Gaussian Process Autoregression for Simultaneous Proportional Multi-Modal Prosthetic Control With Natural Hand Kinematics

Michele Xiloyannis, Constantinos Gavriel, Andreas A. C. Thomik, and A. Aldo Faisal

Abstract— Matching the dexterity, versatility, and robustness of the human hand is still an unachieved goal in bionics, robotics, and neural engineering. A major limitation for hand prosthetics lies in the challenges of reliably decoding user intention from muscle signals when controlling complex robotic hands. Most of the commercially available prosthetic hands use muscle-related signals to decode a finite number of predefined motions and some offer proportional control of open/close movements of the whole hand. Here, in contrast, we aim to offer users flexible control of individual joints of their artificial hand. We propose a novel framework for decoding neural information that enables a user to independently control 11 joints of the hand in a continuous manner—much like we control our natural hands. Toward this end, we instructed six able-bodied subjects to perform everyday object manipulation tasks combining both dynamic, free movements (e.g., grasping) and isometric force tasks (e.g., squeezing). We recorded the electromyographic and mechanomyographic activities of five extrinsic muscles of the hand in the forearm, while simultaneously monitoring 11 joints of hand and fingers using a sensorized data glove that tracked the joints of the hand. Instead of learning just a direct mapping from current muscle activity to intended hand movement, we formulated a novel autoregressive approach that combines the context of previous hand movements with instantaneous muscle activity to predict future hand movements. Specifically, we evaluated a linear vector autoregressive moving average model with exogenous inputs and a novel Gaussian process ($\mathcal{GP}$) autoregressive framework to learn the continuous mapping from hand joint dynamics and muscle activity to decode intended hand movement. Our $\mathcal{GP}$ approach achieves high levels of performance (RMSE of 8°/s and $\rho = 0.79$). Crucially, we use a small set of sensors that allows us to control a larger set of independently actuated degrees of freedom of a hand. This novel undersensored control is enabled through the combination of nonlinear autoregressive continuous mapping between muscle activity and joint angles. The system evaluates the muscle signals in the context of previous natural hand movements. This enables us to resolve ambiguities in situations, where muscle signals alone cannot determine the correct action as we evaluate the muscle signals in their context of natural hand movements. $\mathcal{GP}$ autoregression is a particularly powerful approach which makes not only a prediction based on the context but also represents the associated uncertainty of its predictions, thus enabling the novel notion of risk-based control in neuroprosthetics. Our results suggest that $\mathcal{GP}$ autoregressive approaches with exogenous inputs lend themselves for natural, intuitive, and continuous control in neurotechnology, with the particular focus on prosthetic restoration of natural limb function, where high dexterity is required for complex movements.

Index Terms— Neuroprosthetics, robotic hand, decoding, autoregression, EMG, MMG, Gaussian process, proportional control, neurotechnology.

M. Xiloyannis, C. Gavriel, and A. A. C. Thomik are with the Brain and Behaviour Laboratory, Department of Bioengineering and Department of Computing, Imperial College London, London SW7 2AZ, U.K.

A. A. Faisal is with the Brain and Behaviour Laboratory, Department of Bioengineering and Department of Computing, Imperial College London, London SW7 2AZ, U.K., also with the Data Science Institute, Imperial College London, London SW7 2AZ, U.K., and also with the MRC London Institute of Medical Sciences, London W12 0NN, U.K. (e-mail: a.faisal@imperial.ac.uk).

I. INTRODUCTION

THE ability to dextrously control our hand is key to our evolution and at the centre of almost all skilled object interaction in daily life [1], [2]. It is thus not surprising that loss of hand functionality is severely debilitating and considered more serious than the loss of lower limb function [3]. Currently available commercial devices are typically constrained by limited functionality in terms of degrees of freedom, range of motion, sensory feedback and intuitive control [4]–[6]. Despite the mechanical complexity of the human hand, the bottleneck of modern prostheses lies not only in improving mechatronic design [7], but also in the human-machine interface required for robustly translating the user's intention into a suitable robotic action. This makes prosthetics cumbersome and limits their perceived effectiveness in daily life – limiting users' ability to accept the prostheses as part of their body vs using it as a tool [8]. This explains why despite significant improvements of prosthesis acceptance achieved in the last three decades [6], a main reason for their rejection and

low adoption involves dissatisfaction with functionality and control [9]. Low adoption rates are the result of users switching after initial use of bionic prosthetics to mechanical or even cosmetic prostheses.

In upper arm prosthetics one typically uses sensors to record muscle signals from the residual limb or nerve. These signals are processed and translated into commands for a set of independently controllable actuators. Originally [10], prostheses were controlled using on-off control (or bang-bang control) which is a simple, yet reliable way of actuating a prosthetic device with elementary commands (open vs close hand). With prosthetic mechatronics becoming more sophisticated, and enabling us to actuate multiple individual joints per finger, more sophisticated control approaches were sought. A common compromise was to combine discrete selection of one hand movement across multiple possible ones, with then a direct proportional control of that motion.

We believe that the current challenges amputees face in accepting and retaining use of robotic prostheses can be greatly improved through interfaces that allow continuous, simultaneous and intuitive control of the artificial hand at the level of individual finger joints. Continuous prosthetic control, also referred to as proportional control, does not imply that the relationship between control input and controller output is strictly proportional in the mathematical sense, only that it must be essentially continuous – in fact we explore both linear and nonlinear methods here. To be clear, the term proportional relates to the decoders – the forward path of control, we do not want to confuse it with the term "proportional" in the feedback control system sense (e.g. the "P" in PID control).

Continuous control can allow nuanced movement of the prosthetic hand, just as in the human hand and even enables users to retain or develop their personal movement flourish. To achieve proportional control, one must move away from the typical discrete classification approach and try to map residual muscle activity to joint movements in a continuous fashion. This approach has been investigated in wrist kinematics [11], [12] and hand dynamics [13] in test conditions. The rationale of these studies was to learn the physiological mapping between EMG activity and joint dynamics/kinematics [14] from a set of training signals. It has been shown that modelling this mapping in a linear fashion can yield good prediction accuracies, however, more sophisticated non-parametric statistical methods were suggested to outperform linear models (but incur a higher computational cost) [12]. Despite this progress, inferring the intended movement of the hand and fingers directly from these sensors in the daily setting is still an open challenge. Continuous control eliminates delays of having to switch between a series of discrete grasps or poses. However, switching between multiple discrete grasps or poses has been helpful to users, as it allows them to explicitly configure the hand to the context. This context functionality has to be replaced in some way – we will demonstrate here a suitable approach that operates automatically.

While myoelectric, i.e. electromyography (EMG) based control has long been established in prosthetic control for 70 years [10], simultaneous proportional myoelectric control
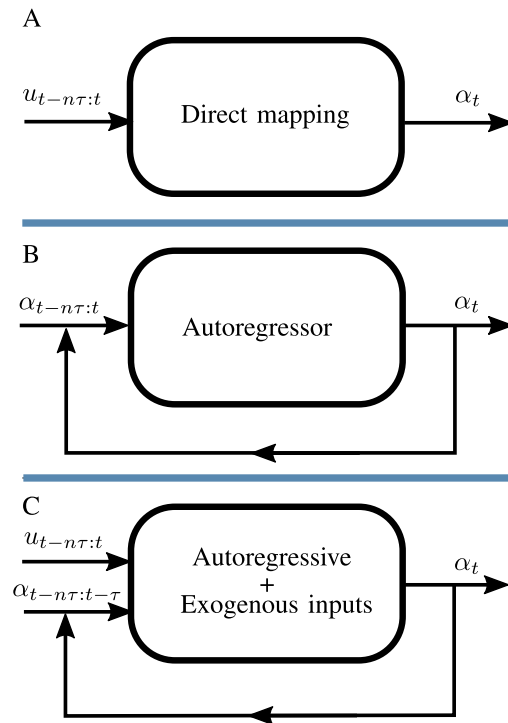

Fig. 1. Overview between different approaches of time series regression as applied to the problem of predicting changes to finger joint angles (α) from muscle activity (u). (A) Direct mapping takes input from present and past muscle activity to decode joint action, (B) Autoregressive relationship takes input from n past joint actions to decode joint action and (C) the Autoregressive model with Exogenous Inputs integrates present and past muscle activity and past joint movements to predict joint action.

systems are far more recent, with studies being reported only in the last 5 years. Historically prosthesis control has moved from mapping continuous myoelectric signals to discrete on-off control, to more sophisticated multi-way classifiers, to combining discrete classifiers of hand movement modes and their simple linear control to the full joint mapping from muscle activity to hand movement [15].

It is worth to structure existing strategies for prosthetic control along the nature of the decoding problem, which we can structure broadly into 3 approaches (see correspondingly Fig. 1.A, B, C). Most simultaneous proportional control approaches map some time window of neuromuscular data to the output of a hand movement. This way of approaching the decoding problem is exemplified by our overview graphic in Fig. 1.A. One is predicting changes to finger joint angles (α) at time t from muscle activity (u) by finding a direct mapping of the form $f(u) = \alpha$, where u may involve a time window or multiple time windows of muscle activity. Most studies focus on "learning" such direct mapping functions from data, typically using single [16] or hierarchically organised [17], [18] neural networks. Moreover, dimensionality reduction and projection mapping methods have been used to harness opportunities arising from high density EMG recordings with many sensors [19], [20], where low-dimensional representations of the muscle signals are optimised for maximal discriminability between movements types. We note that all these prosthetic control frameworks are normally

oversensored, i.e. the ratio of independently controllable robotic degrees of freedom to sensor readings is typically smaller than one.

While the hand has many more independent degrees of actuated freedom than current prosthetics, the question arises if the natural hand uses these degrees of freedom in normal daily life fully independently? In previous work, we showed that both the dominant hand [21] (see also [22], [23]) and the non-dominant hand [24] had joints that were highly correlated – when captured in daily life activities. Thus, our over 20 finger joints are controlled by our brains so as to be effectively embedded on a lower dimensional manifold: four to five dimensions (principal components) explained 80%-90% of the variability in the movement data. Moreover, we have previously found that within the first 1000 ms of action initiation (from a neutral pose) the hand shape configures itself towards different tasks that occur many seconds away. This enabled us to reliably predict the action intention from the initial small changes in hand pose occurring in the first few hundred milliseconds of movement [21]. We then demonstrated that a probabilistic latent variable model can be used to infer the movements of missing limbs from the movements of the residual limbs to control, for example, a prosthetic half-hand in real-time [21], [24]. This previous work thus demonstrated that correlation patterns across joints, i.e. the spatial configuration of the hand, is highly structured in natural tasks and can be harnessed for prosthetic control.

If hand joint configurations are correlated across the hand, what about the temporal correlation structure of natural hand movements? We previously showed using 1. statistical time series segmentation [25], 2. switching motion primitive approaches [26], and 3. sparse coding (compressed sensing) methods [24] that the recent past of ongoing hand movements is highly predictive of the continuation of a movement. Moreover, these compact temporal representations of movement coincided with fundamental grasps and we were able to directly map these to control robotic hand grasps [27]. Thus, natural hand movements enable prediction of future hand movements from past hand movements alone without needing neuromuscular data. We can mathematically interpret this second approach as a problem of time series prediction, where the history of the movements reflects the inherent dynamics of the time series, i.e. the most likely motions of the hand in our case. This inherent dynamics is mathematically captured by continuous autoregressive models, i.e. $f(\alpha_{t-n\tau}, \ldots, \alpha_{t-\tau}) = \alpha_t$ (see Fig. 1.B). Crucially, the complex coordination of multiple joints is simplified, as their correlated movement is the result of the inherent dynamics of the current action. This second approach is simultaneous and proportional by design and can map the movement of a few joints to the movements of the whole hand.

The challenge to the second approach is, of course, that while ongoing hand movements (e.g. reaching for a cup) are in themselves correlated and allow us to predict reasonably well how a grasp is completed, there will always be events within a series of actions where the user makes a decision about a new action that cannot be predicted from the statistics of the past (e.g. after grabbing a cup one may grab a tea pot or a sandwich). This requires explicit control inputs from the user to disambiguate the next action. These control inputs are typically muscle activity patterns in the residual limb [10] or result from retargeted residual nerves [28].

We explore here a combination of the traditional first approach and the unconventional second approach – a third way. We propose a novel framework (see Fig. 2) for information extraction from muscle-related signals which allows continuous, intuitive and naturalistic control of upper limb prostheses. Our new approach combines the non-linear spatio-temporal correlation structure of natural hand movements with muscle-driven control signals to exploit the best of both worlds. Our framework aims not only to learn the relation between muscle activity and hand kinematics, but should also exploit the natural autocorrelation of everyday hand movements to achieve a higher performance than pure control signal decoding strategies (first approach). Thus, we are proposing implementing this framework using tractable non-linear probabilistic models of autoregressive processes with exogenous inputs u (muscle signals) of the nature: $f(\alpha_{t-n\tau}, \ldots, \alpha_{t-\tau}, u_{t-n\tau}, \ldots, u_{t-\tau}) = \alpha_t$ (see Fig. 1.C).

We advance the hypothesis that, in contrast to previous applications, with this approach we should be able to use fewer sensors than independently actuated hand joints and yet achieve high decoding performance across all these joints. Specifically we will test if we are able to decode reliably 11 finger and thumb joint movements from 5 EMG sensors on the forearm.

II. METHODS

A. Experimental Setup

Six healthy male subjects, aged 23-28, gave written informed consent prior to participation in the study. Participants were seated on a chair fitted with a custom elbow rest support. Their elbow rested on a flat surface, the forearm was horizontal and the hand was positioned so that the axis of abduction/adduction of the wrist was parallel to the surface. In order to prevent activation of the extensor *carpi radialis longus* (responsible for the abduction of the wrist) during the experiment, the wrist was also rested on a customised support.

Subjects were instructed to grasp and manipulate 7 objects (example hand postures are rendered in Fig. 4), each object 12 times during two different sessions involving 6 repetitions each. The 7 object manipulations were: grasping a ball, grabbing a book, drinking from a bottle, grasping a door handle, operating a computer mouse, holding a mobile phone and using a pen. To allow for comparison with other studies, we used objects and interactions from a recent taxonomy of hand postures [30].

We adopted our low-cost system for MMG signal acquisition [29] and merged it with myoelectrical signals recorded using a commercially available EMG amplifier [Brain Vision LLC, Morrisville, NC]. We paired the signal recorded from the muscle activity with kinematic measurements of finger movements during everyday tasks. We used this data to train a linear and a non-parametric regression model capable of continuously predicting the angular velocity of 11 joints of
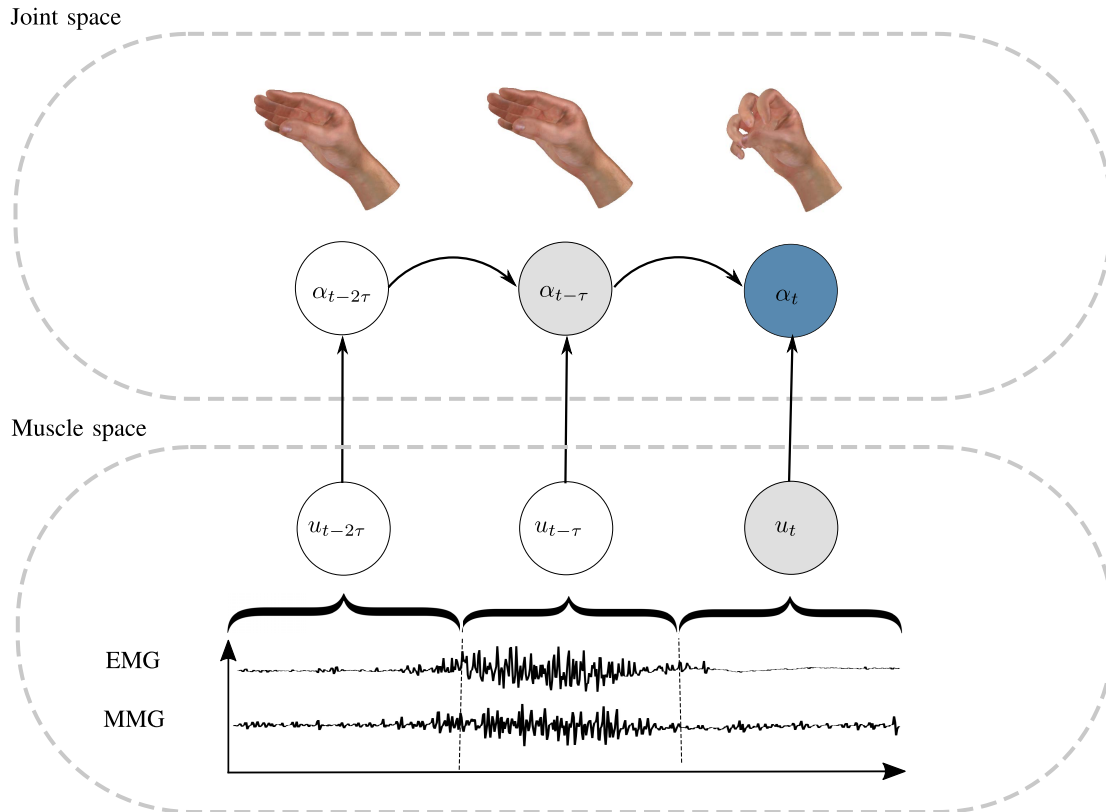
Fig. 2. Illustration of the probabilistic framework (graphical model) used to model the dynamic integration of muscle-activity and hand-joint kinematics. Continuous decoding of muscle signals to predict (prosthetic) hand joints velocities treats the joint velocities as a weakly stationary autoregressive process (joint space). See Tab. II for the details of the system configuration.

the fingers of the hand given the muscle activity and an initial condition.

At the beginning of each trial, the object name and time stamp were sent to a serial port via the ActiChamp amplifier, these were subsequently used for data segmentation. Subjects were also asked to apply their maximum extension and flexion force of their fingers at the beginning and at the end of the experiment. The root mean square (RMS) value of these segments was used to estimate the Maximum Voluntary Contraction (MVC), necessary for normalisation of the signal in the pre-processing phase.

We recorded muscle activity with 5 EMG and 4 MMG channels positioned at the bellies of muscles (* denotes both EMG and MMG measurements) from the posterior (1: *Extensor Digitorum Communis** (EDC), 2: *Extensor Digiti Minimi** (EDM), 3: *Extensor Carpi Radialis* (ECR)) and anterior (1: *Flexor Digitorum Superficialis** (FDS), 2: *Palmaris Longus** (PL)) compartment of the forearm as shown in Fig. 3.(a-b). The channel sites were located by palpating the respective muscles. The MMG sound chambers were positioned between the two electrodes of the same EMG channel (see Fig. 3.(c-f)). A summary of the monitored muscles is shown in Table I. Simultaneously, we measured finger movements using a CyberGlove I [CyberGlove Systems LLC, San Jose, CA]. This is a sensorised glove which measures the angle of all joints within the hand from resistive changes in stretch sensitive sensors placed over each joint. The hand

TABLE I
MONITORED MUSCLES

Function	Muscle	Abbreviation
Extension	1.Extensor Digitorum Communis	EDC
	2.Extensor Digiti Minimi	EDM
	3.Extensor Carpi Radialis	ECR
Flexion	1.Flexor Digitorum Superficialis	FDS
	2.Palmaris Longus	PL

TABLE II
MONITORED HAND JOINTS

Finger	Joint	Abbreviation
Thumb	Carpo-Metacarpal	T_{CMC}
	Metacarpal-Phalangeal	T_{MCP}
	Interphalangeal	T_{IP}
Index	Metacarpal-Phalangeal	I_{MCP}
	Interphalangeal	I_{IP}
Middle	Metacarpal-Phalangeal	M_{MCP}
	Interphalangeal	M_{IP}
Ring	Metacarpal-Phalangeal	R_{MCP}
	Interphalangeal	R_{IP}
Little	Metacarpal-Phalangeal	L_{MCP}
	Interphalangeal	L_{IP}

joints monitored for the purpose of this study are listed in Table II.

B. Data Processing

Prior to further analysis, we preprocessed our data in the following fashion using Matlab (MathWorks Inc., Natick, MA):
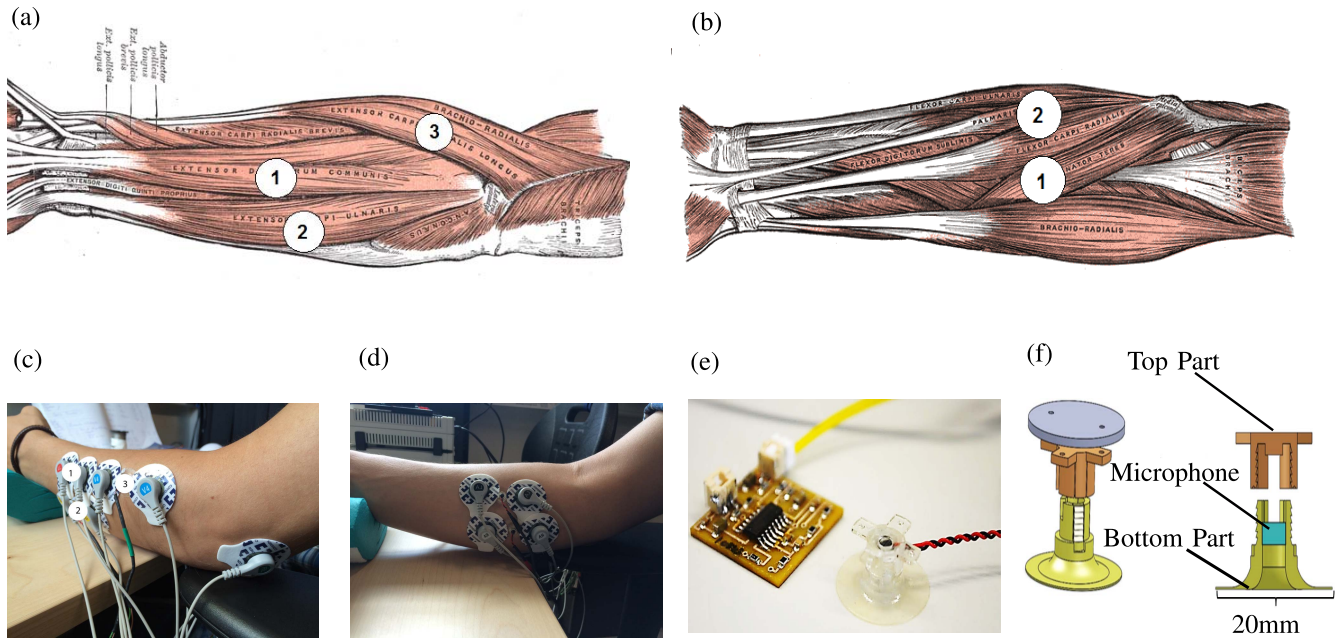
Fig. 3. Monitored muscles in the anterior and posterior compartment of the forearm and sensor placement details. (a + c) Posterior compartment of the forearm and positioning of the electrodes. Each MMG microphone was located between the two corresponding EMG electrodes (in a bipolar configuration) of the same muscle. Channels on the muscles are numbered as in the text. (b + d) Anterior compartment of the forearm and positioning of the EMG and MMG channels. Channels on the muscles are numbered as in the text. (e) Printed circuit board for amplification, rectification and filtering of the MMG signal generated by muscle during contraction. (f) The mechanical signal is transduced by a miniature microphone embodied in a 3D printed resonance chamber). Details of the MMG sensor design are described in [29]. Illustration used in a + b adapted by anonymous author from Wikimedia Commons.

1) Segmentation of the data of different tasks using the time stamps. The resulting segments contained data from different trials of the same task.
2) High-Pass filtering of the EMG ($f_c = 10$Hz) and of the MMG signal ($f_c = 8$Hz) using a 5^{th} order Butterworth filter to remove movement artefacts.
3) Low-Pass filtering of the signals ($f_c = 150$Hz) using a 5^{th} order Butterworth filter to remove high frequency noise.
4) A 50 Hz comb filter to remove power-line interference.
5) Normalisation by Maximum Voluntary Contraction (MVC). This constrains the input signals for the regressor to the interval [0, 1].
6) Downsampling of EMG and MMG signals to the Cyberglove rate, i.e. 138Hz. This procedure was necessary to align the glove data with the muscle-related signals.

The data acquired from the CyberGlove were smoothened using a first-order Savitzky-Golay filter with a running window of 23 consecutive samples to remove the discontinuities induced by the A/D converter. The window width was chosen to lie within the peak of the hand movement autocorrelation function width of half maximum ($\bar{\tau}_{WHM} = 362$ ms, see also Fig. 4). This ensures that smoothing only involved highly autocorrelated samples and is little affecting the overall dynamical changes of the system.

We then differentiated the position data to obtain angular velocity ($\dot{\alpha}$) which is statistically more stationary in natural hand movements, has a more Gaussian distribution and is known to be closely related to motor commands [32] of each joint's flexion/extension movement.
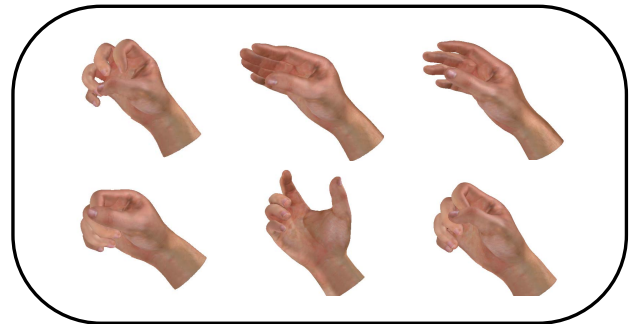

Fig. 4. Example hand poses of subject hand configuration during designated dynamic and static object-manipulation tasks involing objects of different shapes and sizes. Example reconstruction of instantaneous hand pose from top left to bottom right: grasping a ball, grabbing a book, drinking from a bottle, grasping a door handle, operating a computer mouse, holding a mobile phone and using a pen. Images created by rendering the Cyberglove finger configuration joint configuration data using the LibHand library [31].

C. Generalised Time Series Prediction Model for Prosthetics

We implemented a model for continuous prediction of the angular joint velocities of the flexion/extension joints of the hand by expressing the angular velocity of the joints at time t, $\dot{\alpha}_t$, as a function of the angular velocity at time $t - \tau$, $\dot{\alpha}_{t-\tau}$, and the current muscle-activation vector $\mathbf{u_t}$. We are thus treating the joint angular velocity as a wide-sense stationary autoregressive sequence that is influenced by external factors, namely, in our case, the forearm muscle activation pattern. The angular velocity of the joints at time t can thus be

expressed as:

$$\dot{\alpha}_t = f(\dot{\alpha}_{t-\tau}, \mathbf{u_t}) + \epsilon \qquad (1)$$

We seek to learn the function f under the assumption of Gaussian distributed noise $\epsilon \sim \mathcal{N}(0, \sigma_n)$.

To compare linear and non-linear approaches we can rephrase the goal of decoding hand motion from muscle activity as a probabilistic inference, as illustrated in Fig. 2. The current hand configuration is modelled as a random variable θ_t, a 11-dimensional state vector of the 11 independently controlled finger joint angles. The variable θ_t is conditionally dependent on the previous time steps $t - \tau$ (where τ is the sample interval of the system), and the current muscle activation pattern $\mathbf{u}_t$. The activation pattern $\mathbf{u}_t$ is here a $5 + 4$ dimensional vector of muscle activity recorded by EMG and MMG sensors on the forearm (muscle space). We interpret this in our graphical model as the sequence of finger joint motions being caused by muscle activity patterns.

The time-lag τ plays a fundamental role in determining the accuracy and the computational demand of the algorithm. Its value is strictly related to the dynamics of the process that we are trying to describe. In order to identify a reasonable range of time-lags to use in our model, we analysed the autocorrelation function of the angular velocity $\dot{\alpha}$ on the data collected using the Cyberglove:

$$R_{\dot{\alpha}\dot{\alpha}}(\tau) = \sum_{\forall n \in N} \dot{\alpha}(n)\dot{\alpha}(n - \tau). \qquad (2)$$

We found an average, over joints and subjects, Width of Half-Maximum (WHM) of the autocorrelation function at $\bar{\tau}_{\text{WHM}} = 362$ ms, as shown in Fig. 5. To estimate the impact of the time–lag on the model's performance, we trained a Gaussian Process ($\mathcal{GP}$) and a Vector Auto–Regressive Moving Average model with Exogenous Inputs (VARMAX) with time-lags $\tau = 87$ ms, 175 ms, 350 ms, 500 ms. Both models were trained and tested using a Leave-One-Out cross-validation on the 6 tasks shown in Fig. 4.

D. VARMAX Model

We used a VARMAX model to learn a linear version of the function f in Equation 1. This model makes the assumption of linearity between the Auto–Regressive (AR) terms, a Moving Average (MA) term describing the contributions of the innovations ϵ_i and the exogenous inputs $\mathbf{u}$:

$$\dot{\alpha}_t = c + \underbrace{b\mathbf{u}_t}_{\text{EX Inputs}} + \epsilon_t + \underbrace{\sum_{i=1}^{p} \varphi_i \dot{\alpha}_{t-i}}_{\text{AR Contribution}} + \underbrace{\sum_{j=1}^{q} \theta_i \epsilon_{t-i}}_{\text{Innovations}}. \qquad (3)$$

We used the *Dickey-Fuller test* [33] to characterise the stationarity of the time-series and verified that the MA and AR parts of the model were *invertible* and *stable*, respectively. The parameters $\Delta = (c, b, \varphi, \theta)$ were fitted to the training data using Maximum Likelihood (ML) estimate.

We performed the forecasting stage by iteratively predicting the state at a time-lag τ ahead and using it as a new input for the next iteration. The starting point was set to $\dot{\alpha}|_{t=0} = 0$ for each joint.
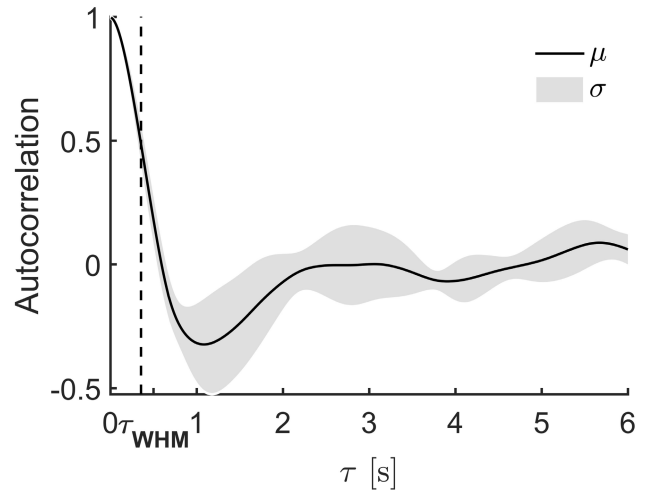

Fig. 5. Human hand movements while performing everyday grasping tasks have a strong spatio-temporal correlation structure as revealed by autocorrelation analysis. The plot shows the grand-average computed over subjects and joints of the hand of the autocorrelation function $R_{\dot{\alpha}\dot{\alpha}}(\tau)$ of angular joint velocities, with one standard deviation σ (grey area) around the mean μ. The mean autocorrelation function shows a width of half maximum at $\bar{\tau}_{\text{WHM}} = 362$ms.

E. Gaussian Process Autoregression With Exogenous Inputs

The VARMAX approach learned a linear approximation to the function f in Equation 1 and we therefore investigate next a non-linear version, which may perform better. To this end, we used a Gaussian Process or $\mathcal{GP}$ model, which is formally equivalent to an artificial neural network with infinitely many neurons in the hidden layer [34]. One can think of a Gaussian Process as defining a distribution over functions: instead of producing single prediction values of f, the $\mathcal{GP}$ model produces a full predictive distribution $p(f)$. Thus, we can not only choose how to evaluate $p(f)$, e.g. as mean, mode or peak of the probability distribution of possible outcomes, but also obtain an implicitly measure of uncertainty that the model has about its own prediction. We can therefore define a loss function over the predictive distribution, and use Bayesian decision theory to perform decoding in an optimal way – so as to minimise the risk of wrongly decoding a hand movement, for example when the signal is ambiguous as to the intended motion.

In $\mathcal{GP}$s, the function f is treated as a latent variable, again under the assumption of normally distributed noise ϵ, and the model returns a process, i.e. a distributions over functions of the form:

$$f(\mathbf{x}) \sim \mathcal{GP}(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')) \qquad (4)$$

fully defined by a mean $m(\mathbf{x})$ and a covariance function $k(\mathbf{x}, \mathbf{x}')$, where $\mathbf{x}$ is the input vector, which we defined as $\mathbf{x} = \{\dot{\alpha}_{t-\tau}, \mathbf{u_t}\}$ (see also Fig. 6).

This Bayesian approach to regression yields to a posterior distribution over the non-linear regression function f. We computed the posterior by combining the likelihood function $P(\mathcal{D}_{\dot{\alpha}}|f, \theta)$ (with θ being a vector of hyperparameters) of the sensory observations $\mathbf{x}$ and our prior over
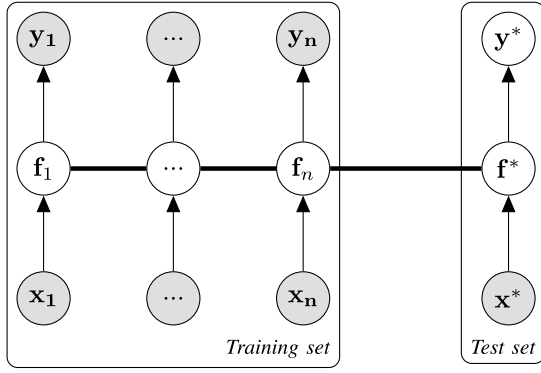
Fig. 6. Probabilistic graphical model of our Gaussian Process decoder. Grey circles indicate observable variables and empty circles represent variables that have to be inferred. In our study × are muscle signals and y joint angles.

possible regression functions $P(f)$. Assuming that the noise of our observations is independent and identically distributed, the posterior distribution over the regression functions has the form:

$$P(f|\mathcal{D}_{\dot{\alpha}}, \theta) \propto \underbrace{P(f|0, K(\mathbf{X}, \mathbf{X}|\theta))}_{\text{Prior}} \underbrace{\prod_{n=1}^{N} p(y_n|f_n, \theta)}_{\text{Likelihood}} \quad (5)$$

where the $\mathcal{GP}$ prior has zero mean and a covariance matrix $K(\mathbf{X}, \mathbf{X}|\theta)$) derived from the covariance function, $k(x, x'|\theta)$.

We encoded the assumption of linearity in the covariance function, using a composite linear kernel with additive noise of the form:

$$k(\mathbf{x}, \mathbf{x}'|\Theta) = \underbrace{\sigma_f \exp\left(\frac{1}{2}\frac{(\mathbf{x}-\mathbf{x}')^2}{\ell^2}\right)}_{\text{Noise}} + \underbrace{\sum_{d=1}^{D} \sigma_l^2 \mathbf{x}\mathbf{x}'}_{\text{Linearity}} \quad (6)$$

with D the dimensionality of the input vector and $\Theta = \{\sigma_f, \ell, \sigma_l\}$ a vector of parameters which defines the properties of the covariance function. The mean function was chosen to be a constant zero function. Fitting of the hyperparameters Θ was done by minimizing the negative log-marginal likelihood by means of a conjugate gradient numerical method [35].

The forecasting stage was done by initialising the state at time zero, $\dot{\alpha}|_{t=0} = 0$ and iteratively predicting the state at a time-lag τ ahead, using a value drawn from the returned Gaussian distribution as an input for the next iteration. A computational limitation of $\mathcal{GP}$ is that their computational demand and memory requirements grow with as $O(n^3)$ and $O(n^2)$ respectively, being n the number of training cases, making their actual implementation feasible only for small data sets. Because of the high computational demand of $\mathcal{GP}$s, we used the Fully Independent Training Conditional (FITC) approximation [36] to solve the inference phase. In brief: To overcome the computational limitations we can introduce new latent variables, known as pseudo-inputs, which do not appear in the predictive distribution, but do significantly affect the goodness of fit. Using these methods, the relation between the training latent function values $\mathbf{f}$ and the test-cases $\mathbf{f}^*$ can be
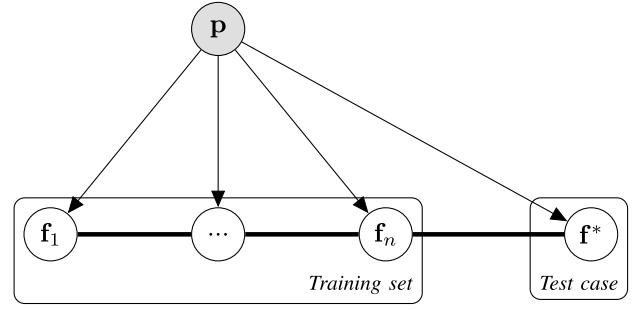

Fig. 7. Graphical model of the relation between the pseudo-inputs p, the training latent functions values $\mathbf{f} = [f_1, \ldots, f_n]^T$ and the test function value $\mathbf{f}^*$.
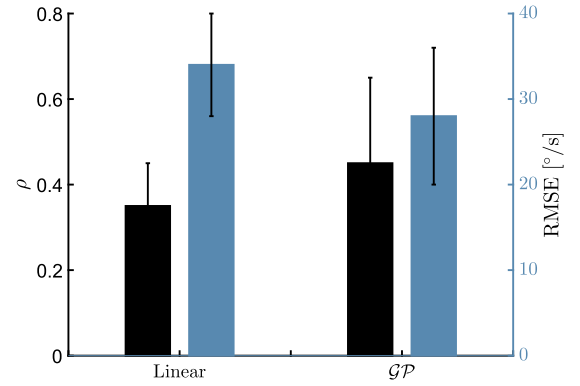

Fig. 8. Direct mapping evaluation: Performance of a direct linear decoder and the $\mathcal{GP}$ model in a muscle-to-joint mapping without autoregressive components, i.e. not taking into account history joint movements. The bar graph shows the average, over joints and subject, of the correlation coefficient ρ (black bar in left axis) and RMSE (blue bar on right axis), between the actual and predicted angular joint velocities. This configuration was evaluated to allow comparison with previous studies that do not use autoregressive approaches (see Table III).

thus formalised:

$$p(\mathbf{f}, \mathbf{f}^*) = \int p(\mathbf{f}, \mathbf{f}^*|\mathbf{u})p(\mathbf{p})d\mathbf{p} \approx \int p(\mathbf{f}|\mathbf{p})p(\mathbf{f}^*|\mathbf{p})p(\mathbf{p})d\mathbf{p}. \quad (7)$$

This equation is represented in the graphical model in Fig. 7. The computational demand of the model now scales with $O(nm^2)$.

We chose the number of pseudo-inputs to be $N/150$, where N is the number of training inputs. The hyperparameters Θ where initialised randomly. K-means clustering of the training input space was used to initialise the position of pseudo-inputs. We then optimised both values by minimising the negative log-marginal likelihood.

F. Evaluation

Using data-driven methods to machine learn a control framework from experimental data, e.g. during calibration of the prosthetic to the end-user, requires a number of caveats to ensure a fair and clean comparison of methods, and to predict how well a given technology will generalise to individual patients that are not part of the research and development process. We highlight here a number of challenges and caveats that we encountered in the field and how we addressed them.
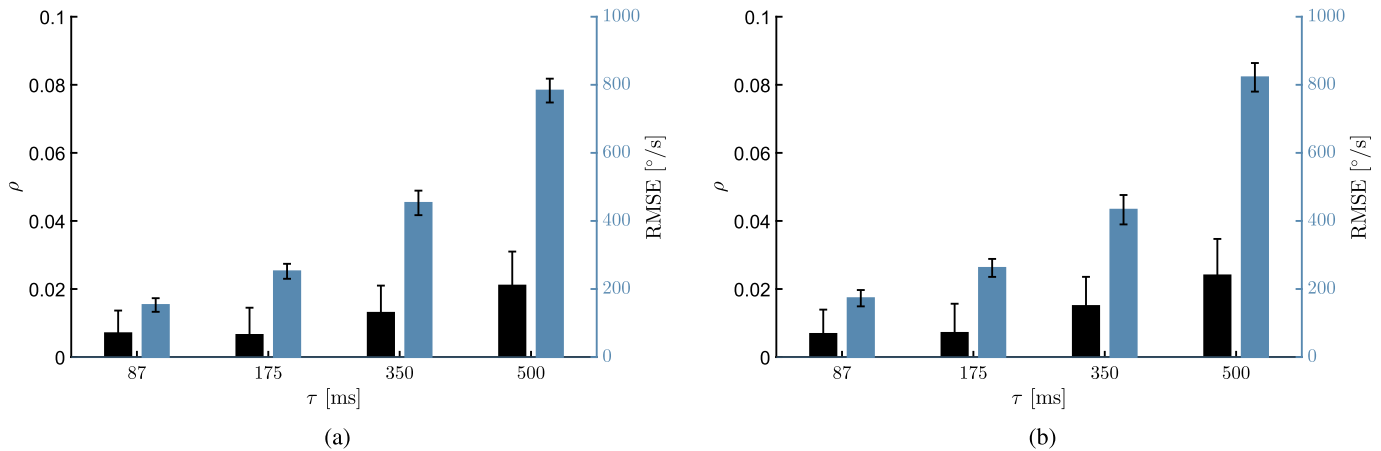
Fig. 9.   Evaluation of the performance of the purely linear (a) and non-linear (b) autoregressive model for varying time-lag τ. These models do not consider muscle activity, and predict the intended hand movement from the context of past hand dynamics. All plots display the average, over joints and subject, of the correlation coefficient ρ (black bars, left y-axis) and RMSE (blue bars, right y-axis), between the actual and predicted angular joint velocities, obtained performing a Leave-One-Out cross-validation over tasks. (a) Performance of the VARMAX model, for varying time-lags τ, with no exogenous inputs. In this case the hand configuration is inferred, at each step, only from its previous state and the muscle activity is ignored. The performance is, as expected, very poor (note the scale) as no muscle activity is used. (b) Performance of the $\mathcal{GP}$ model in the same configuration for varying time-lags τ. Error bars show mean $\pm$ standard deviation (SD).
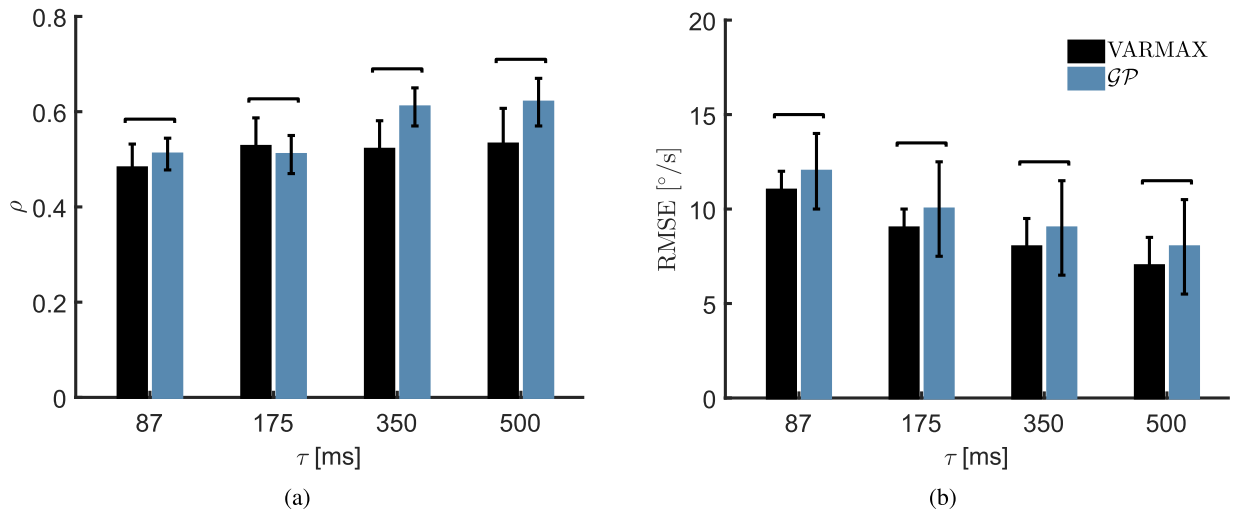

Fig. 10.   Comparison of the performance of the linear (black bars) and $\mathcal{GP}$ (non-linear, blue bars) (b) autoregressive models that incorporate muscle activity as exogenous inputs. The VARMAX (linear) and $\mathcal{GP}$ (non-linear) model use the combined EMG and MMG signals as exogenous inputs as a function of time-lag τ. All plots display the average, over joints and subject, of the correlation coefficient (a) and RMSE (b) between the actual and predicted angular joint velocities. Obtained performing a Leave-One-Out cross-validation over tasks. The performance is about one order of magnitude better in correlation and almost two order of magnitude lower than the purely autoregressive dynamics (see previous figure). No significant difference was found between the VARMAX and $\mathcal{GP}$ (p-value of a two-sample t-test > 0.05). Bars show mean $\pm$ standard deviation (SD).

1) Target Quantities of the Decoder: In prosthetics, a number of target quantities to have been used to determine what the decoder should control. This includes the configuration of the hand, i.e. joint angles of the prosthetic, or torques produced by the motor. The challenge with joint angles is that it is difficult to generalise results obtained from joint angle data using isometric and dynamic movements of the hand, e.g. EMG patterns will look very different for isometric contractions (e.g. holding an object) versus an unrestricted closing motion of the hand where there is no resistance. Torque-based approaches are difficult to implement, because it is complicated to measure force-related data (e.g. end-point forces) in the hand during natural tasks. Therefore, we used the angular joint velocity,

which mechanically speaking captures the impulse of the fingers, and operates both across isometric and dynamic movements. This, combined with the autoregressive approach, gives the system also flexibility to learn whether to tune itself to more force-like signals (by setting the autoregressive weights with different signs, so as to approximate differentiation by finite differences) or position-like signals (by setting the autoregressive weight to positive weights, so as to approximate integration by finite sums).

2) Separation of Training and Test Data: There are a number of methods that can be used to evaluate, train and test the performance of prosthetic controllers that are machine learned from experimental data. Normally, each controller is "fitted"
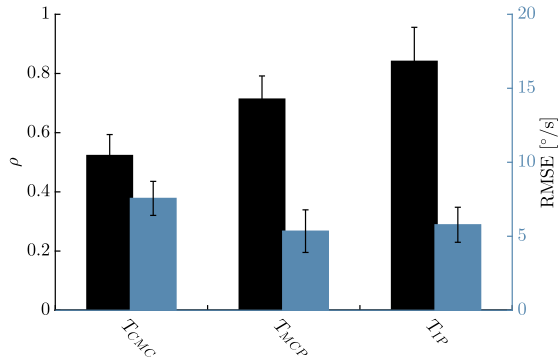
Fig. 11. $\mathcal{GP}$ autoregression with exogenous inputs performance on the thumb (3 joints). Pearson correlation coefficient ρ (black bars and left axis) and RMSE [°/s] (blue bars and right axis) between actual and predicted finger movement. Bar show mean $\pm$ standard deviation (SD).

Fig. 12. Prediction performance on test data (correlation ρ between actual and predicted data) for our *GP* autoregressive model with exogenous inputs, *when training and testing from same task were allowed.* Average performance is highlighted as dotted line as average correlation of 0.79. Bar show mean $\pm$ standard deviation (SD). This elsewhere common procedure was not used for our other reported results due to the concerns of overfitting and reflecting an apparently better performance on test data than is warranted (see main text for details).

to an individual user based on one part of the (training) data from that one user, and then evaluated using another part of the (testing) data from the same user. This is an essential process in any regression or classification setting, otherwise one will risk overfitting and not be able to observe how well a system operates under novel conditions.

Surprisingly, it is often unclear in the existing literature, how this separation of training and test data is implemented. Three general scenarios are possible with very different consequences for prediction quality and performance. Typically some form of EMG and hand kinematic data is collected across a set of tasks or movements. The first approach is to randomly select a subset of samples from the data to build up a training and test data set. This is however problematic, as in this case the time-series and the correlations within a task make training and test data dependent. An improvement would be to divide the data into long chunks of contiguous time (e.g. the first half and the second half of the recording). Alternatively, the subject performs some set of defined tasks, which are often repeated to increase the amount of data collected. Therefore, a second approach is used in several studies where the same task or movement is repeated multiple times, and some repetitions are used for training and others for testing data. We consider this problematic, because these data share the same statistical cause (the same task), and training and test data are thus only conditionally independent. Thus, the consequence will be that most likely both correlation measures and Root Mean Square Error (RMSE) will appear much better than they are. A statistically more sound approach is to randomly select data from a set of tasks designated training tasks, and a different set of tasks designated as source of testing data. We thus do not mix training and testing across repetitions of the same task. This approach enables us to appropriately judge if a decoder well operates in unknown conditions without overfitting. This last method is the one we have applied throughout the results reported here.

3) Cross-Validation: We tested the robustness of our results using cross-validation, which is a model validation technique for assessing how the results of our analysis will generalise
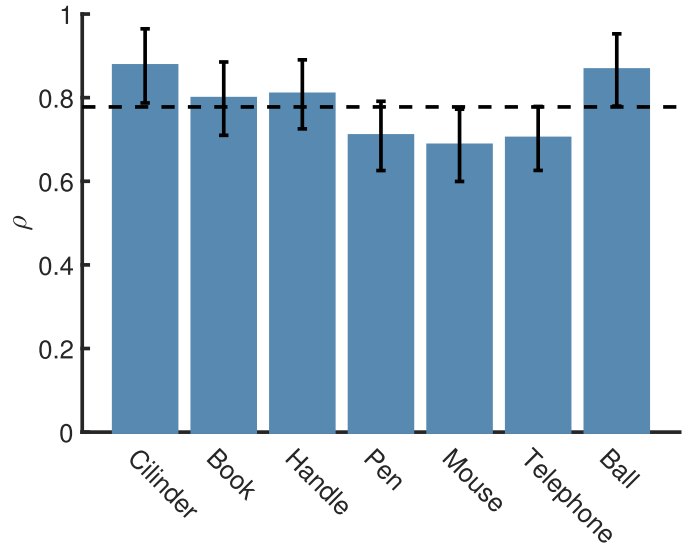
to other data. We used this method because our goal is prediction, and we want to estimate how accurately a predictive model will perform in practice. One round of cross-validation involves partitioning a sample of data into complementary subsets, performing the analysis on one subset (called the training set), and validating the analysis on the other subset (called the validation set or testing set). To reduce variability, multiple rounds of cross-validation were performed using different partitions, and the validation results were averaged over the rounds. This enabled us to assess how well the same control framework performs across all tasks. Moreover, it gave us estimates of the uncertainty we have about the value performance measures for a given control framework.

4) Evaluation of Autoregressive Models That Use Past Data to Predict Future Data: Finally, a key aspect of our approach is to use (non-linear) autoregressive techniques with exogenous inputs. This means that our model relates the current value of the angular joint velocities to both past values of the angular joint velocities and current and past values of the driving exogenous muscle signals. Evaluating the performance of the autoregressive component can be evaluated in two-different ways with important consequences for the expected performance. Two modes of evaluation arise because the system uses past values of angular joint velocities. These can be either the ones the system predicted in the past or the true angular velocities of the joints as obtained from the training and testing data. The former method, will use predictions to update predictions and thus over time diverge more due to the accumulation of noise in the prediction sequence. The latter method will have tight error bounds as past inputs will be based on actual data.

TABLE III

OVERVIEW AND COMPARISON OF PREVIOUS WORK AND RESULTS PRESENTED HERE. ABBREVIATIONS: sEMG SURFACE ELECTROMYGRAPHY, iEMG INTRAMUSCULAR ELECTROMYGOGRAPHY, MMG MECHANOMYOGRAPHY. "IDENTICAL TASKS" TRAINING AND TESTING DATA WERE NOT SEPARATED ACROSS TASK TO EVALUATE PERFORMANCE. "DIFFERENT TASKS" TRAINING AND TESTING DATA WAS SEPARATED BY TASKS TO EVALUATE PERFORMANCE. "DYN" EVALUATION USED MAINLY DYNAMIC MOTIONS (FREE MOVEMENT, PASSIVE OBJECT HOLDING). "ISO" EVALUATION USED MAINLY ISOMETRIC FORCE TASKS (e.g. PRESSING AN OBJECT). "DYN & ISO" EVALUATION USED BOTH DYNAMIC MOTIONS AND ISOMETRIC FORCES

Study	Controlled quantity hand [SI unit] & (DoF)	Sensors	$\frac{\text{DoF}}{\text{Sensor}}$	Model	Performance Pearson's ρ	Train/Test	Task
[19]	Wrist flex/ulnar [°] (2)	6 sEMG	0.33	Neural Net	0.94	identical	dyn
[17]	Wrist pronation [°] (1)	7 iEMG	0.14	Time Delay Neural Net	0.67	identical	dyn
[13]	Wrist flex/ulnar [°] (2)	192 sEMG	0.01	Linear/Non-Lin Direct Map	$0.89 - 0.95$	identical	dyn
[14]	Grasp force [N] (1)	7 sEMG	0.14	Linear Direct Map	$0.89 - 0.97$	identical	iso
[12]	Wrist flex/ulnar,Grasp [°] (3)	6 iEMG	0.5	Linear Direct Map	0.89	identical	dyn
[18]	Finger joint angles [°] (18)	8 sEMG	2.25	Time Delay Neural Network	N/A (RMSE=15°)	identical	dyn
Here	Finger joint velocities [°/s] (11)	5 sEMG + 4 MMG	1.22	GP AR with Exo Input	0.79 (RMSE=8°/s)	identical	dyn & iso
Here	Finger joint velocities [°/s] (11)	5 sEMG	2.2	GP AR with Exo Input	0.67 (RMSE=14.8°/s)	identical	dyn & iso
Here	Finger joint velocities [°/s] (11)	5 sEMG + 4 MMG	1.22	GP AR with Exo Input	0.62	different	dyn & iso
Here	Finger joint velocities [°/s] (11)	5 sEMG + 4 MMG	1.22	Linear AR with Exo Input	0.54	different	dyn & iso
[26]	index+thumb joint angles [°] (6)	9 joint angles (middle, ring + little finger)	0.67	Non-Lin Direct Map	0.47	different	dyn & iso

III. RESULTS

We present in the following section the results of our study to evaluate novel ways of decoding intended hand movements from hand kinematics and muscle activity. We measured the desired changes to finger joint angles (α) as target quantities that we wish to predict from muscle activity (u). We present the results following the approaches displayed in Fig. 1: The direct mapping approach learns a relationship between muscle activity to changes in finger joint configuration (see Fig. 1.A) and is equivalent to many existing approaches for continuous regression of hand configuration from muscle signals. We evaluated the available information in the muscle-activity alone to predict finger movements using both a linear decoder and a $\mathcal{GP}$ decoder (see Fig. 8). The $\mathcal{GP}$ based approach performs better than the linear decoder approach. These direct mapping results set a benchmark for our autoregressive framework and we compared these further below with a detailed comparison with previous studies in Table III.

Next, we investigated to which extent hand movements are predictable by themselves, e.g. one can intuitively predict how a hand motion is going to continue. We captured this inherent predictability of hand movements mathematically using the autoregressive relationship between past finger joint configuration and present finger joint configuration (see Fig. 1.B). We found that human finger movements during daily life hand use show a clear linear autocorrelation structure in time (see Fig. 5) that extends over several seconds. This model predicts future movements based on past finger movements, as it captures the averaged natural dynamics of the hand. For any autoregressive system we need to determine the degree of lag (τ) that the system should look into the past to predict the future. To this end, we run the linear autoregressive model (VARMA model), i.e. VARMAX with no exogenous inputs ($\mathbf{u}_t \in \emptyset$) for varying lags (87 ms to 500 ms). These lags were chosen so as to cover the area of support of the empirical autocorrelation function main peak (Fig. 5). Fig. 9 shows the results of the model run in the aforementioned configuration in terms of Root Mean Square Error (RMSE) and Pearson's correlation coefficient (ρ) between the 11 actual and predicted finger angular joint velocities. The values were averaged across joints, subjects and trials to reflect the overall performance of

the models for varying τ. In Fig. 9 the RMSE (right axis) and the correlation coefficient ρ (left axis) indicate, as expected, a very low performance. The performance in terms of RMSE decreases for increasing lags, while the correlation between predicted and actual hand movements improves with lags size. The results are similar for both the linear and non-linear model. Evaluating the autoregressive approach enables us, however, to set a lower limit to the performance any muscle-driven decoder should achieve.

A. (Gaussian Process) Autoregression With Exogenous Inputs

The third approach (see Fig. 1.C) is our proposed framework which combines direction mapping and autoregressive prediction – effectively placing ongoing muscle activity into the context of ongoing hand movements. This autoregressive model with exogenous inputs captures the natural dynamics of the hand (autoregressive component) as "perturbed" or driven by ongoing muscle activity (exogenous component). We present in the following the evaluation and characterisation of this approach. In analogy to our purely autoregressive model, we measured the lag-dependent prediction performance of the linear autoregressive model with exogenous inputs and $\mathcal{GP}$ based version. Fig. 10 shows the performance of both the VARMAX and the $\mathcal{GP}$ model with a 9-dimensional input space, comprising the 9 muscle-signals from the EMG and MMG sensors ($\mathbf{u}_t \in \mathbb{R}^9$). Both models show an increase in the prediction accuracy for bigger τ values even beyond the correlation length estimated from the data, whilst no significant difference was found between the $\mathcal{GP}$ and VARMAX performance.

The performance of both our models is better than the linear or $\mathcal{GP}$-based direct mapping by about 30%, and the RMSE is about 50%-70% lower across lags. The performance is for all lags and for both the VARMAX and the $\mathcal{GP}$ model better than the purely autoregressive system, by an order of magnitude, in terms of correlation and almost two orders of magnitude in terms of RMSE. More importantly, the $\mathcal{GP}$ based approach produces better correlated results than the linear autoregression. Thus, the linear and $\mathcal{GP}$-based autoregression with exogenous inputs performs better than both direct and
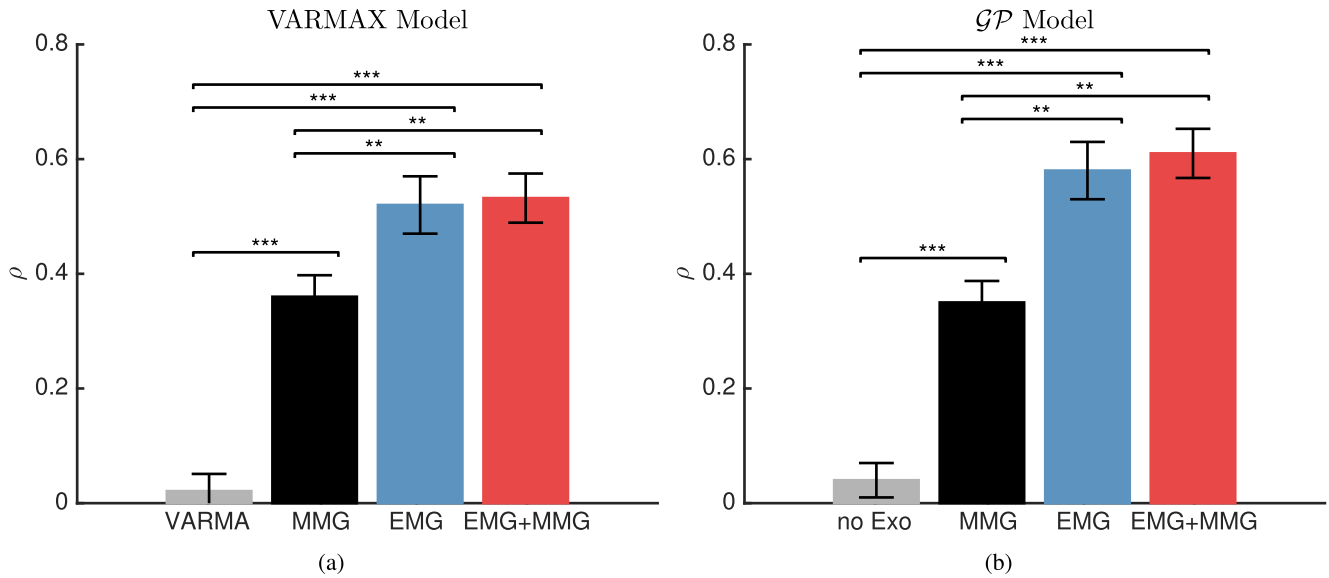
Fig. 13. Contribution of sensory modalities in the linear autoregressive model with exogenous inputs (a) and the *GP* framework (b). Both frameworks yield to similar results: the lower-limit to decoder performance of the decoder is set by only the correlation of natural hand movements, as determined by fitting the VARMA model (no exogenous inputs in the *GP* model), left-most bar in both figures. Including MMG signals as exogenous inputs, or just EMG signals or both MMG and EMG yields large jumps improvements in performance. A significant difference was found between the case with no-exogenous-inputs model and all other configurations, and between the MMG configuration and the EMG and EMG+MMG. ** denotes $p < 0.001$ and *** denotes $p < 0.0001$ of a two-sample t-test with $\alpha = 0.05$. Bar show mean $\pm$ standard deviation (SD).

purely autoregressive methods. Both models are able to control more independent degrees of freedom in the finger actuators (11) than we have sensors to read out muscle activity (9). The ratio of independent Degrees of Freedom (DoF) controlled to sensors is 1.22.

We investigated the underlying mechanisms of this prediction by analysing the optimised values of $\Delta = (c, b, \varphi, \theta)$ and $\Theta = \{\sigma_f, \ell, \sigma_l\}$ for various lags. For the VARMAX model, it was found that with increasing lag τ, the contribution from the autoregression term φ decreases while the opposite happens for the moving average term θ. We found similar results for the GP model with an increasing contribution of EMG-MMG signals $\mathbf{u}_t$, and reduced contribution of previous state $\alpha_{t-\tau}$. This trend could be explained by the natural smoothness of human movements [37], which causes consequent values of angular velocity close in time to be highly correlated. For small values of τ, the model thus relies nearly entirely on the previous state.

Globally, the VARMAX and the GP models did not show significant difference for a fixed time-lag (p-value of a two-sample t-test > 0.05), and both methods proved to be stable over time, with no significant change in prediction quality during the ≈ 20 mins of the experiment.

Fig. 14 shows the predicted joint velocities of the hand for three consecutive ball-grasping tasks, as estimated by the GP, with focus on the index MCP joint during a single repetition. GPs return, for each estimated point, a Gaussian distribution whose standard deviation can be interpreted as a measure of the uncertainty of the predicted value. This feature can be used for risk-based control of neuroprosthetic devices (see below).

The best performing model of our framework is the use the non-linear GPs to combine muscle activity and autoregressive inputs of hand posture with large lags. This model produces a good degree of correlation between actual and predicted hand velocities $\rho = 0.61$ and very low RMSE (8°/s).

The GP-based approach is, by design, predicting each individual finger joint independently of the other joints. Crucially, we find (see Fig. 11) that the thumb's individual joints are highly predictable in terms of correlation ($\rho = 0.82$) and RMSE, which is important given the opposable thumb being the most independently controlled finger of the hand [24].

1) Generalisation Performance: We have to be able to judge how well a decoder will operate in unknown conditions when the user wants to generate a novel set of movements. To assess this generalisation performance we have separate training data from testing data, this is done to avoid overfitting and obtaining spuriously high performance reports. We found in the literature that generalisation performance is assessed with a problematic lack of separation between training and test data (see Methods for details). This occurs when the same task or movement are repeated multiple times and some repetitions are used for training and others for testing data. This is problematic because these data share the same statistical cause (the task or movement), and training and test data are thus only conditionally independent. Hence, we pursued a statistically more rigorous approach by randomly selecting data from a set of tasks designated training tasks, and a different set of tasks designated as source of testing data. We thus do not mix training and testing across repetitions of the same task.

To illustrate the importance of this distinction we computed the performance of our best performing GP autoregressive model with exogenous inputs. We used two methods to separate test and training data: 1. training using separate training and test tasks (our recommended method), and 2. training where we allowed data from different repetitions of the same task to be used for training and testing (see Fig. 12). The performance of the second training approach is considerably higher ($\rho = 0.79$ vs $\rho = 0.61$), the level depending on which task's repetitions were used for training and testing. In general, the test data performance is upward biased by
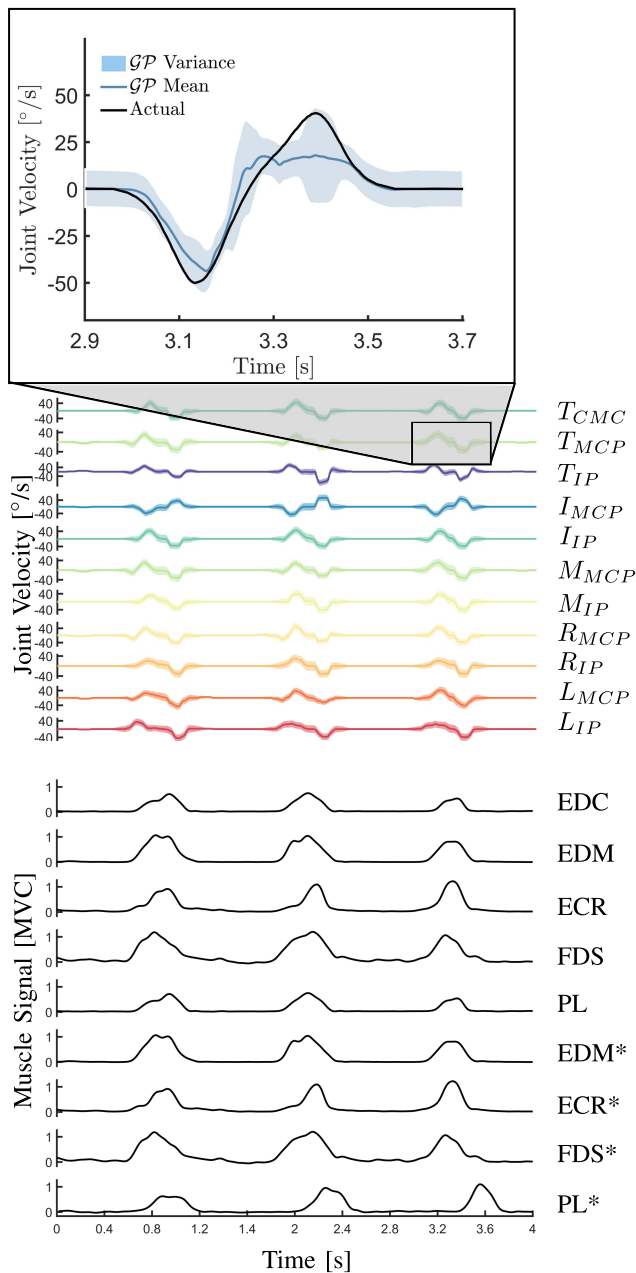
Fig. 14. Towards risk-based control. Muscle signals and decoded joint velocities of the *GP* model for three consecutive repetitions of grasping and squeezing a ball. Bottom set of plots shows the 5 EMG and 4 MMG signals as fraction of maximum voluntary contraction (MVC). Middle set of plots shows the predictive distribution of the on-line decoded joint angular velocity set to control 11 independent finger joints. The inset shows the actual (black) versus predicted (blue) values of the thumb MCP joint. We plot in the inset the uncertainty of the system in its belief of the prediction, as the standard deviation of the *GP* framework's predictive distribution (blue area). ∗ denotes muscle signals from MMG sensors. Refer to Tables II and I for muscles and joints acronyms.

30% when using the second method. These results highlight how repetitions of specific tasks, here grasping a cylinder or a ball, provide considerably upward bias to the test performance values. We report the differences in test performance and the degree of method of separating between training and test data when comparing our results to previous work in the literature (as these used this approach, see Tab. III, column training/testing) and report the more rigorous method (which

results in overall lower performance values) throughout our results otherwise. In addition, any rigorous approach for evaluation has to be complemented by cross-validation to estimate the variability in performance reports. This enables us to carry out statistical tests to compare different approaches and highlights for example that performance differences between our VARMAX and $\mathcal{GP}$ approach that appear when comparing average performance values, can actually be small, when factoring in the variability across reported performance measures.

2) Evaluating Multiple Sensor Modalities: In our multi-modal approach, we combined directly both EMG and MMG data. How important or relevant is each sensor modality for our reconstruction error? We evaluated the relative contribution of these two exogenous modalities. While the $\mathcal{GP}$ approach forms a complex non-linear representation of the multiple modalities, we can use the linear model to get a sense of the individual contributions of the modalities. Fig. 13a shows the results obtained from the VARMAX model, in terms of the correlation coefficient between the actual and predicted velocities, as a result of using MMG signals as exogenous inputs (black bar), EMG signals as exogenous inputs (blue bar) and both (red bar). A statistical analysis (two-sample t-test with $\alpha = 0.05$) revealed a significant improvement from the autoregressive model with no exogenous inputs (VARMA) to the ones with exogenous inputs (p<0.001). Including EMG signals also proved to be significantly better than MMG in improving the performance of the model. Similar results were found for the $\mathcal{GP}$ model (Figure 13(b)).

3) Enabling Risk-Based Control: While standard approaches in prosthetic control decode a specific action, our $\mathcal{GP}$-based approach computes the probability distribution of a whole family of possible actions. We can interpret this distribution of actions as a measure of uncertainty that we have in the users intended action, e.g. whether they movement outputs are tightly distributed around a specific joint configuration or more widely spread around. This enables the controller to estimate the uncertainty it has about exposing a specific action to the user. In Fig. 14 we show the on-line muscle activity received (bottom plots) and the predictions of the $\mathcal{GP}$ framework (middle plots). The $\mathcal{GP}$ model computes a distribution over possible predictions, the model's predictive distribution. I.e. the $\mathcal{GP}$ model returns here a mean value and a standard deviation of the prediction for each instant in time (we show the uncertainty as shaded area denoting one standard deviation around the mean prediction). In the inset Fig. 14, one can see how the uncertainty of the $\mathcal{GP}$ model about its own predictions is larger where the predicted value is far from the actual one. This is a desired outcome, potentially more important than just predicting with a smaller error, because the prosthetic controller is aware of the error it may be making. This enables a prosthetic controller to have fewer unknown unknowns to operate with and enables the notion of risk-based control.

IV. DISCUSSION

The recent introduction of novel multifunction hands, as well as new control paradigms, increase the demand for advanced prosthetic control systems, leading us to present a novel probabilistic framework. We follow the taxonomy of

Fougner *et al.* [15] to place our prosthetic control approach into the field's context. In the context of this taxonomy, we present here a proportional multi-modal simultaneous multi-dimensional proportional control framework. Moreover, we introduce novel aspects of risk-based control and under-sensored control. Our control approach is *proportional*, as a user can map a continuous sensory input interval by varying his control input within a corresponding continuous interval. We use a *simultaneous multi-dimensional* control approach, in that we are controlling multiple independent degrees of freedom (DoF) of the system (individual finger joints) at the same time. This is in contrast to established approaches that combine a classifier that selects which hand movement to execute together with a mapping function (regression function) that actuates that hand movement in a proportional way. To be clear, controlling opening and closing of the hand as a whole may affect multiple mechanical DoF of the hand mechanics, but has only one independent degree of freedom in terms of the user's ability to control it.

We control more independent degrees of freedom than we have independent sensors (11 actuated joints vs 9 sensors, ratio $\frac{11}{9} = 1.22$), while previous work typically has a ratios well below 0.5. Thus, unlike most previous work, our approach enables *undersensored* control, where we are controlling more independent degrees of freedom than we have sensors. This striking result is possible because we exploit the temporal correlation structure of natural hand movements to provide a context for the few muscle signal sources we use. This allows our framework to disambiguate among possible motions.

Multi-modal control implies here that we combine several sensory modalities together, i.e. joint angle configuration of the prosthetic, surface/intramuscular electromyography (s/iEMG) and mechanomyography (MMG). While EMG has long been established in prosthetic control for 70 years [10], mechanomyography has been only recently used before for classification driven prosthetic control approaches [38]. We used our microphone-based surface MMG sensors [29] that enable recording muscle activity without need for electrical contact. This makes day-long recordings of muscle activity feasible and with a signal-to-noise ratio comparable to EMG [39]. We reported that MMG recordings are stable enough to perform reliable muscle force prediction [40], [41]. Modern powered prostheses typically use surface electromyography (sEMG) for neuromuscular monitoring and prosthetic control. Alternatively, recent studies have proven the efficacy of mechanomyography (MMG), EMG's mechanical counterpart [29], [39], which reflects both motor unit recruitment [42] and global motor unit firing rate [43]. While EMG has been successfully used for classification of individual finger movement [44] and for fine actions robotic control [45] in the laboratory environment, prosthetic control mostly involves binary classification to decide whether to open or close the hand. Recently, MMG studies have demonstrated similar results [46], suggesting that MMG signals are valuable alternatives or complements for neuroprosthetic control. We evaluate how the combination of MMG and EMG signals affects performance and find little benefit in including MMG signals to the EMG signals. However, we showed that using MMG signals alone provides a reasonable but lower performance compared to EMG, but EMG signals alone eclipsed the need for additional MMG signals. Considering the MMG is typically factor 10-100 cheaper per channel than EMG [29], we believe that current MMG sensors are probably better suited for discrete control, rather than simultaneous, proportional control of many degrees of freedom. Independent of the specific outcome of MMG performance, the use of a probabilistic inference model to infer joint kinematics for prosthetic controls has a fundamental conceptual advantage – namely that it allows to combine multiple physical modalities (and physical units) of input signals in a principles manner. This is because conventional approaches such as neural networks and also our VARMAX method produce a mapping between the values (irrespective of units) and so change in units (e.g. microvolt range for EMG is very different from centivolt range of MMG) produce scale-dependent changes that normally have to be carefully normalised to map across modalities (e.g. when combining joint angle and voltage inputs). In contrast probabilistic approaches, such as our $\mathcal{GP}$, map values to a probability distribution of latent variables and thereby establishes relationships between these probability distributions. Thus, probabilities are mapped to probabilities, thereby enabling physical scale-invariance and invariance across physical modality (physical units).

A. Comparison to Previous Work

Placing our results in the context previous work requires making comparisons across very different studies – in absence of an established publicly accessible benchmark data set or physical system for evaluating proportional control. We therefore focussed on selected recent work that published details of their data and their performance evaluation procedure (please refer to Tab. III throughout this section) to compare our work to.

Careful selection of the performance evaluation measure for control algorithms is important, as these use the performance as the automatic optimisation criterion for the system. We focussed on decoding joint velocities, which are highly conserved across users [24] and are within users more stationary than other quantities [32], thus promise to generalise better to novel users. Using joint velocities as a target modality enables us to capture actuator impulse, enabling a compromise between controlling hand pose (useful in dynamic tasks, e.g. passively hold an apple) and hand force (useful in isometric force tasks such as opening a bottle). This contrasts with previous work that either only considered hand pose (joint angles) [11], [12], [16]–[18], [25] or a grasp force [13].

The RMSE we report for GP autoregression is averaged across all joints, the degree of error between actual and predicted joint angular velocity is very small 8°/s. This RMSE is equivalent to the angular velocity of a watch hand that takes 45 seconds to complete a full 360° circle and comparable to the natural drift velocity of the resting, passive hand. However, the root mean square error (RMSE) is not necessarily a good measure of performance for a prosthesis control system, because one may confound a low RMSE with a poor ability to follow the dynamics of the system (see [15] for review,

see also [47], [48]). Moreover, as other proportional control studies looked at different modalities e.g. force or joint angle it is difficult for us to compare raw RMSE across different physical units. Joint angles capture little about the dynamics of the hand e.g. do not account well for the fact that pose-based control induces transients as the hand moves from one configuration to the intended configuration, which can make it difficult to control hands fluidly in closed loop. We conclude that we need to compare the instantaneous Pearson correlation coefficient $\rho = \sqrt{R^2}$ of the predicted and actual signal, which indicates how well the predictions perform and co-vary with the data (see column "Perf"). ρ is unit-less and thus allows a somewhat better comparison across systems. Our correlation coefficient for predicting 11 finger joint movements is within 10% to 15% of current state-of-the art approaches focussing on wrist control [12], [16], [18]. We consider finger decoding the harder problem as the hand has more degrees of freedom and these often involve more, smaller and deeper muscles in the forearm than those of the wrist and thus have poorer signal quality.

Comparison of these performance numbers is further complicated by the nature of the performance evaluation (if it is published in sufficient detail). We show that there is a 25% boost in correlation measures of test performance depending on whether one omits separating test and training data so that data from the same task (but different repetitions) is used for both settings (see column Train/Test, report "identical") or whether training and testing data is clearly separated across tasks (see column Train/Test, report "different"). The latter is statistically the more sound measure to evaluate test data performance as it speaks to the ability of a framework to generalise to novel settings (and potentially novel users) – and is the one measure we applied throughout our analysis. The appropriate separation of test and training data needs to be paired with a cross-validation approach to estimate the uncertainty in performance measures. All our results presented here report variability of performance measures, but we did not include this variability in Tab. III) as aside from [24] no uncertainty estimates were published in the tabled literature. Finally, another confounding factor that blurs the difference in performance across studies is that previous work focusses on assessing proportional control in either dynamic movement tasks, i.e. dynamic, free, unimpeded movement of the hand or passive hold of a light object (see column Task - report "dyn"), or on isometric tasks (pressing a force sensor, e.g. [13]) (see column Task - report "iso"). In contrast, our set of data combines both dynamic tasks and isometric object manipulation tasks that arise from interaction with daily life objects. This approach reflects our need and aim to evaluate performance across daily life scenarios and the use of joint angular velocity control – it is difficult to extrapolate how well other methods would generalise across this natural spectrum of tasks. We suggest that focussing on free movement tasks alone or isometric force tasks alone will result in inflated performance reports than when evaluating a more naturalistic combination of the two as we have done here.

Moreover, there are large differences in the number of independently controlled degrees of freedom (DoF), as sometimes a single variable is decoded (e.g. grasp force) while in other cases multiple-actuators are decoded individually. We focussed on actuating 11 individual finger joints independently. The sensors used are very different, encompassing surface EMG electrodes, high-density EMG arrays but also invasive measures of muscle activity (see column "sensors"). Our approach is undersensored, as we control not only a much larger ($>$3-fold) set of independent DoF than most previous approaches, but also have a higher ratio degrees of freedom controlled independently to sensors (electrodes) ($\frac{11}{5} = 2.2$ or $\frac{11}{9} = 1.22$ if we include MMG with the EMG). In contrast, previous work had ratios significantly smaller than 1 (see column "DoF/Sensor ratio"). Reference [17] used comparable levels of undersensored control to our work presented here, however a correlation coefficient was not reported and the RMSE of 15° was estimated from their publication, as no summary RMSE was reported and per joint error was reported as fraction of (unspecified) joint ranges (to obtain the value we assumed 90° joint range for all joints (which is a low value) and averaged across the joints). Also their work considered only dynamic free movement tasks without isometric contractions resulting from object manipulation. Our work controls joint velocities (RMSE 8°/s) and not absolute joint configuration as [17], and thus over a period of roughly 2 seconds our control error would accumulate to their reported RMSE of 15°, however in this time the end-user would have multiple time steps time to compensate for pose configuration error. We thus, conclude that when considering the large number of simultaneously controlled degrees of freedom and the degree of undersensored control for fluid full-finger hand control, GP autoregression operates twice as well as previous comparable work.

Previous approaches use some form of direct-mapping paired most frequently with neural networks which take into account present (and past) muscle activity, while we developed a probabilistic description of the inference problem (cf. Fig. 2) and used an $\mathcal{GP}$ autoregressive with exogenous inputs to solve the problem. Some work performed direct mapping of signals after muscle activity signal feature extraction [11], [13], [18], these features implicitly included time-series element of the muscle signal input. Other work [16], [17] used for example explicitly time-delayed inputs into neural networks to incorporate the history of muscle activities. In contrast here we operate with the history (i.e. context) of the hand movement itself (autoregression), and thus our approach takes into account both past joint kinematic context and muscle activity to perform decoding. Thus our work should be also compared to [21], [25], which exploit directly the conserved spatio-temporal structure of natural hand kinematics to predict from the configuration of fingers 3-5 the configuration of the (missing) thumb and index finger. That work is applied to partial hand prosthetics and operates without having to use muscle activity, while here we reconstruct full hand configuration from knowledge of the spatio-temporal structure and muscle activity, we achieve an over 30% higher correlation coefficient here than reported before. Proportional control approaches focussed on different control quantities, such as wrist control and open/closure of the hand (as a whole) or force of grasp

(see column "controlled quantity"). In contrast, we enabled here detailed fluid control of the individual fingers and especially the thumb, for which we report an average correlation coefficient of $\rho = 0.72$ and RMSE of 7°/s.

B. Risk- and Context-Aware Control

The nervous system and movements are inherently noisy [49], [50]. We use here probabilistic Bayesian methods to process this neurobehavioural data in a principled way to operate with this uncertainty. We presented a novel framework for prosthetic control, framed in the language of Bayesian Decision Theory that reinterpreted the prosthetic decoding problem as a recursive time series prediction model with external inputs. We compared both linear approaches and non-linear methods, based on Gaussian Process ($\mathcal{GP}$) regression. While $\mathcal{GP}$ regression has not been used for direct mapping problems before, it is generally considered equivalent in terms of performance to conventional artificial neural network approaches, which have been extensively used before in direct mapping of prosthetic control. In fact Gaussian Processes are formally equivalent to a neural network with an infinite number of units in the hidden layer [34], [51], [52]. Using $\mathcal{GP}$ respression allows us now to extend the prosthetics taxonomy [15] by the notion of risk-based control. $\mathcal{GP}$ regression provides a principled method for us to obtain the decoders uncertainty in interpreting the user intention. We propose that risk-based control enables us to use the systems own understanding of how uncertain, given the data and the calibration, the decoding of a specific hand-movement can be. This allows the controller to make decisions about whether to go ahead with an unmodulated movement of the hand or choose a modulated movement (e.g. uses smaller amplitudes or forces to avoid dangerous situations where the user suffers from unintended actions). These unique benefits of probabilistic inference over standard neural network based approaches for decoding come at a cost. $\mathcal{GP}$ regression is numerically complex during training (which can be however performed offline) and is also somewhat more complex to evaluate than neural network based methods. We found that these potential concerns do not form an issue for two reasons. 1. we proposed here the use of approximation methods to simply and reduce the complexity of inferring the $\mathcal{GP}$ for prosthetic control and 2. constant advances in embedded processing, have seen $\mathcal{GP}$ processes already successfully deployed in real-time activity recognition in smartphone-based settings using state-of-the-art low-power embedded processors (e.g. [53], [54]).

While better quality and more sensor signals, e.g. from high density electrode arrays, will improve our system's performance further, we demonstrated the power of using the inherent structure of natural movements with recordings from a few sensors, to achieve *undersensored* control, i.e. enabling a larger independent degrees of freedom to be controlled from a smaller number of sensors. This will become increasingly important as robotic prosthetic hands increase in sophistication, enabling not only dextrous control of grasping but also allowing us to perform in-hand object manipulation tasks: The EthoHand [7], for example, is capable of writing a text message on a smart phone or rolling balls inside its palm [7]. This mechatronic capability is enabled by novel thumb designs that enable complex pointing of the thumb, and hence which requires corresponding decoding of thumb action, as we demonstrated here. We showed that the EthoHand can render the complex spatio-temporal structure of natural human hand movements to an unprecedented extent [27] and here we used the very same spatio-temporal structure to boost decoding performance by providing context to the muscle activity recorded.

We propose that more detailed and higher-level context information will further boost the ability to decode correctly ambiguous and noisy sensor signals. This is enabled by current developments of haptic sensing algorithms (Haptic SLAM) that can reconstruct and recognise from finger joint kinematics alone the object shape that is manipulated by the artificial hand [55], [56], but also come from sensory feedback of the hand. We will aim to move beyond low-level kinematics, as used here, to higher order structure of hand-based actions (e.g. reach, grasp, pull), tasks (e.g. drink from cup) and ultimately overall user intentions (e.g. having breakfast) [24]. This high-level information will be eventually available from smart home environments but could also be inferred directly from the activity dynamics with information accessible to a prosthetic hand itself. These developments need to be supported by careful designed performance evaluations, that cleanly separate training and test data, evaluate decoding across both dynamic and isometric tasks taken from daily life. Ultimately, the best performance measures should be the ability, agility and satisfaction with which an actual end-user can perform complex tasks, as demonstrated in [20]). In conclusion we demonstrated how statistical machine learning approaches can yield in a purely data-driven way a method that efficiently combines contextual information with multi-modal data to boost the ability of end-users to control an artificial hand's finger better approaching the dexterity of the human hand.

REFERENCES

[1] A. Faisal, D. Stout, J. Apel, and B. Bradley, "The manipulative complexity of lower Paleolithic stone toolmaking," *PLoS One*, vol. 5, no. 11, p. e13718, 2010.

[2] E. Hecht *et al.*, "Acquisition of Paleolithic toolmaking abilities involves structural remodeling to inferior frontoparietal regions," *Brain Struct. Function*, vol. 220, no. 4, pp. 2315–2331, 2015.

[3] R. Beasley, "General considerations in managing upper limb amputations," *Orthop. Clin. North Amer.*, vol. 12, no. 4, pp. 743–749, 1981.

[4] J. Belter and J. Segil, "Mechanical design and performance specifications of anthropomorphic prosthetic hands: A review," *J. Rehabil. Res.*, vol. 50, no. 5, p. 599, 2013.

[5] E. Biddiss, D. Beaton, and T. Chau, "Consumer design priorities for upper limb prosthetics," *Disab. Rehabil. Assist. Technol.*, vol. 2, no. 5, pp. 346–357, 2007.

[6] E. A. Biddiss and T. T. Chau, "Upper limb prosthesis use and abandonment: A survey of the last 25 years," *Prosthetics Orthotics Int.*, vol. 31, no. 3, pp. 236–257, 2007.

[7] C. Konnaris, C. Gavriel, A. A. C. Thomik, and A. A. Faisal, "Ethohand: A dexterous robotic hand with ball-joint thumb enables complex in-hand object manipulation," *IEEE Biorobotics*, vol. 6, pp. 1–6, Sep. 2016.

[8] T. R. Makin, F. de Vignemont, and A. A. Faisal, "Engineering embodiment: Neurocognitive constraints and requirements as design principles for human augmentation technology," *Nature Biomed. Eng.*, vol. 1, no. 1, p. 0014, 2017.

[9] K. Østlie, I. M. Lesjø, R. J. Franklin, B. Garfelt, O. H. Skjeldal, and P. Magnus, "Prosthesis rejection in acquired major upper-limb amputees: A population-based survey," *Disab. Rehabil. Assist. Technol.*, vol. 7, no. 4, pp. 294–303, 2012.

[10] R. Reiter, "Eine neue elektrokunsthand," *Grenzgebiete der Medizin*, vol. 1, no. 4, pp. 133–135, 1948.

[11] L. H. Smith, T. A. Kuiken, and L. J. Hargrove, "Linear regression using intramuscular EMG for simultaneous myoelectric control of a wrist and hand system," in *Proc. 7th Int. IEEE/EMBS Conf. Neural Eng.*, Apr. 2015, pp. 619–622.

[12] J. M. Hahne *et al.*, "Linear and non-linear regression techniques for simultaneous and proportional myoelectric control," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 22, no. 2, pp. 269–279, Feb. 2014.

[13] C. Castellini, A. E. Fiorilla, and G. Sandini, "Multi-subject/daily-life activity EMG-based control of mechanical hands," *J. Neuroeng. Rehabil.*, vol. 6, no. 1, p. 1, 2009.

[14] D. Farina *et al.*, "The extraction of neural information from the surface EMG for the control of upper-limb prostheses: Emerging avenues and challenges," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 22, no. 4, pp. 797–809, Jul. 2014.

[15] A. Fougner, O. Stavdahl, P. J. Kyberd, Y. G. Losier, and P. A. Parker, "Control of upper limb prostheses: Terminology and proportional myoelectric control—A review," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 20, no. 5, pp. 663–677, May 2012.

[16] C. L. Pulliam, J. M. Lambrecht, and R. F. Kirsch, "EMG-based neural network control of transhumeral prostheses," *J. Rehabil. Res. Develop.*, vol. 48, no. 6, p. 739, 2011.

[17] F. Sebelius, L. Eriksson, C. Balkenius, and T. Laurell, "Myoelectric control of a computer animated hand: A new concept based on the combined use of a tree-structured artificial neural network and a data glove," *J. Med. Eng. Technol.*, vol. 30, no. 1, pp. 2–10, 2006.

[18] J. L. Nielsen, S. Holmgaard, N. Jiang, K. B. Englehart, D. Farina, and P. A. Parker, "Simultaneous and proportional force estimation for multifunction myoelectric prostheses using mirrored bilateral training," *IEEE Trans. Biomed. Eng.*, vol. 58, no. 3, pp. 681–688, Mar. 2011.

[19] D. Yatsenko, D. McDonnall, and K. S. Guillory, "Simultaneous, proportional, multi-axis prosthesis control using multichannel surface EMG," in *Proc. 29th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, Sep. 2007, pp. 6133–6136.

[20] S. Amsuess, P. Goebel, B. Graimann, and D. Farina, "A multi-class proportional myocontrol algorithm for upper limb prosthesis control: Validation in real-life scenarios on amputees," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 23, no. 5, pp. 827–836, Sep. 2015.

[21] J. J. Belić and A. A. Faisal, "Decoding of human hand actions to handle missing limbs in neuroprosthetics," *Frontiers Comput. Neurosci.*, vol. 9, no. 27, 2015, doi: 10.3389/fncom.2015.00027.

[22] M. Santello, M. Flanders, and J. F. Soechting, "Postural hand synergies for tool use," *J. Neurosci.*, vol. 18, no. 23, pp. 10105–10115, 1998.

[23] J. N. Ingram, K. P. Körding, I. S. Howard, and D. M. Wolpert, "The statistics of natural hand movements," *Experim. Brain Res.*, vol. 188, no. 2, pp. 223–236, 2008.

[24] A. A. C. Thomik, S. Fenske, and A. A. Faisal, "Towards sparse coding of natural movements for neuroprosthetics and brain-machine interfaces," in *Proc. IEEE Neural Eng. (NER)*, vol. 7. Sep. 2015, pp. 938–941.

[25] A. A. C. Thomik, D. Haber, and A. A. Faisal, "Real-time movement prediction for improved control of neuroprosthetic devices," in *Proc. IEEE Neural Eng. (NER)*, vol. 6. Sep. 2013, pp. 625–628.

[26] D. Haber, A. A. C. Thomik, and A. A. Faisal, "Unsupervised time series segmentation for high-dimensional body sensor network data streams," in *Proc. IEEE Wearable Implant. Body Sensor Netw. (BSN)*, vol. 11. Apr. 2014, pp. 121–126.

[27] C. Konnaris, A. A. C. Thomik, and A. A. Faisal, "Sparse eigenmotions derived from daily life kinematics implemented on a dextrous robotic hand," in *Proc. IEEE Int. Conf. Biomed. Robot. Biomechatronics (BioRob)*, vol. 6. Sep. 2016, pp. 1–6.

[28] T. A. Kuiken *et al.*, "Targeted muscle reinnervation for real-time myoelectric control of multifunction artificial arms," *JAMA*, vol. 301, no. 6, pp. 619–628, 2009.

[29] S. Fara, C. S. Vikram, C. Gavriel, and A. Faisal, "Robust, ultra low-cost MMG system with brain-machine-interface applications," in *Proc. 6th Int. IEEE/EMBS Conf. Neural Eng. (NER)*, Nov. 2013, pp. 723–726.

[30] N. Kamakura, M. Matsuo, H. Ishii, F. Mitsuboshi, and Y. Miura, "Patterns of static prehension in normal hands," *Amer. J. Occupat. Therapy*, vol. 34, no. 7, pp. 437–445, 1980.

[31] M. Šarić, *Libhand: A Library for Hand Articulation version 0.9.*, 2011.

[32] E. Todorov and Z. Ghahramani, "Analysis of the synergies underlying complex hand manipulation," in *Proc. 26th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc. (IEMBS)*, vol. 2. Sep. 2004, pp. 4637–4640.

[33] D. A. Dickey and W. A. Fuller, "Distribution of the estimators for autoregressive time series with a unit root," *J. Amer. Statist. Assoc.*, vol. 74, no. 366a, pp. 427–431, 1979.

[34] J. Bernardo *et al.*, "Regression and classification using Gaussian process priors," in *Bayesian Statistics*, vol. 6.

[35] C. Rasmussen and C. Williams, "Gaussian processes for machine learning," in *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2006.

[36] J. Quiñonero-Candela and C. E. Rasmussen, "A unifying view of sparse approximate Gaussian process regression," *J. Mach. Learn. Res.*, vol. 6, pp. 1939–1959, Dec. 2005.

[37] Q.-C. Pham, H. Hicheur, G. Arechavaleta, J.-P. Laumond, and A. Berthoz, "The formation of trajectories during goal-oriented locomotion in humans. II. A maximum smoothness model," *Eur. J. Neurosci.*, vol. 26, no. 8, pp. 2391–2403, 2007.

[38] J. Silva, T. Chau, and A. Goldenberg, "MMG-based multisensor data fusion for prosthesis control," in *Proc. IEEE 25th Annu. Int. Conf. Eng. Med. Biol. Soc.*, vol. 3. Apr. 2003, pp. 2909–2912.

[39] C. Gavriel and A. A. Faisal, "A comparison of day-long recording stability and muscle force prediction between BSN-based mechanomyography and electromyography," in *Proc. 11th Int. BSN Conf*, 2014, pp. 69–74.

[40] S. Fara, C. Gavriel, C. S. Vikram, and A. A. Faisal, "Prediction of arm end-point force using multi-channel MMG," in *Proc. 11th Int. Conf. Wearable Implant. Body Sensor Netw.*, Sep. 2014, pp. 27–32.

[41] M. Xiloyannis, C. Gavriel, A. A. C. Thomik, and A. A. Faisal, "Dynamic forward prediction for prosthetic hand control by integration of EMG, MMG and kinematic signals," in *Proc. 7th Int. IEEE/EMBS Conf. Neural Eng. (NER)*, vol. 7. Sep. 2015, pp. 611–614.

[42] K. Akataki, K. Mita, and M. Watakabe, "Electromyographic and mechanomyographic estimation of motor unit activation strategy in voluntary force production," *Electromyogr. Clin. Neurophysiol.*, vol. 44, no. 8, pp. 489–496, 2004.

[43] T. W. Beck *et al.*, "Does the frequency content of the surface mechanomyographic signal reflect motor unit firing rates? A brief review," *J. Electromyogr. Kinesiol.*, vol. 17, no. 1, pp. 1–13, 2007.

[44] F. Tenore, A. Ramos, A. Fahmy, S. Acharya, R. Etienne-Cummings, and N. V. Thakor, "Towards the control of individual fingers of a prosthetic hand using surface EMG signals," in *Proc. 29th Annu. Int. EMBS Conf.*, Aug. 2007, pp. 6145–6148.

[45] B. Crawford, K. Miller, P. Shenoy, and R. Rao, "Real-time classification of electromyographic signals for robotic control," in *Proc. AAAI*, 2005, pp. 523–528.

[46] N. Alves and T. Chau, "Uncovering patterns of forearm muscle activity using multi-channel mechanomyography," *J. Electromyogr. Kinesiol.*, vol. 20, no. 5, pp. 777–786, 2010.

[47] L. Hargrove, Y. Losier, B. Lock, K. Englehart, and B. Hudgins, "A real-time pattern recognition based myoelectric control usability study implemented in a virtual environment," in *Proc. 29th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc.*, Apr. 2007, pp. 4842–4845.

[48] B. Lock, K. Englehart, and B. Hudgins, "Real-time myoelectric control in a virtual environment to relate usability vs. accuracy," in *Proc. Myoelectr. Symp.*, 2005, pp. 1661–1667.

[49] A. A. Faisal, S. B. Laughlin, and J. A. White, "How reliable is the connectivity in cortical neural networks?" in *Proc. IEEE IJCNN*, vol. 2. Sep. 2002, pp. 1661–1667.

[50] A. A. Faisal, L. P. Selen, and D. M. Wolpert, "Noise in the nervous system," *Nature Rev. Neurosci.*, vol. 9, no. 4, pp. 292–303, 2008.

[51] R. M. Neal, "Bayesian learning for neural networks," Ph.D. dissertation, Univ. Toronto, Toronto, ON, Canada, 1995.

[52] D. J. MacKay, "Gaussian processes—A replacement for supervised neural networks?" Cavendish Lab., Cambridge Univ., Cambridge, U.K., Tech. Rep., 1997.

[53] H. Shin, Y. Chon, and H. Cha, "Unsupervised construction of an indoor floor plan using a smartphone," *IEEE Trans. Syst., Man, Cybern. C, Appl. Rev.*, vol. 42, no. 6, pp. 889–898, Jun. 2012.

[54] L. Zhang, J. Liu, H. Jiang, and Y. Guan, "SensTrack: Energy-efficient location tracking with smartphone sensors," *IEEE Sensors J.*, vol. 13, no. 10, pp. 3775–3784, Oct. 2013.

[55] F. M. P. Behbahani, R. Taunton, A. A. C. Thomik, and A. A. Faisal, "Haptic SLAM for context-aware robotic hand prosthetics—Simultaneous inference of hand pose and object shape using particle filters," in *Proc. IEEE Neural Eng. (NER)*, vol. 7. Sep. 2015, pp. 719–722.

[56] F. M. P. Behbahani, G. S. Buxarrais, and A. A. Faisal, "Haptic slam: An ideal observer model for Bayesian inference of object shape and hand pose from contact dynamics," *Eurohaptics*, vol. 6, pp. 1–6, Apr. 2016.

Michele Xiloyannis, photograph and biography not available at the time of publication.

Constantinos Gavriel, photograph and biography not available at the time of publication.

Andreas A. C. Thomik, photograph and biography not available at the time of publication.

A. Aldo Faisal, photograph and biography not available at the time of publication.