NimbRo-OP2X: Adult-sized Open-source 3D Printed Humanoid Robot

Grzegorz Ficht, Hafez Farazi, André Brandenburger, Diego Rodriguez, Dmytro Pavlichenko, Philipp Allgeuer, Mojtaba Hosseini, and Sven Behnke

Abstract-Humanoid robotics research depends on capable robot platforms, but recently developed advanced platforms are often not available to other research groups, expensive, dangerous to operate, or closed-source. The lack of available platforms forces researchers to work with smaller robots, which have less strict dynamic constraints or with simulations, which lack many real-world effects. We developed NimbRo-OP2X to address this need. At a height of 135 cm our robot is large enough to interact in a human environment. Its low weight of only 19 kg makes the operation of the robot safe and easy, as no special operational equipment is necessary. Our robot is equipped with a fast onboard computer and a GPU to accelerate parallel computations. We extend our already opensource software by a deep-learning based vision system and gait parameter optimisation. The NimbRo-OP2X was evaluated during RoboCup 2018 in Montreál, Canada, where it won all possible awards in the Humanoid AdultSize class.

I. INTRODUCTION

Using humanoid robots in daily life scenarios has been the motivation for many research and development efforts. More than 40 years ago, the Waseda robot [1] started the quest for building humanoid robots that could coexist in our environment. Since then, much work has been done in this direction, but we still do not have any humanoid robot capable of coping with the uncertainty and complexity of real-life scenarios. From a technical point of view, the difficulties arise not only from the lack of actuation capacity for the generation of dexterous movements to interact with the environment, but also from deficits on the cognitive level to understand the surrounding and plan tasks accordingly.

Humanoid robotics has recently seen progress in terms of hardware, with the introduction of multiple new and capable platforms. The technological advancements made by these platforms however, are met with drawbacks. The precision technology used and their complex design, both require that the robot is operated by an experienced user with safety equipment such as gantries to minimise possible risk. Many of these advanced platforms are not sold to other researchers. If they were offered, their price tag would be stiff. Due to the closed-source nature of such projects, only a selected few have access to these platforms. This forces other research groups to work with smaller robots in scaled-down environments. Due to the scale difference, the results may not always translate into real-world applications. Another option is to use simulations, which may lack faithful reproduction of real-world effects.



Fig. 1. NimbRo-OP2X in comparison to other platforms, in terms of height and weight. From left to right: Honda Asimo, NimbRo-OP2X, Toyota T-HR3, Boston Dynamics Atlas, DLR-TORO. Out of these platforms, only the NimbRo-OP2X is openly available to any research group.

In this paper, we introduce NimbRo-OP2X, which is a 135 cm tall, humanoid robot with a completely 3D printed structure and deep learning capabilities. NimbRo-OP2X is a substantial improvement over our previous NimbRo-OP and NimbRo-OP2 robots. By incorporating new intelligent actuators and a GPU-enabled computing unit into the design, we make NimbRo-OP2X applicable in more general use cases. In contrast to similarly sized platforms (Fig. 1), NimbRo-OP2X is very light and easy to operate, while having considerable torque output to efficiently walk and perform full-body motions. By using widely available manufacturing techniques and open sourcing both, the hardware and software, we aim to foster accelerated humanoid research in real-world applications.

II. RELATED WORK

For several years, research on motion generation for humanoid robots was carried out mainly by a few companies or universities because of the large costs associated with platform development. The most prominent examples include the HRP2 [2] and Honda's Asimo [3]. The situation did not change only until recently, when the rising interest in humanoid robotics led to the emergence of new humansized platforms. Robots such as Petman [4], DLR-TORO [5], NASA Valkyrie [6] or the recently introduced Toyota T-HR3 [7] are still very expensive or, more importantly, these robots cannot be purchased by other research groups. In 2017, we introduced a larger open-source robot [8], that in comparison with other platforms could be considered as lowcost thanks to its 3D-printed structure and commercial offthe-shelf actuators. The robot proved to be capable enough to win the RoboCup 2017 soccer tournament in its debut [9].

Shortcutting the limiting factor of hardware availability,

All authors are with the Autonomous Intelligent Systems (AIS) Group, Computer Science Institute VI, University of Bonn, Germany. Email: ficht@ais.uni-bonn.de. This work was partially funded by grant BE 2556/13 of the German Research Foundation (DFG).

 TABLE I

 NIMBRO-OP2X SPECIFICATIONS.

Туре	Specification	Value
General	Height & Weight	135 cm, 19 kg
	Battery	4-cell LiPo (14.8 V, 8.0 Ah)
	Battery life	20–40 min
	Material	Polyamide 12 (PA12)
РС	Mainboard	Z370 Chipset, Mini-ITX
	CPU	Intel Core i7-8700T, 2.7-4.0 GHz
	GPU	GTX 1050 Ti, 768 CUDA Cores
	Memory	4 GB DDR4 RAM, 120 GB SSD
	Network	Ethernet, Wi-Fi, Bluetooth
	Other	$8 \times \text{USB}$ 3.1, $2 \times \text{HDMI}$, DisplayPort
Actuators	Total	34 × Robotis XM-540-W270-R
	Stall torque	12.9 Nm
	No load speed	37 rpm
	Control mode	Torque, Velocity, Position, Multi-turn
Sensors	Encoders	12 bit/rev
	Joint torque	12 bit
	Gyroscope	3-axis (L3G4200D chip)
	Accelerometer	3-axis (LIS331DLH chip)
	Camera	Logitech C905 (720p)
	Camera lens	Wide-angle lens with 150° FOV

advancing the state of the art can be done by means of simulation [10]. However, simulators still have several challenges to face until they will provide comparable results to reality. On the other hand, simulations can be performed without any cost, risk, and—most importantly—fast. In our recent work [11], simulation and real-robot experiments were combined for gait optimisation of smaller robots. In this paper, we extend this work to our new larger robot, with more complex gait sequences, to achieve a self-tuned gait that is able to withstand strong impacts.

Deep learning currently plays a key role in enhancing the cognitive abilities of robots. It has made its way to other research areas in robotics such as: speech recognition [12], motion planning [13], learning dynamics [14], and others. Semantic segmentation addresses most of the perception needs for autonomous robots. Due to heavy computational requirements, many deep learning based approaches for semantic segmentation are not suitable for robotics applications. Romera at al. [15] recently proposed "ERFNet" which is ideal for lightweight computers like Jetson TX1. ERFNet is an encoder-decoder architecture for real-time semantic segmentation tasks. We hope that open-sourcing an AdultSize low-cost humanoid platform, equipped with a lightweight GPU for parallel computing, will foster research not only limited to motion planning but deep learning as well.

III. HARDWARE DESIGN

The NimbRo-OP2X builds on the ideas of providing an open-source, larger-sized humanoid robot to the research community at a fraction of the cost of similarly sized platforms. At a height of 135 cm and approximately 19 kg of weight, it aims to provide a powerful middle-ground between smaller low-cost robots, and larger cost-prohibitive



Fig. 2. Redesigned hip joint utilising 3D printed double helical gears. CAD model (left) in comparison to the real joint (right).

platforms. To achieve this, we utilise a combination of 3D printing, commercially available actuators, parallel kinematics and external gearing. Not only does this reduce cost, but also complexity and production time. The total time consumed to produce a working robot (including the design stage and software adaptation) was only three months. Apart from screws, nuts, bearings, actuators, and electronics, the entirety of the robot is 3D printed using Selective Laser Sintering (SLS). To support the weight of the robot with inexpensive actuators, parallel kinematics (in the leg pitch joints) and gear transmissions (in the leg roll and yaw joints) have been used. In total, the robot uses 34 actuators for 18 joints in a similar kinematic structure as the NimbRo-OP2, which was a starting point for our work. With the NimbRo-OP2X, not only have we increased the capabilities, but the presented design improvements contributed to an estimate cost reduction of 17 %. Price for the NimbRo-OP2X parts is in a similar range to that of the popular stationary dual-arm research platform Baxter. A summary of the main technical specifications can be found in Table I.

A. Intelligent Actuators

Although the appearance of the robot may resemble that of the NimbRo-OP2 [8], it is in fact a complete redesign that incorporates several upgrades. This redesign was dictated by the choice to use a new series of intelligent actuatorsthe Robotis Dynamixel XM-540. Some of the more notable physical improvements include a 29% increase in output torque, metal casing which provides better heat dissipation and higher durability in combination with a redesigned gearbox. What truly separates the two robots in terms of actuation, are the new features provided. We have modified our open-source software and firmware to take full advantage of the new hardware. By redesigning our joint control strategy to a more abstract representation, it is now feasible to utilise the new possibilities, such as joint torque sensing and control. An added feature is the ability to effortlessly connect sensors onto the Dynamixel bus, without having to develop additional electronics or software. Up to three analog or digital devices can be connected to a single actuator and operated using the standard protocol. Paired with the 3D printed nature of the design, the platform can be adapted to many applications.

B. Parallel Computing

The current trend in computing is shifting towards machine learning techniques, which have been successfully applied to some of the hardest tasks in selected research areas with high efficiency. The most prominent examples include computer vision and control systems-both of which strongly relate to humanoid robotics. The NimbRo-OP2 already provided substantial computing power, however the limited space in the trunk did not allow for easy GPU integration without redesign. We have increased the available space and fitted a Mini-ITX form-factor mainboard with a standard GPU inserted into the available PCIe slot. Combining the newest Intel Core i7-8700T (6 cores with 12 threads) processor and the Nvidia GTX 1050 Ti (768 CUDA cores) GPU allowed us to achieve significant performance at a low TDP of 105 W. Power is supplied to the computer through a wide-input industrial 250 W M4-ATX on-board power supply. Future upgrades of the computing unit pose no problem and can be done at the component level, as every part is standardised. This improvement reflects our design goal of providing a highly capable and customisable platform.

C. Design Optimisations

Minimising the effort required to produce and assemble a large, customisable humanoid robot has been our goal ever since our collaboration with igus GmbH on their Humanoid Open Platform [16]. The NimbRo-OP2 was a considerable advancement in this regard, however there were some issues that were left unaddressed. Despite 3D printing the whole structure of the robot in PA12 nylon, external gearing required to support the robot weight in selected joints still had to be milled from brass. The production of these custom gears was costly and time consuming, which made it a bottleneck in the whole production process. With the advancement of 3D printing materials and technology, it was possible to design gears that are in every way superior to their brass counterparts.

We decided to use double helical gears over spur gears, due to their higher torque density and smoother operation. Subtractive manufacturing of such gears from metal would be very expensive, as they cannot be cut with simple gear hobs. This is not a limitation when 3D printing. When designing a helical gear, one needs to select a helix angle ψ , as it is not standardised. Typical values start at 15° and do not exceed 45° —the most common ones being 15, 23, 30, and 45° [17]. Higher values provide smoother operation at the cost of higher axial load. The helix angle also influences the gear diameter d—which in standard spur gears is the product of the module m and the number of teeth z. In helical gears, this equation is divided by $\cos \psi$. We decided to use a ψ of 41.4096° , with a module *m* of 1.5. With this combination, the diameter is always double the teeth number, as $\cos \psi =$ 0.75. Not only does this provide smooth operation, but allows to design the transmission to be compatible with spur gears in standard modules of 1 and 2. As the two halves of a double helical gear produce counteracting axial forces, the



Fig. 3. Architecture of our open-source ROS framework.

axial load factor is mitigated. The gears have been printed using SLS technology from the igus I6 material, which was designed specifically for gears. Gears made from it have a low friction coefficient and do not require lubrication. Thickness has been selected to be 14 mm (7 mm for each helical half) to increase tooth overlap. Despite increasing the thickness almost threefold (from 5 mm in the NimbRo-OP2), the plastic gears are on average only 54% of the brass gear weight. The achieved end-result is an elegant, smooth, light and inexpensive transmission that is fast to produce. An example of the gear used in the hip joint can be seen in Fig. 2.

Optimising the design was one of the greater challenges with this platform. New features provided by the actuators and computing unit came at a cost of increased weight. The additional 1.4 kg meant a hypothetical increase of 8% in robot weight. We were able to redesign the leg parts to be lighter by almost 0.5 kg, while making them structurally stronger at the same time. NimbRo-OP2 utilised quite simple shapes and flat surfaces. This decreased weight, but made the legs prone to slight bending. Legs in the NimbRo-OP2X are rounder, slightly narrower and have dedicated cable pathways—all of which contribute to the overall durability and torsional resistance of the structure. Offsetting the weight increase enabled us to take full advantage of the torque benefits provided by the new actuators.

IV. SOFTWARE DESIGN

Capable hardware is essential for reliable performance, although it is the software that truly defines the capabilities of the robot. The software of the NimbRo-OP2X is heavily based on ROS [18]. It is the result of our continued efforts aimed at providing the research community with an easy to use, and modular framework. Five years since its creation in 2013 [19], it still continues to expand on functionality. Some of the more notable features include: a Gazebo simulator, complete vision pipeline for robot soccer, tunable gait as well as a motion designer and player. A simplified schematic of the framework can be seen in Fig. 3. The source code is accessible online [20]. With the NimbRo-OP2X, we focused



Fig. 4. Visual perception network architecture. Like original ResNet architecture, each convBlock is consist of two convolutional layers followed by batch-norm and ReLU activation layers. Note that residual connections in ResNet are not depicted here.

on developing robust vision software with deep learning methods and optimising the tunable gait parameters. Previously developed features along with an in-depth description of the framework's inner workings can be found in [21], [8].

A. Visual Perception

NimbRo-OP2X perceives the environment using a Logitech C905 camera which is equipped with a wide-angle lens and an infrared cut-off filter. With the introduced ability to utilise parallel computing, we are able to supersede our previous approach to vision [22], by using a deep neural network in conjunction with some post-processing. We highlight the parallel computing capabilities of the new platform in the context of our demonstration scenario of robot soccer. The presented vision system can work with different viewing angles, brightnesses, and lens distortions.

Our visual perception architecture is illustrated in Fig. 4. We follow the encoder-decoder approach of recently proposed segmentation models like SegNet [23], V-Net [24], and U-Net [25]. To achieve real-time computation and to minimize the number of parameters in the network, we designed the decoder part to be shorter than the encoder part. This design choice has the drawback that the output of the network has less spatial information than the full resolution

input. We mitigate this effect by using a subpixel centroidfinding method in the post-processing step. In order to minimize the training time, we lower the number of samples in the dataset, and leverage big datasets like ImageNet, we used a pre-trained ResNet-18 model for our encoder part. As shown in Fig. 4, we removed the last fully connected and GAP layers in the ResNet-18 model. Similar to the U-Net model, we also have lateral connections between encoder and decoder parts to directly provide high-resolution details to the decoder. In total, we have 23 convolutional layers and four convolution-transpose layers.

The network is trained on the following object classes: ball, goal post, and robots. Instead of using the full segmentation loss, similar to SweatyNet [26], we use a Gaussian blob around the center of the ball and the bottom-middle point of the goal post and robots. As center point detection is enough for objects mentioned earlier, instead of full semantic segmentation loss, mean squared error is used as the loss function. To find a good learning rate, we followed the approach presented by Smith [27]. We have also tested the recently proposed AMSGrad optimizer [28], but did not see any benefit in our application, over the employed Adam. To train the network as fast as possible, we first downsample the training images and train for 50 epochs while freezing the weights on the encoder part. Freezing the encoder part during the first epochs quickly adjusts the randomly initialized weights in the decoder part, that is inspired by the work of Brock et al. [29] and by the paper from Yosinski et al. [30]. During the next 50 epochs, all parts of the models rae trained. Finally, full-sized images are used to train the model for 50 more epochs. For the pre-trained model, a lower learning rate is used. In total, the whole training process with around 3000 samples takes less than 40 minutes on a single Titan Black GPU with 6 GB memory. Some of the training samples were taken from the ImageTagger library [31] which have annotated samples from different angles, cameras, and brightness. Although the network produced very few falsepositives (around 1% for the ball), we were able to reduce this value by utilising inference time augmentation. Two samples of the testset are depicted in Fig. 5. Extracting object coordinates is done with minor post-processing, since the output of the network is blob-shaped. After thresholding each output channel, we apply morphological erosion and dilation to eliminate negligible responses. Finally, we compute the object center coordinates from the achieved contours.

After detecting game-related objects, we filter them and project each object location into egocentric world coordinates. These coordinates then are further processed in the behaviour node of our ROS-based open-source software for decision making. To alleviate projection errors due to the slight differences between the CAD model and real hardware, we calibrate the camera frame offsets, using the Nelder-Mead [32] Simplex method.

Through the improvements made on our vision system, we were able to perceive a FIFA size 5 ball up to 7 m with an accuracy of 99% and less than 1% of false detections. Goal posts can be detected up to 8 meters with 98% precision



Fig. 5. Object detection. Left column: A captured image from the robot in the test set. Middle column: The output of the network with balls (cyan), goal posts (magenta), and robots (yellow) annotated. Right column: Ground truth.

and with 3% false detection rate. Robots are detected up to 7 meters with a success rate of 90% and a false detection rate of 8%. For line and field detections, we are using non-deep learning approaches [22], but are working on adding two more channels to the network output and use a single network for all detections. The complete vision pipeline including a forward-pass of the network takes approximately 20 ms on the robots hardware.

B. Sample-efficient Gait Optimisation

Perceiving the environment is only the first step to interacting with it. An essential skill for humanoid robots is walking. The gait of the NimbRo-OP2X is based on an open-loop pattern generator that calculates joint states based on a gait phase angle, whose rate is proportional to the desired step frequency [33]. The phase angle is responsible for generating arm and leg movements such as lifting and swinging. We have built around this approach and incorporated corrective actions based on fused angle feedback [34], [35], [36].

One of the drawbacks of this method, is having to tune the gait for each and every robot, as they can differ not only in structure but have slight mechanical differences between theoretically identical robots. We remedy this by using optimisation to find proper values for the Fused Feedback controller used by the gait. To minimise hardware wear-off, this optimisation does not only take place in the real world, but highly exploits information gained through the included Gazebo simulator. This approach has been previously applied by Rodriguez et al. [11] and is now utilised on the NimbRo-OP2X robot.

Since we pursue the goal of optimising a given parameter set in a sample-efficient manner, the proposed method uses *Bayesian Optimisation*. This allows to efficiently trade off exploration and exploitation. It requires to specify a Kernel function K and a parametrisation of a Gaussian Process (GP). Both of these are responsible for encoding essential information about the optimisation environment. We use a kernel that is comprised of two terms. The first one relates to the simulation performance k_{sim} , while the other is an error term k_{ϵ} depicting the difference between simulations and the real world. This results in:

$$k(\mathbf{a}_{\mathbf{i}}, \mathbf{a}_{\mathbf{j}}) = k_{sim}(\mathbf{x}_{\mathbf{i}}, \mathbf{x}_{\mathbf{j}}) + k_{\delta}(\delta_{i}, \delta_{j})k_{\epsilon}(\mathbf{x}_{\mathbf{i}}, \mathbf{x}_{\mathbf{j}})$$
(1)

for an augmented parameter vector $\mathbf{a_i} = (\mathbf{x_i}, \delta_i)$, where δ signals whether the evaluation has been performed in simulation or on the real system. An indication if both experiments were in the real world is k_{δ} , which returns a higher correlation in that case. This way, it is possible to address potentially complex, non-linear transformations between the simulated sequences and the real-world performance. To perform the optimisation process in a data-efficient manner, the approach utilises *Entropy* as a way of measuring the information content. The next evaluation point is chosen with respect to the maximal achievable change in Entropy [37], biased by weight factors trading off simulation and real-robot experiments (see Fig. 6).

For both kernels, we select the Rational-Quadratic Kernel, since it has proven to be suitable in our previous work [11].

The proposed cost function is based on the proportional Fused Feedback e_P , which employs the Fused Angle deviations and thus can be considered a stability measure [11]. Furthermore, to reduce the impact of noise, we examine the cost of the sagittal(α) and lateral(β) planes separately. Since higher gain values lead to higher torques, we also penalise the magnitude of the parameter using a logistic function ν . This results in the cost functions:

$$J_{\alpha}(\mathbf{x}) = \int_0^T \|e_{P\alpha}(\mathbf{x})\|_1 dt + \nu(\mathbf{x})$$
(2)



Fig. 6. At each iteration the algorithm trades off performing a simulation or a real-world experiment. This decision is led by the expected change of Entropy, making the algorithm particularly sample-efficient.

$$J_{\beta}(\mathbf{x}) = \int_0^T \|e_{P\beta}(\mathbf{x})\|_1 dt + \nu(\mathbf{x}), \qquad (3)$$

where e_P is seen as a function of the used gait parameters **x**.

An evaluation of the cost function is done through a complex test sequence consisting of the same, pre-defined combination of forward, backward and sideways gait commands. To receive more reproducible results in simulations, we average the cost of N = 4 simulated runs.

In the described example, we optimise the Arm Angle corrective actions in the sagittal direction. More precisely, we restrict our search to the P and D gains. The method can be similarly applied to other controllers. To preserve the robot's hardware, we restrict the permitted number of real-world evaluations to 15, which is comparable to the time it takes an expert to hand-tune the parameters. These



Fig. 7. The GP posterior of the Arm Angle optimisation is depicted. The blue dots resemble evaluations in simulation, while real experiments are shown in red. A black dot indicates the estimated minimum. The mean cost estimate is displayed (teal mesh), together with its standard deviation (grey mesh). To increase visibility, the upper standard deviation is not shown.

15 runs were accomplished after a total of 161 iterations, thus including 146 simulated gait sequence evaluations. To assess the resulting parameters, the deviations of the new parameters were compared to the old parameters in 5 evaluation sequences each. The results of the optimisation not only produces a qualitatively more convincing walk, but also leads to a reduction of the Fused Angle deviation by approximately 18 %.

The GP posterior of the optimisation is depicted in Fig. 7. Although it might seem that the simulations only had a small influence on the result, they were of utmost importance in early iterations [11]. This is underlined especially by the fact, that the robot did not fall a single time during realworld evaluations, which proves that the GP model is able to utilise simulation data to rule out inferior parameters. This stems from the high costs associated with unstable simulated runs.

Altogether, our approach has proven to be beneficial with the NimbRo-OP2X, and yielded successful parameterisations for the used gait. Thus, it allows for easy operation even for an inexperienced user. It can be a feasible alternative to the time-consuming hand-tuning performed by an expert, which wears off the robot's hardware.

V. EVALUATION

An evaluation of the NimbRo-OP2X was performed during the RoboCup 2018 competition in Montréal, Canada. In the RoboCup Humanoid League AdultSize class, robots autonomously compete in one vs. one soccer games, two vs. two drop-in games, and four technical challenges that test specific abilities in isolation. The soccer games are performed on a 6×9 m artificial grass field, which makes walking challenging. Perceiving the environment and localisation are also challenging, due to the ever changing lighting conditions. In the competition, the NimbRo-OP2X performed outstandingly by winning all out of the four possible awards, including the Best Humanoid Award.

In the main tournament, our robot played a total of six games, including the quarter-finals, semi-finals and finals (see Fig. 10). An additional five drop-in games were played, where two vs. two mixed teams were formed and collaborated during the game. This resulted in 220 minutes of official play time, during which 66 goals were scored and only 5 were conceded.

The optimised gait was not only fast (approx. 0.5 m/s) but also robust, as NimbRo-OP2X never fell during walking in free space or while dribbling the ball. Very strong collisions with much heavier robots made the robot fall over four times.

One of the biggest contributions to the success, was our vision system utilising deep learning. The ball detections were robust up to a distance of 7 m, where the perceived ball was only a few pixels in size. Robot detections allowed for efficient fights for the ball, which often resulted in the ball being dribbled around the opponent. This was a decisive factor during the final, where the NimbRo-OP2X won 2:0 against Sweaty from Hochschule Offenburg, Germany.



Fig. 8. NimbRo-OP2X withstanding a push from the front. The weight is marked in red. After receiving the push, the robot performs corrective actions to recover stability.

The opponent played very defensively and was efficient at dribbling and handling the ball.

Performance of the vision and gait was also tested during the technical challenges. In the Push Recovery challenge, the goal is to withstand pushes from the back and from the front. The disturbance is applied by a 3 kg heavy pendulum swinging on a 1.5 m long string. The weight hits the robot at the height of the center of mass. In order to perform the push, the weight is retracted from the robot by a distance d and then released freely. NimbRo-OP2X was able to successfully withstand both frontal and backward pushes with $d = 0.8 \,\mathrm{m}$. The optimised arm swing corrective action helped substantially-its effect can be seen in Fig. 8 and Fig. 9. A different challenge required the robot to score a goal from a moving ball which is passed by a human player from the corner of the field. Our fast image processing and capable hardware allowed the robot to score a goal in approximately 3 seconds, counting from when the ball was passed.

VI. CONCLUSIONS

In this paper, we have described the newly developed hardware and software aspects of the NimbRo-OP2X. The technical specifications and provided features make it a powerful middle-ground for the research community between smaller, inexpensive robots and larger cost-prohibitive platforms. Not only is NimbRo-OP2X large enough to interact with a human environment, its relatively light weight makes it easy and safe to operate. The minimalistic 3D printed design allows for easy maintenance as well as user modifications and improvements. By demonstrating efficient use



Fig. 9. Push Recovery trial at RoboCup 2018. Optimised corrective arm action in response to expected angle deviation in the sagittal plane.

of its hardware and new deep-learning capabilities, the robot was able to outperform its competition at RoboCup 2018. We plan on releasing the 3D-printable hardware components with a Bill of Materials, as well as updating our open-source software [20] by October 2018.

VII. ACKNOWLEDGEMENTS

This work was partially funded by grant BE 2556/13 of the German Research Foundation (DFG).

REFERENCES

- [1] I. Kato, "Development of wabot 1," *Biomechanism*, vol. 2, pp. 173–214, 1973.
- [2] H. Hirukawa, F. Kanehiro, K. Kaneko, S. Kajita, K. Fujiwara, Y. Kawai, F. Tomita, S. Hirai, K. Tanie, T. Isozumi *et al.*, "Humanoid robotics platforms developed in HRP," *Robotics and Autonomous Systems*, vol. 48, no. 4, pp. 165–175, 2004.
- [3] Y. Sakagami, R. Watanabe, C. Aoyama, S. Matsunaga, N. Higaki, and K. Fujimura, "The intelligent Asimo: System overview and integration," in *Intelligent Robots and Systems (IROS), IEEE/RSJ International Conference on*, vol. 3, 2002, pp. 2478–2483.
- [4] G. Nelson, A. Saunders, N. Neville, B. Swilling, J. Bondaryk, D. Billings, C. Lee, R. Playter, and M. Raibert, "Petman: A humanoid robot for testing chemical protective clothing," *Journal of the Robotics Society of Japan*, vol. 30, no. 4, pp. 372–377, 2012.
- [5] J. Englsberger, A. Werner, C. Ott, B. Henze, M. A. Roa, G. Garofalo, R. Burger, A. Beyer, O. Eiberger, K. Schmid *et al.*, "Overview of the torque-controlled humanoid robot TORO," in *Humanoid Robots* (*Humanoids*), 14th IEEE-RAS International Conference on, 2014, pp. 916–923.



Fig. 10. NimbRo-OP2X during RoboCup 2018. Left: Performing a kick. Right: Competing for the ball.

- [6] N. A. Radford, P. Strawser, K. Hambuchen, J. S. Mehling, W. K. Verdeyen, A. S. Donnan, J. Holley, J. Sanchez, V. Nguyen, L. Bridg-water *et al.*, "Valkyrie: Nasa's first bipedal humanoid robot," *Journal of Field Robotics*, vol. 32, no. 3, pp. 397–419, 2015.
- [7] Toyota Global Newsroom. (2017, Nov) Toyota Unveils Third Generation Humanoid Robot T-HR3. [Online]. Available: https: //newsroom.toyota.co.jp/en/download/20110424
- [8] G. Ficht, P. Allgeuer, H. Farazi, and S. Behnke, "NimbRo-OP2: Grown-up 3D Printed Open Humanoid Platform for Research," in 17th IEEE-RAS Int. Conference on Humanoid Robots (Humanoids), 2017.
- [9] G. Ficht, D. Pavlichenko, P. Allgeuer, H. Farazi, D. Rodriguez, A. Brandenburger, J. Kuersch, and S. Behnke, "Grown-up NimbRo Robots Winning RoboCup 2017 Humanoid AdultSize Soccer Competitions," *RoboCup 2017: Robot World Cup XXI, Lecture Notes in Computer Science*, 2018.
- [10] A. Farchy, S. Barrett, P. MacAlpine, and P. Stone, "Humanoid robots learning to walk faster: From the real world to simulation and back," in *Int. Conf. on Autonomous Agents and Multi-agent Systems (AAMAS)*, 2013.
- [11] D. Rodriguez, A. Brandenburger, and S. Behnke, "Combining Simulations and Real-robot Experiments for Bayesian Optimization of Bipedal Gait Stabilization," *RoboCup International Symposium, Montreal, Canada*, 2018.
- [12] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath *et al.*, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal Processing Magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [13] C. Finn and S. Levine, "Deep visual foresight for planning robot motion," in *Robotics and Automation (ICRA), IEEE International Conference on*, 2017, pp. 2786–2793.
- [14] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, p. 529, 2015.
- [15] E. Romera, J. M. Alvarez, L. M. Bergasa, and R. Arroyo, "Erfnet: Efficient residual factorized convnet for real-time semantic segmentation," *IEEE Transactions on Intelligent Transportation Systems*, vol. 19, no. 1, pp. 263–272, 2018.
- [16] P. Allgeuer, H. Farazi, M. Schreiber, and S. Behnke, "Child-sized 3D Printed igus Humanoid Open Platform," in 15th IEEE-RAS Int. Conference on Humanoid Robots (Humanoids), 2015.
- [17] K. Gopinath and M. Mayuram, "Lecture notes in machine design," 2010.
- [18] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Ng, "ROS: An open-source robot operating system," in *ICRA Workshop on Open Source Software*, 2009.
- [19] M. Schwarz, J. Pastrana, P. Allgeuer, M. Schreiber, S. Schueller, M. Missura, and S. Behnke, "Humanoid TeenSize Open Platform NimbRo-OP," in 17th RoboCup International Symposium, 2013.
- [20] NimbRo. (2015, Oct) igus Humanoid Open Platform ROS Software. [Online]. Available: https://github.com/AIS-Bonn/humanoid_op_ros
- [21] P. Allgeuer, M. Schwarz, J. Pastrana, S. Schueller, M. Missura, and S. Behnke, "A ROS-based software framework for the NimbRo-OP humanoid open platform," in 8th Workshop on Humanoid Soccer Robots, Int. Conference on Humanoid Robots, 2013.

- [22] H. Farazi, P. Allgeuer, and S. Behnke, "A monocular vision system for playing soccer in low color information environments," in 10th Workshop on Humanoid Soccer Robots, IEEE-RAS Int. Conference on Humanoid Robots, 2015.
- [23] V. Badrinarayanan, A. Kendall, and R. Cipolla, "SegNet: A deep convolutional encoder-decoder architecture for image segmentation," arXiv:1511.00561, 2015.
- [24] F. Milletari, N. Navab, and S.-A. Ahmadi, "V-Net: Fully convolutional neural networks for volumetric medical image segmentation," in *3D Vision (3DV), 2016 Fourth International Conference on.* IEEE, 2016, pp. 565–571.
- [25] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [26] F. Schnekenburger, M. Scharffenberg, M. Wülker, U. Hochberg, and K. Dorer, "Detection and localization of features on a soccer field with feedforward fully convolutional neural networks (fcnn) for the Adultsize humanoid robot Sweaty," in *Proceedings of the 12th Workshop* on Humanoid Soccer Robots, IEEE-RAS International Conference on Humanoid Robots, Birmingham, 2017.
- [27] L. N. Smith, "Cyclical learning rates for training neural networks," in Applications of Computer Vision (WACV), IEEE Winter Conference on, 2017, pp. 464–472.
- [28] S. J. Reddi, S. Kale, and S. Kumar, "On the convergence of Adam and beyond," in *International Conference on Learning Representations* (*ICLR*), 2018. [Online]. Available: https://openreview.net/forum?id= ryQu7f-RZ
- [29] A. Brock, T. Lim, J. M. Ritchie, and N. Weston, "Freezeout: Accelerate training by progressively freezing layers," *arXiv preprint* arXiv:1706.04983, 2017.
- [30] J. Yosinski, J. Clune, Y. Bengio, and H. Lipson, "How transferable are features in deep neural networks?" in Advances in neural information processing systems, 2014, pp. 3320–3328.
- [31] N. Fiedler, M. Bestmann, and N. Hendrich, "Imagetagger: An open source online platform for collaborative image labeling," in *RoboCup* 2018: Robot World Cup XXII. Springer, 2018.
- [32] J. A. Nelder and R. Mead, "A simplex method for function minimization," *The Computer Journal*, vol. 7, no. 4, pp. 308–313, 1965.
- [33] S. Behnke, "Online trajectory generation for omnidirectional biped walking," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2006.
- [34] M. Missura and S. Behnke, "Self-stable omnidirectional walking with compliant joints," in 8th Workshop on Humanoid Soccer Robots, Humanoids, 2013.
- [35] P. Allgeuer and S. Behnke, "Fused Angles: A representation of body orientation for balance," in *Int. Conf. on Intelligent Robots and Systems* (*IROS*), 2015.
- [36] —, "Omnidirectional bipedal walking with direct fused angle feedback mechanisms," in 16th IEEE-RAS Int. Conference on Humanoid Robots (Humanoids), 2016.
- [37] P. Hennig and C. Schuler, "Entropy search for information-efficient global optimization," *Journal of Machine Learning Research*, vol. 13, no. Jun, pp. 1809–1837, 2012.