Human-Centric Computing

Jan M. Rabaey, *Fellow, IEEE*

***Abstract*— With the world around us rapidly becoming smarter, an extremely relevant question is how "we humans" are going to cope with the onslaught of information coming at us. One plausible answer is to use similar technologies to evolve ourselves and to equip us with the necessary tools to interact with and to become an essential part of the smart world. Various wearable devices have been or are being developed for this purpose. However, their potential to create a whole new set of human experiences is still largely unexplored. To be more effective, functionality cannot be centralized and needs to be distributed to capture the right information at the right place. This requires a human Intranet, a platform that allows multiple distributed input–output and information processing functions to coalesce and form a single application. In addition, it needs the capabilities to understand, interpret, reason, and act on the obtained data under diverse and changing conditions, and to do so in concert with the human body and its computer, the brain. To this effect, this article explores the concept of human-centric computing, an approach that aspires to create a symbiotic convergence between biological and physical computing.**

***Index Terms*— Body-sensor networks, machine intelligence, neuro-inspired computing, wearable computers.**

I. Introduction

THERE is no doubt that the world around us is getting a lot smarter quickly. With the advent of sensor networks, the Internet of Things (IoT), Cyberphysical Systems, and Swarms, virtually every single component of our daily living environment is being equipped with sensors, actuators, and processing, all of which connected wirelessly into a network that soon will count trillions of sensors [4], [5]. Even more, that network is changing as we speak and is becoming a lot more *dynamic*. Early IoT incarnations were "static," that is, sensory nodes were placed in fixed locations and transmitted the collected data to the Cloud for processing and interpretation. In the emerging IoT++ world, nodes are attached to robots, drones, vehicles, and humans that are moving freely around and are evolving from pure data-gathering entities toward closed-loop sense–interpret–actuate systems, enabling distributed autonomous behavior.

All this leads to the question, what will be the role of us, humans, in this smart new world? Indeed, many concerns regarding information overload, employment, and artificial intelligence (AI) dominance have been raised over the last few years and have created a nonnegligible pushback against technology adoption. A more constructive approach, however, may be to consider an alternative option: embrace that very same technology to make us humans more adapted and fine-tuned to that smarter world. An article in the *National Geographic Magazine* in April 2017 titled "How Humans Are Shaping Their Own Evolution" is stating it very succinctly [6]: "Like other species, we are the products of millions of years of adaptation. Now we're taking matters into our own hands." To do so, one needs to reconsider the relationship between humans and computing. Rather than considering the human with her biological brain as being totally separate from the physical computing world (the cloud, IoT), she should become enmeshed with it. In other words, computing should become human centric.

A quick search through Wikipedia yields the following description: "Human-centered computing (HCC) studies the design, development, and deployment of mixed-initiative human–computer systems. It (…) is closely related to human–computer interaction and information science" [7]. We believe that this definition is still very much based on a traditional human–computer interaction model. True HCC aspires to create a symbiotic convergence between biological and physical computing.

To reflect on what this may mean, let us analyze the core tasks the mammalian brain performs. At its highest level, it is trying to maximize a reward function that can be quite complex, but that, in essence, boils down to a simple "survive, prosper, reproduce" mantra. To accomplish these goals, it observes the external world through a set of sensory channels. All sensory data reaches the brain in the form of tiny patterns that need be assembled in a larger picture on spatial and temporal scales that are relevant to behavior [8]. Interpretation of these patterns, based on the past or learned experiences, leads to decisioning and action—that is, coupled to precise motor actions (move, grab, speak, etc.). All of these forms a tight feedback loop that is continuously being optimized to ensure the best possible outcome(s) (Fig. 1). The optimization process is constrained by the availability of resources, most importantly energy: "Is an action and its result worth the energy that was spent, or would a different action have yielded a better return?"

This is, however, only half of the story. Another essential function of the brain is to ensure the health and well functioning of the (internal) system itself. To support rich external behaviors, the human/animal needs specialized internal organs and functions that help to provide energy to the body where and when it is needed, to store energy for future use, repair broken functionality, and protect against

Manuscript received November 10, 2019; accepted November 12, 2019. Date of current version December 27, 2019.

The author is with the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley, CA 94720 USA, and also with the CTO of the CSTO Division, IMEC, 3001 Leuven, Belgium (e-mail: jan@eecs.berkeley.edu).

Color versions of one or more of the figures in this article are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TVLSI.2019.2956529

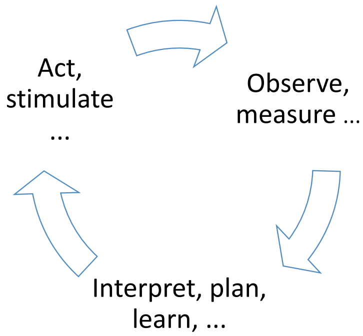

Fig. 1. Inherent feedback loop in human computing.

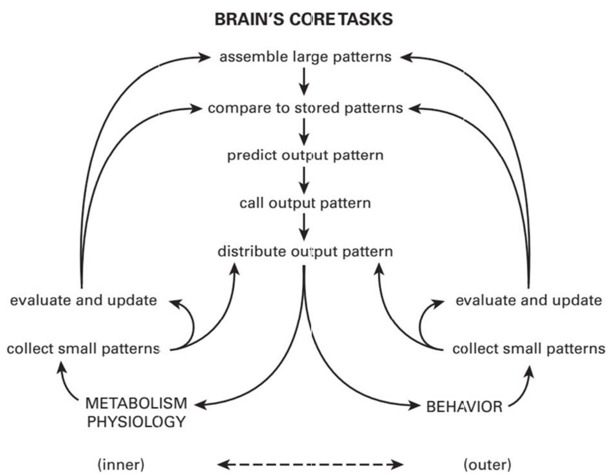

Fig. 2. Inner and outer loops perform the same broad tasks. Note that they also couple to serve each other (from [8]).

infectious agents and parasites. The brain acts as the effective regulator for these functions, most often in a fully automatic self-regulatory mode. To do so, it actually executes a feedback loop that is virtually identical to one for the external tasks (Fig. 2), including sensing/observation, the construction of large patterns from small ones ("what time of the day is it?", "Is food be coming soon?", etc.), as well as the making of decisions and actioning (which could, for instance, be the release of hormones). In fact, both feedback loops are coupled to some degree, and serve each other.

Given these high-level observations, we can start imagining how technology and "physical" computing can help to complement and even augment the biological system. Today already, a number of wearable devices are on the market that allow us to peer into the operation and well-being of our body. We can measure our heart rate, blood oxygenation, blood sugar levels, and our brain waves. In most cases, this translates into a "monitoring only" functionality with no closed-loop feedback. The latter is, however, bound to happen sooner or later. We can imagine a network of devices worn on, implanted in, or moving through the human body to perform a full "introspection," measuring a broad range of biophysical and biochemical parameters. The obtained information can then be interpreted and acted upon using a broad spectrum of measures such as electrical, magnetic, or optical stimulation, and/or

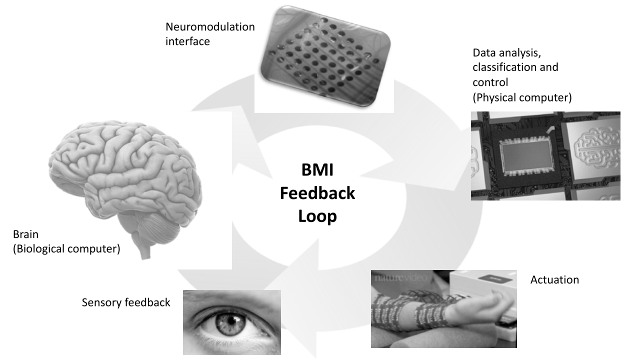

Fig. 3. BMI as a closed-loop feedback system, including both a biological and a physical computer in the loop.

precision drug delivery. As such, it acts as a complement and an extension of the internal loops of our neural system.

The same holds for the external loops. Numerous devices have been developed and used over the years to help address malfunction or degradation of an individual's sensory systems. Examples of such are the hearing aid and the artificial retina. Similarly, prostheses, exoskeletons, and voice synthesizers help to restore functionality to missing or failing motor functionality. Yet, the opportunities go far beyond. Wearable and mobile devices may offer a broad selection of sensory input channels that exceed, supplement or augment our human sensory spectrum, and, as such, provide enhanced "extrospection." This information is passed to an embedded electronic processor, which translates the information into something that maps onto the traditional human input channels, through which it is passed to the brain. The result can be a more efficient, an enhanced, or a truly novel experience providing improved awareness. Augmented-reality (AR) glasses are a perfect example of such superimposing extra information on our visual input channels. Sonification aims to map vision data onto auditory signals, hence, cross-linking between sensory pathways [9]. The GPS function in the smartphone or watch provides us with instant locationing and path planning capabilities. Before its arrival, we spent a large amount of brain resources to those tasks. An extra and distinguishing feature is that the embedded processor(s) can be coupled to the Cloud through a wireless link, offering not only more processing power but also advanced learning and model-building capabilities, merging the experiences of many individuals rather than one single human.

The inner and outer loops of the brain are coupled and serve each other. Taking one extra step forward, we can imagine a direct coupling between the biological (the brain) and the physical computers operating around the human body through brain–machine interfaces (BMIs). These allow for a more direct reading of intent or feedback of outcome, resulting in a more effective and efficient overall system performance. A schematic of the interplay between the two computer domains is illustrated in Fig. 3. Experimental data has shown that the BMI performs better and adapts faster when both systems are allowed to "coadapt" [10].

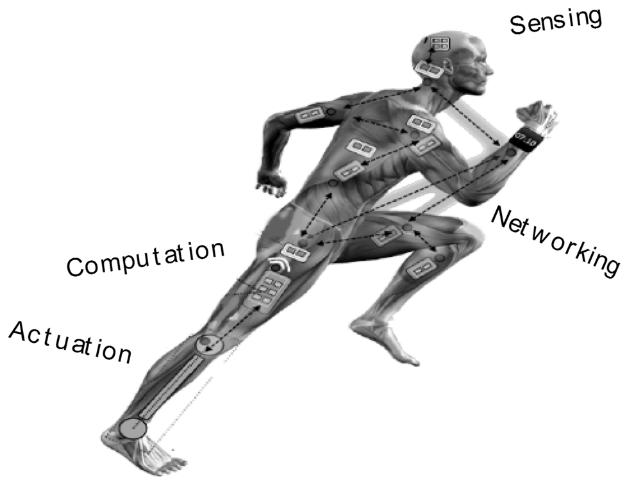

Fig. 4. Human Intranet. Picture courtesy of Y. Khan.

Given this motivation for "human-centric" computing, the rest of the article will be devoted to an analysis of what it will take to make it happen, an identification of the major challenges and barriers in providing intelligence at a human scale and a perspective on possible solutions. A simple case study will help to put it all in perspective.

II. Human Intranet

The human-centric computing paradigm presents a system vision. For example, disease would be treated by chronically measuring bio-signals deep in the body, or by providing targeted, therapeutic interventions that respond on demand and *in situ*. To gain a holistic view of a person's health, these sensors and actuators must communicate and collaborate with each other. The challenge is scale: first in size toward the nanomorphic cell and second in number, toward hundreds of individual motes, each with a unique location and function.

As a framework to enable this, we envision a Human Intranet (Fig. 4): an open, scalable platform that seamlessly integrates an ever-increasing number of sensors and actuators, computation, storage, communication, and energy nodes located on, in, or around the human body and acting in symbiosis with the functions provided by the body itself [11]. Openness is an essential property if the Human Internet is to evolve and to thrive based on the creativity of many contributors—as has been the case in any IT platform such as the PC, the smartphone, and the Internet.

A number of key properties need to be met for humans to accept and embrace a network of nodes immersed around and in their body. Just to name a few: unobtrusive and assimilated, efficient and unburdensome, robust and safe, secure and private, and, most of all, ethical. In addition, the network should easily adapt or be adaptable to changing needs and applications. While some of these properties can be addressed by technology, other concerns may require regulations, policies, or engagement rules to be adopted.

A. Unobtrusive and Unburdensome

Technology scaling is most definitely the largest enabler for the Human Intranet concept. Over less than two decades wireless sense-and-compute nodes have scaled by more than three orders of magnitude in size and energy efficiency. While smaller transistor sizes (Moore's law) have surely helped, most of the scaling in the node size is due to "More than Moore" (MtM) scaling: advanced packaging and integration techniques that allow diverse technologies, such as sensors or energy supplies, to be assembled together with more traditional silicon devices for computation and memory in a 3-D form factor [12]. As a result, fully integrated "smart dust motes" of 1 mm^3 in size, as originally envisioned in 1997 [13], have now been demonstrated [14]. Yet, to be truly unobtrusive, implanted nodes should be even smaller, and extra orders of magnitude of scaling should be pursued. In an ideal scenario, physical motes may become equivalent in size to a biological cell, improving the information flow and reducing the rejection by the body. This requires innovation and creativity. The "neural dust" [15] concept is an example of such.

For nonimplanted nodes worn on the body, conformity and unobtrusiveness are essential factors. We have witnessed a rapid progress over the last few years in the development of flexible devices for sensing, communication, energy harvesting, and energy storage, using techniques such as thin film and printing technologies as well as various forms of smart fibers and patches.

The evolution of BMIs over the last decades serves a perfect example of the trend toward unobtrusiveness. Current state-of-the-art technologies, as used in clinical trials on human patients, still rely on rigid electrode arrays implanted into the cortex and connection to data acquisition and processing hardware via through-the-skull connectors and cables [16], [17]. These systems definitely are not amenable to out-of-clinic daily life utilization and also have a limited operational lifetime. In the past years though, enormous progress has been made in the development of flexible and conformal electrode arrays, low-power data acquisition and neuromodulation ICs and wireless connectivity, and power delivery through the skull. Some of these technologies are now finding their way into medical devices from the major manufacturers and as such will open the door for a much broader set of usage scenarios. Even further down the road, innovative technologies such as "neural dust" [15] will open the door for long-term 24∗7 deployment and may help BMI to transcend from the medical to the consumer space.

B. True Integration of Energy and Information Distribution

Energy sparsity is one of the central challenges in the construction of the Human Intranet. While some nodes (called hubs[1]) may have a sizable energy reservoir in the form of batteries or energy-harvesting capability, others may have zero

[1] The evolved smartphone would be a perfect example of a hub: among others, it provides a bridge to the surrounding world with a diverse set of broadband communication capabilities (5G and beyond), while providing huge computation and data storage capacity. In addition, it serves as an energy reservoir for the rest of the Intranet.

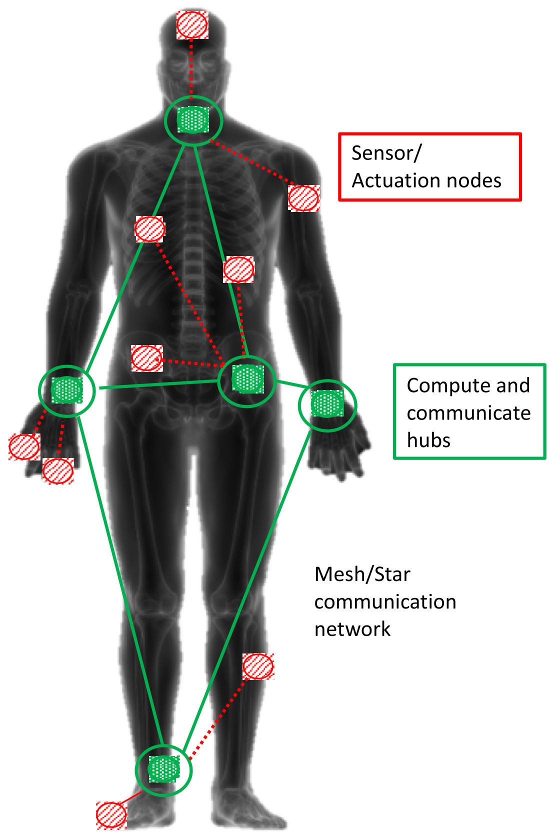

Fig. 5. Data and energy communication skin around the body.

storage and require energy to be provided remotely when information is requested. Paralleling, in a sense, the human nervous and arterial systems, network nodes must collaborate to form a communication "skin" in and around the body (Fig. 5). This could be in the form of a hierarchical and adaptive mesh, delivering energy and/or information over heterogeneous physical links (wired and wireless, electromagnetic, resistive, capacitive, inductive, acoustic, and/or optical [18]), the choice of which is determined by local context: best meeting the needs of the active applications while optimizing the usage of the scarce resources such as energy.

C. Inherently Fail-Safe Operation

Given the often life-critical nature of its applications, it is of essence that basic or partial functionality of the Human Intranet is retained under all circumstances, even if resources fail or are insufficient, when the system goes in overload conditions, or when it is under denial of service attacks. Fail safety must be built in from the ground up and should be an inherent property of both basic components as well as their compositions. Approaches to address this include baselining and safe modes, exploitation of redundancy, and adaptivity and reconfiguration.

D. Secure and Safe

Given the very personal nature of the information being acquired and transmitted as well as the potentially life-threatening effects of indiscriminate or malicious actuation/stimulation, any solution should provide rock-solid privacy and security guarantees. We envision an integrated set of mechanisms such as unique biomarkers and adaptive cloaking [19], jointly labeled as the Human Firewall [20] to ensure that private data, as circulated in the network, is protected and secure, while protecting the network from external intrusions. At the core of any such framework should be solid and mandatory authentication and encryption technologies. When implemented in custom hardware, state-of-the-art encryption can be performed at pJ/bit (similar to what it costs to transmit a bit), while the direct connection to the human body opens the door for bio-parameter-based authentication.

E. Global and Distributed System Intelligence

The Human Intranet operates in a dynamic world subject to both slow evolution and extremely fast changes both in the surrounding environment and in the Intranet itself in activity, conditions, composition, and resource needs and availability. Therefore, the Human Intranet should be constructed as an adaptive and evolutionary system that combines local decision making with centralized global learning and optimization performed in hub nodes (or in the cloud). This approach, in which intelligence is both global and distributed, is essential to deal with issues of latency and single points of failure, while avoiding the trap of many distributed entities with limited knowledge trying to address a global issue. As these considerations are at the core of the human-centric computing paradigm, we delve a bit deeper into what it will take to bring AI to the Human Intranet.

III. AI AT THE EDGE

There are strong incentives to move the processing of sensor data and the generation of actuation patterns for motor functions as close as possible to the end points. Similar arguments hold for closing the feedback loop locally rather than relying on remote processing in the Cloud.

1) *Energy Efficiency:* is the first and foremost consideration. Communication is expensive and wasteful, especially if raw data streams are involved. Nature has figured this out many millions of years ago. A leading principle that has guided the evolution of the brains can be quoted as follows: "Send only what is needed, and send it at the lowest possible rate" [8]. This design principle has, for instance, guided the architecture of the visual and olfactory systems in animals and humans, in which the extraction of the small patterns (features) is performed inside or right next to the sensory arrays ("in-sensor computing"). On a more global scale, "where to perform a given function" is determined by the tradeoff between the relative costs of computation and (wireless) communication, as illustrated in Fig. 6. Since transmitting a bit may be equivalent to many thousands of gate operations, it makes perfect sense to reduce the communication rate by introducing some local processing. At the same time, the computations that can be performed are extremely constrained by the

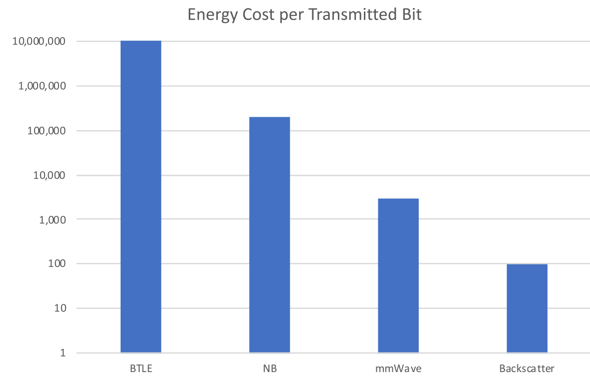

Fig. 6. Typical energy per transmitted bit for low-power wireless technologies normalized to the energy cost of an FO4 NAND gate in a 22-nm CMOS technology (~1 fJ/operation): 1) BT_{LE} at 1 Mbps; 2) Narrowband TX at 10 Mbps; 3) Directional 60-GHz millimeter-wave at 1 Gbps; and 4) Backscatter at 1 Mbps.

available energy. Just as a point of reference, the energy available from a small Zinc-Air Hearing Aid battery is about 200 J (which is equivalent to transmitting 20 GB of data over a BT_{LE} link). This is an order of magnitude lower than what is available in an Apple Watch (~3000 J) and two orders lower than the iPhone X (~40 000 J). Hence, innovative approaches toward "AI at the Edge" are necessary, especially for the wearable and implantable devices.

Some other concerns beyond energy are influential in determining the composition of the system architecture.

2) *Latency:* While many of the regulatory loops in the mammalian body do not require a quick response, a number of the external loops including visual and voice communication and the resulting motor actions certainly do. For those, latencies larger than 10 ms translate into severe degradation in performance and potentially failure. Getting guarantees of this order is extremely hard for Cloud-in-the-Loop solutions.
3) *Robustness and Safety:* As many of the Human Intranet applications directly involve humans, safety is a crucial concern. Solely relying on wireless links to remote computational processing leads to systems that are vulnerable and prone to failures. Some form of autonomy is essential for systems to operate reliably, especially in the presence of failures or major disturbances.
4) *Security and Privacy:* Human-centric computing brings technology as close as possible to the essence of the human being. Data gathered from our body reveals our physical and mental health. Keeping this data local is far more secure than sending it to a remote server. It also helps to preserve privacy. Similarly, inadvertent or malicious actuation of motor functions can be life threatening. Having a firm firewall around the human body and keeping most essential processing inside of it can help to prevent malicious attacks (similar to the body's immune systems) [20].

All this clearly motivates the need for ultralow power (<1 mW) distributed processing. Some of that processing relates to simple tasks such as data acquisition, signal conditioning, and feature extraction, all of which should be located close to the sensors. However, understanding, reasoning, and decisioning require more advanced functionality, a major fraction of which would be learning based. In the last few years, we have witnessed some major advances in the realization of learning-based computing for embedded applications. This has helped to increase energy efficiency to the range of tera operations per second (TOPS)/Watt for deep-net accelerator functions (e.g., [21]). Progress was made using a number of architectural and circuit innovations, including approximate and statistical computing [22], analog processing [23], and in memory computing [24]. These results surely are impressive and help bring AI to the Edge (as described in detail in a keynote presentation at ISSCC in 2019 [25]). However, there are some major concerns. All these deep-net accelerators focus on inference only. Training (learning) requires different operations, is complex and compute hungry, and requires massive amounts of data. It is still far from human-like intelligence, which is based on continuous learning and improvement, learning by equivalence, and very often requires little data to get started. One can argue that similar considerations hold for human-centric computing. Operating in an ever-changing environment, novel situations or changing conditions occur that require quick response without a lot of background training. Similarly, each human is different—hence, there is no one-size-fits-all solution. Customization of the processing for a single human often leads to simpler and more effective solution (see [26]). In short, AI when applied to human-centric computing has to be **plastic** and be capable of continuous evolution. This may require some fundamental rethinking of how computation is performed.

IV. Computing With Patterns

Before selecting any approach to accomplish the above-defined goals, it is worthwhile to peruse the complete AI landscape first. The domain of AI covers any approach that mimics human intelligence, that is, programs that sense, reason, act, learn, and adapt. It is very rich, and spans a broad range of techniques and models (Fig. 7). Many of these approaches fall under the machine-learning (ML) header—those are systems that perform a specific task without using explicit instructions, relying on patterns and inferences instead. One prominent ML approach is the field of Bayesian ML, in which quantities of interest are treated as random variables, and one draws conclusions by analyzing the posterior distribution over these quantities given the observed data.

The neuro- or brain-inspired computing approach forms another important class of ML algorithms. It refers to computational models and methods that are based on abstractions and models of the mechanisms and topologies of the brain. Prominent examples of this class of AI are the artificial neural nets (ANNs), with deep learning networks as the most widespread representative [27]. The latter includes architectures such as deep belief networks, convolutional NNs, and recurrent NNs. These approaches have been immensely successful over a broad range of specific tasks that range from playing

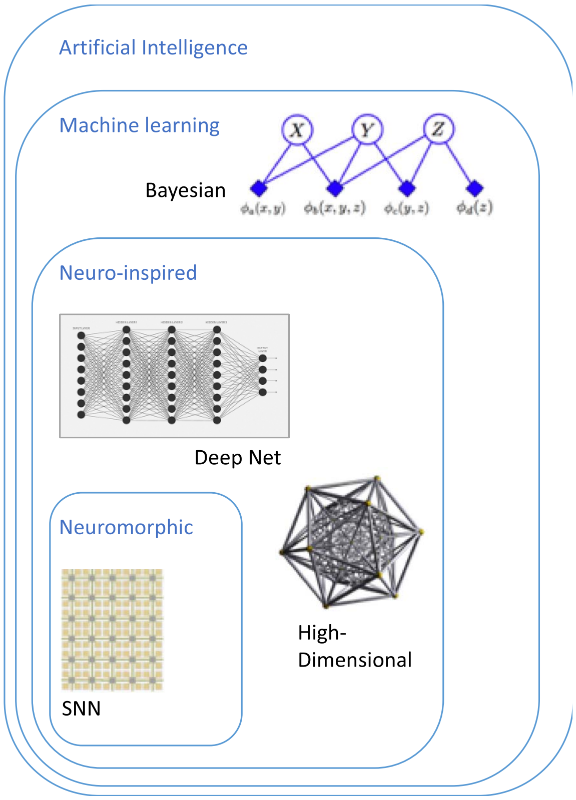

Fig. 7. Coarse classification of AI approaches.

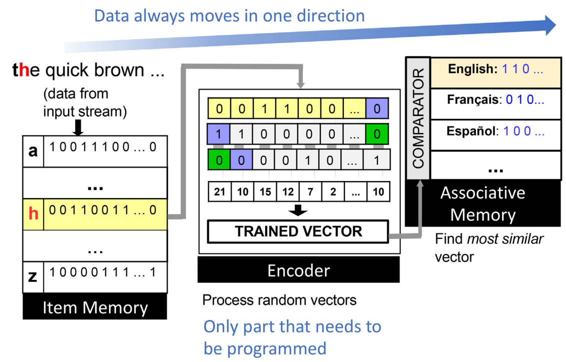

Fig. 8. Data flow in HD processor using language recognition as an example. The stream of input characters is mapped into a HD representation first. Temporal patterns are captured by the encoder, which ultimately creates a single pattern that represents that piece of text. The AM stores the patterns and/or performs search for similar patterns.

games such as Go over facial recognition to autonomous driving. It builds around an abstract model of a "neuron," an interconnect topology, and overlaying learning and inference mechanisms. It is fair to state that the breakthrough of ANNs, after many false starts, was ultimately triggered when huge data sets and fast parallel computing platforms became available.

Finally, the field of neuromorphic computing [28] covers a range of analog, digital, mixed-mode analog/digital very large-scale integration (VLSI), and software systems that implement models of neural systems (for perception, motor control, or multisensory integration). While it technically belongs to the domains of ML and neuro-inspired computing, its inspiration is to build physical computing systems that mimic the operation of the brain in a bottom-up fashion. As such, it presents more of a computing architecture than a computational model. Prominent examples of commercial implementations of neuromorphic computers are the IBM TrueNorth processor [29] and the Intel Loihi processor [30]. The spiking NN (SNN) is one class of neuromorphic NN that has received a lot of attention. Its event-driven executional model makes it particularly attractive for low-energy realizations [31].

It is apparent that any integrated human-centric computing system will combine a selection of AI techniques. Simple classification is often perform best using an ML approach such as a support-vector machine (SVM). Classification or recognition of more complex objects in an image are done very well with a deep NN. Yet, even the broad ensemble of techniques described above does little to address some of the identified needs for human-centric computing, that is the need for plasticity, quick and continuous learning, and dealing with novelty, while at the same time pushing the energy boundaries. Inspiration from the biological brain may help. As illustrated in Fig. 2, the brain operates on *pattern*s. From the sensor arrays, small patterns are extracted that are then gradually composed into larger patterns. While the former happens at the sensor interface, the latter is performed more centrally in the cortex. After comparing these with stored patterns, output behaviors are determined again as patterns, which are then gradually broken down into smaller patterns that regulate motor function. Patterns can be represented in many ways, such as phase, time, or frequency relations between numerous neurons. It is interesting to note that aforementioned approaches such as ANNs and SNNs actually can be considered to performing operations on patterns.

This observation has given rise to a different approach to neuro-inspired computing called high-dimensional computing (HDC) [32]. Rather than using numbers, HDC operates on very long ($D > 1000$) random vectors (that are most often binary) to represent patterns. It builds on the observation that long randomly chosen vectors are (almost) orthogonal. Any deviation from orthogonality means that there is a relation between the patterns encoded in the vectors. To match, it is sufficient that two vectors are "similar" (measured using some norm, such as the Hamming distance). HDC first maps incoming sensory data into high-dimensional (HD) spaces and then proceeds to encode temporal and spatial information by algebraic operations on these vectors. The resulting pattern can be stored in memory, or can be compared against patterns that were already observed (Fig. 8). This can be repeated, creating patterns at different levels of abstraction.

This computational approach comes with some interesting properties, which make it very attractive for ultralow-power operation while supporting plasticity: 1) it is statistical and robust against errors and variations; hence, it can operate at low signal-to-noise ratios; 2) as it is based on algebraic operations on vectors, it provides transparency and allows reasoning about the results; 3) it learns quickly—often a few

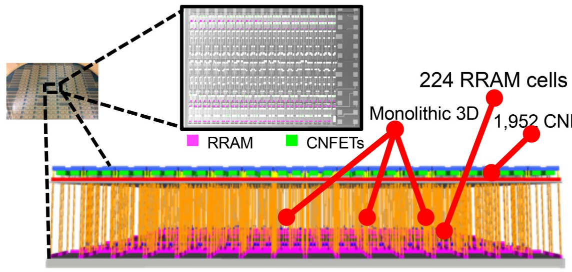

Fig. 9. Monolithic 3-D-integrated HD processor, combining logic and nonvolatile memory (Adopted from [33]).

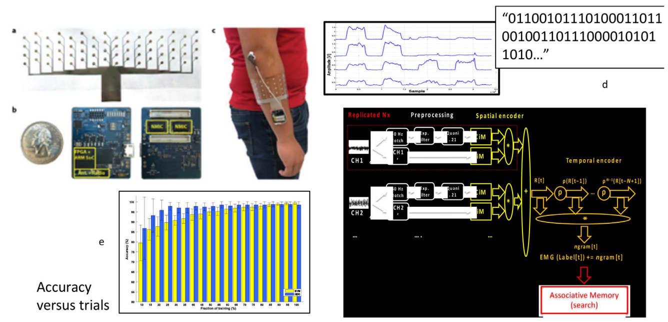

Fig. 10. EMG-based gesture recognition system. (a) EMG signals are obtained from flexible sensor array providing 64 channels. (b) Signals are processed and classified in a small processor system and (c) placed next to the electrodes on the arm. After some preprocessing, sensor data is mapped into an HD space and fed to encoder to capture spatial and temporal information. (d) Patterns representing various gestures are stored and compared gains in an AM. (e) Only a few trials are needed to get high accuracy.

training sessions are sufficient; 4) it can learn incrementally and continuously; 5) it computes in superposition—meaning that data is superimposed on top of each other and can be queried with a single operation; and 6) the central element is the associative memory (AM), which performs search and match. Association is an important part of the mammal brain as well.

Some of the properties mentioned above make HDC a great candidate for realization in deeply scaled technologies supporting true 3-D stacking of devices: robustness against failure and variation; low SNR operation; systolic nearest-neighbor communication only; and in-memory computing (for the AM). A true 3-D realization of an HD processor using carbon-nanotube devices stacked on resistive random access memory (RRAM) memory cell was demonstrated at ISSCC in 2018 (see Fig. 9). While only a limited prototype, it goes a long way in demonstrating the long-term feasibility of advanced AI devices at nanoscales and ultralow power [33].

The above hopefully serves to show that there is a huge opportunity still in the exploration of computational approaches inspired by the mammal brain to create computational modules that meet the tight constraints imposed by human-centric computing, while at the same time exploiting the most advanced approaches semiconductor technology got to offer.

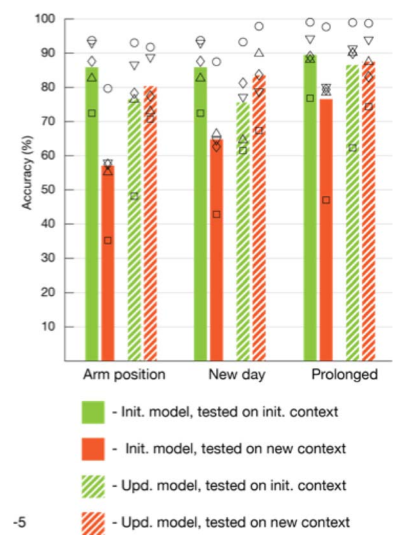

Fig. 11. Gesture recognition accuracy drops when conditions are changed. A short retraining and merging the new experiences with the old restores the accuracy under both conditions. (Accuracy measured over 21 gestures and multiple subject.)

V. Case Study

To bring together the many aspects of human-centric computing and the Human Intranet, as they converge into a single system, we present an example of an electromyography (EMG)-driven hand-gesture recognition system. Accurate recognition of hand gestures is crucial to the functionality of smart prosthetics and other modern human–computer interfaces. Many ML-based classifiers use EMG signals as input features, but they often misclassify gestures performed in different situational contexts (changing arm position, reapplication of electrodes, etc.) or with different effort levels due to changing signal properties.

To combat this variability, an end-to-end system was designed using a large-area, high-density sensor array, an integrated acquisition system, and a robust classification algorithm (Fig. 10) [34]. The EMG electrodes were fabricated on a flexible substrate and interfaced to a custom wireless device for 64-channel signal acquisition and streaming. HD computing was used for the processing of the EMG features using a one-shot learning approach. The HD algorithm is tolerant to noise and electrode misplacement and can quickly learn *in situ* from few trials without gradient descent or back propagation: an average classification accuracy of 96.64% for five gestures is obtained, with only 7% degradation when trained and tested across different days. The system maintains this accuracy when trained with only three trials of gestures; it also demonstrates comparable accuracy with the state of the art when trained with one trial.

This architecture easily supports incremental and continuous learning, as all the learned data is stored in the AM. When conditions change (for instance, the sensor array moves, or the arm is held in a different position), new patterns can be added to the AM, or merged with existing patterns, maintaining accuracy over all the experiences, as illustrated in Fig. 11.

VI. Perspectives

This article presented a vision on how physical and biological computing are on a convergence path. The realization of this vision to its full extent will take many more decades. It requires innovation on many fronts including interfaces, processing, communications, and energy provisioning. Today, though, we already see instances emerge that give a glimpse of what may be possible in the long term: with the advent of the smart watch, wearables are becoming common place; wearable medical devices such as smart patches are making their inroads; hearing aids are starting to integrate multiple sensors and incorporate AI [35]; AR/voltage regulator (VR) glasses help to transform the way humans see the world around them; and BMIs have been demonstrated in laboratory and clinical settings. At the same time, these systems are mostly standalone and *ad hoc* and do not profit from the wealth of information that could be available in a shared platform. It is only when an open platform such as the Human Intranet is adopted that we will see the full power of human-centric computing, serving as a full complement to our brain.

With this evolution come many questions and concerns that need be addressed upfront if one wants to avoid a severe societal backlash. It is extremely likely that a majority of the technology challenges identified above will be resolved in the coming decade(s). However, this is only a necessary and not a sufficient condition for human-centric computing and the Human Intranet to be readily and broadly adopted. Technology that extends or modifies human capability is not always readily accepted. Even more, it is often met with fierce resistance. While this is a less of an issue when the technology addresses severe illness or health issues (the inner loops), the spread of technologies that complement or augment the human capabilities (the outer loops) raise many ethical questions that need be addressed. Hence, there is no doubt that policies and operational rules should be at the core of the human-centric computing paradigm. The same holds for the ways personal data is processed, handled, and shared. Before rushing a technology to the marketplace, we, as technologists, should reflect on the many ways that technology can be misused and incorporate means to deflect.

Acknowledgment

The author would like to thank the many contributions made by his colleagues and his students to the vision and the ideas presented in this article. He is very indebted to the many funding agencies that made this research possible over the years. These include the Defense Advanced Research Projects Agency (DARPA), Semiconductor Research Corporation (SRC) [under the Focus Center Research Program (FCRP), the StarNet program, and the JUMP program], NSF, and the Philippines–California Advanced Research Institutes (PCARI). He would also like to thank the member companies of the Berkeley Wireless Research Center (BWRC). The material contained in this article builds on keynote presentations at major conferences, such as the IEEE Microwave Week (IMS) 2016 [1], VLSI Design Automation and Test (DAT)/Technology, Systems and Applications (TSA) symposia 2018 [2], and the IMEC International Technology Forum (ITF) 2018 [3]. Its presented ideas result from research sponsored by SRC/DARPA under the TerraSwarm and CONIX programs, DARPA Microsystems Technology Office (MTO) and Biological Technologies Office (BTO), NSF, and the member companies of BWRC.

References

[1] J. M. Rabaey, “The human intranet—Where swarms and humans meet,” in *Proc. Design, Autom. Test Eur. Conf. Exhib.*, Grenoble, France, Mar. 2015, pp. 637–640.

[2] J. M. Rabaey, “Homo Technologicus,” in *Proc. Int. Symp. VLSI Technol., Syst. Appl. (VLSI-TSA)*, Hsinchu, Taiwan, Apr. 2016, p. 1.

[3] J. Rabaey, “The dawn of true human-centric computing,” presented at the IMEC ITF Conf., Antwerp, Belgium, May 2018.

[4] E. A. Lee *et al.*, “The swarm at the edge of the cloud,” *IEEE Design Test*, vol. 31, no. 3, pp. 8–20, Jun. 2014.

[5] R. Collins. (Dec. 2015). “Roadmap to trillion sensors forks,” *EE Times.* [Online]. Available: https://www.eetimes.com/document.asp?doc_id=1328466#

[6] D. T. Max, “How humans are shaping our own evolution,” *Nat. Geogr. Mag.* Apr. 2017.

[7] *Human-Centered Computing*. Accessed: Apr. 2018. [Online]. Available: https://en.wikipedia.org/wiki/Human-centered_computing

[8] P. Sterling and S. Laughlin, *Principles of Neural Design*. Cambridge, MA, USA: MIT Press, 2018.

[9] M. Banf and V. Blanz, “Sonification of images for the visually impaired using a multi-level approach,” in *Proc. 4th Augmented Hum. Int. Conf.*, Mar. 2013, pp. 162–169.

[10] A. L. Orsborn, S. Dangi, H. G. Moorman, and J. M. Carmena, “Closed-loop decoder adaptation on intermediate time-scales facilitates rapid BMI performance improvements independent of decoder initialization conditions,” *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 20, no. 4, pp. 468–477, Jul. 2012.

[11] J. M. Rabaey, “The human intranet—Where swarms and humans meet,” *IEEE Pervasive Comput.*, vol. 14, no. 1, pp. 78–83, Jan./Mar. 2015.

[12] H. De Man, “Ambient intelligence: Gigascale dreams and nanoscale realities,” in *Proc. ISSCC*, Feb. 2005, pp. 29–35.

[13] B. A. Warneke *et al.*, “An autonomous 16 mm^3 solar-powered node for distributed wireless sensor networks,” in *Proc. IEEE Sensors*, Jun. 2002, pp. 1510–1515.

[14] G. Templeton. (May 2013). *Smart Dust: A Complete Computer That’s Smaller Than a Grain of Sand*. Extreme Tech. [Online]. Available: https://www.extremetech.com/extreme/155771-smart-dust-a-complete-computer-thats-smaller-than-a-grain-of-sand

[15] D. Seo *et al.*, “Wireless recording in the peripheral nervous system with ultrasonic neural dust,” *Neuron*, vol. 91, no. 3, pp. 529–539, Apr. 2016.

[16] J. P. Donoghue, “Connecting cortex to machines: Recent advances in brain interfaces,” *Nature Neurosci.*, vol. 5, pp. 1085–1088, Oct. 2002.

[17] P. P. Harris. (Apr. 2011). *Braingate Gives Paralysed the Power of Mind Control*. The Guardian. [Online]. Available: https://www.theguardian.com/science/2011/apr/17/brain-implant-paralysis-movement

[18] G.-Z. Yang, *Body Sensor Networks*. London, U.K.: Springer-Verlag, 2014.

[19] S. Gollakota, H. Hassanieh, B. Ransford, D. Katabi, and K. Fu, “They can hear your heartbeats: Non-invasive security for implantable medical devices,” in *Proc. ACM SIGCOMM*, 2011, pp. 2–13.

[20] G. Slack, “The last firewall,” in *Berkeley Engineer Magazine*, vol. 5. 2014. [Online]. Available: https://engineering.berkeley.edu/sites/default/files/docs/BENG-Spring2014.pdf

[21] B. Moons, R. Uytterhoeven, W. Dehaene, and M. Verhelst, “14.5 envision: A 0.26-to-10 TOPS/W subword-parallel dynamic-voltage-accuracy-frequency-scalable convolutional neural network processor in 28 nm FDSOI,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2017, pp. 246–247.

[22] N. R. Shanbhag, N. Verma, Y. Kim, A. D. Patil, and L. R. Varshney, “Shannon-inspired statistical computing for the nanoscale era,” *Proc. IEEE*, vol. 107, no. 1, pp. 90–107, Jan. 2019.

[23] D. Bankman *et al.*, "An always-on 3.8 μJ/86% CIFAR-10 mixed-signal binary CNN processor with all memory on chip in 28-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 54, no. 1, pp. 158–172, Jan. 2019.
[24] N. Verma *et al.*, "In-memory computing: Advances and prospects," *IEEE Solid State Circuits Mag.*, vol. 11, no. 3, pp. 43–55, Jun. 2019.
[25] H.-J. Yoo, "1.2 intelligence on silicon: From deep-neural-network accelerators to brain mimicking AI-SoCs," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2019, pp. 20–26.
[26] A. Rahimi, P. Kanerva, L. Benini, and J. M. Rabaey, "Efficient biosignal processing using hyperdimensional computing: Network templates for combined learning and classification of ExG signals," *Proc. IEEE*, vol. 107, no. 1, pp. 123–143, Jan. 2019.
[27] Y. LeCun, "1.1 deep learning hardware: Past, present, and future," in *IEEE ISSCC Dig. Tech. Papers*, Feb. 2019, pp. 12–19.
[28] C. Mead, "Neuromorphic electronic systems," *Proc. IEEE*, vol. 78, no. 10, pp. 1629–1636, Oct. 1990, doi: 10.1109/5.58356.
[29] S. Ravindran. *Building a Silicon Brain*. Accessed: May 2019. [Online]. Available: https://www.the-scientist.com/features/building-a-silicon-brain-65738
[30] (Oct. 17, 2017). *Intel Unveils Loihi Neuromorphic Chip, Chases IBM in Artificial Brains*. [Online]. Available: https://www.aitrends.com/future-of-ai/intel-unveils-loihi-neuromorphic/
[31] A. Balaji, F. Corradi, A. Das, S. Pande, S. Schaafsma, and F. Catthoor, "Power-accuracy trade-offs for heartbeat classification on neural networks hardware," *J. Low Power Electron.*, vol. 14, no. 4, pp. 508–519, 2018.
[32] P. Kanerva, "Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors," *Cogn. Comput.*, vol. 1, no. 2, pp. 139–159, 2009.
[33] T. F. Wu *et al.*, "Brain-inspired computing exploiting carbon nanotube FETs and resistive RAM: Hyperdimensional computing case study," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, Feb. 2018, pp. 492–494.
[34] A. Moin *et al.*, "An EMG gesture recognition system with flexible high-density sensors and brain-inspired high-dimensional classifier," in *Proc. IEEE Int. Symp. Circuits Syst. (ISCAS)*, Florence, Italy, May 2018, pp. 1–5.
[35] D. Wang, "Deep learning reinvents the hearing aid," *IEEE Spectr.*, vol. 54, no. 3, pp. 32–37, Dec. 2016.

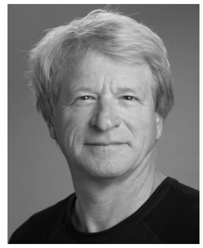

Jan M. Rabaey (S'80–M'83–SM'92–F'95) was a Research Manager with IMEC, Leuven, Belgium, from 1985 to 1987. He is currently the Founding Director of the Berkeley Wireless Research Center (BWRC) and the Berkeley Ubiquitous SwarmLab. In 2019, he holds the CTO of the System-Technology Co-Optimization (STCO), Division of IMEC. He is also a Professor of the Graduate School, Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, Berkeley CA, USA, after being the holder of the Donald O. Pederson Distinguished Professorship for over 30 years. He has made high-impact contributions to a number of fields, including advanced wireless systems, low-power integrated circuits, mobile devices, sensor networks, and ubiquitous computing. His current research interests include the conception of the next-generation distributed systems, as well as the exploration of the interaction between the cyber and the biological world.

Prof. Rabaey is a member of the Royal Flemish Academy of Sciences and Arts of Belgium. He has received honorary doctorates from Lund (Sweden), Antwerp (Belgium), and Tampere (Finland). He has been involved in a broad variety of startup ventures, including Cortera Neurotechnologies, of which he was a co-founder. He was a recipient of major awards, among which the IEEE Mac Van Valkenburg Award, the European Design Automation Association (EDAA) Lifetime Achievement Award, the Semiconductor Industry Association (SIA) University Researcher Award, and the SRC Aristotle Award. He served as the Division Chair for the Electrical Engineering at Berkeley.