The Future of Bionic Limbs: The Untapped Synergy of Signal Processing, Control, and Wireless Connectivity

Federico Chiariotti[1,2], Pranav Mamidanna[3,4], Suraj Suman[2,5], Čedomir Stefanović[2], Dario Farina[4], Petar Popovski[2], and Strahinja Došen[6]

[1]Department of Information Engineering, University of Padova, Padua, Italy (chiariot@dei.unipd.it)
[2]Department of Electronic Systems, Aalborg University, Denmark (cs@es.aau.dk. petarp@es.aau.dk)
[3]I-X Center for AI in Science, Imperial College London, U.K. (p.mamidanna22@imperial.ac.uk)
[4]Department of Bioengineering, Imperial College London, U.K. (d.farina@imperial.ac.uk
[5]Department of Electrical Engineering, IIT Patna, India (ssu@iitp.ac.in)
[6]Department of Biomedical Engineering, Aalborg University, Denmark (sdosen@hst.aau.dk)

The flexibility and dexterity of human limbs rely on the processing of a vast quantity of signals within the sensory-motor networks in the brain and spinal cord, distilled into stimuli that govern the commands and movements. Hence, the use of assistive devices, such as robotic limbs or exoskeletons, is critically dependent on the processing of a large number of heterogeneous signals in order to mimic natural movements. This paper provides a panoramic overview of the three paradigms for the control of bionic limbs based on mechatronic technology. Two of them have already been established in the literature, while the third one, advocated by this paper, is an emerging approach, enabled by the latest developments in connectivity and computation. In the *first* paradigm, the bionic limbs rely on conventional control and are directly reconnected to the human sensory-motor system, which requires a large signal processing bandwidth. The *second* paradigm is based on semi-autonomous limbs, endowed with context-aware processing and certain decision capability. Following the advances in wireless connectivity and cloud/edge processing, this paper introduces a *third* paradigm of connected limbs. A connected limb goes beyond replacement and towards augmentation of human motor skills, offering the possibility for remote interactions and interventions from medical professionals, as well as updating the limb based on large-scale learning. Besides the generic use cases for connected limbs, we present early prototypes and results, along with the associated signal processing challenges that emerge in distributed settings with communication

latency and transmission impairments. The overall conclusion is that the concepts and uses of connected limbs promise an interesting new era for this assistive technology.

1. Introduction

The human ability to perform a wide range of movements smoothly and effortlessly is truly remarkable. Even simple daily tasks involve the coordinated and controlled activation of many degrees of freedom using a highly redundant set of actuators. Despite this complexity, natural movements are executed effortlessly, mostly without conscious thinking. This is achieved by sensory-motor networks in the brain and spinal cord, which implement distributed and parallel processing of incoming stimuli to generate motor commands. In healthy humans, there is a bidirectional connection between the neural controller - the central nervous system - and the limb muscles, via the spinal cord. The controller sends motor commands to the spinal circuits that transmit them to the limb muscles to initiate movements and receives feedback from the limb via the spinal cord. In a recent anatomical study of nerve specimens from human cadavers [1], it was found that ~350,000 nerve axons emerge from the spinal cord to innervate the upper limb. Each second, tens of impulses carry information by traveling through each of these nerve axon fibers, resulting in an extremely large bidirectional communication bandwidth between the limb and the central nervous system. Mechatronics has substantially advanced in the last decades to a level where wearable assistive devices such as robotic limbs and exoskeletons can now mimic their biological counterparts in dexterity, size, and shape, including high precision distributed sensing capacity [2]. However, current human-machine interfaces (HMIs) cannot decode enough commands and/or provide sufficient feedback to allow the users to effectively exploit the rich functionality of the modern bionic limbs [3]. Hence, the limited transmission bandwidth of human-machine interfaces is the main barrier for assistive devices to be truly integrated into the human body.

This paper provides a perspective on the signal-processing approaches that have been instrumental in evolving the technology of bionic limbs. It covers three paradigms for designing bionic limbs beyond the early mechanical devices. The *first paradigm* is characterized by bionic limbs with conventional control, in which the limbs are directly reconnected to the human sensory-motor system. Here the control intelligence remains entirely with the human subject, while the controller is responsible for converting the human-generated signals into commands actuated by the prosthesis. The role that signal processing plays in conventional control is to combine various information acquisition modalities and offer an effective increase in information processing capability. The *second paradigm* introduces a certain level of decision-making and intelligence in the prosthetic limb, promoting them into semi-autonomous systems. Here the controller of the limb performs context-aware processing and partially relieves the burden of cognitive control from the user. Semi-autonomous limbs rely on signal processing for multimodal data fusion, such as, e.g., the decoding of electromyography (EMG) for volitional control and the analysis of RGB images and/or point clouds for autonomous decision-making. These first two paradigms have been extensively investigated and reviewed [3], [4], [5]. In this paper, we therefore do not review in detail these approaches but offer innovative and original formulations and perspectives that encompass the core aspects of the field. The *third paradigm* advocated in this paper is a novel approach in which the prosthesis technology leverages ubiquitous wireless connectivity and the abundance of computing power residing in the edge/cloud to upgrade the existing and offer new functionalities. Indeed, while the first two paradigms have classically focused on the replacement of natural limbs, connected limbs offer not only replacement but also augmentation of human motor skills. For instance, a connected limb can offload part of the computational or cognitive load to a powerful edge/cloud server, or a large number of limbs can use the connectivity to share data and train a common machine-learning model. We present early prototypes of such systems and discuss several novel use cases brought by the connected limbs and elaborate upon their signal processing challenges in distributed settings, dealing with latency and error-prone data transmissions. Figure 1 shows

the three paradigms and summarizes their fundamental features, which are then explained in more detail in the following sections.

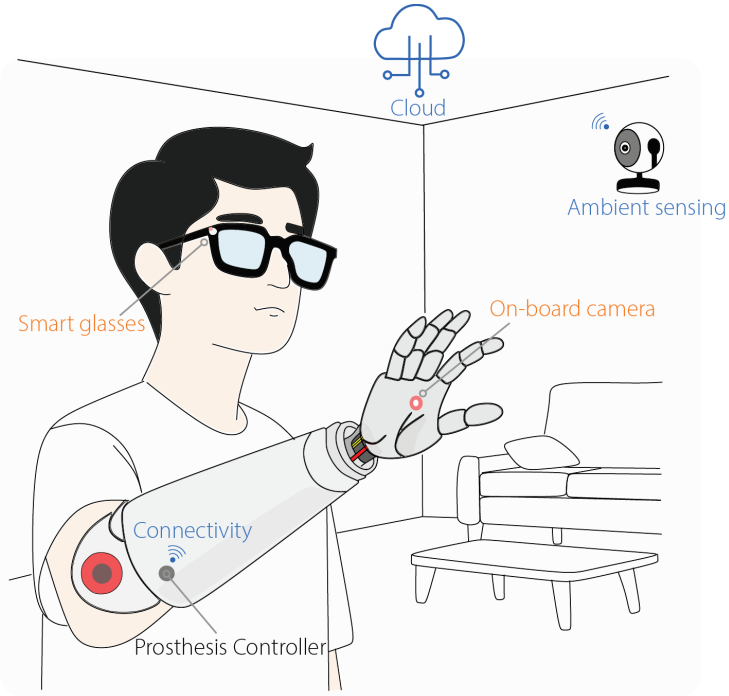

	Conventional	Semi-autonomous	Connected
Control responsibility	Solely on user	Shared with agent (local learning)	Shared with agent (decentralized learning)
Sensing capability	Bio-signals	Bio-signals On-board sensing (cameras, touch sensors)	Bio-signals On-board sensing Ambient sensing
Processing capability	On-board	On-board, edge	Cloud, edge
Learning	Offline, on-board	Offline, on-board	Offline, on-board, Decentralized
Control bandwidth	Low	Medium	High
Cognitive load	High	Low	Low
Challenges	Intention decoding	Shared autonomy Coadaptation Embodiment	Shared autonomy Wireless latency Decentralized learning Embodiment

Fig. 1 The summary of the three paradigms for prosthesis control. In conventional control (gray), the prosthesis controller is a decoder that translates user motion intention into prosthesis commands. A semi-autonomous prosthesis (orange) is equipped with additional sensors (e.g., a camera in the prosthesis or smart glasses) that allow the device to perceive the environment and perform some tasks automatically.

Enhancing these systems with connectivity (blue) enables exploiting Cloud/Edge resources and the Internet of Things to enhance the connected limbs with radically new capabilities.

The paper is organized as follows. Section 2 summarizes the state of the art in conventional control, emphasizing recent approaches that can be used to increase information throughput between the human user and the assistive device. Section 3 introduces the concept of semi-autonomous control, the structure of an ideal controller, and describes several representative prototypes. Sections 4 and 5 present the core contributions of this paper, introducing the next step in the development of smart bionic limbs that are enhanced with connectivity. Radically novel functions enabled by Edge/Cloud access are presented in Section 4, while Section 5 analyses the communication requirements and the most significant trade-offs in real-time operation, as well as describes approaches that can be used to investigate user experience when interacting with smart connected systems. Section 6 concludes the manuscript and summarizes the focus points for further research and development.

2. Conventional control

Ideally, the control of a prosthetic limb by the user should be intuitive and effortless, through the establishment of a bidirectional interface that replicates the natural sensory-motor control. Such an interface would enable seamless feedforward transfer of neural commands from the user to their bionic limbs and feeding back the sensory data translated into neural impulses to the central nervous system to elicit sensations. In this approach, the artificial controller is “reduced” to a component that records signals from the human user and translates those signals into commands for the prosthesis, while the decision-making (system intelligence) is fully allocated to the human [6]. Such allocation of responsibilities seems natural, but it also entails a high cognitive burden of control, as the user needs to generate commands for each and every motion of the system. Mathematically, the control problem in conventional control takes the following form:

$$Minimize: \quad J(x, x_{goal}) \tag{1}$$
$$Subject\ to: \quad \dot{x}_{full} = f(x_{full}, u; \phi) \tag{1a}$$
$$n_{full} = \pi(\tilde{x}, x_{goal});\ \ \tilde{x} = x + \epsilon_f \tag{1b}$$
$$u = \ D(\hat{n}; \theta);\ \ \hat{n} = n + \epsilon_m \tag{1c}$$

where x_{full} is the state of the bionic limb, x_{goal} is the desired state that the user wants to attain and $J(\cdot)$ is a loss function (e.g., the difference between the current and desired prosthetic hand configuration). Expression (1a) is the dynamic model of the bionic limb, where u are the control signals generated by the artificial controller (human motion intention decoder, D) and ϕ is a set of parameters. The expression (1b) models the prosthesis user, who generates control signals n_{full} that originate in the brain, using a control policy π that depends on the goal and the state of the bionic limb (e.g., activate muscles to rotate and/or close the prosthesis to reach the desired hand configuration). However, the user does not have access to the full state x_{full} of the artificial system but only to a subset of states x that can be estimated from the incidental (sound, vision) and artificial sensory feedback (haptic stimulation) that the user receives, with uncertainty determined by the feedback noise ϵ_f. Similarly, not all neural signals generated by the user can be observed and, therefore, decoder D (parameterized by θ) estimates in expression (1c) the prosthesis control signals from the subset n of the full set of neural control signals n_{full} that can be measured using biosensors with a measurement noise ϵ_m. The cost function and control policy are generally user and task-specific since each user perceives and reacts differently to the uncertainties involved.

As indicated in the expression (1) and shown in Fig. 2, the challenge in conventional control of bionic limbs is that it implies dimensionality reduction in both states ($x_{full} \rightarrow \tilde{x}$) and control signals ($n_{full} \rightarrow \hat{n}$), and consequently, the number of commands that can be decoded is typically lower than the number of measured biological signals and controllable degrees of freedom in the bionic limb. The most common approach to implement D is to record electrical muscle activity from the residual limb via surface electromyography (EMG). Using simple processing (D comprises rectification and filtering), the stochastic EMG signals are

converted into smooth and proportionally modulated control signals for prosthesis motors (direct control). The limitation of simple direct control is that the number of commands that can be decoded is equal to the number of muscles that the user can independently control in an intuitive manner, and this approach fails when more degrees of freedom need to be controlled.

However, multiple functions can be decoded by classification of signal patterns (in this case, D is a pattern classifier), rather than direct control of the actuators [7]. In principle, the motor intention associated with a specific task corresponds to EMG signal patterns different from those of other tasks. EMG signals can therefore be projected in a feature space and classified into a large number of intended tasks. While EMG signal classification provides an intuitive way to select commands, the way in which the commands are classified and transmitted to the prosthesis is inherently sequential, while natural control involves the concurrent activation of degrees of freedom. Such simultaneous control can be replicated at a basic level by employing regression methods (now, D is a regressor), which estimate multiple simultaneous control signals from the computed features. However, the reliability of current algorithms for simultaneous control decreases dramatically with the number of degrees of freedom. In addition, the clinical translation of classification and regression is plagued by several factors prominent in daily life, such as the change in muscle patterns due to muscle fatigue, electrode shift, and changing limb position, as well as the need to collect data to train the classifier/regressor. Consequently, pattern classification systems have been brought to the market only recently, despite decades of research, while the use of regression-based control is still confined to research prototypes.

More recent developments in conventional control [8] were directed towards increasing the number and quality of signals that can be recorded from the human body (increasing $dim(n)$ and decreasing ϵ_m) and/or improving the decoder D (increasing $dim(u)$ and decreasing estimation uncertainty). Some examples of the former are, for instance, novel sensing approaches, such as sonomyography, and invasive recording

techniques [9], while deep learning, high-density EMG decomposition into neural impulses via blind source separation, or the use of neuromusculoskeletal modeling are the representative of the latter. In addition, recent efforts are directed toward improving methods to provide artificial sensory feedback [10] in terms of the number (increasing dim($\tilde{x}$)) and quality (decreasing ϵ_f) of feedback variables that can be conveyed to the prosthesis user [11], [12].

While these approaches are promising in terms of increasing the control and feedback bandwidth, they are based on recording a large quantity of data that needs to be processed in real time, requiring significant computational power. However, this task is too complex for the embedded hardware of the current bionic limbs, which are constrained in terms of computational and data storage resources. Similarly, surgical techniques are always associated with the risk of potential complications, they are not widely accessible, and some users might be reluctant to undergo surgery.

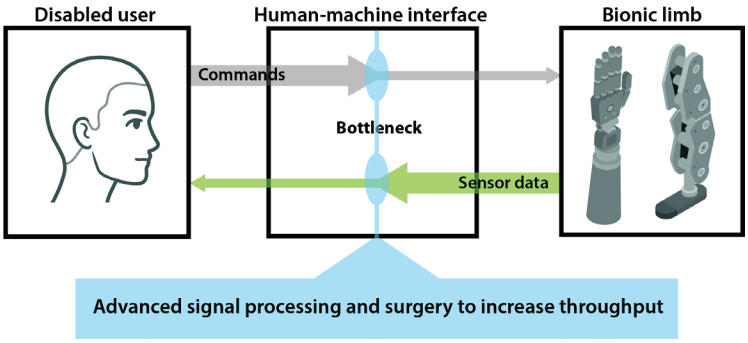

Fig. 2 Human-machine interface is the real bottleneck to achieving effective control of bionic limbs. Humans can generate many commands (thick grey arrow) but only a fraction can be decoded and translated into limb commands (thin grey arrow). Similarly, bionic limbs are equipped with many sensors (thick green arrow) but only limited information can be transmitted as feedback to the user (thin green arrow). Advanced signal processing and surgical techniques can be used to increase the information throughput, but they are not easy to translate into clinical applications.

3. Semi-autonomous bionic limbs

As explained above, the main bottleneck in the development of efficient bionic limbs is the difficulty in establishing a reliable, high-bandwidth, bidirectional connection between the biological and artificial system. An alternative, radically different approach to address this challenge is to make the prosthetic limbs smarter (Fig. 1, semi-autonomous control). To this aim, the system is equipped with (i) advanced sensors to perceive the environment as well as user behavior, and (ii) powerful data processing capabilities to analyze and interpret this information and to make autonomous decisions on how to implement high-level goals of the user [4]. This is a different conceptualization of an assistive system, promoting the robotic controller from a simple decoder of the user intent within the conventional leader-follower scheme, into an intelligent agent collaborating with the user in accomplishing functional tasks. Within this approach, the user transmits information only about a high-level goal (e.g., grasping a bottle) while the details of the movements (e.g., specific approach trajectory, grasp type, and size) are planned and executed by the prosthesis controller, thereby decreasing the cognitive burden of control. In line with the formalism we used to model conventional control, semi-autonomous control can be formulated as:

$$\begin{aligned}
&Minimize: && J\left(x, x_{goal}\right) && (2)\\
&Subject\ to: && \dot{x}_{full} = f\left(x_{full}, u; \phi\right) && (2a)\\
& && u = g(u_v, u_a) && (2b)\\
& && n_{full} = \pi_v\left(\tilde{x}, x_{goal}\right);\ \tilde{x} = x + \epsilon_f && (2c)\\
& && u_v = D(\hat{n}; \theta_D);\ \hat{n} = n + \epsilon_m && (2d)\\
& && u_a = \pi_a\left(x_{full}, \hat{x}_{goal}; \theta_\pi\right);\ \hat{x}_{goal} = E\left(x_{full}, s; \theta_E\right) && (2e)
\end{aligned}$$

where, in (2b) the final control command u is determined as a function g of the volitional commands issued by the human u_v and the automatic controller u_a. From (2d), the way u_v is generated remains the same as in the conventional control. According to (2e), the automatic controller estimates the desired state $\hat{x}_{goal}$ through a learnable mapping E parameterized by θ_E, and using extrinsic sensors s and the current state of the limb x_{full}. Subsequently, it generates the control signal u_a through a learned policy π_a parameterized by

θ_π. Note that the artificial controller has access to the full state x_{full} (contrary to the user, who perceives a subset of "observable" states) and that both agents (user and artificial controller) define their own goals (x_{goal} and $\hat{x}_{goal}$), which are then used to drive their decisions. The important aspect of the semi-automatic control approach is that the requirements on the decoder $D(\hat{n}; \theta_D)$ can be substantially relaxed, as most of the control signals can be supplied automatically ($dim(u_a) > dim(u_v)$) and/or automatic commands can be used to refine those estimated voluntarily. This approach can also overcome the intrinsic bottleneck of the decoder in terms of control dimensionality, since those signals that cannot be estimated from the biological measurements can be relegated to the automatic controller. Indeed, a recent publication [13] demonstrated simultaneous control of 4 degrees of freedom of a prosthesis, which is well beyond what can be achieved using conventional approaches presented in Section 2. Finally, the negative impact of limited feedback delivered to the user (x_{full} reduced to $\tilde{x}$) can be decreased as well because the control responsibility is shared between the two agents. The use of shared control to accommodate high-density tactile data for the control and feedback in robotic hands is discussed in detail in [14].

The development of such smart bionic limbs can profit significantly from adopting and integrating insights from "general" robotics, where the sensing technology as well as the methods for automatic reaching and grasping have been developing for decades. However, the context is still different and, therefore, an automatic translation would not work. In robotics, the artificial controller has full control of all the DoFs of a robot, whereas in prosthetics, some DoFs are controlled by the user's body. For instance, a transradial amputee is responsible for moving and orienting the prosthesis in space, while the prosthesis controller drives the wrist and hand DoFs.

Finally, although at first glance, integrating an "electronic brain" into a prosthesis can look unnatural, a more thorough consideration reveals that, functionally, this is not that different from biological motor control. Grasping in non-disabled people is performed mostly subconsciously and the subject is normally

aware only of the high-level goal. The person decides to grasp an object while the movement unfolds "automatically", which is also the main idea of the semi-autonomous approach to the control of bionic limbs.

The anatomy of a semi-autonomous controller

The structure of an ideal controller of a smart prosthesis is shown in Fig. 3 while two representative prototypes are depicted in Fig. 4. The MYO-PACE system [13] demonstrates the full potential of semi-autonomous control but uses external components and the setup is therefore cumbersome, whereas KIT prototype [15] shows that the critical components of a smart prosthesis can be integrated into the hand leading to a powerful yet self-contained device. In general, the critical functions for the emergence of semi-automatic control are pervasive sensing (artificial extero- and proprioception, s in eq. (2)), cognitive-like processing and decision-making $E(x_{full}, s; \theta_E)$, and fusion $g(u_v, u_a)$ of automatic $\pi_a(x_{full}, \hat{x}_{goal}; \theta_\pi)$ and volitional $D(\hat{n}; \theta_D)$ control.

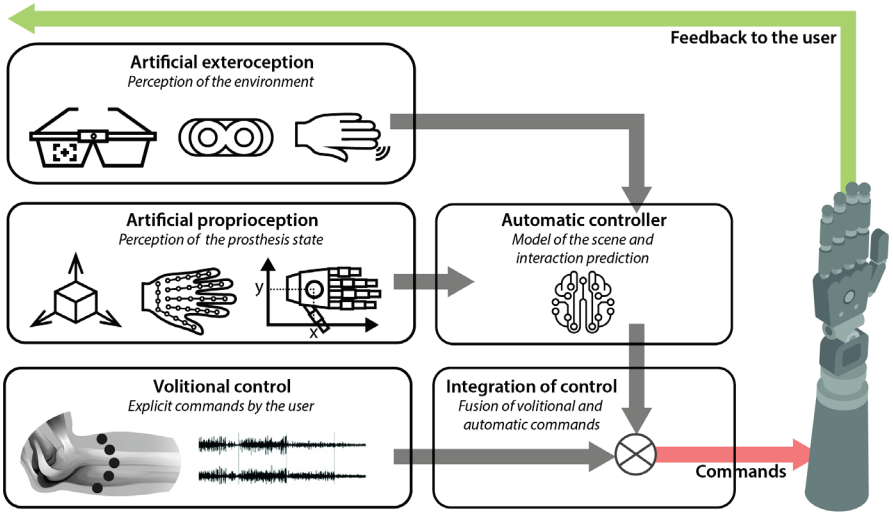

Fig. 3 The structure of an ideal controller of a semi-automatic bionic limb. The information about the environment and the prosthesis state provided by the artificial exteroception and proprioception modules are analyzed using

cognitive-like processing to estimate the likely point of interaction with the environment and the properties of the target object. Based on this, the automatic controller generates commands for the prosthesis, which are fused with the volitional control of the user, and then transmitted to the bionic limb. Finally, to establish bidirectional control, the prosthesis state is conveyed as feedback to the user via tactile, auditory, or visual stimulation.

Proprioception is the biological sense of limb position and movement. Humans can close their eyes and still know where their hands and arms are in space. Analogously, the role of the *artificial proprioception module* is to estimate the position and orientation of the prosthetic limb in space as well as its internal configuration. The latter is relatively simple to obtain, as most modern prostheses are sensorized with encoders measuring joint angles. Depending on the hand dexterity, this includes wrist position, hand aperture, and/or joint angles of individual fingers. The orientation of the limb can be measured by placing inertial measurement units (IMUs).

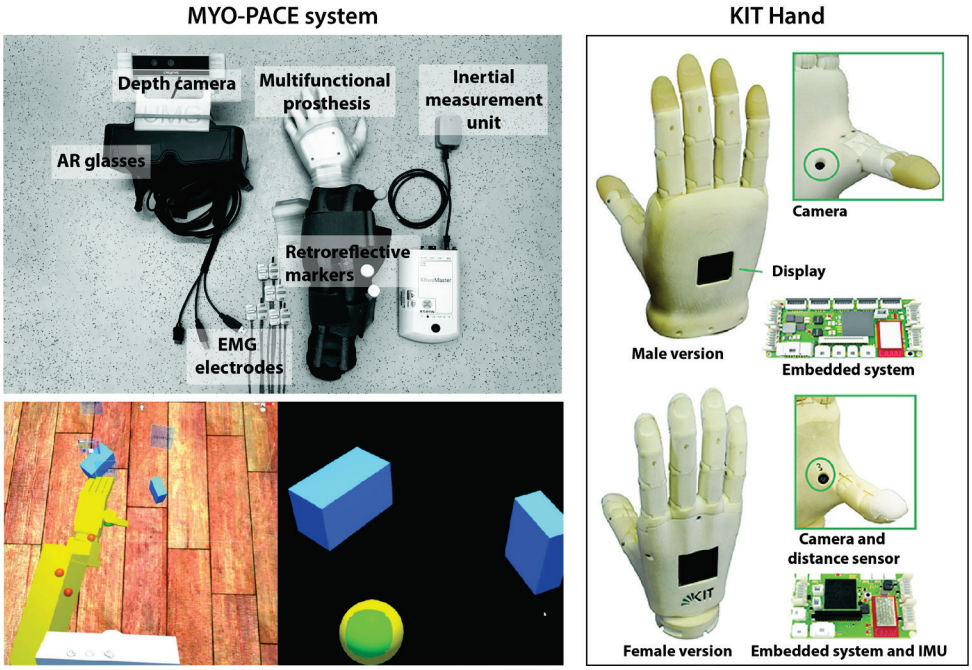

Fig. 4 Two representative prototypes of semi-autonomous limbs: the MYO-PACE system (left) [13] and the KIT hand [15] (right). *The MYO-PACE system includes a depth sensor to capture the scene, an inertial measurement unit and*

retroreflective markers to track the prosthesis position, and augmented reality glasses to provide feedback to the user. This allows for a complete reconstruction of the scene including multiple objects and the prosthesis (bottom-left). The KIT hand is a compact self-contained system with an integrated camera, a display for user feedback, and a powerful embedded system for onboard processing.

These sensors are not part of the standard design but can be easily integrated into a prosthesis. The most challenging aspect is estimating the position of the limb in space. In the MYO-PACE system [13], for instance, the position is estimated through sensor fusion (Kalman filtering) by combining the estimate obtained from inertial measurement units with visual tracking of the 3 retroreflective markers placed on the bionic limb.

Exteroception refers to perceiving the external environment by biological senses. Analogously, the role of the *artificial exteroception module* is to detect the environment in the vicinity of a prosthesis. This could be achieved by equipping the prosthesis with tactile and vision sensors. An RGB camera [16] and/or depth sensors [17] can be mounted on the prosthesis or the user. The latter configuration requires wearing an additional external component (e.g., a camera in Google Glasses) but provides a broader perspective (i.e., user versus hand point of view). The sensor data are processed using computer vision methods to estimate the structure of the scene in front of the user/prosthesis and object properties (shape and size) [13], or directly recognize the desired grasp type using machine learning [16], [18].

The *decision-making and cognitive processing module* fuses the information from artificial proprio- and exteroception modules to estimate the user's intention and the likely point of interaction between the prosthesis and the environment. For example, if the prosthesis moves toward an object and comes close to it, the MYO-PACE [13] hand assumes that this is the target that the subject intends to grasp and then activates predefined grasping strategies to automatically preshape the prosthesis.

However, bionic limbs cannot function fully autonomously. As humans are in the loop, it is necessary that they feel to be in control. Ideally, the user does not need to know when and which functions are activated automatically; instead, the impression should be that the bionic limb follows their intentions and does what they want. Therefore, the last step in the control pipeline is to fuse the decisions of the automatic controller with the volitional input by the user. Any of the approaches described in Section 2 can be used to estimate the voluntary commands, most often, from the recorded EMG. The automatic and volitional control streams can be integrated using different schemes:

$$g_1 = \begin{cases} u_v & \text{if } c(u_v, x_{full}) \\ u_a & else \end{cases} \quad (3)$$

$$g_2 = [u_v; u_a] \quad (3a)$$

$$g_3 = wu_v + (1-w)u_a \quad (3b)$$

where, g_1, g_2, and g_3 represent three distinct approaches to integrate volitional and automatic control. Control schemes under g_1 switch the control between the two streams based on criteria c that depends on u_v and x_{full}. For instance, the human could initiate the switch by producing a particular command. A brief activation of extensor muscles can indicate that the user wants to open the hand and grasp the object in front of the prosthesis. On the other hand, one could concatenate both streams (g_2), such that the human controls certain DoFs, while the remaining DoFs are controlled automatically. Lastly, the two streams could be weighted by their respective uncertainties to generate the optimal control command such as in g_3. To implement this approach, we need both volitional and automatic control methods that can estimate not only commands but also their associated uncertainty (e.g., Gaussian process regression and Bayesian deep learning [19]). Importantly, the last two integration approaches have not yet been implemented within working prototypes.

Finally, as shown in Fig. 3, to close the control loop, the system provides feedback to the user. The feedback can be delivered via tactile stimulation (haptics), but also using wearable visual displays, for instance, augmented reality glasses [20].

Fig. 4 shows the prototypes of upper limb smart prostheses, but the same controller structure (Fig. 3) is valid for the lower limb prostheses and exoskeletons [5]. In the case of microprocessor-controlled devices, the voluntary control is mechanical (user swinging the limb) while EMG can be used in the systems with powered knee and ankle. The prototypes in the literature used computer vision to assess the type of terrain and automatically configure the prosthesis parameters (e.g., increasing knee damping for walking downslope) [21], [22].

4. Connected bionic limbs

Given the advances in wireless systems in terms of rate, latency, reliability, and number of connected devices, the natural step in the evolution of smart prostheses is to enhance them with pervasive connectivity that will allow them to actively communicate with other Internet of Things (IoT) devices as well as access Cloud-based resources. This will endow the connected bionic limbs with an unprecedented level of environmental awareness and virtually unlimited computational power. Early indicators of this vision have already started to appear, and the proposed concept reflects the developmental trends in the industry, where prosthetic systems are increasingly equipped with wireless connectivity.

Indeed, equipping a prosthesis with a Bluetooth connection is becoming a standard option in commercial devices. This allows a trained prosthetist to connect to a prosthesis from their laptop and configure its parameters. Newer systems are delivered with a mobile phone app that can be employed by the user to configure some high-level functions (e.g., select a set of active grasps) or calibrate the pattern classification controller. Wu et al. [23] developed a framework for prosthesis research that allows streaming the data from the prosthesis sensors to an external processing unit (mobile phone or a laptop) as well as to a commercial Cloud server. The server provided data storage and a graphical interface so that the data could be visualized but also processed by a remote user (experimenter). During experimental validation, the participants employed a prosthesis to perform a pick-and-place task while the recorded data were transmitted to the

server and then used offline to analyze system performance and retrain the classifier, which was then downloaded back to the prosthesis. Shatilov et al. [24] presented a system in which a 3D-printed dexterous prosthesis communicated with a mobile companion that was in turn connected to a Cloud server. In this case, the connectivity was used not only to transmit the data to the Cloud for storage and offline processing (as in Wu et al. [23]), but also directly for online control using computational offloading. The prosthesis user could select between a local classifier implemented on the mobile companion and a remote classifier running on the Cloud server. The local classifier was automatically activated in case the connection to the server was lost. The study demonstrated several benefits of computational offloading, for instance, lower overall latency and decreased power consumption.

The general framework conceptualizing a connected bionic limb enhanced with Edge and Cloud computing is shown in Fig. 5. Sensor data recorded from a wearable system, user, and environment are uploaded to a Base Station (BS) with Mobile Edge Computing (MEC) capabilities. The MEC-enabled BS implements data fusion and computes control and feedback signals, which are then transmitted back to the connected limb, where the commands are executed and feedback delivered to the user. The use of network edge infrastructure provides a powerful online control loop with low latency. Meanwhile, the functions that are not time critical (e.g., data logging) and those that require more significant processing power, such as the training of large machine learning models, are performed using Cloud computing and data storage services, which are much more powerful than MEC nodes, but have a longer communication Round-Trip Time (RTT) due to their physical distance from the user. Ideally, the delay due to communication with the MEC BS will be imperceptible to the user, who will have the impression that the processing is performed locally. Overall, the proposed solution implements closed-loop semi-autonomous control, in which low-level control decisions, made by the MEC BS to enact the user's high-level commands, would be effectively transparent to the user. Finally, the connected limb also implements a local controller that runs independently of the connectivity, providing a graceful performance degradation if the connection is lost.

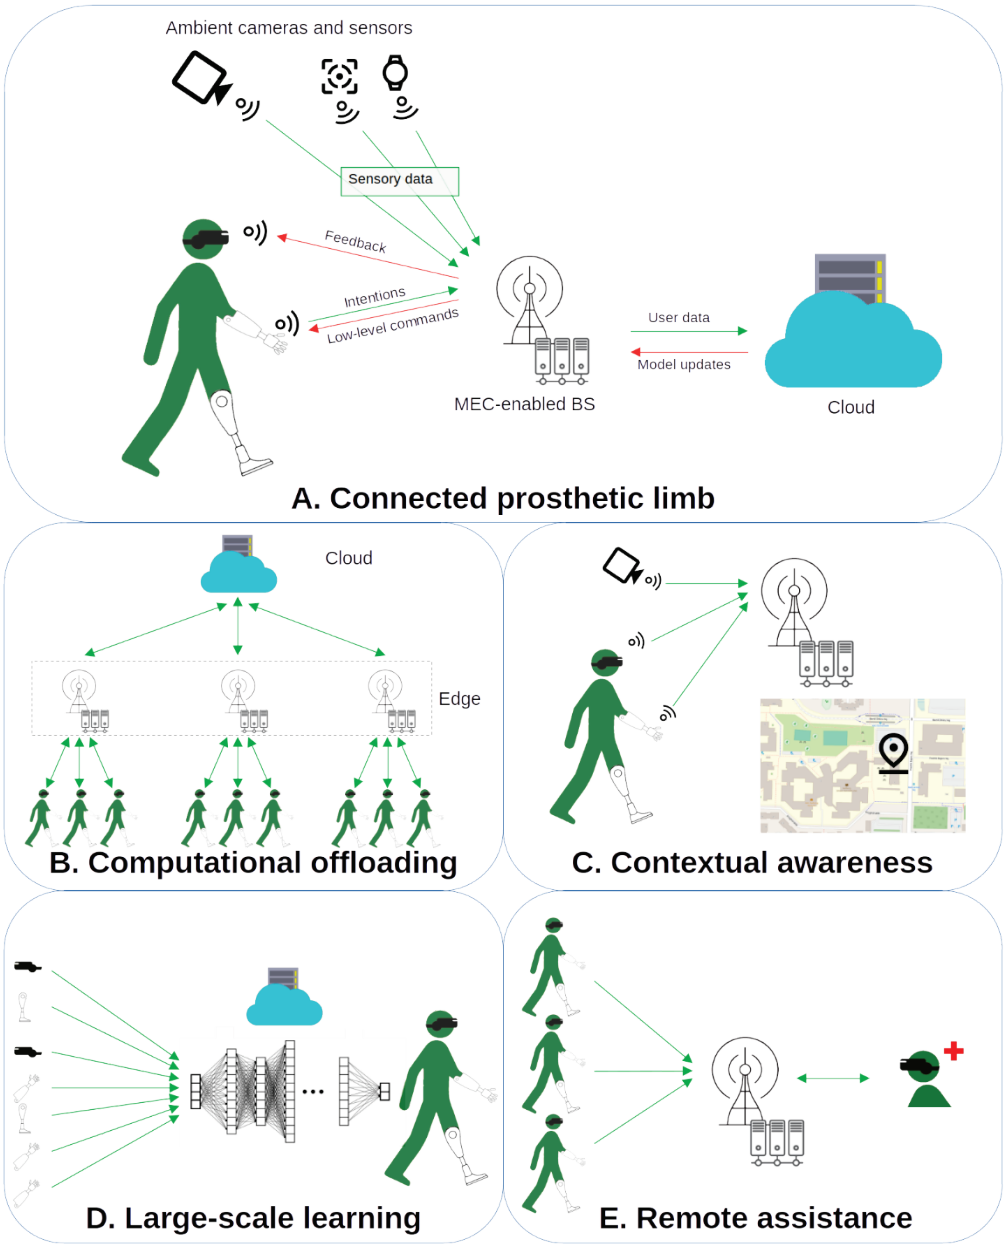

Fig. 5 (A) Schematic of a connected assistive prosthetic limb. The limb is aided by ambient sensors to expand its perception, and computation of commands is performed by the MEC-enabled BS, based on a machine learning model that uses contextual data and user inputs. The data are also uploaded to the Cloud, which can update and train the learning model to personalize it and improve the limb's functionality. (B-E) Key new features of connected assistive robots are explained further in the text.

Once the processing and storage are offloaded to the edge, assistive devices will be empowered with an abundance of data, sensor fusion, advanced signal processing, and data mining capabilities. Moreover, the control and feedback methods will not be hard-coded in the device but instead run at the network edge, where they get updated or even replaced transparently to the user and without any change to the device hardware. The bidirectional connection to the Cloud will therefore equip the bionic limbs with capabilities not available hitherto. However, there are some engineering challenges to support this vision, particularly at the network level: the computational and communication aspects need to be considered jointly for each service, and the latency and transparency requirements can be extremely stringent. In the following, we will explain the main advantages and trade-offs of the paradigm in detail.

4.1 Computational offloading

Computational offloading is the key factor in the proposed architecture: necessary steps for semi-autonomous control of bionic limbs described in Section 3 require significant processing power to be performed in real-time. Artificial exteroception requires advanced sensor fusion, often including high-dimensional data sources such as video, and the decision-making and cognitive processing module needs significant computing power to understand the proprioceptive and exteroceptive inputs, infer the user's intentions from this context, properly shape the low-level control of the prosthetic limb, and generate the appropriate user feedback. Such processing cannot be performed onboard bionic limbs, which are wearable devices with limited computational resources and battery power. Indeed, most prototypes presented so far in the literature (e.g., MYO-PACE in Fig. 4) used dedicated lab computers to implement the required processing and only the final commands (after integration) were transmitted to the prosthesis. On the other hand, commercial Edge devices [25] are significantly more powerful than bionic limb embedded controllers and do not face the same battery and space constraints. The physical distance between Cloud data centers and the end users may be an issue when offloading these operations, reducing the available computation time due to the impact of the communication RTT on the latency budget: the rise of MEC represents a

middle ground between local and Cloud processing that can provide the required performance and balance communication and computation needs in the latency budget. Compression can also help make communication more efficient, leaving more time for processing: for instance, in the case of EMG, pre-processed features at intermediate layers of the neural network, or even the outputs of different phases of the regression, classification, and control pipeline can be processed locally and then transmitted to the Edge, with different advantages and drawbacks.

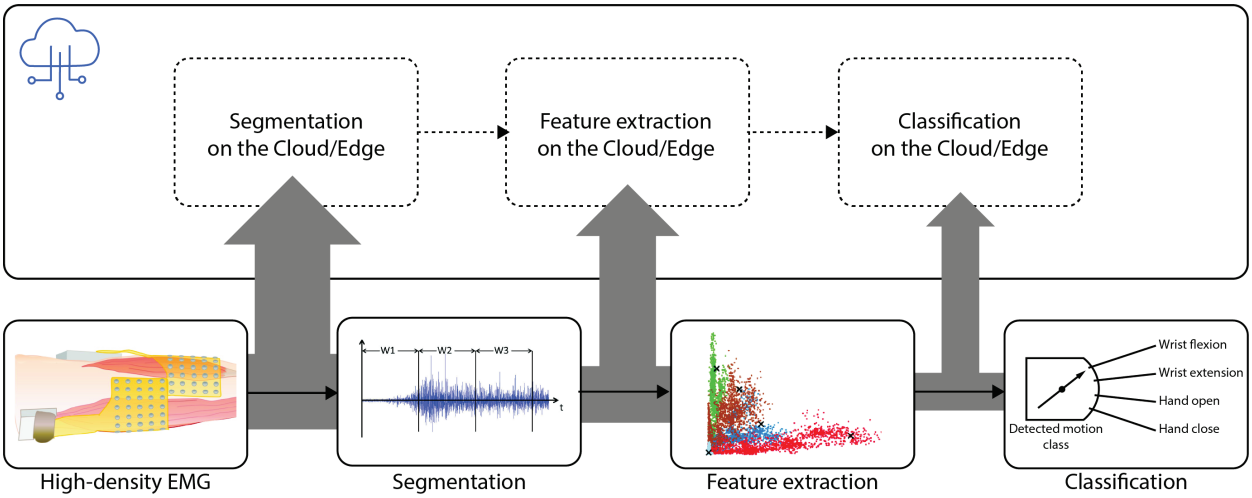

Fig. 6 The pattern classification pipeline comprises EMG recording, segmentation of raw data into (non)overlapping windows, computation of features from the data segments, and classification. Each of these steps (together with the subsequent steps) can be offloaded to the Cloud/Edge and this entails a tradeoff as discussed in the text. The thickness of the lines indicates the amount and rate of the data generation at each step of the pipeline. (the snippets in the segmentation and classification blocks are adapted from [7] *and* [6], *respectively)*

For instance, as shown in Fig. 6, the pattern classification pipeline for prosthesis control comprises data segmentation, feature extraction, and classification steps, and "inserting" a remote connection between any of the two blocks entails a trade-off. The data transmission later in the chain decreases the required communication bandwidth but also limits the flexibility of the remote Cloud controller. For instance, if the prosthesis sends raw EMG signals (individual samples), the remote controller can decide how to segment

the data, which features to compute, and which classifier to train to decode the user's intention. On the other hand, if the EMG signals are segmented and the features computed locally on the prosthesis, the amount of data that needs to be transmitted is substantially reduced but the only decision at the remote side concerns the classifier selection. If the classification is implemented using deep learning, the choice is simplified because the input in this case is the segmented signal (rather than the features). Nevertheless, the advantage of transmitting the raw data is that the remote controller can still choose the segmentation parameters (e.g., the length of the window and the amount of overlap).

Signal processing challenges

The shared control objective formulated in Section 3 then becomes more complex, as we introduce a latency component explicitly, as the previous model contains it only implicitly within the system dynamics. If computational and communication latencies are considered, decisions are made on delayed inputs and are implemented by the limb with a further delay. We can consider the overall latency T as the sum of the perception latency T_p, which represents the time between something happening in the real world and the MEC BS becoming aware of it, and the action latency T_a, which represents the time between the MEC BS receiving sensory data and the implementation of the decided action by the connected limb.

A key performance indicator is the overall latency of the control loop: there are 4 steps that need to be accomplished in a short time to maintain the illusion of instantaneous control for the user, i.e., maintain the transparency of the system:

1. The sensor signals are collected from the bionic limb and compressed for transmission;
2. The compressed data must be transmitted through the network to the MEC or Cloud server;
3. The received signal must be decoded by the server and processed to generate: (i) a command signal for the prosthetic limb, i.e., commands for all its actuators that implement the user's high-level intentions into fine-grained motion, and (ii) the feedback for the user, to be provided using methods mentioned in Section 2 and 3 (e.g., smart glasses or tactile stimulation);

4. The command and the feedback need to be transmitted back to the smart prosthesis.

Steps 1 and 2 contribute to T_p, while steps 3 and 4 to T_a.

We can then rewrite the control equations as follows:

$$\begin{aligned}
&Minimize: && J\left(x, x_{goal}\right) && (4)\\
&Subject\ to: && \dot{x}_{full} = f\left(x_{full}, u; \phi\right) && (4a)\\
& && u(t) = g\left(u_v(t - T_a), u_a(t - T_a)\right) && (4b)\\
& && n_{full} = \pi_v\left(\tilde{x}, x_{goal}\right); \tilde{x} = x + \epsilon_f && (4c)\\
& && u_v(t - T_a) = D(\hat{n}(t - T); \theta_D); \hat{n} = n + \epsilon_m && (4d)\\
& && u_a(t - T_a) = \pi_a\left(x_{full}(t - T), \hat{x}_{goal}(t); \theta_\pi\right);\ \hat{x}_{goal}(t) = E\left(s, x_{full}(t - T); \theta_E\right) && (4e)
\end{aligned}$$

The equations have similar interpretations as in (1) and (2). However, the result at the action execution time t depends on sensory information with an age T, while actions are computed by the MEC or Cloud server after a delay T_p, i.e., T_a seconds before the actual action. We can note that the automatic controller can partially compensate for the latency by considering it in its objective function and adapting its actions accordingly, but only if said latency is predictable. The latency components may even be different for u_v and u_a, as the two might be computed in different locations, e.g., volitional commands can be processed on the Edge, while automatic commands can be processed on the Cloud. We can then analyze the components of the latency, along with the inherent trade-offs and limitations.

The constraining factor in steps 2 and 4 is the channel capacity C of the bottleneck link in each direction, and the corresponding latency is directly proportional to the length L of the transmitted data block. The uplink latency T_{UL}, i.e., the communication latency in step 2, can then be formalized as:

$$T_{UL} = \frac{L}{C} + \tau,$$

where τ represents the computation and processing delay across all links in the connection. This term is independent of the packet length L, and significantly affects the decision on where to offload the computing task: while powerful Cloud servers can significantly reduce the computation time in step 3, both the number of links and the physical distance covered by them inevitably increase with respect to a MEC server co-

located with the cellular Base Station (BS), and consequently, so does the total delay τ. On the other hand, the first term of the sum remains the same, as it depends on the bottleneck capacity C: in modern networks, the access link, i.e., the one between the user and the BS, is usually the bottleneck, making this term the same in both cases. The same considerations can be made for the downlink latency T_{DL} in step 4.

On the other hand, the T_{UL} term is affected by the compression in step 2: the raw EMG signal, for instance, might have a significant size if recorded using a high-density electrode, slowing down transmission if it is transmitted without any compression. As we discussed above, the first drawback of compression is in the accuracy of the final control, as the control policy depends on the noisy signals $\hat{n}$, and increasing compression also makes the noise terms more relevant. Minimizing T_p by compressing the signal can then have negative effects on control accuracy, and the trade-off must be carefully considered, as even advanced machine learning algorithms might not be able to compensate for lower-quality inputs. Finally, there is another issue: compression is a computationally intensive operation, and, as discussed before, the whole reason for offloading computation tasks is the limited power of the prosthetic limb's onboard embedded processor.

The design of the compression, feature extraction, and learning algorithms must be performed jointly to properly control these trade-offs: a more complex neural network will probably require adaptations, such as network compression and pruning, to run on the Edge in real-time, and will tilt the latency budget towards computation, requiring more efficient compression as well. A further step would be a joint optimization of application and network parameters, which may become necessary to support higher demands in terms of transmitted data within the stricter latency bounds imposed by the processing steps: in this case, the network needs to explicitly support the transmission of the data, following the Tactile Internet approach. The analysis of these trade-offs and validation of the system parameters will be performed in Section 5.

4.2 Contextual awareness

The vision for a semi-autonomous bionic limb is intimately tied to connectivity: while proprioception can easily be accomplished with the limb's onboard sensors, exteroception is limited by the physical properties of the prosthetic limb. Firstly, visual sensors mounted directly on a prosthetic limb have a limited perspective, and may be blocked by the movements of the limb itself; having a wider perspective, such as the one that can be provided by head-mounted cameras, e.g., in smart glasses, or even a stereoscopic vision of the environment from multiple points of view, which can be provided by appropriately placed ambient sensors and cameras, would provide a wealth of context information that can improve performance [23]. Other types of ambient sensors would also help the limb's exteroception; e.g., radar sensing exploiting communication infrastructure [26] can provide additional information on the location and movement of people and objects, while other wearable sensors [27] can enrich the context on the user's state and intentions. Several of the artificial exteroception module's functions can benefit from this type of additional information, and the continuous inflow of sensor data can be fused and integrated on the Cloud to build permanent models of the environment that can gradually grow and adapt as the user navigates their surroundings. The gradual creation of a comprehensive model of the user's living space and other commonly visited environments will increase the intelligence of the system over time, improving its performance in the most common tasks. For instance, an exoskeleton or lower limb prosthesis could build a precise map of the terrain commonly visited by the user (e.g., location and size of the stairs, obstacles to avoid, etc.), or a hand prosthesis can maintain a database of household objects with their geometrical models and associated grasping strategies. Likewise, the prosthetic limb could be configured to provide information to and interact with devices in the environment. For example, a home appliance can identify itself to a hand prosthesis, which can configure a set of grasps required to interact with the device, or automatic doors and traffic lights can detect slow-moving exoskeleton users and adapt to let them through safely.

Signal processing challenges

Naturally, the explosion of data sources and rich contextual channels, such as video or depth maps, needs to be processed in real-time to maintain transparency: this is a challenge that was already discussed in Section 4.1. The fusion of sensor data from different sources and on different timescales for the reconstruction of a 3D environment is an active research topic [28], and even self-localization may be a complex problem if different sensors are dynamically included in the fusion process [29]. This challenge requires both Bayesian reasoning and the application of modern machine learning tools, as the three-dimensional map needs to include information from past interactions as well as less frequent data from battery-constrained sensors [30].

The transmission of the high-dimensional data from ambient sensors and cameras also increases the complexity of the joint communication and computation optimization described in Section 4.1: the exteroceptive inputs need to be transmitted to the MEC-enabled BS by different sources, which may use different types of connections, e.g., proprietary IoT technologies and standards. Furthermore, transmissions of the data from ambient sensors may take away resources for the transmission of data from the semi-autonomous limb, reducing the capacity of this connection and tightening the latency budget. The use of approaches that consider the Value of Information (VoI) of different data sources, which is itself dependent on the context and on spatial and temporal correlations between measurements [31], can significantly improve performance but requires an intelligent orchestration of data acquisition from the artificial exteroception module. Beyond simply integrating data, the module will be required to actively decide which data to consider and how obtaining it from wirelessly connected ambient sensors can benefit the quality of the user's control performance and experience using the prosthetic limb.

4.3 Large-scale machine learning

Reliable and trustworthy data on the usage of assistive systems in daily life are currently missing, but it could play a critical role in driving the development of clinically useful assistive technology. Such data can be easily collected if the assistive systems are fully connected, as proposed in our framework. While privacy

considerations must always be kept in mind, this represents a possible solution to empower large-scale learning: the processing power of the Cloud [32] can be exploited to perform non-time-sensitive tasks such as the training of new and updated models, which can be kept and shared among users. These generalized models could provide, for example, the most common grip types, wrist orientations, and grasping forces that amputees employ when performing activities of daily living. If the information from onboard sensors is combined with data from the environment, providing localization and interaction, this could also allow categorizing the use cases according to the context. The creation of personalized models that transfer knowledge from more general ones and adapt to individual users' preferences and needs is a requirement [33] due to the different usage patterns, activities, and preferences of bionic limb users. This can be achieved with Federated Learning (FL) [34], where the local models would exploit longitudinal data to adapt to specific usage patterns and users' habits. If there is a particular activity that the user often repeats, this pattern could be detected so that the system proactively supports relevant function sequences. This way, machine learning-based control can benefit not just from the user's experience, but from the aggregated experience of the whole user population. The local models can then be trained on the limb itself or on the Edge, preserving users' privacy by not transmitting any personal data to the limb manufacturers, while the models themselves can be integrated into a larger general model. The optimization of this process is a well-studied problem in FL [35], which can be solved using traditional optimization.

Signal processing challenges

The efficiency considerations for this aspect of the framework are less demanding, as it does not have to deal with strict transparency and control performance guarantees. On the other hand, large-scale machine learning on health data exposes users to security and privacy risks. The MEC architecture can be vulnerable to Distributed Denial of Service (DDoS) attacks from malicious devices [36], or even to man-in-the-middle attacks if parts of the network infrastructure are compromised. Furthermore, since attacks have effects in

the real world, the standard security paradigms are not sufficient, and we need to consider stronger defense and mitigation techniques: assistive systems should not only be protected from attacks but there should be mechanisms to detect when an attack is happening and prevent any harm from coming to the user. On-board capability becomes a fundamental requirement in this case, as it can mitigate the impact of cyberattacks by simply switching to local control, reducing the semi-autonomous functionality, and temporarily causing the user some discomfort to prevent attackers from causing more significant damage or injury. Privacy is another issue: the risks of attack are also privacy risks, as malicious actors in the network can get access to extremely sensitive medical data [37]. Furthermore, the collection of data and training of machine learning models can be a significant privacy risk: the gathered data are extremely sensitive, and both classical and innovative security techniques need to be deployed. FL might help mitigate this risk [38], since in FL, each user equipment (either the robotic limb itself or a node at the edge of the network) computes its own updates to the learning model, which are then propagated in a compressed form to the central server that performs the integration and creates the general model. This technique avoids direct transmission of input data, mitigating the impact of leaks, and it has already been proposed as a privacy-conscious way to perform data analysis on electronic health records [39]. Although even this kind of distributed model is not completely immune to privacy leakage [40], there are techniques to mitigate the problem, restoring full privacy at a small performance cost.

4.4 Data logging and remote intervention

Continuous data logging on the Cloud/Edge will allow online monitoring of the status of bionic limbs while they are used in everyday environments (e.g., at home). Logging and access to such data can provide new insights about users' behavior and preferences, critical for the further development of the technology that can address real user needs.

The data collected from the connected assistive systems can be used to assess the state of the device, as well as the level of user activity. Such information would enable early diagnostics and contact with the user when required [41]. The online streaming could be used to provide remotely guided training in the relevant environment (e.g., the user's home and/or work). This approach would be particularly relevant for the early adoption of assistive technology, with the potential to decrease the current high rejection rates [42]. The streamed data can help ensure the safety of older or disabled people, either in their own home or in assisted living facilities, by increasing the accuracy of remote alert systems that can make medical personnel aware of potential health emergencies [43]. More importantly, the automatic detection of dangerous situations can alert medical personnel if the user needs assistance while using the device. For example, if the user experiences difficulties using the assistive system, they can connect to the online service center, where an expert is available to provide help and, if necessary, take over control. For example, if a user of an exoskeleton falls, it could be difficult or even impossible for them to recover the system and stand up. However, a remote assistant that receives the data from the device has a complete insight into the exoskeleton state and the surroundings, thus being likely able to maneuver the system upright, or at least into a safer configuration before the arrival of physical assistance.

In this case, the real-time requirements are much less stringent, as the transparency requirement is dropped: the human response is measured in a scale of seconds to minutes, and movements in the guided stage need to be slowed down to avoid potential harm to the user. In these cases, offloading to the Cloud does not represent a significant challenge, as the RTTs between the user and the Cloud and between the Cloud and the hospital or emergency center are well within the requirements.

5. Framework validation

In this section, we introduce an illustrative example of a connected bionic limb that integrates state-of-the-art control, sensing, and haptic feedback technologies. This system is then used to estimate the amount of

data that needs to be transmitted through the communication interface between the limb and the Cloud/Edge (up and downlink), conducting a feasibility analysis for the latency requirements under different system settings.

5.1 Communication and computing system requirements

As described in Section 4.1, the main communication challenge is in the optimization of the latency, which needs to be considered over the closed loop: the time between a user expressing an intention and the associated limb response and tactile and visual feedback needs to maintain the illusion of real-time. In a typical Edge scenario, the propagation and network processing times are usually short and do not contribute significantly to the latency budget. Firstly, we present a possible estimate of the data requirements for a bionic limb using visual and tactile feedback, along with mechanical, EMG, and depth map-type sensing. The estimated requirements for the uplink (UL) from the limb to the BS and the downlink (DL) from the BS to the user, used for sensory data and feedback respectively, are listed in Table 1.

Table 1: Data generated by a realistic artificial limb.

Interface	System	Channels	Data	Bitrate	Position
Volitional control	Upper/lower limb prosthesis	2-3000	EMG, 2 kHz	192 Mb/s (UL)	Covering the whole limb with conductive pads every 2-3 mm
Electro-tactile feedback	Upper/lower limb prosthesis	2-3000	Tactile, 2 kHz	192 Mb/s (DL)	Same interface used for recording the EMG
AR feedback	Epson Moverio BT-300	720p	Video, 30 fps	87.2 Mb/s (DL)	Smart glasses
Mechanical sensing	Artificial skin	1-2000	Tactile, 2 kHz	128 Mb/s (UL)	Covering the prosthesis
Context sensing	Intel Real Sense Depth Camera D435	1080p	Video, 30 fps	196 Mb/s (UL)	Anywhere

The total required bitrate to transmit uncompressed data is approximately $R = 600$ Mb/s for the uplink and $R=300$ Mb/s for the downlink. While modern wireless networks can achieve these bitrates, this often

requires dedicated resources and favorable channel conditions, with a high risk of generating significant queuing delays that can impair performance and break the system's transparency requirement.

The computational constraints are also significant: if we consider an ideal scenario with a mmWave link with a capacity $C = 2$ Gb/s in either direction, the computational power ψ_{BS} of the Edge device [44] still represents a bottleneck. If we denote the computational complexity of the task as ω, we then have the following expression for the overall latency T:

$$T = \frac{R}{C} + \frac{\omega}{\psi_{BS}}.$$

In Fig. 7, we analyze the real-time performance of the state-of-the-art NASNet [45] Convolutional Neural Network (CNN) architecture, which requires approximately 24 billion operations, and the more compact Xception network [46], which has a lower accuracy but requires 8 billion operations.

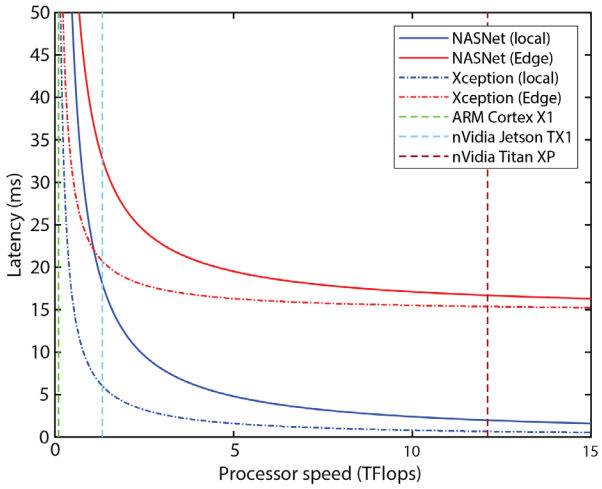

Fig. 7 Overall latency as a function of computing power in the considered example. Edge processing requires an additional delay to transmit the computation inputs and results, but local processing can only have a limited processor speed due to space and energy constraints on the assistive device.

If we consider an embedded ARM Cortex X1 chip, an nVidia Jetson TX2 embedded computer, commonly used in state-of-the-art MEC setups, and the nVidia Titan XP GPU with an Intel i7 processor, only the latter can run the Xception CNN in less than 10 ms. A realistic connected assistive system would then have a latency of about 20 ms for Xception and 30 ms for NASNet, which is narrowly within the acceptable range for a seamless user experience. Communication represents a significant part of the latency budget, requiring about 15 ms even under ideal conditions.

We also mention the power requirements for the computations: the power consumption of the nVidia Jetson to run CNN models is between 5 and 8 W [44], within the manageable range for a MEC-enabled BS. On the other hand, more powerful server GPUs, such as the nVidia Titan require about 20 times more power, as the total consumption as measured in [47] is between 80 and 180 W depending on the GPU configuration. However, even Edge devices would put a significant strain on prosthetic device batteries: a constant 5 W draw to run inference and context adaptation on the bionic limb itself would completely deplete a standard 5000 mAh cell phone battery, working at the standard 5 V, in about 5 hours. Less powerful embedded chips, such as the ARM Cortex X1, would increase the battery lifetime to acceptable levels but do not provide acceptable real-time performance, as Fig. 7 clearly shows. In real use cases, wireless channel conditions are often far from ideal, and the slow adoption of mmWave technology, which has still not reached commercial viability, constrains the maximum data rate. Performing local compression of sensed information before transmitting it can significantly reduce the communication cost, allowing operation over more limited wireless links without significant delay increases. However, compression is itself an operation that requires some time: the bionic limb, whose computational capabilities are limited, takes some time to compress sensory information, as well as to decompress feedback if it is also transmitted in a compressed form, as would be necessary for video feedback, following the procedure described in Section 4.1.

Denote the compression and decompression delay as τ_c and the compression rate, which represents the ratio between the raw data size and the compressed size, as η. The overall latency T in a system with local compression is then given by:

$$T = \frac{R}{\eta C} + \tau_c + \frac{\omega}{\psi_{BS}}.$$

The compression delay is also often directly proportional to η, so we can consider the following optimization problem for a given channel capacity and limb computational power ψ_L:

$$\eta^* = argmin_\eta \frac{R}{\eta C} + \frac{\eta}{\psi_L}.$$

This minimization problem highlights an important trade-off between the time required to communicate the data, which is reduced by a more efficient compression, and the computation required by the complexity itself.

An additional problem is represented by the possible presence of multiple users accessing the network: these might be multiple bionic limbs or normal network traffic. Naturally, if the capacity of a network is shared among several users, the resources available to each one are reduced proportionally. Another significant issue is caused by the effect of the uncertainty on the wireless channel conditions, which must be considered: the channel capacity might vary stochastically over time, leading to a wide difference between average and worst-case performance. We then need to select an acceptable reliability $\rho \in [0,1)$ to perform the optimization: for example, if we select $\rho = 0.95$, we can tolerate an excessive latency on 5% of frames. Reliability is a system design parameter, which depends on the Quality of Experience (QoE) targets and the intelligence of the processing algorithms.

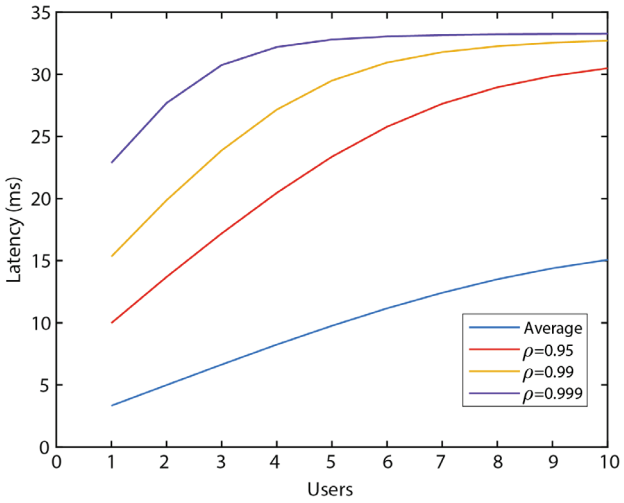

Fig. 8 Cumulative Distribution Function (CDF) of the UL latency with a compression rate of 10 (leading to a generated rate of 60 Mb/s) and an average channel capacity of 300 Mb/s.

The analysis must then consider the full Cumulative Distribution Function (CDF) of the communication latency, shown in Fig. 8 for different numbers of active users under realistic conditions. By considering a simple system in which users get an equal share of the channel capacity, we can analyze the average and worst-case latency. Even in this simple case, the need for proper dimensioning of the system is clear, and the worst-case performance is significantly different from the average: as the number of users in the system increases, the average latency degrades gracefully, while the worst-case latency grows significantly faster. It is important to remember that any increase in the communication latency shrinks the computation latency budget by an equivalent amount, running into the hardware limitations of the MEC processors. Consider Fig. 9, which shows the results with compression: assuming a computationally efficient compression and decompression process, only requiring 5 ms on both sides, average performance only degrades slightly: the average latency for a BS equipped with the nVidia Jetson TX2 processor is about 20 ms for Xception and 40 ms for NASNet. On the other hand, worst-case performance, computed using the assumptions above, is significantly worse: while the 35 ms necessary for Xception are still within the acceptable range, NASNet requires more than 50 ms, violating the transparency requirement and breaking the instantaneity illusion.

We explored the trade-off between average-sense and worst-case optimization, jointly considering compression and communication settings, in our recent work [48], which provides a deeper mathematical characterization.

There is also another issue: while we can optimize latency considering a worst-case scenario, the controller needs to become even more complex to compute $u_a(t)$ with a stochastic latency, as it needs to consider any possible delay between observation and action and its effect on control performance. This carries important implications from the control point of view and can open the way for significant new research in this direction.

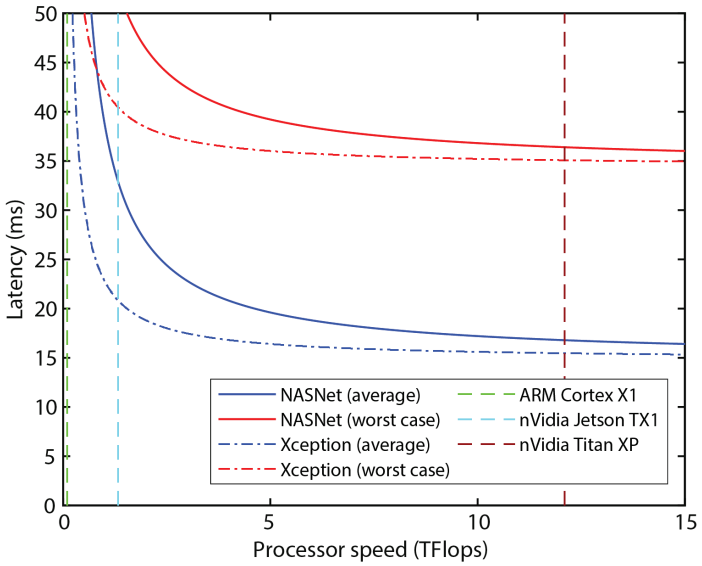

Fig. 9 Overall latency for Edge processing as a function of computing power under uncertainty. The worst-case scenario considers $\rho = 0.95$. *The average channel capacity is 300 Mb/s, with a single user, and the compression efficiency is 90%, with an additional compression and decompression delay of 5 ms.*

5.2 Quality of Experience analysis

Cloud-connected bionic limbs can only become a reality if they satisfy stringent real-time latency requirements that do not interfere with users' quality of experience. Therefore, it is crucial to consider the tradeoff between the quality of user experience and (closed loop) latency when designing the

communication system requirements. However, not much is known about how users react to the impairments in the wireless link. Nevertheless, this point can be addressed using the theory and methods of manual control cybernetics (MCC) [49]. MCC aims to understand and model human behavior when controlling dynamic systems using different command and feedback interfaces, and a common experimental paradigm is to ask the subjects to control a system so that its output tracks the desired reference trajectory. By applying system identification methods to the experimentally collected data, a model of the human controller (e.g., transfer function) can be estimated.

In recent experiments [50], we applied this approach by extending the typical human-in-the-loop setup with the model of a wireless communication channel (Gilbert-Elliott model combined with experimentally measured latency profiles), as shown in Fig. 10. We then used this framework to analyze the difference in quantitative performance as users experienced scenarios where the control and feedback commands underwent simulated wireless impairments. Participants were instructed to perform a compensatory tracking task that involved manipulating a cursor on the screen using a joystick. To achieve this task, they received feedback on the tracking error (reference minus generated trajectory) either visually or through the vibrotactile medium. Both the outgoing joystick commands and incoming feedback commands were subject to simulated wireless impairments based on real-life measurements of round-trip data transmission over an LTE-BS. Namely, the commands were subject to latency, with random fluctuations around the latency, and packet loss according to the measured latency profiles. We observed that performance degraded significantly for a low-load scenario corresponding to a round-trip latency of 178 ms, and tactile feedback led to further worsened performance. Therefore, the tradeoffs involved in maintaining user experience are also affected by the type of feedback involved in the teleoperation system, or the bionic limb, apart from the amount of data. However, fluctuations around the mean latency and packet loss did not affect the performance.

While we investigated quantitative performance degradation, issues related to the equally important subjective experience of control were not considered. These could be measured using questionnaires or tasks that specifically investigate users' psychophysical aspects such as the sense of agency. While psychometric measures relating to this issue are still under investigation, such measures also play a key role in identifying where to allocate communication resources.

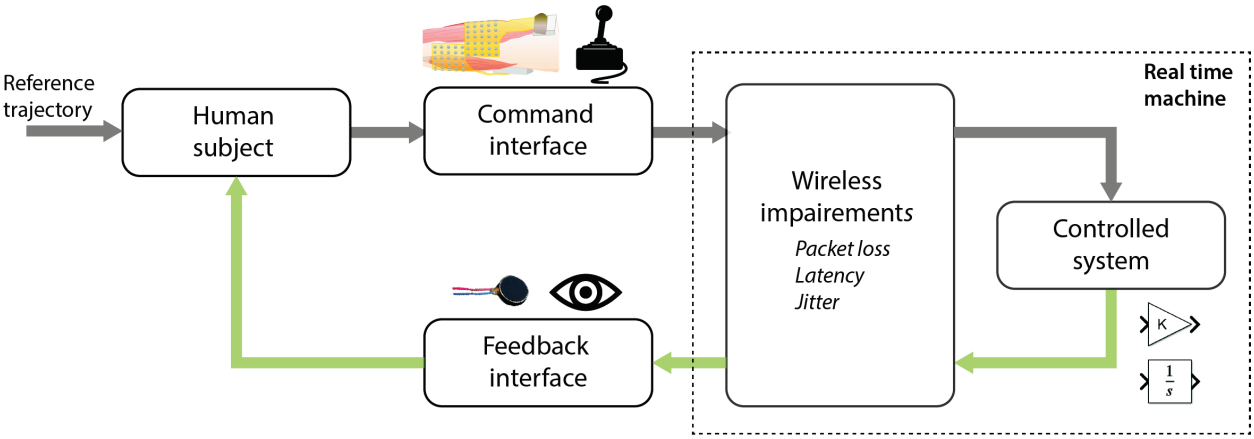

Fig. 10 Closed loop control test bench to experimentally assess the impact of wireless impairments on the user experience and performance. The blocks inside the dashed box are simulated on a real-time machine while the other blocks are real systems. The simulation allows testing a range of parameters that characterize the connection between the limb and the user through the Cloud/Edge in combination with different models of bionic limb response (position or velocity control), command interfaces (EMG or joystick), and feedback types (visual or tactile).

6. Conclusions and future work

This is a very exciting period in the development of bionic limbs. New signal processing approaches and revolutionary surgical techniques are being proposed and tested, leading to substantial improvements in human-machine interfacing. However, although these are indeed important steps in establishing a natural connection between the users and their advanced bionic limbs, the human-machine interface is likely to remain a bottleneck in the many years to come. This is due to the sheer complexity and scale of the human

nervous system, which pose major challenges when trying to understand and decode the neural code of movement (commands) and sensation (feedback). The present manuscript starts by acknowledging these points but then also proposes how to effectively bypass these challenges by enhancing bionic limbs with intelligence and automatic functioning, boosted by the vast computational resources of the Cloud and MEC. Pervasive connectivity, especially the emerging 5G systems, can be used to augment the sensing and processing capabilities of bionic systems, paving the way for radically novel features in the area of assistive technologies: comprehensive contextual awareness, data logging, large-scale learning, and adaptation, as well as remote monitoring and intervention. The proposed approach has the potential to fundamentally change the applicability, acceptance, and perception of assistive robotic systems within society. Nevertheless, connected bionic limbs are in the early stage of development and many open questions remain to be addressed, including but not limited to 1) understanding of the user experience and tolerance with respect to wireless impairments, 2) allocation of functions between local and Cloud/MEC processing to ensure robust operation, 3) identification of the most useful functions to support by the Cloud/MEC, 4) design of the smart limbs with integrated sensing and connectivity interfaces, 5) clinical testing of the connected limbs out of the lab to assess utility and acceptance, especially related to the potential use of external components (e.g., smart glasses) and the impact of automatic control on the embodiment, and 6) data protection and privacy. Despite the potential of cloud-connected bionic limbs, reproducing the dexterity and flexibility of human motion is likely to remain a major challenge (e.g., humans might grasp similar objects in different ways depending on the context). Nevertheless, access to the vast computational and storage resources of the Cloud/Edge can enable a continuously learning prosthesis that can gradually capture the nuances in the user's intentions. Importantly, although the present manuscript focused on upper and lower-limb robotic prosthetics, most of the concepts can be equally well applied to arm and leg exoskeletons.

7. Acknowledgement

This work was supported by the Villum Investigator Grant "WATER" from the Velux Foundation, Denmark (Petar Popovski); by the European Research Council (ERC) under the Synergy Grant Natural BionicS (810346) (Dario Farina); and by the Independent Research Fund Denmark through project CLIMB (Čedomir Stefanović and Strahinja Došen).

The outline of the sketch in Figure 1 was derived from an AI-generated image created using Microsoft Image Creator, built on DALL-E 3. The resulting image was then modified in Inkscape, where it was vectorized and additional changes were made to the scene. Some elements of Fig. 5 were freely used and adapted following the appropriate license terms. The Red Cross logo and neural network example were obtained from Wikimedia Commons (https://commons.wikimedia.org/wiki/File:Redcrosslogo.png, https://commons.wikimedia.org/wiki/File:Example_of_a_deep_neural_network.png) under a CC-BY-SA-4.0 license. The map was obtained from the OpenStreetMap Foundation under the Open Data Commons Open Database License (ODbL). Some icons for sensors and nodes came from the Google Material Symbols library (https://fonts.google.com/icons) under an Apache 2.0 license. The Cloud and server icons were obtained from VectorPortal (https://vectorportal.com/vector/file-transfer-via-cloud/29094, https://vectorportal.com/vector/server-icons/31669) under a CC-BY-4.0 license. Some elements in Figs. 3 and 10 were purchased from https://www.colourbox.com/ and some elements in Fig. 5 were purchased from https://www.istockphoto.com/ and https://www.shutterstock.com/.

References

[1] B. Gesslbauer, L. A. Hruby, A. D. Roche, D. Farina, R. Blumer, and O. C. Aszmann, “Axonal components of nerves innervating the human arm,” *Ann Neurol*, vol. 82, no. 3, pp. 396–408, Oct. 2017, doi: 10.1002/ana.25018.

[2] A. Chortos, J. Liu, Z. Bao, Z. B.-N. Materials, and U. 2016, “Pursuing prosthetic electronic skin,” *Nat Mater*, vol. 15, no. 9, p. 937—950, Oct. 2016, doi: 10.1038/nmat4671.

[3] D. Farina *et al.*, “Toward higher-performance bionic limbs for wider clinical use,” *Nat Biomed Eng*, Oct. 2021, doi: 10.1038/s41551-021-00732-x.

[4] W. Guo, W. Xu, Y. Zhao, X. Shi, X. Sheng, and X. Zhu, “Toward Human-in-the-Loop Shared Control for Upper-Limb Prostheses: A Systematic Analysis of State-of-the-Art Technologies,” *IEEE Trans Med Robot Bionics*, vol. 5, no. 3, pp. 563–579, Aug. 2023, doi: 10.1109/TMRB.2023.3292419.

[5] M. Tschiedel, M. F. Russold, and E. Kaniusas, “Relying on more sense for enhancing lower limb prostheses control: a review,” *J Neuroeng Rehabil*, vol. 17, no. 1, p. 99, Dec. 2020, doi: 10.1186/s12984-020-00726-x.

[6] A. D. Roche, H. Rehbaum, D. Farina, and O. C. Aszmann, “Prosthetic Myoelectric Control Strategies: A Clinical Perspective,” *Curr Surg Rep*, vol. 2, no. 3, p. 44, Oct. 2014, doi: 10.1007/s40137-013-0044-8.

[7] M. Ortiz-Catalan, R. Brånemark, and B. Håkansson, "BioPatRec: A modular research platform for the control of artificial limbs based on pattern recognition algorithms," *Source Code Biol Med*, vol. 8, no. 1, p. 11, Oct. 2013, doi: 10.1186/1751-0473-8-11.

[8] A. Marinelli *et al.*, "Active upper limb prostheses: a review on current state and upcoming breakthroughs," *Progress in Biomedical Engineering*, vol. 5, no. 1, p. 012001, Jan. 2023, doi: 10.1088/2516-1091/acac57.

[9] S. Dosen, "Toward self-contained bidirectional bionic limbs with high information throughput," *Sci Robot*, vol. 8, no. 83, 2023, doi: 10.1126/scirobotics.adk6086.

[10] S. J. Bensmaia, D. J. Tyler, and S. Micera, "Restoration of sensory information via bionic hands," *Nat Biomed Eng*, Oct. 2020, doi: 10.1038/s41551-020-00630-8.

[11] Y. Abbass, M. Saleh, S. Dosen, and M. Valle, "Embedded Electrotactile Feedback System for Hand Prostheses Using Matrix Electrode and Electronic Skin," *IEEE Trans Biomed Circuits Syst*, vol. 15, no. 5, 2021, doi: 10.1109/TBCAS.2021.3107723.

[12] Y. Abbass, S. Dosen, L. Seminara, and M. Valle, "Full-hand electrotactile feedback using electronic skin and matrix electrodes for high-bandwidth human-machine interfacing," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 380, no. 2228, 2022, doi: 10.1098/rsta.2021.0017.

[13] J. Mouchoux, S. Carisi, S. Dosen, D. Farina, A. F. Schilling, and M. Markovic, "Artificial Perception and Semiautonomous Control in Myoelectric Hand Prostheses Increases Performance and Decreases Effort," *IEEE Transactions on Robotics*, vol. 37, no. 4, 2021, doi: 10.1109/TRO.2020.3047013.

[14] L. Seminara *et al.*, "A hierarchical sensorimotor control framework for human-in-the-loop robotic hands," *Sci Robot*, vol. 8, no. 78, Oct. 2023, doi: 10.1126/scirobotics.add5434.

[15] P. Weiner, J. Starke, S. Rader, F. Hundhausen, and T. Asfour, "Designing Prosthetic Hands With Embodied Intelligence: The KIT Prosthetic Hands," *Front Neurorobot*, vol. 16, Mar. 2022, doi: 10.3389/fnbot.2022.815716.

[16] G. Ghazaei, A. Alameer, P. Degenaar, G. Morgan, and K. Nazarpour, "Deep learning-based artificial vision for grasp classification in myoelectric hands," *J Neural Eng*, vol. 14, no. 3, p. 36025, Oct. 2017, doi: 10.1088/1741-2552/aa6802.

[17] M. N. Castro and S. Dosen, "Continuous Semi-autonomous Prosthesis Control Using a Depth Sensor on the Hand," *Front Neurorobot*, vol. 16, 2022, doi: 10.3389/fnbot.2022.814973.

[18] C. Shi, D. Yang, J. Zhao, and H. Liu, "Computer Vision-Based Grasp Pattern Recognition With Application to Myoelectric Control of Dexterous Hand Prosthesis," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 28, no. 9, pp. 2090–2099, Sep. 2020, doi: 10.1109/TNSRE.2020.3007625.

[19] B. Zhong, H. Huang, and E. Lobaton, "Reliable Vision-Based Grasping Target Recognition for Upper Limb Prostheses," *IEEE Trans Cybern*, pp. 1–13, 2020, doi: 10.1109/TCYB.2020.2996960.

[20] M. Markovic, H. Karnal, B. Graimann, D. Farina, and S. Dosen, "GLIMPSE: Google Glass interface for sensory feedback in myoelectric hand prostheses," *J Neural Eng*, vol. 14, no. 3, 2017, doi: 10.1088/1741-2552/aa620a.

[21] B. Zhong, R. L. da Silva, M. Li, H. Huang, and E. Lobaton, "Environmental Context Prediction for Lower Limb Prostheses With Uncertainty Quantification," *IEEE Transactions on Automation Science and Engineering*, vol. 18, no. 2, pp. 458–470, Oct. 2021, doi: 10.1109/TASE.2020.2993399.

[22] K. Zhang *et al.*, "Environmental Features Recognition for Lower Limb Prostheses Toward Predictive Walking," *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 27, no. 3, pp. 465–476, Oct. 2019, doi: 10.1109/TNSRE.2019.2895221.

[23] H. Wu, M. Dyson, and K. Nazarpour, "Internet of Things for beyond-the-laboratory prosthetics research," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 380, no. 2228, Jul. 2022, doi: 10.1098/rsta.2021.0005.

[24] K. A. Shatilov, D. Chatzopoulos, A. W. T. Hang, and P. Hui, "Using Deep Learning and Mobile Offloading to Control a 3D-printed Prosthetic Hand," *Proc ACM Interact Mob Wearable Ubiquitous Technol*, vol. 3, no. 3, pp. 1–19, Oct. 2019, doi: 10.1145/3351260.

[25] J. Shao and J. Zhang, "Communication-Computation Trade-off in Resource-Constrained Edge Inference," *IEEE Communications Magazine*, vol. 58, no. 12, pp. 20–26, Dec. 2020, doi: 10.1109/MCOM.001.2000373.

[26] J. Pegoraro, F. Meneghello, and M. Rossi, "Multiperson Continuous Tracking and Identification From mm-Wave Micro-Doppler Signatures," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 59, no. 4, pp. 2994–3009, Apr. 2021, doi: 10.1109/TGRS.2020.3019915.

[27] F. Khoshmanesh, P. Thurgood, E. Pirogova, S. Nahavandi, and S. Baratchi, “Wearable sensors: At the frontier of personalised health monitoring, smart prosthetics and assistive technologies,” *Biosens Bioelectron*, vol. 176, p. 112946, Mar. 2021, doi: 10.1016/j.bios.2020.112946.

[28] M. P. Muresan, I. Giosan, and S. Nedevschi, “Stabilization and Validation of 3D Object Position Using Multimodal Sensor Fusion and Semantic Segmentation,” *Sensors*, vol. 20, no. 4, p. 1110, Feb. 2020, doi: 10.3390/s20041110.

[29] T.-M. Nguyen, M. Cao, S. Yuan, Y. Lyu, T. H. Nguyen, and L. Xie, “VIRAL-Fusion: A Visual-Inertial-Ranging-Lidar Sensor Fusion Approach,” *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 958–977, Apr. 2022, doi: 10.1109/TRO.2021.3094157.

[30] J. Li *et al.*, “Deep learning in multimodal remote sensing data fusion: A comprehensive review,” *International Journal of Applied Earth Observation and Geoinformation*, vol. 112, p. 102926, Aug. 2022, doi: 10.1016/j.jag.2022.102926.

[31] F. Chiariotti, A. A. Deshpande, M. Giordani, K. Antonakoglou, T. Mahmoodi, and A. Zanella, “QUIC-EST: A QUIC-Enabled Scheduling and Transmission Scheme to Maximize VoI with Correlated Data Flows,” *IEEE Communications Magazine*, vol. 59, no. 4, pp. 30–36, Apr. 2021, doi: 10.1109/MCOM.001.2000876.

[32] T. K. Rodrigues, K. Suto, H. Nishiyama, J. Liu, and N. Kato, “Machine Learning Meets Computation and Communication Control in Evolving Edge and Cloud: Challenges and Future Perspective,” *IEEE Communications Surveys and Tutorials*, vol. 22, no. 1, pp. 38–67, Jan. 2020, doi: 10.1109/COMST.2019.2943405.

[33] B. Paaßen, A. Schulz, J. Hahne, and B. Hammer, “Expectation maximization transfer learning and its application for bionic hand prostheses,” *Neurocomputing*, vol. 298, pp. 122–133, Jul. 2018, doi: 10.1016/j.neucom.2017.11.072.

[34] T. Li, A. K. Sahu, A. Talwalkar, and V. Smith, “Federated Learning: Challenges, Methods, and Future Directions,” *IEEE Signal Process Mag*, vol. 37, no. 3, pp. 50–60, May 2020, doi: 10.1109/MSP.2020.2975749.

[35] B. Kar, W. Yahya, Y. D. Lin, and A. Ali, “Offloading Using Traditional Optimization and Machine Learning in Federated Cloud-Edge-Fog Systems: A Survey,” *IEEE Communications Surveys and Tutorials*, vol. 25, no. 2, pp. 1199–1226, 2023, doi: 10.1109/COMST.2023.3239579.

[36] J. Ni, X. Lin, and X. S. Shen, “Toward Edge-Assisted Internet of Things: From Security and Efficiency Perspectives,” *IEEE Netw*, vol. 33, no. 2, pp. 50–57, Mar. 2019, doi: 10.1109/MNET.2019.1800229.

[37] A. Paranjothi, M. S. Khan, and M. Nijim, “Survey on Three Components of Mobile Cloud Computing: Offloading, Distribution and Privacy,” *Journal of Computer and Communications*, vol. 05, no. 06, pp. 1–31, 2017, doi: 10.4236/jcc.2017.56001.

[38] T. S. Brisimi, R. Chen, T. Mela, A. Olshevsky, I. Ch. Paschalidis, and W. Shi, “Federated learning of predictive models from federated Electronic Health Records,” *Int J Med Inform*, vol. 112, pp. 59–67, Apr. 2018, doi: 10.1016/j.ijmedinf.2018.01.007.

[39] S. Truex *et al.*, “A Hybrid Approach to Privacy-Preserving Federated Learning,” in *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security*, New York, NY, USA: ACM, Nov. 2019, pp. 1–11. doi: 10.1145/3338501.3357370.

[40] Z. Wang, M. Song, Z. Zhang, Y. Song, Q. Wang, and H. Qi, “Beyond Inferring Class Representatives: User-Level Privacy Leakage From Federated Learning,” in *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*, IEEE, Apr. 2019, pp. 2512–2520. doi: 10.1109/INFOCOM.2019.8737416.

[41] S. F. Atashzar, J. Carriere, and M. Tavakoli, “Review: How Can Intelligent Robots and Smart Mechatronic Modules Facilitate Remote Assessment, Assistance, and Rehabilitation for Isolated Adults With Neuro-Musculoskeletal Conditions?,” *Front Robot AI*, vol. 8, Apr. 2021, doi: 10.3389/frobt.2021.610529.

[42] S. Salminger *et al.*, “Current rates of prosthetic usage in upper-limb amputees – have innovations had an impact on device acceptance?,” *Disabil Rehabil*, pp. 1–12, Oct. 2020, doi: 10.1080/09638288.2020.1866684.

[43] Z. Jeddi and A. Bohr, “Remote patient monitoring using artificial intelligence,” in *Artificial Intelligence in Healthcare*, Elsevier, 2020, pp. 203–234. doi: 10.1016/B978-0-12-818438-7.00009-5.

[44] P. Kang and J. JO, “Benchmarking Modern Edge Devices for AI Applications,” *IEICE Trans Inf Syst*, vol. E104.D, no. 3, pp. 394–403, Mar. 2021, doi: 10.1587/transinf.2020EDP7160.

[45] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le, “Learning Transferable Architectures for Scalable Image Recognition,” in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Jun. 2018, pp. 8697–8710. doi: 10.1109/CVPR.2018.00907.

[46] F. Chollet, “Xception: Deep Learning with Depthwise Separable Convolutions,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, Jul. 2017, pp. 1800–1807. doi: 10.1109/CVPR.2017.195.

[47] J. Guerreiro, A. Ilic, N. Roma, and P. Tomas, “Modeling and Decoupling the GPU Power Consumption for Cross-Domain DVFS,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 30, no. 11, pp. 2494–2506, Nov. 2019, doi: 10.1109/TPDS.2019.2917181.

[48] S. Suman, F. Chiariotti, Č. Stefanovic´, S. Došen, and P. Popovski, “Statistical Characterization of Closed-Loop Latency at the Mobile Edge,” *IEEE Transactions on Communications*, vol. 71, no. 7, pp. 4391–4405, Jul. 2023, doi: 10.1109/TCOMM.2023.3277523.

[49] M. Mulder *et al.*, “Manual Control Cybernetics: State-of-the-Art and Current Trends,” *IEEE Trans Hum Mach Syst*, vol. 48, no. 5, pp. 468–485, Oct. 2018, doi: 10.1109/THMS.2017.2761342.

[50] S. Suman *et al.*, “Closed-loop Manual Control with Tactile or Visual Feedback under Wireless Link Impairments.” *IEEE Trans Haptics* (under review).