NICOL: A Neuro-inspired Collaborative Semi-humanoid Robot that Bridges Social Interaction and Reliable Manipulation

MATTHIAS KERZEL¹, PHILIPP ALLGEUER¹, ERIK STRAHL¹, NICOLAS FRICK¹, JAN-GERRIT HABEKOST¹, MANFRED EPPE², and STEFAN WERMTER.¹

¹Knowledge Technology, University of Hamburg, Hamburg, Germany (e-mail: {matthias.kerzel,philipp.allgeuer,erik.strahl,stefan.wermter}@uni-hamburg.de) ²Data Science Foundations Institute, Hamburg University of Technology, Hamburg, Germany (e-mail: manfred.eppe@tuhh.de)

Corresponding author: Matthias Kerzel (e-mail: matthias.kerzel@uni-hamburg.de).

The authors gratefully acknowledge partial support from the German Research Foundation DFG under projects CML (TRR 169), IDEAS/MoReSpace (402776968), LeCAREbot (433323019), and the EU projects SECURE and SOCRATES. The authors thank Phuong Nguyen for his contributions to the NICOL robot and Sergio Lanza for his feedback on the manuscript.

ABSTRACT Robotic platforms that can efficiently collaborate with humans in physical tasks constitute a major goal in robotics. However, many existing robotic platforms are either designed for social interaction or industrial object manipulation tasks. The design of collaborative robots seldom emphasizes both their social interaction and physical collaboration abilities. To bridge this gap, we present the novel semi-humanoid NICOL, the *Neuro-Inspired COLlaborator*. NICOL is a large, newly designed, scaled-up version of its well-evaluated predecessor, the Neuro-Inspired COmpanion (NICO). NICOL adopts NICO's head and facial expression display, and extends its manipulation abilities in terms of precision, object size and workspace size. To introduce and evaluate NICOL, we first develop and extend different neural and hybrid neuro-genetic visuomotor approaches initially developed for the NICO to the larger NICOL and its more complex kinematics. Furthermore, we present a novel neuro-genetic approach that improves the grasp-accuracy of the NICOL to over 99%, outperforming the state-of-the-art IK solvers KDL, TRACK-IK and BIO-IK. Furthermore, we introduce the social interaction capabilities. Overall, this article presents for the first time the humanoid robot NICOL and, thereby, with the neuro-genetic approaches, contributes to the integration of social robotics and neural visuomotor learning for humanoid robots.

INDEX TERMS Neuro-genetic visuomotor learning, neuro-robotics, humanoid robotics

I. INTRODUCTION

What are the main prerequisites for a collaborative robot to act and work alongside humans? On the one hand, the robot needs to be able to perform precise and reliable object manipulation in a large workspace, but on the other hand, the robot also needs to be able to interact with human coworkers on a social level. This requires auditory communication as well as the understanding and synthesis of both social cues and non-verbal communication like gestures and facial expressions.

Social interaction is one of the crucial factors for intuitive human-robot collaboration and the ability of robots to learn from humans [1]–[3]. Gaze, as a non-verbal social cue, facilitates shared attention, which can signal future actions or draw attention towards an object or location [4]. Likewise, facial expressions can be used to give feedback or warn about dangers. Speech output also offers an intuitive way to request aid and instructions, or to suggest a course of action. There exist only very few robots, however, that integrate social interaction with reliable manipulation of adult-scale objects while remaining affordable. A robot is desirable that bridges this gap (see Fig. 1).

The Pepper robot by Softbank Robotics, for instance, is an affordable semi-humanoid social robot with rudimentary physical object manipulation abilities. It features a variety of body language skills and communication abilities, and can be programmed to understand certain social cues. However, even though it is equipped with arms and



FIGURE 1. The NICOL robot is an affordable platform that bridges the gap between social robotics and reliable object manipulation in a large workspace.

hands, it is not designed to grasp or lift but to gesture. A good example of an affordable robot designed for physical object manipulation is the Franka Research 3 by Franka Emika [5]. It can perform high-precision grasping in a large workspace, but it does not feature any built-in social interaction abilities.

The iCub robot is a good example of a humanoid robot that features both reliable object manipulation (for smaller objects) and social interaction skills [6] and is, therefore, an excellent research platform for developmental robotics and Human-Robot Interaction (HRI). The iCub is arguably among the most successfully engineered robots that combine these skills, but it is also very expensive and hard to afford for small research groups and companies with only little funding (e.g., start-ups). Hence, to bridge the gap between affordable social robots like Pepper, and affordable robots designed for object manipulation like the Franka Research 3, we present the novel adult-sized semihumanoid robot NICOL, the *Neuro-Inspired COLlaborator* (see Fig. 2).

The social interaction capabilities of NICOL are based on the well-evaluated design of its predecessor, the humanoid child-sized Neuro-Inspired COmpanion (NICO) [7]. Like the NICO, the new NICOL has an articulated head that can express stylized facial expressions and focus its gaze on interaction partners or relevant elements of the shared workspace. While NICO has limited lowprecision object manipulation abilities, NICOL can manipulate objects using two Robotis OpenManipulator-P arms with adult-sized five-fingered Seed Robotics RH8D hands attached. It has a workspace similar to an adult person and can handle everyday objects.

In addition to the hardware, NICOL also features a software framework that is directed towards two groups of users—the roboticists that wish to work at a low level with a high degree of control, and the users from different backgrounds, like artificial intelligence, machine learning, cognitive modelling and human-robot interaction, that require a simpler and more abstract interface. Low-level control is delivered via the Robot Operating System (ROS) and provides a rich control interface with great customizability and extensibility. Manipulating and developing within a ROS workspace often requires many non-trivial software cross-dependencies and requires a certain level of expertise to navigate, use and maintain. For users with little background in robotics, a more lightweight system is desired. Furthermore, such users often demand a lightweight simulated version of the robot for performing rapid iterations of parallelised deep reinforcement learning experiments. NICOL meets the demands of both groups-roboticists and AI researchers. Its software offers a full ROS interface, but users can also connect to a front-end Python interface running on top of ROS, which enables them to directly use the robot without having to explicitly deal with ROS. In addition, we provide a simulated virtual version of NICOL.

In this article, we first place the NICOL robot in context by examining existing related platforms (Section II). This demonstrates how our robot fills the gap between social robotics and object manipulation in an affordable manner. The next three sections then cover our three main contributions:

- Section III presents the hardware and software concept for NICOL, the *Neuro-Inspired COLlaborator*. Herein, we also introduce the social interaction capabilities of NICOL, including face and emotion generation capabilities as well as auditory and visual capabilities.
- Section IV presents a thorough evaluation of NICOL's kinematics. We adapt and extend different neural and hybrid neuro-genetic visuomotor approaches initially developed for the NICO to the larger NICOL and its more complex kinematics. We show the challenges of neural end-to-end learning for complex kinematics and successfully evaluate a hybrid neuro-genetic approach.
- Section V introduces a novel neuro-genetic visuomotor learning approach that improves the learned grasping accuracy of NICOL to over 99%, and contributes to the growing research on neuro-robotic approaches that are suitable for transfer between platforms with different morphologies [8]. Here, we also present experiments to show that the grasp-accuracy of the NICOL improves to over 99%, outperforming the state-of-the-art inverse kinematics (IK) solvers KDL, TRACK-IK and BIO-IK.

Overall, this article presents for the first time the NICOL robot, which serves as an affordable research platform that integrates social robotics and humanoid manipulation, and as a testbed for neural visuomotor learning.

II. RELATED WORK

A. RELATED ROBOTIC PLATFORMS

NICOL combines the capabilities of dexterous robotic manipulation and social interaction, and so has different types

Robot	Туре	Humanoid Body	Humanoid Head	Facial Expressions	Gestures	Grasping	Workspace	Price Range
UR-5	robot arm	×	×	×	×	1	adult-sized	mid
KUKA-DLR Arm	robot arm	×	×	×	×	1	adult-sized	mid
Franka Research 3	robot arm	×	×	×	×	1	adult-sized	low
Furhat	robot head	×	✓	✓	×	×	none	mid
iCub head	robot head	×	✓	✓	×	×	none	mid
ROBOTIS OP3	infant-sized robot	 ✓ 	×	×	 Image: A set of the set of the	×	tiny	low
Nao	infant-sized robot	✓	✓	×	1	1	tiny	low
iCub	child-sized robot	✓	 Image: A set of the set of the	✓	 Image: A set of the set of the	1	child-sized	high
NimbRo-OP	child-sized robot	✓	✓	×	×	×	child-sized	mid
Pepper	child-sized robot	(✓) upper body	✓	×	1	×	child-sized	mid
Рорру	child-sized robot	✓	✓	✓	1	1	child-sized	low
NICO (ours)	child-sized robot	✓	✓	✓	 ✓ 	1	child-sized	low
Talos	adult-sized robot	✓	 ✓ 	×	 Image: A set of the set of the	1	adult-sized	high
PR2	adult-sized robot	×	×	×	×	1	adult-sized	high
Sawyer	adult-sized robot	×	×	✓	×	1	adult-sized	mid
ARMAR-6	adult-sized robot	(✓) upper body	×	×	1	1	adult-sized	high
TOMM	adult-sized robot	(✓) upper body	×	×	×	1	adult-sized	high
Fetch	adult-sized robot	×	×	×	×	1	adult-sized	high
R1	adult-sized robot	(✓) upper body	×	✓	1	1	adult-sized	high
Sophia	adult-sized robot	(✓) upper body	✓	✓	 ✓ 	×	adult-sized	high
NICOL (ours)	adult-sized robot	(✓) upper body	 Image: A second s	\checkmark	 Image: A set of the set of the	1	adult-sized	mid

TABLE 1. Robotic platforms for social interaction and object manipulation. Price ranges are estimated; a low price range is below 10.000 EUR, a medium price range is below 100.000 EUR and a high price range is above this. Few platforms are designed for both social interaction and physical tasks while still being affordable. NICOL fills the gap of a robotic platform with both an upper humanoid body and humanoid head, capable of facial expressions, gestures and object manipulation in an adult-sized workspace at an affordable price tag.

of robots as related platforms.

There are many robotic manipulators for industry and research. Some common platforms are the UR-5 [9] and its related designs from Universal Robots, as well as the KUKA-DLR Lightweight Robot arm [10], and the Franka Research 3 [11]. These robotic arms have at least six degrees of freedom (DoF) to ensure solvable inverse kinematics for general 6-dimensional poses. These manipulators and their various end-effectors are primarily designed to handle objects and lack the social interaction capabilities that enable intuitive learning from, or teaching by, humans.

On the other side of the spectrum are platforms designed for social interaction with no, or limited, manipulation capabilities. For instance, the iCub [6] is available as a stand-alone 3-DoF head that can display facial expressions and perform gaze shifts. The Furhat [12] robot head is another example and can project animated or recorded faces for interaction purposes.

Infant-sized or toy-sized humanoids form another type of social platform. Their small size is advantageous for research, they have an affordable price, and they are inherently safe due to low motor speeds and weight. Wellknown examples are the Softbank Robotics NAO (formerly Aldebaran) and the ROBOTIS OP3 (descendant of the DARwIn-OP) [13]. While infant-sized platforms often have humanoid manipulators, they are too small to manipulate adult-sized items or reach them in domestic-scale environments.

The next larger category of humanoids is child-sized. A popular platform for developmental research is the iCub [6]. Its size resembles a child of about 90 cm, it is well-actuated with 53 DoF, features human-like hands and gaze shifts, and can display stylized facial expressions. NICOL

VOLUME *, 2016

is significantly more affordable and simple than the iCub, and is thereby less challenging to modify and/or customize for particular needs. The NimbRo-OP family of robots [14]-[16] are designed for the RoboCup soccer league, and focus on bipedal walking with arms for balance as well as getting up, but not on manipulation. As such, they do not even have hands that can be actuated. Softbank's Pepper robot [17] features a humanoid torso on top of an omnidirectional wheeled platform. It has 20 DoF in the upper body, but its arms and hands are designed for gesturing, not grasping. Poppy is an open-source 3D-printed robot designed for education, artists, and scientists [18]. It is completely customizable and comes with a display to show facial expressions and other information. Poppy can perform bipedal locomotion, but in its primarily advertised configuration, it features only very limited grasping functionalities. The concept or our NICO robot is quite similar to that of Poppy. NICO is an open platform for researching neuro-robotic models for human-robot interaction, as well as visuomotor learning [7]. Its head adapts the open iCub design, and, in contrast to Poppy, integrated LED arrays display stylized facial expressions. Furthermore, NICO has quite sophisticated grasping abilities to manipulate small objects, provided by two 6-DoF arms with fully functional child-sized anthropomorphic hands produced by Seed Robotics.

Adult-sized humanoids with manipulation abilities can generally handle real-life objects. However, such platforms are often prohibitively expensive and difficult to maintain—like PAL's Talos [19], which was introduced at a price of about 1 million euros—or their design is too non-humanoid for social interaction, like the Atlas initially developed by Boston Dynamics, and the PR2 from Willow Garage. Some adult-sized platforms lack object manipulation abilities entirely, like the Hanson Robotics Sophia robot. The ARMAR-6 from the KIT [20] is designed for complex collaborative tasks in industrial environments, but it neither features a human-like face nor can it display facial expressions. Likewise, the TOMM [21] and the onearmed Fetch robot [22] focus on manipulation and wheeled mobility, but not on social cues. The R1 from the IIT [23] is a wheeled platform with grippers that can display an animated face on a screen. A related design is the Rethink Robotics Sawyer, the successor of the well-known humanoid Baxter robot. It features a tablet for displaying animated eyes, but it only has a single end effector and is not humanoid.

Table 1 summarizes our analysis of robotic platforms. It is evident that there is a gap in available platforms that are designed for both social interaction and adult-sized object manipulation tasks. The few platforms that fulfil these two criteria are prohibitively expensive. We propose the open NICOL design to address this gap in the state of the art. In this article, we show that the NICOL is capable of social interaction tasks by extrapolating from previous work on the NICO, its "smaller brother". NICO and NICOL share the design of their head and facial expression mechanism(s), but importantly, the latter possesses a significantly increased arm strength and workspace size, allowing for more real object scenarios.

III. NICOL SEMI-HUMANOID ROBOTIC PLATFORM

NICOL is designed for research on robots that learn from and collaborate with humans. It can serve as a platform for integrating neuro-robotic visuomotor models combined with affective interaction, joint attention, and shared perception. NICOL, depicted in Fig. 2, consists of three main components: two manipulators with anthropomorphic hands, and the head with a facial expression display. NICOL adapts the head design of the child-sized humanoid NICO and combines it with two Robotis manipulators with five-fingered hands. NICOL can naturally collaborate with human partners on physical tasks, thanks to its ability to manipulate adult-size domestic objects and to use nonverbal cues like gaze or facial expression.

A. SOCIAL INTERACTION CAPABILITIES OF NICOL

The interaction capabilities of NICOL are based on the design of NICO. Here, we summarize the main findings and studies. The seven facial expressions of NICO, which are now shared by NICOL, are evaluated by Churamani et al. [24]. In a study with twenty participants from eleven different countries, five expressions (neutral, happiness, sadness, surprise and anger) are recognized by the participants with an accuracy of over 75%. The positive effect of the emotion display on the robot's subjective user rating is verified with a Godspeed questionnaire [25]. The freely programmable LED arrays are also used in learning emotion expression via reinforcement learning [26]. Ng et al. [27] used facial expressions in combination with neurocognitive models for social cue recognition and behavior control, and report that a natural language dialogue system benefits from the robot directly looking at the face of its interaction partner. Together with a more personalized conversation, this behavior made participants perceive NICO as more intelligent and likeable. Beik-Mohammadi et al. [28] show that using social gestures and more socially engaging dialogue enhanced the robot's perceived likeability and animacy. Finally, Kerzel et al. [3] show that participants guided by NICO in a visuomotor learning scenario via verbal requests and facial expressions, as compared to instructions by a human experimenter, rated the human-robot interaction as more positive and engaging. This increased engagement can improve the outcome of the learning scenario.

In summary, the studies strongly suggest that the design of NICOL's head, which is adapted from NICO, can effectively create social cues in terms of facial expressions and gaze. They also show that these social cues positively affect human-robot interactions.

B. ARMS AND HANDS

The arms of NICOL consist of two Robotis OpenManipulator-P¹ arms with 6 DoF and a payload of 3 kg. As end-effectors, two SeedRobotics RH8D adult-sized robot hands² with a manipulation payload of 750 g are used. All five fingers in the hand are tendon-operated. Each threesegment finger is operated by a single tendon. Each hand has eight actuated DoF—three in the wrist for rotation³, flexion, and abduction, two DoF in the thumb for abduction and flexion, and one DoF each for index finger flexion, middle finger flexion and combined flexion of the ring and little finger. The tendon mechanism allows the hand to coil around objects without further fine control. The arms of NICOL can reach up to 100 cm and therefore have a workspace slightly larger than that of an adult sitting at a table.

C. HEAD AND AUDIOVISUAL SENSING

The design of the shell of NICO's head is based on a modified version of the open-source iCub design [6]. The design balances human features with enough abstraction to avoid the uncanny valley effect [29]. Stylized facial expressions can be displayed by NICOL using three LED arrays that are placed behind the eyes (two 8x8 LED arrays) and mouth (one 16x8 array). Due to the head's material properties, the individual LEDs form lines shining through the head's material, as shown in Fig. 2. An internal speaker is used to facilitate spoken communication. The head is articulated with two DoFs for pitch and yaw movements. Two See3CAM CU135 cameras with a 4096 x 2160 (4K) resolution are mounted in the eye sockets of NICOL. Their

¹https://www.robotis.us/openmanipulator-p

²https://www.seedrobotics.com/rh8d-adult-robot-hand

³Note that the wrist rotation is redundant and therefore not used.



FIGURE 2. NICOL semi-humanoid platform. Left: NICOL jointly works on a grasping task with an experimenter (in the upper-right corner thereof, different facial expressions of NICOL are shown that can give feedback during the task). Right: NICOL object manipulation and expression in real life (top) and simulation (bottom).

fisheye lens has a field of view of 202 degrees. Two Soundman OKM II binaural microphones are placed at the sides of NICOL's head; due to the absence of head-internal fans or mechanics, the robot's ego noise is very low.

D. TABLE ENVIRONMENT AND SAFETY MEASURES

The table environment measures 100 x 200 cm and is located at a height of 74 cm. The head and arms of the NICOL are mounted to a vertical support at the rear centre of the table. Scaffolding at the corners of the table creates a visible delimitation for the workspace of the robot. A human interaction or collaboration partner is thereby implicitly made aware when actively reaching into the robot's workspace. For additional safety, two emergency shutdown buttons are integrated within the operator's workspace. The scaffolding can further be used to mount external sensors or devices, like cameras or light sources.

E. SOFTWARE AND SIMULATION ENVIRONMENT

The NICOL API is based on the Robot Operating System (ROS) middleware. All functionalities are provided through a modified version of the hardware controller delivered with the Robotis arms, extended to support the Seed Robotics hands, additional sensors, as well as drivers for custom hardware (facial expressions). The API integrates the MoveIt planning framework and provides a Pythonbased client. For prototyping and development, as well as for simulation of real-world scenarios, NICOL is realized in both the Gazebo simulation environment (see Fig. 3) and CoppeliaSim (see Fig. 2). The robot description is based on a URDF model, making it possible to import NICOL in many other simulators and environments.

IV. VISUOMOTOR LEARNING AND STATE-OF-THE-ART IK SOLVERS

To evaluate NICOL as a research platform for neurorobotic visuomotor learning, we first transfer an end-to-end neural visuomotor learning approach from NICO, a smaller humanoid robot, to NICOL. Based on the results, we propose two variations of a modular neuro-hybrid approach. In these modular approaches, we use a neural network for image processing and different IK solver for the kinematic. We evaluate the baseline performance of state-of-the-art IK solvers that are established in the robotics community and suggest a novel hybrid neuro-genetic approach that can handle the added kinematic complexity of a humanoid.

A. STATE-OF-THE-ART IK APPROACHES

Solving IK is a well-known and studied problem in the robotics community. The *MoveIt planning framework* [30], as part of the ROS middleware, offers the required flexibility to utilize different IK solvers for an experimental setup. It enables a plug-in-based configuration of inverse kinematics solvers as well as high-level motion planners. Mainly third-party kinematics libraries, contributed by foundations or private sector organizations are available, and these have become state-of-the-art due to the popularity of the ROS framework. The inverse kinematics and planning functionalities are provided in the form of ROS services and offer a variety of parameterized kinematics calculations.

The *Kinematics and Dynamics Library* (KDL) by Orocos [31] is de-facto the most widely adapted inverse kinematics solution in the ROS ecosystem. It supports kinematic chains with a minimum of 6 DoF. KDL utilizes the well-known pseudo-inverse Jacobian method to determine suitable joint configurations. As the Jacobian matrix holds the partial derivatives between joint space and Cartesian



FIGURE 3. Left: Size comparison between an adult human's hand, NICOL, and NICO. Upscaling the hands of NICO allows NICOL to manipulate everyday objects and tools. Right: NICOL in a virtual Gazebo environment during a dexterous manipulation task and dropping a baseball into a basket.



FIGURE 4. Neural architecture for end-to-end visuomotor learning.

space, the difference between the so-called seed state (typically the initial joint configuration) and the target state can be calculated in Cartesian space. Forward kinematics are used on the seed state to determine the current endeffector pose. A suitable joint configuration for achieving the target Cartesian end effector pose can then be obtained by calculating the error between the initial pose and the target pose, and multiplying it with the inverted Jacobian. The algorithm makes use of the Newton-Raphson method and iteratively minimizes the error between the previous and next target poses.

With the intention of improving the coverage of the robot workspace, TracLabs identified key issues in the KDL framework and proposed *TRAC-IK* in 2015 [32]. Remarkable about the approach is that two IK solvers are executed in parallel. One is a modified reimplementation of KDL called KDL-RR. The method additionally checks for local

optima during the iterative improvement step and moves the next seed far enough away from them. The second solver, SQP-SS, formulates the IK problem as a non-linear sequential quadratic programming optimization problem. TRAC-IK returns the solution of the solver that first terminates with a valid joint configuration. In contrast to the pseudo-inverse Jacobian, SQP-SS is particularly capable of directly handling constraints, such as the joint limits, during the optimization step. This ability is particularly relevant for humanoid robots like NICOL, whose joints are constrained to mimic the human range of motion.

Starke et al. introduced the *Bio-IK* method [33] [34]. The algorithm combines multiple bio-inspired optimization methods, most importantly evolutionary and particle swarm optimization, in order to solve the inverse kinematics problem. In difference to KDL and TRAC-IK, it does not rely on any Jacobian mathematics. The evolutionary



FIGURE 5. Three architectures for visuomotor learning.

algorithm is fundamental to the approach. Every individual in the population corresponds to a joint space robot pose. Momentums are assigned to every individual by hybrid particle swarm optimization, and these momentums are continuously updated during runtime. Besides selection, mutation and recombination, also elitism is used to prevent the deterioration of already-found solutions. Local search is executed on the elites, and simulates mutations on single genes. The classical MoveIt interface only allows setting the pose or position goals. Bio-IK additionally offers extended functionality in a separate service [35], where multiple custom goal types are available, e.g. linear endeffector trajectories or minimal joint displacement configurations. The various custom goal types can be combined into a single request, as each custom goal type is treated as a weighted partial cost function by the algorithm. The service also offers an approximate mode, leading to lower precision in the IK solutions but enabling higher coverage of the workspace.

B. END-TO-END NEUROROBOTIC VISUOMOTOR LEARNING

The supervised neural end-to-end visuomotor approach, first introduced in [36], learns to map a single object's visual input in the robot's workspace to joint configurations for reaching for the object. The approach was evaluated with a grasp success rate of approximately 85% on NICO with 5 DoF in a workspace of size 30x40 cm and a training set of 400 samples. In this paper, we evaluate if the larger workspace and more complex NICOL kinematics influence the learning outcome. The neural architecture is shown in Figure 4. In Section V-B we report experimental results.

C. HYBRID NEURO-GENETIC VISUOMOTOR LEARNING

Genetic algorithms are based on evolutionary selection, recombination and mutation processes of a population of individuals, modelled by their chromosomes [37]. Each chromosome encodes a potential solution to a task. Following Kerzel et al. [38], the chromosomes encode a joint configuration to reach a given pose. During each iteration, these chromosomes, expressed as individuals, are ranked according to their fitness, i.e., the position and orientation error compared to the goal pose. Successful individuals are copied into the next generation, with possible random mutations to their chromosome-encoded joint configuration. Niching is used to preserve diverse chromosomes; on every CPU of the computer system, a different evolution is implemented, thus full utilizing multiple cores with minimal overhead.

While genetic algorithms are adept at escaping local minima, they often have issues further optimizing found solutions. To overcome this limitation, we hybridize the genetic algorithm with gradient-based Sequential Least SQuares Programming (SLSQP). This computationally expensive procedure is only applied to the n-best individuals of the population. While preliminary experiments have indicated that SLSQP is prone to be attracted by local minima, the initialization with solutions found by the genetic algorithm can overcome this shortcoming. Optimized hyperparameters for the genetic algorithm and SLSQP were adapted from [38], where they yielded good results for NICO.

While the hybrid genetic algorithm provides IK solutions, a neural component is utilized to determine the position of the grasp object from visual input. The architecture is shown in Figure 5, experiments are reported in Section V-B.

V. IK AND VISUOMOTOR LEARNING EXPERIMENTS

We first curate a new dataset (Section V-A) for visuomotor learning on the NICOl platform. A first analysis of the dataset allows us to get an insight into NICOL's workspace. In Section V-B, we evaluate how well our neural end-to-end and a hybrid neuro-genetic visuomotor learning approach can be transferred from the child-sized humanoid NICO to the adult-sized NICOL. We evaluate how the significantly increased workspace and the more complex manipulator affect the learning outcome. In Section V-C and V-D, we present a modular neuro-hybrid approach in which we evaluate state-of-the-art IK solvers and a neuro-generic approach based on GAIKPy.

A. DATASET FOR VISUOMOTOR LEARNING ANALYSIS OF NICOL'S WORKSPACE

We collect three datasets⁴, each consisting of 10,000 samples in a virtual environment. The datasets differ in the number of active joints of NICOL, ranging from six to eight. As the joints of NICOL are constrained to mimic the human range of motion, we hypothesize that the additional degrees of freedom beyond six will increase the number of reachable grasp poses. Each sample contains the following: 1) an image from the egocentric perspective of NICOL with a grasp-object placed onto the table on the right side of the workspace in a 100 x 100 cm area; the images are cropped and resized to 63 x 96 pixels. 2) The x- and y-coordinates of the grasp-object on the table. 3) The configuration of NICOLs 8 arm joints that result in a grasp-pose. The grasp pose has the hand pointing forward with the palm touching the grasp-object at the middle of its height, as shown in Fig. 4 (right side). We use a genetic algorithm GAIKPy introduced in [38] for computing the joint configuration for grasping the object that is placed at a random position within the robot's workspace. If no suitable joint configuration can be found with the genetic algorithm, the sample is rejected, and another random object position is generated.

Fig. 6 shows the distribution of all successful grasp samples in NICOL's workspace for the three experimental conditions. Due to the rigid mechanical constraints of our defined grasping pose, the joint configuration and selfcollision avoidance, the area within the defined workspace for which grasp poses can be found increases with the number of active joints. Only with eight active joints, good coverage of the workspace can be achieved. Therefore, we focus our analysis on this dataset.

B. END-TO-END NEUROROBOTIC VISUOMOTOR LEARNING EXPERIMENTS

First, we evaluate a neural end-to-end learning approach using the neural architecture adopted from [36]. The input to the neural network is a 96x64 pixel image from NICOL's perspective of its workspace with a single object placed in it. Depending on the dataset, the output of the network is the seven or eight joint values of NICOL's right manipulator. The hyperparameters were optimized using Hyperopt [39]. The model was trained with an Adam optimizer and a batch size of 35, and the MSE over the joint configuration as a loss function on a random 90-10 split of the entire dataset. Hyperparameter ranges and results after are shown in Table 2, and the resulting architecture is shown in Fig. 4.

We use 10-fold cross-validation to obtain our results. For evaluation, we denormalize the joint output values and calculate the robotic hand's pose using the known forward kinematics of NICOL. The mean position error is 0.42 ± 0.21 meters and the mean orientation error is 7.72 ± 5.88 degrees. Based on previous work [36], [38], we count a grasp as successful if the position error of the resulting pose is < 10 mm and the sum of orientation errors is < 20 degrees. These limits are regularly exceeded. Table 3 summarizes the results with a single-digit grasp accuracy. We conclude that the end-to-end approach cannot be transferred directly to NICOL. We attribute this result to several factors: The workspace of NICOL with 100x100 cm is larger than the workspace of NICO with 30x40 cm; more importantly, the arms of NICOL afford more complex motions and more diverse joint configurations during grasping.

C. MODULAR APPROACH WITH CLASSICAL IMAGE PROCESSING, NEURAL COORDINATE TRANSFORMATION AND IK SOLVERS

Next, we analyse if the low accuracy of neural end-to-end visuomotor learning can be attributed to issues in visually locating the grasp object in the image, transforming the image coordinates to world coordinates or the computation of inverse kinematics. In previous work, we addressed this issue using a modular, hybrid neuro-genetic approach [38]. We adopt this approach to the NICOL and use three modules as shown in Fig. 5: Classical image processing to locate the object in the image, a neural-network to transform the image coordinates into world coordinates and existing state-of-the-art and a novel generic IK solvers. To extract the grasp object's position from the non-downsampled image (4208x3120) using classical image processing, we detect the top and the bottom of the grasp-object with a standard color-based detector, resulting in the x, y position and the radius of the base and top parts of our graspobject. These 6 input parameters are fed into an MLP, which outputs the x- and y-coordinates of the object on the table. The network consists of one dense layer with 180 neurons and ReLU activation; the architecture is based on

⁴The datasets will be released at https://www.inf.uni-hamburg.de/en/ inst/ab/wtm/research/corpora.html



FIGURE 6. Left: Placement of the 10,000 collected samples (green) versus the rejected samples (red) for 6 (left), 7 (middle) and 8 (right) DoF. The images show the continuity of the IK map, i.e., to what extent a smooth IK trajectory ends up as a smooth joint trajectory if each instantaneous pose is individually passed through the IK algorithm.

Parameter	Parameter range	Result 8 DoF
# filters conv ₁	8-32 (step size 8)	32
# filters conv ₂	8-32 (step size 8)	8
kernel size conv ₁	3-5 (step size 1)	3
kernel size conv2	3-5 (step size 1)	5
# dense_1 neurons	16-1024 (step size 16)	256
# dense_2 neurons	16-256 (step size 16)	160
normalisation	[01] or [-11]	[01]
image size	48x32, 96x64, 192x128	96x64
learning rate	[0.01, 0.001, 0.0001, 0.00001]	0.0001

TABLE 2. Ranges and results of hyperparameter optimization of the end-to-end visuomotor architecture for 7 and 8 DoF.

Experiment	DoF	Mean pos. error	Median pos. Error	Mean orient. error	Median orient. error	Grasp accuracy
End-to-end	8	$0.4161 \pm 0.2085 \text{ m}$	0.3908 m	7.72 ± 5.88 deg.	6.01 deg.	0.01 %

TABLE 3. Results for the neural end-to-end approach for 8 DoF. Results indicate that the neural end-to-end approach cannot be transferred to the more complex kinematics of NICOL.

Parameter	Parameter range	Result	
# dense neurons	10-200 (step size 10)	180	
normalisation	[01] or [-11]	[01]	

TABLE 4. Ranges and results of hyperparameter optimization of the neural image-to-coordinate transformation.

[38] and was optimized with Hyperopt [39]. Table 4 shows the optimization ranges and results.

We use 10-fold cross-validation for all reported results. First, we evaluate the error of the transformation. As shown in the first line in Table 5, the resulting mean position error is 0.0039 ± 0.0023 m, and the median position error is 3 mm. Applying the above-established criteria for a successful grasp (obviously without considering orientation) resulting in a grasp accuracy of 98.47%. Next, we evaluate the complete grasp architecture by using IK solvers to compute the robot's joint configuration based on the predicted object coordinates. Table 5 summarizes the results for the state-of-the-art IK solvers KDL, TRACK-IK and BIO-IK as well as for our novel genetic IK solver GAIKPy. Again, we apply the criteria for a successful grasp

(position error is < 10 mm, and the sum of orientation errors is < 20 degrees). We achieve accuracies of 90.52%, 90.29% and 89.60% for KDL, TRAC-IK and BIO-IK. In both conditions, GAIKPy shows the highest accuracy with 98.45% and 92.85%.

D. NOVEL NEURAL IMAGE-TO-COORDINATE TRANSFORMATION AND GENETIC ALGORITHM

To evaluate the limits of neural visuomotor learning, we modify the end-to-end architecture presented above to regress directly the x- and y-coordinates of the object on the table instead of the joint configurations. We then use state-of-the-art IK solvers and GAIKPy to compute a joint configuration that reaches for the output coordinate with a given hand orientation.

We use 10-fold cross-validation to obtain the results in Table 6. Applying the criteria for a successful grasp, 96.51%, 96.2% and 95.71% for KDL, TRAC-IK and BIO-IK. GAIKPy shows the highest accuracy with 99.17%.

E. DISCUSSION

In Section V-A, we show that 8 DoF are needed to cover sufficiently large parts of NICOL's intended workspace. In

Experiment	DoF	Mean pos. error	Median pos. Error	Mean orient. error	Median orient. error	Grasp accuracy
transformation	-	0.0039 ± 0.0023	0.00343	-	-	98.47
KDL	8	0.0229 ± 0.1031	0.00444	0.519 ± 2.924	1.817e-05	90.52
TRAC-IK	8	0.0244 ± 0.1068	0.00445	0.564 ± 3.042	1.630e-05	90.29
BIO IK	8	0.0280 ± 0.1144	0.00448	0.686 ± 3.343	1.987e-04	89.60
GAIKPy	8	0.0067 ± 0.0302	0.00436	0.337 ± 0.687	0.1186	92.85

TABLE 5. Results for the modular approach with Classical Image Processing, Neural Coordinate Transformation and IK solvers for 8 DoF. GAIKPy achieves the highest grasp accuracy of all evaluated IK solvers.

Experiment	DoF	Mean pos. error	Median pos. Error	Mean orient. error	Median orient. error	Grasp accuracy
transformation	-	0.0018 ± 0.0012	0.00158	-	-	99.98
KDL	8	0.0191 ± 0.1034	0.00125	0.487 ± 2.834	1.837e-05	96.51
TRACK-IK	8	0.0209 ± 0.1078	0.00125	0.540 ± 2.979	1.627e-05	96.20
BIO-IK	8	0.0233 ± 0.1130	0.00126	0.625 ± 3.196	2.035e-04	95.71
GAIKPy	8	0.0019 ± 0.0102	0.00122	0.465 ± 1.151	0.1841	99.17

TABLE 6. Results for our novel hybrid approach combining neural image-to-coordinate transformation and genetic IK with GAIKPy for 8 DoF. Again, GAIKPy achieves the highest grasp accuracy.

contrast, non-humanoid industrial robot arms often only need 6 DoF to cover their workspace. We attribute this finding to NICOL design mimicking human motion, resulting in a more humanoid distribution of DoF along the arm and, more importantly, severe constraints on its joints. Next, we show that the increased kinematic complexity of NICOL can not be handled by an end-to-end neural visuomotor learning approach in V-B. However, we successfully apply a hybrid modular approach using image processing and a neural coordinate transformation with an IK solver in V-C. We also demonstrate that NICOL integrates well with the state-of-the-art IK solvers KDL, TRACK-IK and BIO-IK. However, the best results were achieved by using our genetic algorithm-based GAIKPy IK solver. We attribute this finding to the combination of more than 6 DoF in conjunction with joint constraints, posing a challenge for traditional state-of-the-art IK solvers, as shown in Tables 5 and 6. Finally, in V-D, we positively evaluated a novel hybrid approach that uses a convolutional neural network to extract target grasp positions from an image in conjunction with an IK solver. Again, our GAIKPy approach yields the highest grasp accuracy with a success rate of over 99%.

This added accuracy comes at the cost of longer processing times. We give each IK solver 1 second to compute a solution. While GAIKPy fully utilizes this time, the state-of-the-art solvers show a different behaviour: If cases, where they can find a good solution, this often happens much faster, however, they don't utilize the time budget to find suitable solutions for difficult poses.

We argue that in human-robot collaboration, these slower IK solutions are less problematic. For a humanoid robot, collaborating with a human, we face different challenges compared to an industrial robotic arm. For humanrobot collaboration, safe and more importantly, predictable motions are essential. For safety reasons, we need to limit the technical possible maximum joint speed of the humanoid anyway. The slower but more accurate hybrid neuro-genetic approach yields better accuracy scenarios, while still fulfilling the time constraints for collaborative human-robot scenarios.

VI. CONCLUSION

In summary, we introduce NICOL as a novel semihumanoid robotic platform for research in social robotics and physical human-robot collaboration. NICOL is a semihumanoid robot that combines the already positively evaluated social interaction capabilities of the NICO platform [7] with adult-sized manipulators and five-fingered hands for the everyday handling of objects in a robot environment. NICOL is intended for research on robots collaborating with, and learning from, humans in scenarios that require both advanced object manipulation and interaction abilities. NICOL addresses a gap in the state of the art for affordable and open robot designs that enable search in human-robot interaction and social robotics on the one side, but also allow precise grasping of larger objects, as shown in our experiments.

We focus our evaluation of NICOL on inverse kinematics for a reach-for-grasp task for three reasons: First, grasping is an essential ability of humanoid robots that is the foundation of many more complex abilities and can be applied in many experimental scenarios. Secondly, our main incentive to design the adult-sized NICOL as a successor of the child-sized NICO is to create a robotic platform that has a workspace and can handle objects comparable to an adult human, thus being able to be used in physical collaboration scenarios. Thirdly, we have observed that standard IK solvers have issues with kinematic chains that have more than 6 DoF while at the same time having strict constraints on the individual joint limits to mimic the human range of motion. We address these last points with a detailed set of experiments in which we evaluate neural end-to-end learning approaches for visuomotor learning and compare state-of-the-art IK solvers against a novel neural and hybrid neuro-genetic IK approaches. We show that neural end-to-end learning is challenging due to the more complex kinematics of NICOL and overcame this challenge with a neuro-genetic approach. Furthermore, we demonstrate that our novel hybrid neuro-genetic approach outperforms classical IK solvers in terms of accuracy by taking advantage of the relaxed time constraints in humanrobot collaboration scenarios.

In future work, we will apply the grasp learning in a real-world scenario, in which human-robot interaction and visuomotor abilities are combined, and participants aid and are aided by NICOL during a collaborative task.

REFERENCES

- A. L. Thomaz, G. Hoffman, and C. Breazeal. Reinforcement learning with human teachers: Understanding how people want to teach robots. In ROMAN 2006 - The 15th IEEE International Symposium on Robot and Human Interactive Communication, pages 352–357, Sep. 2006.
- [2] Vasant Srinivasan and Leila Takayama. Help me please: Robot politeness strategies for soliciting help from humans. In Proceedings of the 2016 CHI conference on human factors in computing systems, pages 4945–4955. ACM, 2016.
- [3] Matthias Kerzel, Theresa Pekarek-Rosin, Erik Strahl, Stefan Heinrich, and Stefan Wermter. Teaching nico how to grasp: An empirical study on crossmodal social interaction as a key factor for robots learning from humans. Frontiers in Neurorobotics, Jun 2020.
- [4] Bilge Mutlu, Allison Terrell, and Chien-Ming Huang. Coordination mechanisms in human-robot collaboration. In Proceedings of the Workshop on Collaborative Manipulation, 8th ACM/IEEE International Conference on Human-Robot Interaction, pages 1–6. Citeseer, 2013.
- [5] FRANKA RESEARCH 3. http://https://www.franka.de/de/research/, 2023. [Online; accessed 12-March-2023].
- [6] Giorgio Metta, Lorenzo Natale, Francesco Nori, Giulio Sandini, David Vernon, Luciano Fadiga, Claes Von Hofsten, Kerstin Rosander, Manuel Lopes, José Santos-Victor, et al. The iCub humanoid robot: An open-systems platform for research in cognitive development. Neural Networks, 23(8-9):1125–1134, 2010.
- [7] Matthias Kerzel, Erik Strahl, Sven Magg, Nicolás Navarro-Guerrero, Stefan Heinrich, and Stefan Wermter. Nico - neuro-inspired companion: A developmental humanoid robot platform for multimodal interaction. In Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pages 113–120, Aug 2017.
- [8] Justinas Mišeikis, Inka Brijacak, Saeed Yahyanejad, Kyrre Glette, Ole Jakob Elle, and Jim Torresen. Transfer learning for unseen robot detection and joint estimation on a multi-objective convolutional neural network. In 2018 IEEE International Conference on Intelligence and Safety for Robotics (ISR), pages 337–342. IEEE, 2018.
- [9] Universal Robots. UR5 Technical Specifications. Accessed 2020-12-27.[Online]. Available at: http://www.universal-robots.com.
- [10] Rainer Bischoff, Johannes Kurth, Günter Schreiber, Ralf Koeppe, Alin Albu-Schäffer, Alexander Beyer, Oliver Eiberger, Sami Haddadin, Andreas Stemmer, Gerhard Grunwald, et al. The kuka-dlr lightweight robot arm-a new reference platform for robotics research and manufacturing. In ISR 2010 (41st international symposium on robotics) and ROBOTIK 2010 (6th German conference on robotics), pages 1–8. VDE, 2010.
- [11] Franka Emika. Franka research 3. https://www.franka.de/research, 2022. Accessed 2022-12-27.
- [12] Samer Al Moubayed, Jonas Beskow, Gabriel Skantze, and Björn Granström. Furhat: a back-projected human-like robot head for multiparty human-machine interaction. In Cognitive behavioural systems, pages 114–130. Springer, 2012.
- [13] Inyong Ha, Yusuke Tamura, Hajime Asama, Jeakweon Han, and Dennis W Hong. Development of open humanoid platform DARwIn-OP. In SICE Annual Conference 2011, pages 2178–2181. IEEE, 2011.
- [14] Max Schwarz, Julio Pastrana, Philipp Allgeuer, Michael Schreiber, Sebastian Schueller, Marcell Missura, and Sven Behnke. Humanoid teensize open platform NimbRo-OP. In Robot Soccer World Cup, pages 568–575. Springer, 2013.
- VOLUME *. 2016

- [15] Philipp Allgeuer, Hafez Farazi, Michael Schreiber, and Sven Behnke. Child-sized 3D printed igus humanoid open platform. In 15th International Conference on Humanoid Robots (Humanoids), Seoul, Korea, 2015.
- [16] Grzegorz Ficht, Hafez Farazi, André Brandenburger, Diego Rodriguez, Dmytro Pavlichenko, Philipp Allgeuer, Mojtaba Hosseini, and Sven Behnke. NimbRo-OP2X: Adult-sized open-source 3D printed humanoid robot. In 18th International Conference on Humanoid Robots (Humanoids), Beijing, China, 2018.
- [17] Amit Kumar Pandey and Rodolphe Gelin. A mass-produced sociable humanoid robot: Pepper: The first machine of its kind. IEEE Robotics & Automation Magazine, PP:1–1, 07 2018.
- [18] Matthieu Lapeyre, Pierre Rouanet, Jonathan Grizou, Steve N'Guyen, Alexandre Le Falher, Fabien Depraetre, and Pierre-Yves Oudeyer. Poppy: Open source 3d printed robot for experiments in developmental robotics. In International Conference on Development and Learning and on Epigenetic Robotics (ICDL-EpiRob), pages 173–174, 2014.
- [19] Olivier Stasse, Thomas Flayols, Rohan Budhiraja, Kevin Giraud-Esclasse, Justin Carpentier, Joseph Mirabel, Andrea Del Prete, Philippe Souères, Nicolas Mansard, Florent Lamiraux, et al. Talos: A new humanoid research platform targeted for industrial applications. In 2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids), pages 689–695. IEEE, 2017.
- [20] Tamim Asfour, Mirko Waechter, Lukas Kaul, Samuel Rader, Pascal Weiner, Simon Ottenhaus, Raphael Grimm, You Zhou, Markus Grotz, and Fabian Paus. Armar-6: A high-performance humanoid for humanrobot collaboration in real-world scenarios. IEEE Robotics & Automation Magazine, 26(4):108–121, 2019.
- [21] Emmanuel Dean-Leon, Brennand Pierce, Florian Bergner, Philipp Mittendorfer, Karinne Ramirez-Amaro, Wolfgang Burger, and Gordon Cheng. Tomm: Tactile omnidirectional mobile manipulator. In 2017 IEEE International Conference on Robotics and Automation (ICRA), pages 2441–2447. IEEE, 2017.
- [22] Fetch Robotics, 2019.
- [23] Alberto Parmiggiani, Luca Fiorio, Alessandro Scalzo, Anand Vazhapilli Sureshbabu, Marco Randazzo, Marco Maggiali, Ugo Pattacini, Hagen Lehmann, Vadim Tikhanoff, Daniele Domenichelli, et al. The design and validation of the r1 personal humanoid. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 674–680. IEEE, 2017.
- [24] Nikhil Churamani, Matthias Kerzel, Erik Strahl, Pablo Barros, and Stefan Wermter. Teaching emotion expressions to a human companion robot using deep neural architectures. In International Joint Conference on Neural Networks (IJCNN), pages 627–634, Anchorage, Alaska, May 2017. IEEE.
- [25] Christoph Bartneck, Dana Kulić, Elizabeth Croft, and Susana Zoghbi. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. International Journal of Social Robotics, 1(1):71–81, Jan 2009.
- [26] Nikhil Churamani, Pablo Barros, Erik Strahl, and Stefan Wermter. Learning empathy-driven emotion expressions using affective modulations. In Proceedings of the International Joint Conference on Neural Networks (IJCNN 2018), pages 1400–1407. IEEE, IEEE, Jul 2018.
- [27] Hwei Geok Ng, Paul Anton, Marc Brügger, Nikhil Churamani, Erik Fließwasser, Thomas Hummel, Julius Mayer, Waleed Mustafa, Thi Linh Chi Nguyen, Quan Nguyen, Marcus Soll, Sebastian Springenberg, Sascha Griffiths, Stefan Heinrich, Nicolás Navarro-Guerrero, Erik Strahl, Johannes Twiefel, Cornelius Weber, and Stefan Wermter. Hey robot, why don't you talk to me? In Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pages 728–731, Aug 2017.
- [28] Hadi Beik-Mohammadi, Nikoletta Xirakia, Fares Abawi, Irina Barykina, Krishnan Chandran, Gitanjali Nair, Cuong Nguyen, Daniel Speck, Tayfun Alpay, Sascha Griffiths, Stefan Heinrich, Erik Strahl, Cornelius Weber, and Stefan Wermter. Designing a personality-driven robot for a human-robot interaction scenario. In 2019 IEEE International Conference on Robotics and Automation (ICRA), pages 4317–4324, Montreal, Canada, May 2019.
- [29] Masahiro Mori, Karl F MacDorman, and Norri Kageki. The uncanny valley [from the field]. IEEE Robotics & Automation Magazine, 19(2):98–100, 2012.
- [30] David Coleman, Ioan Sucan, Sachin Chitta, and Nikolaus Correll. Reducing the barrier to entry of complex robotic software: a moveit! case study, 2014.

- [31] R. Smits. KDL: Kinematics and Dynamics Library. http://www.orocos. org/kdl.
- [32] Patrick Beeson and Barrett Ames. Trac-ik: An open-source library for improved solving of generic inverse kinematics. In 2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids), pages 928–935, 2015.
- [33] Sebastian Starke, Norman Hendrich, and Jianwei Zhang. A memetic evolutionary algorithm for real-time articulated kinematic motion. In 2017 IEEE Congress on Evolutionary Computation (CEC), pages 2473– 2479, 2017.
- [34] Philipp Sebastian Ruppel. Performance optimization and implementation of evolutionary inverse kinematics in ros. Master's thesis, University of Hamburg, Department of Computer Science, 2017.
- [35] Github tams-group/bio_ik_service: Ros service for calling bioik from python or java. https://github.com/TAMS-Group/bio_ik_service. Accessed: 2021-01-06.
- [36] Matthias Kerzel and Stefan Wermter. Neural end-to-end self-learning of visuomotor skills by environment interaction. In Lintas A., Rovetta S., Verschure P., and Villa A., editors, Artificial Neural Networks and Machine Learning – ICANN 2017, volume 10613 of Lecture Notes in Computer Science, pages 27–34. Springer, Cham, Sep 2017.
- [37] Stephen Marsland. Machine Learning: An Algorithmic Perspective, Second Edition. Chapman & Hall/CRC, 2nd edition, 2014.
- [38] Matthias Kerzel, Josua Spisak, Erik Strahl, and Stefan Wermter. Neurogenetic visuomotor architecture for robotic grasping. In Igor Farkaš, Paolo Masulli, and Stefan Wermter, editors, Artificial Neural Networks and Machine Learning – ICANN 2020, LNCS, pages 533–545. Springer, 2020.
- [39] James Bergstra, Daniel Yamins, and David Cox. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In 30th International Conference on Machine Learning (ICML 2013), pages 115–123, 2013.

...