



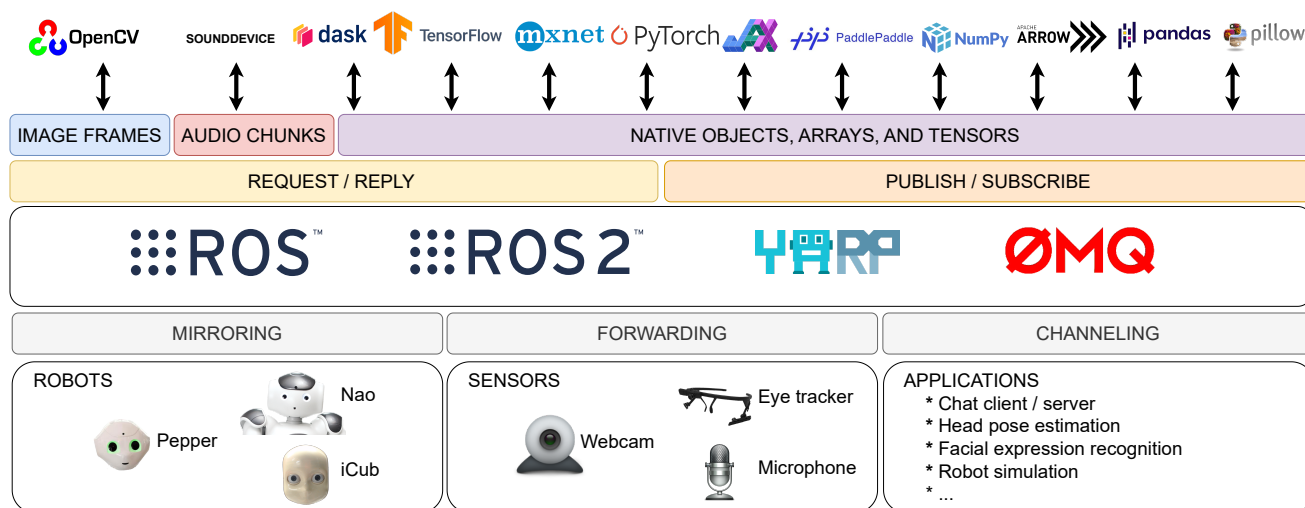
# Wrapyfi: A Python Wrapper for Integrating Robots, Sensors, and Applications across Multiple Middleware

Fares Abawi  
fares.abawi@uni-hamburg.de  
University of Hamburg  
Hamburg, Germany

Di Fu  
di.fu@uni-hamburg.de  
University of Hamburg  
Hamburg, Germany

Philipp Allgeuer  
philipp.allgeuer@uni-hamburg.de  
University of Hamburg  
Hamburg, Germany

Stefan Wermter  
stefan.wermter@uni-hamburg.de  
University of Hamburg  
Hamburg, Germany



**Figure 1: Overview of the Wrapyfi framework. From top to bottom: 1) Data types are encoded or decoded depending on the transmission mode; 2) Encoded objects are prepared for transmission using the Request/Reply or Publish/Subscribe communication pattern; 3) Messages are transmitted through the selected middleware protocol; 4) Messages sequenced according to the communication scheme; 5) Messages exchanged between robots, applications, and sensors.\***

\* The "nine dots" ROS and ROS 2 logos are trademarks of Open Source Robotics Foundation. TensorFlow, the TensorFlow logo, and any related marks are trademarks of Google Inc. The OpenCV logo is a trademark of <https://opencv.org>. The NumPy logo is used in accordance with the NumPy logo guidelines. The pandas logo is used in accordance with the brand and logo guidelines. PyTorch, the PyTorch logo and any related marks are trademarks of The Linux Foundation. The name ZeroMQ and the "OMQ" logo are used in compliance with creative commons license Attribution-NoDerivatives 4.0 International (CC BY-ND 4.0). The logos for Dask, Apache MXNet, paddlepaddle, PIL (Pillow), JAX, and YARP are included with respect to their trademark policies; we acknowledge that these are subject to copyrights, trademarks, or registered trademarks of their respective holders. We do not claim ownership of these copyrights or trademarks. The use of these logos does not indicate endorsement by the trademark or copyright holders, nor does it suggest any affiliation or endorsement by the authors of this work.

## ABSTRACT

Message oriented and robotics middleware play an important role in facilitating robot control, abstracting complex functionality, and unifying communication patterns between sensors and devices. However, using multiple middleware frameworks presents a challenge in integrating different robots within a single system. To address this challenge, we present Wrapyfi, a Python wrapper supporting multiple message oriented and robotics middleware, including ZeroMQ, YARP, ROS, and ROS 2. Wrapyfi also provides plugins for exchanging deep learning framework data, without

additional encoding or preprocessing steps. Using Wrapyfi eases the development of scripts that run on multiple machines, thereby enabling cross-platform communication and workload distribution. We finally present the three communication schemes that form the cornerstone of Wrapyfi's communication model, along with examples that demonstrate their applicability.

<http://software.knowledge-technology.info#wrapyfi>.

## CCS CONCEPTS

• **Software and its engineering** → **Message oriented middleware**; • **Computer systems organization** → *External interfaces for robotics*; • **Theory of computation** → Distributed computing models.

## KEYWORDS

Message oriented middleware, robotics middleware, distributed computing, deep learning frameworks, human-robot interaction



This work is licensed under a Creative Commons Attribution International 4.0 License.

HRI '24, March 11–14, 2024, Boulder, CO, USA  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-0322-5/24/03.  
<https://doi.org/10.1145/3610977.3634943>

**ACM Reference Format:**

Fares Abawi, Philipp Allgeuer, Di Fu, and Stefan Wermter. 2024. Wrapyfi: A Python Wrapper for Integrating Robots, Sensors, and Applications across Multiple Middleware. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24)*, March 11–14, 2024, Boulder, CO, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3610977.3634943>

**1 INTRODUCTION**

Real-time robotic applications require exchanging multimodal data arriving from a variety of sensors. A framework that distributes sensory information across processes is necessary, especially for robot-robot and human-robot interaction [17]. Multiprocess and multithread instances are used to parallelize independent methods. However, such parallelization approaches are limited to single machines and may not be sufficient for applications with a large number of sensors or computationally expensive processing methods. Eventually, this leads to performance bottlenecks on consumer-grade computers. Message oriented and robotics middleware, such as ZeroMQ [10], YARP [16], ROS [21], and ROS 2 [13], were developed to tackle such challenges. Middleware frameworks use communication protocols to exchange data and distribute operations across several machines and nodes [7].

ROS [21] is a middleware commonly used in the robotics community. ROS provides control hardware interfaces, visualization tools, and communication models for many robotic platforms [2]. Its widespread use is a direct result of its early adoption of open source and the vast amount of robotic tools provided by its developers and contributors. However, ROS is scheduled for deprecation in favor of ROS 2 [13]. Many robotic platforms and packages, nonetheless, have not been updated to support this transition yet. Although bridges were developed to enable communication between ROS, ROS 2, and WebSocket, integrating such bridges into working pipelines requires major modifications to the underlying code and its structure. This demands following certain naming conventions and limiting the message types supported, resulting in additional effort. Other middleware designed specifically for certain robotic platforms such as YARP [16] used by the iCub [15] robot, provide interfaces for communicating with ROS [18] as well. However, their usage dictates modifying scripts to accommodate specific message types. This poses a major hurdle for developers wanting to integrate different robots and middleware, as a result, restricting the cross-compatibility of their applications with existing systems.

To improve interoperability between different robotic platforms and reduce reliance on a particular middleware, we have developed the open source Wrapyfi<sup>1</sup> (illustrated in Figure 1) framework, a Python wrapper supporting multiple middleware bindings. Wrapyfi is a simpler alternative to GenoM3 [14]. GenoM3 adopts a model-driven approach and uses templates to define the components and data exchanges across middleware. Since it is specifically developed for Python scripting, Wrapyfi eliminates the need for having to learn another language or to define templates, unlike GenoM3. REMS [26] is a middleware built in Python with simplistic interfaces for educational purposes. Although REMS supports a large set of robots and simulation environments, it does not address interoperability between different middleware operating on them.

<sup>1</sup><https://github.com/fabawi/wrapyfi>

Wrapyfi's decorator-based design integrates easily with existing workflows, prioritizing minimal modifications for improved multi-robot communication. Beyond robotic applications, its adaptability is observed in supporting message oriented middleware, facilitating communication with interfaces that do not necessarily require the additional packages and tools provided by robotics middleware. Deep learning frameworks like JAX [4] and PyTorch [20], support multi-machine parallelization mainly through remote procedure calls. The approaches adopted in distributing models and data differ greatly, including the communication patterns used and the orchestration of communication, having either a single or several controllers. By offering a standard approach for multiple frameworks, and supporting two of the most common communication patterns, namely publish-subscribe and request-reply—also known as the request-response or client-server pattern—Wrapyfi offers greater control over communication dynamics in comparison to each framework's parallelization protocol.

Open Neural Network Exchange (ONNX) [3] is a framework designed to standardize machine learning model representations, offering compatibility with a wide range of deep learning frameworks. However, using ONNX with any framework requires converting the model formats. In contrast, Wrapyfi does not impose such a constraint or bind developers to a specific protocol. Wrapyfi does not only allow for native Python object exchanges but also transports data structures such as arrays and tensors, which are relied upon in deep learning applications. This integration makes Wrapyfi a useful tool for developers, allowing them to take advantage of both robotics and deep learning ecosystems.

**2 DATA TYPES**

Wrapyfi employs a type-aware serialization method that automatically transforms the objects exchanged between script mirrors into a format compatible with the selected middleware. Wrapyfi supports the following data types:

**Native objects, arrays, and tensors.** Wrapyfi allows for the transmission of a variety of data types used in Python. Prior to transmission, these data types are converted into JSON strings to ensure compatibility across different middleware platforms. Wrapyfi supports using NumPy [8] arrays and enables their sharing across mirrored scripts. Moreover, Wrapyfi offers a plugin interface that developers may use to customize the transmission of other types of objects. This feature allows encoding objects as strings, which can eventually be decoded back into their original structure. Wrapyfi comes with built-in plugins for exchanging Arrow [22] vectors, pandas<sup>2</sup> [19] data frames, and Pillow<sup>3</sup> images. It also supports tensors from major deep learning frameworks such as TensorFlow [1], PyTorch [20], MXNet [6], JAX [4], PaddlePaddle [12], and Dask [23]. These plugins make it possible to exchange data between different frameworks and to integrate deep learning models into robotic systems. When specified, the tensors transmitted using Wrapyfi can be mapped to GPUs or CPUs different from the ones specified on a publishing script's end, allowing for the distribution of computationally demanding deep learning models.

<sup>2</sup>pandas version 1 with NumPy as a backend

<sup>3</sup><https://github.com/python-pillow/Pillow>

**Images.** ROS, ROS 2, and YARP provide specialized message types for transmitting images. We use image messages to stream raw monochrome, RGB, and JPEG-encoded images. ZeroMQ does not provide such specialized message structures. Therefore, we make use of the multipart message structure to create an image interface, allowing us to standardize middleware behavior and transmit the image properties to a specified topic.

**Audio chunks.** ROS and ROS 2 do not provide messages structured for audio transmission, so we create custom messages and services to transmit audio along with its properties. The number of audio channels transmitted can vary in size, as long as the audio chunk structure follows the python-sounddevice format<sup>4</sup>. For YARP, we use the existing sound port and transmit the audio as a sequence. Whereas, for ZeroMQ, we transmit a string, encoding the auditory signal along with its properties as a single multipart message.

### 3 COMMUNICATION SCHEMES

Wrapyfi manages script interactions using three communication schemes—Mirroring, Forwarding, and Channeling. Mirroring enables concurrent execution of multiple scripts with synchronized actions. Forwarding creates chains of methods to tunnel arguments and return values across different middleware configurations. Channeling allows for the broadcasting of multiple return values via one method, each using potentially different middleware. Each scheme addresses different challenges in distributed systems.

**Listing 1: Decorated method registering the data type, middleware, topic, connection protocol, and blocking behavior. 'so' passes the first argument (mware) from the method to the decorator. Similarly, 'blocking' passes the keyword argument.**

```
1 class MirrorCls(MiddlewareCommunicator):
2     @MiddlewareCommunicator.register('NativeObject',
3         '$0', 'MirrorCls', '/example/read_msg',
4         carrier='tcp', should_wait='blocking')
5     def read_msg(self, mware, msg='', blocking=True):
6         msg_ip = input('type message:')
7         obj = {'msg': msg, 'msg_ip': msg_ip}
8         return obj,
```

The MiddlewareCommunicator is a Wrapyfi class for establishing communication methods. It implements the register decorator for setting the middleware types, topics, and various communication parameters. Each method set to publish, subscribe, request, or reply should be encapsulated with this decorator. Listing 1 illustrates the use of the register decorator to register a method for YARP middleware communication, specifying object type, middleware, name of the class, YARP port (topic), communication protocol, and whether the method should await a response, which results in blocking the subscribing method until the publisher transmits a message. The read\_msg method obtains user input from one process, allowing all other subscribing processes to acquire user input from a single process invocation.

<sup>4</sup><https://github.com/spatialaudio/python-sounddevice>

**Listing 2: Activating a method in 'publish' mode. When the method is called, its results are returned to the caller and transmitted to the listener.**

```
1 mirror = MirrorCls()
2 mirror.activate_communication(
3     'read_msg', mode='publish')
```

In Listing 2, setting the mode to 'publish' triggers read\_msg upon method call, whereas 'listen' returns the message received over the middleware. These modes enable the establishment of communication following the publish/subscribe pattern. Alternatively, setting the activate\_communication mode to 'request' or 'reply' triggers the request/reply pattern.

**Mirroring.** Mirrors are multiple scripts running concurrently. The scripts share arguments and return values using a predefined communication pattern. The behaviors of all mirrored scripts are identical. However, their methods could either execute functionality in place or acquire their return values from another publisher. By calling read\_msg in Listing 1 using a single publishing script, all subscribing mirrors receive the same return object when invoked as well. Regardless of the communication pattern or blocking behavior, all scripts follow the same pipeline with similar method returns.

**Forwarding.** The forwarding scheme in Wrapyfi enables passing arguments to multiple methods, each with a different middleware setting. This forms a chain of methods, transferring arguments and return values across middleware and topics. Forwarding employs multiple scripts with unique functions, connected by register decorators, making it suitable for creating multi-step processes with several scripts having partial component support. In Listing 3, we demonstrate data transmission between a system without ZeroMQ support and another without Yarp support, using an intermediary system that supports both. The first system dispatches the message using Yarp by invoking send\_yarp. The intermediary system then forwards it using ZeroMQ to send\_zmq. The final system, with Yarp disabled, receives the message via ZeroMQ by listening to send\_zmq. This scheme is needed when strict specifications are required regarding compatibility of software and middleware between systems, as in the case of robots.

**Listing 3: Demonstration of forwarding with two methods each using a different middleware.**

```
1 class ForwardCls(MiddlewareCommunicator):
2     @MiddlewareCommunicator.register('NativeObject',
3         'yarp', 'ForwardCls', '/example/native_yarp_msg',
4         carrier='mcast', should_wait=True)
5     def send_yarp(self, msg):
6         return msg,
7
8     @MiddlewareCommunicator.register('NativeObject',
9         'zeromq', 'ForwardCls', '/example/native_zmq_msg',
10        carrier='tcp')
11    def send_zmq(self, msg):
12        return msg,
```

**Channeling.** In the channeling scheme, Wrapyfi enables broadcasting to multiple middleware by encapsulating a method with

numerous decorators, each corresponding to a return value with its own data type and middleware. This is illustrated in Listing 4, where a method transmits three different data types over varied middleware, such as a Yarp native object message comprising a NumPy image and an audio chunk, a ROS image (OpenCV [5] compatible), and a ZeroMQ audio chunk. This scheme supports the simultaneous reception of different data types. If the environment lacks support for a specified middleware, a None type object is returned. Channeling is especially useful for handling multiple sensory inputs from different sources, allowing selective acquisition and disregard of unnecessary sensory input. This provides a balanced approach between mirroring and forwarding, altering the pipeline based on the returns received from the supported middleware.

**Listing 4: Demonstration of Channeling with one method reading multiple returns of different data types through multiple middleware.**

```


1 class ChannelCls(MiddlewareCommunicator):
2     @MiddlewareCommunicator.register('NativeObject',
3     'yarp', 'ChannelCls', '/example/native_yarp_msg',
4     carrier='mcast', should_wait=True)
5     @MiddlewareCommunicator.register('Image',
6     'ros', 'ChannelCls', '/example/image_ros_msg',
7     carrier='tcp', width='$img_width',
8     height='$img_height', rgb=True, queue_size=10)
9     @MiddlewareCommunicator.register('AudioChunk',
10    'zeromq', 'ChannelCls', '/example/audio_zmq_msg',
11    carrier='tcp', rate='$aud_rate',
12    chunk='$aud_chunk', channels='$aud_chann')
13    def read_mulret_middleware(self,
14    img_width=200, img_height=200,
15    aud_rate=44100, aud_chunk=8820, aud_chann=1):
16        ros_img = np.random.randint(256,
17        size=(img_height, img_width, 3), dtype=np.uint8)
18        zeromq_aud = (np.random.uniform(-1,1, aud_chunk),
19        aud_rate,)
20        yarp_native = [ros_img, zeromq_aud]
21        return yarp_native, ros_img, zeromq_aud

```


## 4 USE CASES

**Facial expression imitation.** Participants exhibit eight facial expressions while sitting in front of two robots, Pepper [25] and iCub [15] as depicted in Figure 2. The robots then imitate the participants' expressions. Pepper represents emotions through color changes, while iCub displays robotic facial expressions. The forwarding scheme in Wrapyfi tunnels interactions between the different system components and middleware configurations, enabling the exchange of visual and facial expression data between the robots and the recognition model [24]. Forwarding manages image acquisition across robots and synchronizes the transfer of facial expressions to and from the model by sequentially invoking each robot's acquisition and action methods.

**Head orientation imitation.** In this example, we imitate a participant's head orientation and eye movements on a simulated iCub [27] as shown in Figure 3. The input coordinates arrive either from a wearable eye tracker [11] fitted with an IMU or a vision-based head pose estimation model [9]. The channeling scheme allows switching between the input sources by specifying the return element propagated to the robot.



**Figure 2: Facial expression imitation on the Pepper and iCub.**



**Figure 3: Head and eye movement imitation using either an IMU-fitted eye tracker or a head pose estimation model.**

## 5 CODE AND USAGE

To install Wrapyfi, compatible middleware, and required interfaces: <https://wrapyfi.readthedocs.io/>.

We additionally provide instructions on running Wrapyfi examples: <https://wrapyfi.readthedocs.io/en/latest/examples.html>.

Tutorials detail the steps needed to run the Wrapyfi use case scripts: <https://wrapyfi.readthedocs.io/en/latest/tutorials.html>.

We also evaluate transmission latency of the Wrapyfi plugins: <https://wrapyfi.readthedocs.io/en/latest/evaluation.html>.

## 6 CONCLUSIONS

Wrapyfi is a framework that simplifies data transfer across different middleware platforms. Two of Wrapyfi's key strengths are the transmission of custom data types and support for multiple middleware. We introduced three communication schemes—mirroring, forwarding, and channeling—each serving a different set of applications. The framework currently supports two common communication patterns: publish-subscribe and request-reply. In future work, we plan to extend Wrapyfi to support more communication patterns that are available in some middleware platforms, such as actions in ROS 2, which are similar to asynchronous request-reply. We also aim to provide interfaces for custom messages and middleware-specific data types. Wrapyfi's modular design permits integrating further middleware, expanding the array of potential applications.

## ACKNOWLEDGMENTS

The authors gratefully acknowledge partial support from the German Research Foundation DFG under project CML (TRR 169).

# REFERENCES

- [1] Martin Abadi et al. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems. In *12th {USENIX} Symposium on Operating Systems Design and Implementation (OSDI)*. 265–283. <https://www.tensorflow.org/>
- [2] ABi (Ed.). 2019. *Open-Source Robotics Projects*. ABi Research.
- [3] Junjie Bai, Fang Lu, Ke Zhang, et al. 2019. ONNX: Open Neural Network Exchange. <http://github.com/onnx/onnx>
- [4] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2018. JAX: Composable transformations of Python+NumPy programs. <http://github.com/google/jax>
- [5] G. Bradski. 2000. The OpenCV Library. *Dr. Dobbs's Journal of Software Tools* (2000).
- [6] Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. 2015. MXNet: A Flexible and Efficient Machine Learning Library for Heterogeneous Distributed Systems. *ArXiv abs/1512.01274* (2015). <https://mxnet.apache.org/>
- [7] Ayssam Elkady and Tarek Sobh. 2012. Robotics Middleware: A Comprehensive Literature Survey and Attribute-Based Bibliography. *Journal of Robotics* 2012 (2012). <https://doi.org/10.1155/2012/959013>
- [8] Charles R. Harris et al. 2020. Array programming with NumPy. *Nature* 585, 7825 (2020), 357–362. <https://doi.org/10.1038/s41586-020-2649-2>
- [9] Thorsten Hempel, Ahmed A. Abdelrahman, and Ayoub Al-Hamadi. 2022. 6D Rotation Representation For Unconstrained Head Pose Estimation. In *IEEE International Conference on Image Processing (ICIP)*. IEEE, 2496–2500. <https://doi.org/10.1109/ICIP46576.2022.9897219>
- [10] Pieter Hintjens. 2013. *ZeroMQ: Messaging for Many Applications*. "O'Reilly Media, Inc.". <https://zeromq.org/>
- [11] Moritz Kassner, William Patera, and Andreas Bulling. 2014. Pupil: An Open Source Platform for Pervasive Eye Tracking and Mobile Gaze-based Interaction. In *Adjunct Proceedings of the ACM International Joint Conference on Pervasive and Ubiquitous Computing (UBICOMP)*. ACM, 1151–1160. <https://doi.org/10.1145/2638728.2641695>
- [12] Yanjun Ma, Dianhai Yu, Tian Wu, and Haifeng Wang. 2019. PaddlePaddle: An Open-Source Deep Learning Platform from Industrial Practice. *Frontiers of Data and Computing* 1, 1 (2019), 105–115. <https://doi.org/10.11871/JFDC.ISSN.2096.742X.2019.01.011>
- [13] Steven Macenski, Tully Foote, Brian Gerkey, Chris Lalancette, and William Woodall. 2022. Robot Operating System 2: Design, architecture, and uses in the wild. *Science Robotics* 7, 66 (2022). <https://doi.org/10.1126/scirobotics.abm6074>
- [14] Anthony Mallet, Cédric Pasteur, Matthieu Herrb, Séverin Lemaignan, and Félix Ingrand. 2010. GenoM3: Building middleware-independent robotic components. In *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 4627–4632. <https://doi.org/10.1109/ROBOT.2010.5509539>
- [15] Giorgio Metta et al. 2010. The iCub humanoid robot: An open-systems platform for research in cognitive development. *Neural Networks* 23, 8–9 (2010), 1125–1134. <https://doi.org/10.1016/j.neunet.2010.08.010>
- [16] Giorgio Metta, Paul Fitzpatrick, and Lorenzo Natale. 2006. YARP: Yet Another Robot Platform. *International Journal of Advanced Robotic Systems* 3, 1 (2006), 8. <https://doi.org/10.5772/5761>
- [17] Youssef Mohamed and Séverin Lemaignan. 2021. ROS for Human-Robot Interaction. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 3020–3027. <https://doi.org/10.1109/IROS51168.2021.9636816>
- [18] Lorenzo Natale, Ali Paikan, Marco Randazzo, and Daniele E Domenichelli. 2016. The iCub Software Architecture: Evolution and Lessons Learned. *Frontiers in Robotics and AI* 3 (2016), 24. <https://doi.org/10.3389/frobt.2016.00024>
- [19] The pandas development team. 2020. *pandas-dev/pandas: Pandas*. <https://doi.org/10.5281/zenodo.3509134>
- [20] Adam Paszke et al. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32 (NeurIPS)*. Curran Associates, Inc., 8024–8035. <https://www.pytorch.org/>
- [21] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y Ng. 2009. ROS: An open-source Robot Operating System. In *IEEE International Conference on Robotics and Automation Workshop on Open Source Software (ICRAOSS)*, Vol. 3.2. IEEE, 5.
- [22] Neal Richardson, Ian Cook, Nic Crane, Dewey Dunnington, Romain François, Jonathan Keane, Dragoş Moldovan-Grünfeld, Jeroen Ooms, and Apache Arrow. 2023. *arrow: Integration to 'Apache' 'Arrow'*. <https://github.com/apache/arrow/>.
- [23] Matthew Rocklin. 2015. Dask: Parallel Computation with Blocked algorithms and Task Scheduling. In *Proceedings of the 14th Python in Science Conference*, Kathryn Huff and James Bergstra (Eds.). 130–136.
- [24] Henrique Siqueira, Sven Magg, and Stefan Wermter. 2020. Efficient Facial Feature Learning with Wide Ensemble-based Convolutional Neural Networks. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence*. AAAI, 5800–5809. <https://doi.org/10.1609/aaai.v34i04.6037>
- [25] SoftBank Robotics Group. [n. d.]. Pepper the humanoid and programmable robot. <https://www.aldebaran.com/en/pepper>
- [26] Yusuke Tanaka and Ankur Mehta. 2022. REMS: Middleware for Robotics Education and Development. *ArXiv abs/2210.05784* (2022).
- [27] Vadim Tikhanoff, Angelo Cangelosi, Paul M. Fitzpatrick, Giorgio Metta, Lorenzo Natale, and Francesco Nori. 2008. An Open-Source Simulator for Cognitive Robotics Research: The Prototype of the ICub Humanoid Robot Simulator. In *Proceedings of the 8th Workshop on Performance Metrics for Intelligent Systems (PerMIS '08)*. ACM, 57–61. <https://doi.org/10.1145/1774674.1774684>