Applying Radical Constructivism to Machine Learning

A Pilot Study in Assistive Robotics

Markus Nowak • German Aerospace Center (DLR), Germany • markus.nowak/at/dlr.de
Claudio Castellini • German Aerospace Center(DLR), Germany • claudio.castellini/at/dlr.de
Carlo Massironi • Istituto di Psicologia e Psicoterapia Interazionista Psicopraxis, Italy • carlo.massironi/at/gmail.com

> Context • In this article we match machine learning (ML) and interactive machine learning (iML) with radical constructivism (RC) to build a tentative radical constructivist framework for iML; we then present a pilot study in which RC-framed iML is applied to assistive robotics, namely upper-limb prosthetics (myocontrol). **> Problem** • Despite more than 40 years of academic research, myocontrol is still unsolved, with rejection rates of up to 75%. This is mainly due to its unreliability – the inability to correctly predict the patient's intent in daily life. **> Method** • We propose a description of the typical problems posed by ML-based myocontrol through the lingo of RC, highlighting the advantages of such a modelisation. We abstract some aspects of RC and project them onto the concepts of ML, to make it evolve into the concept of RC-framed iML. **> Results** • Such a projection leads to the design and development of a myocontrol system based upon RC-framed iML, used to foster the co-adaptation of human and prosthesis. The iML-based myocontrol system is then compared to a traditional ML-based one in a pilot study involving human participants in a goal-reaching task mimicking the control of a prosthetic hand and wrist. **> Implications** • We argue that the usage of RC-framed iML in myocontrol could be of great help to the community of assistive robotics, and that the constructivist perspective can lead to principled design of the system itself, as well as of the training/calibration/co-adaptation procedure. **> Constructivist content** • Ernst von Glasersfeld's RC is the leading principle pushing for the usage of RC-framed iML; it also provides guidelines for the design of the system, the human/machine interface, the experiments and the experimental setups.

> Key words • **Machine learning, interactive machine learning, radical constructivism, assistive robotics, human-machine interaction, co-adaptation.**

Introduction

« 1 » According to Arthur Samuel (1959), *machine learning* (ML from now on) is "the subfield of computer science that […] gives computers the ability to learn without being explicitly programmed." Can radical constructivism say something useful about machine learning, something which would enrich its capabilities, our understanding of it, and possibly shed light on learning *tout court?*

« 2 » First of all, what is machine learning? For the benefit of those readers starting from a realist perspective, let us look at it, at least initially, using a realist language. Samuel's definition is to some extent correct: indeed, ML is an "explicit program," since it runs on computers, and today's computers must still be programmed in the old-fashioned way; but it is a program that observes statistical regularities in the world and matches them against one another.

« 3 » As a direct consequence of this, the output of ML will sometimes not match our expectations, i.e., "it will do the wrong thing," and not as the result of a bug. This, as a realist statistician would put it, "is due to the uncertainty inherent to statistics – one can never be statistically sure that something is true." Or, as a hypothetical realist (the most common sub-type of realism among the ML community) would put it: "statistical truth is only true most of the times." Therefore, a program that searches for statistical similarities in the world will now and then, e.g., deem as similar two things which the researcher defines as not belonging to the same category, and vice versa. This is correct and must be accepted, as opposed to bugs in standard programming, which are always bad and must be eliminated.

« 4 » The only way to "debug" an ML program is to show it more exemplary regularities – to "teach" it something more about the world – to enrich its own model of how the world works – to help it to better organise its own private world according to the researcher's idea of the world.

« 5 » More concretely, ML builds a mathematical function (a "model" from now on) approximating the observable behavior of some variables of an unknown process of interest, given some very basic restrictions on the shape of the model itself, and a set of examples – a set of input data (values of the variables) and corresponding target values to which each datum is associated, sampled from the process itself. This set represents the regularities so far observed during the past behavior of the process. The model, which compactly represents them, can be used to predict the future behavior of the process (an excellent introductory text is Shalev-Shwartz & Ben-David 2014).

« 6 » For instance, an ML model can be built using a set of images acquired from a street camera and corresponding (face-yes/face-no) values, denoting whether an image contains a human face or not. After the model has been built, will it correctly identify new images as containing/not containing a face? Another example: an ML model of the temperature of the Mediterranean Sea can be built using a set of temperature values and the times at which they were observed. Will the temperature of the Mediterranean at specific future times be correctly predicted by the model?

« 7 » Mathematically speaking, the model is built by minimising a cost functional associated to the examples. It is an *optimal fit* of the examples, naturally endowed with the ability to both compactly explain the past target values for each known input datum, and to approximate target values associated with so-far-unseen input data. The model is therefore an attempt to "make sense" of the examples, to "organise" them, to use them in order to predict the future behavior of the process.

« 8 » It obviously follows that the quality of the model (its predictive power) depends on how much the samples collected so far are representative of the behavior of the process both in the past and in the future. So, the answer to the questions posed in §6 is "yes, *provided that a good set of examples was collected in the beginning.*"

« 9 » Notice that the minimisation of a cost functional is a completely mechanical procedure; moreover, no *a priori* physical knowledge about the process to be modelled is, in principle, required – only the ability to draw examples from it. In this sense, an ML model is indeed a machine that "learns without being explicitly programmed" – a softer, perhaps more flexible way of telling our computers what to do, than programming. And the idea is a winning one: ML has recently (at least to some remarkable extent) solved problems that were considered beyond the reach of computers, e.g., form detection in pictures, automated medical diagnosis, speech recognition, content analysis of a text, the game of *Go,* etc. So far, so good.

Is machine learning a radical constructivist business?

« 10 » A great deal of the research in ML seems to suffer from a methodological weakness: machine learning tends to be used as a number-crunching black box, at which to throw as many examples as possible, hoping that it will yield a usable relationship between input data and target values. Too often, scarce attention is paid to the quality, the origin and the meaning of the examples (e.g., Wagstaff 2012). Moreover, examples are considered to be "the reality," rather than being considered artefacts manufactured by the researcher's explicit or implicit choices. The whole procedure suffers from an insufficient awareness of the epistemological problem.[1]

« 11 » This weakness stems, in our opinion, from a widespread *realist* attitude to knowledge and learning, in statistics in general and in ML in particular. A "realist statistician," we can say, assumes that "there is a world out there" and that "we can build a real, even if somewhat rough, model of this world." Once such a model is built, no further changes are needed. In the case of ML, the example set represents knowledge about the world out there, given at the beginning of time, used to predict the future evolution of the target process.

« 12 » In one sentence, ML is so far prevalently *a realist love affair,* for *realist statisticians.* But even a realist statistician (and those who adopt some form of realism) *may* observe that there are indeed many cases in which *this attitude will fail;* in particular, it will fail whenever too few examples are available, e.g., because they are expensive to collect, or if the process of interest is non-stationary, implying that the examples collected at the beginning of time will at some point no longer represent its behavior.

« 13 » Thus, we propose to shift the attitude to ML from realist to *radical constructivist,* as radical constructivism (RC from now on) is defined by Ernst von Glasersfeld (e.g., Glasersfeld 1983, 1995).

« 14 » There are at least four remarks suggesting such a change in paradigm to a realist ML researcher.

« 15 » *In the first place,* let us notice that if we strip the concept of ML to the bare bones, all we are left with (§5) is an agent that tries to *organise perceptual objects*, obtained through specific sensory channels, as best as it can. No physical, chemical, mathematical, ontological, …, knowledge about the process of interest is required. This means that in ML, no knowledge of "external reality" need be assumed. ML deals only with "perceptual" data. This is a very radical-constructivist concept (Glasersfeld 1995: 58f) that we call in short "the construction of experiential reality."

« 16 » *In the second place,* ML is about *matching "perceptual" patterns* – finding regularities among subsets of examples, compactly representing these regularities and using them to predict new target values (§5 again). That is what an ML model does.[2] Not incidentally, matching perceptual patterns is also one of the foundations of RC: "learning as a constructive activity" (Glasersfeld 1983).

« 17 » *Thirdly,* consider again the realist attitude to ML (§11): as opposed to the realist statistician, for the RC statistician indeed "there is a world out there," but as well "we cannot build a *real* model of this world – we can only build a viable representation of it (one of the many possible), useful to do something specific in it" (utilitarianism) *and* in agreement with our pre-conceptions of this world (conceptual coherence). We are continually forced to test the viability of this representation, for our specific purposes and according to our pre-knowledge, through our interaction with the world. The value of an idea of the world is measured in term of fitness to achieve a specific goal *and* (better) fitness against other ideas the subject has about the world, not in term of the correspondence between the idea and a mind-independent reality (Glasersfeld 1995: 68f). We say in this case, that viability is *utilitarianism* plus *conceptual coherence.*

« 18 » *Fourthly,* "pre-knowledge and learning." The realist attitude to ML assumes that knowledge is free from pre-knowledge. The radical-constructivist attitude, as opposed to that, contends that knowledge –

1 | Deep learning coupled with big data represents an unfortunate push in this very direction, albeit a very successful one from a practical point of view.

2 | Actually, *pattern matching* or *pattern recognition* is the old umbrella term for ML.

every possible segmentation of the perceptive field – depends on, and is shaped by, the subject's *pre-definition of what can be seen in the perceptive field*, and that this pre-definition is shaped, in turn, by the interaction the subject has had with the others and with the world (learning). Furthermore, according to RC, the subject does not interact *with the other* (and the other's signals), but only and exclusively *with her perception of the other* (and of the other's signals) and with her previous personal ideas of the other and of the world, since human beings cannot access the "real world" (a mind-independent reality) but only *their perception of the world*. This is a very different model of interaction from the realist one

« 19 » Particularly, during interaction with the others, the subject

a recognizes a specific situation according to her memorized "schemata,"
b performs a specific activity associated with the situation, and
c checks her own specific expectations that that activity should produce a specific previously experienced result.

If this does not happen, the subject is *perturbed* and forced to review her initial sensory elements to find a new structure in these sensory elements and eliminate the perturbation. *All* these processes are presumed to be *subjective* and *internal* to the cognizing agent (again, Glasersfeld 1995: 68f). So, an ML engineer will endow her ML system with an initial simple set of schemata (pre-education) and an engine to apply these schemata, having expectations, possibly to be disconfirmed, and trying to reshape her *sensory material* (learning). This actually is, and we can call it in short, "von Glasersfeld's learning theory."

« 20 » Therefore, an RC statistician engaged in ML would, as opposed to her realist colleague,

a collect examples according to her cultural pre-conception of the world,
b build a temporarily viable model of the world – *viable* according to her own explicit or implicit goals and pre-assumptions/pre-definitions of the world,
c have expectations and check how well the model works, and if the response is not good, she would
d try to reorganize the examples and/or collect new ones, with which to update the model – go back to step (a).

From a (realist) engineer's perspective, this endless loop aims at countering the potential non-stationarity of the process to be modelled.

« 21 » From what we have said so far, it almost appears as if ML already were an RC business. In order to complete the picture though, we also need to enforce the loop outlined in §§19f – we need the ability to *have expectations and update the model at any time,* specifically whenever it does no longer reflect the expectations about the underlying process or the system's goals – whenever its predictive power has become unsatisfactory. Updating a model means changing it in order for it to accommodate old and new knowledge – to accommodate new examples, gathered on demand without the need to obliterate all past knowledge. (Notice that sometimes some of the past knowledge *must* be forgotten, but it is essential not to be *forced* to forget it upon updating!) Model updates must be triggered by some kind of feedback from the world confirming or perturbing the model, perhaps an external agent, able to judge the model's current performance, on the basis of some well-defined purpose.

« 22 » Although little practiced (and even less theoretically studied) in ML literature, this idea already exists and is called, not incidentally, *interactive machine learning* (iML from now on). iML adds to standard ML the possibility of being helped by an external agent, recognising that the predictive power of the current model has become insufficient, and that a new data gathering and model update is required. iML, so far, has been tested in conditions that are particularly hard for standard ML, such as recognising the presence of complex structures in an image: whenever the model failed to correctly categorise an image, a human operator would weigh in, give the system a further example, and request a model update.

« 23 » Interestingly, iML has recently been linked to (non-radical) constructivism by Adavait Sarkar, who claims that

> "the interaction loop of interactive machine learning systems facilitates constructivist learning, as it maximises the interaction between the end-user's experience of the model, and their ideas regarding the model status." (Sarkar 2016: 1472)

However, this is, to the best of our knowledge, the only case so far in which these two fields have talked to each other. This, although iML has been used and implicitly defined in a number of cases (for instance, in Fails & Olsen 2003; Iturrate et al. 2015; Strazzulla et al. 2017). Some recently revamped ML approaches, e.g., recurrent neural networks, can even be viewed as "interactive in nature"; to the best of our knowledge, however, a coherent conceptual framework about interactivity in ML is still missing, and this is where RC can help.

« 24 » In practice, interactivity is enforced through *incrementality*. An incremental ML system is precisely an ML system that allows for updating/downdating its current model. The good news is that, in principle, any standard ML system can easily be turned into an incremental one by storing the examples seen up to now, and whenever a model update is (somehow) triggered, adding the new examples to the old ones, selecting the examples of interest from the new example set, and then re-building the model from scratch using the selected examples only.[3]

« 25 » To some extent iML, as enforced so far in literature, already smells like RC; but this generally remains an intuition of the researcher – there is no adoption of a theory of knowledge and learning as RC. It is *interaction* conceived as a realist scientist can conceive it (sometime as an anti-theoretical scientist can conceive it). Actually, through the interaction, the iML system, in the intention of a realist scientist, builds a "true" model of reality (this way bypassing the problem of a changing reality) simply by "adding input data." We claim that adding to this an epistemological awareness and a more robust learning theory, as offered by RC, will open new paths of research and technological improvement.

« 26 » Our argumentation shows that in the end it will be useful to adopt an RC-framed iML. The tentative framework we sketched above is an attempt at opening a discussion between the RC community and the ML community to enrich our idea of RC-framed iML.

3| This solution can be computationally/memory intensive but there are ways around the problem in the majority of the cases of interest.

Radical constructivist machine learning in action

« 27 » This new point of view of an RC-framed iML raises the question: what is it useful for? An immediate, almost trivial idea (also inspired by the definition of iML in§§22f), is that *human-machine interaction* should be the typical problem area in which ML, and iML, can be empowered by RC.

« 28 » A second, perhaps less immediate way of empowering ML with RC consists in empowering the statistical analytical tools behind ML with the ideas outlined in §§15f, e.g., giving the ML system a "set of schemata" (a sort of "culture"), some pre-selectors to segment its perceptive field, to pre-treat/pre-interpret the information it will crunch and match (to have some "expectations" on the world and the possibility of being "perturbed"). We can call this pathway "crunch *before* match."[4]

« 29 » We talk about human-machine interaction whenever a human subject must guide, teach, control a machine (a robot, a computer, a virtual avatar, etc.) that is endowed with only limited autonomy (Card, Newell & Moran 1983). Here, the standard ML tools at the disposal of the engineer usually fail since, for the machine, modelling human behavior is extremely hard; nevertheless, it is needed to some extent, if one wants to detect the subject's intent, that is, what the subject wants the machine to do. Human behavior is non-stationary, complex, culture- and goal-directed, almost unpredictable in the medium and long run; plus, usable examples from humans can be excruciatingly hard to obtain. All these aspects make the problem of human-machine interaction extremely hard for "realist" ML.

« 30 » So, what is needed in human-machine interaction is a way to constantly "monitor" the desires of the subject, continually gather new examples and learn from her, engage her in a dialog with the expectations of the ML system – the perfect problem for an RC-framed iML system. In addition, in this field we have the almost obvious chance to exploit the judgment of the human as the feedback system/external agent ("the world talks back to the *knowing machine*") mentioned in §§21f, to trigger the perturbation and the model updates (Castellini 2016). The match with the RC concepts of assimilation, scheme theory, accommodation, and equilibration is hereby clear: the ML system must have the capability to be perturbed and to re-equilibrate the perturbation produced by the interaction with the world into its model of the world/of human intent.

« 31 » We claim that RC-framed iML could be a more useful/interesting/sophisticated choice than traditional ML and iML, and especially so whenever dealing with "feedback" in human-machine interaction.

« 32 » We have arrived at this idea in a somehow non-linear way. Namely, the need for iML in human-machine interaction stems from the frustration of the second author of this article, an engineer who has been trying for 10 years to build smart prosthetic arm/hand control systems *(upper-limb myocontrol),* which is a typical case of human-machine interaction. Unhappy with ML-based myocontrol, he recently tried to evolve ML pursuing "a more interactive pathway" (Castellini 2016); while doing so, he faced a new set of problems which called for appropriate conceptual tools. The encounter with the third author of this article, a trained RC psychologist dealing with decision-making models and decision-support systems in the subfield of investing, resulted in the usage of the concepts of *knowledge, learning, communication,* and *feedback* as traditionally developed within RC.

« 33 » Upper-limb myocontrol (Fougner et al. 2012) consists of using muscle activation of the remaining upper limb of an amputated human subject to detect her intention to move and accordingly control a prosthetic arm/hand to perform the desired action quickly, precisely, safely and reliably – in this case ML is used to transform such activity (input data) into control commands (target values). Most such systems are, currently, realist ML systems: a great deal of arm/hand/muscle configurations are gathered initially, a model is built, then its accuracy is tested while the amputated subject tries to control the prosthesis. The few exceptions (e.g., Gijsberts et al. 2014; Hahne, Markovic & Farina 2017; Mathewson & Pilarski 2017) are proving to work in practice. The only commercial solution enforcing ML, namely the *Complete Control* system by CoApt LLC, employs iML in the form of the option to "re-calibrate" the prosthetic control system whenever the user so wishes (Lock et al. 2011; Simon, Lock & Stubblefield 2012; also, personal communication by Blair Lock, CEO of CoApt LLC, 2017).

« 34 » The traditional realist approach to myocontrol still fails after 40 years of research (Jiang et al. 2012; Farina, Jiang & Rehbaum 2014), the main problem being *unreliability:* the inability to guarantee that an arm prosthesis will do exactly what the subject wants, for exactly the length of time she wants. Unreliability can be catastrophic (e.g., prosthetic hand unwantedly releasing the steering wheel while driving) or in the best case "just" frustrating, humiliating and socially unacceptable. As a consequence of this impasse, the acceptance of self-powered prostheses by upper-limb amputees is very limited, with rejection rates of up to 75%. Simply put, state-of-the-art upper-limb prostheses do not work well enough to justify the cost and effort required to use them (Micera, Carpaneto & Raspopovic 2010; Peerdeman et al. 2011; Castellini et al. 2014).

« 35 » Our goal is to show that RC-framed iML could change the situation. Unreliability arises from the innumerable variety of different situations in which a prosthesis must perform a certain action. For instance, maintaining a firm grasp on a rail must be ensured notwithstanding external force disturbances, changes in applied muscle activation, the posture of the arm, etc. Since it is *de facto* impossible to build an example set, at the outset, containing examples of all these situations and all further possible ones, sooner or later realist ML-based myocontrol will fail (Castellini 2016). A spectacular example of this is represented by the outcome of the ARM competition of the 2016 *Cybathlon* – see, e.g., Wolf Schweitzer, Michael Thali & David Egger (2018) for a detailed analysis of the current pitfalls and practical requirements of myocontrol.

« 36 » As opposed to that, existing iML-based myocontrol counters this problem ex-

4| What statistical analysis can gain by adopting the RC perspective – how the "maths" can be used differently – is a very interesting research agenda for the future; notice that nowadays the ML community "teaches a culture or schemata to the system" by choosing an ML algorithm specific to each different task the ML system needs to pursue.

actly thanks to on-demand model updating: whenever a new situation arises in which it fails, the subject "teaches" the system how to cope with it; the system, in turn, readily adapts to the new knowledge (Castellini 2016). Unsurprisingly, iML-based myocontrol already is reported by, at least, Arjan Gijsberts et al. (2014) and Ilaria Strazzulla et al. (2017), where, however, very little is said about the most proficient/natural way to design and enforce the interaction between the system and the user. Another crucial aspect or side-product of iML, namely co-adaptation, is only now being explored (Hahne, Markovic & Farina 2017), yet there is no indication of what theoretical framework could/would optimally guide the design of the interaction interface.

« 37 » Is our claim that RC-framed iML is superior when dealing with "feedback" in human-machine interaction (§31) justified, specifically as far as upper-limb myocontrol (§35) is concerned? Does RC-framed iML enforce better myocontrol than realist ML? This is a subset of our RC research agenda: "Can a knowing subject (here an ML system) that adopts *a non-realist theory of knowledge* do better than a realist one?" To shed some light on this question we have compared an RC-framed iML-based myocontrol system with a traditional ML-based myocontrol system in a pilot experiment involving human subjects.

Experiment

Overview

« 38 » Ideally, two upper-limb myocontrol systems, an RC-framed and interactive one and a non-interactive (RC versus realist) one would be compared during completely unrestricted daily-living usage by two distinct groups of amputated subjects. Here we adopted some simplifications. Firstly, we engaged fifteen intact human subjects only; secondly, we used two 3D hand models displayed on a computer screen instead of tangible prosthetic devices.

« 39 » The experiment as a whole consisted of three sub-experiments, each of which will be from now on referred to as Experiment 0, 1 and 2, respectively. Namely, we compared (Experiment 0) a traditional, non-interactive, *realist* upper-limb myocontrol ML system, with (Experiment 1) a "part-time" RC-framed interactive ML system, a "weakly RC" system, and (Experiment 2) a "full-time" RC-framed interactive ML system, a "fully-fledged" RC system.

ML method

« 40 » Following the motto "keep it as simple as possible, but not simpler than that," attributed to Albert Einstein, all three ML systems employed in the experiment are based upon least-squares regression in the regularised form called *Ridge Regression*, filtered through a non-linear mapping called *Random Fourier Features* (RR-RFF). RR-RFF exists both in "batch" form (i.e., non-incremental and therefore non-interactive) and in incremental form (iRR-RFF – for the mathematical details see Castellini 2016). Notice that iRR-RFF is guaranteed to yield the same optimal model as RR-RFF, whenever the same example sets are used with either method. This enables a fair comparison between an ML and an RC-framed iML system even from an exquisitely mathematical point of view.

Participants

« 41 » Fifteen intact human subjects (5 females and 10 males, age 19–54 years) participated in the experiment. Before the experiment took place, it was clearly explained to each participant, both orally and in writing, that no health risk was involved. Each participant signed an informed consent form. The experiment was previously approved by the internal committee for data protection of the Institution where the experiments took place, and it followed the guidelines of the World Medical Association Declaration of Helsinki. The participants were randomly assigned to one experiment only, so that five participants took part in each experiment.

Experimental setup

« 42 » The experimental setup was common to all three experiments, and consisted of

- a *Myo* bracelet by Thalmic Labs,
- two 3D hand models displayed on a computer screen, and
- a simple voice reproduction/speech recognition system.

« 43 » The Myo bracelet (https://www.thalmic.com) consists of eight uniformly spaced sensors, able to detect the electromyographic signal generated by the muscle activity of the subject's forearm.

« 44 » The 3D hand models realistically mimic the motions of a human wrist and hand. One of the models is white while the other is rendered in skin-like texture; the former (from now on referred to as *the stimulus*) is used to provide visual stimuli to the participants, i.e., it is controlled by the software itself; whereas the latter (from now on referred to as *the prosthesis*) simulates the prosthesis – given the current model, it enforces the predicted motions of the hand and wrist, as evaluated from the data provided by the bracelet.

« 45 » The speech recognition and synthesis system is the one embedded in Microsoft.NET Framework 4.5, able to distinguish a small set of words (in this case, "good"/"bad") pronounced by the subjects, and to utter predefined voice messages, which we configured with a clearly synthetic female voice. All sentences were uttered in the first person and address the participant in the second person (e.g., "you are now going to teach me how I should move").

Experimental protocol

« 46 » Each participant sat comfortably in front of the computer screen, and the bracelet was wrapped around her forearm. She was instructed to hold the forearm vertically, leaning the elbow on the table.

« 47 » A voice message was played. The prosthesis "spoke," explaining to the subject that the "*training session*" (the experiment) was about teaching it – a new kind of hand/wrist prosthesis able to learn – how to properly perform the movements intended by the participant, and that to this aim, a 3D hand model would eventually appear on the screen, representing itself (the prosthesis). The prosthesis *clearly stated* that the participant would not be judged on her performance, but rather that she was going to "show" the prosthesis how each single movement was to be performed, by simply doing it; rather, the prosthesis was to be judged by the participant on its learning ability. Particular care was taken by the prosthesis in asking the participant to be patient and not to get disappointed if it did not correctly execute the required task. After all, concluded the prosthesis, its learning ability was "only in its infancy."

« 48 » We chose this way to communicate to the subject what to do in the experiment, to try and build a psychological context of interaction rules and reciprocal roles, potentially inducing in the subject the construction of positive "emotions" (Harré 1986) toward the "learning" prosthesis. More generally, while designing the experimental protocol, we also tried to take into account the main criticisms raised by postmodern social psychology with regard to the way human subjects are treated in most experimental psychology (and in general in experiments involving human subjects), i.e., that experimenters neglect to offer to (co-construct with) the subject a semiotic definition of the experimental situation meaningful to her, and take into account the meaning of the experimental setting from their own perspective only (Gergen 1978a, 1978b, 1985; Gergen & Gergen 1985; Harré 1979; Harré, Smith & van Langenhove 1995). Figure 1 shows a schematic representation of the experimental setup.

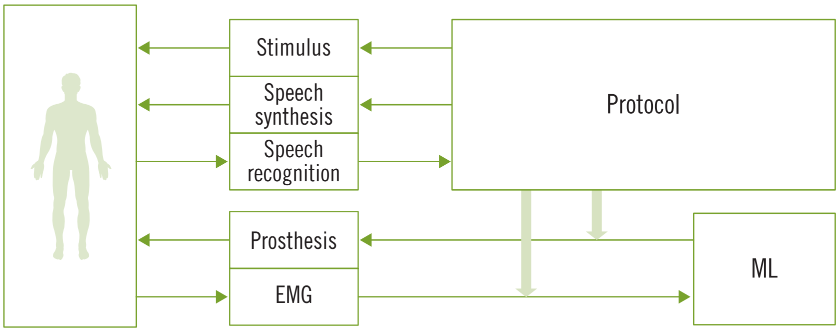

Figure 1 • A schematic depiction of the experimental setup. Subjects interact with the protocol controller via speech recognition, speech synthesis, and by looking at a PC screen on which the stimulus is displayed. The ML method "converts" EMG signals into live configurations of the prosthesis to be shown as well to the subject; while the protocol controller establishes what to display and utter, and when to open/close the flow of information between the subject and the ML method. The protocol controller, together with the ML method, constitute a flexible framework for all three experiments, allowing for different levels of interactivity.

Experiment 0

« 49 » Experiment 0 (the realist machine learning system) consisted of two phases that we will call model building (MB) and model testing (MT).

« 50 » At the beginning of MB, the stimulus was shown on the screen; the prosthesis then explained that "the white hand on the screen" (the stimulus) would now perform a series of hand and wrist movements (tasks), and that the participant should simply mimic what the stimulus was doing with her hand and wrist, as accurately as possible, in order to give the prosthesis a chance to "try and understand" what each movement looked like when seen through the signals it received from the bracelet.

« 51 » Soon afterwards, the stimulus was shown on the screen. A randomised sequence of 30 tasks (6 actions, each action repeated 5 times), was played by the stimulus. The actions were: no-action; wrist flexion; wrist extension; wrist pronation; wrist supination; and hand closing. No voice interaction was provided during this phase. MB would end at the end of this sequence.

« 52 » In this experiment, the ML system had, so to speak, the expectation that all signals it would receive would be "good" signals. *The system would experience no perturbation in the building of its inner vision of the "world."* In other words, each signal was "fitting" with previous signals, and the system was forced to accept all signals as good ones, upon which to build its own "reality" (model).

« 53 » In practice, the model was evaluated in the interval between MB and MT, using the data collected during MB. The evaluation took a few seconds, so that no apparent interruption would be felt by the participant.

« 54 » At the beginning of MT, the forearm of the participant would be hidden from view using an opaque cardboard partition (this is our rough approximation for the subject "wearing" the prosthesis); then the prosthesis would appear on the screen, beside the stimulus. It would then explain that now the stimulus would show a further series of tasks, similar to those that had appeared during MB, that the participant must reproduce those actions, and that the prosthesis would try to understand the signals it received from the bracelet and mimic the action performed by the participant as best it could.

« 55 » It also explained that, after each task had been performed, the prosthesis would verbally ask that the participant evaluate its performance; the participant would then be asked to say "good" or "bad" according to her own judgment.

« 56 » Soon afterwards, a further, randomised sequence of 90 tasks (the same 6 actions as during MB, but in this case each action was repeated 15 times) was played by the stimulus. After each task, the judgment would happen: the prosthesis would ask how it had performed, and the participant would answer "good" or "bad." Figure 2 shows a bird's eye view of the experimental setup while a subject was engaged in Experiment 0, MT.

Experiment 1

« 57 » Experiment 1 (the "part-time" interactive machine learning system) consisted of two phases like Experiment 0. In this case, however, the hand of the participant was hidden behind the cardboard partition already during MB, and the prosthesis would be immediately visualised beside the stimulus (the participants "wore" the prosthesis from the beginning). The same randomised sequence of tasks as in Experiment 0 was played by the stimulus; but in this case, after each task, the participant would be asked by the prosthesis to evaluate its own performance, just like during MT of Experiment 0.

« 58 » If the participant answered "good," the data gathered during the task was directly added to the machine-learning

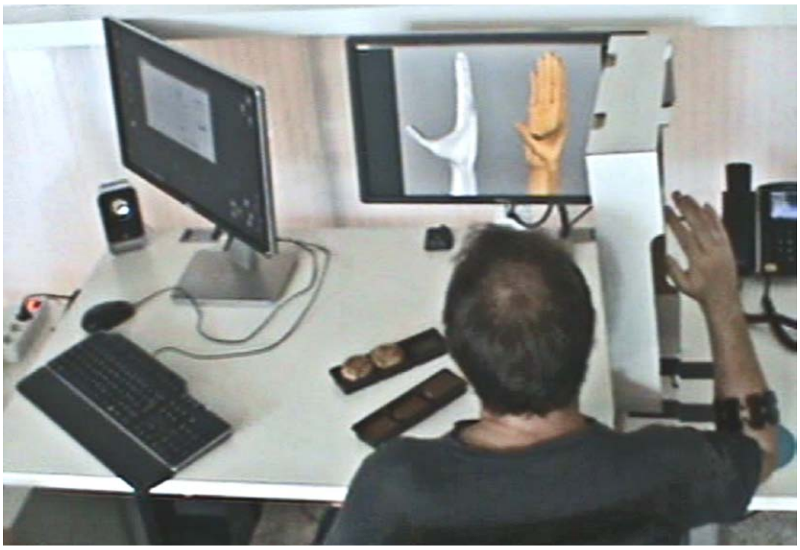

Figure 2 • The experimental setup while a subject performs Experiment 0, MT. The stimulus (white hand) and the "prosthesis" (skin-textured hand) are displayed on the screen; the subject's right arm (wearing the Myo bracelet) and hand are shielded from view using a cardboard partition.

model in order to reinforce the positive result. If the participant answered "bad," she was vocally instructed to perform the task once again, and the data collected during this new instance of the action would be added to the model in order to correct for the previous negative performance. In both cases, the model would be immediately re-evaluated in order to reflect the new acquired data without delay. This way, assimilation and accommodation directly enter the picture of iML: via the external/human feedback.

« 59 » In this case, we can say, the ML system had the expectation that all signals it would receive would be "good" ones, but at the same time it would indeed experience some perturbation (the negative human feedback), so it was forced to not assimilate all signals, but rather to accommodate some specific ones, changing its recognition pattern and building a different scheme (model).

« 60 » MT in Experiment 1 was identical to MT of Experiment 0.

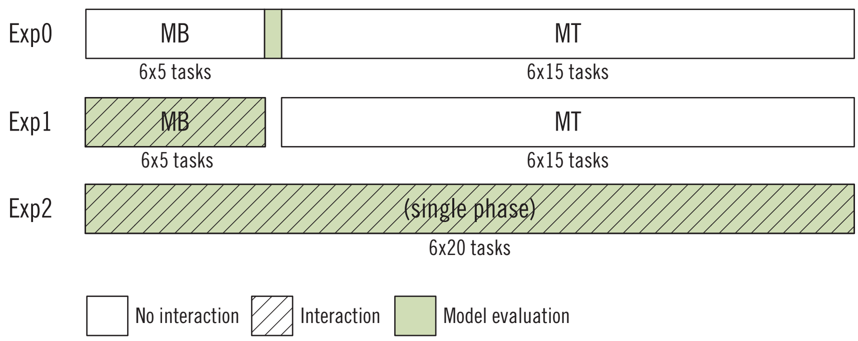

Figure 3 • A graphical representation of the three experiments. Phase MB (Model Building) of Experiment 1 and the entire Experiment 2 are interactive; model generation happens between phases MB and MT (Model Testing) in Experiment 0, during phase MB in Experiment 1 and during the entire experiment in Experiment 2.

« 61 » Substantially, Experiment 1 consisted of *a partially interactive version of Experiment 0:* during MB, the participant would offer the prosthesis some confirmation and some perturbation, therefore helping the prosthesis to better learn the patterns corresponding to the required actions, so that in the end the model would reflect the corrections.

« 62 » Notice that the amount of data used to build the model in Experiment 1 was exactly the same as in Experiment 0 (30 tasks) – what changed was the added interaction with the participant, and consequently, the *possibility for the system to have confirmation or perturbation of its inner world* formed with the data gathered during MB.

Experiment 2

« 63 » Experiment 2 consisted of one phase only, identical to MB of Experiment 1, except that the stimulus would play a randomised sequence of 120 tasks (the same 6 actions as in the previous experiments, but in this case each action was repeated 20 times). As in Experiment 1, the model would be re-evaluated after each task (again, according to the "good"/"bad" judgement of the participant); but for each task, *data gathered during the past five repetitions only of this action were used to build the model.* This ensured that the amount of data used to build the model was, again, the same as in the previous experiments (30 tasks).

« 64 » Experiment 2 consisted therefore of a "*continual learning/feedback*" *version of Experiment 1,* enforcing vocal interaction between the participant and the system at all times.

« 65 » According to the RC learning theory, our ML system *did* have a scheme of the world: it had indeed the expectation that all signals received would be "good" ones, but constantly experienced (*internal*) confirmation and perturbation, thus being forced to modify its own model of the world. The system designed for Experiment 2 is our main conceptual (and practical) attempt at arriving at iML via the RC theory of knowledge and learning, prevalently stressing the RC idea of viability relating to "utilitarianism" (see above, §17).

« 66 » Figure 3 graphically represents the three experiments, while Figure 4 shows flow-charts of each experiment.

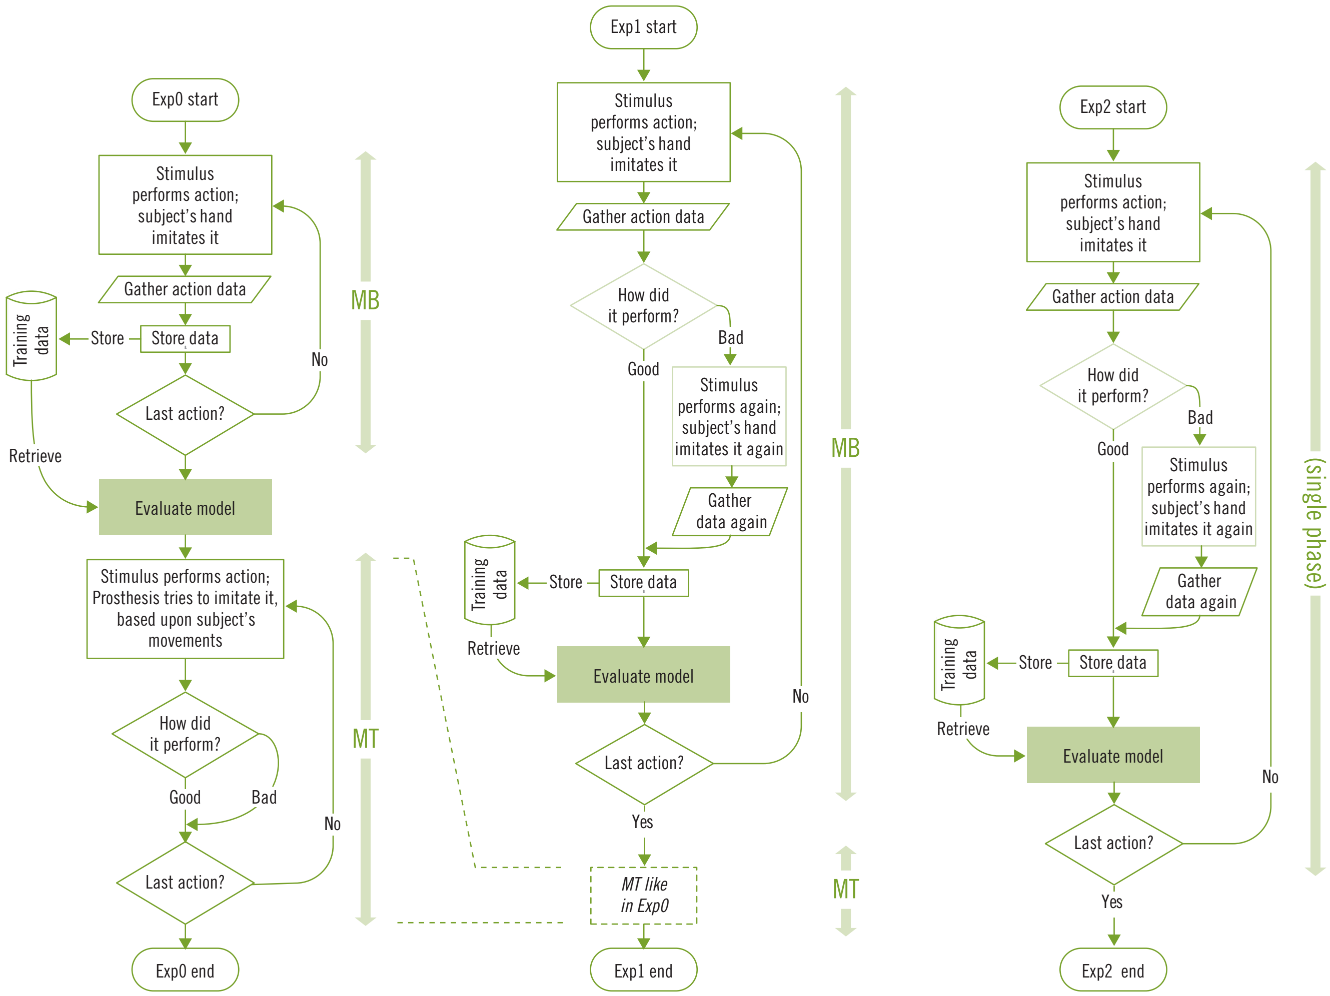

Figure 4 • Flow-charts of the three experiments (from left to right: Experiment 0, 1 and 2). Notice that the single phase in Experiment 2 is identical to MB in Experiment 1. Also notice that model evaluation happens only once during Experiment 0, whereas it happens within an interaction loop in Experiments 1 and 2.

Evaluation measures

« 67 » We wanted to measure which of the three machine learning systems (realist, "part-time" RC-framed interactive and "full-time" RC-framed interactive) could produce a model capable of better understanding the patterns produced by the subject, thereby properly performing (as a prosthesis) the actions that the subject wanted to do. So, we adopted a measure that was objective for the experimenter and the research community of upper-limb myocontrol: the normalised root-mean-squared error (nRMSE) between the position of the stimulus and that of the "prosthesis" during each task – essentially, the discrepancy between the desired position and what the prosthesis manages to do. To evaluate the nRMSE, for each task we considered the last second in which the stimulus was performing the required action, in order to neglect as far as possible any transition effect (i.e., the time the subjects needed to become aware of what was asked of them, and to move their own hand and wrist to the required position).

« 68 » We also wanted to measure which of the three machine learning systems was perceived as the best one by the subjects, so we also adopted a measure that was objective for the subject: the number of poor/good judgements expressed by the subject during the experiment.

« 69 » These two measures described in the two preceding paragraphs make our experimental evaluation akin to the Target Achievement Control test (TAC test, see, e.g., Simon et al. 2011), an assessment test well-known in the myocontrol community; the only remarkable difference is that whether a task is successful or not is left to the participant's judgment.

« 70 » Lastly, we wanted to evaluate the quality of the subject-prosthesis relationship from both the subject's and the machine's point of view – we were interested not only in measuring the performance, but also in checking the reciprocal adaptation. So, first of all we estimated how the signals of each

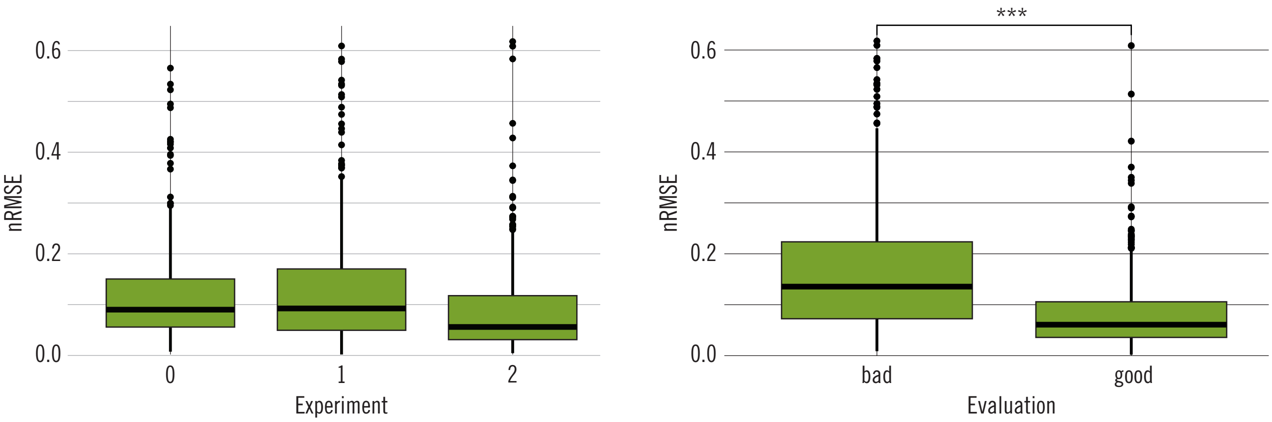

Figure 5 • Left: nRMSE grouped per experiment. Right: nRMSE grouped according to the subjective good/bad judgment. Median values (thick black lines), 25%/75% percentiles ("hinges"), extreme values (larger/smaller than 1.5 times the inter-quartile range) from the hinge (whiskers), and outliers (single dots).

subject changed during the experiments; this was done by evaluating, at each task, Roland Fisher's cluster separateness index (Fisher 1936) for the signal clusters corresponding to the past 30 tasks (one cluster per action, resulting in six clusters). Fisher's index increases the more the clusters are separated, compact and distinct from each other; it represents therefore a measure of improvement in the "quality" of the signals produced by a subject.

« 71 » Moreover, after the experiment, we conducted a semi-structured interview with each subject, focused on

- the quality of the subject-machine learning system interaction,
- the subject's judgment of the system's learning capacity, and
- the fatigue experienced by the subject while teaching the system.

We conducted a qualitative thematic analysis on the semi-structured interview transcripts, through the conventional process of familiarisation with data, generating initial codes, searching for themes among codes, reviewing themes, defining and naming themes, and producing the final report.

Experimental results and analysis

« 72 » In a first round of evaluation, it was determined that subject #13 in Experiment 0 performed exceptionally badly (extremely high nRMSE) while subject #11 in Experiment 1 performed exceptionally well (extremely low nRMSE); data from these two subjects were removed from the analysis as they were considered outliers. To keep the data sets balanced, we also removed one subject's data at random (namely subject #7) from Experiment 2. So, the analysis was based upon data from 4 subjects per experiment.

« 73 » Furthermore, the analysis was conducted on the last 90 tasks only, in order, again, to maintain a balanced dataset, and to avoid considering the inevitable acquaintance effect that each subject went through in the beginning of each experiment (phases MB of Experiments 0 and 1 and first 30 tasks of Experiment 2).

Global statistics

« 74 » Figure 5 shows the global statistics of the experiment. The average nRMSE was 11.61% (SD = 9.23%), 12.81% (SD = 11.39%) and 8.87% (SD = 9.01%), respectively, for Experiment 0, 1 and 2 (left panel, no statistically significant difference was found using a repeated-measures one-way ANOVA test – $F(2, 9) = 3.96$, $p = 0.058$). In order to check whether nRMSE was correlated with the good/bad judgment, we also verified that the nRMSE is on average 7.9% (SD = 6.5%) and 16.76% (SD = 12.48%) in turn, if grouped according to the good/bad judgement (right panel, Welch's t-test yields $t(510.41) = 13.06$, $p < 10^{-4}$). Moreover, the number of good/bad judgments was 230/130 201/159 and 259/101 in turn for Experiment 0, 1 and 2, with a statistically significant difference (the Chi-squared significance test yields $\chi^2(2) = 20.15$, $p < 10^{-4}$).

« 75 » From these results we can say that

- Experiment 2 resulted in an overall slightly better error rate than Experiments 0 and 1, although the high standard deviations reduce the statistical significance of these results;
- Experiment 2 elicited more "good" judgments than Experiment 0, which in turn elicited more than Experiment 1; and
- "good" subjective judgments are positively correlated with lower nRMSE.

All in all, nRMSE values are in line with previous literature obtained from analogous experiments (Gijsberts et al. 2014; Ravindra & Castellini 2014; Connan et al. 2016).

Evolution in time

« 76 » Figures 6 and 7 go into a little more detail, showing the nRMSE and number of good/bad responses for each experiment, subject and task, along time. From these further graphs, we conclude that

- Experiments 0 and 1 produced high values of the nRMSE roughly scattered in Figure 6 all along the course of time (yellowish cells appearing all along the course of the tasks) whereas in Experiment 2 the error seems to settle to lower values in the second half;
- subjects seemed to be much happier (prevalence of "good" judgments) in Experiment 2, especially subjects #3 and #8, than in the other experiments – particularly, subject #4 almost consistently judged the performance as "bad."

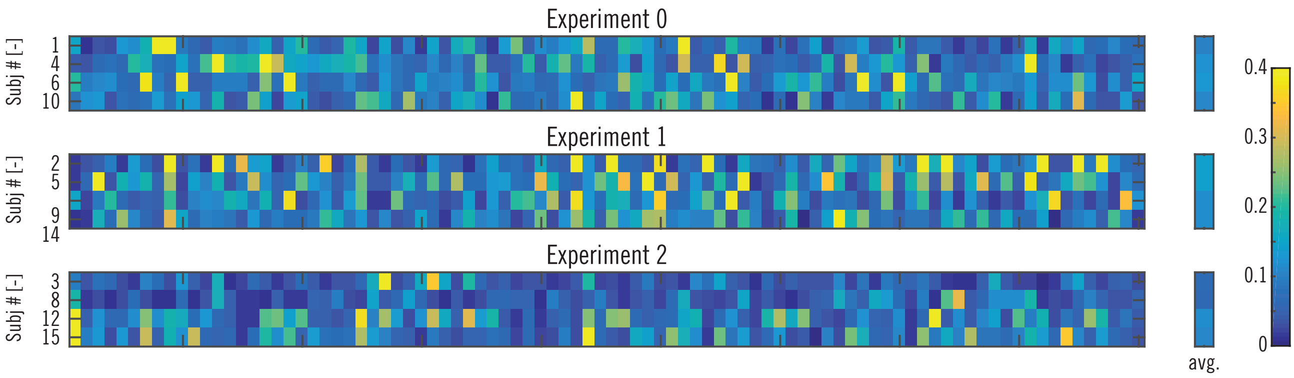

Figure 6 • nRMSE for each experiment, subject and task, as the experiments progressed.

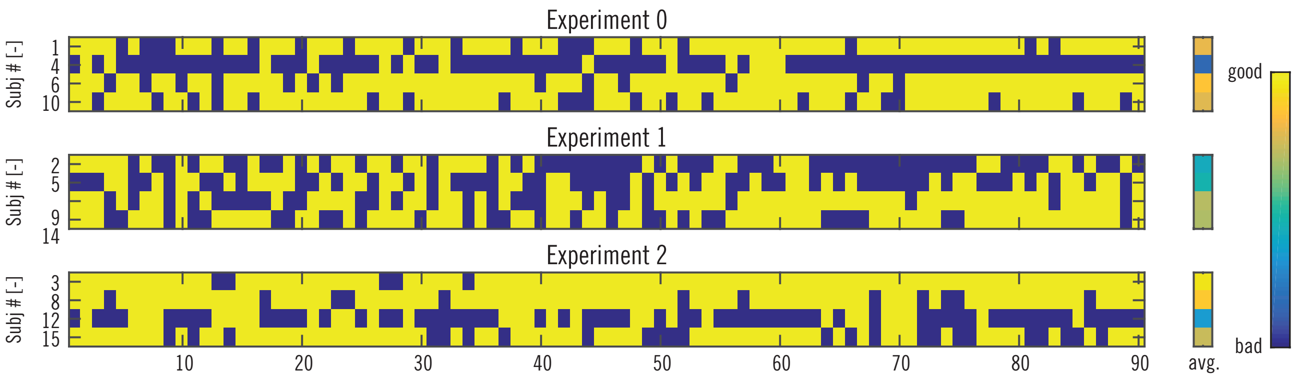

Figure 7 • "Good" (yellow) and "bad" (blue) judgments for each experiment, subject and task, as the experiments progressed.

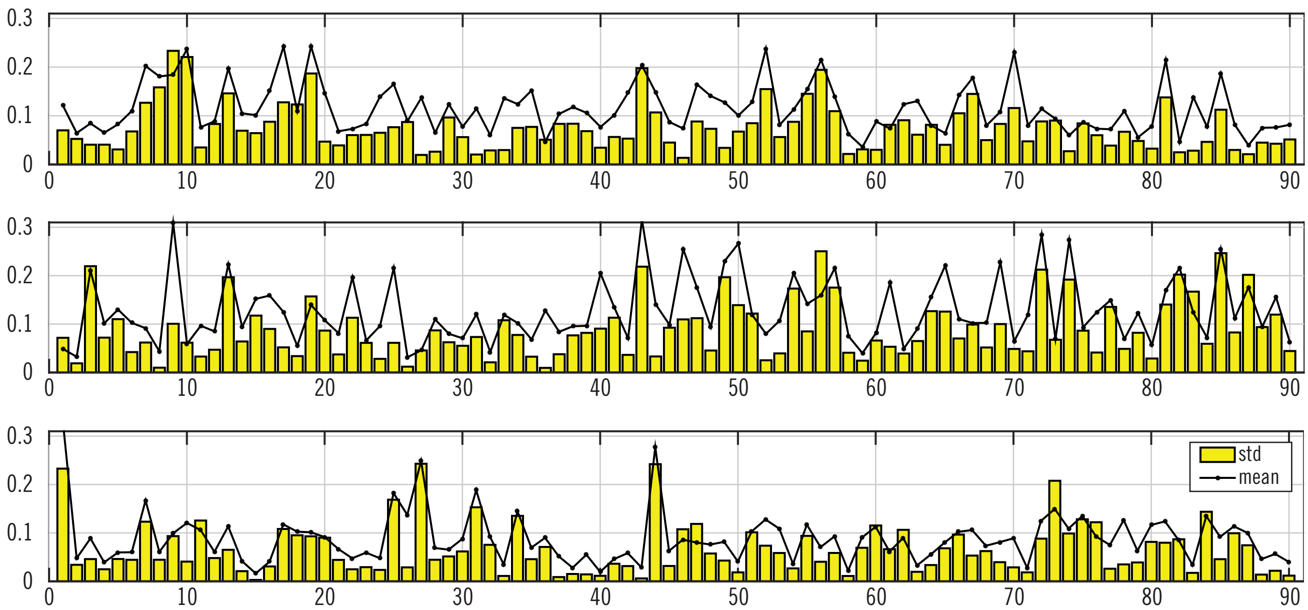

Figure 8 • Mean nRMSE and its standard deviation, averaged across subjects, for each experiment. Top to bottom: Experiment 0, 1 and 2.

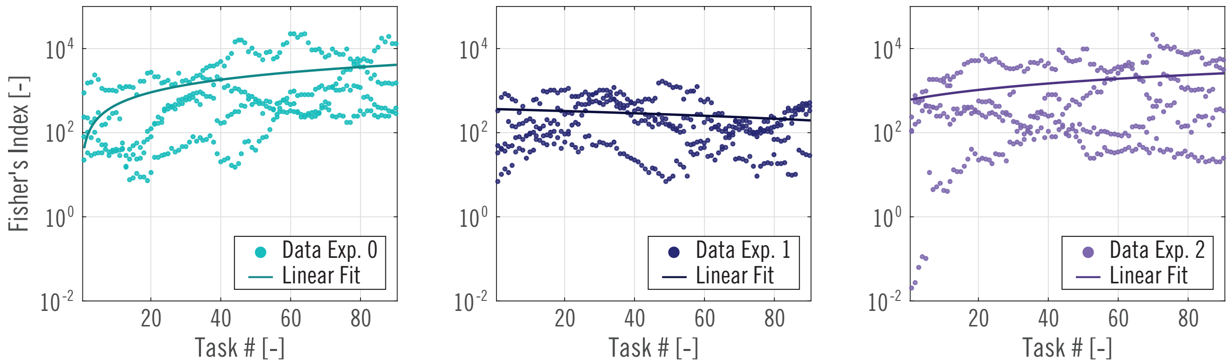

Figure 9 • Fisher's separateness index for each experiment (Experiment 0, 1 and 2 from left to right), plus linear interpolants (since the y-axis is logarithmic, they appear as logarithmic curves).

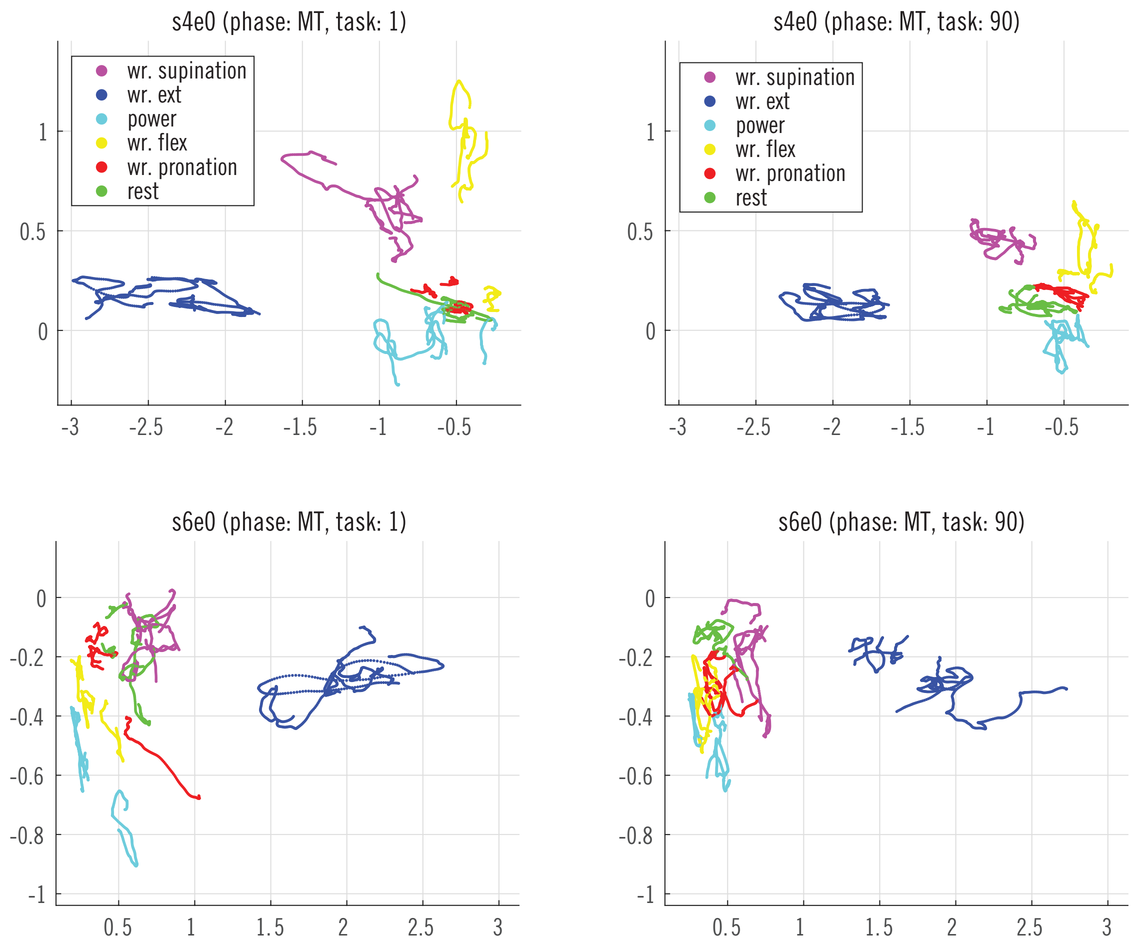

Figure 10 • Change in the signal clusters as two subjects (#4, upper panels and #6, lower panels) progress from task 1 (left column) to task 90 (right column) of Experiment 0. Dimensionality reduction obtained using Principal Component Analysis; the first two principal components retain 90.73% of the signal variance for subject #4 and 95.40% for subject #6. See also "Additional material."

« 77 » Figure 8 shows the temporal evolution of the nRMSE, averaged across all subjects (mean values and standard deviations), which confirms (consider §76 again) that not only the mean values, but also the standard deviations of the nRMSE remain lower in Experiment 2 than in the other two Experiments.

« 78 » Figure 9 shows Fisher's index along time, for all tasks, subjects and experiments. Experiments 0 and 2 elicited, on average, an increase in the separateness of the signal clusters.

« 79 » Lastly, Figure 10 shows 2D-reduced signal clusters obtained from two subjects, #4 and #6, at tasks 1 and 90 of MT in Experiment 0. (These two subjects are chosen as an exemplary good and an exemplary bad subject.) The higher compactness and separateness of the clusters at task 90 (that is, at the end of the Experiment) is apparent, especially for Subject #4. Subject #6 shows poorer cluster separateness, though – five actions appear "lumped" together.

Semi-structured interviews

« 80 » The semi-structured interviews we conducted allow us to conclude, in the first place, that all subjects involved in Experiment 2 *complained about muscle fatigue towards the end,* whereas only one subject not involved in Experiment 2 did. This is due to the increased number of tasks performed, in turn due to the possibility of judging "bad" and potentially having to repeat the previous action at all times. The finding that the nRMSE obtained in Experiment 2 seems not to particularly increase towards the end (Figure 6 and 8, bottom panel), and that its standard deviation remains low (Figure 8, bottom panel), is all the more remarkable.

« 81 » Secondly, no pattern is apparent in the judgments along time (Figure 7); we found that each subject approached the Experiments with seemingly different hopes and expectations. For example, subjects #4 and #12 mostly judged "bad" and both reported posture/muscle discomfort; subjects #2 and #5 judged "bad" quite often, the former reporting difficulty in rating the movements only as good or bad and the latter reporting frustration due to the "continual oscillation" of the prosthesis; lastly, subjects #3, #6 and #8 mostly judged "good," and all reported being "positively impressed" by the progress obtained by the prosthesis in the beginning.

« 82 » It is interesting to note that, upon a closer look at Fisher's index for each subject (not displayed), subjects who mostly judged "good" consistently ended up with higher Fisher's index and vice versa.

General remarks

« 83 » The experimental results shown above let us make a few claims. Given the low number of subjects involved, matching the semi-structured interviews with the experimental results allowed us to add a "layer" of meaning-for-the-subject of what happened during the experiment, offering us clues on how to read the data gathered in the experimental setting. Something particularly useful in a pilot study with a small number of subjects, but also useful in general in experiments involving human beings.

« 84 » The subjective measure of satisfaction, that is the good/bad judgments, is by and large in agreement with the objective one for the experimenter (the nRMSE), as is apparent from Figure 5, right panel: on average, whenever the subjects saw that the prosthesis was "doing the right thing," they judged "good" and vice versa. This shows that the voice and visual interaction was well designed. As opposed to that, the experiment number (0, 1 or 2) turns out to significantly skew the number of good/bad judgments (for instance, Experiment 2 elicited significantly more "good" than "bad" judgments) but *not* the nRMSE: although the error is on average lower for Experiment 2 than for 0 and 1, and lower for 0 than for 1, it is not significantly so. These two remarks seem to somehow collide, but it is not yet clear to us in what sense.

« 85 » There is a significant evolution in time of the subjects' signals (Figures 9 and 10 – see also "Additional material") during Experiments 0 and 2. We speculate that the increase in Fisher's index during Experiment 2 could be due to the concurrent evolution of the subjects and the machines. Notice, however, that during Experiment 0 the ML model was not adapting at all during the MT phase, although some of the subjects involved in Experiment 0 reported that they felt that "the machine was learning."

« 86 » All in all, the "partially interactive" experiment, that is Experiment 1, seemed to produce slightly worse results than the non-interactive one; whereas the "fully interactive" one, Experiment 2, produced slightly better objective results than both the other experiments, and definitely better subjective results – higher satisfaction expressed by the subject.

« 87 » Muscle fatigue seems to have played a significant role in the experiment, which we had not foreseen. Unfortunately, this was mostly the case in Experiment 2 since interaction means more tasks to perform by the subject. Still, the error in Experiment 2 is more "uniform" (lower mean, lower standard deviation) than in the other cases. A further refinement of the experimental design will need to take into account fatigue as an unavoidable problem; but we should at the same time remember that fatigue is one of the factors that make myocontrol a non-stationary modelling problem, that is to say, one of the problems that iML should be better at tackling, in principle.

Conclusion

« 88 » One must admit that, if taken from the point of view of the engineer, the results of the experimental analysis are somewhat disappointing – there is no definite, statistically significant objective improvement (although there is some) when enforcing more interactivity, also conceptualised in line with RC's learning theory. Still, the subjects involved in the experiments generally reported a smoother interaction with the ML system in the case of Experiment 2.

« 89 » Therefore, the missing suggestion that the RC approach gives to the ML practitioner – that of using RC-framed iML – goes farther toward creating a better experience for the user of an upper-limb prosthesis. The application of RC to this problem gives us useful insight into how to design the interactive prosthesis of the future; points us toward more ecological experiments, more deeply embedded in daily life, aiming at enforcing interactivity with the subject at all times, just like it happens with modern gadgets such as, e.g., smartphones. Reciprocal adaptation inspired by RC's learning theory seems to definitely be a factor to be exploited in this field.

« 90 » The main contribution of this work is to reframe iML within radical con-

MARKUS NOWAK

received an MSc degree (Dipl.-Ing.) in mechanical engineering from the Technical University of Munich, Germany. In his master's course of study he focused on medical engineering and numerical mechanics. His master's thesis was carried out in cooperation with the Institute of Robotics and Mechatronics of the German Aerospace Center, Oberpfaffenhofen, Germany. After his masters he continued as a research fellow at the Robotics and Mechatronics Center with a focus on human–machine interfaces for controlling upper-limb prosthetics.

CLAUDIO CASTELLINI

received a Laurea in Biomedical Engineering in 1998 from the University of Genova, Italy, and a PhD in Artificial Intelligence in 2005 from the University of Edinburgh, Scotland. Since 2009 he has been a researcher at the Institute of Robotics and Mechatronics of the German Aerospace Center, Oberpfaffenhofen, Germany, concentrating on human-machine interfaces for the disabled and assistive robotics. He is currently (co-)author of about 85 papers that have appeared in international journals, books and peer-reviewed conferences.

CARLO MASSIRONI

received a Laurea in Clinical Psychology in 1996 from the University of Padua, Italy and a Specializzazione in Cognitive-Interactive Psychotherapy (constructivist and interactionist models in psychotherapy and problem solving) in 2001 from the Istituto di Psicologia e Psicoterapia Interazionista Psicopraxis, Padua, Italy. Since 2003 he has been an associate researcher at the same institute, concentrating on decision making, investment decision making and investment decision support systems. Since 2001 he has worked as an asset manager. He is currently (co-)author of some 10 papers that have appeared in international journals, books and peer-reviewed conferences.

structivism. Although still far from fully tackling the theoretical implications of this idea, in this article we try to show what the potentialities of such a link are. Especially, we expect the marriage between RC and ML to produce, in the near future, a set of guidelines on how to design the ML "statistical engine" and the interaction that is at the core of interactive machine learning: how to reframe interaction and feedback according to RC's learning theory. What should be asked of the human operator, how and when? How should the information so obtained be used? Neither the engineers' community, nor the world of functional assessment can, at this stage, thoroughly answer this question.

« 91 » Extensions to this research should definitely include at least the capability, for an RC-framed ML system, to decide *internally, autonomously* whether a signal is not a good one. This means that the system must trigger *by itself* a perturbation whenever a signal does not fit its conceptualization (model) of the world, thus enforcing the RC idea of *viability* related to *conceptual coherence.*

« 92 » All in all, this article has explored the application of radically constructivist "glasses" to a typical problem in human-robot interaction, and specifically to upper-limb myocontrol. More generally, we have tried to rework, in RC terms, some of the problems faced by ML and iML; we have speculated that the usage of RC-framed iML matches some of the ideas that, among others, von Glasersfeld applied to human learning. Our results suggest that RC-inspired interactivity has the potential to improve human-robot interaction, especially from the point of view of the humans.

Additional material

The dynamic 3D evolution of the clusters from Task 1 to 90 can be seen in two short movie clips at http://constructivist.info/data/13/2/s4e0.gif and http://constructivist.info/data/13/2/s6e0.gif.

Acknowledgements

This work was partially supported by the German Research Foundation's project *Tact_hand: Improving control of prosthetic hands using tactile sensors and realistic machine learning* (DFG Sachbeihilfe CA-1389/1, see http://gepris.dfg.de/gepris/projekt/272314643).

Received: 31 August 2017
Accepted: 26 January 2018

Open Peer Commentaries

on Markus Nowak et al.'s "Applying Radical Constructivism to Machine Learning"

A Radical Constructivist Approach to the Human-Machine Interface

William Craelius
Sichuan Bayi Rehabilitation Center, China, and Rutgers University, USA
craelius/at/soe.rutgers.edu

> Upshot • Thousands of projects aimed at improving the functionality of upper-limb prostheses over the decades have failed to significantly advance the field of assistive robotics. Having been unfamiliar with radical constructivism (RC) so far, I want to see how its approach could contribute, particularly for amputees. Perhaps the most profound insight to be gained from RC is that the prosthesis is the machine to be taught by the user to serve her needs, not the other way around.

« 1 » Current myoelectric (hereafter "myoe") control systems for upper-limb prostheses embody a small repertoire of utilitarian movements that can be executed individually upon user activation of specific muscles in the residual limb. The movement repertoire is severely limited by inadequacy of the user's interface with her prosthesis, i.e., the human–machine interface (HMI). Thus, while the mechanical hardware of modern robotic hands can nearly or completely reproduce human dexterity, prosthetic users cannot, and new control paradigms are urgently needed. Currently available HMIs are non-intuitive, and demand much more mental attention than do natural movements. Typically, prosthetic grasping is triggered by user volition for "wrist flexion" (despite the users having no functional wrist), that produces a particular muscle activation signal in the residuum. In some prostheses, several different similarly pre-programmed tasks may be activated, depending on the user's ability to learn and produce a sequence of the correct muscle activations in her residuum. The practical utility of a prosthesis thus depends on the user's ability to learn not only the right moves by her residuum, but also of her body poses, which are an important part of the motor control loop (Metzger et al. 2012). Functionality also depends upon the situation: relatively good control can be achieved under relatively fixed, static conditions; however, in situations requiring careful calibration of overall body movements (e.g., carrying an egg), it falls short. In general, most activities of daily living that involve manipulation, and certainly any tasks that require dexterity, exceed the capabilities of available HMIs. A common failure that cannot be fixed by the myoe controller is malfunction of the sensors themselves, commonly caused by sweating or dislodgement, which is not the fault of the ML. Alternative sensors of muscle activity, which are less subject to failure, have been demonstrated (Castellini et al. 2014: 22), but are not yet widely adopted. Thus, for several reasons, including some that have nothing to do with the type of ML used, many users of upper-limb prostheses abandon theirs.

« 2 » Myoe controllers, in their most primitive (typical) configuration, direct pre-programmed prosthetic actions upon receiving signals from specific muscles. Muscular activities are statistical events that can be compared against an explicitly pre-programmed value, a paradigm that is considered to be a form of machine learning (ML) (§2). Any muscle signal that exceeds a pre-set amplitude threshold produces a binary "1" input to the prosthesis. In some cases, the user can serially trigger several different prosthetic motions from a pre-programmed repertoire, by executing particular sequences of discrete motions, each of which produces distinguishable muscle contractions. The potential movement repertoire is limited by the skill and patience of the user in producing strings of supra-threshold muscle contractions, and the abilities of the HMI (computer) to compare the input with the pre-programmed threshold. One advantage of the current myoe paradigm is binarization of muscle activations, which provides relatively noise-free, unambiguous examples to serve as ML inputs; this feature, however, is at the expense of information about movement magnitude and force.

« 3 » RC theory requires any "viable" model of the world to be both utilitarian and "conceptually coherent" (Glasersfeld 2005), so we can ask whether the realist model is viable (§17) and otherwise conformable to RC principles. In terms of RC, the realist model embodies *utility*, i.e., it produces at least one useful task fairly well: grasping. With regard to being *conceptually coherent* (§17), however, it fails, since there are now better solutions to the problem. Moreover, the RC idea of "learning as a constructive activity" requires *continual* learning by the prosthesis as it interacts with its user, representing a 180-degree shift from the current myoe control paradigm. The current myoe model, as described in the first two paragraphs above, is based on a "'realist" attitude, which assumes, according to §11, continual performance as programmed, requiring "no further changes." If accurate, this attitude would represent the antithesis of RC principles, however, it may be a bit exaggerated if taken too literally. The need for periodic myoe program adjustments is widely recognized and in practice, and changes are implemented where practical, regardless of what type of controller. In the traditional sense of the word, "viable," however, we must acknowledge the practical viability of realist myoe control, because it serves many thousands of amputees.

« 4 » The RC framework, as elaborated by Markus Nowak, Claudio Castellini and Carlo Massironi, introduces a radically new, and possibly improved, prosthetic control paradigm. The first and most obvious insight from RC is that our present prosthetic model employs a strategy opposite to machine learning: instead of the prosthesis learning the proper responses to the user, *it* acts as the teacher, demanding the correct input from the user (who is the learner) for proper performance. A second insight is the potential pitfalls of a supposed *realist* attitude to statistics. Muscle activation signals are composed of Gaussian noise, generated by a large number of asynchronous motor units, which are variably active for each movement in a sequence. The application of a statistical test to such signals, using fixed decision boundaries, is bound to lead to erroneous decisions, since two identical events, such as sequential movement commands, can be statistically different in their muscular representations. Thirdly, the idea of incrementally updating the controller (incremental ML) is integral to constructivism. This process, iML, was demonstrated in the pilot studies, consisting of constant monitoring and teaching of the prosthesis by the user (§65), and appears to be viable.

« 5 » It is useful to compare the current (realist) myoe model with a hypothetical RC-framed control model. Current myoe systems treat their input as an unequivocal signal, either to be detected or rejected, according to its magnitude. RC insight recognizes that inputs to an ML system exist only as "perceptual objects" that must be organized, without knowledge of their meaning. This is an important reminder that ML inputs represent a "reality" constructed by myoe sensors and thus constitute a noisy estimate of reality, which in our case consists of muscle activations. These perceptual objects must be matched against patterns consisting of objects perceived by an imperfect system. This framework compensates for mistakes and mis-interpretations, by incorporating explicit procedures for correcting them. This ongoing positive feedback tends to promote positive emotions between user and her assistive robot.

« 6 » It is also useful to evaluate the pilot study of incorporating RC principles into prosthetic control (§38 ff). The experiments were elegantly designed and executed, but the results rather disappointing (§88). Here, I critique the experimental design from my interpretation of RC principles.

« 7 » Firstly, training the prosthesis (machine) was done by subjects performing general movements and static positioning of joints related to their hands. While the prosthesis may more easily execute these movements, they do not fit well within the RC framework. The protocol involved no purpose, and lacked motivation. Humans like to perform *tasks*, especially those that are interesting, challenging, and have *utility* (Gorsic et al. 2017). Examples of this phenomenon can be seen in previous studies wherein motor-disabled persons and amputees taught their virtual prosthesis to play and win standard games, such as pegboard (Kuttuva et al. 2005; Yungher & Craelius 2012).

« 8 » A second critique is testing *non-disabled* subjects on the use of an assistive robot. From an RC perspective, this seems *conceptually incoherent*. Persons with motor disabilities may be better teachers of assistive robots than able-bodied persons, as suggested by the two studies cited above. In a study of 12 persons with arm paresis due to brain injury playing a virtual pegboard game wearing a sensor sleeve on the affected arm, a significant improvement in speed of 15% was achieved after 30 trials with their virtual assistive robot; controls, in contrast, showed negligible improvement with their prosthesis (Yungher & Craelius 2012). Qualitatively similar results were found in a smaller study comparing virtual prosthetic teaching by controls with that of persons having upper-limb amputation (Kuttuva et al. 2005).

« 9 » A final critique relates to the anatomical differences between the residuum and the intact limb. Muscles and tendons in the residuum are radically rearranged, and any natural synergies among them are disrupted. Additionally, the typical conical shape of the residuum may be a better substrate for sensor sleeves, which can readily accommodate 32 sensors as opposed to 8 sensors typically applied to sound limbs.

« 10 » Since the experiment with subjects (§55) incorporated their vocal feedback to the prosthesis via speech recognition (SR), it is interesting to compare that technology, perhaps the oldest and most common ML system, with the current myoe ML system. There are four ways in which the systems differ radically from each other:

- myoe systems are necessarily custom-tailored to individual users, whereas SR systems are designed to be universal,
- myoe systems are trained with relatively few examples, whereas SR systems are trained by as many examples as possible,
- SR is inherently interactive, whereas current myoe systems are not, and
- SR employs statistical *predictive* methods, unlike myoe, which does this minimally.

« 11 » At the same time, myoe and SR systems share a major commonality, since both can be considered *recursive translators*. An SR translator interacts with a human who verbalizes an idea as words in one language, interprets those and translates the idea into words in another language. The myoe controller interacts with a human who expresses a desired movement by muscle activations, and interprets those and translates the desired movement as a robotic movement. It is noteworthy that speech recognition was not too long ago ridiculed as "crack-pot" by some, whose anecdote referred to the case when a computer program translated the phrase, "out of mind, out of sight" into Chinese and back to English, and it replied, "blind idiot." We see analogous sorts of anomalies occurring in myoe control, when the prosthesis "goes off the rails," but given progress along the right directions, as exemplified in this pilot study, assistive robotics technology may now be at a developmental stage resembling that of SR several years ago. When prosthetic controllers achieve the same accuracy as speech recognizers, amputees will then be able to enjoy a high degree of dexterity.

William Craelius received his PhD in Biomedical Engineering from Northwestern University, and did Postdoctoral training in Neuroscience at Stanford University. In 1999 his team at Rutgers University developed *Dextra*, the first "multi-finger" controlled prosthesis, which was used by an amputee to play a short piano piece, broadcasted on the BBC. He currently is Professor of Biomedical Engineering at Rutgers University, where he founded the first MS Degree Program in Prosthetics. He is a Distinguished Fellow of the *National Institute for Disability Research and Resources*, and Fellow of the *American Institute of Medical and Biological Engineers*. Webpage: http: //www.dextrahand.org

Received: 23 February 2018
Accepted: 28 February 2018

The EMG Properties Limit Ultimate Classification Accuracy in Machine Learning for Prosthesis Control

Richard F. *ff.* Weir
University of Colorado Denver, USA
richard.weir/at/ucdenver.edu

> Upshot • Machine learning (ML) has been applied in many forms and under many names over the years to the problem of mapping arrays of surface electromyogram (EMG) signals measured on the arm of a person with an amputation and then trying to correlate those signals to the control of multi-degree-of-freedom prosthetic arms. While being intrigued by the idea of the interactive machine learning (iML) component of the study, I am not surprised that iML did not do noticeably better than standard approaches. The issue, as demonstrated by many researchers, is not our ability to do ML but rather the fundamental problem associated with using EMG as the inputs to the ML system and the clinical issues associated with stable acquisition of those signals.

« 1 » Markus Nowak, Claudio Castellini and Carlo Massironi present a flawed argument for the problems of machine learning (ML) in EMG control of multifunctional upper-extremity (UE) prostheses. Also, they present an engineering-/science-centric view of persons with amputations and how they use their prostheses as justification for the use of advanced ML techniques, with little understanding of the *clinical* drivers for the current state of UE prostheses control.

« 2 » In §10 the authors claim that

> "machine learning tends to be used as a number-crunching black box, at which to throw as many examples as possible, hoping that it will yield a usable relationship between input data and target values. Too often, scarce attention is paid to the quality, the origin and the meaning of the examples […]"

I claim that the authors and others are guilty of exactly this. If we stopped to consider the problem a little more, we might be able to achieve a different result.

« 3 » First, we need to consider the population for whom we hope to fabricate hands and arms. Persons with trans-radial (TR) level amputations make up more than 70% of the upper-limb amputee population. For 99% of these individuals this is a unilateral involvement. Of this TR population 80% will use some sort of prostheses. But the implication of the unilateral involvement is that most people with an upper-limb amputation still have a good limb that they will use, over their prosthesis, for most tasks. This makes the barrier to acceptance for UE prostheses very high and means that anything that is perceived as heavy, uncomfortable, bothersome, or a hassle will not be used.

« 4 » Trans-humeral (TH) level abandonment rates tend to be higher because of the need for an elbow, which serves to isolate the hand from the residual limb. Here abandonment rates are at about 50%. This population makes up 10–15% of the UE population. So, when Nowak et al. say in their abstract that "[d]espite more than 40 years of academic research, myocontrol is still unsolved, with rejection rates of up to 75%," this appears to be a hyperbole used to justify their technology. I consider it not accurate as it continues to propagate the myth that the current devices are not useful.

« 5 » Second, we need to consider standard-of-care fitting practices and why we are where we are. The most common standard-of-care myoelectric (EMG control) fitting is a 2-site myoelectric prosthesis for persons with TR loss. Myoelectric systems have been widely accepted as a viable clinical option since the 1980s and in particular the TR myoelectric fitting has found success due to its cosmetic appeal, lack of suspension straps and high grip strength, but system robustness, weight, and cost are still issues (Atkins, Heard & Donovan 1996). A standard 2-site myoelectric system is typically limited to 2 degrees of freedom (DoF) with a co-contraction, or rate, used to switch between DoF because, in general, one can only get 2–3 independent *surface* EMG sites on the residual limb of a trans-radial subject before cross-talk becomes an issue (Ajiboye & Weir 2005). It is the limited number of control sites and the associated limit on the number of controllable DoF that led investigators to explore other means of acquiring and using multi-DoF control schemes such as ML. Users certainly want more DoFs, but not if it is a hassle.

« 6 » Pattern recognition (PR, which is what the field of prosthetics control calls ML) was first explored in the 1960s and 1970s (Herberts et al. 1973; Lawrence & Kadefors 1974; Taylor & Finley 1974) and reinvented in its current form in 1993 (Hudgins, Parker & Scott 1993). Since then it has undergone much development (Englehart & Hudgins 2003; Scheme & Englehart 2011; Sensinger, Lock & Kuiken 2009; Farrell & Weir 2008a, 2008b; Simon, Lock & Stubblefield 2012; Zhou et al. 2007) to get it to a point where it could transition to the clinic. PR algorithms seek to correlate *patterns* of EMG activity with a *given/desired* arm motion or hand posture. EMG signals are measured by an array of myoelectrodes on the residual limb. Other "features" in addition to amplitude are extracted and correlation between the EMG activity pattern and a desired motion is determined by *training* an ML algorithm/classifier with the extracted features for each EMG signal.

« 7 » Currently, a good ML/PR classifier can readily achieve about 90–95% accuracy. During training, patterns of muscle activity are recorded while the user holds a desired posture. Multiple trials are recorded for each posture. This has to be repeated for every desired posture. One can see how, if a large number of different postures are desired with multiple trials for each posture, training the system can get tiresome for the user quickly. Also, as currently implemented, PR is sequential in nature and users are limited to only those postures that they have trained the ML system to recognize. So far, because of this, PR techniques have met with little success outside the laboratory, so why is this so?

« 8 » Returning to the idea that the barrier to success is high for persons with unilateral amputations – the training is a bother that has limited pattern recognition adoption. People want to put on their prosthesis and go. They do not want extended training periods every time they don their arm and they do not want to then have to repeat the training throughout the day as the environment in their socket changes the electrode properties. As a result, it has taken a long time for PR systems to make it to the clinic.

« 9 » Third, we need to consider the nature of the EMG signal used for the control. The raw myoelectric (EMG) signal is a broadly Gaussian random signal whose amplitude increases with muscle contraction level. This signal needs amplification/filtering/integrating/processing to extract the RMS value for use in amplitude-based myoelectric control (Childress & Weir 2004; Parker & Scott 1985). Filtering adds a delay to the system, decreasing system responsiveness, which, if excessive, frustrates the user. Furthermore, EMG control provides no feedback – myoelectric signals are recorded on the surface of the skin and sent out to the motor and nothing deliberate comes back. Incidental feedback in the form of motor whine and socket pressures are used by skilled users. Clinical issues such as motion artifact, skin impedance changes and electrode lift-off also present challenges that must be overcome during the fitting process.

« 10 » When using EMG signals as inputs to the ML algorithms, the random noisy nature of the EMG signals presents difficulties for ML classifiers of choice. Small changes in electrode position can have a dramatic effect on the machine-learning/classification accuracy. Donning and doffing the prosthesis can alter the classification. As the number of degrees of freedom (DoF) to be controlled increases, moving from one posture to another may result in overlapping muscle activity patterns, reducing the ability of the classifier to separate the EMG patterns. In addition, extrinsic factors such as electrode movement, electrode lift-off, changes in skin impedance, or moving to positions outside of the initially trained position or using the prostheses under varying loads or in different positions can all significantly degrade classifier performance (Fougner et al. 2011). This makes it extremely difficult for PR systems to find success with users.

« 11 » Finally, we must understand that the goal of research into ML systems for UL prosthesis control is to build systems that will someday be worn by individuals with limb loss and that given what I said above in §§3f there are a host of clinical issues that will be ultimate drivers of success.

« 12 » The way CoApt, LLC, (Chicago, IL) was able to circumvent these issues and launch the first clinically successful pattern-recognition system was to use an eight-electrode system to control 1 DoF in the hand (no grip patterns) and only 3 DoF in total (hand, wrist, and elbow) (Uellendahl & Tyler 2016; Baschuk et al. 2016). This enables a user to do the "on-the-fly" training using CoApt's prosthesis guided training (Lock et al. 2011; Simon et al. 2011) system, which since users are only controlling 2–3 DoF, does not have onerously long training. In a field that has been locked into only 2 electrodes as standard of care, CoApt's approach of providing a system of 8 integrated myoelectrodes to control 2 DoF is changing how clinicians and researchers are thinking about the provision of myoelectric care.

« 13 » What we see is that it was a knowledge of the field and the population to be fitted, as well as a clinically viable way to allow training by users on the go that enabled the CoApt system to move from the laboratory to the field. The ML algorithm CoApt uses is not sophisticated, just good enough, because it is not the determinant for success. What we see in ML/PR is that by using every available technique (including fuzzy logic, linear discriminant analysis (LDA), principal component analysis, non-negative matrix factorization, self-organizing feature maps, support vector machines, random forests, cepstral constants, neural networks, and multinomial regression) a classification accuracy of about 90–95% can be achieved but not more. So, in the field of prosthetics control, the LDA classifier with a time-domain feature set or auto-regressive constants has become the standard, because it is low cost from a computational perspective, easy to implement, and is as good as anything else. Could the authors expand on the feature set as well as the actual classifier they used? There was a lack of detail on the actual classifier used and no mention of the features used to train the classifier.

« 14 » So, bottom line, the interactive aspect of the iML trial is an interesting concept in the target article, and ought to be a good thing in the long run. Talking to the subject and telling them as the training session progresses that a training movement is "good" or "bad" and then only using the "good" training datasets to build the classifier ought to bias the classifier training dataset to "good" examples. But when I read that ultimately the iML pilot study results did not show much improvement it did not surprise me. It is hard to get beyond the 90–95% classification accuracy rate, since this is most likely a consequence of the poor properties of the EMG signals used as the system inputs. We need to do something different with the EMG signals or integrate them with other types of input signals.

Richard F. ff. Weir, PhD, is Director of the Biomechatronics Development Laboratory at the University of Colorado Denver | Anschutz Medical Campus. The Laboratory's research is focused on the development of advanced prosthetic systems for individuals with limb loss. This research covers all aspects of the problem ranging from neural control and sensing; signal decode and algorithm development; mechatronic design and development; and and novel actuator technologies; to clinical deployment of these systems. We led the development of the implantable myoelectric sensor (IMES) system, which is currently undergoing first-in-human trials. Most recently, we have been exploring optogenetics to non-invasively optically interface with the peripheral nervous system to provide enhanced control and sensory feedback. Dr. Weir has been involved in the field of prosthetics research in one form or another for over 30 years.

Received: 1 March 2018
Accepted: 5 March 2018

Choosing the Right Observables

Peter Cariani
Boston University, USA
cariani/at/bu.edu

> **Upshot** • The stripped-down experimental setup may be missing important sensory proprioceptive and tactile observables that may well be crucial for designing useful, effective, and flexible general-purpose motor prosthetic devices. Because trainable machines cannot by themselves add new observables, designers must foresee which ones are needed.

« 1 » Motor substitution – the replacement of organic effector organs with artificial ones – has long been stymied by the problem of control – how to effectively control the artificial muscles using neural and/or muscle signals that are produced by the human operator. Markus Nowak, Claudio Castellini and Carlo Massironi address the problem of controlling arms and hands via muscle activation signals (myocontrol) observed via electrical sensors (upper arm electromyography (EMG)).

« 2 » Their target article focuses on the human-machine feedback loops in play when one has a human training and operating an adaptive prosthetic device. After first introducing machine-learning schemes, the discussion quickly moves away from the ultimate problem of effective motor substitution and into the stripped-down experimental setup, where an adaptive machine controller is trained to produce a small set of six alternative discrete static wrist-hand positions (§51: rest/no-action, wrist supination, extension, flexion, pronation, and hand-closing). The adaptive controller decides how to move a simulated hand given a particular goal (a target wrist-position category) and the eight-channel EMG output of the *Myo* bracelet (§42), which here is worn by normal subjects with intact upper limbs. The main focus of the target article is on the role of human-machine interactions during different stagings of model building and training phases.

« 3 » Effectively solving the problem of motor substitution will substantially enhance the lives of many people, and I think the limited experiments outlined here are well worth pursuing. Innovations in design strategies and how we think about them also have large ripple effects in other domains, such that bringing constructivist ideas to the design process have implications far beyond prosthetic devices (as I say, *all technology is prosthesis*, in that every technology that is meaningful is some amplification or augmentation of our biological, bodily and mental functionalities).

Motor control under natural vs. experimental conditions

« 4 » In order to understand the experimental setup, I found it necessary to draw schematics that depict the functional organizations of humans with intact upper limbs vs. those of the trainable machines that are considered here (Figure 1). The first order of business when evaluating a system that is designed to operate in non-virtual realms is to examine its goals (the functions it implements, whether for itself or in service of a designer's goals), what its observables (measurements, realized through sensors) and modes of action (realized through effectors) are, and how these are coordinated (how percept-action mappings are determined). These are the basic functionalities of any *purposive, percept-coordination-action system* (Cariani 1989, 2011, 2015). A system is *purposive* by virtue of embedded goals, evaluation mechanisms for assessing whether goals are attained (satisfied) or better performed, and means of directing or steering behavior to better attain goals. For each goal, the steering mechanism consists of a percept-action mapping, i.e., how the system should behave given its sensory inputs (its-current-observed-state-of-its-environment) given that current goal. Such a system has *agency* vis-à-vis that goal if it has the autonomy to pursue attainment of that goal.

« 5 » In the target article, we have two adaptive, purposive percept-coordination-action systems that interact to train each other, namely the human operator and the trainable prosthetic device. This could be seen as two problems of first-order cybernetics: how does the human best give evaluative feedback that trains the machine to recognize different muscle activation patterns (EMG signals), and how does the machine train the human to modify patterns of muscle activations such that it can better classify the signals? Provided that the two systems have enough variety in their responses and in channels that mediate their communications for mutual adaptation to be possible and functionally beneficial, we can also view this as one problem of second-order cybernetics in which the dynamics of the interactions might be crucial.

« 6 » However, whether the order of interactions ultimately matters in improving the quality of prosthetic movements may depend on whether there is room for improvement. If the human user cannot effectively learn to change muscle activation patterns that are observable via the eight channels of EMG or the trainable machine is already exploiting the limited data it has to the fullest, then not much benefit in prosthetic function may be gained from modifying sequences of model building and testing. If I understand Figure 8 correctly, the mean hand-configuration errors (nRMSE), which quantify the similarity of the hand positions produced by trainable classifier with the target hand positions, should be improving with training. However, no such trend in the error metric is seen for any of the three experimental protocols over the course of 90 trials. This could possibly be indicative of a ceiling effect – the classifier rapidly achieves its optimal performance such that further training does not help and also that the potential benefits of modifications of the training-test protocol are hidden. As the authors note (§§80f), there are also additional subjective factors, such as perceived muscle fatigue, smoother interaction, and positive impressions of prosthesis operation that are entirely relevant to patient acceptance and use that may be amenable to improvement by adjusting training and testing protocols.

« 7 » It could be the case that including additional physiological observables would permit higher optimal levels of functioning that could benefit from mutual adaptation. Choice of observables – measurements to be made – is the most important decision to be made in constructing a predictive model, and choice of feature primitives is likewise the most important decision in designing a trainable classifier. In general, choosing

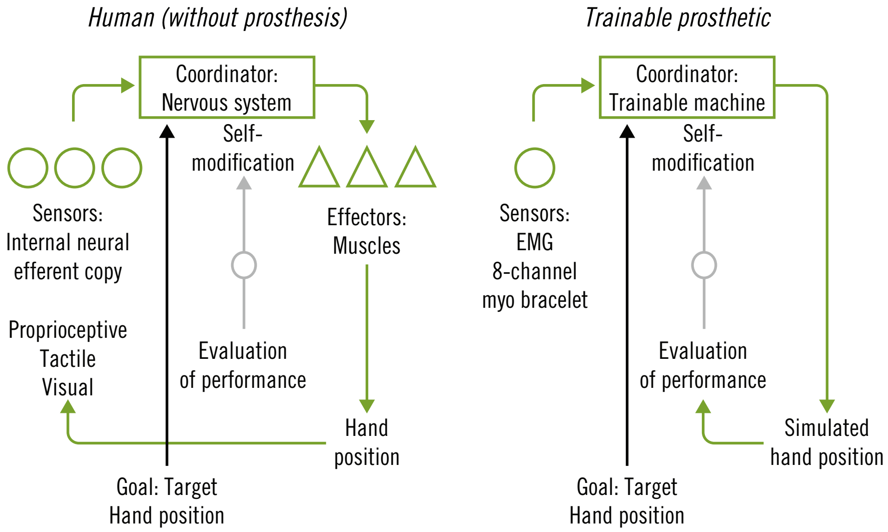

Figure 1 • Functional organization of human upper arm control (left) compared with the trainable prosthetic used in the experiments (right). Sensors that carry out measurement operations are indicated by circles, coordinations that map input sensory states to output decision states by boxes, and effectors that carry out actions by triangles. Arrows indicate causal chains of effects. Black arrow indicates goal directive (task target). Gray arrow indicates evaluative feedback about the efficacy of the last action (measurement of performance) and the operation of modifying percept-action mappings. Note that the experimental setup is highly impoverished in terms of sensory observables for feedback control and possible action states (many vs 6 hand positions).

what measurements to make, however, is an ill-defined, unformalizable, domain-specific problem for which there are no effective procedures other than building new measuring devices and trying them out. In practice, designers think of as many possible relevant observables as they can and then eliminate those feature primitives that yield no benefit. Biological systems have solved this fundamental problem by evolving new sensory receptors, modes of neural coordination, new effectors, and new possibilities for action. Long ago (Cariani 1989), I proposed a class of biologically inspired devices that would adaptively build their own hardware including their sensors, coordinative parts, and effectors such that they would construct their own primitives, thereby solving the problem, in principle at least, of finding the right ones.

« 8 » A machine-learning classifier can only be as good as its feature primitives and a machine-learning controller can only be as good as its set of possible actions. A trainable machine does not have the means of creating new feature or action primitives, such that it is prisoner to the sets of sensors and effectors its designer chooses for it. Without adequate variety in these domains, such systems can learn up to a point, but they will be ultimately constrained by the limitations of their pre-specified sets of features and actions.

« 9 » When tackling a problem in bionics, and especially when troubleshooting why artificial prostheses do not work as well as their biological counterparts, it is useful to first compare the functional organization and operational structure (physiology) of the two systems. My concerns as a physiologist, systems scientist, and cybernetician would be what the experimental setup is leaving out in terms of observables, actions, and feedbacks.

« 10 » In the normal, intact biological case (Figure 1, left), a human (or animal) has a number of sensory channels that provide critical feedback for movement and positioning of limbs. Perhaps most importantly, humans and animals have proprioceptive feedback that provides information about the limb positions and muscle stretch. I once worked on the problem of spinal cord regeneration (Wang et al. 2008), which involved facilitating the regrowth of neural connections between proprioceptive afferents and their associated sensory pathways in the spinal cord. Despite intact muscles and motor neurons, rats deprived of neural signals in forelimb afferents completely lose the use of their forelimbs, but once these connections are restored, functions also return. Humans who have lost their proprioceptive afferents through disease have great difficulty executing movements such as walking, and only through concerted, sustained attentional effort can they learn to use visual feedback to guide their limbs. In many common situations, we can also benefit from tactile feedback. There are also thought to be neural efferent copy signals that provide the brain with copies of the command signals that are activating muscles. As far as I can tell, the present prosthetic setup involves only eight channels of EMG data that would be analogous to using motor command signals or their efferent copies. There is thus visual feedback, which may be adequate for simple, static hand positions, but there is no proprioceptive or tactile feedback, which might be necessary for flexible movements or grasps. If I were involved in the problem of designing a flexible, general-purpose prosthetic device, I would look first to incorporating proprioceptive and tactile feedback signals from artificial hands and arms into prosthetic controllers (Ciancio et al. 2016). Incorporating whole new classes of observables is, of course, a much more formidable task for an experimenter, so it is entirely understandable why the experimental setup reported here would not (yet) include them.

Understanding the experiments

« 11 » The experimental setup in the target article is complicated to the uninitiated and is confusing to sort out, especially if one is more focused on the motor substitution problem than on human-machine interactions. The nature, adequacy, robustness, and informational content of the eight channels of EMG data and the effects of alternative machine-learning algorithms are never spelled out in detail: Are they operating on time-series EMG data? How similar are signals from amputees and non-amputees? How many independent dimensions or distinctions can they convey? Does performance using the *Myo* bracelet data depend critically on the type of machine-learning algorithm that is used? However, these considerations may be critical for interpreting these results and for solving the more general problem of designing prosthetic devices that are going to be practically useful to their users. There are many detailed technical questions that can be asked concerning the transferability of findings from the experimental setup to the practical situations of amputees who will use such devices.

« 12 » The multiple means of evaluation and sequencing of performance and training trials further complicate understanding. Multiple means of evaluative feedback included human subjects seeing their own hands or seeing simulated hands on a computer screen and giving good/bad judgements vs. machine-based distance geometry metrics (nRMSE) of hand configuration similarity. However, was the nRMSE metric used directly in some cases to train the machine or was the training feedback always from the human operator (good/bad) and the nRMSE simply used as a non-subjective (intersubjectively verifiable) measure of the accuracy of the system? (Q1)

« 13 » In §22 *interactive machine learning* (iML) is contrasted with good-old-fashioned *machine learning* (ML) in that a human operator, rather than some completely artificial evaluation process, provides physiological observables (eight channels of upper arm EMG) and feedback to the trainable classifier/controller (as seen in Figure 2 of the target article). In this case, it seems that by far the most important role for the humans in this setup is to provide the EMG patterns (via the *Myo* bracelet cuff on the operator's right arm) that will be classified by the trainable machine to generate simulated hand positions. In this situation we have two adaptive systems, the operator, who may be learning to adjust muscle actions in order to steer the trainable machine to produce more appropriate hand positions, and the trainable machine, which is simultaneously updating its classification of the EMG data based on the evaluative feedback it receives from the user. Given that the evaluations of six simulated hand positions by human trainers are binary decisions (good/bad) concerning the similarity of target and produced (the 3D simulated and visually rendered and displayed images in the figure) hand positions, could the evaluative feedback have been easily replaced by the nRMSE distance-geometry metric? (Q2)

« 14 » On the other hand, perhaps the main rationale for making the human operator give explicit feedback is to focus the operator's attention on the task and to provide greater reward when desired actions are obtained. From the increasing separations of the EMG patterns depicted in Figure 9, the training of the human operator did appear to significantly modify the EMG signals that are picked up from the Myo bracelet. One would think that this greater separation of input signals would cause the system to make fewer confusions that produce classification errors. However, as the authors remark (§88), the effects of training on performance appear to be minimal. The time course of the position-error metric (mean nRMSE) in Figure 8 shows no obvious improvement with training (trial 1 to trial 90). The hand-position separations at the first and last trials for best and worst performers in Figure 10 similarly show little obvious improvement.

« 15 » In summary, it appears that most of the effectiveness of the prosthetic classifier-controller is due to its ability to separate the EMG patterns without the benefit of human evaluative feedback. In these experiments, the human operator is critical in the generation of the EMG patterns but not essential for giving the trainable machine feedback. Nevertheless, I agree with the authors that interactive machine learning, which gives the user control over when and under what circumstances to update the input-output function of the machine, is nevertheless likely to be a promising strategy for prosthetic design.

Realist vs. constructivist approaches to design

« 16 » Some sections of the target article (§§34–37) discuss the effects of epistemology on design. We humans are all self-modifying, self-constructing systems, whereas most of our artificial systems are not. Machine-learning systems, to the extent that they do self-modify and self-construct, do so within much more constrained avenues of possible modifications than we humans and animals do. On the other hand, we can be clearer about what is going on within the trainable machine than we can about what is going on in the minds of its designers. I tend to prefer to talk about the capabilities and limitations of self-constructing vs. non-constructing systems rather than design paradigms, which reside in the heads of human designers, as important as these can be.

« 17 » It is possible to talk in terms of (realist) ontology-based design (§§3–12), where a physical or virtual world with a description that is meant to be complete is first postulated, and then a system within that world is specified to have some sort of effective behavior that fulfills some function. Partial, often statistical, observations of this "god's eye" universe by limited actors are then overlaid onto this postulated world. Three basic types of realism are physical realism, mathematical realism (platonic idealism), and logical realism (propositional objectivism). In realism, an objective world of one sort or another is held to exist independently of any observers, such that realists find it meaningful to talk in terms of "true" knowledge of the details of this world even apart from how one would observe them.

« 18 » An alternative to realism is to take an epistemological approach in which one adopts the perspective of a limited observer-actor. The observer-actor strives to achieve particular ends, such as predicting future events or bringing about particular desirable events), given limited means of observing the world and acting on it. The observer-actor, without needing an explicit ontology or access to any unobserved world-states, forms a (non-referentialist) model for effective prediction and action that then guides expectations and actions. This model is based entirely on tangible observations and evaluations.

« 19 » Although I have no evidence for this assertion, I would think that realist designers would be more inclined to try to design devices directly, from physical principles, whereas constructivists would be inclined towards making devices adaptive and semi-autonomous, such that they construct their own effective means for anticipation and action.

Peter Cariani is an auditory neuroscientist, theorist, and teacher. His main research contributions have involved epistemology of evolutionary robotics, biosemiotics, temporal coding of pitch, and neural timing nets. He is a Senior Research Scientist in the Hearing Research Center at Boston University; Associate Professor, Part-Time at Berklee College of Music; and Clinical Instructor, Part-Time in Otology and Laryngology at Harvard Medical School. He currently teaches courses related to the neuropsychology of music, the performing arts, and consciousness.

Received: 26 February 2018
Accepted: 2 March 2018

Diving Deeply into Radical Constructivism

Marco C. Bettoni
Steinbeis Consulting Centre Knowledge Management and Collaboration, Switzerland
marco.bettoni/at/weknow.ch

> Upshot • Applying radical constructivism to machine learning is a challenge that requires us to dive very deeply into its theory of knowing and learning. We need to clarify its fundamental concepts, if possible, in operational terms. This commentary aims at outlining how this kind of clarification could look in the case of 3 such concepts: (a) the construction of experiential reality; (b) learning as a constructive activity; (c) the viability of conceptual structures.

Introduction

« 1 » One of the major experiences that led Ernst von Glasersfeld to adopt a constructivist way of thinking was his pioneering work in artificial intelligence, starting in 1959 with the machine translation project at the Centre for Cybernetics at the University of Milan, created and directed by Silvio Ceccato (Glasersfeld 1995: 7). Thus, I am rather enthusiastic about the idea of applying von Glasersfeld's theory of knowing and learning to machine learning (ML) and hope that my comments will support the efforts of Markus Nowak, Claudio Castellini and Carlo Massironi in continuing this promising line of research.

« 2 » In the field of assistive robotics for limb amputees, electromyographic signals generated by muscle activity in the remaining upper limb are used as input data for a machine learning (ML) system; the system should then produce control commands for a prosthetic arm/hand accordingly in order to let it perform the desired action (§33).

« 3 » Unfortunately, this so-called upper-limb myocontrol, after 40 years of research, is still failing (§34) with rejection rates of up to 75%. As a means of improving such systems (smart prosthetic arm/hand control systems), the authors of the target article suggest (§32) developing traditional ML to form an *interactive* machine learning (iML), which allows for system updates whenever its actions are unsatisfactory (§§21f). But this poses new problems, which require appropriate conceptual tools, in particular, a coherent conceptual framework about interactivity. This is where the authors anticipate that radical constructivism (RC) could help (§23), especially through its concepts of experiential reality (§15), learning as a constructive activity (§16), viability (§17), assimilation, scheme theory, accommodation and equilibration (§30).

« 4 » The application of RC to iML – so called *RC-framed iML* – for the task of upper-limb prosthesis is expected to provide useful insight into how to design the interactive prosthesis of the future (§89). The authors are convinced that their approach has the potential to improve human-robot interaction. Thus, they propose to shift the attitude towards ML from a realist to a radical constructivist attitude, as defined by von Glasersfeld (§13). They see their draft of an RC-framed iML presented in the target article, as an attempt at opening a discussion between the RC community and the ML community (§26).

« 5 » Applying RC to ML requires us to dive very deeply into radical constructivism and clarify its fundamental concepts. So, I will look at three fundamental concepts used in what the target article calls a "tentative framework" (§26) about "interactivity" (§23) and will try to dive deeper into them.

A | The construction of experiential reality

« 6 » Nowak et al. mention this concept and quote von Glasersfeld (1995: 58f) as a reference where it appears as a section title. I will highlight the essential parts of this section by not only repeating the same formulation but also by reformulating and extending them in my own terms.

« 7 » Humans, as infants and later as adults, can construct the reality they experience for themselves. As infants, humans develop the basic concepts that constitute the essential structure of their individual experiential reality, without needing a specific physical structure to exist in its own right as a *corresponding* structure.

« 8 » For example, let us look at the development of the notion of the "object" in a human infant. In phase 1, the infant coordinates sensory signals recurrently available at the same time in its sensory field (the "locus" of raw material that Immanuel Kant called "the manifold") and establishes by that many different object concepts; these object concepts are like operational routines for constructing the formerly constructed objects of interest again at a later point (a ball, a face, a cat, etc.) whenever suitable sensory components are available. The notion of "object" in general, then, is whatever the mind constructs *as common* to all these routines (a kind of abstract, generalised, operational routine) due to a principle of efficiency, implemented like in perception by means of "preferred paths" or "sequence patterns" (de Bono 1991: 81f; de Bono 1992: 10f). Later, in phase 2, the infant becomes able to run through such operational routines even when no suitable sensory components are available in its sensory field; in this case, the infant executes a conceptual coordination of a previously constructed object; it produces a *re-presentation* (written with the hyphen as a reminder that this term means a repetition, a replay, a re-construction from memory, of a past experience, not a picture of something in a mind-independent world).

« 9 » Thus, I do not agree with Nowak et al. when they say that the agent tries to "*organize perceptual objects*" (§15). Rather, I would avoid both "perceptual" and "objects" and say that the agent "organises a sensory field," conceived as the raw material that Kant called "the manifold," in which there are no objects unless we construct them. And when we have constructed them, I would not assign them to the sensory field but rather to our experiential reality, and there to a process that operates at a higher operational level. It is similar to looking at the skies on a clear night: you can only see an ordered pattern of stars, even a constellation, if you organise the single stars (the signals in your sensory field) by selecting some and connecting them, thus *constructing the pattern* in your mind rather than perceiving it (Glasersfeld 1999: 12; Bettoni & Eggs 2010: 133).

« 10 » Moreover, the essence of a "very radical-constructivist concept" here is not dealing with "'perceptual' data" (§15) but that the "object" as a generic concept, as a conceptual structure (and later many others), is constructed by organising a sensory manifold in many different ways and later by abstracting what is common to these previously constructed conceptual structures.

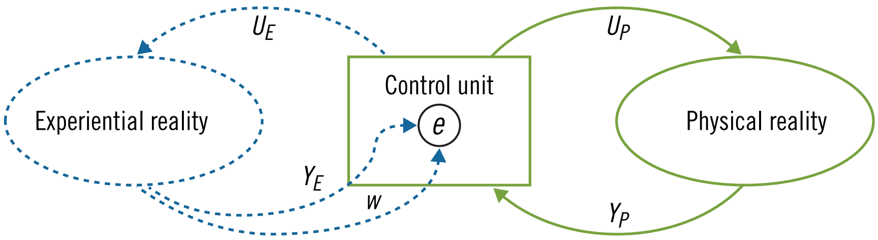

Figure 1 • The cybernetic model of viability: A coupled system of two processes controlled by one control unit. Abbreviations: Y = controlled variable; w = set point variable; e = control deviation; U = manipulated variable; index E = experiential reality; index P = physical reality.

B | Learning as a constructive activity

« 11 » This concept used in the target article (§16) references an early article by von Glasersfeld (1983) of the same title. But I would not say that this early article presents "*matching 'perceptual' patterns*" as a foundation of RC. Since the fundamental epistemological principle of RC is "fit" not "match" ("viability" not "correspondence"), I would suggest avoiding the use of "match" altogether, even when it refers to sensory patterns or conceptual structures and not to pictures of the physical world.

« 12 » An elementary form of learning requires two components (Glasersfeld 1995: 152f):

- something like a memory,
- the ability to compare two signals, a present one and a goal-signal that constitutes a reference value.

Once these requirements are met, the preconditions of inductive learning are satisfied. In the event of a perturbation, all that is further needed for this elementary form of learning to occur is a rule or principle that leads the system to repeat actions that were recorded as successful in its past experience (see also de Bono 1991: 42f), thus reducing or eliminating this kind of new perturbation.

« 13 » Although the interactions the subject has had with the world shape what will be the result of new interactions (§18), the previous knowledge that they provide is not enough for the re-cognition of a certain situation (§19; Glasersfeld 1995: 65). In fact, the sensory field provides vastly more signals than those needed for its segmentation. The organism must therefore always actively select which signals to use in order to construct either a known or a new pattern that will trigger a particular scheme, so that the pattern can be assimilated. How can the agent do this active selection? I agree with von Glasersfeld (1995: 78f) that Ceccato's idea of an attentional system (Ceccato 1964) that produces successive pulses of attention and has the ability to form combinatorial patterns of attentional moments, can provide a model of how the mind actively selects signals in the sensory field. These pulses of attention, which I have called "attentional quanta" (Bettoni 2018), also constitute the operational structure of abstract concepts (Glasersfeld 1995: 167f). Could Ceccato's attentional system also be implemented in the ML system for enabling it to do the needed active selection?

« 14 » Whenever a scheme is activated and the triggered activity does not yield the expected result, the discrepancy between expectation (reference value) and the experienced result creates a perturbation in the system. This perturbation is equivalent to a variation of the input into a controller unit of a control loop with negative feedback (cybernetics, control engineering). It is a novel kind of perturbation; it is not associated with a specific sensory pattern or with a specific scheme and may lead to an accommodation, an adjustment of the scheme or the formation of a new one. In this way, assimilation and accommodation enable an agent to learn.

C | The viability of conceptual structures

« 15 » I agree that we cannot "build a *real* model of this world" (§17) but I disagree with saying that we can "build a viable representation of it" because, again, our conceptual structures cannot be said to represent a real mind-independent world. They merely fit with our own experience and they are viable as means for consistently organising our experience (Glasersfeld 1983).

« 16 » In order to dive deeper into the concept of viability, I suggest making use of the language of cybernetics and control engineering. This allows us to illustrate the concept of viability by means of a *system model* (see Figure 1) where we have *one* control unit that *controls two* process units; it is a very peculiar architecture of a coupled control system with two fundamentally different processes and hence two fundamentally different, but coupled, control loops.

The control loop of physical reality

« 17 » On the right-hand side of the diagram, I differentiate between reality as a physical controlled system or process, the person as its controller and two interactions between these two units: the physical effect of a person on reality (controller output, manipulated variable U_P) and the physical effect of this reality[1] on a person (controller input, controlled variable Y_P).

« 18 » The controlled variable Y_P only affects the person in the form of a manifold (Kant 1966: B 102; Glasersfeld 1995: 40f), i.e., in an unstructured manner. In the diagram, this is indicated by the fact that the arrow ends at the periphery of the control unit and does not penetrate into the inner circle, like the other variables.

The control loop of experiential reality

« 19 » On the left-hand side of the diagram, I differentiate between the experiential world as the entirety of the experiences acquired by a person (her knowledge base) and the person as the controller in the form of a separate unit; this separation is purely heuristic in nature for illustrative purposes. In this model, I also assign to the experiential world the role of a controlled system, but a conceptual (conceptually constructed) rather than a physical controlled system.

« 20 » There are three interactions between these two units here: the conceptual effect of a person's control unit on her experiential world (manipulated variable U_E) and two conceptual effects of the experiential world on the person's control unit. The set point variable w corresponds to the goals, intentions and expectations. The controlled variable Y_E is somewhat more complicated: a person takes the controlled variable Y_P, transforms it into thought content (manipulated variable U_E), seeks to integrate this into her experiential world (assimilation, accommodation etc.) and ends up with the controlled variable Y_E.

« 21 » The control deviation e is formed from a comparison between the set point variable w and the controlled variable Y_E; this produces a binary variable e, which provides information as to whether or not there are any obstacles in the way of pursuing the goals, i.e., whether or not the current state can be deemed *viable*. If the manipulated variable U_P has led to a solution or generates any concepts that are either compatible with existing conceptual structures (lack of contradictions) or in harmony with conceptual structures that others regard as viable, then in the control unit we will obtain $e = 0$, i.e., the current state will be considered viable and will be reinforced.

1| By "physical reality" I mean the world of constraints in which organisms live (Glasersfeld 1983) and by "physical effect" I mean variations in the sensory field due to those constraints.

Conclusion

« 22 » Diving deeper into concepts such as *the construction of experiential reality* and *learning as constructive activity* ensures that the development of an RC-framed iML will be more consistent with RC. Furthermore, due to the central role assigned to interactivity by an iML approach, the double-loop model of viability presented here could become the starting point or foundation for developing the missing "coherent conceptual framework about interactivity" that ML needs (§23). Here the model deals with a *human-world* interaction, where the human is the active agent and the world provides constraints. In ML the roles are swapped: we have to model an *ML-human* interaction, where the ML system is the active agent and the constraints are provided by the human (§30).

Marco C. Bettoni, Prof. emer. in Knowledge Technologies, is currently starting his own consulting business, in collaboration with the Steinbeis Network, focusing on knowledge management, communities of practice and e-collaboration. From 2005 to 2017, he was Director of Research & Consulting at the Swiss distance university Fernfachhochschule Schweiz. From 1977 to 2005, he worked as a researcher, engineer and lecturer with industrial and academic organisations in machine design, engineering education, IT development, artificial intelligence (knowledge engineering) and knowledge management. Homepage: http://www.weknow.ch

Received: 23 February 2018
Accepted: 27 February 2018

A Sociocultural Perspective for Learning Loops

Marco Guicciardi
Università degli Studi di Cagliari, Italy
marco.guicciardi/at/unica.it

> Upshot • I point out that from a sociocultural perspective, repeated experiential interaction loops are not enough for constructing new context-dependent knowledge: the loops must be grounded in specific social practices, which are either culturally or historically situated. Also, to tightly connect human user and interactive machine-learning system, triple-loop learning needs to be used as well as criteria for validating an expectation's confirmation.

« 1 » The improved interaction between users and learning systems in interactive machine learning (iML) needs a better understanding of how end-user involvement impacts the learning process (Amershi et al. 2014). To contribute to the discussion that Markus Nowak, Claudio Castellini and Carlo Massironi have opened, I want to highlight some properties of this interaction.

« 2 » Gregory Bateson (1979: 78) pointed out that one cannot hear the sound of one hand clapping. Likewise, the contributions of the human and the iML system to solving these problems cannot be decoupled. Thus, in iML we have to put the "human into the loop" (Holzinger 2016) to enable what neither a human nor a computer could do on their own.

« 3 » A conventional machine-learning (ML) system can be instructed with ever more examples when learning a stationary process (§12). Human behavior, however, is non-stationary (§29) and biomedical data sets are full of uncertainty and incompleteness (e.g., missing data, noisy data, etc.), which makes the application of conventional ML difficult or even impossible (Holzinger 2016).

« 4 » Since human and iML system are tightly coupled, some form of reflexivity is required to take into account the relationship that includes both elements as a part of it. As Erving Goffman (1974: 85) states, "a reflexive element must necessarily be present in any participant's clearheaded view of events; a correct view of a scene must include the viewing of it as part of it."

« 5 » What kind of relationship do we have to take into account if every description implies an observer who describes it (Foerster 1981: 258)? Obviously, reflexivity is not confined only to our observed interaction between human and iML system, but extends to a different level that also includes us as knowing subjects. For example, in §47ff the authors build a psychological context of interaction rules and reciprocal roles for the human – reassuring that the prosthesis learning ability was "only in its infancy" – which also involves readers who identify with this parental role.

« 6 » We can find another example of reflexivity in §37, where the authors claim that iML, based on radical constructivism, is superior to conventional realist ML. But who is the knowing subject in the experiment described afterwards (§§38ff): the iML system (here the learner) or the human (here the user) who adopts a non-realist theory of knowledge? Or both? The interaction loops increase the opportunities for users to impact the learner and, in turn, for the learner to impact the users (Amershi et al. 2014).

« 7 » Can we assume that the human, while providing feedback to the iML system, is developing a better awareness of her knowledge constructs? Does the iML system build on cultural knowledge thanks to the feedback provided by the human? Can the human and iML system create, together, new knowledge outside a social context where distances, shapes and sizes are culturally defined?

« 8 » An iML system can be conceived as a constructivist system that generates a certain kind of knowledge through experiential interaction loops (Sarkar 2016). This is how the prosthetic hand learns movements, i.e., by means of acquiring correct examples and feedback (§57ff). However, I claim that repeated experiential interaction loops are not enough for constructing new knowledge: the loops must be grounded in specific social practices, which are either culturally or historically situated. These practices are governed by constraints, in which people engage with "objects" or other constructed entities, understood in terms of apparently independent, decontextualized properties.

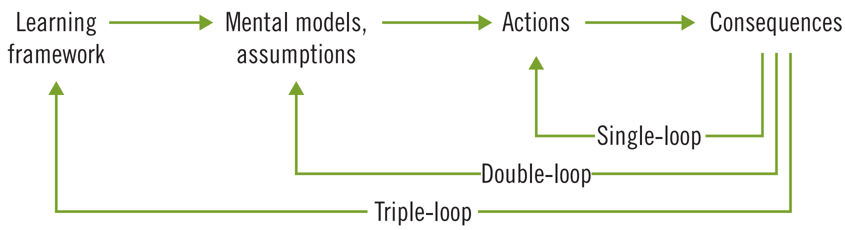

Figure 1 • Learning loops (Modified from "Modes of organizational learning" by Soren Eilertsen and Kellan London, https://le2oa.wikispaces.com/file/view/Modes+of+Organisational+Learning.pdf)

As Peter Berger and Thomas Luckmann (1967: 61) observed, humans are capable of producing a world that they then experience as something other than a human product (see also Packer & Goicoechea 2000).

« 9 » So, assuming a sociocultural perspective, activities are characterized as practices of a community. There are many different ways of moving a hand, because, as Goffman observed:

> "[…] while the substratum of a gesture derives from the maker's body, the form of the gesture can be intimately determined by the microecological orbit in which the speaker finds himself. To describe the gesture, let alone uncover its meaning, we might then have to introduce the human and material setting in which the gesture is made." (Goffman 1964: 133)

« 10 » In order to execute an action such as grasping a cup of tea or repairing a bicycle, in addition to movements and validated procedures, does an iML have to learn something about the "frame" (Goffman 1974), i.e., the human and material setting in which the action has to be executed? (Q1). For the experiment described in §§38ff this means the prosthetic hand's activity requires a rich context where meaning can be negotiated, and understanding can emerge and evolve (Sarkar 2016).

« 11 » In order to establish a tight coupling between a human and her iML system, we need triple-loop learning that is able to transcend single- and double-loop learning:

- Single-loop learning occurs when the system learns new skills and capabilities through incremental improvement: the system assimilates the information that it can already recognize. Errors are detected and corrected by a human agent, who acts without perturbing the system (see also, in §53, the evaluation phase of experiment 0).
- Double-loop learning is reflective and occurs when errors are detected and corrected, and expectations, and/or assumptions are called into question and challenged. As Nowak et al. report in §59ff (Experiments 1 and 2), when dealing with complex, non-programmable issues, the iML system was perturbed. The concept of perturbation refers to a stimulus that does not conform, or gently subverts, the expectations and mental model of the users, forcing them to construct new knowledge in order to accommodate this experience (Sarkar 2016). Error detection still occurs, but the iML system is required to change its assumptions and mental model to try to understand the "connecting structure" (Bateson 1979) that helps to detect these errors.
- Triple-loop learning involves a learning framework where "the subject learns the context of the action and how these actions are connected to the world" (Lutterer 2012). Here we must include the context, because movements and mental processes are formed in and through participation in specific social practices, which can be both culturally and historically situated (Packer & Goicoechea 2000). Learning to move a hand is also, and always, a learning of context (Bateson 1972: 293): activity is dialectically constituted in relation to the setting (Lave 1988: 151).

« 12 » Figure 1 shows how the relationship between the three loops of learning can be depicted. Each successive loop extends beyond the boundary of, and includes, the previous loop.

« 13 » Coming back to the experiments in the target article, positive feedback or confirming a prediction strengthens the experiential reality that the human and iML system are constructing together. In §52, during model building, all the instantiations are supposed to be "good" signals. In other words, the human confirms the expectations of the system, which was building its own "reality." Furthermore, during model testing, whenever a particular prediction concerning an action or reaction of the other turns out to be corroborated by what the other does, this strengthens, in a different loop, the experiential reality and the mental models that both are constructing together.

« 14 » Since prediction is different from explanation, in §59 when the system received negative feedback, I claim that the iML system was perturbed in a twofold manner: because its expectations did not fit and because its assumptions were not confirmed.

« 15 » Likewise, by means of perturbation, the iML system was stimulated by the human to construct new knowledge, for example, some criteria for validating an expectation confirmation, or using our previous terminology, both are construing concurrently a new shared learning framework.

Marco Guicciardi is a psychologist and psychotherapist. He holds a degree in psychology from the University of Padua, Italy. Since 2000 he has been Associate Professor in Psychometric at the University of Cagliari, Sardinia. His research and publications focus on health, sport and exercise psychology. He specializes in the construction and validation of psychological instruments for enhancing well-being. Nowadays, in collaboration with engineers and IT, he is studying prototypes and mobile devices for proactive health, targeted at people with non-communicable diseases as Diabetes, Parkinson's, Alzheimer's, etc.

Received: 20 February 2018
Accepted: 5 March 2018

Are Our Limbs Agents that Need to Estimate Our Intentions?

Martin Flament Fultot
University of Connecticut, USA
martin.flament_fultot/at/uconn.edu

> Upshot • I argue that the authors miss an important distinction between realism and representationalism. Because of this, their diagnosis of the current state of machine learning is valid, but for the wrong reasons. As a consequence, their approach to upper limb prosthetics may not be a step in the right direction.

« 1 » The target article constitutes a positive and very welcome contribution to a problem that has been plaguing human-machine interaction since its inception, namely prosthetics. Although Markus Nowak, Claudio Castellini and Carlo Massironi introduce their reflection as being about machine learning (ML) in general, the particular case of prosthetics is of such significance that the authors' ideas about the latter deserve as much scrutiny as their general concern with machine learning. This commentary will thus focus both on the authors' contention with what they take to be the realist stance towards ML and on the particular study they chose.

« 2 » Nowak et al. start by noting in §10 that ML tends to be used currently as a "number-crunching black box" the function of which is simply to yield useful mappings from input to output according to the designers' interests. I agree with the authors' lament that, very often, the meaning of the mappings and the processes going on in the machine are opaque to the designers, who do not seem to care. A quick and shallow rebuttal to this could be: "So what? There is nothing wrong *a priori* with having a completely instrumental attitude towards a particular computational tool." But the authors go further and they argue that whenever the model fails *in practice* it is indeed because of that theoretical attitude just mentioned, in other words, ML models are being limited because their designers are not paying attention to deeper conceptual issues. Examples of such limitation are a model's failure to produce the expected output or, as in the case of prosthetics, the dramatic and systematic failure to produce a satisfactory coupling between human and machine, as exemplified by the painful figure of 75% rejection rate the authors rightly mention. This part of their assessment and critique seems accurate and there does seem to be a connection between the practical shortcomings of ML and its conceptual foundations. However, the rest of the authors' diagnosis and subsequent suggestion of a solution may not be so accurate.

« 3 » The main contention here is that Nowak et al. conflate two different things in their critique, to wit, realism and representationalism. To be fair, they are not the first ones to do this conflation, which can be traced back at least to Francisco Varela, Evan Thompson and Eleanor Rosch (1991) and perhaps even as far back as Edmund Husserl's phenomenology. In a nutshell, the argument of Nowak et al. is that ML fails because it is designed to naively attempt to build a statistical model of an external world, but for any modeling of "reality" to be accurate, the sample input – the examples to which the model is exposed – needs to be exceedingly large. We can already notice here that any reasoning that reaches this conclusion is faced with a choice point. We can either blame the realist attitude of believing that there is an external world that the model needs to reflect, *or* we can blame the very attempt to *model* such a reality. Perhaps, and this option is rather ignored by the authors, there *is* a "naive" external reality, but the right approach to learning and knowledge in artificial intelligence, at least if the goal is to approach human performance or to make interaction with humans possible, is not that of trying to *represent* the world through a model, but rather to *fit* the world. Just like the woodpecker's beak does not represent the tree – the beak is definitely not a *model* of the tree, yet it is a perfect *complement* to the tree for the purposes of the woodpecker (e.g., drilling a hole and catching termites) – ML systems could benefit from not attempting to represent their targets but to fit them in some meaningful way. The problem, in short, is not trying to represent *the world*, but rather trying *to represent* the world.

« 4 » Interestingly, Nowak et al. seem to have a tacit understanding of this problem, but I suggest that, because they fail to discern between realism and representationalism, they aim their criticism at the wrong target. For instance, following Ernst von Glasersfeld they claim §17 that "the value of an idea of the world is measured in terms of fitness to achieve a specific goal […], not in terms of the correspondence between the idea and a mind-independent reality." But we can readily see in this statement that if the value of the idea resides in its ability to contribute to the achievement of a specific goal, then the content of the idea itself must be about achieving that goal. *That*, I suggest, should be the reason why the authors contrast fitness with "correspondence" and *not* the belief that there is indeed a mind-independent reality. Again, the woodpecker is successful in drilling a hole in the tree and catching insects for its nourishment, not because it refuses to treat the tree as "bird-independent," but because the beak and the muscular forces it applies to the tree adequately fit the latter's material properties for the purpose of drilling a hole in it. Bird and beak are in *direct* contact with the tree, they are both complementary and fit each other as mathematical duals (Gibson 1979; Shaw, Kugler & Kinsella-Shaw 1990). It is not the tree, or the external world for that matter, that needs to go, but mediational states between it and the subject (or learning machine).

« 5 » Furthermore, if we give up realism, it is virtually impossible to make sense of what the meaning of constructivist "perturbations" to the system are. If perturbations are mind-dependent and do not come from an external reality, then what is the system adapting to? But a deeper question actually addresses the authors' own concern about the meaning of what ML models do and the origin of that meaning. Perturbations must have an independent origin at least partially if they are to be meaningful to the system and effective in driving it towards improved performance. But according to Nowak et al., following their understanding of radical constructivist theory, "[*a*]*ll* these processes" must be subjective and internal – including the perturbation §19 (emphasis original). This idea is not, however, that an idealist mind is generating ideal perturbations to its own ideal perception. The issues with such forms of idealism are well known. We must rather interpret that wherever the perturbations come from, they are being shaped, interpreted, idealized somehow by the subjective agent, and they only make sense as perturbations to the agent from that subjective perspective, product of her own making (or construction). Internal conceptual schemata – the subject's pre-knowledge – are the usual posit since Kant, although they have not been without detractors (see, e.g., Donald Davidson's well-known 1974 paper). But then again, even these schemata must have an origin and we cannot posit more regressing subjects and their own schemata to account for them. It is revealing that the ethologist does not have this problem. Bird's beak and tree form a closed system, they evolved together and interact straightforwardly perturbing each other. No pre-knowledge or schemata are needed, Jakob von Uexküll's seemingly constructivist concept of Umwelt notwithstanding (Uexküll 2010).

« 6 » Thus, it seems that the authors' move towards radical constructivism as a solution to the current state in ML is a right step, although in the wrong direction. It is a move towards more reliance on representation as intermediate subjective constructs, and that could be precisely what is crippling progress in statistics-based ML. In the following I will address the point that, as a consequence of the authors' failure to identify representation as the origin of the issues faced by ML, their proposal for prosthetics misses the target too.

« 7 » It is very surprising to find no mention at all of embodiment in the target article. Yet upper-limb prosthetics constitutes a proverbial problem of embodiment. The challenge is to make an external object part of the patient's *body*, as the lost limb used to be. But notice that our healthy limbs are the exact opposite of an autonomous subjective agent trying to construct an internal model of ourselves where all the interactions with us are "user-independent." Such an idea is actually strikingly counter-intuitive. It is one thing to acknowledge that, because of the limitations of prosthetics that need to be coupled to a body *ex novo*, unlike our limbs, which grew with us and have interacted with us since our fetal stage, one needs to adapt the prosthetic limb to the body in a very short time and thus some form of ML, interactive or otherwise seems necessary as a practical necessity. It is another thing, however, to approach this problem, which is actually an unfortunate contingency, by making the disconnect between limb and body even deeper, yet that is precisely what the radical constructivist approach to prosthetic adaptation appears to imply. The ideal goal would be to be able to *grow* a new limb, as salamanders do, and let the interactions between neural, muscular and bony tissues adapt to one another during the growth process. The end result would be an *embodied* limb, one that is part of the subject, directly coupled to all the other limbs, nervous cells, etc., and certainly not anything that resembles a *separate* agent that interacts with us through intermediate mental schemata.

« 8 » Moreover, some of the negative consequences of Nowak et al.'s second experiment, namely the patient's painful fatigue, are a cruel reminder that the perturbations a subject needs to deal with are quite "real," for lack of a better term. The prosthetic is a massive body, and the earth is pulling on it through gravity. These are external constraints that the learning process taking place on the side of the prosthetic limb simply cannot anticipate or cope with. From its agential, subjective point of view, it is all a matter of guessing the agent's intentions, constructing a model, a predictive schema of EMG patterns and adapting to its constraints. Little does it know that there is a concrete living being on the other side struggling to produce movements by generating the right muscular contractions against torques and discomfort. But these muscular contractions not only need to deal with gravitational forces, they are also not *meant* to serve as signals for an ML-based prosthetic limb to interpret – we do not move our healthy limbs through vicarious muscular contractions and high-amplitude electric potential at the surface of our muscles. We move our healthy limbs by fitting our intentions to the external non-muscular forces, but our intentions are embodied and include the limb itself as well as the external force fields (Merleau-Ponty

1962). By applying the authors' version of a radical constructivist ML solution to prosthetic limbs, this embodiment is lost, and what was a direct, high-bandwidth, logically shallow flow from intention to motor performance in the healthy coupling is interrupted by representations (Neumann 1958; Haugeland 1998; Dreyfus 2002). The result is qualitatively the same as it was with so-called "realist" ML, because both the latter and what the authors take to be a radical constructivist ML remain representational. Restoring the functionality of a lost limb, given the loss of its deeply intimate coupling with the rest of our body and nervous system, is extremely difficult. However, progress will hardly be made by establishing even more subjective discontinuities between patient, prosthetic limb and world. On the contrary, what is needed is a blurring of the boundaries between them as much as possible and an acknowledgement of the overwhelming influence of material constraints such as external forces, which are more than meaningless perturbations, and can also be an active part of the coordination of movement (Bernstein 1967).

Martin Flament Fultot holds a MA in Philosophy from the University of Paris IV Sorbonne. He is currently a PhD candidate also at the Center for the Ecological Study of Perception and Action at the University of Connecticut and a member of the Science, Norms, Decisions research team in France in association with the CNRS. His research interests include the philosophy of cognitive and behavioral sciences, behaviorism, ecological psychology, enactivism, self-organization, dynamical systems, robotics, artificial intelligence, embodied cognition, biology and thermodynamics.

Received: 23 February 2018
Accepted: 25 February 2018

Authors' Response

Radical Constructivism in Machine Learning: We Want More!

Markus Nowak, Claudio Castellini & Carlo Massironi

> Upshot • Our commentators' very constructive criticisms point out a number of weaknesses in the design of our experiment, and offer insight into how such weaknesses might have led to the poor results of the experiments. We summarise the suggestions, which point in a few precise directions, and outline how we will try to implement them in the near future.

« 1 » While there seems to be general agreement among our commentators that framing interactive machine learning (iML) in radical constructivism (RC) is a valid contribution to research in machine learning (ML) and could improve the situation of myocontrol, they also point out that the experiment might have been too simple, and for that reason it has failed to fully prove our point. They all suggest improvements to both the setup and the experimental design, and it is fascinating to note that similar suggestions, or at least suggestions that hint at similar directions, come from researchers who are experienced in very different fields. We feel inspired and will continue on this path, extending it further.

As simple as possible, but not simpler

« 2 » Our working assumption was that rethinking myocontrol through RC would give hints and suggestions for the design of the upper-limb prosthesis of the near future. This intuition stemmed from our own previous work in incremental/interactive learning applied to myocontrol, for which RC seems to provide an ideal theoretical framework, as it is rich in ideas to be applied in practice. So, we put in place a simple experiment with three variants, each variant being "more radical constructivist" than the previous one, hoping that the ideas based upon RC would have led to a clear improvement in performance.

« 3 » This has turned out to hardly be the case. So, was our working assumption correct? The consensus emerging from the open peer commentaries that we received seems to indicate that it was. But at the same time, the general feeling is that the experiment was too simple with respect to medical applications, and that this excess simplicity has blurred the distinctions between the three experiments almost completely. Once again, Albert Einstein's motto, which we have cited already in §40 of our target article, is key: if you simplify too much you lose sight of the problem you are trying to solve, and you might end up finding an inadequate solution.

Did we simplify too much?

« 4 » **William Craelius**'s statement "The experiments were elegantly designed and executed, but the results rather disappointing" (§6) is a perfect synopsis of our work. His first criticism refers to the lack of purpose in the tasks and therefore the lack of motivation for the participant. Although we fully agree with his remedy of shaping the tasks in a game-like manner, some studies, such as one by Ludger van Dijk et al. (2016), point out that improvements achieved in games using control modalities based on muscle signals do not necessarily transfer to improvements in prosthetic control. **Craelius** himself, with reference to Yungher & Craelius (2012), states that "a significant improvement in speed of 15% was achieved after 30 trials with their virtual assistive robot; controls, in contrast, showed negligible improvement with their prosthesis" (§8). Therefore, we chose tasks or actions that reflect what a user would do with her prosthesis, although a repetitive task might lack excitement for her.

« 5 » **Craelius** issued two more criticisms. Firstly, he pointed out that there was a lack of disabled subjects in our study. This is undeniably a shortcoming, as is the low number of participants, which we will amend in the mid-term future. Secondly, Craelius stated that the residuum can accommodate a larger number of sensors than a sound limb. In our experience, though, the opposite seems to be the case. Still, we agree that an increased number of sensors can lead to richer information about muscle activity. When it comes to placement, our approach is to cover the whole circumference of the residuum rather than to target specific muscles. Doing so is our way of dealing with the mentioned radical muscle and tendon rearrangement and disruption of synergies among them.

« 6 » This last criticism is shared by Richard Weir. His lament against the non-patient-centric view of the ML community (§§1f), of which we are part and parcel, is totally well founded. For starters, electromyography has a number of well-known downsides, and the scientific community has been suggesting for almost a decade now that novel ways to detect muscle activity in a residual limb should be conceived and tested (Castellini 2014). We ourselves are active in this field so we could not agree more on this, but this was not the focus of this work. Anyway, we acknowledge that using too few sensors (possibly of the wrong kind) would inevitably make sophisticated interaction useless – a view that all our commentators seem to share. Improving the sensors should be synergistically coupled with the RC-framed approach to iML. Furthermore, the lack of feedback is an issue that is present in the entire field of myoelectric control. So far, no clinical system (besides body-powered hooks) provides relevant feedback to the wearer. As yet, we ourselves have investigated this topic very little.

« 7 » Again, we fully agree with Weir (§13) that the "ML algorithm […] is not the determinant for success"; we rather argue that the ML method, whatever it is, needs to possess certain characteristics – at least incrementality, which leads to interactivity. The pilot study presented here makes intentional use of a standard method, briefly mentioned in §40 of our target article. We solely used the low-pass filtered rectified amplitude of the signals. (For further details on the method we refer the interested reader to Gijsberts et al. 2014.)

« 8 » Furthermore, our work does not in any way challenge the effectiveness of the standard two-sites-of-residual-activity myoelectric system widely used in clinics. Pattern recognition (PR) potentially solves the issue of switching commands required in two-sites control, leading to "natural" control. From our experience, not having to rely on these commands would be a very welcome advancement for the prosthesis wearers; however, PR comes at a price, one of which being the lengthy initial calibration process. But the work that we present in our article is aimed at tackling exactly this problem: we want to eliminate the need to train all possible actions, in all required postures, for several times at once, in the beginning.

« 9 » This is exactly where interactivity leads to a better combined performance of wearer and prosthesis, following the imperative: "do not collect *more* data, rather collect *better* data." In Experiment 2 for example, we started with an empty model (no training data at all) and updates only occurred when, and if, required. Repeating all gestures in all postures, as mentioned by Weir ("moving from one posture to another may result in overlapping muscle activity patterns," §10), is exactly what the RC-framed iML should avoid.

« 10 » All in all, the picture starts to emerge – myocontrol could be a paradigmatically *holistic* problem: either you solve all its aspects at once, or you will not be able to solve it at all. Therefore, all suggestions we received (design more engaging tasks; improve the sensors; improve interaction; and give feedback) need to be taken into account collectively. Peter Cariani, starting from his conceptual background and his research agenda as a physiologist, gets to a similar conclusion. He compares the physiology of motor control under natural vs. experimental conditions and suggests enriching the human-machine interaction by adding more *observables, actions,* and *feedbacks* available to both the machine and the human, such as incorporating proprioceptive and tactile feedback signals from artificial hands and arms into prosthetic controllers. The grand goal is that of reproducing the wealth of bidirectional flow of information taking place in intact subjects: a large number of sensory channels, significant proprioceptive feedback, tactile feedback, copies of the command signals that are activating muscles.

« 11 » Regarding Cariani's questions about the usage of the normalised root mean squared error: the nRMSE was used only as an *a posteriori* inter-subjective measure of the accuracy of the system and it played no role whatsoever in the selection of the training data. The training feedback was always and only that provided by the subject (Q1). We confirm that in this case the feedback could have easily been replaced by a threshold posed on the nRMSE itself (Q2), which is what usually is done in the field of myocontrol – for instance when using the Target Achievement Control test as described by Ann Simon et al. (2011). Actually, from the point of view of the engineer, this is a very unusual characteristic of our experimental protocol: to employ a subjective judgment to determine whether a task was successful or not, instead of an inter-subjectively verifiably measure. We ourselves have used such measures in the past.

« 12 » Particularly fascinating in Cariani's commentary is the idea that the initial choice of observables, actions and feedbacks determines a cognitive "cage" in which the ML system is trapped – and since we have had no chance so far to design a ML system that evolves its own sensors and actuators, the cage remains as it is for the rest of the experiment and plays a key role in its outcome.

« 13 » There is an unfortunate practical implication of Cariani's idea: any prosthetic system endowed with insufficient hardware will *never* get to a satisfactory level of integration and performance, no matter how smart the ML method and/or the interaction schema is – one more hint at the holistic nature of myocontrol. Things are made even worse by the extremely high acceptance threshold in the field, as pointed out by Weir (§3), who, by the way, also touches upon this "cannot-neglect-the-hardware" conundrum when he says:

> "It is the limited number of control sites and the associated limit on the number of controllable DoF [degrees of freedom] that led investigators to explore other means of acquiring and using multi-DoF control schemes such as ML. Users certainly want more DoFs, but not if it is a hassle." (§5)

« 14 » Cariani's view is that human-machine interaction can be seen as two *adaptive, purposive percept-coordination-action* *systems*, in which the data collection should be even more dependent on the evaluative feedback to the machine than in our simple experiment. Here too, we could not agree more.

Attention and culture

« 15 » More ideas and suggestions, particularly focussing on the interaction, are to be found in the remaining commentaries. Starting from an exquisitely radical constructivist perspective and adopting our research agenda as a working hypothesis, **Marco Bettoni** offers some useful linguistic/conceptual suggestions and two operational models. **Bettoni** suggests operating a conceptual switch from organising *perceptual* objects to organising a *sensory* field and then, after having constructed a conceptual object, assigning the object not to the sensory field but rather to the "experiential reality" (a higher operational level). We definitely agree with his suggestions, particularly with the request to avoid using the word "match" altogether, "even when it refers to sensory patterns or conceptual structures and not to pictures of the physical world" (§11).

« 16 » From a sociocultural perspective, **Marco Guicciardi** argues that human-machine interaction loops must be grounded in specific social practices, culturally and historically situated (repeated experiential interaction loops are not enough). In designing our experiments, we have already tried to enrich the socio-cultural dimension with respect to a classical experimental setting by working on the dimension of meaning, and on the mutual roles of human and machine: designing meanings for the interaction and inventing reciprocal roles for human and machine.

« 17 » Still, **Guicciardi** goes even further, suggesting giving the iML system the capacity to grasp the "frame" of the interaction (the rich context), to uncover the meaning (culturally and historically situated) of a gesture. His suggestions are twofold: we should give the iML system the capacity to grasp the sociocultural context of the interaction; at the same time, we should enable the participant to be more aware of the iML system's knowledge constructs; and both should have the capacity to create together new knowledge outside the original social context where distances, shapes and sizes are culturally defined. In order to provide at least a partial answer to his Q1, one initial move in this direction is to diversify and enhance the sensor modalities available to the iML system – not only to have more sensors of a specific kind, but also more different kinds of sensors relating to different kinds of data. For instance, there could be feedback from the device, environmental information, a more articulated dialogue between the prosthesis and the participant and a skilled way of extracting information from it. In this sense, it is likely that the more data, the better, provided that the iML system is able to discern the relevant information from that which is irrelevant.

As salamanders do?

« 18 » Finally, **Martin Flament Fultot** starts from a non-representationalist conceptual background, *à la* "intelligence without representation" (Brooks 1991), mixing behaviourism and Maurice Merleau-Ponty's phenomenology. **Flament Fultot**'s research agenda is extremely different from ours:

> "The ideal goal would be to be able to grow a new limb, as salamanders do, and let the interactions between neural, muscular and bony tissues adapt to one another during the growth process. The end result would be an embodied limb." (§7)

As he clearly states, he is not interested in trying to build "anything that resembles a *separate* agent that interacts with us through intermediate mental schemata" (ibid), as we, instead, are.

« 19 » For instance, in §3, **Flament Fultot** suggests trying not to represent the world through a model, but rather to *fit the world just like the woodpecker's beak fits the tree*. But what is meant by *representation* here? In ML it is customary to do away with this concept by blurring the distinction between a representation and, for example, the weights of a neural network. These two positions do not clash with each other, rather they start from two completely different definitions of a representation. For instance, we agree with **Flament Fultot** about stressing the concept of fitting, but we are definitely not interested in equipping the iML system with the capacity to build a "true representation" of the world. Rather, our ideal iML system should just organise its sensory field to build its experiential reality (see **Bettoni**'s commentary and our response above).

« 20 » Surprisingly, there is a final point of strong agreement between **Flament Fultot** and us, and this is the concept of embodiment, or more precisely *having the prosthesis feel like a part of the patient's body*. This concept is slowly finding its way in the human-robot-interaction community, too, due to the intuition that control will improve as the user embodies the prosthesis. Such embodiment can only be realised via technologies that are not yet in sight, including extreme mechatronic dexterity, detailed feedback with sensory substitution, and close-to-perfect myocontrol. Given the current state of the art, for this experiment we have instead chosen to make the prosthesis a better, friendly, more responsive, *tool/buddy*, but in the future an upper-limb prosthesis will be used like a pair of glasses: don it and it works fine, doff it and go to sleep, don it again the next morning and it will work again just like yesterday – see **Weir**'s remark at §8. The road to embodiment is still very long, but we view our attempt as a small step towards that goal.

Received: 11 March 2018
Accepted: 12 March 2018

Combined References

Ajiboye A. B. & Weir R. F. ff. (2005) A heuristic fuzzy logic approach to EMG pattern recognition for multifunctional prosthesis control. IEEE Transactions on Neural Systems and Rehabilitation Engineering 13(3): 280–291.

Amershi S., Cakmak M., Knox W. B. & Kulesza T. (2014) Power to the people: The role of humans in interactive machine learning. AI Magazine 35(4): 105–120.

Atkins D., Heard D. C. Y. & Donovan W. H. (1996) Epidemiologic overview of individuals with upper-limb loss and their reported research priorities. Journal of Prosthetics and Orthotics 8: 2–11.

Baschuk C. M., Hoeun P., Katzenberger L. M., Latour D. A., Passero T. & Tompkins E. J. (2016) Utilization of pattern recognition with patients initially contraindicated for myoelectric control of upper-limb prostheses. In: Proceedings of the first international symposium on innovations in amputation surgery and prosthetic technologies, Chicago IL: 79–81.

Bateson G. (1972) Steps to an ecology of mind. Chandler, San Francisco. Reprinted in 2000 by The University of Chicago Press.

Bateson G. (1979) Mind and nature: A necessary unit. E. P. Dutton, New York.

Berger P. & Luckmann T. (1967) The social construction of reality: A treatise in the sociology of knowledge. Anchor, Garden City NY.

Bernstein N. (1967) The coordination and regulation of movements. Pergamon Press, Oxford.

Bettoni M. (2018) Aufmerksamkeitsquanten: Bausteine autonomen und kreativen Denkens. In: Hug T., Mitterer J. & Schorner M. (eds.) Proceedings of the Ernst von Glasersfeld Centenary Conference. Innsbruck University Press, Innsbruck, in press.

Bettoni M. & Eggs C. (2010) User-centred knowledge management: A constructivist and socialized view. Constructivist Foundations 5(3): 130–143. ▸ http://constructivist.info/5/3/130

Brooks R. A. (1991) Intelligence without representation. Artificial Intelligence 47: 139–159. ▸ http://cepa.info/4059

Card S. K., Newell A. & Moran T. P. (1983) The psychology of human-computer interaction. Lawrence Erlbaum, Hillsdale NJ.

Cariani P. (1989) On the design of devices with emergent semantic functions. Unpublished PhD Dissertation, State University of New York at Binghamton, Binghamton, New York.

Cariani P. (2011) The semiotics of cybernetic percept-action systems. International Journal of Signs and Semiotic Systems 1(1): 1–17. ▸ http://cepa.info/2534

Cariani P. (2015) Sign functions in natural and artificial systems. In: Trifonas P. P. (ed.) International handbook of semiotics. Springer, Dordrecht: 917–950.

Castellini C. (2016) Incremental learning of muscle synergies: From calibration to interaction. In: Bianchi M. & Moscatelli A. (eds.) Human and robot hands: Sensorimotor synergies to bridge the gap between neuroscience and robotics. Springer, New York: 171–193.

Castellini C., Artemiadis P., Wininger M., Ajoudani A., Alimusaj M., Bicchi A., Caputo B., Craelius W., Dosen S., Englehart K., Farina D., Gijsberts A., Godfrey S., Hargrove L., Ison M., Kuiken T., Markovic M., Pilarski P., Rupp R. & Scheme E. (2014) Proceedings of the first workshop on peripheral machine interfaces: Going beyond traditional surface electromyography. Frontiers in Neurorobotics 8: 22.

Ceccato S. (1964) A model of the mind. Methodos 61: 3–75.

Childress D. S. & Weir R. F. ff. (2004) Control of limb prostheses. In: Smith D. G., Michael J. W. & Bowker J. H. (eds.) Atlas of amputations and limb deficiencies: Surgical, prosthetic and rehabilitation principles. Third edition. American Academy of Orthopaedic Surgeons (AAOS) Rosemont IL: 175–198.

Ciancio A. L., Cordella F., Barone R., Romeo R. A., Bellingegni A. D., Sacchetti R., Davalli A., Di Pino G., Ranieri F., Di Lazzaro V., Guglielmelli E., Zollo L. (2016) Control of prosthetic hands via the peripheral nervous system. Frontiers in Neuroscience 10: 116.

Connan M., Ruiz Ramírez E., Vodermayer B. & Castellini C. (2016) Assessment of a wearable force- and electromyography device and comparison of the related signals for myocontrol. Frontiers in Neurorobotics 10: 17.

Davidson D. (1974) On the very idea of a conceptual scheme. Proceedings and Addresses of the American Philosophical Association 47: 5–20.

de Bono E. (1991) I am right, you are wrong. Penguin Books, London.

de Bono E. (1992) Serious creativity. Harper Collins, New York.

Dreyfus H. L. (2002) Intelligence without representation – Merleau-Ponty's critique of mental representation: The relevance of phenomenology to scientific explanation. Phenomenology and the Cognitive Sciences 1(4): 367–383. ▸ http://cepa.info/4592

Englehart K. & Hudgins B. (2003) A robust, real-time control scheme for multifunction myoelectric control. IEEE Transactions on Biomedical Engineering 50(7): 848–854.

Fails J. A. & Olsen D. R. (2003) Interactive machine learning. In: Proceedings of the 8th international conference on Intelligent user interfaces (IUI '03) ACM, New York NY: 39–45.

Farina D., Jiang N. & Rehbaum H. (2014) The extraction of neural information from the surface EMG for the control of upper-limb prostheses: Emerging avenues and challenges. IEEE Transactions on Neural Systems and Rehabilitation Engineering 22(4): 797–809.

Farrell T. R. & Weir R. F. ff. (2008a) A comparison of electrode implantation and targeting on pattern classification accuracy for prosthesis control. IEEE Transactions on Biomedical Engineering 55(9): 2198–2211.

Farrell T. R. & Weir R. F. ff. (2008b) Analysis window induced controller delay for multifunctional prostheses. In: Proceedings of the University of New Brunswick's Myoelectric Controls/Powered Prosthetics Symposium (MEC '08) Fredericton, New Brunswick, Canada: 13–15 August 2008: 225–228.

Fisher R. A. (1936) The use of multiple measurements in taxonomic problems. Annals of Human Genetics 7(2): 179–188.

Foerster H. von (1981) Notes on an epistemology for living things. In: Foerster H. von, Observing systems. Intersystems Publications, Seaside CA: 258–271.

Fougner A., Scheme E., Chan A. D. C., Englehart K. & Stavdahl O. (2011) Resolving the limb position effect in myoelectric pattern recognition. IEEE Transactions on Neural Systems and Rehabilitation Engineering 19(6): 644–651.

Fougner A., Stavdahl Ø., Kyberd P. J., Losier Y. G. & Parker P. A. (2012) Control of upper limb prostheses: Terminology and propor-

tional myoelectric control – A review. IEEE Transactions on Neural Systems and Rehabilitation Engineering 20(5): 663–677.
Gergen K. J. (1978a) Experimentation in social psychology: A reappraisal. European Journal of Social Psychology 8: 507–527.
Gergen K. J. (1978b) Toward generative theory, Journal of Personality and Social Psychology 36(11): 1344–1360.
Gergen K. J. (1985) The social constructionist movement in modern psychology. American Psychologist 40(3): 266–275. ▸ http://cepa.info/4538
Gergen K. J. & Gergen M. (eds.) (1985) Historical social psychology. Lawrence Erlbaum, Hillsdale NJ.
Gibson J. J. (1979) The ecological approach to visual perception. Houghton, Boston.
Gijsberts A., Bohra R., Sierra González D., Werner A., Nowak M., Caputo B., Roa Garzón M. A. & Castellini C. (2014) Stable myoelectric control of a hand prosthesis using non-linear incremental learning. Frontiers in Neurorobotics 8: 8.
Glasersfeld E. von (1983) Learning as constructive activity. In: Bergeron J. C. & Herscovics N. (eds.) Proceedings of the 5th Annual Meeting of the North American Group of Psychology in Mathematics Education, Vol. 1. PME-NA, Montreal: 41–101. ▸ http://cepa.info/1373
Glasersfeld E. von (1995) Radical constructivism: A way of knowing and learning. Falmer, London. ▸ http://cepa.info/1462
Glasersfeld E. von (1999) How do we mean? A constructivist sketch of semantics. Cybernetics and Human Knowing 6(1): 9–16. ▸ http://cepa.info/1503
Glasersfeld E. von (2005) Thirty years radical constructivism. Constructivist Foundations 1(1): 9–12. ▸ http://constructivist.info/1/1/009
Goffman E. (1964) The neglected situation. American Anthropologist. 66(2): 133–136.
Goffman E. (1974) Frame analysis. An essay on the organization of experience. Harper & Row, New York.
Gorsic M., Cikajlo I., Goljar N. & Novak D. (2017) A multisession evaluation of an adaptive competitive arm rehabilitation game. Journal of Neuroengineering and Rehabilitation 14: 128.
Hahne J., Markovic M. & Farina D. (2017) User adaptation in myoelectric man-machine Interfaces. Scientific Reports 7: 4437.
Harré R. (1979) Social being: A theory for social psychology. Blackwell, Oxford.
Harré R. (ed.) (1986) The social construction of emotions. Blackwell, Oxford.
Harré R., Smith J. & van Langenhove L. (eds.) (1995) Rethinking psychological methods. Sage, London.
Haugeland J. (1998) Having thought. Harvard University Press, Cambridge MA.
Herberts P., Almström C., Kadefors R. & Lawrence P. (1973) Hand control via myoelectric patterns. Acta Orthopaedica Scandinavica 44: 389–409.
Holzinger A. (2016) Interactive machine learning for health informatics: When do we need the human-in-the-loop. Brain Informatics 3: 119–131.
Hudgins B., Parker P. & Scott R. N. (1993) A new strategy for multifunction myoelectric control. IEEE Transactions on Biomedical Engineering 40(1): 82–94.
Iturrate I., Chavarriaga R., Montesano L., Minguez J. & del R. Millán J. (2015) Teaching brain-machine interfaces as an alternative paradigm to neuroprosthetics control. Scientific Reports 5: 13893.
Jiang N., Dosen S., Müller K. R. & Farina D. (2012) Myoelectric control of artificial limbs: Is there a need to change focus? IEEE Signal Processing Magazine 29(5): 148–152.
Kant I. (1966) Kritik der reinen Vernunft. Edited by I. Heidemann. Reclam, Stuttgart. Edition A originally published in 1781, edition B in 1787.
Kuttuva M., Flint J., Burdea G. & Craelius W. (2005) Manipulation practice for upper-limb amputees using virtual reality. Presence 14: 175–182.
Lave J. (1988) Cognition in practice: Mind, mathematics and culture in everyday life. Cambridge University Press, Cambridge UK.
Lawrence P. D. & Kadefors R. (1974) Classification of myoelectric patterns for the control of a prosthesis. In: Herberts P., Kadefors R., Magnusson R. & Petersen I. (eds.) The control of upper-extremity prostheses and orthoses. C. C. Thomas, Springfield IL: 190–200.
Lock B. A., Simon A. M., Stubblefield K. & Hargrove L. J. (2011) Prosthesis-guided training for practical use of pattern recognition control of prostheses. In: MEC 11 Raising the Standard. Proceedings of the 2011 MyoElectric Controls/Powered Prosthetics Symposium. University of New Brunswick, Fredericton CA: 61–64.
Lock B., Simon A. M., Stubblefield K. & Hargrove L. (2011) Prosthesis-guided training for practical use of pattern recognition control of prostheses. In: Parker P. A. (ed.) MEC'11 Raising the standard. UNB, Fredericton: 61–64.
Lutterer W. (2012) Deutero-learning. In: Seel N. M. (eds.) Encyclopedia of the sciences of learning. Springer, New York: 939–941.
Mathewson K. W. & Pilarski P. M. (2017) Reinforcement learning based embodied agents modelling human users through interaction and multi-sensory perception. In: AAAI spring symposium on interactive multisensory object perception for embodied agents, March 27–29, Stanford University, USA: 477–481.
Merleau-Ponty M. (1962) Phenomenology of perception. Routledge and Kegan Paul, New York.
Metzger A. J., Dromerick A. W., Holley R. J., Lum P. S. (2012) Characterization of compensatory trunk movements during prosthetic upper limb reaching tasks. Archives of Physical Medicine and Rehabilitation 93: 2029–2034.
Micera S., Carpaneto J. & Raspopovic S. (2010) Control of hand prostheses using peripheral information. IEEE Reviews in Biomedical Engineering 3: 48–68.
Neumann J. von (1958) The computer and the brain. Yale University Press, New Haven CT.
Packer M. J & Goicoechea J. (2000) Sociocultural and constructivist theories of learning: Ontology, not just epistemology. Educational Psychologist 35(4): 227–241.
Parker P. A. & Scott R. N. (1985) Myoelectric control of prostheses. Critical Reviews in Biomedical Engineering 13(4): 283–310.
Peerdeman B., Boere D., Witteveen H., in 't Veld R. H., Hermens H., Stramigioli S., Rietman H., Veltink P. & Misra S. (2011) Myoelectric forearm prostheses: State of the art from a user-centered perspective. Journal of Rehabilitation Research and Development 48(6): 719–738.
Ravindra V. & Castellini C. (2014) A comparative analysis of three non-invasive human-machine interfaces for the disabled. Frontiers in Neurorobotics 8: 24.
Samuel A. L. (1959) Some studies in machine learning using the game of checkers. IBM Journal of Research and Development 3(3): 210–229.

Sarkar A. (2016) Constructivist design for interactive machine learning. In: Proceedings of the 2016 CHI conference extended abstracts on human factors in computing systems (CHI EA '16) ACM, New York: 1467–1475.

Scheme E. & Englehart K. (2011) Electromyogram pattern recognition for control of powered upper-limb prostheses: State of the art and challenges for clinical use. Journal of Rehabilitation Research and Development 48(6): 643–659.

Schweitzer W., Thali M. J. & Egger D. (2018) Case-study of a user-driven prosthetic arm design: Bionic hand versus customized body-powered technology in a highly demanding work environment. Journal of NeuroEngineering and Rehabilitation 15: 1.

Sensinger J. W., Lock B. A. & Kuiken T. A. (2009) Adaptive pattern recognition of myoelectric signals: Exploration of conceptual framework and practical algorithms. IEEE Transactions on Neural Systems and Rehabilitation Engineering 17(3): 270–278.

Shalev-Shwartz S. & Ben-David S. (2014) Understanding machine learning: From theory to algorithms. Cambridge University Press, New York.

Shaw R. E., Kugler P. N. & Kinsella-Shaw J. (1990) Reciprocities of intentional systems. In: Warren R. & Wertheim A. H. (eds.) Resources for ecological psychology: Perception and control of self-motion. Lawrence Erlbaum Associates, Hillsdale NJ: 579–619.

Simon A. M., Hargrove L. J., Lock B. A. & Kuiken T. A. (2011) Target achievement control test: Evaluating real-time myoelectric pattern-recognition control of multifunctional upper-limb prostheses. Journal of Rehabilitation Research & Development 48(6): 619–628.

Simon A. M., Lock B. A. & Stubblefield K. A. (2012) Patient training for functional use of pattern recognition-controlled prostheses. Journal of Prosthetics and Orthotics 24(2): 56–64.

Simon A. M., Lock B. A., Stubblefield K. A. & Hargrove L. J. (2011) Prosthesis-guided training increases functional wear time and improves tolerance to malfunctioning inputs of pattern recognition–controlled prostheses. In: MEC 11 Raising the Standard. Proceedings of the 2011 MyoElectric Controls/Powered Prosthetics Symposium. University of New Brunswick, Fredericton CA: 65–69.

Strazzulla I., Nowak M., Controzzi M., Cipriani C. & Castellini C. (2017) Online bimanual manipulation using surface electromyography and incremental learning. IEEE Transactions on Neural Systems and Rehabilitation Engineering 25: 227–234.

Taylor D. R. & Finley F. R. (1974) Multiple-axis prosthesis control by muscle synergies. In: Herberts P., Kadefors R., Magnusson R. & Petersen I. (eds.) The control of upper-extremity prostheses and orthoses. C. C. Thomas, Springfield IL: 181–189.

Uellendahl J. & Tyler J. (2016) A case series study of pattern recognition for upper-limb prosthesis control. In: Proceedings of the first international symposium on innovations in amputation surgery and prosthetic technologies, Chicago IL: 68.

Uexküll J. von (2010) A foray into the worlds of animals and humans: With a theory of meaning. Translated by J. D. O'Neil. University of Minnesota Press, University of Minnesota Press, Minneapolis MN.

van Dijk L., van der Sluis C. K., van Dijk H. W. & Bongers R. M. (2016) Task-oriented gaming for transfer to prosthesis use. IEEE Transactions on Neural Systems and Rehabilitation Engineering 24(12): 1384–1394.

Varela F. J., Thompson E. & Rosch E. (1991) The embodied mind: Cognitive science and human experience. MIT Press, Cambridge MA.

Wagstaff K. L. (2012) Machine learning that matters. In: Langford J. & Pineau J. (eds.) Proceedings of the 29th international conference on machine learning (ICML-12) Omnipress, Madison WI: 529–536.

Wang R., King T., Ossipov M. H., Rossomando A. J., Vanderah T. W., Harvey P., Cariani P., Frank E., Sah D. W., Porreca F. (2008) Persistent restoration of sensory function by immediate or delayed systemic artemin after dorsal root injury. Nature Neuroscience 11(4): 488–496.

Yungher D. & Craelius W. (2012) Improving fine motor function after brain injury using gesture recognition biofeedback. Disability and Rehabilitation: Assistive Technology 7(6): 464–468.

Zhou P., Lowery M. M., Englehart K. B., Huang H., Li G., Hargrove L., Dewald J. P. A. & Kuiken T. A. (2007) Decoding a new neural machine interface for control of artificial limbs. Journal of Neurophysiology 98(5): 2974–2982.