Showcase 5 - Final demonstrator, in-hand object manipulation performed by robotic hand
(D3.14 - SGA3)

Figure 1: Image of the Shadow Dexterous Hand in the process of manipulating a cube.

Project Number:	945539	Project Title:	HBP SGA3

Document Title:	Showcase 5 – Final demonstrator, in-hand object manipulation performed by robotic hand
Document Filename:	D3.14 (D83) SGA3 M42 RESUBMITTED 231208.docx
Deliverable Number:	SGA3 D3.14 (D83)
Deliverable Type:	Demonstrator
Dissemination Level:	PU = Public
Planned Delivery Date:	SGA3 M42 / 30 SEP 2023
Actual Delivery Date:	SGA3 M42 / 18 SEP 2023 (resubmitted 08 Dec 2023)
Author(s):	Mario SENDEN, UM (P117) Tonio WEIDLER, UM (P117)
Compiled by:	Mario SENDEN, UM (P117)
Contributor(s):	
WP QC Review:	Jeannette BOSCHMA, UM (117), Anita VAN OERS, UM (117)
WP Leader / Deputy Leader Sign Off:	Rainer GOEBEL, UM (P117)
T7.4 QC Review:	N/A
Description in GA:	Simulation of robotic hand performing in-hand object manipulation.
Abstract:	This report presents the final demonstrator of the Human Brain Project, a simulation of in-hand object manipulation by an anthropomorphic robotic hand, the Shadow Dexterous Hand. Utilising deep learning, robotics, and neuroscientific insights, we developed a biologically constrained recurrent convolutional neural network (RCNN) using our newly developed open-source modelling library, AngoraPy. The RCNN, trained through reinforcement learning on high-performance computing clusters, simulates complex hand movements. The project successfully combined neuroscience and artificial intelligence to understand and simulate the human brain's coordination of in-hand object manipulation, demonstrating superior performance and faster convergence than non-biological counterparts do. The results provide valuable insights into the brain's sensorimotor transformations necessary for goal-oriented action execution.
Keywords:	Frontoparietal Network, Robotic Hand, In-hand Object Manipulation, Recurrent Convolutional Neural Network (RCNN), Open-Source Library, Goal-Driven Deep Learning
Target Users/Readers:	General public, neuroscience community, HPC community.

Table of Contents

Table of Figures

History of Changes made to this Deliverable (post Submission)

Date	Change Requested / Change Made / Other Action
18 SEP 2023	Deliverable submitted to EC
	Resubmission with specified changes requested in Review Report Main changes requested: • The document is poorly written in some parts: the Introduction includes copied sentences from later sections thus fails to adequately contextualise the work within HBP or EBRAINS. • Convolutional neural network implementation and development of AngoraPy for goal-directed deep learning is also used in D3.8, so this overlap should be eliminated (either here or in D3.8). • The same CNN is biologically inspired with respect to primate cortical connectivity, but with little details on the functional effect of such architectural specificity is reported - this must be updated. • Demarcation of the new work after M21 should be clearly described. • The claimed insights obtained on the sensorimotor transformations of the brain should be described. • No peer-reviewed publications (or any other references) are cited, so they must be provided in the revised version of the document.
	Revised draft sent by WP to PCO. Main changes made, with indication where each change was made: • Improved writing (all Sections) • Removed duplicate sentences (Section 1) • Added publications (see Section 3.6) • Added clear demarcation of the new work after M21 (see Sections 2 and 3) • Added description of insights obtained on the sensorimotor transformations of the brain (see Section 2)
	Revised version resubmitted to EC by PCO via SyGMa

1. Introduction

The intricate skill of in-hand object manipulation engages a network that spans sensory, association, and motor regions of the brain. Despite considerable progress, our understanding of how the brain coordinates dexterous movements remains incomplete. Showcase 5 addresses this challenge by harnessing the capabilities of deep learning, robotics, and neuroscientific knowledge.

The frontoparietal hand-action network executes sensorimotor transformations that are essential for performing goal-oriented actions. While decoding studies have provided important insights on the neural representations related to dexterity, they are ill-suited for characterizing the neural transformation performed by the frontoparietal network. To gain a complete understanding of dexterous actions, it is imperative to not only identify the representations carried by individual brain regions but also to examine computations that take place between these regions. The primary scientific objective of Showcase 5 is to shed new light on the transformations occurring in the frontoparietal hand-action network during in-hand object manipulation.

To achieve this scientific objective, Showcase 5 utilizes and expands upon goal-driven modelling (GDM). Deep neural networks (DNNs) combine brain inspired computational principles with unprecedented efficiency in solving perception tasks. This makes DNNs highly suitable for revealing the types of representations and transformations that may underlie complex, high-level functions of biological systems. As such, GDM seeks to develop biologically plausible DNNs to formulate new hypothesis about brain functionality. However, GDM, has long been solely applicable to the neuroscience of perception. Its advantages naturally extend to sensorimotor processing, but its adoption is hindered by the engineering and computational challenges that arise when networks are in closed loop with an external environment. Training such networks is difficult using typical supervised learning techniques that require labelled datasets. Reinforcement-learning (RL) constitutes an alternative as it allows for online data generation by exploring the task environment to optimise the network based on collected rewards. The methodological objective of Showcase 5 is to develop a framework that enables neuroscientists to design and train biologically inspired deep neural networks of sensorimotor functions.

The methodological goal was realized in the form of AngoraPy, a dedicated Python library unifying the design and training of goal-driven models and the construction of anthropomorphic tasks under the hood of an integrated, easy-to-use API. It targets neuroscientists who want to build biologically inspired, yet functionally performant models of the sensorimotor system without requiring them to have expert knowledge in RL or high-performance computing (HPC).

AngoraPy allowed us to significantly advance on the scientific goal of gaining a deeper understanding of the frontoparietal hand-action network. Using AngoraPy, we trained a recurrent convolutional neural network (RCNN) to perform an in-hand object manipulation task with an anthropomorphic robotic hand. The architecture of the RCNN is based on functional and anatomical knowledge of the frontoparietal network in the primate brain involved in visually guided hand and arm movements. Training data was gathered from the agent's experience within a simulated environment. The representations transformations emerging in the trained architecture have subsequently been investigated using tools from neuroscience and machine learning to provide new insights and hypotheses as to how the brain coordinates complex hand movements.

Showcase 5 relied heavily on the HBP's EBRAINS research infrastructure. Training in-hand object manipulation as reported here requires several years of simulated experience. Generating this experience and optimizing upon it is computationally expensive and requires extensive parallelization, for which the use of HPC clusters (SC6) is essential.

As we move forward, our primary focus will be on the continuous development and enhancement of AngoraPy. Our aim is to establish this toolkit as an open-source software that can facilitate research on the brain's sensorimotor system. In the period immediately following the HBP, we plan to dedicate our efforts towards building a robust community around AngoraPy. By fostering a community of users and contributors, we aim to ensure that AngoraPy continues to evolve and adapt to the changing needs of the neuroscience research landscape.

The work carried out in Showcase 5 contributes to the HBP Scientific Area of "Modelling". It uses goal-driven modelling to capture the comprehensive interplay between perception, cognition, and action in an ecologically valid manner through a model whose neurocomputational strategies automatically arise from parameter optimisation.

Communities that should be interested in this work include neuroscientists, cognitive scientists, roboticists, and AI researchers. This is because the project provides a novel approach to understanding complex brain functions, offers insights into the application of deep learning in neuroscience, and presents a new method for training artificial agents to perform complex tasks.

The downstream applications of Showcase 5 span across various domains, including neurohealth and robotics. The insights gained from this research could be instrumental in developing advanced strategies for individuals who have lost their motor skills due to neurological disorders or injuries. For instance, the principles learned from the project could be applied to design intelligent prosthetics that mimic natural hand movements, thereby improving the quality of life for these individuals. Furthermore, the work could contribute to the development of advanced robotic systems capable of complex, human-like manipulation tasks. The insights gained from this research could inform the design of robotic hands that can perform intricate tasks in various fields, such as manufacturing, healthcare, and service industries.

2. Scientific Motivation

Understanding the neurocomputational principles underlying human sensorimotor control is key to elucidating the mechanisms governing the human sensorimotor system. One prominent area of investigation concerns human dexterity, which involves highly precise movements of the hand and fingers. Showcase 5 is dedicated towards designing bioinspired models that can execute complex dexterous movements. To achieve this, showcase 5 employs deep learning to emulate biological brain regions and their functions. Deep neural networks provide exceptional effectiveness in solving ecologically valid tasks, and therefore have the potential to uncover the representations and computations that drive the complex, high-level functions of biological systems.

During the M21 review, Showcase 5 presented an early iteration of the dexterity model, demonstrating the coordination of basic hand movements, such as touching the thumb with a specific finger and performing finger-opposition sequences. Since then, we have developed a more sophisticated model that integrates both the macro- and mesostructure of the frontoparietal and pericentral cortical networks and is trained on the much more complex task of manipulating objects in the hand. The bioinspired model goes beyond current state-of-the-art A.I. models in terms of performance, and it offers valuable insights into the sensorimotor representations and transformations that occur within the brain.

Analysis of the biologically inspired model indicates that the Primary Motor Cortex (M1) and Premotor Cortex (PMC) in our model encode reward information. Furthermore, the model PMC is involved in action planning, while model M1 primarily executes current actions. These findings correspond with neuroimaging studies, thus validating the model. Furthermore, our analyses indicate that information becomes increasingly integrated and abstract as it travels from sensory and prefrontal regions to motor regions. This presents a novel hypothesis about the hierarchical nature of neural processing where complex, raw sensory data is progressively distilled into fewer, key components required for efficient movement planning.

3. Technical Specification

3.1 Goal-Driven Modelling of Sensorimotor Systems

Goal-driven modelling traditionally requires substantial amounts of labelled data and has as such been confined to the neuroscience of perception where such data is readily available. However, the ability of deep neural networks to combine brain inspired computational principles with functional performance can also be relevant for developing models of cognitive and sensorimotor processing. In particular, sensorimotor processing involves closed-loop systems where brain models act (via a body) on the environment and vice versa. Although efforts along this line exist, supervised learning under such conditions is difficult, as data is costly to collect and anatomically specific to the (potentially non-human) individuals recorded.

Reinforcement learning presents a potent alternative to supervised learning that forgoes labelled datasets. Instead, RL algorithms collect data online by exploring the task environment to optimise model parameters based on collected rewards. A downside of the RL approach is, however, its computational cost and substantial engineering effort. For this reason, the scientific community can strongly benefit from modelling libraries through which neuroscientists circumvent much of the implementation efforts and that simultaneously alleviate the computational costs in a distributed, efficient design. Showcase 5 developed a framework for this purpose. Following the M21 review meeting, this framework was consolidated in the form of AngoraPy (Weidler et al., 2023, Frontiers in Neuroinformatics; https://github.com/ccnmaastricht/angorapy; https://doi.org/10.5281/zenodo.7770180), a dedicated Python library unifying the design and training of goal-driven models and the construction of anthropomorphic tasks under the hood of an integrated, easy-to-use API. It targets neuroscientists that want to build biologically inspired but functional models of the sensorimotor system without requiring them to have expert knowledge in RL or high-performance computing.

3.2 Neural & Anatomical Model

We applied AngoraPy to study human dexterity through the lense of in-hand object manipulation. The human hand is controlled by a large-scale frontoparietal and pericentral network (FPN) mapping sensory information to desirable motor commands. This network, alongside the hand's anatomical capabilities, enables the human's unmatched dexterous agility. We modelled the hand as an anthropomorphic robotic body based on the Dexterous Shadow Hand, and the FPN as a deep neural network implementation biologically constrained by structural connectivity and regional cell densities.

3.2.1 Anatomical Model

A simulated MuJoCo implementation of the Dexterous Shadow Hand constitutes the anatomical model of the motor plant the neural model of this demonstrator controls. This anatomical model comprises 24 joints with 20 degrees of freedom and thus poses a complex challenge of fine motor control. Control is applied to positional actuators directly moving the joints. Muscles were not modelled to avoid additional complexity and achieve better performance. Figure 1 shows a rendered image of the simulated robot, which is holding the cube it needs to manipulate.

3.2.2 Network Model

We implement the frontoparietal and pericentral network controlling the robotic hand as a deep recurrent convolutional neural network. The architecture is biologically inspired through constraints to pathways and bottlenecks. Partaking regions have been identified from previous work revealing their involvement in manual motor control. Pathways were sourced from structural connectivity data extracted by siibra from the Jülich Atlas. Computational bottlenecks manifest themselves in the number of units per modelled area. An area's model unit count was determined as a relative mapping to the area's cortical cell count. To calculate cell counts, densities were extracted from von Economo and cortical volumes from the Jülich Atlas. Figure 2 depicts the resulting model.

3.3 Manipulation: An Ecologically Valid Dexterous Task

The demonstrator at hand concerns the dexterous task of in-hand object manipulation. We adapted the task definition introduced by OpenAI. The model must rotate a uniquely faced cube (see Figure 1) into a randomly sampled target orientation within 8 seconds. If the target is reached, a new goal is sampled and the agent can accumulate more reward, up to 50 times. If the target is not reached within the time limit, or the cube is dropped, the episode ends, and no more reward can be collected.

The reward we use to incentivise a human-like solution to this problem is the sum of (i) a constant punishment incentivising speed, (ii) a progression term rewarding the rotational distance made up for within a timestep, (iii) a punishment for dropping the cube, (iv) a bonus for successfully achieving a target, and (v) an energy usage cost reflected by the squared force applied in the joints.

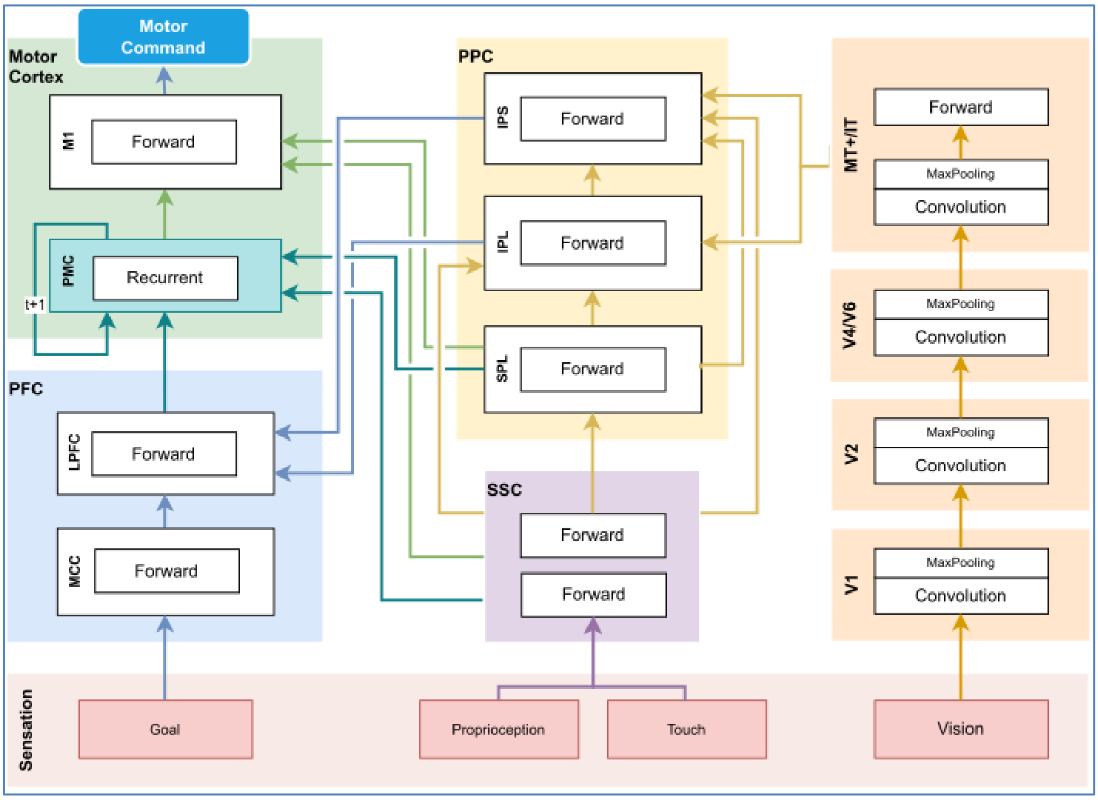

Figure 2: Biologically constrained architecture implementing the frontoparietal network.

3.4 Training

The behavioural policies adopted by our model of the frontoparietal network (depicted in Figure 2) are trained in a reinforcement learning setup using proximal policy optimisation (PPO). PPO alternates between two phases: the gathering of experience based on its current policy and the updating of the policy based on the previously collected data. Experience is represented by the state of the environment, the action the agent decided to execute in this state and the reward it received. Based on an estimate of the average value of every state and the actual reward received, advantageous actions are fostered, whereas disadvantageous decisions are suppressed. Our implementation of this setup efficiently trains arbitrary neural network models with potentially multimodal input (vision, proprioception and touch). Since M21, this framework has been further improved to meet the challenge of in-hand object manipulation. On the technical side, we most notably developed a new parallelisation system for both the gathering of experience and the optimisation (Figure 3). This new, native MPI implementation integrates smoothly with the remainder of the software without requiring additional third-party packages. Additionally, it is tailored towards the setup of EBRAINS' HPC cluster CSCS, on which we train the policies. Together, this enables substantially faster training on an arbitrary number of nodes, a crucial requirement for the complexity of the tasks tackled in Showcase 5. To integrate the visual modality, the system now communicates all modalities of states in separate streams, allowing any possible architecture to process them flexibly. Furthermore, policy and value functions are now trained asynchronously. Whereas the policy only receives humanoid sensory inputs, the value function can access additional high-level information such as the object's rotational and positional velocity. Thereby, the value function can make more precise predictions, which is crucial for effective learning, but the model retains its plausibility, since the dedicated parts of the network are omitted after training. Lastly, we shifted from a continuous, beta-distributed to a discrete, multi-categorical action space. While this substantially increases the dimensionality of the output space the model projects to, the predicted distribution is less complex. One should note that even using continuous actions, the agent learns reasonable policies. However, the performance in a multi-categorical distribution cannot be matched in reasonable training times. Regardless, since we directly predict joint angles and thereby already omit a biologically plausible translation from motor commands to actions, this does not harm the validity of our approach.

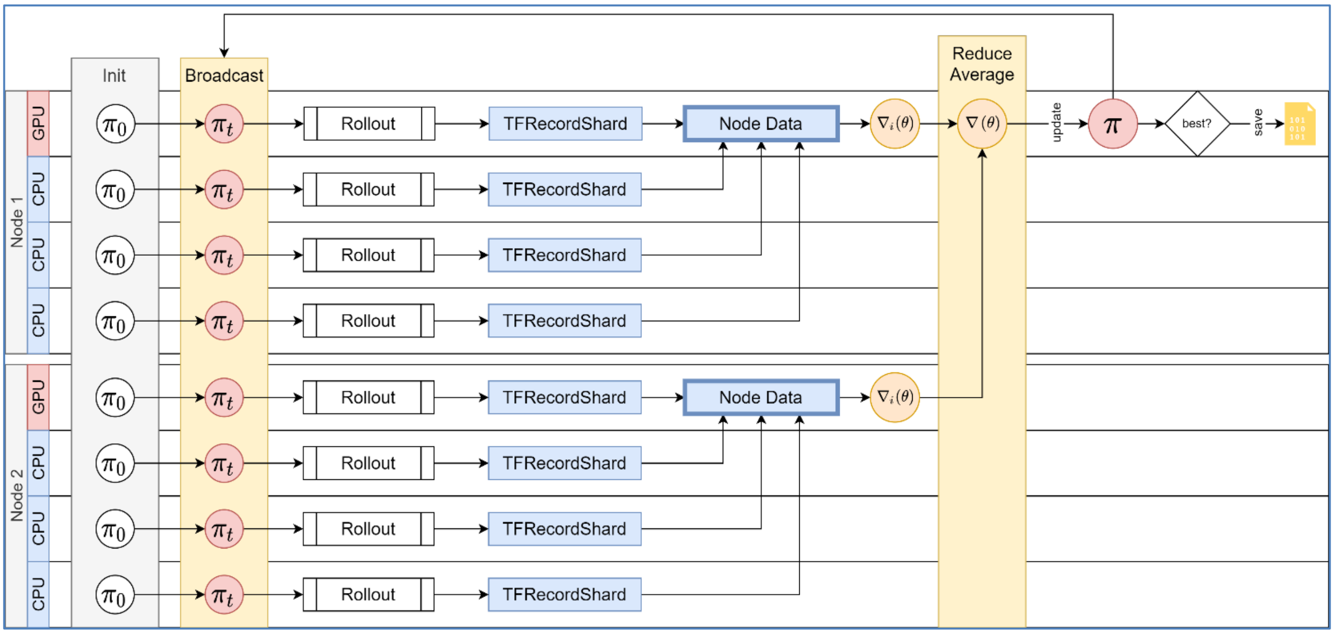

Figure 3: Integrated Parallel Setup of Both Gathering and Optimisation During Training

3.5 Performance & Analysis

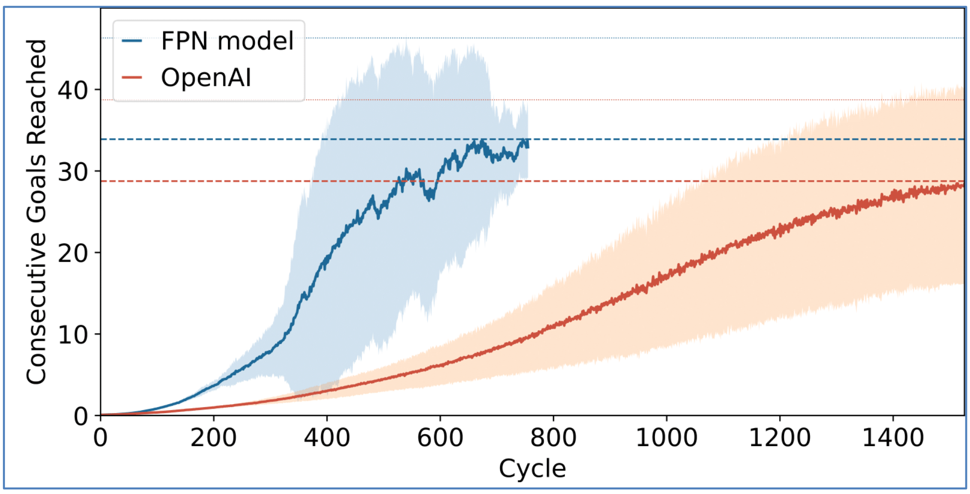

Figure 4: Learning curve of the FPN model developed for this demonstrator

We use AngoraPy to train both our FPN model and a baseline based on the architecture introduced by OpenAI. Agents are trained on EBRAINS Fenix Infrastructure (specifically CSCS Piz Daint) distributed over 32 nodes on which 384 CPU workers collect experience, and 32 GPUs optimise the network. After training on approximately 1.5 years of simulated experience, the FPN model is capable of chaining on average 35 goals. As Figure 4 shows, the brain-inspired model developed for this showcase significantly outperforms the baseline in convergence speed.

3.6 Publications

Weidler T, Goebel R and Senden M (2023). AngoraPy: A python toolkit for modeling anthropomorphic goal-driven sensorimotor systems. Front. Neuroinform. 17:1223687. doi: 10.3389/fninf.2023.1223687

Weidler, T., & Senden, M. (2023). AngoraPy - Anthropomorphic Goal-Oriented Robotic Control for Neuroscientific Modeling (Version 0.9.0). doi:10.5281/zenodo.6636482

Weidler, T., Goebel, R., & Senden, M. (2023). Synergizing Anatomy and Function: A Goal-driven Model of Frontoparietal Dexterous Object Manipulation. In *Conference on Cognitive Computational Neuroscience 2023* (pp. 147-150).

4. Relations to EBRAINS

The work greatly relies on HBP's EBRAINS research infrastructure. Training in-hand object manipulation as reported here requires several years of simulated experience. Generating this experience and optimising upon it is computationally expensive and requires extensive parallelisation, for which the use of high-performance computing (HPC platforms, SC6) clusters is essential. Furthermore, the design of the RCNN is guided by the human brain atlas (SC2) to identify brain regions involved in complex hand movements, as well as their interconnectivity and local cell profile. This benefits from the development of the siibra atlas client, which gives direct programmatic access to data organised within the EBRAINS human brain atlas. It supports the multilevel character of this atlas, which defines cytoarchitectonic maps in multiple reference template spaces at different spatial resolutions (namely the MNI Colin, ICBM152 asymmetric as well as the BigBrain microscopic space). It links them with complementary maps related to brain function, connectivity and fibre architecture.

Conversely, Showcase 5 informs the design of the supporting framework, in terms of specifications. We continuously communicate with the development team of the Neurorobotics Platform (NRP, SC4) to provide insights on the MuJoCo physics engine and its relevance for such training tasks, as well as essential requirements for transferring simulation of the RCNN onto the NRP, such as accelerated simulation time and the possibility to programmatically launch several instances of the NRP for deployment on high performance computing clusters. In fact, future versions of the NRP are now projected to include MuJoCo as an optional native simulator, which is further facilitated by the simulator's recent move to a free-to-use model.

Finally, AngoraPy has been incorporated into the corpus of EBRAINS tools. Standardised as a modular workflow, other researchers will be able to leverage the software without understanding in detail the underlying mechanisms.

Crucially, this demonstrator motivates the emergence of infrastructure workflows that span a meaningful range of complementary services, from data services (SC2) to embodiment and function-related services (SC4), with direct support from dedicated compute resources (SC6). No such combination of interoperable services is to be found outside of EBRAINS. Demonstrating the way EBRAINS thus empowers us to break new ground in cognitive computational neuroscience - and its practical application - contributes to establishing the infrastructure's Unique Value Proposition.

5. Looking Forward

As we move ahead, our foremost objective is to continually develop and enhance AngoraPy. Our goal is to establish this toolkit as a freely accessible resource to facilitate a comprehensive understanding of the brain's sensorimotor system. By improving and expanding AngoraPy, we aspire to create a novel computational and theoretical framework that enables a holistic study of the brain's sensorimotor system.

Following the Human Brain Project (HBP), our focus will be on constructing a strong AngoraPy community. We aim to establish this community as a cooperating platform wherein members can contribute to enriching the toolkit. Members can do this through comprehensive documentation provisioning, additional showcases and examples development, and the introduction of new functionalities to AngoraPy.

By fostering a user and contributor community, our goal is to ensure that AngoraPy continuously evolves and adapts to the changing needs of the neuroscience research landscape. This approach features collaboration, inclusivity, and versatility to enhance the utility of AngoraPy, pushing forward our understanding of the brain's sensorimotor system.