INTUITIVE OPTIMAL PROSTHESIS TUNING THROUGH USER TUNED COSTS OF EFFORT AND ACCURACY

Anjana Gayathri Arunachalam, Kevin B. Englehart, and Jon W. Sensinger
Institute of Biomedical Engineering, University of New Brunswick

ABSTRACT

Clinicians and prosthesis users care about the practical attributes of movement like the physical effort required by the user, the response time of the device, the reliability, and the accuracy of the movement performed. But the calibration parameters of a prosthetic device are relatively abstract and do not directly correspond to those quantities that the users and the clinicians inherently care about. Here, we propose an intuitive tuning technique that allows clinicians to tune prostheses based on the things that end-users actually care about. We use well-established engineering techniques (optimal control) to determine the set of best possible solutions for different relative preferences of the user. This required optimizing the problem for multiple objectives, (effort, time, reliability, and accuracy) to compute the best tuning parameters for a wide range of trade-offs. By solving this optimization problem, the complexity of the relationship between the performance and the prosthesis parameters can be implemented as a mapping procedure, and thereby hidden from the user. This simplifies the calibration process and allows clinicians or users to intuitively customize the device for their individual needs.

INTRODUCTION

Biological movement and motor coordination can be thought of as optimization tasks that minimize the cost of effort and time while maximizing the reward obtained from performing the movement [1], [2]. The costs are mathematical representations of quantities that the brain tries to minimize when generating any motor command to move our body. Several different cost functions like effort, metabolic energy, and endpoint variance have been used to describe specific movements. This inconsistency in literature actually suggests that humans optimize a combination of different costs [3] and simply change their cost priorities to perform different tasks. The composite costs of effort, accuracy, reliability and time sufficiently describe how we consistently coordinate our joints to perform different tasks [4]. This cost preference also changes from person to person. Suppose we ask a group of people to write by hand the entire abstract of an article- some might care more about the time spent on the task and write as fast as they could, while others might care more about the reliability of the outcome and don't mind spending a few extra minutes.

Clinical motivation

Let's take the example of driving a car. We usually care about things like fuel efficiency, comfort, safety, and the dynamic response of the car. The input parameters like the steering force, powertrain characteristics, suspension control, two-wheel and four-wheel drive modes can be adjusted to reflect our personal priorities in terms of the things we care about when driving. An experienced driver might be able to easily tune these parameters to get the desired response. But for new drivers, this could be a daunting task. The driving mode options provided by most car manufacturers nowadays, simplifies this task for both new and experienced drivers. The different modes like the eco mode, comfort mode or the sport mode speak the language of the user and directly convey the information in terms of things they care about.

Likewise, both clinicians and end-users of prosthesis care about the costs of effort, time, accuracy, and reliability incurred by the users when making a movement with a prosthetic device. The input parameters for a myoelectric prosthesis are abstract quantities like the device gain, amplifier thresholds, and control mapping paradigms like proportional position or velocity control. But unlike the example with the cars, the clinicians are not provided with a set of "driving mode" settings that simplifies the relationship between the cost space in which they care about and the input parameter space in which they work.

Moreover, there is an inherent trade-off in the balance of these cost preferences. That is, we cannot have the best of both worlds and improve both the speed and the accuracy of our movements simultaneously. When larger control signals are produced to make faster movements, the multiplicative nature of the noise in our myoelectric signals deteriorates the accuracy and the reliability of the movement. The device gain parameter should hence be adjusted to a sweet spot that best reflects the user preference. But there is another catch, some combinations of the input parameters always produce results that are worse for all possible user preferences. For example, proportional velocity control always performs better than position control in terms of both the costs of effort and reliability. Conventional tuning techniques force the clinicians to tune a limited number of abstract parameters that do not directly reflect the cost space that they care about, and make it much harder to tune for the optimal set of parameters for a given user.

So when tuning a prosthetic device, it will be beneficial to avoid the sub-optimal input parameters to ensure that the user gets the best experience. Optimization methods or optimal control techniques can be used to identify the optimal input parameters for each individual user preference. This approach makes the tuning procedure much more intuitive for the clinicians as the abstract device input parameters can be computed and set by an algorithm that optimizes based on the costs that users inherently care about.

Background

We care about a variety of things when making a movement. As we have multiple objectives that we wish to optimize, we need to perform multi-objective optimization (MOO) to find the best set of tuning parameters for each user's personal preference. For a MOO, no single solution can be best with respect to all the conflicting objectives and we have several optimal solutions instead. The optimal solutions are those in which we cannot further reduce the cost of one objective function without increasing the cost of another. These optimal solutions are called the Pareto solutions, or the Pareto set [5].

There are two broad strategies for obtaining these Pareto solutions for MOO problems.

1) The first method is to scalarize the different objectives and to repeatedly solve for the entire range of cost- preferences.

   In terms of our prosthesis tuning example, this would mean that the different objective functions of effort, time, reliability, and accuracy get added up with relative weights that represent the user's preference for the different cost functions. The computation is then repeated for every user's individual preference.

   For example, if the total cost (J) is represented as a sum of two independent costs (J_1 and J_2):

   (i) $J = J_1 + J_2$, represents an equal preference for the two costs.

   (ii) $J = J_1 + 10\,J_2$, represents that the user cares about 10 times as much about the second cost when compared to the first one.

2) The second method is to find multiple Pareto-optimal solutions in a single run, without any prior information about the relative preference of the different costs.

   For our prosthesis tuning problem, this would mean that we compute a set of optimal tuning parameters without any reference to the user's individual preference for the different costs. This method can be represented mathematically as:

   J = [J1, J2], where we simultaneously try to optimize for both the cost functions J1 and J2.

The first strategy expresses the user's preference in terms of simple relative weights that change the optimum tuning parameters for the individual. But the tuning parameters must be recomputed every time the user's preference changes. The interpretation of relative weights also becomes incorrect when the multiple objective functions are not normalized appropriately [6]. The second strategy has the advantage of solving multiple optimal tuning conditions in a single simulation run, but it is not possible to incorporate the user preference into the algorithm. Our third option is optimal control which best approximates how humans move their joints and control human-machine interfaces [7], but the disadvantage is that this method requires a single objective function. Hence, we decided to blend the three ideas to get multiple meaningful Pareto solutions for the tuning parameters such that they can be saved in a look-up table to avoid re-evaluation.

Technical challenge

There are infinite solutions that satisfy the Pareto optimality condition and form the Pareto front. Ideally, we would like to obtain the optimal tuning parameters for a finite number of points on the Pareto front. These Pareto solutions can be saved in a look-up table that can aid with prosthesis tuning. Due to the inherently nonlinear nature of the human movement cost functions, an equally spaced set of relative weights does not produce a uniform Pareto set. Figure 1 shows an example in which only the costs of effort and accuracy of the movement were considered, for a first-order dynamic model of the prosthetic device. In this case, the Pareto solutions found are clustered towards one end of the Pareto front in which the effort cost is much lower than that of the accuracy cost. This means that a step-change in the user's preference for the two costs will lead to a larger change in the cost of effort when compared to the cost of accuracy. This also indicates that the problem is not accurately normalized (and that it requires nonlinear normalization mapping).

To describe the priorities meaningfully in terms of the cost, we need to flip the problem on its head and obtain a set of evenly distributed Pareto points that correspond to specific user priorities. This will allow us to bin the different regions of the Pareto front into "tuning modes" similar to the driving mode options provided by car manufacturers.

In addition to conveying the tuning information in the language of the user, this technique allows us to entirely avoid normalization. In this article, we propose a gradient-based approximation technique that can be used to produce an evenly distributed set of points on the Pareto front.

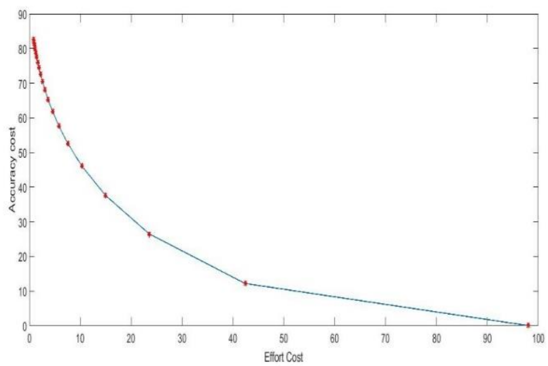

Figure 1: Pareto front for a set of cost preferences linearly spaced between 0.01 and 10. The red asterisks indicate the location of the obtained Pareto set.

METHODS

The main aim of this study is to determine if a uniform set of Pareto solutions can be obtained for a generic optimal control problem that has multiple objective functions. Optimal control guarantees the best solution, but it requires a single objective function, which is at odds with our attempt to enable users engage with multiple objectives. To solve this problem, we strategically assign weights across the multiple objectives to enable them to be considered as a single objective (which can then be solved using optimal control). Because this scalar composite cost can be solved using optimal control, it is guaranteed to land on the Pareto front, but where it lands on the Pareto front depends on the weights we choose. In order to strategically assign those weights to ensure an evenly distributed set of Pareto front, we use a crowding metric to decide where on the Pareto front we would like to land next, and then estimate the weights that should get us in that ballpark using a gradient-based approximation technique. This process is further described in the section on the Gradient-approximation technique.

To demonstrate the feasibility of the blended optimal control and MOO based approach, we use a simple optimal control example that relates to our ultimate aim of clinical prosthesis tuning. The example is that of a simple human-machine interface that is used to perform a target reaching task. In order to keep the problem simple and retain the multi-objective nature of human movement, only the two contrasting objectives of effort and accuracy were used to compute the cost incurred to the user. The human-machine interface model contained two parts: the human component that modelled the user's motor commands, and the prosthesis component that performed a reaching movement in response to the user's control signal. The system was assumed to be deterministic and the potential uncertainties in the control signal and the environment were not modelled.

The human component of the model produces an optimal control signal (u) that minimizes the composite cost of effort and accuracy to the user. The machine or the prosthesis component was simulated using standard zero, first or second-order dynamic system models. The tuning parameters of the prosthesis were not optimized in this study to reduce the number of optimization parameters. The duration of the movement was fixed to be a single time step for all conditions. The model was entirely implemented in MATLAB (Release 2016a, The Mathworks, Inc., Natick, MA.)

Cost Function and User Priority

The cost of effort was defined as the squared control signal u that represents the magnitude of the myoelectric signal from the user.

$$J_u = u^2 \tag{1}$$

The cost of accuracy penalizes based on the error between the target (represented as g) and the movement endpoint.

$$J_a = [g - x(p)]^2, \tag{2}$$

where $x(p)$ is the position of the simulated device at the end of the movement and the final time p is set to one for all the simulations without loss of generality.

The total cost is a weighted sum of the accuracy and the effort costs and is represented by:

$$J = \alpha . J_u + J_a, \tag{3}$$

where α shows the user's relative preference for the two costs. A large value of α shows that the user would rather minimize their physical effort even if it means that they don't reach the target accurately. A small α value indicates that the user prioritizes the endpoint accuracy and won't mind spending more effort. For a perfectly normalized set of costs, an α value of one will indicate that the user cares about the two costs equally. But as the magnitudes of the cost values are highly task-dependent, normalization was not performed for our system.

Gradient-based approximation technique

The purpose of this algorithm boils down to two simple things – selecting the next point on the Pareto front that needs to be populated and landing there by computing the required α value. The distance (Mahalanobis form) between consecutive points was used as a measure of crowding and the next point was selected to ensure an even distribution of points on the Pareto front. We need to compute the desired user preference level or the α value to land at these points and a simple gradient approximation technique was used to achieve this. The required user preference value was calculated using the following equation.

$$\propto_{Desired} = \frac{\left(J^i_{Desired} - J^i_{Current}\right)}{\frac{\partial J^i}{\partial \alpha}\big|_{J^i_{Current}}} + \propto_{Current} \quad , \qquad (4)$$

where J^i corresponds to the individual costs of the effort or accuracy objective functions, and $\frac{\partial J^i}{\partial \alpha}\big|_{J^i_{Current}}$ refers to the sensitivity of the individual cost with respect to changes in α, computed at the current Pareto solution. These sensitivities are obtained by perturbing the α values at the different Pareto solutions obtained and are essentially just numerical approximations of the gradient at those points.

The two extreme points on the Pareto front that correspond to slopes 0.01 and 100 were picked heuristically by tuning the α values for the given task. After this, the "selection" and the "landing" algorithms were used iteratively to obtain a uniformly distributed Pareto set.

RESULTS

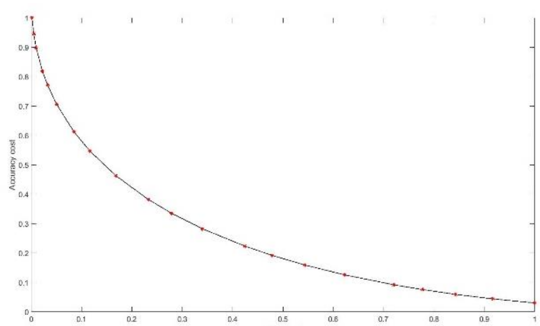

Figure 2: Uniform Pareto front obtained using the proposed algorithm

In order to ensure that a uniform Pareto front can be produced using our algorithm, a variety of scenarios were tested. The prosthesis was modelled as a standard zero, first, or second-order dynamic system. The proposed algorithm was able to successfully produce a uniform distribution on the Pareto front for all three cases. Figure 2 shows an example of the Pareto solutions found for a first-order dynamic model. These results show that the algorithm is generalizable and can be applied to optimize the tuning parameters for a variety of different user preferences.

DISCUSSION

The aim of this study was to determine if we can generate sufficient Pareto solutions for a human-machine interaction model, such that we can adequately describe different user preferences. Our simulations demonstrate the feasibility of the proposed method and show that it is robust for a variety of scenarios. The concept of relative tuning that we have described in this article could allow intuitive prosthesis calibration in terms of quantities that the clinicians and the patients care about. It will also permit the device to be tuned by the user. The user-tunability function can allow them to optimally perform vastly different tasks like painting and yard work, which requires them to change their personal cost priorities.

Limitations

As the intention of this article was to understand if a uniform set of Pareto solutions can be formed, the mathematical model was simplified to reduce the computational complexity of the problem. For example, the system was assumed to have no noise and the tuning parameters of the prosthesis model were not optimized.

REFERENCES

[1] E. Summerside, R. Shadmehr, and A. A. Ahmed, "Vigor of reaching movements: reward discounts the cost of effort.," *J. Neurophysiol.*, pp. 1–30, 2018.

[2] B. Berret and F. Jean, "Why Don't We Move Slower? The Value of Time in the Neural Control of Action," *J. Neurosci.*, vol. 36, no. 4, pp. 1056–1070, 2016.

[3] B. Berret, E. Chiovetto, F. Nori, and T. Pozzo, "Evidence for composite cost functions in arm movement planning: An inverse optimal control approach," *PLoS Comput. Biol.*, vol. 7, no. 10, 2011.

[4] R. Shadmehr and S. Mussa-Ivaldi, *Biological Learning and Control*. 2012.

[5] K. Deb and D. Kalyanmoy, *Multi-objective optimization using evolutionary algorithms*. John Wiley & Sons, 2001.

[6] R. T. Marler and J. S. Arora, "The weighted sum method for multi-objective optimization: New insights," *Struct. Multidiscip. Optim.*, vol. 41, no. 6, pp. 853–862, 2010.

[7] E. Todorov, "Optimality principles in sensorimotor control," *Nature Neuroscience*. 2004.