

Audio-based human health data gathering (automatic auscultation) through Machine Learning: a quantum leap in biometric data collection and analysis at scale.

Whitepaper prepared for Decorte Future Industries Ltd (12176817), by Erika Bondareva (CTO at Decorte Future Industries; PhD Machine Learning, University of Cambridge) and Dr Roeland Decorte (Founder and CEO at Decorte Future Industries).

1. Introduction

This whitepaper discusses the technology being developed by Decorte Future Industries Ltd, which to the authors' best knowledge is the first development of a commercial device-agnostic engine for automated human health data gathering through audio for holistic health analysis.

In the commercial sphere, audio-based biometric extraction has previously been deployed in a scattered way, generally in highly specific contexts to address specific conditions or pathologies. In the academic sphere, automated audio-based health data gathering has recently seen a boom in publications and studies proving its efficacy, but this has not yet filtered through to the commercial sphere. The technology being developed by Decorte Future Industries is inspired by, though independent from, the research of its members at the University of Cambridge.

In this whitepaper, by gathering both academic and commercial sources, we compare and contrast automated audio-based health data gathering to existing methodologies currently in operation in the commercial field.

The open-ended nature of automated audio-based health data gathering, powered by signal processing and machine learning, allows a single low-cost sensor (microphone) to collect a wealth of health metrics. This leads to a dramatic reduction in cost for multimodal biometric tracking, and therefore a quantum leap in humanity's ability to continuously collect human body data. The absence of the need for continuous skin contact – a prerequisite for, for example, ECG- and PPG-based methods widely utilised in the wearables space – further allows audio-based methodologies to overcome the main issues of lack of comfort and obtrusiveness. While ECG- and PPG-based methods can mostly be applied exclusively to a limited space of cardiovascular

metrics, thus representing a dead-end for multimodal sensing, audio allows continuous extraction and expansion of many different types of biometrics.

The contents of this whitepaper, as well as any specifics regarding any other aspect of the technology and methodology, are copyright and protected property of Decorte Future Industries. Please do not reproduce, copy or share without permission from the company.

2. Market need, and present absence of accurate human health data gathering through wearables

Despite the ongoing boom in wearables (Fig. 1), the current offering of consumer-focused wearable devices **consistently fails to yield meaningful health data**, whether they are wrist-worn devices, or clothing-based. This means that the user or a third party, such as a clinician, are generally unable to take action based upon this data. If a user commits to continuously wearing a wearable device, it is highly desirable that the collected data is “actionable” – meaning that it is of a sufficient quality and variety for the data to be useful to the wearer and third parties, such as clinicians, for body data analysis and/or health management.




Figure 1. A boom in wearables has led to what is essentially a doubling of sales between 2018-2019.

Smuck (University of Stanford) et al., in their 2021 review of “the emerging clinical role of wearables: factors for successful implementation in healthcare” (published in Nature’s Partner Journal, Digital Medicine), write that “[t]o date, **most commercially available wearables are**

limited in scope, tracking one or two health-related variables, and have yet to produce accurate measurement of many markers of health status that they attempt to assess.”¹

Most methods of biometric collection are focused on “one sensor-one metric”, or, at best, “one sensor-two metrics” systems. In order for the data to become actionable, there is a **need for a much greater amount of metrics and cross-referencing**. Scaling “one sensor-one metric” systems to achieve such a holistic overview of health, however, is impractical: it would necessitate multiple wearables, implying scaling costs and complexity for both users and providers. **As such, with some minor exceptions, wearables are currently not used in a meaningful way in any major nation’s healthcare system.**

Nevertheless, “[t]o the extent that wearables overcome these limitations, they hold much promise towards expanding the clinical repertoire of patient-specific measures, and they are considered an important tool for the future of precision health.”²

In fact, if the issues regarding the accuracy and variety of metrics, as well as the high cost of data gathering, can be solved, the impact on society and on the world’s healthcare systems could be immense: the UK Government Office for Science, in its “*Future of Ageing*” review notes that “[i]n just one such application, such technology could potentially **cut more than \$2 trillion a year**: patients with conditions such as heart disease and diabetes could be monitored through [...] attached sensors, which can transmit readings and alert the patient, nurses and physicians when vital signs indicate an impending problem, thus avoiding crises.”³

A breakthrough in both the accuracy and variety of human health data gathered by mobile sensors, whether in wearables, intelligent clothing, or elsewhere, would, it has been noted, especially be felt in **three main sectors: government, healthcare and care, and sports**. Hanuska (University of Berkeley, California) et al. write that “companies who develop products that use big data to identify trends and can use the predictive analytics to improve the performance, health, and safety of larger populations, will achieve market success. These commercial applications, particularly in professional athletics, military, and healthcare have significant market potential.”⁴

¹ Smuck, M., Odonkor, C.A., Wilt, J.K. et al. (2021), *The emerging clinical role of wearables: factors for successful implementation in healthcare*, npj Digit. Med. 4, 45, p.1, citing Butte, N. F., Ekelund, U. & Westerterp, K. R. (2012) *Assessing physical activity using wearable monitors: measures of physical activity*. Med. Sci. Sports. Exerc. 44, S5–S12; Schrack, J. A. et al. (2016) *Assessing daily physical activity in older adults: unraveling the complexity of monitors, measures and methods*, J. Gerontol. A Biol. Sci. Med. Sci. 71, 1039–1048; Freedson, P., Bowles, H. R., Trojano, R. & Haskell, W. (2012) *Assessment of physical activity using wearable monitors: recommendations for monitor calibration and use in the field*. Med. Sci. Sports. Exerc. 44, S1–S4.

² Smuck, M., Odonkor, C.A., Wilt, J.K. et al. (2021), “*The emerging clinical role of wearables: factors for successful implementation in healthcare*”, npj Digit. Med. 4, 45, p.1.

³ UK Government Office for Science (2015) “*Future of Ageing: How are attitudes and behaviours to the ageing process changing in light of new media and new technology?*”, p.15.

⁴ Hanuska, A. et al. “*Smart clothing market analysis*”, Sutardja Center for Entrepreneurship & Technology, University of Berkeley, p.41.

In answer to the market problems listed above, the current healthcare landscape, and in order to cater to explicitly expressed market need, Decorte Future Industries developed its cloud-based Sonus engine, which can be accessed by any third party hard- and software, that uses audio, rather than traditionally employed technologies in wearables such as ECG and PPG.

3. Why audio?

Auscultation, which is the process of listening to the human body for diagnostic purposes, is among the **oldest diagnostic techniques**, dating back, according to some sources, to ancient Egypt. For the past couple of centuries, listening to the human body, using a stethoscope, has been a swift and cost-efficient way of assessing anything from respiratory patterns, digestive wellbeing to cardiovascular health.⁵

Auscultation is slowly becoming phased out from clinical practice in developed countries as auscultation-based diagnostics is an extremely difficult skill to master⁶ while there are many advanced (albeit much more expensive) diagnostic technologies at clinicians' disposal – that is, if they have access to these in the context of a centralised healthcare centre, e.g. a hospital. And yet, despite its gradual disappearance from centralised medicine over the last decade, when combined with the most recent technologies and algorithms in machine learning and signal processing, this **extremely low cost but highly effective** method can enable the first continuous mass human health data gathering, overcoming the challenges faced by current technologies for continuous human health data monitoring.

Mass human health data gathering requires automatisations, and in order to get the variety and quality needed to achieve actionable data (see above), while keeping the costs low and the devices **intuitive and comfortable**, auscultation proves to be the most promising way forward. In fact, the quality of data gathering and analysis achieved in academic studies through automated auscultation, powered by signal processing and machine learning, has been so high that it has **surpassed even that of in-person diagnoses** by trainee doctors.^{7, 8, 9}

We believe there are four main reasons why audio is an extremely attractive sensing modality to solve the problems currently seen in the remote human body- and health-data gathering and wearables markets. Audio-based monitoring is:

⁵ Sarkar, M. et al. “*Auscultation of the respiratory system*”. In: *Annals of Thoracic Medicine* 10.3 (2015), pp. 158–168.

⁶ Salvatore Mangione. “Cardiac auscultatory skills of physicians-in-training: a comparison of three English-speaking countries”. In: *The American Journal of Medicine* 110.3 (2001), pp. 210–216.

⁷ Paschalis A. Bizopoulos and Dimitris Koutsouris. “Deep learning in cardiology”. In: *CoRR* (2019).

⁸ A. K. Dwivedi, S. A. Imtiaz, and E. Rodriguez-Villegas. “Algorithms for automatic analysis and classification of heart sounds—A systematic review”. In: *IEEE Access* 7 (2019), pp. 8316–8345.

⁹ It is good to note, however, that our audio-based continuous health monitoring does not currently focus on making diagnoses, while possessing the capability to do so. It cannot and should not (presently) remove physicians from human health management altogether. Rather, in its early stages, focus is on providing an unprecedented wealth of data to stakeholders such as clinicians, users, governments and businesses to base decisions and diagnoses upon.

1. **Low-cost, integrating with basic hardware** – microphones are widely present and integrated in most devices, have a cost that is orders of magnitude lower than other biometric or diagnostic technologies, and rely on sensing technology that has been around since the 19th century, manufacturing of which has been perfected over decades.
2. **Non-invasive and safe** – relying simply on audio collection, there is no need for manipulation of the user’s body; audio collection cannot by itself cause any discomfort, nor does it require any change in the user’s routines.
3. **Easy to use** – it does not rely on any external factors, and the location of application becomes significantly less important as the signal processing methods become more advanced and powerful.
4. **Open-ended/scalable** – audio can be used for detecting a multitude of pathologies and intermediate states, tracking biometrics and pathology progression, as well as tracking psychological state. In effect, it offers a single source – a “digital twin” of the human body – that can be mined over and over again for different biometrics.

4. Comparison to the most common method of human health data-gathering through wearables and other remote mobile solutions (ECG and PPG)

The advantages of using audio-based techniques for human data gathering in intelligent clothing can be best characterised in contrast to existing methods. The most wide-spread sensing modality in the space of wearables is photoplethysmography (PPG), which is a light-based sensor commonly used in smartwatches for tracking heart rate.

Smartwatches are worn on a wrist, which, in terms of distance, is as far as one can be from the heart. Naturally, given the distance from the heart to the wrist, the signal attenuates and becomes **highly susceptible to noise** from a variety of sources – especially muscle movement. As a result, the heart rate data obtained from a PPG sensor in its raw form tends to be best when the user is still – which is highly **impractical** when accounting for modern lifestyles. It is commonly accepted that the “wrist is one of the worst places to accurately obtain a heartrate”.⁹

Advanced signal processing is needed to account for the noise, and even with most recent advances in signal processing and machine learning the **accuracy of PPG remains limited**. Furthermore, as PPG relies on light reflecting from the blood vessels in the skin, **contact with the skin is crucial**, and the signal is lost when the sensor loses contact with the skin, which is a very common occurrence. Another effect of this is that the colour of the wearer’s skin can occasionally impact accuracy levels of readings, with significant attention paid in recent national

⁹ Cecilia Mascolo, Professor in Mobile Systems, University of Cambridge, Co-Director Centre for Mobile, Wearable Systems and Augmented Intelligence, in <https://www.bcs.org/content-hub/turing-talk-2021-sounding-out-wearable-and-audio-data-for-health-diagnostics/>, consulted 04/08/2021.

and international press to the fact that PPG-based devices can return less accurate results when worn by people of colour.¹⁰

In contrast, audio based sensors do not rely nearly as heavily on location, and the flexibility in location allows placing the sensors as close to the heart as possible, meaning the signal is much less susceptible to motion noise.

Data gathering solutions such as fitness straps, smart clothing and medical [holter monitors](#), almost exclusively rely on embedded electrocardiogram (ECG) electrodes. However, ECG electrodes have a number of caveats. A typical ECG exam in a clinical setting requires the placement of ten electrodes in very precise locations all over the patient's body. It requires data collection to take place in complete stillness, and the electrodes to all have a thin gel coating to increase the skin conductivity, ensuring high quality signal. This is known as a 12-lead ECG and is a gold standard in medicine. Naturally, virtually none of the commercial tracker bands or smart wear on the market achieves similar coverage, or have 10 electrodes embedded. Instead, most solutions rely on a single, 2 or 3-lead ECG, which is useful for heart rate and heart rate variability detection, but is **limited for medical diagnostics**.

In addition, skin contact is a significant challenge. As mentioned before, gel on the electrodes is an essential part of the ECG exam due to the importance of good signal conductivity. However, having gel-based electrodes in wearables or intelligent clothing is unfeasible – in addition to it being technologically challenging – as gel of the types needed to guarantee such conductivity naturally dries out over time. To combat this issue, most devices employ dry electrodes. However, the **accuracy of dry electrodes is significantly lower** than of their gel-based counterparts.

Being aware of this issue, but without alternatives available to them, multiple companies in the space have confirmed their products **rely on user perspiration**, where sweat acts as the conductor of electrical signals. This, however, effectively limits the useability and accuracy of such solutions to the contexts of sports or other perspiration-inducing activities, as no reliable reading can otherwise be achieved. Furthermore, even in those contexts, such a method is only effective from the moment significant perspiration is produced, which rarely lines up with the moment the perspiration-inducing activity is begun - therefore losing significant amounts of data.

A major hurdle for mass adoption of ECG-based straps and wearables is that they almost exclusively employ “e-textiles” or “intelligent textiles”, which are textiles where conductive wires have been woven into the yarn in order to send across ECG signals. These are **notoriously expensive to produce, requiring multiple complex layers**. However, the very high cost of such e-textiles is only the second most significant blocker for their mass adoption: next to their price tag, due to the presence of metals and unavoidable oxidation processes, the **washability of**

¹⁰ Bent, B., Goldstein, B. A., Kibbe, W. A., & Dunn, J. P. (2020). Investigating sources of inaccuracy in wearable optical heart rate sensors. NPJ digital medicine, 3, 18.

such textiles has still not been fully achieved. This means that virtually all such technologies produced, whether by startups or tech giants such as Google (see e.g. Project Jacquard), have both been expensive and not reliable washable - effectively preventing widespread adoption in wider consumer markets.

In comparison, in December 2021, **Decorte Future Industries succeeded in achieving 99.6% accuracy** collecting heart-related metrics using exclusively audio from a standard microphone, **compared to data collected simultaneously by ECG** at an exponentially higher cost.




Figure 2. Comparison of data gathered simultaneously from an ECG device strapped to the body, and audio - yielding a 0.4% average error percentage, and a 0.07% average error per 10 second window..

5. Conclusion

Enabling a quantum leap in automated human health data gathering has become the ultimate goal for us at Decorte Future Industries Ltd. Offering companies, governments and healthcare professionals across the world access to audio-based sensing allows us to do just that. We automate perhaps the oldest form of holistic health data gathering, leveraging its unique affordability, safety and - most importantly - ability to capture an extremely wide variety of biomarkers, spanning from cardiovascular and respiratory health to mental wellbeing and mood prediction. Machine learning and signal processing, as well as the ubiquitousness of microphones in modern society, is now beginning to allow us to bring this tried and tested method into the 21st century, and deploy it on a scale never seen before.

About the authors

Erika Bondareva is CTO at Decorte Future Industries, as well as a PhD candidate at the University of Cambridge. Her doctoral work is on “Machine learning on wearables for efficient cardiovascular health diagnostic and progression through audio signals.” She holds an MRes in Sensor Technologies and Applications from the University of Cambridge and an MEng in Biomedical Engineering from the University of Glasgow.

Dr Roeland Decorte is Founder and CEO at Decorte Future Industries. He holds a BA/MA, MPhil and PhD in a codebreaking discipline, all from the University of Cambridge, where he also taught for seven years. He designed the initial technology as well as sensor infrastructure underlying the company’s solutions.

Further reading on the potential of audio-based health data gathering

[BCS, The Chartered Institute for IT “Turing Talk 2021: Sounding out wearable and audio data for health diagnostics”, \(2021\).](#)

[J.L.Kröger., O.HM. Lutz, P. Raschke “Privacy Implications of Voice and Speech Analysis – Information Disclosure by Inference”. In: M. Friedewald et al. “Privacy and Identity Management. Data for Better Living: AI and Privacy. Privacy and Identity” \(2019\). IFIP Advances in Information and Communication Technology, vol 576. Springer.](#)

Further reading on clinical applications of audio

[A. K. Dwivedi, S. A. Imtiaz, and E. Rodriguez-Villegas. “Algorithms for automatic analysis and classification of heart sounds–A systematic review”. In: IEEE Access 7 \(2019\), pp. 8316– 8345.](#)

[Renard Xaviero Adhi Pramono, Stuart Bowyer, and Esther Rodriguez-Villegas. “Automatic adventitious respiratory sound analysis: A systematic review”. In: PLOS ONE 12.5 \(May 2017\), pp. 1–43.](#)

[E. Ambrosini et al. “Automatic speech analysis to early detect functional cognitive decline in elderly population”. In: Annu Int Conf IEEE Eng Med Biol Soc. \(2019\).](#)

[Declan Bray et al. “Assessing motility through abdominal sound monitoring”. In: Annual International Conference of the IEEE Engineering in Medicine and Biology Society. Chicago, IL, USA: IEEE, \(1997\), pp. 2398–2400.](#)

P. McColgan et al. “*Evaluation of the clinical utility of a carotid bruit*”. In: QJM: An International Journal of Medicine 105.12 (2012), pp. 1171–1177.