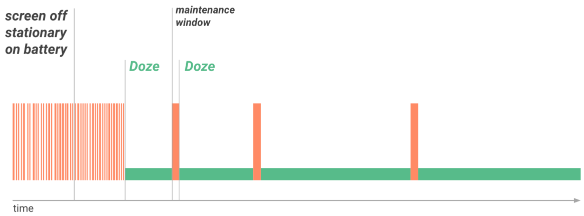


FIGURE 2: Android Doze mode. The OS puts the device in a low-power mode where network accesses, alarms, notifications, and processing are deferred until occasional *Maintenance* periods. (src: developer.android.com/training/monitoring-device-state/doze-standby)

defer certain alarms set through `AlarmManager`. This is not to say that apps would not be able to do any periodic background tasks – these tasks could still be performed during the so-called *Maintenance* periods. By forcing all apps to use common maintenance periods, the OS would ensure that more time is left for the phone to save energy by remaining in the power-saving mode. The transition to *Doze* mode became more aggressive over time. While initially limited to situations where the screen is off and the device is static, with Android 7.0 Nougat a switch to *Doze* mode would also happen when the phone was on the move.

*Doze* for the first time explicitly takes the energy optimization related to background tasks away from developers and delegates it to the OS. Here I should also mention *App Standby*. This mode is, in a way, an individual app's *Doze*. Infrequently used apps, subject to certain exceptions, are put in this mode, where their background and network activity is restricted.

The above changes had a profound effect on mobile sensing research apps. These applications often rely on periodic background sensing and processing, and would hardly ever provide enough added value to their users to become frequently used apps. Thus, they would in the best case be hampered by the *Doze* mode, which would interfere with the strict periodicity of the sensing envisioned by the researchers, and in the worst case, would end up in the lowest of the *App Standby* buckets (introduces in Android 9.0 Pie) reserved for the apps that are almost never opened, and on which significant restrictions in terms of background processing are imposed.

## ROAD AHEAD

The above is just a short excerpt of the changes that affected the feasibility of what used to be the standard mobile sensing practices in smartphones.

Digging deeper, a whole gamut of individual vendors' tweaks have imposed additional restrictions on background processing in particular[7]. Anyone trying to re-implement apps from the golden age of smartphone sensing is bound to face insurmountable obstacles in the shape of imprecise timing of task execution, restricted access to many sensors, and the need for substantial involvement from the user side. Several efforts, notably [12] and [13], provided practical solutions engineered to go around the mobile OS restrictions of that time, yet remained vulnerable to future restriction tightening. Thus, it is unsurprising that, according to my analysis of research trends, the most common reaction appears to be to simply give up on the smartphone sensing. Yet, the most pervasive personal computing device – and the uniquely versatile research tool, unlikely to be paralleled by any technology in the near future – should not be discarded before every attempt is made to salvage its former research appeal. Furthermore, with the loss of an easy-to-use data collection tool, we potentially lose the ability to reproduce past research findings. Thus, in the remainder of this article I ponder on a few ways in which smartphone sensing can be made compliant with the shifting trends of personal data perception and smartphone utility.

## RETHINKING THE DATA EXCHANGE PLAYFIELD

The balance of power between users, whose data is sensed, and companies or researchers, who would benefit from this data, is traditionally uneven, as the lack of social expectations, norms, and legislation related to personal data collection has provided a carte blanche for blanket data harvesting without the need to compensate the source. The perception of data ownership is changing, not in small part due to GDPR, CCPA, and similar regulations, which to smartphone users give more leverage when it comes to negotiations on how their data is used. This change in the balance of power opens opportunities for new data exchange paradigms.

Evidation Health[8], for instance, offers a platform that allows health researchers to obtain data collected by participants' smartphones and wearables. Evidation mobile app integrates mobile sensing, experience sampling method (ESM) querying, and sustains long-term data collection through a system of monetary rewards that users obtain for the data they consensually share. The financial burden of participation rewards falls on the research institutions using the data.

Behind the scenes, Evidation and similar frameworks require tremendous effort from the developers' side, since acquiring participants' data while complying to increasingly tighter sensing and background processing restrictions becomes highly challenging. However, the generality of the *data trading* concept may make it attractive for OS vendors to integrate such a marketplace in their products. As a result, the sensing related to this could be exempt from some of the restrictions detailed in the previous two sections. A drawback of such integration, however, is that it could potentially further entrench the role of OS vendors as data gatekeepers. Indeed, commodifying data sensing in the above manner could be based on, for instance, Apple's Health app, which already consolidates the smartphone – smartwatch sensing ecosystem.

At the other end, the idea of individuals retaining ownership of their data is not new, and a privacy-preserving architecture called Personal Data Vaults (PDVs) was originally proposed in 2010 [14]. The most likely reason for PDVs remaining in the realm of academia was the lack of incentive to actually give users control of their data. With the new era of data value and privacy awareness, we may expect this to change.

## EMBRACING PRIVACY

The pioneering mobile sensing app – CenceMe – involved activity sensing from the background, periodic sensing of the microphone, even a phone's camera taking photos at random moments and sending them to a server [15]. Even if reimbursed, it is improbable that today's smartphone users would be willing to supply such data. The humanity's awareness of privacy issues related to digital data has evolved, and the permissions that users have to grant to an app collecting the above data are the evidence of this.

One avenue towards more privacy-sensitive sensing is to simply acquire less sensitive data. For instance, less accurate location obtained by network triangulation is deemed less sensitive than accurate GPS location, and in certain situations can still be sufficient for answering research questions posed. Such modality substitution is not always possible, but is worth considering, as avoiding sensitive data sampling, among other things, increases the potential pool of research study participants.

Differential privacy (DP) enables privacy-preserving sharing of information pertaining to a group of users by

[7] https://dontkillmyapp.com/ keeps track of barriers to background processing on different device models and discusses possible circumnavigations around these barriers.

[8] https://evidation.com/

slightly perturbing individual data, yet preserving the properties of interest in the aggregated data [16]. DP could ensure privacy guarantees, while simultaneously providing the same value as the original data. Having in mind that DP is already used by Apple for QuickType predictive keyboard, for example, it is surprising that DP-based sensing is not readily available for iOS mobile app developers to use.

Privacy can be further protected by preventing the sensed data from leaving the device. For instance, instead of transferring raw GPS coordinates to the server for the analysis, researchers can construct data clustering models on the device and report back the processed information, such as the amount of time a user spent at home, with the location of the home remaining hidden[9]. Android is already providing such infrastructure for human activity recognition, as the built-in classifier internally samples the sensors and returns one of six activity categories without revealing raw sensor data at all.

Constructing machine learning (ML) models from the sensed data is often performed once data is aggregated on a server. Requiring that the data remains local renders the model construction much more challenging. One way of tackling this is federated learning (FL), a paradigm where ML models are constructed in a distributed manner by a group of devices, where each device trains a model using its own dataset. Only trained parameters of the model are aggregated at a server and no data ever leaves the device on which it was collected. FL is already harnessed by Google and Apple for training the Gboard predictive keyboard and Siri voice assistant, respectively. OS-level support for FL is currently missing, yet independent research efforts, such as Flower[10], already enable smartphone-based FL.

## TRADING SENSING FOR PROCESSING

The separation of data collection from data processing is evident in many mobile sensing studies conducted in the last 15 years [15], [8], [7]. While this is convenient, as such separation allows for different data processing methods to be applied and different research questions to be posed on the same dataset, it fails to minimize the collected data. Precision, temporal resolution, and the overall number of data samples can all be reduced, if data is collected with particular processing pipeline in mind. For example, on-device Kalman filtering could provide medium to long-term estimates of the target phenomenon while masking individual measurements [17].

Further avoidance of fine-graine sensing can be achieved with Compressive sensing (CS), a signal processing technique that allows the reconstruction of the original signal from samples taken at rates signficantly lower than the Nyquist rate [18]. In case of sparse signals, perfect reconstruction can be achieved with the number of samples proportional to $K \log(N/K)$, where $K$ is the number of non-zero components in the $N$-dimensional space of all possible signals. Many real-world signals are naturally sparse. For instance, human voice occupies only certain frequencies, pixels in a camera image usually convey a certain structure, etc. In practice, CS can reduce the number of samples needed for reconstruction by a few orders of magnitude, compared to conventional sampling. CS requires non-uniform sampling, thus it naturally fits the processing paradigm imposed by Android's *Doze* mode (Figure 2), where the uniformity of task execution, necessary for traditional sampling, is inherently broken. Furthermore, recently proposed techniques for high-level inference from CS data reduce not only the amount of data that is sampled, but also obviate the need for sending data to the server for signal reconstruction [19].

## MOBILIZING OS VENDORS' SUPPORT

Fine-grain in-context permission querying, location sensing aggregated across multiple apps, assuring that device sleep periods are uninterrupted, and similar features introduced in mobile OSs over the last fifteen years, have improved the smartphone experience for end-users. At the same time, as most of these features hamper the usability from the researchers' point of view, we should acknowledge that modern OSs have hugely benefited from findings brought by early experimental mobile sensing apps[11].

To continue the proliferation of the smartphone platform, OS vendors should ensure that the platform remains a viable tool research. Doing this without hurting the end-user experience or compromizing security is challenging, yet, many of the functionalities conceptualized in this section could coexist with features of a modern user-oriented OS. For instance, differential privacy-based sensing could be supported at API level and access to such sensing could be guarded by a new set of easier-to-obtain permissions; similarly, an

[9] I use this in the InterruptMe mobile app, for example [8]
[10] https://flower.dev

[11] The immediate impact of research on real-world ubiquitous computing products has been observed at least since the early nineties [20].

API call could enable researchers to supply functions that would be calculated on the phone, and only the processed result will be reported; periodic sampling compressive sensing-based sampling could also be supported by APIs that work in conjunction with the Doze mode; sensed data could be kept in PDVs and permissions would be needed to access the data, not initiate sensing, while on-device processing, e.g. in form of federated learning, could be subject to more lenient permissions.

Yet, even without significant changes, certain applications might be exempt from some of the OS restrictions. We have already witnessed this in 2020 with (Google/Apple) Exposure Notification (GAEN) framework[12]. GAEN, among other functionalities, enabled long-running periodic bluetooth scanning, something which is extremely difficult, if not impossible, to implement in a regular Android/iOS mobile application. GAEN runs as a system service on the smartphone with a special permission by Google/Apple, thus is not subject to the same restrictions as a regular app. Other apps can harness the service, only if authorized by Google/Apple. We can envision a similar arrangement for research apps – sensing can be controlled by system services and approved apps would get access to the data. This, however, just like GAEN, could be controversial, as deciding whether an app is indeed a legitimate research app puts tremendous power in OS vendors' hands, due to ethical, legal, political, and other consequences that such a decision may carry, especially since malicious spyware apps could find their way to users' phones.

Finally, Android already has Developer options (by default hidden) where the system behavior can be configured, so to facilitate application debugging. A similar Research mode could be implemented. When put in such a mode, the smartphone would be less restrictive on long-running background processing, would allow, for instance, background location sensing, and so on. While such a solution would not turn six billion phones to research devices, it could be useful for small-scale studies where researchers supply a dozen or so phones to study participants.

## CONCLUSIONS

For years the smartphone represented the most versatile, the most ubiquitous tool for mobile sensing. It served as a foundation for research in areas as diverse as natural resource conservation, earthquake monitoring, and transport planning.

However, the clash between a wider usability of the device and its ability to support (unrestricted) personal data acquisition has lead major OS vendors to gradually restrict programmability of the device. This has already been shown to have restricted the ease of use of the smartphone as a research tool [13], [12]. The analysis presented here indicates that, perhaps as a consequence, and in combination with external factors, such as tightened data collection legislation and COVID-19 pandemics that restricted the opportunities for smartphone-based user studies, the smartphone's predominance among research platforms in mobile computing[13] is waning. Yet, I believe that by revisiting the methods we use for sensing and processing the data, and through close collaboration of OS vendors, legal authorities, the research community, and users, whose privacy and data value will be taken into account, we can bring the smartphone its former glory and set grounds for further scientific advancements through mobile sensing.

## REFERENCES

1. G. Miller, "The smartphone psychology manifesto," *Perspectives on psychological science*, vol. 7, no. 3, pp. 221–237, 2012.
2. J. K. Laurila, D. Gatica-Perez, I. Aad, O. Bornet, T.-M.-T. Do, O. Dousse, J. Eberle, M. Miettinen *et al.*, "The mobile data challenge: Big data for mobile computing research," in *Proc. Mobile Data Challenge Workshop (MDC) in conjunction with Pervasive*, Newcastle, UK, 2012.
3. S. Arora, V. Venkataraman, A. Zhan, S. Donohue, K. M. Biglan, E. R. Dorsey, and M. A. Little, "Detecting and monitoring the symptoms of parkinson's disease using smartphones: A pilot study," *Parkinsonism & related disorders*, vol. 21, no. 6, pp. 650–653, 2015.
4. F. Ben Abdesslem, A. Phillips, and T. Henderson, "Less is more: Energy-efficient mobile sensing with senseless," in *Proceedings of the 1st ACM Workshop on Networking, Systems, and Applications for Mobile Handhelds*, Barcelona, Spain, August 2009.
5. P. Welke, I. Andone, K. Blaszkiewicz, and A. Markowetz, "Differentiating smartphone users by app usage," in *ACM UbiComp*, Heidelberg, Germany, September 2016.

[12] https://www.google.com/covid19/exposurenotifications/

[13] The scope of the analysis presented in this paper, however, prevents us from claiming that the popularity of the smartphone as a *data collection tool*, especially in social sciences [1], is waning.

6. S. Zuboff, *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. Profile books, 2019.
7. L. Canzian and M. Musolesi, "Trajectories of depression: unobtrusive monitoring of depressive states by means of smartphone mobility traces analysis," in *ACM UbiComp*, Osaka, Japan, September 2015.
8. V. Pejovic and M. Musolesi, "Interruptme: Designing intelligent prompting mechanisms for pervasive applications," in *ACM UbiComp*, Seattle, WA, USA, September 2014.
9. P. Zhou, Y. Zheng, and M. Li, "How long to wait? predicting bus arrival time with mobile phone based participatory sensing," in *ACM MobiSys*, Low Wood Bay, Lake District, United Kingdom, June 2012.
10. A. Madan, M. Cebrian, D. Lazer, and A. Pentland, "Social sensing for epidemiological behavior change," in *ACM UbiComp*, Copenhagen, Denmark, September 2010.
11. R. Wang, F. Chen, Z. Chen, T. Li, G. Harari, S. Tignor, X. Zhou, D. Ben-Zeev, and A. T. Campbell, "Studentlife: assessing mental health, academic performance and behavioral trends of college students using smartphones," in *ACM UbiComp*, Seattle, WA, USA, September 2014.
12. A. González-Pérez, M. Matey-Sanz, C. Granell, and S. Casteleyn, "Using mobile devices as scientific measurement instruments: reliable android task scheduling," *Pervasive and Mobile Computing*, vol. 81, p. 101550, 2022.
13. O. Petter, M. Hirsch, E. Mushtaq, P. Hevesi, and P. Lukowicz, "Crowdsensing under recent mobile platform background service restrictions: A practical approach," in *Workshop on Human Activity Sensing Corpus and Applications (with UbiComp)*, London, UK, September 2019.
14. M. Mun, S. Hao, N. Mishra, K. Shilton, J. Burke, D. Estrin, M. Hansen, and R. Govindan, "Personal data vaults: a locus of control for personal data streams," in *ACM CoNEXT*, Philadelphia, PA, USA, December 2010.
15. E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng, and A. T. Campbell, "Sensing meets mobile social networks: the design, implementation and evaluation of the cenceme application," in *ACM SenSys*, Raleigh NC USA, November 2008.
16. C. Dwork, "Differential privacy," in *Automata, Languages and Programming*, M. Bugliesi, B. Preneel, V. Sassone, and I. Wegener, Eds. Berlin, Germany: Springer Berlin Heidelberg, 2006, pp. 1–12.
17. Z. Chen, H. Zou, H. Jiang, Q. Zhu, Y. C. Soh, and L. Xie, "Fusion of wifi, smartphone sensors and landmarks using the kalman filter for indoor localization," *Sensors*, vol. 15, no. 1, pp. 715–732, 2015.
18. D. L. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
19. A. Machidon and V. Pejović, "Enabling resource-efficient edge intelligence with compressive sensing-based deep learning," in *ACM Computing Frontiers*, Turin, Italy, May 2022.
20. M. Weiser, "The computer for the 21st century," *Scientific American*, vol. 265, no. 3, pp. 94–104, September 1991.


## ACKNOWLEDGMENTS

This work was supported in part by Slovenian Research Agency Projects "Context-Aware On-Device Approximate Computing" (J2-3047), "eXtreme Learning on the Edge through Context-Aware Federation and Approximation" (N2-0393), and core funding P2-0098 and was conducted during the author's sabbatical at Lancaster University, UK. The author would like to thank Octavian Machidon for his comments on the early version of the manuscript.

**Veljko Pejović** is an associate professor at the Faculty of Computer and Information Science, University of Ljubljana, Slovenia. His research interests include mobile computing, HCI, and resource-efficient computing. Dr. Pejović received his Ph.D. degree in computer science from the University of California Santa Barbara, CA, USA. Contact him at veljko.pejovic@fri.uni-lj.si.