Morphological Computation: The Good, the Bad, and the Ugly

Keyan Ghazi-Zahedi[1] Raphael Deimel[2] Guido Montúfar[1] Vincent Wall[2] Oliver Brock[2]

***Abstract*— In many robotic applications, softness leads to improved performance, robustness, and safety, while lowering manufacturing cost, increasing versatility, and simplifying control. The advantages of soft robots derive from the fact that their behavior partially results from interactions of the robot's morphology with its environment, which is commonly referred to as morphological computation (MC). But not all MC is good in the sense that it supports the desired behavior. One of the challenges in soft robotics is to build systems that exploit the morphology (good MC) while avoiding body-environment interactions that are harmful with respect to the desired functionality (bad MC). Up to this point, constructing a competent soft robot design requires experience and intuition from the designer. This work is the first to propose a systematic approach that can be used in an automated design process. It is based on calculating a low-dimensional representation of an observed behavior, which can be used to distinguish between good and bad MC. We evaluate our method based on a set of grasping experiments, with variations in hand design, controller, and objects. Finally, we show that the information contained in the low-dimensional representation is comprehensive in the sense that it can be used to guide an automated design process.**

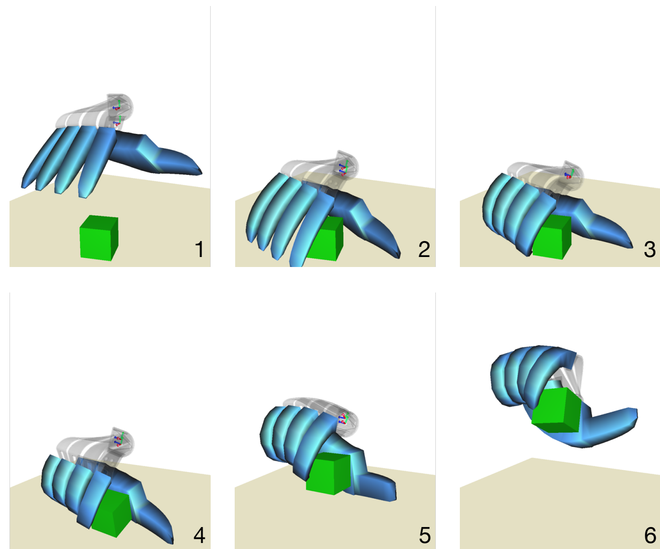

Fig. 1. Typical RBO Hand 2 grasp motion used to identify good and bad MC. The grasp shown here is an instance of good MC, which means that the compliance of the hand contributed to a firm grasp with a very simple controller. Bad MC can be observed, if the object is not tightly held which leads an increased interaction of the hand and the object, e.g., by slowly slipping out of the hand. The goal of this paper is to differentiate the two and make the results applicable in an automated design process.

I. Introduction

Soft robotics is a successful branch of robotics. In many applications, softness leads to improved performance, robustness, and safety, while lowering manufacturing cost, increasing versatility, and simplifying control [1], [2]. In spite of these advantages, there currently is no systematic method for exploiting the benefits of softness in robot design. At the moment, human designers rely on experience and intuition to design competent soft robots.

The advantages of soft robots derive from the way their behavior is generated. As with traditional robots, the behavior of soft robots is affected by the control commands a robot receives. However, this control-based behavior is modified through compliant interactions of the robot with its environment. These compliant interactions adapt the behavior to a particular context, without the need for explicit control. It is therefore important to note that the behavior of soft robots is *not* exclusively the result of control, it partially results from interactions of the robot's morphology with its environment. This latter part of the robot's behavior, stemming from interactions, is referred to as morphological computation (MC) [3], [4].

[1] Information Theory of Cognitive Systems, Max Planck Institute for Mathematics in the Sciences, Germany

[2] Robotics and Biology Laboratory, Technische Universität Berlin, Germany

We gratefully acknowledge financial support by the European Commission (SOMA, H2020-ICT-645599) and the German Priority Program DFG-SPP 1527 "Autonomous Learning".

But not all MC is good. The interactions between a soft robot and the environment can also be harmful, for example, if it un-does what control accomplished or simply causes failure. We call this *bad* or *ugly* MC. Of course, MC can be good if the control-based behavior is modified in a favorable way (these informal definitions of good, bad, and ugly MC will be stated more precisely in Sec. II). To illustrate this with an example from soft manipulation: If MC leads to the adaptation of a soft hand to the shape of an object that results in a good grasp (see Fig. 1), we call that *good* MC. If the compliance of the fingers lead to a less firm grasp, we consider this *bad* MC. Both forms of MC describe hand-object interactions, but only the former is desirable, while the latter is to be avoided.

The automated design of soft robots must minimize bad MC and maximize good MC, relative to a particular task. In this paper, we propose for the first time a method to identify good and bad MC from observed behavior. If the observed behavior can be represented in some high-dimensional space, our method identifies sub-spaces associated with good MC and sub-spaces associated with bad MC. Such a criterion is a first and important step towards a quantitative design of soft robots.

We apply our method to soft grasping using an anthropomorphic robot hand based on pneumatic soft continuum actuators, known as RBO Hand 2, [5]. Compared to a rigid manipulator, the design of a soft manipulator that is able to safely grasp a variety of objects is less obvious. To visualize this point, one can image a balloon slowly being filled with air. If otherwise unconstrained, the balloon will expand almost equally in all directions. For RBO Hand 2's fingers, this would be an undesired behavior. Hence, a thread was carefully wrapped around the fingers in such a way that it allows and expansion of the dorsal (outer) side of the fingers, while it suppresses an expansion on the ventral (inner) side [5]. As a result, the hand closes when the air pressure is increased. Stated otherwise, certain degrees of freedom (DOF) of the fingers were restricted, while others where retained. The method proposed in this paper is designed to automatically detect which DOFs should be suppressed and which should be enhanced in order to achieve a desired behavior with minimal control.

The method consists of two steps. First, the components of the behavior that can only be attributed to the physical hand-object interaction are extracted from the data. Second, the covariance of the DOFs is calculated. We show, that this form of dimensionality reduction can be used to distinguish between good and bad MC in a way that is applicable in an automated design process. The method is not limited to soft manipulation but can be applied to soft robotics in general.

The next section presents MC in more detail, together with a quantification for it. The quantification allows a more formal understanding of good and bad MC. This is followed by a presentation of our method, the experimental results, and a discussion, before this work closes with conclusions.

II. Good vs. Bad Morphological Computation

It has been shown that the physical properties of the morphology (and their interaction with the environment) can perform functions that are normally attributed to the brain, and thereby, reduce the required computational complexity significantly [3], [6], [7]. This is known as morphological computation (MC) [4], [3].

Good MC are body-environment interactions which contribute to the desired behavior in a way that reduces the required controller complexity. Bad MC are body-environment interactions which make the desired behavior more difficult to maintain (e.g., object slipping out of the hand due to the hand's compliance). We may define ugly MC as body-environment interactions, which are neither clearly good or bad. This will be discussed in more detail below (see Sec. VII).

A. Quantifying Morphological Computation

We proposed an information-theoretic quantification of MC in [8] that was successfully evaluated on muscles models [9]. In what follows, we will describe the concept behind the measure used in this work based on the causal model of the sensorimotor loop (SML) [10]. The left-hand side of Fig. 2 displays a schematics which shows the interaction of the agent's controller, sensors, actuators, and environment. It is important to note here, that we use the term *world* to capture the body and environment (which is in alignment with the agent-environment distinction used in reinforcement learning). The right-hand side of Fig. 2 shows the causal graph for the SML of a reactive system, in which the random variables W, S, A, and W', refer to the current world state, current sensor state, current actuator state, and finally, the next world state (a discussion about non-reactive systems with respect to MC is found in [8]). In the context of this work, the world states W and W' are defined as the Cartesian coordinates of each simulated joint (see Sec. V), and the action state are the four air pressure channels that control the motion of the hand. The causal graph is completed by the three kernels α, β, and π, which represent the world dynamics, sensor model, and policy. We denote a random variable with a capital letter X, its instance with a lowercase letter x, and its alphabet with a calligraphic letter $\mathcal{X}$. Hence, the random variable X takes values $x \in \mathcal{X}$.

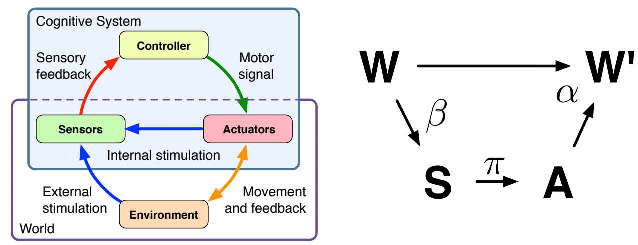

Fig. 2. Left: Schematics of the sensorimotor loop (SML), which shows the interaction between the controller/brain, actuators, sensors, and environment (adapted from [11]); right: Causal model of the SML. W, S, A, W' refer to the world, sensor, and actuator state of the current time step and the world state of the consecutive time step

Let us assume that the world state at a given time step does not have any influence on its state at the next time step. This means, that in the diagram shown in Fig. 2, the causal link between W and W' is not present. For this case, we would argue that there is no MC, because the robot's state is fully determined by the controller's action. The quantification $\mathrm{MC_W}$ uses the Kullback-Leibler divergence $D_{KL}(p(w'|w,a)||p(w'|a))$ (which corresponds to the conditional mutual information $I(W';W|A)$) to measure how much the observed behavior differs from the case in which we assume that the world does *not* influence itself. It is formally given by:

$$\mathrm{MC_W} = \sum_{w',w,a} p(w',w,a) \log_2 \frac{p(w'|w,a)}{p(w'|a)}. \quad (1)$$

A full discussion of the measure and algorithm can be found in [8], [9]. We can now describe good and bad MC more formally in the next section.

B. Formal discussion of good and bad MC

The Eq. (1) can be rewritten in the following form

$$\mathrm{MC_W} = I(W';W|A) = H(W'|A) - H(W'|W,A). \quad (2)$$

This shows that $\mathrm{MC_W}$ is maximal, if the conditional entropy $H(W'|A)$ is maximal (the influence of the second term is discussed below), which means that the uncertainty about the next world state W' given an action A must be large. In the context of this work, maximizing $H(W'|A)$ could result from a soft manipulator which changes its state in an uncontrolled fashion, e.g. a hand which vibrates every time it moves. This is something that should clearly be avoided. The second term restricts this to some extent by forcing the reduction of next world state's uncertainty given the previous world and action states, but the argument still holds in principle. Not all behaviors that maximize the conditional entropy $H(W'|A)$ while minimizing $H(W'|W, A)$ will be beneficial, i.e. good in all scenarios.

The measure given in Eq. (1) is used in the analysis of our method below. It is formalized for discrete systems only. Hence, the data need to be pre-processed before $\mathrm{MC_W}$ can be applied. This is discussed next.

C. *Calculating* $\mathrm{MC_W}$ *on the recorded data*

Estimating information theoretic quantities based on limited samples in high dimensional spaces is a challenging problem, subject of ongoing investigations, and beyond the scope of this work. To account for this problem in our case, we chose to estimate MC only based on the Cartesian coordinates of the finger tips. Eq. (1) is defined for discrete systems, which means that a binning is required. We empirically determined 300 bins for each coordinate and 10 bins for the controller commands (see [9] for discussion).

III. Identifying Good and Bad MC

The goal of this work is to find a method to systematically identify DOFs that contribute to a high grasp success (good MC), i.e. that allow for a high versatility with a simplified control and to distinguish them from bad MC (those DOF that impede good grasping). In other words, we want to enhance compliance that contributes to a good grasp and reduce compliance that reduces the grasp success. A scenario is an automated design process in simulation, in which a morphology is evaluated based on grasp attempts of several objects. In a purely evolutionary setting, the morphological parameters, e.g. the stiffness of the DOFs, would be open to random modifications. Soft manipulators inherently have many DOFs, which renders such an approach barely practical. Instead, a method to extract the characteristics of a grasp behavior is required that can be used to guide an automated optimization process.

In the following sections, we first motivate covariance matrices as a method of dimensionality reduction and a way to extract the characteristics of motions. The calculations are performed on pre-processed data that only contain the hand's motion which results from the interaction of the hand with the object (see Sec. III-B). To avoid artifacts, only a portion of the recorded grasps was used, which is described the last segment of this section.

A. *Characterizing behaviors by their covariance*

This section describes covariance matrices as a method of dimensionality reduction that captures the characteristics of grasps, and hence, can be used to distinguish between good and bad MC. This is motivated by other work (e.g. [12]), which showed that covariance matrices can be used to obtain a meaningful hierarchical clustering of, e.g. humanoid behaviors. In this particular example, different crawling behaviors were clustered closer together based on the covariances. The distance between clusters reflected the difference in the behaviors, e.g. crawling vs. climbing out of a pit.

The hypothesis is that large covariance coefficients which are shared among experiments with high grasp success relate to useful compliance and should be reinforced during the design process. Likewise, large coefficients shared among unsuccessful grasps should be suppressed, because they indicate harmful compliance. We will discuss this in more detail further below (see Sec. VII). We first describe how we calculate the covariance matrix of a grasp, before we discuss how the data are pre-processed. The covariance matrix of a data set $B \in \mathbb{R}^{T \times N}$ is given by

$$C(B) = (c_{ij})_{i,j=1,2,\dots,N} = \frac{1}{T}\sum_{t=1}^{T}(B_{t,i} - \bar{B}_i)(B_{t,j} - \bar{B}_j),$$

where $B_{t,j}$ refers to the entry in the t-th row and j-th column of B, and $\bar{B}_j = \frac{1}{T}\sum_{t=1}^{T} B_{t,j}$ is the mean value of the j-th column. In the context of this work, the matrix B contains the x, y, z values of each coordinate frame of the simulated hand (see Fig. 3) over time. The index j refers to each coordinate, hence $B_{t,1}, t = 1, \dots, T$ is the recorded data for x_1 for an entire grasp. This means that the covariance coefficient c_{36} contains the covariance of the first and second frame's z coordinate. Hence, a large positive coefficient c_{36} would mean that the movements along the z-axis of the first and second frame should be highly correlated for a successful grasp. In an automated setting, the coupling of these two frames' movements would be enforced by increasing the stiffness along these DOFs.

B. *Extracting motions that result from MC*

We are interested in the interaction of a soft manipulator with an object. Therefore, we pre-process the data before calculating the covariance matrices in the following way. For each combination of a RBO Hand and controller (see below, Sec. V), we first record the prescriptive behavior, which is the hand's movement without any graspable object present in the scene. We refer to this data set as B_p. The recording of the grasp itself is denoted by B_g. The element-wise difference of these two behaviors describes movements of the coordinate frames that are the result of the hand's interaction with the object and it is denoted by $B = B_g - B_p$. Hence, the covariance matrix $C(B)$ contains information about correlations of movements that only relate to the soft manipulator's compliance.

C. Avoiding artifacts

A potential concern is that all recordings in which the object was dropped early or not grasped at all, will not differ significantly from the prescriptive behavior, and hence $B = B_g - B_p$ will be mostly zero. This can be avoided if only a fraction of the time steps are taken into account, which is why we evaluated different time frames in our analysis. We chose to present the results based on 75 time steps (excluding the 10 approaching time steps, see Sec. V) to avoid a clustering based on grasp success vs. grasp failure as well as a clustering only reflecting the object's shape.

The next section describes the clustering method that was used to visualize the similarities of different grasps based on their covariance.

IV. Visualizing clusters of grasp behaviors

We used t-SNE [13] to obtain a two dimensional representation of the covariance matrices C for visualization. This method, t-Distributed Stochastic Neighbor Embedding, constructs pairwise similarities of the input data and visualizes them in n dimensions, where $n = 2$ is chosen in this work. These visualizations capture the local structure of the data, while also revealing global structure such as the presence of clusters at several scales. The results section will show the obtained clusters colored with the grasp success (explained next) and $\mathrm{MC_W}$ (see Sec. II-A), but this information was not used during the clustering itself.

A. Determining grasp success

The grasp success is determined by the average distance of the object to the hand during the last 10 time steps of each recorded behavior. The distance is measured between the object's geometric center and the coordinate frame located at the first frame of the second finger (see Fig. 3). The object size is then subtracted from the measured distance to ensure comparability of the results for all objects. It was discussed above, that the covariance matrices were only calculated on a subset of the recorded data. This is not the case for the grasp distance, which was always estimated on the last 10 time steps of each full recording.

The next section discusses how the data were acquired.

V. RBO Hand Grasp Simulations

We used a simulation of the RBO Hand in order to obtain a large data set of fully observed grasping behaviors (see Fig. 1). Recent improvements to simulation algorithms [14] enable us to simulate complete grasp attempts in near real-time. The simulator is implemented with the SOFA framework [15] and relies on its Compliant module [14]. In this setup, soft hands are modeled as a tree of Cosserat beams (i.e. kinematic chains with ball joints), to which a collision surface is attached [16]. Fig. 3 (right-hand side) shows the actual simulation model. Large trihedra indicate the links (beam elements), purple trihedra the joint location between two adjacent links. The surrounding wire-frame is the collision mesh, which is attached to the links via linear blend skinning and follow the motion of the actuator's "backbone". Mechanical parameters are computed using a recently published model [5].

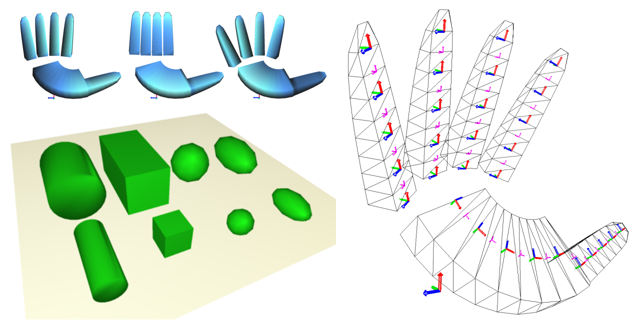

Fig. 3. Left: different hand morphologies and the set of objects used during the experiments; right: Illustration of the model used to simulate soft hands: The trihedra indicate discrete links, adjacent ones are connected by passively compliant ball joints (purple trihedra) in between. The surface mesh is attached to the frames, and is used to compute collision and model contact.

Using simulation makes it especially easy to obtain motion data, e.g. fingertip frame motion. In addition, simulations can be run in parallel, which results in a much larger data set to conduct the investigation on. Another advantage of using simulation is that the hand morphology can be changed easily. For this paper, we created two variations of the RBO Hand 2 (see Fig. 3, left-hand side) and simulated them as well as the original hand. These three hands were combined with three distinct motion primitives that implement variations of the surface-constrained grasp (see Fig. 1). This yields nine hand-controller combinations to compare. In the morphology domain, the spread between the four fingers was varied, and in the control domain the roll angle of the wrist during the grasping motion was varied in $\{-15°, 0°, 15°\}$.

A. Simulation data

In total, the simulated hand consists of 32 coordinate frames (see Fig. 3). The dynamics of each finger (and thumb) are modeled by five coordinate frames. The remaining 7 coordinate frames are distributed along the palm and used to control the wrist and arm motion. For each frame, the simulator records the pose, of which only the 3D positional data are used in this work. The underlying assumption is that the orientation of the coordinate frames can be reconstructed from the x, y, z coordinates of consecutive frames. Furthermore, as we are interested in analyzing grasping, we transformed all coordinate frames into the wrist frame, which means that the state of the hand is given by 31 coordinate frames, and hence, by a vector $b_t \in \mathbb{R}^{93}$.

We recorded the data for different hand designs, controllers, objects, and object's initial position (see Fig. 3 and previous section). The data include nine different hand-controller configurations, eight different objects, and finally, 27 different initial positions (x, y, θ) for each object, resulting in total 1944 recorded behaviors.

Each grasp is recorded for 300 time steps with a step width of 0.01 s and can be divided into four phases: 1. approaching of the object (hand moves downwards for 10 time steps),

2. grasping (30 time steps), 3. lifting (30 time steps), and finally, 4. evaluation of the grasp stability (final time steps). Hence, each grasp is captured in the matrix $B_g = \mathbb{R}^{300\times 93}$.

VI. Results

We present the covariance matrices of the grasps clustered with t-SNE (see Sec. IV) and colored by grasp success, morphological computation, object type, and object initial position. This allows to understand what kind of information is stored in the covariance matrices and if it can be used to distinguish between good and bad MC. For the sake of clarity and because of space restrictions, we decided to only plot and discuss the results for one RBO Hand and one controller as the results hold for every other combination too.

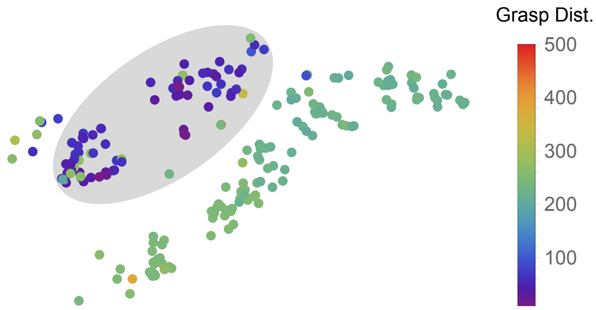

(a) Clustering of grasp behaviors, colored by grasp distance. Smaller values (blue) correspond to better grasping The instances within the gray ellipse are those of interest, with high grasp success and high MC.

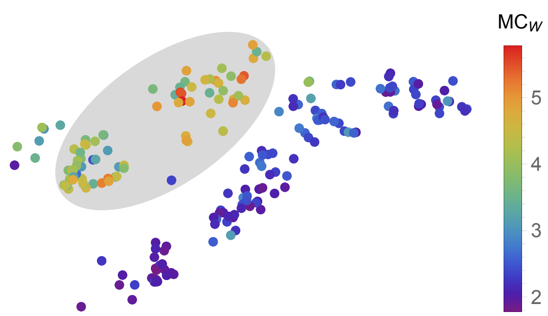

(b) Clustering of grasp behaviors, colored by $\mathrm{MC_W}$. Larger values (red) relate to better grasping. The instances within the gray ellipse are those of interest, with high grasp success and high MC.

Fig. 4. Clustering of grasp behaviors, colored by grasp success (a) and $\mathrm{MC_W}$ (b)

The main results are shown in Fig. 4. Fig. 4(a) shows the clusters colored by grasp success and Fig. 4(b) shows the clusters colored by $\mathrm{MC_W}$. By comparing the two plots, it is seen that the two large clusters resulting from the covariance matrices, can be very well explained by good grasping with high MC (highlighted with gray ellipse in background) and unsuccessful grasps.

To verify that the clustering cannot be equally well explained by the object shape or object initial position, we also plot the clusters colored by these two parameters. Fig. 5 shows that neither the object's shape nor its initial position can be used to fully explain the clustering. It also seems that the two sub-clusters of the cluster with high grasp success and high MC can partly be explained by the object's shape. This means that, as can be expected, the covariance matrices not only contain information about correlations of DOFs that contribute to good grasping with high MC, but also some information about the grasped object. The object's initial position (see Fig. 5) is least informative in describing the two major clusters.

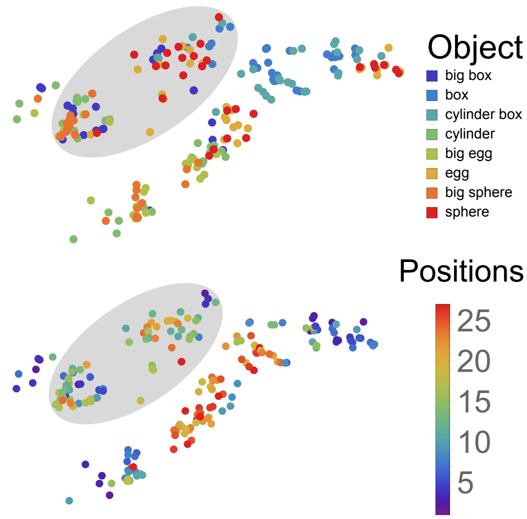

Fig. 5. Top: This figure shows that there is some form of distinction possible based on the shape of the grasped object. Yet, the objects do not explain successful vs. unsuccessful grasping and high vs. low MC (compare with Fig. 4(a) and Fig. 4(b)); bottom: This figure shows that the clustering cannot be explained by the object's initial position.

VII. Discussion

Fig. 4 and Fig. 5 showed that a clustering based on the covariance matrices of the difference between grasping and its corresponding prescriptive behavior leads to clusters which are best explained by high grasp success and high MC (based on the quantification $\mathrm{MC_W}$). The local relation of the instances reflect the similarities of the corresponding covariance matrices, which is what was referred to as sub-spaces of good and bad MC in the introduction.

If we look at representative examples of covariance matrices from sub-spaces with high density, two conclusions can be drawn. First, the matrices have a regular structure, which means that there is a high regularity in how the DOFs interact. Second, successful grasps have stronger positive covariance coefficients and weaker negative coefficients compared to less successful grasps. *Ugly* MC can now be defined as large coefficients, which are similar for successful and less successful behaviors, and hence, are not conclusive. To summarize, we can identify coefficients that relate to *good*, *bad*, and *ugly* MC in clusters with high density. Together with the regularity of the matrices, this suggests that the covariance coefficients can be used in an automated design process to guide modifications of the morphology (which is discussed next).

A. Covariance coefficients can guide an automatic design process

Positive coefficients correspond to DOFs which increase and decrease together. Negative coefficients correspond to DOFs which act reversely, i.e. if one increases, the other decreases. The latter could correspond to a finger movement, in which one of the coordinate frames moves upwards (positive z-movement) while the other coordinate frame moves downwards (negative z-movement). This would be an example of compliance that un-does what the controller tried to achieve. As the hand closes (e.g. negative z-movement of one the finger's coordinate frame) the "harmful" compliance of the finger (e.g. positive z-movement of another coordinate frame) prohibits a firm grasp.

A comparison the covariance matrices reveals that stronger negative coefficients correspond to less successful and stronger positive correlations correspond to successful grasping. Given the density of the sub-spaces (see Fig. 4), it should be possible to identify positive coefficients which are most dominant over all successful grasps. The related DOFs should be enforced, e.g. by increasing the stiffness between them. Strong positive coefficients that are dominant in all unsuccessful grasps mean that the relative motion of the related DOFs should be softened. Negative coefficient can be used analogously. Hence, the dimensionality reduction based on covariance matrices can support a systematic modification of the morphological properties (e.g. stiffness) of a soft robot in an automated design process.

The final section concludes this work and gives an outlook on currently ongoing work.

VIII. Conclusions

Currently, the success of a soft robot design relies on the expertise and intuition of its designer. The reason is that, up to this point, there is no systematic way to analyze how the compliance of a soft robot contributed to a desired behavior. Such a systematic approach is the first required step in an automated design process of soft robots. This work is the first to propose such a method and to discuss how it could potentially be used to guide an automated design process.

We described physical processes which result from the interaction of the soft materials with the environment and are beneficial as good morphological computation. Naturally, in soft robots, there are body-environment interactions that are harmful, which means that they might render the actions sent to the robot worthless. We referred to this form of compliance as bad morphological computation. Hence, for an automated design process of soft robots, we need a systematic way to identify good and bad morphological computation based on observations of the robot's interaction with its environment.

For this purpose, we conducted a series of grasp experiments, with variations of RBO Hand 2 shape, object, object initial position, and controller. We showed that the covariance matrices calculated on the difference of grasping to prescriptive behavior contain information that allows to distinguish between good and bad morphological computation based on observations alone. We discussed how the covariance coefficients relate to the compliance of the corresponding degrees of freedom and how the coefficients can be used to guide an automated design process.

The next step is to evaluate the proposed method in an automated process to design a soft manipulator for a specific task. To this point, we are able to identify good and bad MC, i.e. to distinguish DOFs which support the desired behavior form those which harm the success. The final step is to use this information in a simulated set-up to modify the parameters of, e.g. a simulated hand. In particular, the covariance coefficients of grasps in clusters with high density (small differences between the C-matrices), that also have high MC and a good grasp, will be used to modify the softness/stiffness of RBO Hand 2's DOFs. This is the topic of currently ongoing research.

References

[1] A. Verl, A. Albu-Schäffer, O. Brock, and A. Raatz, Eds., *Soft Robotics: Transferring Theory to Application*. Springer, 2015.

[2] J. Hughes, U. Culha, F. Giardina, F. Guenther, A. Rosendo, and F. Iida, "Soft manipulators and grippers: A review," *Frontiers in Robotics and AI*, vol. 3, p. 69, 2016.

[3] R. Pfeifer and J. C. Bongard, *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, MA: The MIT Press (Bradford Books), 2006.

[4] C. Paul, "Morphology and computation," in *Proceedings of the International Conference on the Simulation of Adaptive Behaviour*, Los Angeles, CA, USA, 2004, pp. 33—38.

[5] R. Deimel and O. Brock, "A novel type of compliant and underactuated robotic hand for dexterous grasping," *The International Journal of Robotics Research*, vol. 35, no. 1–3, pp. 161–185, 2016.

[6] D. Polani, "An informational perspective on how the embodiment can relieve cognitive burden," in *Artificial Life (ALIFE), 2011 IEEE Symposium on*, 2011, pp. 78–85.

[7] G. Montúfar, K. Ghazi-Zahedi, and N. Ay, "A theory of cheap control in embodied systems," *PLoS Comput Biol*, vol. 11, no. 9, p. e1004427, 09 2015.

[8] K. Zahedi and N. Ay, "Quantifying morphological computation," *Entropy*, vol. 15, no. 5, pp. 1887–1915, 2013.

[9] K. Ghazi-Zahedi, D. F. Haeufle, G. F. Montufar, S. Schmitt, and N. Ay, "Evaluating morphological computation in muscle and DC-motor driven models of hopping movements," *Frontiers in Robotics and AI*, vol. 3, no. 42, 2016.

[10] N. Ay and K. Zahedi, "On the causal structure of the sensorimotor loop," in *Guided Self-Organization: Inception*, ser. Emergence, Complexity and Computation, M. Prokopenko, Ed. Springer, 2014, vol. 9, pp. 261–294.

[11] R. Pfeifer, M. Lungarella, and F. Iida, "Self-organization, embodiment, and biologically inspired robotics," *Science*, vol. 318, no. 5853, pp. 1088–1093, 2007.

[12] G. Martius, R. Der, and N. Ay, "Information driven self-organization of complex robotic behaviors," *PLoS ONE*, vol. 8, no. 5, p. e63400, 05 2013.

[13] L. van der Maaten and G. Hinton, "Visualizing high-dimensional data using t-SNE," *The Journal of Machine Learning Research*, vol. 9, no. 85, pp. 2579–2605, 2008.

[14] M. Tournier, M. Nesme, B. Gilles, and F. Faure, "Stable Constrained Dynamics," *ACM Trans. Graph.*, vol. 34, no. 4, pp. 132:1–132:10, 2015.

[15] J. Allard, S. Cotin, F. Faure, P.-J. Bensoussan, F. Poyer, C. Duriez, H. Delingette, and L. Grisoni, "SOFA - an open source framework for medical simulation," in *Medicine Meets Virtual Reality*, Long Beach, California, Etats-Unis, 2007, pp. 13–18.

[16] R. Deimel, "Soft robotic hands for compliant grasping," Ph.D. dissertation, Technische Universität Berlin, Berlin, 2017.