



Stefano CANALI, Alessandro FALCETTA, Massimo PAVAN,
Manuel ROVERI, Viola SCHIAFFONATI

Dipartimento di Elettronica, Informazione e Bioingegneria
Politecnico di Milano

 POLITECNICO DI MILANO



To Personalize or Not To Personalize? Soft Personalization and the Ethics of ML for Health

An interdisciplinary effort



Manuel Roveri

Design and
development of
TinyML systems



Massimo Pavan



Alessandro Falcetta



Viola Schiaffonati

Ethical concerns
and values



Stefano Canali

Outline

- **IoT** increasingly spreading into the domain of **medical** and **social care (H-IoT)**
- Specific **risks** and **ethical issues** emerging from this use now widely discussed in current literature
- **Lack** of a **systematic discussion** on the role of **ML** in **H-IoT** and its **impact** on **ethical concerns**
- Case: **glucose-monitoring** to raise **alerts** of critical situations and **four different scenarios** with a different role for **ML**
- Mapping **ethical concerns** into these **different scenarios**
- **Personalization**: the new hope in **TinyML**
- Conclusion and research directions

IoT and health

Internet of Things (IoT) in Healthcare: Benefits, Use Cases and COVID Impact

By Stuart Rauch

Last updated on Dec 12, 2022

1443



HEALTHCARE TECHNOLOGY

IIOT: THE INTERNET OF THINGS

IoT in Healthcare: 15 Examples of Internet of Things Healthcare Devices and Technology

7 examples of how the internet of things is facilitating healthcare

In a new age of remote healthcare, how is the internet of things enabling new forms of medical treatment, understanding and care?



Innovate using the vast, connected universe



IEEE
SPECTRUM

FOR THE TECHNOLOGY INSIDER | 01.20

*WHAT TO LOOK FOR IN THE COMING YEAR

- Autonomous Fighter Jets
 - Wafer-Scale Chips
 - Drone Delivery
- Exascale Computing
- Robot Farm Hands
- A New Generation of Mars Landers

and more...


Top Tech 2020*

Triumphs & Turning Points

A SPECIAL REPORT



INTERNET OF EVERYTHING BY STACEY HIGGINBOTHAM



MACHINE LEARNING ON THE EDGE

▶ IN FEBRUARY, a group of researchers from Google, Microsoft, Qualcomm, Samsung, and half a dozen universities will gather in San Jose, Calif., to discuss the challenge of bringing machine learning to the farthest edge of the network, specifically microprocessors running on sensors or other battery-powered devices.

The event is called the Tiny ML Summit (ML for “machine learning”), and its goal is to figure out how to run machine learning algorithms on the tiniest microprocessors out there. Machine learning at the edge will drive better privacy practices, lower energy consumption, and build novel applications in future generations of devices.

As a refresher, at its core machine learning is the training of a neural network. Such training requires a ton of data manip-

ulation. The end result is a model that is designed to complete a task, whether that’s playing Go or responding to a spoken command.

Many companies are currently focused on building specialized silicon for machine learning in order to train networks inside data centers. They also want silicon for conducting inference—running data against a machine learning model to see if the data matches the model’s results—at the edge. But the goal of the Tiny ML community is to take inference to the smallest processors out there—like an 8-bit microcontroller that powers a remote sensor.

To be clear, there’s already been a lot of progress in bringing inference to the edge if we’re talking about something like a smartphone. In November 2019, Google open-sourced two versions of its machine learning algorithms, one of

which required 50 percent less power to run, and the other of which performed twice as fast as previous versions of the algorithm. There are also several startups such as Flex Logix, Greenwaves, and Syntiant tackling similar challenges using dedicated silicon.

But the Tiny ML community has different goals. Imagine including a machine learning model that can separate a conversation from background noise on a hearing aid. If you can’t fit that model on the device itself, then you need to maintain a wireless connection to the cloud where the model is running. It’s more efficient, and more secure, to run the model directly on the hearing aid—if you can fit it.

Tiny ML researchers are also experimenting with better data classification by using ML on battery-powered edge devices. Jags Kandasamy, CEO of Latent AI, which is developing software to compress neural networks for tiny processors, says his company is in talks with companies that are building augmented-reality and virtual-reality headsets. These companies want to take the massive amounts of image data their headsets gather and classify the images seen on the device so that they send only useful data up to the cloud for later training. For example, “If you’ve already seen 10 Toyota Corollas, do they all need to get transferred to the cloud?” Kandasamy asks.

On-device classification could be a game changer in reducing the amount of data gathered and input into the cloud, which saves on bandwidth and electricity. Which is good, as machine learning typically requires a lot of electricity.

There’s plenty of focus on the “bigger is better” approach when it comes to machine learning, but I’m excited about the opportunities to bring machine learning to the farthest edge. And while Tiny ML is still focused on the inference challenge, maybe someday we can even think about training the networks themselves on the edge. ■

POST YOUR COMMENTS AT
spectrum.ieee.org/tinymachinlearning-jan2020

ILLUSTRATION BY Dan Page


20 | JAN 2020 | SPECTRUM.IEEE.ORG

Move (intelligent) processing as close as possible to data generation units ...


- ✓ Increase autonomy
- ✓ Reduce decision-making latency
- ✓ Reduce transmission bandwidth
- ✓ Increase energy-efficiency
- ✓ Security and Privacy
- ✓ Incremental/Adaptive Learning
- ✓ Ecosystem of units

- Low computing ability
- Constraints on energy
- Constraints on memory (RAM/FLASH)
- Complexity in design and development
- Strong connection between HW, SW and ML

Arduino



STM32




Domain	STM32 L1 Series	STM32F4 Series
	Ultra Low-Power	High-Performance
Flash Memory (kB)	32 to 512	64 to 2048
RAM Memory (kB)	4 to 80	32 to 320
CPU	ARM® Cortex®-M3	ARM® Cortex®-M4
Frequency (MHz)	32	84 to 180
Supply Voltage (V)	1.65 to 3.6	1.71 to 3.6
Supply Current (μA)	0.28 (0.28) to 230	1.1 (140) to 282

Wearable technologies



Big promises...

Personalization, monitoring, assistance



Not just promises...



The screenshot shows a web browser displaying the New York Times website. The article is titled "Just How Accurate Are Fitbits? The Jury Is Out" by Mike McPhate, dated May 25, 2016. The article is categorized under "PERSONAL TECH". The main text discusses the accuracy of activity trackers, mentioning a study by plaintiffs in a class-action lawsuit against Fitbit. The study found that the pulse-monitoring technology used in Fitbit's Surge and Charge devices was "highly inaccurate during elevated physical activity." The article also mentions that researchers from California State Polytechnic University, Pomona, had 43 subjects wear the devices as they ran, jogged, and jumped rope, among other activities, and then compared the readings with those of an electrocardiogram.

PERSONAL TECH

Just How Accurate Are Fitbits? The Jury Is Out

By MIKE McPHATE MAY 25, 2016



Fitbit Charge HR Tony Cenicola/The New York Times

Many users of [activity trackers](#) have always harbored suspicions: How accurate are these things?

A handful of tests by [journalists](#) and researchers have tried to bring clarity to the issue. Results, alas, have been mixed.

The [latest study](#), released by the plaintiffs [in a class-action lawsuit](#) against Fitbit, found that the pulse-monitoring technology used in the company's wrist-bound Surge and Charge devices was "highly inaccurate during elevated [physical activity](#)."

Researchers from California State Polytechnic University, Pomona, had 43 subjects wear the devices as they ran, jogged and jumped rope, among other activities, and then compared the readings with those of an electrocardiogram.

flipkart OFFERS ZONE

YOUR ONE-STOP SHOP FOR EXCITING NEW OFFERS

flipkart SHOP NOW »


Datafication, surveillance, profiling, ...

Spazzolini elettrici intelligenti, il test di : l'IA c'è, ma non serve

di **Massimiliano Di Marco** - 08/11/2018 09:45

4





FINDING ONE FACE IN A MILLION

A new benchmark test shows that even Google's facial recognition algorithm is far from perfect

Helen of Troy may have had the face that launched a thousand ships, but even the best facial recognition algorithms might have had trouble finding her in a crowd of a million strangers. The first public benchmark test based on 1 million faces has shown how facial recognition algorithms from Google and other research groups around the world still fall well short of perfection.

Facial recognition algorithms that had previously performed with more than 95 percent accuracy on a popular benchmark test involving 13,000 faces saw significant drops in accuracy when taking on the new MegaFace Challenge. The best performer, Google's FaceNet algorithm, dropped from near-perfect accuracy on the five-figure data set to 75 percent on the million-face test. Other top algorithms dropped from above 90 percent to below 60 percent. Some algorithms made the proper identification as seldom as 35 percent of the time.

"MegaFace's key idea is that algorithms should be evaluated at large scale," says Ira Kemelmacher-Shlizerman, an assistant professor of computer science at the University of Washington, in Seattle, and

the project's principal investigator. "And we make a number of discoveries that are only possible when evaluating at scale."

Finding a million faces matter because facial recognition algorithms inevitably face such challenges in the real world. People increasingly trust these algorithms to correctly identify them in security verification scenarios, and law enforcement may also rely on facial recognition to pick

The most popular benchmark until now has been the Labeled Faces in the Wild (LFW) test created in 2007. LFW includes 13,000 images of just 5,000 people. Many facial recognition algorithms have been fine-tuned to the point that they scored near-perfect accuracy when picking through the LFW images. Most researchers say that new benchmark challenges have been long overdue.

"The big disadvantage is that [the field] is saturated—that is, there are many, many algorithms that perform above 95 percent on LFW," Kemelmacher-Shlizerman says. "This gives the impression that face recognition is solved and working perfectly."

With that in mind, University of Washington researchers raised the bar by creating the MegaFace Challenge using 1 million Flickr images of 690,000 unique faces that are publicly available under a Creative Commons license.

The MegaFace Challenge forces facial recognition algorithms to do verification and identification, two separate but related tasks. Verification involves trying to correctly determine whether two faces presented to the facial recognition algorithm belong to the same person. Identification involves trying to find a matching photo of the same person among a million "distractor" faces. Initial results on algorithms developed by Google and four other research groups were presented at the IEEE Conference on Computer Vision and Pattern Recognition on 30 June. (One of MegaFace's developers also heads a computer vision team at Google's Seattle office.)

The results presented were a mix of the intriguing and the expected. Nobody was surprised that the algorithms' performances suffered as the number of distractor faces increased. And the fact that algorithms had trouble identifying the same person at different ages was a known problem. However, the results also showed that algorithms trained on relatively small data sets can compete with those trained on very large ones, such as Google's FaceNet, which was trained on more than 500 million photos of 10 million people.

For example, the FaceNet algorithm from Russia's N-TechLab performed well on certain tasks in comparison with FaceNet, despite having trained on 18 million photos of 200,000 people. The SIAT MMLab algorithm, created by a Chinese team under the leadership of Yu Qiao, a professor with Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences, also performed well on certain tasks.

Nevertheless, FaceNet has so far performed the best overall. It delivered the most consistent performance across all testing.


The huge drops in accuracy when scanning a million faces matter because facial recognition algorithms inevitably face such challenges in the real world. People increasingly trust these algorithms to correctly identify them in security verification scenarios, and law enforcement may also rely on facial recognition to pick suspects out of the hundreds of thousands of faces captured on surveillance cameras.

The most popular benchmark until

IEEE Spectrum
Aug, 2016


Open issues

- **Privacy** and **low security** standards (Elkhodr, 2011)
- **User access** and **ownership** of data (Hummel et al., 2020; Hummel & Braun, 2020)
- **Datafication** of private spaces and **privatization** of medical services (Ishmaev, 2020)
- **Reuse** of medical data for other purposes and **limits of informed consent** (Mittelstadt & Floridi, 2016)



A (real-world) example: glucose-monitoring for diabetes

Glucose-monitoring to raise alerts of critical situations



Personal data



Alert/alarm




AI Model



Which are the technological actors involved in the processing?

How to get a personalized model for glucose-monitoring?

Glucose-monitoring to raise alerts of critical situations



AI Model



Scenario #1: Inference on Cloud and pre-defined model



Scenario #2: Inference on Cloud and personalized model



Scenario #3: Inference on device (TinyML) and predefined model



Scenario #4: Inference on device (TinyML) and personalized model (On-device TinyML)



Different solutions, different issues

	Inference on Cloud	Inference on Device
Pre-defined model	<ul style="list-style-type: none">• Traditional solution• Personal data must flow through the internet• Less effective model• “Connected” device	<ul style="list-style-type: none">• “TinyML for inference”• Personal data remain on the device• Less effective model• Intelligent tiny devices
Personalized model	<ul style="list-style-type: none">• More advance solution• Personal data must flow through the internet• More effective model• “Connected” device	<ul style="list-style-type: none">• “TinyML for training”• Personal data remain on the device• Personal model• Intelligent tiny devices• Requires a powerful device

Ethical concerns

ORIGINAL PAPER

Ethics of the health-related internet of things: a narrative review

Brent Mittelstadt^{1,2,3} 

Risks of Internet enabled devices

Sensitivity of health-related data

Impact on the delivery of healthcare

Abstract The internet of things is increasingly spreading into the domain of medical and social care. Internet-enabled devices for monitoring and managing the health and well-being of users outside of traditional medical institutions have rapidly become common tools to support health-care. Health-related internet of things (H-IoT) technologies increasingly play a key role in health management, for purposes including disease prevention, real-time tele-monitoring of patient's functions, testing of treatments, fitness and well-being monitoring, medication dispensation, and health research data collection. H-IoT promises many benefits for

health and healthcare. However, it also raises a host of ethical problems stemming from the inherent risks of Internet enabled devices, the sensitivity of health-related data, and their impact on the delivery of healthcare. This paper maps

the main ethical problems that have been identified by the relevant literature and identifies key themes in the on-going debate on ethical problems concerning H-IoT.

A review of the ethical issues (Mittelstadt, 2017)


- ***Devices***: inherent risks of Internet enabled devices
 - E.g., personal **privacy** and **social isolation**
- ***Data***: sensitivity of health-related data
 - E.g., **data sharing** and **autonomy**
- ***Practices***: impact on the delivery of healthcare
 - E.g., **social isolation**

Mapping these issues into our scenarios

- On Cloud vs. on device



- Predefined vs. personalized model



Mapping ethical concerns

- Same concerns in our different scenarios but different degrees (shades of grey)
- Very serious (darker), serious (medium), less serious (lighter)

Ethical issues for H-IoT (Mittelstet, 2017)	Inference on the Cloud with predefined model (I)	Inference on the Cloud with personalized model (II)	Inference on device with predefined model (III)	Inference on device with personalized model (IV)
Privacy	Constant monitoring and surveillance	Constant monitoring and surveillance	Constant monitoring and surveillance	Constant monitoring and surveillance
	Data control	Data control	Data control	Data control
	Data security	Data security	Data security	Data security
Autonomy	Freedom & independence	Freedom & independence	Freedom & independence	Freedom & independence
Consent	Limits of informed consent	Limits of informed consent	Limits of informed consent	Limits of informed consent
	Limits of anonymization	Limits of anonymization	Limits of anonymization	Limits of anonymization
Ownership and data access	Data access	Data access	Data access	Data access
	Transparency of results	Transparency of results	Transparency of results	Transparency of results
	Sharing of benefits	Sharing of benefits	Sharing of benefits	Sharing of benefits
Social isolation	Use IoT devices as replacement of social interactions	Use IoT devices as replacement of social interactions	Use IoT devices as replacement of social interactions	Use IoT devices as replacement of social interactions
decontextualization of health and well-being	Simplification of parameters	Simplification of parameters	Simplification of parameters	Simplification of parameters
	Lack of integration of complex and contextual information	Lack of integration of complex and contextual information	Lack of integration of complex and contextual information	Lack of integration of complex and contextual information
'Good' care and user well-being	Quality of care delivered through H-IoT	Quality of care delivered through H-IoT	Quality of care delivered through H-IoT	Quality of care delivered through H-IoT
Risks of non-professional care	Privatization, marketing, targeting	Privatization, marketing, targeting	Privatization, marketing, targeting	Privatization, marketing, targeting

Privacy - Constant monitoring and surveillance

- Concerns on **constant surveillance**
 - E.g. **More data** collected to develop personalized models
 - **Very serious** in the case of **personalized models** rather than in the case of predefined models (serious)
 - E.g. **Sharing** of personal data on the cloud with third parties
 - **Very serious** in the case of **personalized models** rather than in the case of predefined models (serious)

Ethical issues for H-IoT (Mittelstdat, 2017)	Inference on the Cloud with predefined model (I)	Inference on the Cloud with personalized model (II)	Inference on device with predefined model (III)	Inference on device with personalized model (IV)
Privacy	Constant monitoring and surveillance	Constant monitoring and surveillance	Constant monitoring and surveillance	Constant monitoring and surveillance

Privacy - Data control and security

- Concerns on **data control**
 - E.g. **Accessibility** and **transparency** of cloud services
 - **Very serious** on cloud rather than on device (less serious)
- Security of **personal data** (breaches)
 - **Very serious** on cloud rather than on device (less serious)

Ethical issues for H-IoT (Mittelsttat, 2017)	Inference on the Cloud with predefined model (I)	Inference on the Cloud with personalized model (II)	Inference on device with predefined model (III)	Inference on device with personalized model (IV)
Privacy	Data control	Data control	Data control	Data control
	Data security	Data security	Data security	Data security

Lesson learned

- Articulating the **current (simplistic) narrative**
 - **More privacy and less issues** moving from cloud to **device** (Tiny-ML)
- The **benefits** of our analysis (different scenarios)
 - **Not only** a matter of **cloud vs device**, but also of **personalized models**


Beyond privacy: new issues

- Mitigation of privacy issues also by technical means (e.g., HE) but new emergent issues for personalized models
 - Transparency of results and accessibility of encrypted models
 - Quality and evaluation of efficacy
 - Responsibility of evaluation

Ethical issues for H-IoT (Mittelstet, 2017)	Inference on the Cloud with predefined model (I)	Inference on the Cloud with personalized model (II)	Inference on device with predefined model (III)	Inference on device with personalized model (IV)
Privacy	Constant monitoring and surveillance	Constant monitoring and surveillance	Constant monitoring and surveillance	Constant monitoring and surveillance
	Data control	Data control	Data control	Data control
Autonomy	Data security	Data security	Data security	Data security
	Freedom & independence	Freedom & independence	Freedom & independence	Freedom & independence
Consent	Limits of informed consent	Limits of informed consent	Limits of informed consent	Limits of informed consent
	Limits of anonymization	Limits of anonymization	Limits of anonymization	Limits of anonymization
Ownership and data access	Data access	Data access	Data access	Data access
	Transparency of results	Transparency of results	Transparency of results	Transparency of results
Social isolation	Sharing of benefits	Sharing of benefits	Sharing of benefits	Sharing of benefits
	Use IoT devices as replacement of social interactions	Use IoT devices as replacement of social interactions	Use IoT devices as replacement of social interactions	Use IoT devices as replacement of social interactions
decontextualization of health and well-being	Simplification of parameters	Simplification of parameters	Simplification of parameters	Simplification of parameters
	Lack of integration of complex and contextual information	Lack of integration of complex and contextual information	Lack of integration of complex and contextual information	Lack of integration of complex and contextual information
'Good' care and user well-being	Quality of care delivered through H-IoT	Quality of care delivered through H-IoT	Quality of care delivered through H-IoT	Quality of care delivered through H-IoT
Risks of non-professional care	Privatization, marketing, targeting	Privatization, marketing, targeting	Privatization, marketing, targeting	Privatization, marketing, targeting

New issues and trade-offs: glucose monitoring

- **Encrypted personalized models** raising concerns for **transparency** - impossible to inspect the trained models without secret keys
 - **Burden** on the **user** of checking results of the models and their **quality** and **reliability**
- Personalized models running the risk of **learning** as “normal” **situations** that are instead **harmful** for the health of the users: in the case of TinyML the model is trained directly on device (unpredictable evolution of the learning)
 - **Experts** increasingly **less involved**



Privacy and personalization

- **Interplay of different values** connected to the uses in different ways: need for **balance** and **trade-off**
- **Epistemic and ethical values:** privacy, autonomy, justice, fairness, transparency, beneficence, non-maleficence, ...

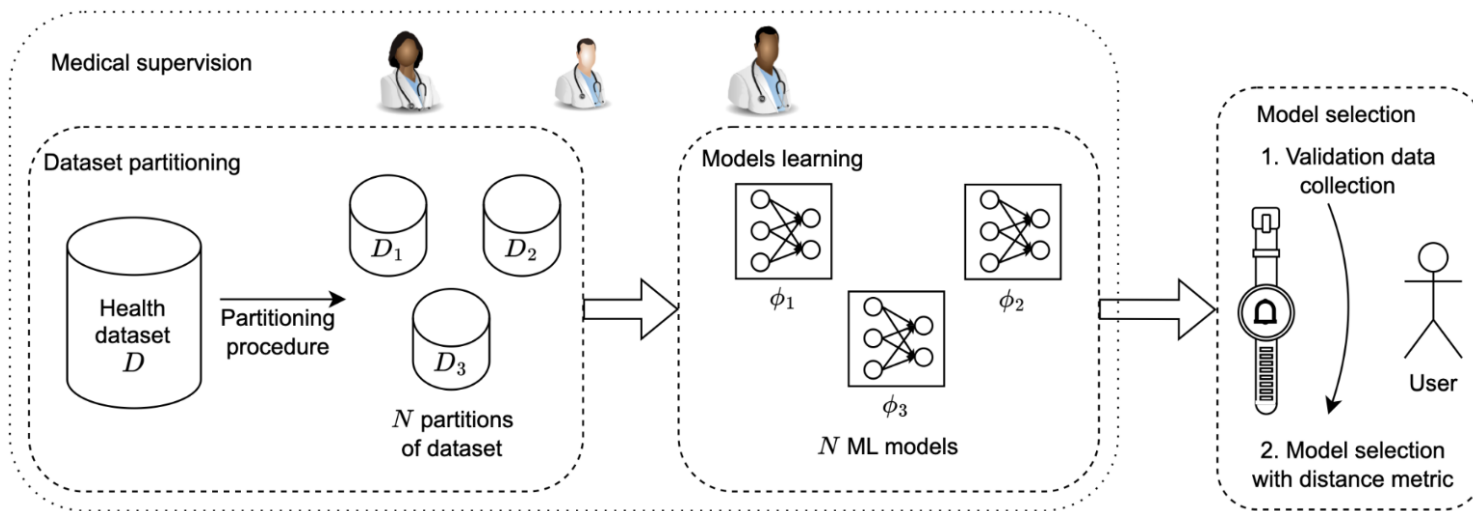
How to address these issues?

The proposed approach: soft personalization

- Need to **control** the learning and evolution of **personalized models** respecting **ethical values** and **epistemic requirements** necessary in the biomedical context
 - **Privacy** and **security** remaining central
- **Epistemic requirements:** models not overestimating and properly classifying problematic states as anomalies
- **Ethical values:** models not harming users and not problematic from principles of non-maleficence and beneficence

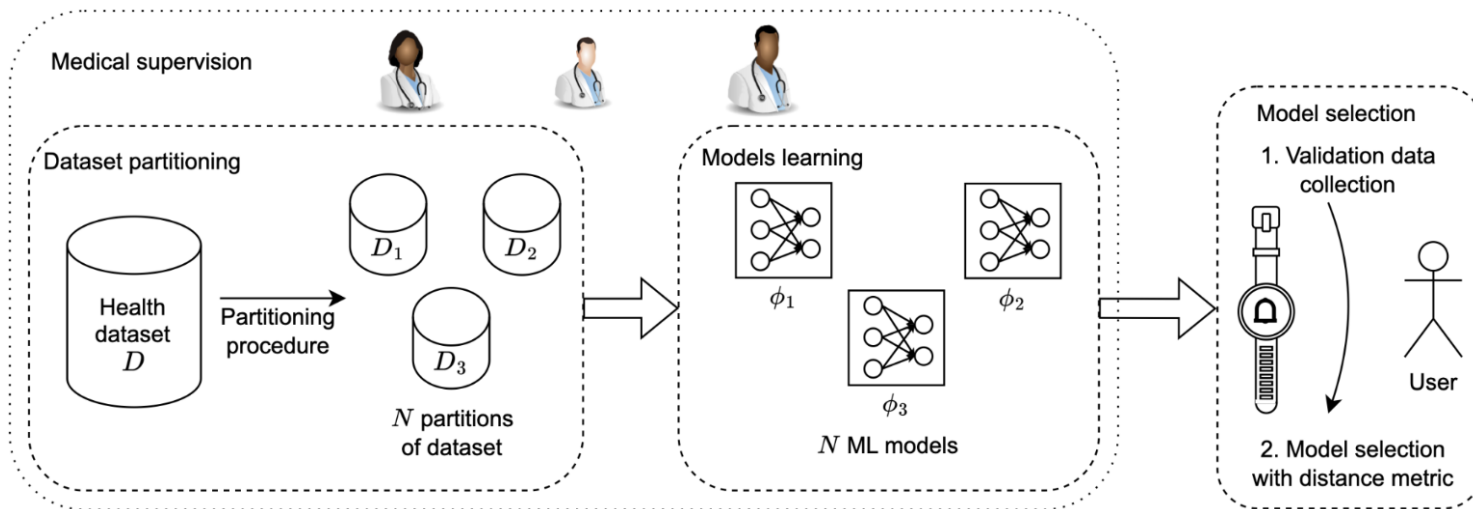
Soft personalization: the idea

- In our case of IoT and ML for glucose
 - Development of **different models** starting from personal data and more
 - **Selection and evaluation** of models by **users and experts**
 - **Less over-fitting** and control with more **conservative and specific thresholds** (quality, representativeness)
 - Greater guarantees of **reliability** (beneficence, non-maleficence)



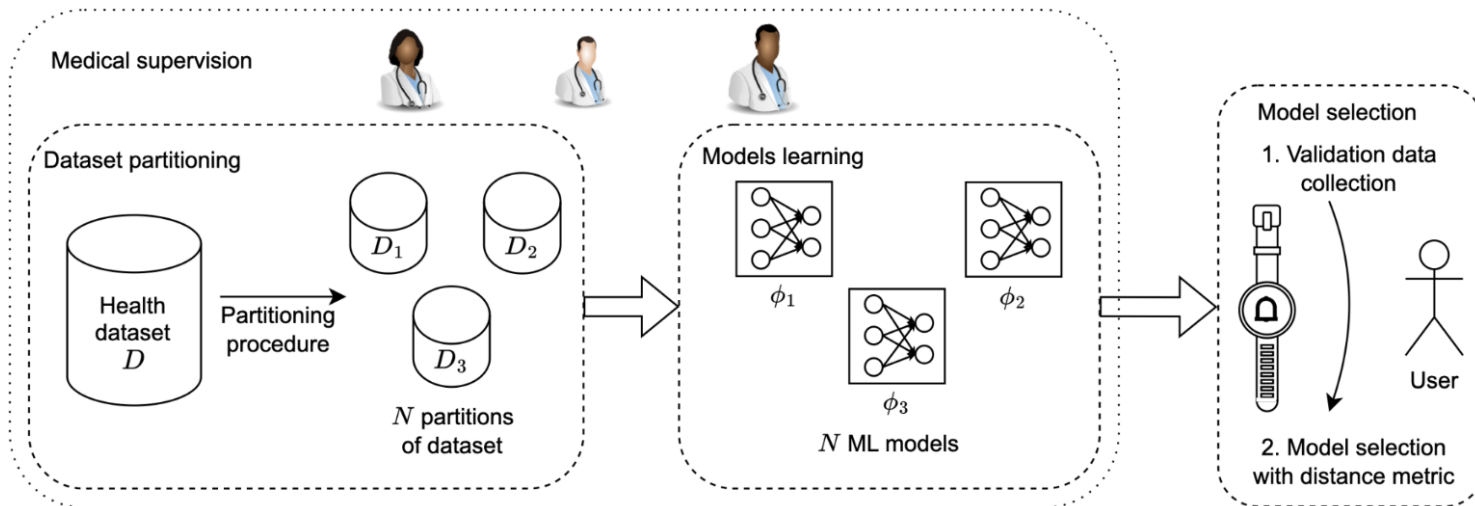
Soft personalization: the limitations

- Some limitations remain
 - **Data quality and responsibility on user**
 - **Collaboration, education and user burden of checking results**
 - **Pathologies and lifestyles evolving over time**



Soft personalization: trade-offs

- Impossible to solve these issues only by technical means
- Always a matter of **trade-offs**
 - Overall **efficacy** and **quality** and reliability (e.g. disease with no specific pattern)
 - Choice of **specific values** to respect (privacy, accuracy)
 - **Exclusion of others** (general effectiveness, transparency)



Conclusions and future directions

- **Promises** straddling IoT, ML and health, but also **significant problems** and **open issues** (privacy, personalization, control)
- Development and **innovative algorithmic solutions** for **privacy** preservation and their limitations
- **Soft personalization** as a way to **mitigate issues** and go in specific value directions, but **various trade-offs** and issues remain **open**

