The SLS-generated Soft Robotic Hand –

An Integrated Approach using Additive Manufacturing and Reinforcement Learning

Arne Rost
Corporate Research
Festo AG & Co. KG
73734 Esslingen, Germany
rost@de.festo.com

Stephan Schädle
Institute for Artificial Intelligence
University of Applied Sciences Ravensburg-Weingarten
88250 Weingarten, Germany
schaedst@hs-weingarten.de / s@schaedle.info

Abstract – To develop a robotic system for a complex task is a time-consuming process. Merging methods available today, a new approach for a faster realization of a multi-finger soft robotic hand is presented here. This paper introduces a robotic hand with four fingers and 12 Degrees of Freedom (DoFs) using bellow actuators. The hand is generated via Selective Laser Sintering (SLS), an Additive Manufacturing method. The complex task execution of a specific action, i.e. the lifting, rotating and precise positioning of a handling-object with this robotic hand, is used to structure the whole development process. To validate reliable functionality of the hand from the beginning, each development stage is SLS-generated and the targeted task execution is trained via Reinforcement Learning, a machine learning approach. Optimization points are subsequently derived and fed back into the hardware development. With this Concurrent Engineering strategy a fast development of this robotic hand is possible. The paper outlines the relevant key strategies and gives insight into the design process. At the end, the final hand with its capabilities is presented and discussed.

Soft Robotic Hand; Reinforcement Learning; Additive Manufacturing; SLS; Concurrent Engineering; Compliant; LearningGripper;

I INTRODUCTION

Soft robots are one kind of intrinsically safe service robots, which are able to operate in close human proximity. Flexible structures and elastic actuation concepts with a variable stiffness are a main requirement for such a system. One way to generate the required structures is to additively generate the robot parts and to integrate predefined areas for different functional tasks, e.g. for flexible force transmission, soft actuation and surface contact sensing. The Bionic Handling Assistant (BHA) is an example for such a robotic system [1]. Selective Laser Sintering is one method to additively generate these required parts quickly derived from existing 3D-CAD-data [2].

Solving a manipulation task involves, even for humans, non-trivial sensory motor skills and high levels of adaptation. Diverse approaches for multi-finger robot hand designs like [3] and [4] are using electric DC motors for joint actuation and pre-programmed control algorithms for task execution. Some approaches are using self-reconfiguring strategies [5, 6]. Reinforcement Learning is a meth- od of the machine learning domain to teach a robot system what to do [7]. The learning process is not supervised, i.e. the robot has to learn suitable actions in a certain hand configuration without any human guidance. In our case these actions are movements of the twelve laser-sintered bellow actuators of the hand, shown in Fig. 1 and Fig. 4, which are supposed to move four fingers with 3 DoFs each. Every performed action is rewarded and the learning system using trial and error will reinforce actions leading to a predefined goal, e.g. lifting an object. Every action that drops this object is penalized and will be avoided the next time the same gripper configuration is reached. After several hours of learning the desired system is adaptive on any desired goal orientation or target object shape.

The key content of this paper is to combine above mentioned Additive Manufacturing with Reinforcement Learning to speed up the development of a multi-fingered soft robotic hand with bellow actuators.

II CHALLENGES AND REQUIREMENTS

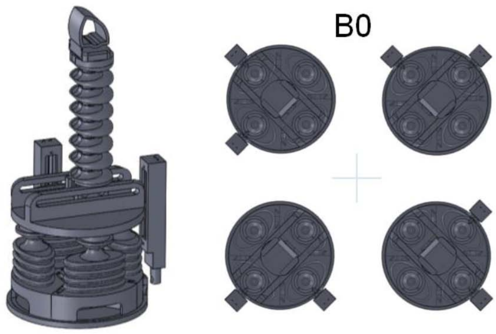

Fig. 1 Left: starting with finger structure F1,
Right: starting with a generic base layout B0 = B0-F1

There are several challenges for the development of a soft robotic hand. The key challenge is to identify the fitting functional structure for a desired task execution. In our case, the desired task is to be able to lift an object with the fingers of a robotic hand and to be able to rotate and orientate this object to reach a defined end position. The movement here should have a human-like pattern. This requirement excludes certain hand layouts and predefines a hand layout to start with. For our scenario we need a minimum of four fingers to be able to hold the object with two fingers and at the same time rotate the object with the other two fingers. An additional challenge is to identify the size and position of these fingers in relation to each other. The desired handling object has a major influence on this requirement. Since we are interested to show a scenario similar to a human rotating a fruit, e.g. an apple, the starting object is a sphere and has a diameter of 100 mm and a mass of no more than 250 grams. The next challenge is to identify the best kinematic structure for each finger. As a starting point we use a finger design close to the BHA, as shown in Fig. 1.

Within Finger F1, force transmission from fingertip to finger base and actuation of the finger are combined in one functional structure. One bendable bellow actuator is designed to represent the upper finger and four bellow actuators are moving the base of this actuator. As a next functional requirement four fingers have to be aligned with 90 degree offset relative to each other facing to a central point to be able to use two fingers to hold the object while the other two are able to rotate the object. A successful learning strategy can only be implemented, if additional challenges are addressed: The repeatability and absolute accuracy of the positioning of all degrees of freedom needs to be within certain tolerable limits, so that the learned strategies are executable. Furthermore the integrated sensors need to be able to measure the kinematic movement as precise as possible without overloading the learning system with irrelevant data. With no known soft robotic hand with SLS-generated bellow actuators using Machine Learning being realized before, we decided to follow a Top-Down approach. We predefined the final system to have a force closure grasp of the handling object and to have the ability to hold the object in the air with only two fingers. To realize a robotic hand with four fingers, a force closure grasp capability of up to 250 grams and the ability to know the current state of its gripping object, sequential robot development procedures would have been too time-consuming for our needs. A parallel concurrent design approach was therefore chosen to be able to realize this system in a given time-period of ten months.

III METHODS USED FOR AN INTEGRATED APPROACH

To address these challenges the methods Concurrent Engineering (CE), Additive Manufacturing (AM) and Reinforcement Learning (RL) were combined into an integrated approach, as shown in Fig. 2. Concurrent Engineering (CE) is a development method to speed up tasks in a shared simultaneous process [8]. Additive manufacturing (AM) parts may be manufactured using different ways, e.g. Fused Deposition Modelling (FDM), Selective Laser Sintering (SLS), Laminated Object Manufacturing (LOM) or Digital Light Processing (DLP), as listed and described in [2]. Fig. 3 shows a SLS-Machine used to generate the parts used for this approach.

The relevant Machine Learning (ML) method within this approach is Reinforcement Learning (RL), which uses a model free temporal difference tabular learning algorithm (Q, SARSA) to learn basic manipulation strategies [9]. We used a hierarchical reinforcement approach learning basic manipulation skills on the first level [10]. On the second level an orientation of the gripping object is learned [7]. A gripper simulation can help tuning the learning algorithms but also helps to identify required hardware changes. Moreover ML is ideal for system evaluation, because the strength of the learning is the initially random actuation in every reachable state of the system, so e.g. a constructive weak spot will be found quickly during usage.

Fig. 3 SLS-Machine for Additive Manufacturing (AM)

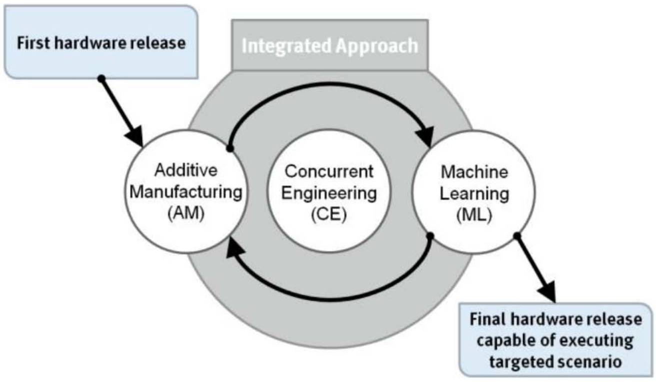

Fig. 2 Integrated Approach using AM, CE and ML

IV CONCURRENT ROBOTIC ENGINEERING APPROACH

Starting within the Additive Manufacturing domain a modular, soft robotic hand design was chosen. Five major hardware versions of the hand base structure were generated using this process. The generic Version B0, 1st-follow up Version B1, 2nd-follow up Version B2, 3rd-follow up Version B3 and 4th-follow up Version B4. The first iteration was done without machine learning phase: a simple connective base was designed to fix the four loose fingers onto a modular platform B1, see also Fig. 1. The generated hand hardware can be seen from above and in a 3D-view in Fig. 6. In parallel seven major hardware releases of the finger structure F1 to F7 were developed, as shown in Fig. 9. Base and fingers combined, nine major soft robotic hand hardware releases were generated, as listed in Tab. 1.

Number	Base structure	Finger structure
#1	B0	F1
#2	B1	F2
#3	B1	F3
#4	B2	F4
#5	B2	F5
#6	B3	F4
#7	B3	F5
#8	B4	F6
#9	B4	F7

Table I. Nine major hardware releases of robotic hand

A. Hardware Phase with Base B1

Four SLS-generated fingers F1, one in each corner, are used as a first setup for the machine learning process. Each finger is easily exchangeable via a connection interface. The work space of this 1st learning version was designed to handle a sphere of 100 mm diameter. B1 has no socket in the middle to rest the handling object. It uses quick actions valves to pressurize SLS-generated bellow actuators, as shown in Fig. 4. The handling object is a ball made of Styrofoam. Each finger has a force sensor placed underneath a solid state hinge to detect a fingertip contact. Four bellow actuators are working within the finger base to move the 1st and 2nd DoFs of the finger base together. The movement of the finger base is detected via two linear potentiometers on the outside. The upper finger bending is detected via a flexible force sensor inside each finger, which changes its resistance if bent.

B. Machine Learning Phase 1

Main objective: To get started. The RL-Algorithms are learning using a reward system. A reward is the reaction of a performed action. Defining this reward, the learning agent tries to fulfill the given goal with a smaller penalty. In the presented case the agent should learn to rotate a lifted sphere. If the sphere has fallen out of the gripper, the highest penalty is applied. A rotation in which the sphere is touched by at least two fingers yields a maximum reward. Fig. 5 shows the different numbers of manipulation steps until the sphere drops using different hardware configurations. It is obvious, that the socket is able to improve the number of manipulation actions. Due to its weight, the sphere is tending to skid down. Therefore the manipulation could be improved inserting a socket. The sphere then would be relifted instead of falling out the gripper. The whole learning process could be improved this way.

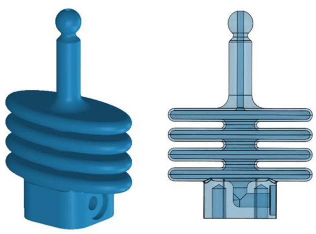

Fig. 4 SLS-generated bellow actuators for Soft robotics, here: version for 1st and 2nd DoF of finger F7

C. Hardware Phase with Base B2

A socket in the middle of the hand was derived within MLP 1 as a necessity to efficiently learn better strategies. The sensors of the 1st and 2nd DoF of each finger had to be decoupled to work properly to keep measuring accuracy within demanded limits. The form, function and elasticity of the 1st fingertip design were insufficient and had to be redesigned. A lifting of the sphere was still not possible. To reduce the number of valves and additional components and to reduce system weight, two of the four bellow actuators within the finger base were omitted and replaced with springs as antagonists. In addition, a ball joint in the middle of each finger was replaced by a new kinematic finger base structure, which decouples the two base joints. New proportional valves replaced the old quick action valves to be able to reach a higher overall accuracy. A new handling object, a 100 mm solid sphere, was designed and SLS-generated. The linear potentiometers used within finger base were replaced with two rotatory potentiometers, which measured the 1st and 2nd DoF directly.

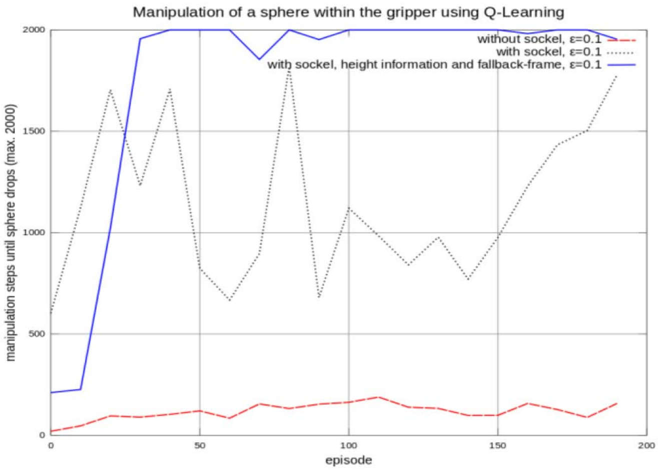

Fig. 5 Hardware releases in comparison: Base B1, B2 and B4

D. Machine Learning Phase 2

Main objective: For further development of the base structure and improvements of the finger structure. The sphere now may be relifted instead of falling out the gripper, but there is no closed-loop feedback about the current state of the height of the ball. Having additional state information if the object is lifted could result in a much better relift of the sphere, see also Fig. 5. The new valves and sensors are enabling a far better position control in the 1st and 2nd DoF. The increased friction at the finger-tip leads to the ability to lift heavier spheres, i.e. including an IMU-sensor. The form of the finger-tip is evaluated in a simulation to apply e.g. rolling along the object, as shown in Fig. 7. It turns out, that a flexible half-sphere as a finger-tip has significantly better lifting abilities.

E. Hardware Phase with Base B3

A height sensor within the socket in the middle was derived within MLP 2 as helpful to speed up the learning process. The sensor to detect the bending of the upper finger was placed on the outer side of the finger in a neutral bending zone. A mechanism to change the socket height for up to 15 mm was integrated. In parallel the fingers were further adapted. For B2 and B3 finger design F5 was developed. The surface is coated with a rubber like substance to increase grip between the handling object and the fingertip. The handling object itself was cut into two parts to be able to integrated holes for releasing left-over powder from the SLS-generative process. This enabled the design of a lighter sphere. A finger base connector to the outside world, e.g. to connect with a robot arm, was integrated, as shown in Fig. 8.

F. Machine Learning Phase 3

Main objective: For optimization of the whole hand structure. The height information can be used to differentiate between the lifted and the socket position of a handling object. During the learning process the sphere is not centered above the socket. In many cases the sphere is falling out of the gripper due to a roll motion emerging because of too strong forces during the manipulation operation. An additional plane as a fallback frame would allow a re-lift, even if the sphere has fallen out of the gripper. The integrated force sensor within the finger-tip is used to report the actual force. Too strong contact forces are getting a negative reward within the learning phase.

G. Hardware Phase with Base B4

With this hardware version a fallback-frame for the learning process was integrated, as shown in Fig. 10. The frame is able to catch the

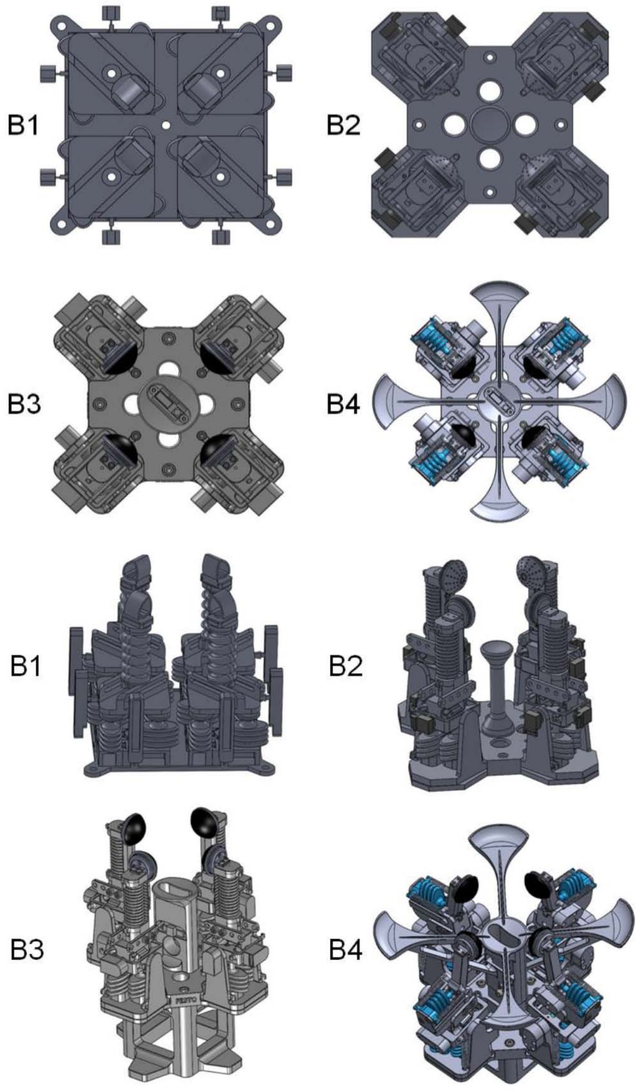

Fig. 6 Base hardware versions; Top: from above; Bottom: 3D-view

handling object if the learning process is in its very beginning stages and it is able to speed up the learning process. The structure of the upper finger was redesigned in F6 to be able to do more movement towards the center of the hand. In parallel the flexible sensor of the 3rd DoF was substituted with a rotatory potentiometer. This version introduced a fully integrated internal wiring. This wiring reduced possible unwanted external friction effects. With version F7 the final magnetic joint sensors for all three DoFs of each finger were integrated. The handling object was finally adapted and to visualize the target of a possible rotation action a brand name was written onto the surface.

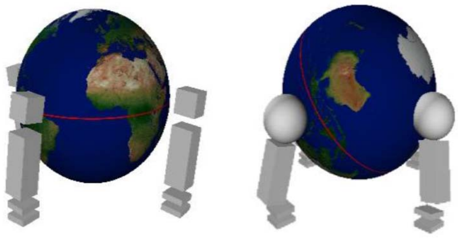

Fig. 7 Simulation of different finger-tips with ODE-Physics and RL-Glue experiments [11]

H. Machine Learning Phase 4

Main objective: Learning to rotate and lift an object into a user-given end position. The finger evolution took place in parallel to the hand base structure evolution, as shown in Fig. 8. In this phase, the final finger version F7 is used together with the final base B4 to learn the given task. Fig. 10 shows this hardware-version B4-F7 during assembly. Solving the required task, the learning is expanded to a hierarchical RL approach using the previously learned manipulation skills achieving a targeted sphere configuration, as it is explained in [7].

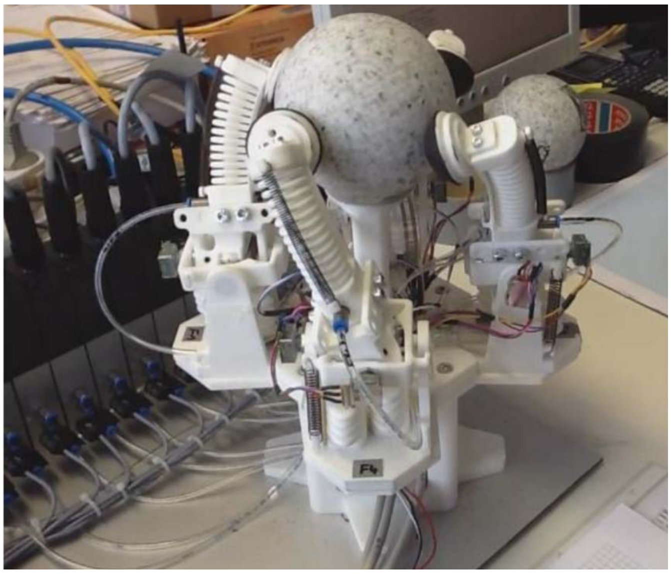

Fig. 8 Hardware release #7: B3-F5 gripping and lifting

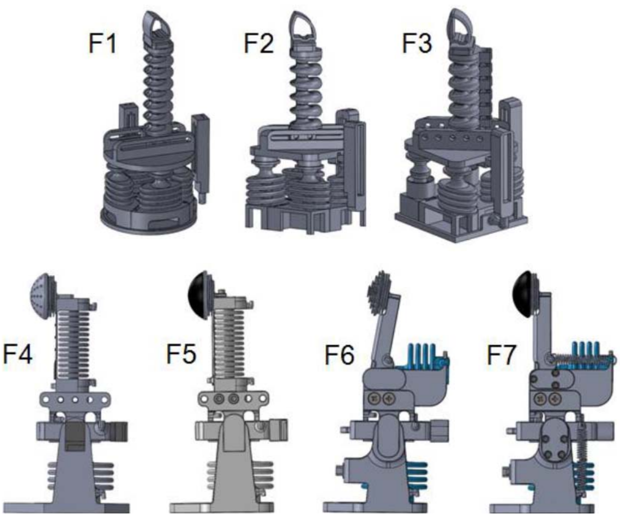

Fig. 9 Seven finger-releases F1 to F7

Fig. 10 SLS-generated fallback-frame and a partly-assembled final hand B4-F7

V FINAL MULTI-FINGER SOFT ROBOTIC HAND

The resulting final hardware version for the “LearningGripper” is a combination of the base structure B4 and finger structure F7. A detailed front view and a side view of finger structure F7 is shown in Fig. 11, while Tab. 2 names all relevant components. Finger F7 has an internal cable wiring, magnetic joint sensors to determine the joint angles, three compliant bellow actuators, to actuate each joint and three springs to act as antagonists for these bellow actuators. The internal wiring ends at a RS232-connector interface, which allows a fast exchange of each finger.

Number	Description
#1	Fingertip
#2	Finger
#3	Force sensor within fingertip
#4	Finger bellow actuator 3rd DoF
#5	Finger spring 3rd DoF
#6	Magnetic sensor 3rd DoF
#7	Internal cable guiding
#8	Finger spring 2nd DoF
#9	Magnetic sensor 1st DoF
#10	Magnetic sensor 2nd DoF
#11	Finger bellow actuator 1st DoF
#12	Finger bellow actuator 2nd DoF
#13	Finger spring 1st DoF
#14	Cable connector
#15	Hand base connector

Table II. Description of finger components
The fingertip uses a flexible force sensor to measure a surface contact, as shown in detail in Fig. 12.

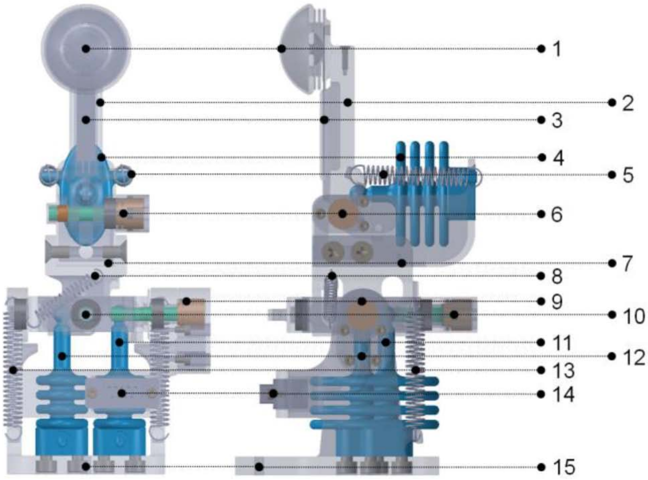

Fig. 11 Front view and side view of finger F7; according numbers are to be found in Tab. II.

The contact zone shown on the right side of Fig. 12 is a combination of an elastic silicone material and a rigid inner structure with barbs to support the silicone, inspired by a human finger. An integrated bellow-like structure works as a solid state hinge and allows a hemisphere on the backside of the contact-zone to relay the contact forces of the fingertip to a force sensor.

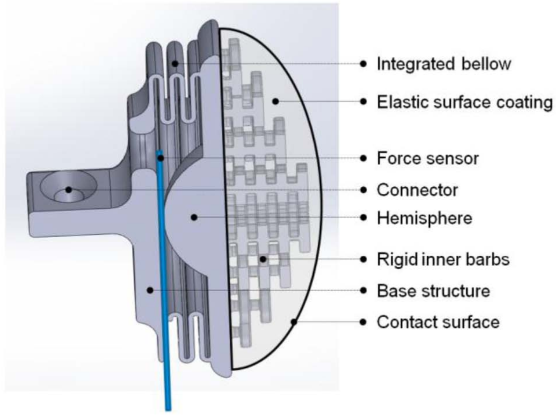

Fig. 12 Side view of the SLS-generated fingertip of finger F7

Base B4 has four modular finger docks, each to be opened with a single screw. It has a socket in the middle to catch a handling object and within this socket a height sensor to determine the object lifting height. In addition is has a fallback frame for the very early learning stages. Except for the screws, the bearings, the sensors, and the springs, all parts of the hand are SLS-generated. The bellow actuators are pressurized with up to 3.5 bar. Fig. 13 shows the finished hand.

Fig. 13 Final hardware B4-F7 holding a handling object with two fingers

Fig. 14 shows this camera-ready final setup with two fully operational hands, which was shown on “Hannover Messe 2013” industry trade fair. Using air pressure to drive the bellow actuators of each hand, the force closure grasp of the final multi-finger soft robotic hand is compliant and allows a very close human-robot-interaction without additional safety features. A person may easily take out and re-insert the handling object to try out the system performance.

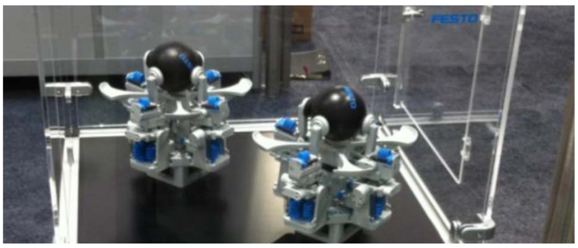

Fig. 14 Exhibition system with two “LearningGrippers”

VI CONCLUSION AND FUTURE WORK

The modular multi-finger soft robotic hand is in operation. A new integrated approach for a parallel development of the task execution and the hardware structure was shown. The generated robotic hand is capable of lifting and rotating a given object, e.g. a sphere or a cube, into a desired final position. It learns the steps necessary to execute a desired task via Reinforcement Learning. Still, there are some points left for future work. One is the still necessary adaptation of some of the parameters of a learned strategy if a modular finger is exchanged. Another direction for future activities is the optimization of the transfer-learning between two systems. One system learns a strategy, but the task execution is on a slightly different system. These differences occurring due to different inner-system-frictions in the actuators shall be learned in a more efficient way to reproduce a stable and robust manipulation. Orienting fruits for an industrial pick-and-place robot is a possible future usage scenario.

Video:
http://www.youtube.com/watch?v=JLuPbAMxoSU

REFERENCES

[1] Festo AG & Co. KG / IPA, "Bionic Handling Assistant", http://www.festo.com/rep/en_corp/assets/pdf/BHA_en.pdf, Esslingen, Germany, 2011

[2] Breuninger, J., Becker, R., Wolf, A., Rommel, S. and Verl, A., "Generative Fertigung mit Kunststoffen – Konzeption und Konstruktion für Selektives Lasersintern", Berlin, Heidelberg, Springer, 2013

[3] Chen, C.-H., "Virtual Model for a Multi-Finger Robot Hand Design", 4th International Conference on Computational Intelligence, Communication Systems and Networks (CICSyN) 2012, Phuket, Thailand, 2012, p.176-180

[4] Davis, S., Tsagarakis, N.G and Caldwell, D.G., „The Initial Design and Manufacturing Process of a Low Cost Hand for the Robot iCub", 2008 8th IEEE-RAS International Conference on Humanoid Robots, Daejeon, Korea, 1-3 December 2008, p.40-45

[5] Neubert, J., Cantwell, A.P., Constantin, S., Kalontarov, M., Erickson, D. and Lipson, H., "A robotic module for stochastic fluidic assembly of 3D self-reconfiguring structures", 2010 IEEE International Conference on Robotics and Automation (ICRA), Anchorage, USA, 3-7 May 2010, p.2479-2484

[6] Sciliano, B., Khatib, O., "Handbook of Robotics", Berlin, Heidelberg, Springer, 2008

[7] Schädle, S. and Ertel, W., "Dexterous manipulation using hierarchical reinforcement learning", In: Autonomous Learning workshop, IEEE International Conference on Robotics and Automation 2013, Karlsruhe, Germany, 2013

[8] Bullinger, H.-J. (Publ.); Warschat, J.: "Forschungs- und Entwicklungsmanagement: simultaneous engineering, Projektmanagement, Produktplanung, rapid product development" Stuttgart: Teubner, 1997

[9] Sutton, R.S. and Barto, A.G, "Reinforcement learning: An introduction", Cambridge, MIT Press, 1998

[10] Dietterich, T.G. "Hierarchical reinforcement learning with the MAXQ value function decomposition", 2000, Journal of Artificial Intelligence Research, no. 13, pp. 227-303

[11] Tanner, B. and White, A., "RL-Glue: Language-independent software for reinforcement-learning experiments", 2009, The Journal of Machine Learning Research, vol. 10, pp.2133-2136