

Automated Co-Design of Soft Hand Morphology and Control Strategy for Grasping

Raphael Deimel

Patrick Irmisch

Vincent Wall

Oliver Brock

Abstract—To leverage soft hands to their full potential for grasping, we propose to design their morphology and control signals together. Considering both parameter domains makes it easier and faster to find solutions compared to fixing parameters of either domain. Additionally, the approach scales well to high-dimensional parameter spaces, which is a precondition to make automated co-design useful for soft hands. We further present an efficient simulator for simulating grasps with soft hands which is based on the SOFA framework and enables us to simulate more than a million grasps per day. These two complementary improvements promise a boost in the development of competent soft hands and their control in the future.

I. INTRODUCTION

Grasping performance of robotic hands generally increases if the hand is able to comply to the shape of a grasped object and to robustly exploit contact with the environment [1]. An effective way of achieving such compliance in robotic hands is through the use of inherently soft materials [2]–[4]. But this advantage comes at a price. As the grasping behavior now results both from the control signal and from the hand’s interaction with the environment, the hand design problem gets more complicated. The performance of an inherently soft hand is the result of explicit control *combined* with the deformations that result from interactions between hand, object, and environment. This implies that to design a competent soft-bodied grasping system we must consider control strategies as well as the hand’s morphology. To complicate matters further, the specification of compliant structures requires many design parameters, resulting in a very large design space.

We propose to *co-design* both hand morphology and control strategy for soft hands. For a given grasping problem, co-design enables us to find suitable matches between morphology parameters and control parameters. The goal is to have a controller that leverages the abilities of the morphology *and* to have a morphology that complements the abilities of control. To achieve this, we must search the combined parameter spaces for high quality solutions. Optimization in such a high-dimensional and nonlinear space is difficult. We present a formulation of the optimization that implicitly provides an abundance of solutions in parameter space, making it feasible to co-design many parameters.

To be able to co-design soft hands quickly, we develop a simulator capable of evaluating the task success of points in

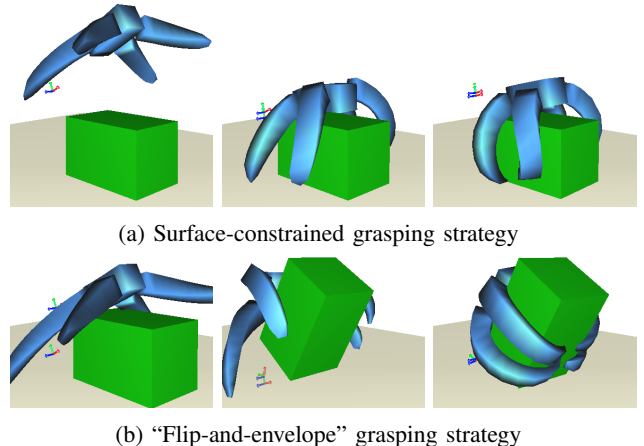


Fig. 1: Co-design creates soft hands and matching control signals that work together to generate reliable grasping behavior on a given task. Co-design can be used to adapt a known grasping strategy such as in (a), but it may also get creative and invent new ones, such as in (b).

the cross-domain parameter space, which represent particular system configurations. The simulator can simulate soft hands and their mechanical interactions efficiently, enabling us to evaluate 5.6 million grasp experiments for this paper.

In our experimental evaluation, we apply co-design to a grasping problem to quantify the improvements achievable over designing in a single domain (control, morphology) only. We will show that co-design outperforms the decoupled design problem in terms of success rate and convergence rate. We demonstrate and analyze the co-design method in a low-dimensional and a high-dimensional parameter space.

II. RELATED WORK

Co-design for grasping touches upon three different areas of related work: soft hand design, co-design as a method, and soft-body simulation.

a) Soft Hands and Grippers: Soft hands and grippers have recently received significant attention [5]. They promise to support reliable and robust grasping [1], allowing the hardware to take over tasks from control, such as establishing many redundant contacts [6] and keeping contacts stable [7].

b) The Co-Design Method: The development of methods for co-designing the body and the brain has been mentioned as one of the frontiers of current soft robotics research [8]. An example of co-design of soft robots focuses on the evolution of walking robots [9]. Co-design in itself is

All authors are with the Robotics and Biology Laboratory, Technische Universität Berlin, Germany

We gratefully acknowledge financial support by the European Commission (SOMA, H2020-ICT-645599).




Fig. 2: Illustration of the investigated gripper, its parameters (left) and controller parameters (right)

not a new idea; it has been proposed under various labels, including co-evolution and fully automated design [10]. In the context of grasping, however, automated co-design has not been applied yet. Existing research optimizes control and morphology separately. Some grippers are optimized for grasping [7], [11]–[13] but do not consider the control part of the problem. Vice versa, automated design in the control domain is a very active field of research, such as learning dynamic movement primitives [14], optimizing motion primitives with cross-entropy [15], and optimizing pre-grasp manipulation [16]. But the hand morphology is always assumed to be immutable.

c) Simulation of Soft Hands: In the context of co-design, simulation offers a versatile method to assess the quality of a combination of a specific hand and its control, but until recently it has not been computationally efficient enough for automated design. Improvements in simulation algorithms [17] enables simulating complete grasp attempts at about $0.1 \times$ real-time speed on a single CPU core [18]. These algorithmic improvements are implemented in the SOFA simulator framework [19] and used for real-time modeling of actuation effects [20]. For modeling the main component of pneumatic soft hands, Cosserat beams (i.e. chains of ball joints) have proven to be a useful abstraction [4], [21], [22]. Their main disadvantage is the need of a conversion from actuator shape to joint actuation ratios and stiffness matrices. A popular grasp simulator for non-soft (i.e. classical rigid-bodied) hands is GraspIt! [23]. While GraspIt! is orders of magnitude faster than the simulator presented in this paper, crucially it does not model soft and dynamic interactions between hand and object.

III. EXAMPLE DOMAIN

We want to develop a co-design method and validate it in the context of a well-defined example domain, as optimization should always be relative to a particular task. We consider the problem of designing a soft gripper and a control strategy to pick up an object from a table. We chose a gripper design that has not been built before (shown in Fig. 2), in order to avoid the influence of prior knowledge on the experiment. The gripper is assembled from PneuFlex actuators [2], which are a specific type of a pneumatically actuated soft continuum actuator. The gripper is assumed to be mounted on a robot arm with a passively compliant wrist. The controller commands arm position and the air pressure in the PneuFlex actuators. Trajectories are defined by linearly




Fig. 3: The Simulation consists of distinct phases: approach-and-grasp, lift, disturb, and un-grasp. The function above indicates the costs assigned to a simulation if the object is lost at the given time.

interpolated key frames. For many grasping scenarios, an initial grasping strategy can be imitated from observing humans or from published examples [1]. Because of this, we can avoid the need to evolve gripper and strategy from scratch and can focus on adapting existing concepts. In our case, we used the surface-constrained-grasp [1] strategy as the starting point. The goal in this setup is to determine both shape and trajectory parameters that lead to a robust and reliable grasp.

IV. CO-DESIGN OF CONTROL AND MORPHOLOGY

Co-design is understood as the optimization of a robotic system across at least one domain boundary, in this paper it is the boundary between control and morphology. The goal of co-design is to adapt structures on both sides of the domain boundary to each other, so that the resulting system performance improves. The representations in both domains should be chosen such that it includes as many parameters as possible that potentially help with solving the task. Optimization works better with a smaller parameter space though, so the representation should not include parameters where we already know that they have little or no effect on the task (e.g. the color of the gripper). Finally, the optimization criterion for co-design should encode that we do not need to find *the* optimal system, but just one out of many that solves the given task sufficiently well. For grasping this can be achieved by e.g. assigning minimum costs once a grasp is obtained. Inclusion of many parameters and a cost function with many potential solutions probably are the two crucial components to make co-design for soft hands feasible. Furthermore, any solution that we find needs to be robust against small perturbations, both in hand morphology (to make manufacture feasible) and in control (to make perception and planning feasible). Therefore, we strongly favor hand-controller combinations in *regions* of success, over ones in narrow local optima. Conveniently, this simplifies the optimization problem, as we can do larger exploration steps, which in turn means that we can co-design with more parameters.




Fig. 4: Model for simulating soft hands: Large trihedra indicate links, small, purple ones indicate joint locations in between. The triangular mesh is used for detecting collision and applying contact forces.

V. EFFICIENT SIMULATION OF PNEUMATIC SOFT HANDS

Co-design does not rely on simulation per se, but simulation can accelerate exploration of the design space. The task of simulation is to provide a versatile and holistic method to measure the quality of a given set of parameters in the co-design parameter space. For grasping with soft hands that effectively requires a full-fledged physics simulation of the complete grasping attempt, as large-scale deformation leads to complex, nonlinear mechanical interactions whose outcomes we are not able to describe and predict analytically. This section presents a simulator which is complementary to the co-design method, but nevertheless is a crucial component to render the co-design method computationally tractable.

Our soft-hand simulator is based on the SOFA framework [19] and its accompanying Compliant module [17]. Scenes are composed of soft hand models, objects based on meshes, and auxiliary environmental objects such as a simple table surface. The simulation model of the hand is constructed from a parametric description at runtime, which enables us to easily modify individual morphology parameters. In addition to the physics simulation, the simulator also executes a controller which determines actuator air pressures and wrist motion during the simulation. For this paper, the simulated scenes consist of a soft hand, a controller, a single object, and a table surface. For modeling soft hands, we employ a multi-model approach, where each PneuFlex actuator [2] is modeled as both a Cosserat beam and a collision hull. By separating the actuator into a beam model and a collision hull, we achieve a relatively low number of mechanical degrees of freedom while still realizing the main deformation modes of the surface. Figure 4 illustrates the concrete simulation model for the PneuFlex actuator and the gripper. The locations of individual links are indicated by the large trihedra. Adjacent links are connected by ball joints, whose locations are indicated by the smaller trihedra. Each ball joint is modeled as a 3D rotary spring that emulates the passive compliances of the adjacent actuator segments, which are estimated from the local cross section geometry of the actuator [3], [18]. The wire frame indicates the collision surface, which is used to detect contact forces and to apply interaction forces to. Each mesh vertex follows the motion of the two links closest to it. For the given simulation we use linear blend skinning (as implemented

by the SkinningMapping component of SOFA). It enables a crude but reasonable approximation of the actuator surface.

A soft hand is composed of several fingers with usually one actuator each. In order to assemble actuators into a soft hand, we connect their base link to a common wrist frame (X_A in Fig. 2) with a stiff joint, which in turn is connected to an arm frame (X_D in Fig. 2) via a passively compliant wrist joint. The gripper structure is shown in Fig. 4 and resembles a star-shaped kinematic chain. In all simulations we discretized to ten links per finger, resulting in a total of 246 mechanical degrees of freedom for the gripper. During simulation, arm position X_D and air pressures p_i are commanded by a controller. In this paper, we use a keyframe-based open-loop controller to define the actuation pattern.

The phases of a single simulated grasp attempt are illustrated in Fig. 3. First, the controller moves the gripper towards the table. Then, it inflates the fingers to grasp. After lifting back up to the initial height, the controller performs a series of increasingly stronger shaking motions by rotating the arm frame X_D in order to test the grasp quality. Finally, the controller deflates all fingers to release the object. In total each grasp attempt lasts nine seconds in simulation time. The simulation is able to run at about $0.1 \times$ real-time on a standard desktop (Intel® Core™ i5-6600K CPU @ 3.50GHz) at 10 ms time steps. This enables us to simulate more than a million grasp attempts per day on a moderately-sized compute cluster with 2000 CPU cores.

VI. EXPERIMENTS

In the introduction, we laid out two main claims for co-design. First, we claimed that it speeds up finding solutions, as with soft hands control strategy and hand morphology strongly influence each other. Second, we claimed that careful selection of representation and optimization criteria enables co-design with many parameters. To corroborate these claims, we compare co-design against strategies that only optimize in one domain in an example task explained in Sec. III. We conducted two experiments, the first one in a two-dimensional parameter space, allowing for visual interpretability of the results. The second experiment extends the parameter space to 23 dimensions to validate the findings with a more realistic number of parameters, and to evaluate the scalability of the proposed method in terms of computational costs. We do not, however, evaluate the transferability into real hardware, which is an important but separate topic.

1) *Implementation of Optimization:* Co-design is not dependent on a specific optimization method. For the experiment we chose the particle filter optimization method [24], as it is easily parallelizable, requires few assumptions, and able to represent multi-modal distributions. With this method every particle represents a concrete set of values for all parameters of both domains to optimize. A particle filter is then used to iteratively approximate the high-dimensional grasp success distribution which is defined by the simulated scenes and the particular cost function used. Optimized hands and control patterns can then be obtained by sampling from the filter’s particle set. This formalization is identical to the




Fig. 5: Illustration of the steps of the optimization method performed for each episode

particle swarm optimization method [25], except that we do not bias the particles’ motions.

Fig. 5 illustrates the steps for each optimization episode. The optimizer starts with a seed particle and a set of parameter ranges as input. The optimizer moves particles randomly in parameter space (step 1) and simulates a number of grasping trials for each particle in parallel. It then evaluates the cost function (step 2) and assigns a corresponding weight to each particle (step 3). In the last step, the particle set is resampled using stochastic universal sampling (step 4). To move particles, we use a normal distribution with zero mean and variance $\sigma = 0.05 \times \text{parameter range}$. This hyperparameter balances local vs. global exploration. The cost function defining grasp success in the experiments is illustrated in Fig. 3. After a grasp attempt, disturbances to the arm orientation are applied, which progressively increase in amplitude (see the video attachment for details). The costs decrease linearly with the duration the object center stays close (<80 mm) to the hand center, defined by the averaged position of all links of the gripper. The simulation concludes with opening the hand to drop the object, to catch errors in collision detection and constraint resolution. If the simulation fails for some reason, it is excluded from the costs evaluation. After determining the cost of several simulations, the weight of each particle is assigned to $weight = (\alpha + \frac{1}{trials} \sum_i^{trials} cost_i)^{-1}$. The hyperparameter $\alpha = 0.2$ discourages overfitting to the most recent episode.

2) *Grasping Scene and Parameterization*: The gripper and its parameters are illustrated in Fig. 2. It consists of four PneuFlex actuators acting as fingers around a fixed palm. The investigated parameters of the morphology domain are height h_i , length x_i and rubber hull thickness d_i of each finger i . Additionally, the spread between two pairs of fingers can be adjusted with the angle γ . Parameters in the control domain are the nominal air pressure p_i during grasping, and the pose of the arm frame X_D after the approach phase. The hand-controller combinations are tested on three object shapes (cuboid, cylinder, sphere), whose parameters were varied in order to avoid overfitting. Gaussian noise was added at each trial to object position ($\sigma_{pos} = 5mm$), orientation ($\sigma_{ori} = 5^\circ$), weight and size ($\sigma_w, \sigma_s = 5\%$). The parameters to optimize

Morphology domain		
Parameter	2D	23D
x_i	10 mm to 180 mm	10 mm to 180 mm
h_i	20 mm	15 mm to 25 mm
d_i	5 mm	2 mm to 7 mm
γ	45°	0° to 110°
Control domain		
Parameter	2D	23D
X_D roll, yaw	0°	-60° to 60°
X_D pitch	15°	-60° to 60°
X_D z (vertical)	-60 mm to 200 mm	-50 mm to 150 mm
X_D x	-20 mm	-60 mm to 60 mm
X_D y	0 mm	-60 mm to 60 mm
p_i	120 kPa	0 kPa to 150 kPa
		$p_i < d_i \cdot 80kPa/2mm$

TABLE I: Parameter ranges for the gripper’s morphology and controller’s motion primitive during optimization

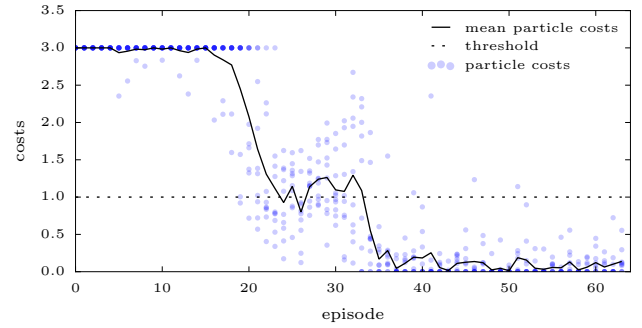



Fig. 6: Example of the cost evolution of one optimization run in 2D parameter space: Dots indicate the evaluated costs of individual particles and the line indicates the average costs for each episode. The threshold indicates the level used to compare convergence speeds in Fig. 7b and 9b.

are illustrated in Fig. 2 and listed in Table I. The parameter limits were chosen so that a gripper could actually be built, the controller can actually be executed, and the approximations of the simulation model remain physically plausible. The pressure was limited to avoid excessive curvature, which would wear out the actuators. It would also cause a strong flattening of the collision mesh, which too would make the simulation itself increasingly unrealistic.

3) *Evaluation Measures*: The progress of a typical optimization run is visualized in Fig. 6 as a scatter plot of particle costs and cost average for each episode. After each completed optimization run, we evaluated the quality of the final particles, and how fast successful particles were found. For the two-dimensional case (Fig. 7) we evaluated 42 grasps to estimate the quality of a particle, and for the 23-dimensional experiment (Fig. 9) 78. Optimization was judged successful, if at least one final particle achieved a perfect grasp success ($cost < 0.01$) on the test. We additionally quantified the convergence speed of each individual optimization run by counting the number of episodes from which on the average particle cost stayed below the threshold $c_{convergence} = 1.0$ for the remaining episodes.



(a) Ratio of succeeded optimizations to all (b) Median convergence speed with upper and lower quartiles

Fig. 7: Experimental results for optimizing two parameters in a single domain (s.d.), with different numbers of domain switches (1-32), and using concurrent (conc.) optimization strategies

4) *Co-Design vs. Single-Domain Optimization*: In the two-parameter experiment, we compare eight different optimization strategies which also reflect common modes of development for robot systems. On one end of the spectrum, we optimize in a *single domain* (s.d.) only, and keep the parameters of other domain fixed. On the other end of the spectrum, we co-design in both domains *concurrently* (conc.). In between, we optimize in one parameter domain during each episode, but switch domains between episodes for a total of (1,2,4,8,16 or 32) times during the whole optimization. Each optimization strategy uses a set of 10 particles, optimizes for 64 episodes and evaluates 6 grasp attempts per particle and episode. Initial seed particles were drawn using rejection sampling ($cost > 2.0$) in order to avoid spending computation time on trivial optimization runs where the initial particle set already contains a solution. To reduce inter-subject variance, the same seed particles were used across all optimization setups. To avoid bias, optimization runs were started equally often with optimizing either of the two domains.

Fig. 7 shows the statistics over 30 runs per optimization strategy, totaling 1,000,800 individual simulations. Fig. 7a shows the success rate, while Fig. 7b shows the median episode after which the particle set converges towards low cost. When optimizing in one domain only, probability of success is poor, but it quickly increases when interleaving both domains. Fig. 7b also shows that switching more often yields a strong boost in convergence speed. Success probability and convergence speeds consistently improve, and co-design performs best.

In addition to the statistical results, in the two-dimensional parameter space we can visualize the actual cost landscape to gain qualitative insights. Fig. 8 shows the average costs for a grid of 60×60 parameter values over 36 grasp attempts each. Please note that these results were not used for the optimization. The white and light gray areas indicate high success rates and therefore good hand-controller combinations. Two distinct areas of low costs are indicated. They correspond to two different grasping strategies that the automated co-design found, which are also shown in Fig. 1. The first strategy





Fig. 8: Cost landscape of the two-parameter experiment: The optimization finds two areas that correspond to distinct grasping strategies. An example of region (a) is shown in Fig 1a, an example of (b) in Fig. 1b.

(a) is a surface-constrained-grasp (top-grasp), which was the strategy the motion primitive was originally intended for. But the second strategy (b) operates in a completely unexpected manner. The object is contacted from the top, flipped to expose the underside, and then wrapped by the long fingers to secure an enveloping grasp. This outcome demonstrates the potential of co-design to solve tasks with novel grasping strategies.

5) *Scalability to Many Parameters*: The evaluation of the two-dimensional experiment shows that co-design is faster and more successful than iterations of single-domain approaches. For real-world application to soft hands, however, co-design must be able to perform automated design in high-dimensional parameters spaces. In a second experiment, we therefore increased the number of variable parameters to 13 in the morphological domain and 10 in the control domain. To reduce total computational costs, we only tested the most interesting optimization variants, *single domain*, 1 switch, and *concurrent*. The rest of the experiment was set up the same as the two-dimensional one, but used a set of 60 particles and optimized for 128 episodes. Each optimization strategy was run 30 times and results were averaged. In total, we simulated 4,536,000 grasps for this experiment.

The results are shown in Fig. 9. Co-design again outperformed all other tested strategies. The success rate of the 1-switch strategy increased considerably, but co-design still finds solutions faster. We additionally looked at the average amount of simulations required for convergence, which is 6120 simulations vs. 420 for the two-dimensional parameter space. The computational costs therefore increased about 15 times, while the number of parameters increased 11.5 times. This provides some indication that computational costs in co-design increase slower than the exponential increase of the design space. We believe the cause of this behavior to be



(a) Ratio of succeeded optimizations to all for each setup (b) Median convergence speed with upper and lower quartiles

Fig. 9: Statistics for optimizing 23 dimensions in total, for single domain(s.d.), one switch(1), and concurrent optimization(conc.)

the heuristics from Sec. IV for formulating the cost function and problem representation.

VII. LIMITATIONS

The results obtained with co-design are promising, but the example domain addressed here was relatively simple. Future applications should demonstrate co-design based on more complex and diverse grasping scenarios. The optimization algorithm used was chosen for its simplicity. It could be replaced by algorithms that leverage the information obtained from the expensive samples more efficiently.

This paper does not assess the size of the simulation gap between simulated behavior and that of real soft hands. Substantial efforts have been made though, to validate the actuator models against real actuators [18] to improve plausibility of the simulation results. The compliance of soft hands additionally limits the effect of small errors in morphology on contact forces, further accommodating the transferability of solutions to real systems. Nevertheless, we expect a subsequent optimization with real hardware to be prudent and necessary to account for the simulation gap.

VIII. CONCLUSION

We investigated the feasibility of co-designing morphology of soft hands and their control strategies for grasping. We proposed guidelines for selecting parameter spaces and optimization criteria that support co-design in high-dimensional space. Complementing the co-design method, we presented an efficient simulator for soft hands which enabled us to simulate a total of 5.5 million grasp attempts for this paper.

We showed in a quantitative evaluation that co-design consistently outperforms optimization limited to only one domain, even when we switch between domains iteratively during optimization. This behavior was observed with both two and 23 varied parameters. An analysis of the computational costs further suggests that the effort to co-design scales favorably, much slower than the expected exponential increase in the number of free design parameters. These results open up the prospect of efficiently customizing soft grippers and strategies for specific tasks, in order to provide a cost effective and high-performance alternative to engineered fixtures and general-purpose robot hands.

REFERENCES

- [1] C. Eppner, R. Deimel, J. Álvarez-Ruiz, M. Maertens, and O. Brock, "Exploitation of environmental constraints in human and robotic grasping," *IJRR*, vol. 34, no. 7, pp. 1021–1038, 2015.
- [2] R. Deimel and O. Brock, "A Compliant Hand Based on a Novel Pneumatic Actuator," in *ICRA*, 2013, pp. 2047–2053.
- [3] —, "A novel type of compliant and underactuated robotic hand for dexterous grasping," *IJRR*, vol. 35, no. 1-3, pp. 161–185, 2016.
- [4] A. Marchese, K. Komorowski, C. Onal, and D. Rus, "Design and control of a soft and continuously deformable 2D robotic manipulation system," in *ICRA*, 2014, pp. 2189–2196.
- [5] 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, 2016.
- [6] J. Amend, E. Brown, N. Rodenberg, H. Jaeger, and H. Lipson, "A Positive Pressure Universal Gripper Based on the Jamming of Granular Material," *IEEE Trans. Robot.*, vol. 28, no. 2, pp. 341–350, 2012.
- [7] L. U. Odhner, L. P. Jentoft, M. R. Claffee, N. Corson, Y. Tenzer, R. R. Ma, M. Buehler, R. Kohout, R. D. Howe, and A. M. Dollar, "A compliant, underactuated hand for robust manipulation," *IJRR*, vol. 33, no. 5, pp. 736–752, 2014.
- [8] D. Rus and M. T. Tolley, "Design, fabrication and control of soft robots," *Nature*, vol. 521, no. 7553, pp. 467–475, 2015.
- [9] J. Hiller and H. Lipson, "Automatic design and manufacture of soft robots," *IEEE Trans. Robot.*, vol. 28, no. 2, pp. 457–466, 2012.
- [10] J. Pollack, H. Lipson, P. Funes, S. Ficici, and G. Hornby, "Coevolutionary robotics," in *Proceedings of the First NASA/DoD Workshop on Evolvable Hardware*, 1999, pp. 208–216.
- [11] M. Ciocarlie and P. Allen, "Data-driven optimization for underactuated robotic hands," in *ICRA*, May 2010, pp. 1292–1299.
- [12] A. M. Dollar and R. D. Howe, "The Highly Adaptive SDM Hand: Design and Performance Evaluation," *IJRR*, vol. 29, no. 5, pp. 585–597, 2010.
- [13] —, "Joint coupling design of underactuated hands for unstructured environments," *IJRR*, vol. 30, no. 9, pp. 1157–1169, 2011.
- [14] S. Niekum, S. Osentoski, G. Konidaris, and A. Barto, "Learning and generalization of complex tasks from unstructured demonstrations," in *IROS*, 2012, pp. 5239–5246.
- [15] A. Gupta, C. Eppner, S. Levine, and P. Abbeel, "Learning dexterous manipulation for a soft robotic hand from human demonstrations," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct 2016, pp. 3786–3793.
- [16] J. King, M. Klingensmith, C. Dellin, M. Dogar, P. Velagapudi, N. Pollard, and S. Srinivasa, "Pregrasp Manipulation as Trajectory Optimization," in *Robotics Science and Systems (RSS)*, 2013.
- [17] 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.
- [18] R. Deimel, "Soft robotic hands for compliant grasping," Ph.D. dissertation, Technische Universität Berlin, Berlin, 2017.
- [19] 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, MMVR 15, February, 2007*, Long Beach, California, Etats-Unis, 2007, pp. 1–6.
- [20] C. Duriez, E. Coevoet, F. Largilliere, T. Morales-Bieze, Z. Zhang, M. Sanz-Lopez, B. Carrez, D. Marchal, O. Goury, and J. Dequidt, "Framework for online simulation of soft robots with optimization-based inverse model," in *2016 IEEE International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAP)*, 2016, pp. 111–118.
- [21] F. Renda, V. Cacucciolo, J. Dias, and L. Seneviratne, "Discrete Cosserat approach for soft robot dynamics: A new piece-wise constant strain model with torsion and shears," in *IROS*, 2016, pp. 5495–5502.
- [22] F. Renda, M. Cianchetti, M. Giorelli, A. Arienti, and C. Laschi, "A 3D steady-state model of a tendon-driven continuum soft manipulator inspired by the octopus arm," *Bioinspiration & Biomimetics*, vol. 7, no. 2, p. 025006, 2012.
- [23] A. Miller and P. Allen, "Graspit! a versatile simulator for robotic grasping," *IEEE Robot. Autom. Mag.*, vol. 11, no. 4, pp. 110–122, 2004.
- [24] B. Liu, S. Cheng, and Y. Shi, *Particle Filter Optimization: A Brief Introduction*. Springer International Publishing, 2016, pp. 95–104.
- [25] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *IEEE International Conference on Neural Networks*, vol. 4, 1995, pp. 1942–1948 vol.4.