Improving Reliability of Myocontrol Using Formal Verification

Dario Guidotti, Francesco Leofante, Armando Tacchella, and Claudio Castellini

***Abstract*—In the context of assistive robotics, myocontrol is one of the so-far unsolved problems of upper-limb prosthetics. It consists of swiftly, naturally, and reliably converting biosignals, non-invasively gathered from an upper-limb disabled subject, into control commands for an appropriate self-powered prosthetic device. Despite decades of research, traditional surface electromyography cannot yet detect the subject's intent to an acceptable degree of reliability, that is, enforce an action exactly when the subject wants it to be enforced. . In this paper, we tackle one such kind of mismatch between the subject's intent and the response by the myocontrol system, and show that formal verification can indeed be used to mitigate it. Eighteen intact subjects were engaged in two target achievement control tests in which a standard myocontrol system was compared to two "repaired" ones, one based on a non-formal technique, thus enforcing no guarantee of safety, and the other using the satisfiability modulo theories (SMT) technology to rigorously enforce the desired property. The experimental results indicate that both repaired systems exhibit better reliability than the non-repaired one. The SMT-based system causes only a modest increase in the required computational resources with respect to the non-formal technique; as opposed to this, the non-formal technique can be easily implemented in existing myocontrol systems, potentially increasing their reliability.**

***Index Terms*—Myocontrol, prosthetics, electromyography, assistive robotics, formal verification, satisfiability modulo theories.**

I. Introduction

DESPITE decades of joint academic research in the fields of assistive robotics, machine learning and human-robot interaction, a widely accepted self-powered upper-limb prosthesis is still not available in the clinics [1]. The ideal upper-limb prosthesis is modular (i.e., applicable to all possible upper-limb amputations), human-looking, light-weight, silent and dexterous, enabling the amputated person to recover most of the lost physiological functions [2]. The most advanced attempts in this direction, the *DEKA Arm* and its commercial counterpart, the *Luke* prosthetic arm (www.mobiusbionics.com/luke-arm), already enforce some of these characteristics from a mechatronic point of view — whether they are effective in daily living is currently being assessed [3].

Manuscript received August 8, 2018; revised November 26, 2018; accepted January 9, 2019. Date of publication March 5, 2019; date of current version April 8, 2019. *(Corresponding author: Dario Guidotti.)*

D. Guidotti, F. Leofante, and A. Tacchella are with the Department of Informatics, Bioengineering, Robotics and Systems Engineering, University of Genoa, 16126 Genoa, Italy (e-mail: dario.guidotti@edu.unige.it).

C. Castellini is with the DLR—German Aerospace Center, Institute of Robotics and Mechatronics, 82234 Weßling, Germany.

This paper has supplementary downloadable material available at http://ieeexplore.ieee.org, provided by the author.

Digital Object Identifier 10.1109/TNSRE.2019.2893152

Still, one of the biggest hurdles, if not the biggest one, is how to reliably and naturally let the patient control such devices. It is widely accepted [4] that the main way forward is to apply machine learning (ML) to biological signals, typically surface electromyography (sEMG, [5]), and associate muscle activation patterns to control commands (torque, force, position, velocity) to be issued to the prosthetic device. The concept of natural, simultaneous and proportional (s/p) control [6] constitutes a wishlist for this kind of man-machine interfaces: they must be able to activate the prosthesis as the subject desires, continually, smoothly and effectively. Still, the statistical nature of ML, together with the diversity of human signals, and the unreliability of the physical interface connecting a prosthesis to the subject's body and embedding the sensors turn this problem, in principle an easy one, into an extremely hard one in practice [1].

Properly gathering data to build a ML model enforcing good s/p myocontrol is already a challenge. Batch learning seems to be inappropriate [7] and is being replaced by incremental learning coupled with an appropriate interaction strategy [8]–[10]. But even assuming that a good data set has been gathered in batch fashion, several issues might arise whenever the model of muscle activation blatantly does not reflect the subject's intent; this would be the case, e.g., of a non-linear model *predicting low activation values* (that is to say, weak grasping forces exerted by the prosthesis) *for high values of the sEMG signals.* An example can be found in Figure 2. Consider the curve labelled "Original model": a subject using it to grasp an object would actually cause the prosthesis to release and drop it (predicted activation lower than 1) while increasing her activation past 2.5 to stabilize the grasp. (Multiple evidence is present in literature, e.g. [11]–[13] showing that grasp stabilsation implies in many cases an increase in the global force applied to objects.) In this case, gathering more data from the subject to amend the model is not desirable, as she would need to apply a large amount of force — a task which could lead to muscle strain, fatigue and frustration. Instead, it would be very desirable to mechanically amend the model to have it avoid mismatches such as the one described above.

In order to partially fix this problem, in this paper we propose to couple a standard ML-based s/p myocontrol system with a *Satisfiability Modulo Theories (SMT) solver* in order to iteratively repair and improve the model generated by the ML method, until the intent mismatch is mitigated. As shown for the first time in [14], by leveraging the expressive power of the theories supported by SMT technology, we can represent the ML model as well as a property encoding the desired behaviour (in this case, that the prediction remains higher than 1), as a Boolean combination of arithmetic constraints which can be efficiently reasoned upon by state-of-the-art solvers such as Z3 [15], MathSAT [16] or dReal [17]. Not only we can establish algorithmically whether a model satisfies a given property, but in cases where the property is not satisfied, we can use the SMT solver to iteratively *repair* the ML model. As a consequence of the formal approach, the repaired model is *mathematically guaranteed* to predict values higher than 1 for high values of sEMG, therefore better matching the subject's intent.

SMT solvers [18] are a family of algorithmic procedures used to solve formal verification problems. An SMT solver typically determines the satisfiability of a first-order logic formula expressed in a theory of interest such as, e.g., the theory of lists, arrays, bit vectors or integer arithmetic; in our case, real numbers with transcendental functions. To this aim, given an input formula Φ, an SMT solver first builds its Boolean abstraction $\overline{\Phi}$ by replacing each constraint with Boolean variables $A, B, C, \ldots$:

$$\begin{aligned} \Phi &: \underbrace{x \geq y}_{A} \wedge (\underbrace{y > 0}_{B} \vee \underbrace{x > 0}_{C}) \wedge \underbrace{y \leq 0}_{\neg B} \\ \overline{\Phi} &: \quad A \quad \wedge (\quad B \quad \vee \quad C \quad) \wedge \quad \neg B \end{aligned}$$

where, e.g., $x, y \in \mathbb{R}$. Subsequently, a Boolean Satisfiability solver enumerates all satisfying assignments to $\overline{\Phi}$; if at least one such assignment is found which is also consistent in the underlying theory, then a satisfying solution is found for Φ; otherwise, the formula is unsatisfiable.

To check whether the idea could work in practice, we have engaged 18 intact subjects in two online goal-reaching experiments. Three myocontrol systems were compared in both experiments, namely (a) a standard myocontroller; (b) a myocontroller whose model was repaired naively according to a simple heuristics (PDO), yielding no guarantee of correct behaviour; and (c) a myocontroller which was repaired using SMT. Both experiments were instances of the Target Achievement Control (TAC) test for prosthetic hands [19]: the first, as usual, consisted of a set of online goal-reaching tasks, in which a 3D hand model needed to be held in a specific configuration for a determined amount of time. The second experiment was carefully designed to both check that such a mismatch is a major problem, and that repairing the models can solve it to a large extent: we artificially lowered the predicted activation values, inducing the subjects to apply more force while trying to reach the target. This would simulate the above-mentioned attempt to stabilize the grasp, potentially inducing the wrong control behaviour and thereby making the target very hard to reach.

The experimental results indicate that repairing the models is effective in both cases, and, especially in the second experiment, employing a repaired machine increased the task success rate from 14.58% to about 46%, due to a significantly increased reachability of the targets. The computational price to pay for this improvement is an added duration of the model building phase of about 15s (SMT) and about 8s (PDO). No added time is required while predicting, meaning that the repaired systems affords a higher reliability while leaving the online performance unhampered. Remarkably, the technique retains a large degree of generalizability, so that it could also be employed to enforce more desirable properties such as, e.g., reducing the inter-activation interference.

A. Related Work

To the best of our knowledge, there have been no attempts so far at manipulating models obtained in the context of myocontrol. In [20]–[22] we have already proposed to "dope" the dataset of a myocontrol system with synthetic data obtained by linearly combining preexisting sEMG patterns, in order to be able to predict combined activations of multiple degrees of freedom; but this procedure, although rigorously defined and experimentally validated, has no mechanical component and enforces no mathematical guarantee. On the other hand, several works have been proposed that leverage automated reasoning to verify, and possibly repair, machine learning artifacts. In [23] a method is proposed to verify that models learned with Support Vector Regression [24] always provide a bounded response as long as supplied inputs lie within acceptable operational parameters. SMT solving has also been applied to verify properties of different classes of neural networks such as shallow networks [14], Deep Neural Networks (DNN) [25] and more general classes of DNNs [26] — see [27] for a recent account on the subject. To the best of our knowledge this is the first time that a machine learning artifact is actually deployed in a safety-critical setting after being verified and repaired with SMT.

II. Background

A. Myocontrol and Intent Mismatch

Natural, simultaneous and proportional myocontrol is an instance of (multi-variate) regression as intended in the ML lingo: using a set of $i = 1, \ldots, N$ observations $\mathbf{x}_i \in \mathbb{R}^d$, each one paired with a target value $\mathbf{y}_i \in \mathbb{R}^m$, build an approximant function $f(\mathbf{x}) : \mathbb{R}^d \rightarrow \mathbb{R}^m$ which best fits the set of observations / target values, call it $S = (X, Y)$ with $X \in \mathbb{R}^{N \times d}$ and $Y \in \mathbb{R}^{N \times m}$, and offers the best generalisation power on so-far unseen data. Each observation consists of d features evaluated from a set of sensors and denotes the muscular activation corresponding to an action (e.g., wrist flexion, power grasp, etc.); each associated target value, in turn, is a vector of m motor activation values (currents, torques, ...) for a prosthetic device and corresponds to the desired action as enacted by the device itself. The approximant f is an "intent detector" for an upper-limb prosthesis wearer: whenever the subject's muscles are activated to enforce a specific action, the prosthesis should perform it.

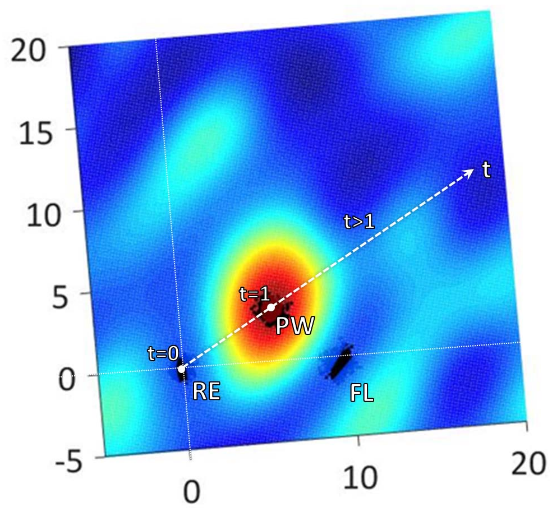

Fig. 1. A typical dataset S, obtained after gathering observations for three actions (black dots; rest, RE; power grasp, PW; wrist flexion, FL); the colour of the heat map denotes the approximant for power grasping, f_{PW}. Values of the input space lying on the straight line $\overline{RE} + (\overline{PW} - \overline{RE})t_{PW}$ roughly denote power grasping with increasing strength.

As is now customary (see, e.g., [28], [29], but also the "prosthesis-driven calibration" enforced in the *Complete Control* myocontrol commercial package by CoApt Engineering, www.coapt.com), in practice S is built by gathering, for each desired action, an adequate number of observations recorded while the subject is stimulated to perform it; each such observation is then coupled with the target value enforcing the action by the prosthetic device. For instance, the subject is asked to power grasp ("make a fist"); once the experimenter verifies that the signals have reached a stable pattern, well distinct from the baseline, a representative amount of observations is recorded and associated to (synthetic) target values denoting maximal activation of all fingers. This methodology is called *on-off goal-directed training*. As a result of this, at the end of the data gathering phase, S consists of one or more observation clusters for each action considered, coupled with adequate target values (refer to [30] for a detailed description of this process). If properly built out of S, f will smoothly and timely activate every motor of the prosthesis whenever required; minimal and maximal muscular activations, as gathered from the subject, will correspond to minimal and maximal motor activations; moreover, under plausible assumptions, intermediate activation values too will be correctly predicted in a monotonically-increasing fashion (examples of this can be found in, e.g., [8] and [30]).

Consider Figure 1, showing a 2D-reduced exemplary S containing three observation sets, RE, PW and FL, gathered in turn while the subject was resting, making a power grasp and flexing the wrist. The Figure also shows, as a heat map, the function f_{PW} corresponding to power grasp, as obtained after building the model using a standard non-linear regression ML method. We assume that the straight line $\overline{RE}+(\overline{PW}-\overline{RE})t_{PW}$ with $t_{PW} \geq 0$ denotes increasing and coordinated activation of the muscles used to power grasp. For $t_{PW} = 0$, that is around $\overline{RE}$ (where by $\overline{C}$ we denote the average of the set C), the subject is at rest; as t_{PW} increases, she starts power-grasping; and she reaches the maximum activation value (the one which she produced during the data gathering phase, around $\overline{PW}$) for $t_{PW} = 1$. Taking into account the natural adaptation of the subject to the system [10], [29], [31], such an activation function could reasonably accommodate the power grasp of a prosthetic wrist for $0 \leq t_{PW} \leq 1$.

Unfortunately, it is also clear that, whenever the subject activates her own muscles to a higher degree than that represented by $\overline{PW}$ (she "grasps with more force"), the activation response is no longer adequate: the value of f_{PW} *decreases* for $t_{PW} > 1$ and it will reach zero already for $t_{PW} \approx 2$. In practice, the subject tries to increase the grasping force but the hand applies *less* of it, almost surely leading to a drop of the grasped object, or worse. We now turn to the problem of trying and "repairing" the model f to avoid this behaviour.

B. Verifying Safety of a ML Model Using SMT

Verification of ML models such as f via SMT involves reasoning over a set of arithmetic constraints expressed over real numbers that provide a rigorous mathematical description of the behaviour of the model. Given a logical formula expressing the behaviour of the model to be verified, call it Φ, we pursue the goal of proving that $\Phi \supset \Sigma$, where the operator $\supset$ denotes logical *implication* and Σ encodes the property that "the predicted activation is always higher than 1 over a specific manifold of interest". More formally, let us assume that a non linear model f has been built out of a data set $S = (X, Y)$ (entirely recorded from the subject during an initial data-gathering phase), which in this specific case takes the form $f(\mathbf{x}) = \mathbf{w}^T\phi(\mathbf{x})$ where $\mathbf{w} \in \mathbb{R}^D$ and $\phi : \mathbb{R}^d \to \mathbb{R}^D$ non-linearly maps d-dimensional observations onto D-dimensional vectors in a *feature space*. To each action A considered while gathering S we then associate one straight line such as the one described above for power grasp, and one parameter $t_A \in \mathbb{R}$. Given the above assumption, that moving along a line t_A denotes performing A with increasing/decreasing force, to avoid low activation values for high sEMG values for A we require that, whenever $t_A \geq 1$, $f(t_A) \geq 1$, that is, in first-order logic terms,

$$\forall t_A.\ t_A \geq 1 \ \supset\ \mathbf{w}^T\phi(t_A) \geq 1.$$

(With a slight abuse of notation, we denote by $\phi(t_A)$ the evaluation of ϕ over the input points lying on the straight line associated to t_A. More rigorously one should write $\phi(\overline{RE} + (\overline{A} - \overline{RE})t_A)$.) So, the logical language in which we must express the above formula requires the usage of constraints containing transcendental functions which make the problem undecidable. In order to tackle undecidability, two different approaches can be pursued. The first possibility is to build a conservative abstraction of the learned model (referred to as *concrete*) that does not include non-linearities and provides an over-approximation of the concrete model. In this way, when the SMT solver proves that the response of the abstract model cannot exceed a stated bound as long as the input values satisfy given preconditions, we can certify that the concrete model enjoys the same property. However, if a counterexample is found then it is either an artifact of the abstraction (being an over-approximation, the abstraction allows for more behaviors than the concrete model), or a true counterexample proving that the concrete model is not safe with respect to the target property. If the counterexample originates from the former, then refinement procedures [32] need to be put in place to tighten the abstraction.

The other possibility is to directly encode the model including the above mentioned non-linearities and resort to *incomplete* decision procedures, as is the case for the solver dReal [17] used for our experiments. In particular, dReal implements a δ-complete decision procedure [33]: Given an SMT formula and a positive rational number δ, the solver determines either that the formula is unsatisfiable, or that its δ-*weakening* is satisfiable.[1]

C. Using SMT to Prevent Dropping Objects

In practice, rather than proving the δ-validity of $\Phi \supset \Sigma$, one hopes to prove the δ-*unsatisfiability* of its negation, $t_A{\geq}1 \ \wedge \ \mathbf{w}^T\phi(t_A) < 1$. If, as opposed to this, at least one value t_A^U can be found which δ-satisfies the formula (an *Unsafe point*), then an activation lower than one will be predicted whenever the muscle activity of the subject reaches $t_A = t_A^U$. In this case we need to generate a new model f' for which $f'(t_A^U) > 1$, then check for δ-satisfiability again. Once the formula is finally declared δ-unsatisfiable (no unsafe points can be found anymore), the resulting f is guaranteed to behave correctly, at least in the sense defined by Σ and within the approximation defined by δ.

To build a new f once a t_A^U is found we proceed as follows: we first generate a set of new observations $X' \sim \mathcal{N}(t_A^U, \sigma_A)$ where $\sigma_A = \text{stdv}(A)$. We then associate to each observation in X' a target value $\mathbf{y}_A$ consisting of a 0 for each of the m motors not involved in a and of a value $y_A = y_0 + c_A(\overline{RE} + (\overline{A} - \overline{RE})t_A^U)$ for each motor involved in a. Here y_0 is experimentally estimated and c_A is the slope of the straight line connecting $(\overline{RE}, 0)$ and $(\overline{A}, 1)$. Intuitively, such target values are an attempt at correcting f_A so that it looks like a linear response for $t_A > 1$. Lastly, we generate a new model out of $S \cup (X', \mathbf{y}_A)$ and repeat this loop until no more unsafe points can be found. A pseudocode of this procedure is visibile below.

w ← **buildModel**(S)
unsafe ← True
while unsafe **do**
  for each action A **do**
    [unsafe,t_A^U] ← δ-**satisfiable**($t_A \geq 1 \wedge \mathbf{w}^T\phi(t_A) < 1$)
    if unsafe **then**
      $(X', Y') \leftarrow (\mathcal{N}(t_A^U, \sigma_A), \mathbf{y}_A)$
      w ← **buildModel**($S \cup (X', Y')$)
    end if
  end for
end while

Figure 2 shows a typical run of the above mentioned algorithm for the action Power Grasp.

[1]The δ-weakening of a formula φ is defined as its numerical relaxation. For instance, the δ-weakening of $x = 0$ is $|x| \leq \delta$.

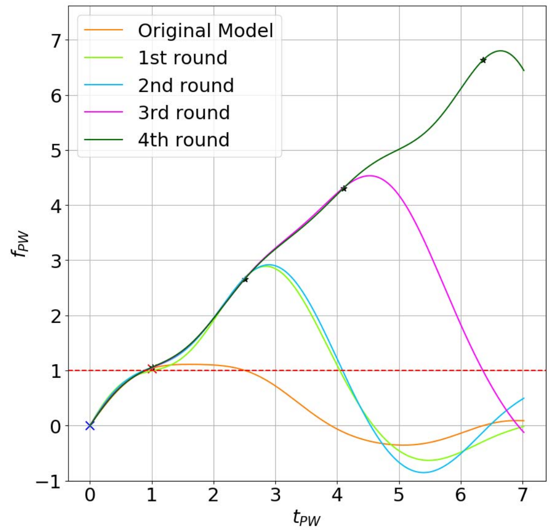

Fig. 2. Successive f_{PW} as returned by a typical run of the SMT repairing algorithm. At each round a new unsafe point is found, a corresponding cluster is added to S and a new f_{PW} is built. Notice how f_{PW} becomes "more and more linear" for $t_{PW} > 1$ as the algorithm progresses.

III. Materials and Methods

We designed two experiments involving human subjects. Both experiments are instances of the Target Achievement Control (TAC) test for prosthetic hands [19], in which a physiologically plausible target configuration of the upper limb is visually presented to the subject, who is then asked to have a virtual upper limb match the stimulus. In our case, each subject needed to reach each goal using either a standard myocontroller, one repaired using the SMT approach described in the previous Section, or one repaired using a simple non-exhaustive unsafe-point search method in place of the SMT system. This method, that we employed as a simpler "baseline" alternative to SMT, works as follows: for each action A and straight line considered, $\overline{RE} + (\overline{A} - \overline{RE})t_A$, we check whether $f_A(t_A) < 1$ for a discrete set of points evenly spaced by 0.001 in the interval $1 < t_A < t_A^{MAX}$, where t_A^{MAX} is computed so that it doesn't exceed the limits of the sensors. We call this alternative method Physiologically Driven Optimization (PDO). Notice that, of course, once PDO yields a positive result, i.e., no more unsafe points, there is no formal guarantee that this is true.

The first experiment was a plain instantiation of the TAC test, with a tolerance threshold of 15% for the goal to be reached, required target dwelling time of 1.5s and timeout of 15s. In the second experiment, the values of f_A were reduced by a factor of 0.75 : the subjects would be stimulated to reach maximal activation and, to this aim, they would reach $t_A = 1$; but they would only see $f_A = 0.75$, and therefore they would then increase their force to $t_A > 1$.

A. Experimental Setup

The experimental setup was common to all experiments, and consisted of a *Myo* bracelet by Thalmic Labs and two 3D hand models displayed on a computer screen. The Myo bracelet (once available at www.thalmic.com, now no longer in production) consists of eight uniformly spaced sensors, able

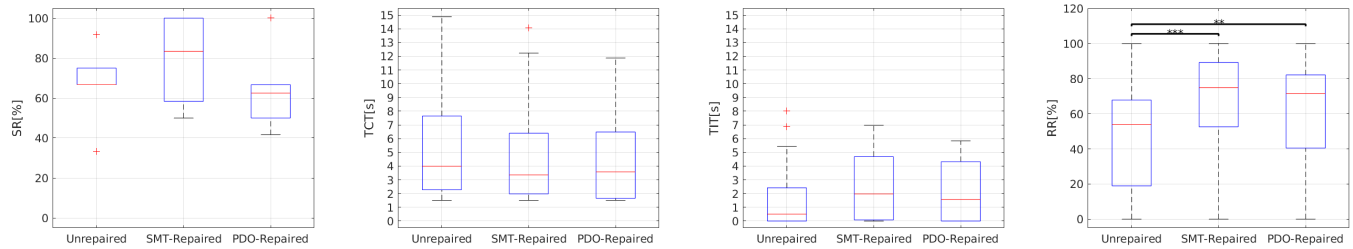

Fig. 3. Success rate (SR), time to complete task (TCT), time in target (TIT), and reaching rate (RR) for the first experiment, for each of the three systems considered.

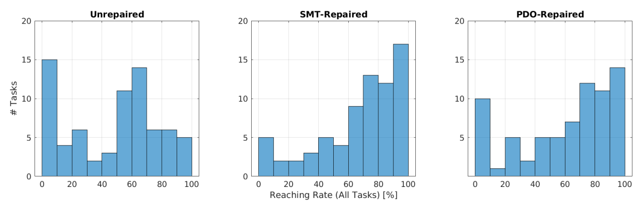

Fig. 4. Frequency diagram (histogram) of the reaching rate, computed on all tasks of the first experiment, for each of the three systems considered.

to detect the electromyographic signal generated by the muscle activity of the forearm. The 3D hand models realistically mimic the motions of a human wrist and hand. One of the models is used to provide the visual stimulus, i.e., it is controlled by the software; the other one enforces the predicted motions of the hand and wrist as evaluated from the data provided by the bracelet and using either the unrepaired system, the SMT-repaired or the PDO-repaired ones.

B. Participants

Six intact human subjects (1 female and 5 males, age 21-32 years) participated in the first experiment whereas twelve (4 females and 8 males, age 23-27 years) participated in the second one. 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 World Medical Association's Declaration of Helsinki.

C. Experimental Protocol

Each participant was assigned a distinct sequence of TAC tasks (i.e., hand/wrist configurations to mimic on the screen) using either the unrepaired, SMT or PDO system. The sequence of both the learning machines and the tasks was randomized for each participant in order to achieve counterbalancing. The number and type of task was the same for each participant.

In both experiments, each participant sat comfortably in front of the computer screen, and the bracelet was wrapped around her/his forearm. She/He was instructed to hold the forearm with an angle of 45 degrees, leaning the elbow on the armrest. We told the participant that she/he would be required to undergo a training session for the prosthesis and then that she/he would be asked to face a series of 36 tasks, during which she/he would need to guide the prosthesis in following the stimulus, using her/his own movement, as close as possible. At the beginning of training session, the stimulus was shown on the screen; we then explained that the stimulus would perform a series of movements (tasks), and that the participant should simply mimic what the stimulus was doing with her own arm. The data-gathering phase followed: each action (rest, RE; wrist flexion, FL; wrist extension, XT; power grasp, PW) was played once by the stimulus, and while the participant followed the movement of the stimulus, the observations were collected. After the end of the data gathering, a second hand model would appear on the screen, and the subject was asked to have this model match the previous one.

D. Data Processing and Intent Detection

The observations collected during the model building phase of the experimental protocol from the Myo bracelet were rectified and mildly low-pass filtered with a 2nd order Butterworth filter (cutoff 1Hz), then directly used together with the target values to train an instance of Ridge Regression with Random Fourier Features (RR-RFF). RR-RFF has been already used to enforce incremental, non-linear s/p myocontrol (see [8] and [9], where more details about the method can be found). Here, it suffices to say that this method enforces an approximant function $f(\mathbf{x}) = \mathbf{w}^T\phi(\mathbf{x})$ of the form described in Subsection II-B, in which ϕ consists of cosines weighted through randomly-sampled frequencies, inducing a finite-dimensional approximation of a Gaussian kernel [34]. RR-RFF is fast both in the evaluation of the model and in prediction, can be made incremental and is bounded in space. It can also be viewed as an instance of Linear Regression employing a specific finite(D)-dimensional kernel. To control the 3D hand model on the screen, the predicted activation values were capped in in the interval [0, 1] and directly used to control the positions of the virtual joints of the model. A movie attached to this paper shows excerpts of a typical experiment.

IV. Experimental Results

A. Performance Measures and Statistics

For each task, four measures of performance were reported of: the *Success Rate (SR)*, i.e., the fraction of tasks successfully completed; the *Time to Complete the Task (TCT)*, that is the time it took to complete asuccessful task; and the *Time In Target (TIT)*, total time the participants managed to stay in the target although the task was unsuccessful. Additionally, in order to check whether a failure was really due to unreachability of the target, for each task we monitored the *Reaching*

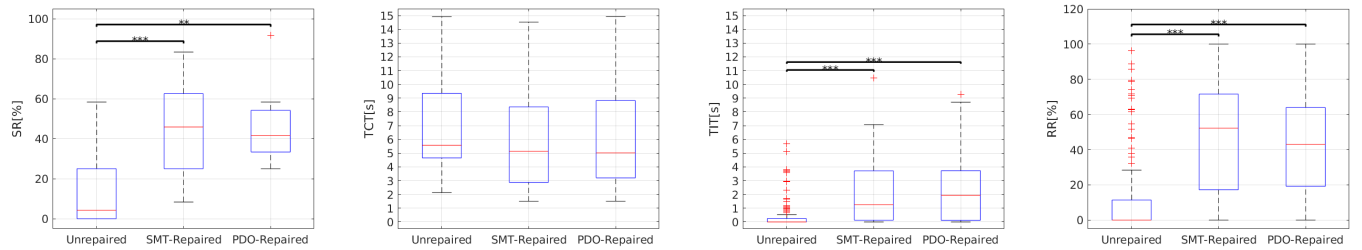

Fig. 5. Success rate (SR), time to complete task (TCT), time in target (TIT), and reaching rate (RR) for the second experiment, for each of the three systems considered.

Rate (RR), defined as the fraction of time during which the output of f_a was higher than the minimum acceptable value; in practice, the percentage of time the subject actually managed to reach the desired activation value. Lastly, we report the *Time To Repair (TTR)*, that is time required to complete the repair process, whenever SMT or PDO were used. In the boxplots the central mark indicates the median and the edges of the box denote the 25th and 75th percentiles; the whiskers extend to the extreme data points not considered outliers; the outliers are plotted individually using the '+' symbol. Statistically significant difference between sample distributions is evaluated using Wilcoxon's Signed Rank test or the Mann-Whitney U test (for paired and unpaired samples in turn), since all distributions were non-Gaussian. Effect sizes are computed using Cohen's d Coefficient. The TCT is evaluated over successful tasks only, whereas the TIT for failed ones only.

B. First Experiment

Figure 3 shows the above-mentioned measures of performance for each of the three systems considered (unrepaired, SMT-repaired and PDO-repaired) obtained by the subjects during the first experiment.

Although the SR is relatively higher when SMT is used, there is no significant difference among the three systems; notice, though, that SMT allows a SR of 100% to be reached, and that the mean SR in this case is quite high ($79.17\% \pm 23.42\%$). The times (both TCT and TIT) are quite similar to one another (again, no significant difference), whereas the RR is significantly higher for unrepaired/SMT ($p < .001, d = 0.71$) and unrepaired/PDO ($p < .01, d = 0.46$). This was to be expected, since a repaired f offers to the subject a larger success interval. Figure 4 confirms this impression: the distribution of the RR is shifted rightwards for SMT and PDO with respect to the unrepaired system (more so for the SMT than for the PDO), denoting that it was much easier in those cases to remain in the target, irrespective of the success in the tasks.

The TTR was $15.32s \pm 4.58s$ for SMT and $10.72s \pm 3.96s$ for PDO.

C. Second Experiment

The picture becomes even clearer if we turn to the results of the second experiment. Consider Figure 5, analogous to Figure 3 but for the second experiment.

Due to unreachability of the targets, the success rate is much lower then in the first experiment ($14.58\% \pm 21.06\%$) for the unrepaired system, and significantly better than this for SMT ($46.53\% \pm 22.88\%$ with $p < .001, d = 1.45$) and for PDO ($45.83\% \pm 18.29\%$ with $p < .01, d = 1.58$).

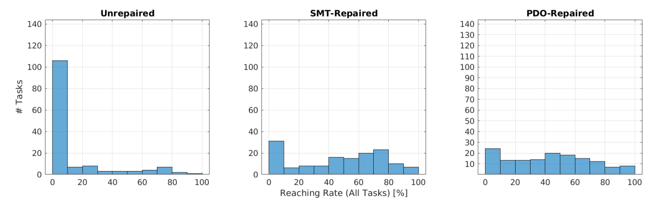

Fig. 6. Frequency diagram (histogram) of the reaching rate, computed on all tasks of the second experiment, for each of the three systems considered.

The TIT is significantly higher for SMT and PDO ($0.45s \pm 1.08s$ for unrepaired, $2.13 \pm 2.31s$ with $p < .001$ for SMT and $2.30s \pm 2.38s$ with $p < .001$ for PDO) whereas the TCTs appear uniform. Lastly, the reaching rates for unrepaired, SMT and PDO (in turn, $12.76\% \pm 23.84\%$, $46.79\% \pm 30.26\%$ and $42.59\% \pm 27.87\%$) are significantly higher for unrepaired/SMT ($p < .001, d = 1.25$) and unrepaired/PDO ($p < .001, d = 1.15$).

Figure 6 reveals that it is much harder to reach the required activation whenever using the unrepaired machine in this second experiment (distribution highly skewed leftwards), whereas, again, using a repaired machine yields distributions which are much less skewed.

Lastly, the TTR was $16.7s \pm 10.02s$ for SMT and $14.12s \pm 13.418s$ for PDO.

V. Discussion and Conclusions

The standard ML-based approach to myocontrol is based upon classification of sEMG patterns. For the past decade, however, s/p myocontrol has been advocated as a better alternative, its main advantage being that it enforces a continuous, infinite manifold of reachable prosthesis configurations, therefore giving to the subject a better control and immersion experience. Moreover, it is deemed that small errors in the intent detection in s/p myocontrol would have a less dramatic effect on the control itself, whereas crossing a decision boundary in classification can lead to catastrophic effects. Notice that proportional (non-simultaneous) myocontrol has been in use since the 60s in the classic two-sensor control system, and that the idea already appears in [35] and [36]; interestingly, the advent of classification caused the loss of proportionality already enforced in the previous approaches. Anyway, here we have employed such an s/p myocontrol system thanks to simultaneous regression on all motors of a virtual prosthetic hand.

Of course, s/p myocontrol also presents a number of added difficulties; in this work we tackled one of them, a non-intuitive behaviour appearing whenever muscle activation is increased to stabilise a grasp but, as a result, the prosthesis applies less force than before. Data gathering in regions of high activation in the input space can be problematic, so we have rather used a formal procedure, SMT, to try and "repair" ML models to mathematically ensure that the predicted activation remains high in such regions. We have also compared SMT with a simpler repairing technique called PDO, which yeidls no guarantee of safety but is easier to implement and faster to execute. The experiment results clearly show that (a) such intent mismatch *does* occur, and can even become ubiquitous in specific conditions; and that (b) SMT/PDO repairing can effectively prevent it.

In particular, consider the results of the first experiment, which enforced the standard condition of using a myocontrol system. Here, using a repaired system versus a standard one improves the success rate. The phenomenon is not statistically significant, possibly also due to the relatively small number of subjects involved. SMT seems to be slightly better than PDO, and it enables a few subjects to achieve *all* tasks (SR=100%). TCT and TIT are very similar to one another, denoting that repaired machines enforce the same performance as the unrepaired one as far as timings are concerned. As opposed to this, the significantly higher RR shows that it was easier for the subjects to actually reach the desired goal, and stay within it, using the repaired systems. The frequency diagrams confirm that *reachability improves once a repaired system is used*. Using a repaired machine added 10-15 seconds to the model-building phase.

The results of the second experiment, in which we "dampened" the predicted activation, thereby inducing the subjects to apply more force to reach the targets, are on the same line and go even further. Here the SR for the unrepaired machine is dramatically low but gets significantly better thanks to the repairing. In this experiment, using a repaired machine added 14-17 seconds to the model-building phase. Notice that in this case a statistically significant difference in the TITs appears, which are much higher for SMT and PDO — actually close to zero for the unrepaired machine (consider also the results and related histograms of the RR). This means that in this experiment, without repairing, *whenever a task fails* (recall that TIT is evaluated on failed tasks only) *it does so because of the wrong behaviour of the model, which predicts low activation values*. In other words, it is close-to impossible for the subjects to reach $f_A > 1$ even though they manage to travel past $t_A = 1$. Of course some subjects could make it, due to a favourable instantiation of the f_A during the data gathering phase. As in all ML, here too randomness can play a crucial role; repairing via SMT can also be seen as a way to smoothen the uncertainty introduced by ML, while retaining its good properties of adaptation to each subject. So we conclude that repairing is effective in improving intent detection, and that this can be achieved at the price of spending a few seconds in the beginning of an experiment. Consider again the movie clip attached to this paper to get an idea of the advantage brought over by repairing the models.

As far as the comparison between SMT and PDO is concerned, no statistically significant difference can be noticed, although in the first experiment SMT yields a better SR; the TTRs are similar in the two experiments, which is not surprising since the only difference between the two experiments was in the online testing phase, i.e., *after* the model were built, and on similar datasets. Taking into account that PDO cannot guarantee that the desired condition will be enforced, one should prefer SMT. Of course, we cannot give any indication of how differently SMT and PDO would perform on more complex problems, e.g., with more actions or more observations and/or in true daily-living conditions; but we expect the complexity of SMT to increase dramatically as the problem gets harder — as opposed to that, this should be no problem for PDO. There is indeed a trade-off then, between employing a possibly slower method with a guarantee (SMT) or a possibly faster one which would only work well in practice. Such a trade-off must be analyzed on a case-by-case basis. Moreover, PDO is an easy method to implement.

One remark about the encoding of the wrong model behaviour. We have made the assumption that the straight line ($\overline{A} - \overline{RE}$) encodes, for an arbitrary action A, the enforcement of A with increasing force; this assumption is physiologically justified (see. e.g., [37]–[41]) but, obviously, only to some extent. Extreme flexion or muscle fatigue would seriously invalidate it and the input signals would no longer lie on the straight line. So far we have observed very little of this phenomenon, but it will need to be taken into account in the general case. Actually, given that in RR-RFF the non-linear basis functions are cosines, that is smooth ones, enforcing the safety condition as we have done probably buys the system some extra safety: if the subject even moves slightly away from the straight line, the f_A should be not so different.

Perspectives / Future work

The approach presented is, *per se*, very general, since the logical language enforced in dReal [17] can encode safety of RR-RFF-based myocontrol with respect of numerous interesting properties Σ as presented in Subsection II-B. One example is that of actually forcing a monotonically increasing behaviour for values of $0 \leq t_A \leq 1$. Another even more interesting possibility is that of trying and eliminating *action interference*, another undesired behaviour of myocontrol, in which while trying to perform an action, another one gets unwillingly performed — for example, the wrist unwillingly pronates while the subject only tries to power grasp, which could lead to unwanted effects. To this aim, we should probably try and verify a more complex property (or set of properties) for each straight line.

All in all, we may not as yet claim that the approaches presented in this work are applicable in daily-living scenarios. The TAC test, although online and successfully tested on 18 subjects, is no representative of daily-living activities in which movement artifacts, added weights, fatigue, electrode displacement, etc. usually play a determinant role. Patients must be involved, and possibly interactive learning [10] must be used to enable on-the-fly corrections to the ML model to improve the overall reliability — this is our next planned experiment. Notice, anyway, that at least the PDO approach is easy to implement and, as long as the ML method's performance does not crucially depend on N (as is the case of RR-RFF), won't alter training and prediction times.

Acknowledgments

The authors would like to thank Dr. Bernhard Weber and all members of the USIG (User Studies Interest Group) of the DLR, for the extremely useful discussion on the user study and the related statistics.

References

[1] C. Castellini *et al.*, "Proceedings of the first workshop on peripheral machine interfaces: Going beyond traditional surface electromyography," *Frontiers Neurorobot.*, vol. 8, p. 22, Aug. 2014.

[2] S. M. Engdahl, B. P. Christie, B. Kelly, A. Davis, C. A. Chestek, and D. H. Gates, "Surveying the interest of individuals with upper limb loss in novel prosthetic control techniques," *J. NeuroEng. Rehabil.*, vol. 12, no. 1, p. 53, 2015.

[3] B. N. Perry, C. W. Moran, R. S. Armiger, P. F. Pasquina, J. W. Vandersea, and J. W. Tsao, "Initial clinical evaluation of the modular prosthetic limb," *Frontiers Neurol.*, vol. 9, p. 153, Mar. 2018.

[4] I. Vujaklija, D. Farina, and O. C. Aszmann, "New developments in prosthetic arm systems," *Orthopedic Res. Rev.*, vol. 8, pp. 31–39, Jan. 2016.

[5] R. Merletti, A. Botter, C. Cescon, M. A. Minetto, and T. M. M. Vieira, "Advances in surface EMG: Recent progress in clinical research applications," *Critical Rev. Biomed. Eng.*, vol. 38, no. 4, pp. 347–379, 2011.

[6] N. Jiang, J. Vest-Nielsen, S. Muceli, and D. Farina, "EMG-based simultaneous and proportional estimation of wrist/hand kinematics in uni-lateral trans-radial amputees," *J. Neuroeng. Rehabil.*, vol. 9, no. 1, p. 42, Dec. 2012.

[7] C. Castellini, *Human Robot Hands: Sensorimotor Synergies to Bridge Gap Between Neuroscience Robotic*. New York, NY, USA: Springer, 2016, pp. 171–193.

[8] A. Gijsberts *et al.*, "Stable myoelectric control of a hand prosthesis using non-linear incremental learning," *Frontiers Neurorobot.*, vol. 8, no. 8, Feb. 2014.

[9] I. Strazzulla, M. Nowak, M. Controzzi, C. Cipriani, and C. Castellini, "Online bimanual manipulation using surface electromyography and incremental learning," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 25, no. 3, pp. 227–234, Mar. 2017.

[10] M. Nowak, C. Castellini, and C. Massironi, "Applying radical constructivism to machine learning: A pilot study in assistive robotics," *Constructivist Found.*, vol. 13, no. 2, pp. 250–262, Mar. 2018. [Online]. Available: http://constructivist.info/13/2/250.nowak

[11] J. R. Flanagan, M. K. O. Burstedt, and R. S. Johansson, "Control of fingertip forces in multidigit manipulation," *J. Neurophysiol.*, vol. 81, no. 4, pp. 1706–1717, Apr. 1999.

[12] R. S. Johansson, J. L. Backlin, and M. K. O. Burstedt, "Control of grasp stability during pronation and supination movements," *Experim. Brain Res.*, vol. 128, nos. 1–2, pp. 20–30, Sep. 1999.

[13] G. P. Slota, M. S. Suh, M. L. Latash, and V. M. Zatsiorsky, "Stability control of grasping objects with different locations of center of mass and rotational inertia," *J. Motor Behav.*, vol. 44, no. 3, pp. 169–178, May 2012.

[14] L. Pulina and A. Tacchella, "An abstraction-refinement approach to verification of artificial neural networks," in *Proc. CAV Int. Conf. Comput. Aided Verification*, 2010, pp. 243–257.

[15] L. M. de Moura and N. Bjórner, "Z3: An efficient SMT solver," in *Proc. TACAS Int. Conf. Tools Algorithms Construction Anal. Syst.*, 2008, pp. 337–340.

[16] A. Cimatti, A. Griggio, B. J. Schaafsma, and R. Sebastiani, "The MathSAT5 SMT solver," in *Proc. TACAS Int. Conf. Tools Algorithms Construction Anal. Syst.*, 2013, pp. 93–107.

[17] S. Gao, S. Kong, and E. M. Clarke, "dreal: An SMT solver for nonlinear theories over the reals," in *Proc. CADE Int. Conf. Automated Deduction*, 2013, pp. 208–214.

[18] C. W. Barrett, R. Sebastiani, S. A. Seshia, and C. Tinelli, "Satisfiability modulo theories," in *Handbook Satisfiability*. Amsterdam, The Netherlands: IOS Press, 2009, pp. 825–885.

[19] A. M. Simon, L. J. Hargrove, B. A. Lock, and T. A. Kuiken, "The target achievement control test: Evaluating real-time myoelectric pattern recognition control of a multifunctional upper-limb prostheses," *J. Rehabil. Res. Develop.*, vol. 48, no. 6, pp. 619–628, 2011.

[20] M. Nowak and C. Castellini, "Wrist and grasp myocontrol: Simplifying the training phase," in *Proc. ICORR Int. Conf. Rehabil. Robot.*, 2015, pp. 339–344.

[21] M. Nowak, B. Aretz, and C. Castellini, "Wrist and grasp myocontrol: Online validation in a goal-reaching task," in *Proc. IEEE Int. Symp. Robot Hum. Interact. Commun.*, Aug. 2016, pp. 132–137.

[22] M. Nowak and C. Castellini, "The LET procedure for prosthetic myocontrol: Towards multi-DOF control using single-DOF activations," *PLoS ONE*, vol. 11, no. 9, pp. 1–20, Sep. 2016.

[23] F. Leofante and A. Tacchella, "Learning in physical domains: Mating safety requirements and costly sampling," in *Proc. Int. Conf. Italian Assoc. Artif. Intell.*, 2016, pp. 539–552.

[24] A. J. Smola and B. Schölkopf, "A tutorial on support vector regression," *Statist. Comput.*, vol. 14, no. 3, pp. 199–222, Aug. 2004.

[25] G. Katz, C. W. Barrett, D. L. Dill, K. Julian, and M. J. Kochenderfer, "Reluplex: An efficient SMT solver for verifying deep neural networks," in *Proc. Int. Conf. Comput. Aided Verification*, 2017, pp. 97–117.

[26] X. Huang, M. Kwiatkowska, S. Wang, and M. Wu, "Safety verification of deep neural networks," in *Proc. Int. Conf. Comput. Aided Verification*, 2017, pp. 3–29.

[27] F. Leofante, N. Narodytska, L. Pulina, and A. Tacchella. (2018). "Automated verification of neural networks: Advances, challenges and perspectives." [Online]. Available: https://arxiv.org/abs/1805.09938

[28] D. S. Gonzàlez and C. Castellini, "A realistic implementation of ultrasound imaging as a human-machine interface for upper-limb amputees," *Frontiers Neurorobot.*, vol. 7, P. 17, Oct. 2013.

[29] J. M. Hahne, S. Dóhne, H. J. Hwang, K. R. Müller, and L. C. Parra, "Concurrent adaptation of human and machine improves simultaneous and proportional myoelectric control," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 23, no. 4, pp. 618–627, Jul. 2015.

[30] G. Patel, M. Nowak, and C. Castellini, "Exploiting knowledge composition to improve real-life hand prosthetic control," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 25, no. 7, pp. 967–975, Jul. 2017.

[31] M. A. Powell and N. V. Thakor, "A training strategy for learning pattern recognition control for myoelectric prostheses," *J. Prosthetics Orthotics*, vol. 25, no. 1, pp. 30–41, Jan. 2013.

[32] E. M. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith, "Counterexample-guided abstraction refinement," in *Proc. Int. Conf. Comput. Aided Verification*, 2000, pp. 154–169.

[33] S. Gao, J. Avigad, and E. M. Clarke, "δ-complete decision procedures for satisfiability over the reals," in *Proc. Int. Joint Conf. Automated Reasoning*, 2012, pp. 286–300.

[34] A. Rahimi and B. Recht, "Uniform approximation of functions with random bases," in *Proc. 46th Annu. Allerton Conf. Commun., Control, Comput.*, Sep. 2008, pp. 555–561.

[35] S. Jacobsen, *Control Systems for Artificial Arms*. Cambridge, MA, USA: MIT Press, 1973. [Online]. Available: https://books.google.de/books?id=a7VwnQEACAAJ

[36] S. C. Jacobsen and R. W. Mann, "Control systems for articial arms," in *Proc. IEEE Conf. Syst., Man., Cybern.*, Sep. 1973, pp. 298–303.

[37] F. J. Valero-Cuevas, "Predictive modulation of muscle coordination pattern magnitude scales fingertip force magnitude over the voluntary range," *J. Neurophysiol.*, vol. 83, no. 3, pp. 1469–1479, 2000.

[38] B. Poston, A. D.-D. Santos, M. Jesunathadas, T. M. Hamm, and M. Santello, "Force-independent distribution of correlated neural inputs to hand muscles during three-digit grasping," *J. Neurophysiol.*, vol. 104, no. 2, pp. 1141–1154, 2010.

[39] A. de Rugy, G. E. Loeb, and T. J. Carroll, "Muscle coordination is habitual rather than optimal," *J. Neurosci.*, vol. 32, no. 21, pp. 7384–7391, May 2012.

[40] J. He, D. Zhang, X. Sheng, S. Li, and X. Zhu, "Invariant surface EMG feature against varying contraction level for myoelectric control based on muscle coordination," *IEEE J. Biomed. Health Inform.*, vol. 19, no. 3, pp. 874–882, May 2015.

[41] A. Al-Timemy, R. Khushaba, G. Bugmann, and J. Escudero, "Improving the performance against force variation of EMG controlled multifunctional upper-limb prostheses for transradial amputees," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 24, no. 6, pp. 650–661, Jun. 2016.