Showcase 6: Release of closed-loop sensorimotor demonstrators (D3.3 – SGA3)

Release of a set of functional, closed-loop demonstrators, connecting developed cognitive architectures to relevant physical agents, demonstrating the performance of a range of sensorimotor tasks.

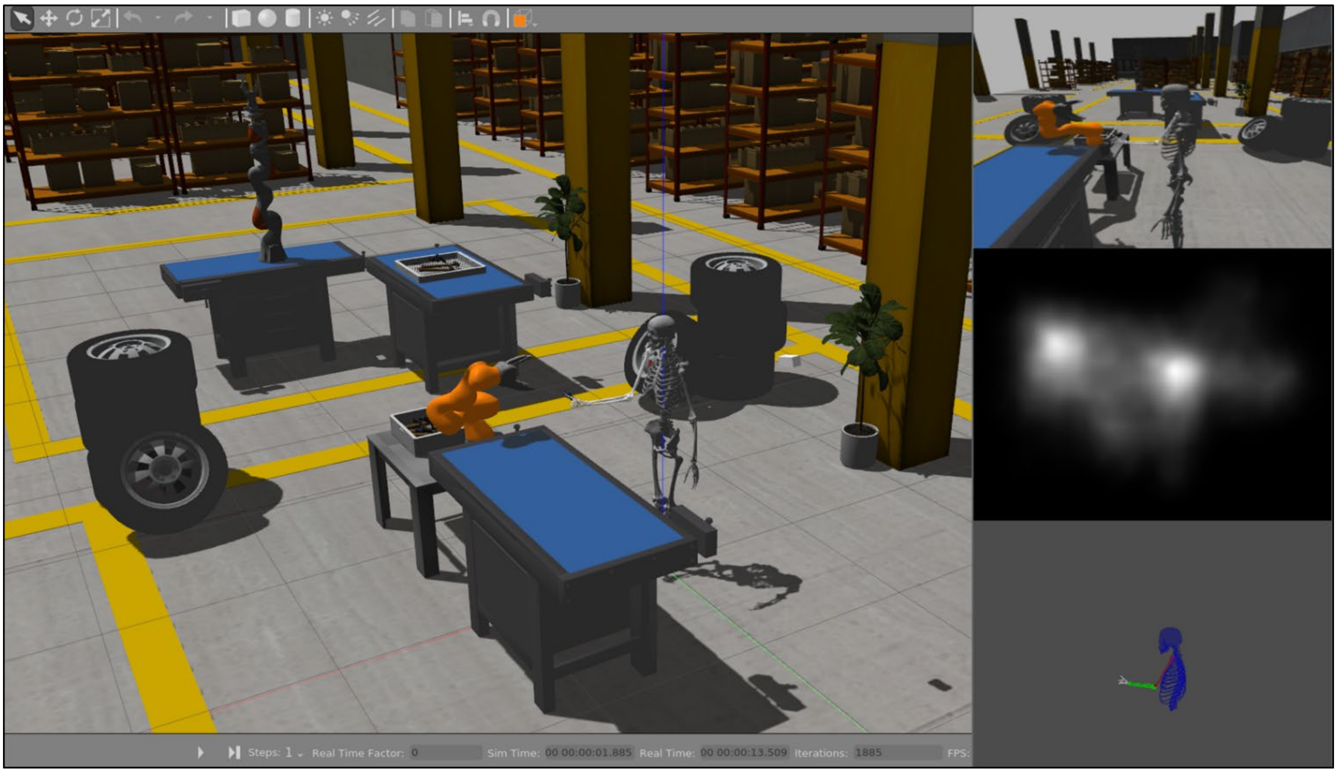

Figure 1 Simulation demonstrator; scene rendered (left), camera viewpoint (top right), computed salience map (mid right), and segmentation masks (bottom right).

Project Number:	945539	Project Title:	HBP SGA3

Document Title:	Showcase 6: Release of closed-loop sensorimotor demonstrators
Document Filename:	D3.3 (D22) SGA3 M21 ACCEPTED 220520.docx
Deliverable Number:	SGA3 D3.3 (D22)
Deliverable Type:	Demonstrator
Dissemination Level:	PU = Public
Planned Delivery Date:	SGA3 M21 / 31 Dec 2021
Actual Delivery Date:	SGA3 M21 / 23 Dec 2021; accepted 20 May 2022
Author(s):	Yannick Morel (UM, P117)
Compiled by:	Yannick Morel (UM, P117)
Contributor(s):	Alban Bornet (EPFL, P134), Gregory Francis (EPFL, P134), Alice Bruel (EPFL, P134), Ignacio Abadía (UGR, P66), Jörg Conradt (KTH, P39), Jens Pedersen (KTH, P39), Juan Romero Bermudez (KTH, P39), Thomas Nowotny (UoS, P106), James Turner (UoS, P106), Garibaldi Pineda Garcia (UoS, P106), Egidio Falotico (SSSA, P49), Francesco Iori (SSSA, P49), Gojko Perovic (SSSA, P49), Michael Zechmair (UM, P117), Raphael Stolpe (UM, P117)
WP QC Review:	Rainer Goebel, (UM, P117), Sacha van Albada (JUELICH, P20), Victoria Neumann (TUM, P56)
WP Leader / Deputy Leader Sign Off:	Rainer Goebel (UM, P117)
PCO QC Review:	N/A
Description in GA:	Release of a set of functional, closed-loop demonstrators, connecting developed cognitive architectures to relevant physical agents, demonstrating the performance of a range of sensorimotor tasks.
Abstract:	Deliverable D3.3 entails a number of demonstrators, composed of functional neural models, supporting a range of sensorimotor functions in closed loop. They are structured to address practical problems in the area of robotics and automation, specifically on the topic of human-robot interactions. The present document provides context to these developments, a summary of models implemented and results achieved by M21, and instructions on how to access the demonstrators. It further draws a number of perspectives, describing the targeted M42 demonstration and discussing additional emerging collaborations of relevance.
Keywords:	Functional neural models, functional architectures, embodiment, simulation, robotics, real-world systems, real-time systems, automation and control, neuromorphic computing.
Target Users/Readers:	Robotics community, automation, control and AI community, computational neuroscientists, cognitive neuroscientists.

Table of Contents

Table of Figures

1. Context

The work conducted within Work Package 3 (WP3) is structured around a number of integrative demonstrators, which motivate contributions from and active collaboration of most tasks (and corresponding areas of expertise) in the WP. These demonstrators address a range of functions, typically starting from lower abstraction sensorimotor loops, progressing towards cognitive skills, including planning and decision-making. They operate at different levels of description and biological plausibility, from detailed spiking models (emphasising functional description) to more abstract rate-based representations (emphasising functional performance). These demonstrators implement a functional modular approach, helping prototype a supporting modular cognitive framework, in close collaboration with EBRAINS Service Categories (SCs). They emphasise embodiment, and consider implementation on NeuroMorphic Computing (NMC) hardware where relevant.

Outline of the Showcase Demonstrator: Neural Architecture Supporting Safe Human-robot Interactions

The demonstrator discussed in the following lies at the more functional, less biologically plausible end of the spectrum. It addresses a cobotics scenario, in which a human is physically interacting with a robotic arm. The ambition consists in being able to demonstrate by M42 (safe) collaborative assembly, with a degree of task complexity comparable to that found on a factory assembly line (e.g. in automotive). Development of the demonstrator is being pursued starting from lower levels of abstraction (with aspects related to physics, embodiment, action, and perception), with a demonstration of sensorimotor loops by M21 (as illustrated by the demonstrator that the present document accompanies). Functions involved in this demonstration include perception, motor control, and planning, their development primarily conducted by contributors from Task 3.4. This demonstrator is to be built upon to include a number of cognitive functions, to be demonstrated in the M42 Deliverable D3.6: Closed-loop demonstrators addressing advanced cognitive and sensorimotor functions. In this perspective, the specifics of the demonstrator were selected to motivate contributions from a variety of participants in WP3. In particular, the considered scenario emphasises the need for robust scene understanding (investigated in Task 3.6, Task 3.1), with the ambition of providing safety guarantees for the human worker, without having to rely on substantial limitations of the robotic systems' speed of movement (thus preserving productivity). Developing the ability of the artificial agent to provide support in the assembly task motivates the integration of forms of hierarchical planning (Task 3.7), breaking down the overall assembly task in elementary steps. The ability to manipulate a wide variety of parts to assemble motivates the development of solutions to facilitate learning of grasping affordances (Task 3.4). The emphasis on safety also motivates the development of measures anticipating possible risks encountered by the human worker. This may be pursued by attempting to anticipate collisions and shocks (considering respective kinematics of robot and human movements), but also the mechanical stress applied to the muscle-tendon complexes of the worker engaged in the task (considering dynamic exchange of efforts between the robotic arm and human arm, for example). This can be pursued by extracting informative contextual cues from employed models and available sensors, and exploiting a model of working memory (as investigated in Task 3.7) to match emerging contexts with appropriate reactions from the artificial agent. Technical integration of the demonstrator is conducted in close collaboration between Task 3.4 (leading aspects related to embodiment), Task 3.1 (with relevance to the development of the supporting architecture framework), and Task 3.10 (software support). It is made possible by the support from partners in infrastructure WPs, in particular providing support on aspect related to modelling (SC3), embodiment (SC4), High Performance Computing (HPC) and NMC (SC6). The demonstrator discussed hereafter is chiefly developed and implemented on EBRAINS in a simulated scenario. A reduced-scope counterpart of this scenario, focusing on key functions and features, is being implemented on physical systems, relying on NMC hardware (in particular SpiNNaker, in collaboration with Tasks 5.8 and 5.10) to allow real-time operation. This is accompanied by an emphasis on functional Spiking Neural Networks (SNNs); with support from Task 3.3, in particular on aspects related to learning for SNNs.

2. Sensorimotor Demonstration

In the following, we summarise the developments conducted by M21 on sensorimotor aspects, describe the manner in which they come together to allow demonstration of sensorimotor loops, discuss the different scientific and technical problems addressed, and articulate the relation between the work presented and ongoing development of EBRAINS, in particular as pertains to functional modelling and embodiment SCs.

2.1 Specifications of the M21 Demonstration

Work conducted on motor control explores the complementarity between functional cerebellar and spinal cord models, for the purpose of motion control of upper limbs. This work builds upon previous investigations in the use of cerebellar models to address the overall control problem, and on spinal cord models to support rhythmic movements, typically for lower limbs. The cerebellar model supports motor learning and adaptation. It takes the form of a functional SNN that integrates the main cerebellar synaptic plasticity mechanism (Spike-Timing-Dependent Plasticity, STDP). It features five cerebellar layers: Mossy Fibres (MF), Granule Cells (GC), Purkinje Cells (PC), Climbing Fibres (CF), and Deep Cerebellar Nuclei (DCN). Inputs are applied through the MF (sensory signals) and CF (instructive signal). The MFs project the sensory information to GCs, which project through their axons (Parallel Fibres, PF) to PCs. PCs also receive excitatory inputs from CFs. At the PF-PC connection, STDP is balanced through Long-Term Potentiation (LTP) and Long-Term Depression (LTD), driven by the instructive signal from CFs. DCN neurons drive the final cerebellar output, receiving excitatory input from MFs and CFs, and inhibitory input from PCs (the only inhibitory connection in the cerebellar loop). The model used is presented in (Abadía, Naveros and Garrido, et al. 2019) and (Abadía, Naveros and Ros, et al. 2021). The spinal cord model relies on rate coding, following an approach similar to that in (Tsianos, Goodner and Loeb 2014). It reproduces the stretch reflex and reciprocal inhibition between antagonists for each muscle based on Prochazka's spindle rate. The former provides regulation of muscle length, while the latter limits co-contraction. The cerebellar and spinal cord models, together, control movements of a musculoskeletal model of a human upper limb, with a muscle model adapted from (Thelen 2003). The motion control problem addressed is that of trajectory generation for waypoint manoeuvring. In the M21 demonstration, the musculoskeletal model includes two Degrees of Freedoms (DoFs, single rotations at the shoulder and elbow) and seven muscle-tendon complexes. The cerebellum model is implemented using the Event-Driven Look-Up Table (EDLUT) simulation tool discussed in (Naveros, Luque, et al. 2014), (Naveros, Garrido, et al. 2017). The spinal cord model is implemented using Python with the FARMS library. The musculoskeletal model implementation was adapted from OpenSim and integrated in the Bullet physics engine (same physics engine used in the NRP). The trajectories followed are designed with consideration for natural human movement characteristics, as discussed in (Plamondon, et al. 1993). The cerebellum model is trained to adjust produced motor commands so that the effective, produced trajectory matches the prescribed one, as discussed in (Medina 2019). The integration of cerebellar, spinal cord, and musculoskeletal models in the demonstrator allows simulation of natural upper limb movements, including both kinematics and dynamics. The latter is of special interest for the considered demonstration scenarios, making it possible to explore exchange of efforts between the robotic system and a human worker's upper limb (robotic arm model and skeletal model shown in Figure 2, left). Such interactions may occur, for instance, in situations in which there is a handover from human to robot (or conversely), or when both human and robot act upon the same parts being assembled.

The work performed on perception primarily addresses the visual modality. It builds upon previous developments in Co-Design Project 4 (CDP4) on visuo-motor integration, combining specific models developed in CDP4 with a deep predictive coding model to pursue segmentation of objects of interest in the scene. More specifically, the saliency map computation model discussed in (Kroner, et al. 2020) is combined with a predictive coding model adapted from that in (Lotter, Kreiman and Cox 2017). The Predictive neural Network (PredNet) exploits the salience information to direct attention (in terms of prediction error computation, weighing prediction error as a function of computed salience of the corresponding part of the frame) towards more informative parts of the frame.

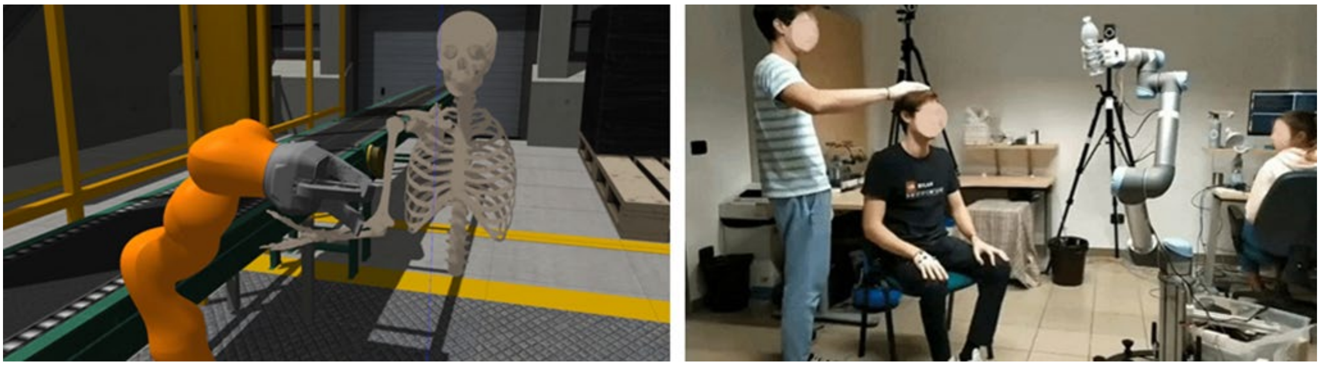

Figure 2 Close-up of robotic arm and gripper (left), motion capture setup to train planning (right).

Further, the PredNet model is adjusted to include connection time delays, as described in (Hogendoorn and Burkitt 2018), to elicit emergence of a hierarchical representation of temporal features. The predictive network is trained to achieve next-frame prediction (unsupervised learning). The model is extended with a decoder, trained to infer segmentation masks from the information contained in the PredNet's latent variables (see Figure 1, camera frame top right, salience map middle right, segmentation masks bottom right). Specifically, the decoder combines the different pooling levels of the latent variables using the up-sampling structure discussed in (Long, Shelhamer and Darrell 2015). The scene rendered by the demonstrator's simulation model is used to automate generation of training data. The combined saliency computation, PredNet model and decoder is used to detect and localise relevant segments of the human model (in particular, its arms). The information obtained allows inferring the location of these segments within the frame and within the scene. The localisation of the human arms is used to inform movements of the robotic arm, allowing movement coordination so that robot and human model are able to interact in a productive, useful manner (e.g. exchanging tools from human hand to robotic gripper, interacting together with parts being assembled). Work has begun on implementation of a complementary sensing modality, exploiting capacitive measures to infer proximity of the human, as described in (Schlegl, et al. 2013). The strategy under consideration involves relying on the visual modality to establish the overall scene structure, and relying on the proximity detection to ascertain the position of the human's hand and arm relative to that of the robotic manipulator. To that end, a capacitive sensor model was implemented in the simulation model, and an array of such sensors has been integrated in the robotic arm's segment supporting the end-effector. Note that measures from a single such sensor have been shown informative enough to allow approximate regulation of the gripper's position relative to that of the human hand (shown in Figure 2, left). The same PredNet and decoder structure as that used for vision has been adapted to infer the pose of the human arm relative to that of the robotic arm. Current results demonstrate adequate performance in terms of relative position estimation, with improvements still required from the relative attitude estimation. Networks used for perception were implemented using Python with TensorFlow.

In complement to the above, work was performed to develop the planning function necessary to guide movements of the robotic arm. In particular, efforts have targeted direct interactions between the human worker and the robotic system. Special attention is afforded to the process of handing over tools or relevant assembly parts (from the robot to the human and conversely), which remains an active area of investigation in robotics, as discussed in (Ortenzi, et al. 2021). The objective is to inform movements of the robotic arm with human expectations, in such a manner that interactions (e.g. when handing over an object) feel natural and comfortable to the worker. Such qualities are expected to lead to a quantifiable, positive impact on productivity, for instance, as measured by speed of completion of a given collaborative task, but also with consideration for the mental charge and fatigue of the worker. The planning model integrated in the M21 demonstrator relies on a Long Short-Term Memory (LSTM) model to predict movements of the human, and on a complementary Artificial Neural Network (ANN) to assess the rendezvous point for the handover. The information produced by these models informs a Dynamic Movement Primitives (DMP) framework prescribing robot movements. The approach allows performing the considered handover task without requiring a priori definition of the expected handover location. Further, the solution provides robustness to external perturbations such as an unexpected, external push on the human hand, or unpredictable, erratic hand movements from the human. Networks involved were implemented in Python. Training

Figure 3 Closeup of robotic arm, gripper, and human model (left), example of handover (right).

relied on data sets collected using a motion capture setup (shown in Figure 2, right). Recordings of a representative set of free human hand motions were used to train the human movement prediction model (LSTM), while recordings of human-to-human handovers were used to train the rendezvous point predictor (ANN). Networks were trained using gradient descent. Learning hyper-parameters were selected using Bayesian Optimisation and the Asynchronous Successive Halving Algorithm (ASHA scheduling).

The functional models described above were integrated into a single simulation model designed to be forward-compatible with the upcoming version of the NeuroRobotics Platform (NRP, developed in SC4). The simulation is anchored by a Gazebo model, providing rendering of the considered scene, with physics simulated using the Bullet physics engine. Bullet was extended to include a novel contact-resolution model (based on models of material deformation), to provide physically faithful, computationally stable resolution of contacts. This development proved necessary as off-the-shelf solutions (typically, constraint-based, such as those natively integrated in Bullet) struggle to reliably describe efforts produced by closed kinematic chains, such as those encountered in grasping. Bullet was also extended by the integration of a muscle-tendon model, as described in the paragraph on motor control. The robotic arm model considered is a Kuka iiwa, the gripper is the Robotiq 3-Finger Adaptive Gripper (see Figure 3, left). Functional integration of aforementioned neural models was performed through a lightweight Python signal coordination layer, built upon insights gathered from the NRP's Integrated Behavioural Architecture (IBA). Design of this communication layer was conducted in coordination with SC4 developers (Task 5.10), with insights provided by contributors to Task 3.1 on cognitive architectures, and benefitted from discussions with contributors to the Infrastructure Voucher on Cognitive Architecture for Therapy Robots and Avatars (CATRA), which addresses comparable problems.

The integrated simulation model illustrates the ability to inform movements of the robotic arm with information gathered by the perception function. It also demonstrates the capacity to simulate exchange of efforts between the human arm model and the robotic arm. This is demonstrated through three distinct simulation scenarios. In the first such scenario, the human (represented by a skeleton model) is firmly gripping a cylindrical rod. The robotic arm is also made to grip the rod, and attempts to pull it away from the human. This demonstrates the ability to rely on the visual perception to infer the location of the rod in the scene, allowing the robot to reach and grasp it. It also illustrates the ability to describe effort exchanges between the (relatively stiffer) robotic arm and (more compliant) human arm model, with the connected multibody systems forming a closed kinematic chain. Compliance of the gripper's fingers was adjusted to show slippage and release of the object by the robotic system upon crossing a given effort threshold. The second simulation scenario builds upon the first one to demonstrate a simulated handover. The human arm brings a tool within the workspace of the robot; vision allows inferring the position of the tool. Based on this information, the robotic planner adjusts movements of the robot to come within grasping range of the tool. The tool is then grasped (see Figure 3, right), brought to a storage bin, and released. The third simulation scenario provides a proof of concept illustrating the ability of the physics simulation model and combined sensorimotor functions to support simulation of an assembly task. The task

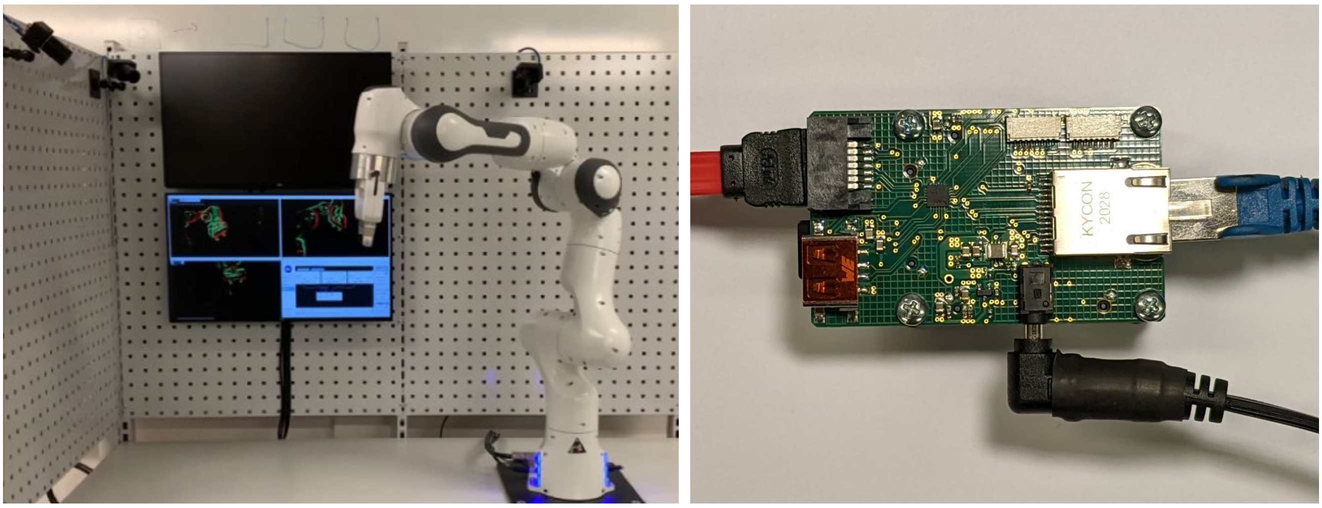

Figure 4 Physical robotic setup (left) and developed SpiNNaker I/O interface board (right).

considered involves affixing an interior panel onto the frame of a car door. The robotic arm is made to collect the panel from a resting position, and brings it into position, aligned with the direction required for clasping onto the assembly. The human worker then pushes the panel into the car-door frame.

Complementary to the simulation model discussed above, the same problem of supporting safe interactions between a robotic system and a human worker was investigated in a real-world setting, in a physical demonstrator (see Figure 4, left). Real-world developments being more effort-intensive than simulation models, the scope of the work considered is more limited than that in simulation, focusing on key technical and scientific challenges. A deliberate choice was made to investigate different approaches and technologies in this physical setup, as opposed to constraining ourselves to directly implementing (parts of) the models investigated in simulation. Specifically, this has allowed us to cover a greater overall scope, making the most of relatively limited resources. The physical demonstrator has focused on functional spiking technology, in particular using Dynamic Vision Sensors (DVS camera) for vision, and targeting implementation on SpiNNaker for real-time operation. Dedicated hardware was developed to achieve the input/output bandwidth (between real-world sensor and SpiNNaker system) required to support real-time operation, in particular allowing inflow of events from DVS cameras to SpiNNaker (Figure 4, right). Similarly, low-abstraction level software developments on SpiNNaker were required to implement operations necessary to support visual processing (e.g. convolution operations). This work was conducted in close collaboration between model developers (Tasks 3.4, 5.10), experts in learning for SNNs (Task 3.3), and core SpiNNaker personnel (Task 5.8), exemplifying the co-design process driving EBRAINS developments. The developed technology was exploited to implement event-based visual segmentation of objects of interest. The model developed builds upon that discussed in (Ronneberger, Fischer and Brox 2015), its development facilitated by the Norse library (Pehle and Pedersen 2021). Three distinct data sets were assembled to support training. The first data set was generated in simulation (using the NRP), collecting frames together with events for a set of tools in a given static scene. The second set considered a number of reference 3D-printed tools (e.g. hammer, wrench, or screwdriver) made to move in space, with events collected by stereo Davis346 event cameras. A third data set considers a reference shape (circular pattern), with events collected by three DVXplorer event cameras. The data was used to train SNNs to achieve segmentation, in real-time, on SpiNNaker. The model was trained with backpropagation through time, using the SuperSpike surrogate gradient method discussed in (Zenke and Ganguli 2018). The segmentation information is used to infer the object's location, which in turns is used to inform movements of a robotic arm.

2.2 Scientific and Technological Problems Addressed

The work conducted addresses a wide range of problems in computational neuroscience, but also in automation, control, robotics, and Artificial intelligence (AI). The central motivation consists in investigating research problems that EBRAINS' tools and services on modelling, embodiment, and NMC for real-time applications allow one to consider. These efforts enable informing ongoing infrastructure developments, highlighting specific needs using concrete examples. Each area of investigation also addresses a number of scientific problems as discussed below.

In particular, the work conducted on motion control allows exploring functional complementarity between spinal cord and cerebellum for non-cyclical movements, building upon previous efforts investigating embodied functional performance of standalone cerebellar models for motor control. Specifically, the implementation allows interrogating assumptions that have informed development of the considered models, specifically in terms of separation of functional concern (i.e. respective responsibility of the different models to support functions required to achieve the overall motor behaviour). In detail, the approach enables verifying improvements (e.g. facilitated, faster training) when the spinal cord fulfils its expected functions (e.g. muscle coordination), freeing the cerebellum from the corresponding functional burden. More generally, this perspective is similar to that pursued within Task 3.2 on functional scaffold models (work discussed in Deliverable 3.4), where part of the work focuses on composing models to describe the range of functions involved in sensorimotor loop. Active areas of collaboration are emerging in this respect, with functional models involved in the presently discussed demonstrator providing connections to embodiment for models investigated in Task 3.2. Further, from a dynamical system and control theory perspective, activities pursued on motor control aspects explore the merit of bio-inspired controller architectures, along an avenue of investigation that has proved particularly fruitful in pursuing legged and swimming robotic locomotion over the past decade. Specifically, the general approach relied on the use of a functional spinal model to support DoF coordination for locomotion, leaving the motion control problem to be addressed by a complementary control algorithm, whose output's role becomes comparable to that of descending signals. This perspective is here explored for upper limbs, with relevance to control of robotic arms, but can be generalised to address the broader problem of control in latent spaces. The spirit of this perspective involves the development of simplifying changes of variables, known to be of special import for instance in the control problem for non-minimum-phase nonlinear systems.

Work on perception aims to build upon HBP developments to implement solutions for robust scene understanding. The approach followed involves exploring merit (from a performance point of view) in this perspective of functions previously developed to implement a functional model of human vision processing, and of combining them with complementary deep networks. Of special interest is the notion of gainfully exploiting scene-representative information contained in latent variables emerging from trained deep networks. In addition, the work considered may explore the possibility of supplementing purely visual information with additional contextual cues, characterising relations in a broad sense between objects composing the scene. The objective in this respect involves going beyond label assignment to detected, classified objects, to include consideration of their role within the task at hand (e.g. relation between the detected tool and the current operation foreseen at the current stage of the assembly task underway), and additional contextual constraints (such as safety considerations in the situation of human-robot interactions).

The work conducted on planning aspects considers a problem in robotics that remains largely open, revolving around natural interactions between human and machine. In addition to addressing the aforementioned scientific questions, the technology developed with this demonstrator is expected to prove of relevance in several additional areas, most explicitly in industrial robotics, directly addressing some of the challenges that remain intrinsically related to the notion of Industry 4.0, in which the central promise of man-machine synergies has remained largely unfulfilled.

2.3 Relation to EBRAINS

The work involved relies on the HBP's Research Infrastructure and actively contributes to its co-design. In particular, it directly builds upon the NRP's IBA to support functional integration of a wide range of neural models with embodiment. The specific framework developed to support the demonstrator relaxes a number of requirements of the IBA (related to the corresponding version of the NRP) to facilitate integration. It was designed in collaboration between the technical engineering and scientific coordination task (Task 3.10), contributors to Tasks 3.1 and 3.4 with expertise on development of functional cognitive architectures, and contributors to the CATRA infrastructure Voucher. Further, the work conducted contributes to defining requirements in terms of required tools and services to support this type of research. This is in particular the case for aspects related to learning. Training of the functional models considered in the demonstrator (including event-based vision models, intended to be applied to real-world problems) makes use of simulation-generated data sets. Generation of such data sets can be supported by exploiting the type of digital embodiment afforded users by the NRP. The demonstrator provides a meaningful number of distinct motivating examples in this respect, helping establish what functionalities are useful for the platform to support. This involves aspects related to experiment scripting (for training-data generation), and deployment of parallel simulation instances on High-Performance Computing (HPC) platforms (SC6). A similar co-design approach is pursued in the use of NMC (SpiNNaker, SC6) to support real-time computation for vision (and in the middle term, motor control), with special emphasis on advancing the state of interfaces (in terms of usability and bandwidth) between SpiNNaker boards and physical systems (sensors and actuators). This work is enabling in nature, allowing materialising the potential of NMC (in terms of computation speed and energy consumption) for real-time, real-world visual processing. Relevance of this work extends beyond the considered robotic manipulation scenarios. It is of special interest for visual processing in embedded systems, providing attractive alternatives to currently used visual processing solutions in mobile systems. Prospective applications include unmanned vehicles (drones), but also driving automation. This development provides EBRAINS a unique capability in relation to NMC, contributing to its unique value proposition. Finally, the work conducted contributes to the development of closed-loop demonstrators showcasing the type of research made possible by combining functional neural models with embodiment, which corresponds to the specific value proposition of the NRP. It has motivated the development of a library of functional modules, which prospective users may train to support a range of sensorimotor and planning functions. As such, the work performed will expand the portfolio of content available for the NRP. We will investigate, in the near term, the opportunity of hosting and making discoverable such models on the Knowledge Graph.

2.4 How to access the demonstrator

The simulation demonstrator (including all relevant models) can be downloaded from the following location, https://drive.ebrains.eu/d/11cf313d0a724aa2807e/. Please refer to the document titled *Installation and Execution.odt* (found at the above link) for the requisite steps involved. A short video descriptive of the status of the physical demonstrator can be found at the following location, https://drive.ebrains.eu/f/b4dca37faeef42a99361/?dl=1. The simulation demonstrator is in the process of being ported to the current version of the NRP.

3. Perspectives

In the following, we discuss upcoming work, building upon the results achieved by M21. We distinguish specific contributions to the showcase demonstrator, which define specifications for the M42 showcase, and activities that extend beyond the scope of the demonstrator, including collaborations emerging from the work conducted.

3.1 M42 Showcase Specifications

Work on motor control will extend the number of skeletal DoFs considered (and corresponding required muscle-tendon complexes), to achieve a broader range of movements (at the shoulder, elbow, and wrist), extending the arms' workspace. Work on validation of the model will be conducted. To that end, electromyography (EMG) data is being collected to investigate the efficacy of the developed spinal cord model in reproducing human muscle recruitment profiles, for a battery of arm trajectories. Further, the developed models will be used to investigate to which extent the learning of movements is assisted by spinal circuits. Performance in this respect will be quantified in terms of speed of convergence and movement accuracy. Similarly, the contribution of spinal circuits to perturbation rejections will be investigated, with performance quantified in terms of perturbation force magnitude and resulting trajectory deflection. Activities on aspects related to perception will explore the development of multimodal perception models, combining visual and (capacitive) proximity perception. Work in this respect has begun, in collaboration with contributors to Showcase 4. An approach similar to that used in this demonstrator is investigated, building upon the Multimodal Predictive Coding Network model (MultiPredNet) discussed in (Pearson, et al. 2021). Similarity between considered problems greatly facilitates the collaboration. In particular, each included capacitive sensor in the present model fulfils a role qualitatively comparable to that of a given whisker in the Showcase 4 model. In addition, performance of the perception model developed will be benchmarked in comparison to alternative solutions found in the literature, with a special emphasis on quantifying performance improvements (in accuracy and computational burden) obtained from the use of salience computation and consideration of temporal patterns. The model developed for motion planning of the robotic arm will be extended to include functions of collision avoidance. Benchmarking will be performed to assess performance in terms of quantifiable metrics (e.g. task completion time or robotic arm idle time), but also accounting for qualitative aspects. These include subjective impression of safety and comfort of the human worker during interactions. Data collection in this perspective has begun. More broadly, the Showcase Demonstrator will be extended in the direction of greater abstraction, integrating cognitive functions in collaboration with relevant tasks in the WP. This includes the implementation of a hierarchical planning model, building upon sensorimotor functions to address an assembly task (expanding upon the initial proof of concept, shown in Figure 5). In addition, a model of working memory will be implemented to adjust behaviour of the artificial agent in response to contextual cues. Such cues will characterise human-robot collision risks (based on kinematic information) and stress applied to the human's muscle-tendon complexes. The robotic arm's behaviour will be altered to mitigate detected risks. Including support for flexible planning and working memory functions will lay the foundations for external users to adapt the presently described showcase demonstrator to address alternate tasks of interest to them.

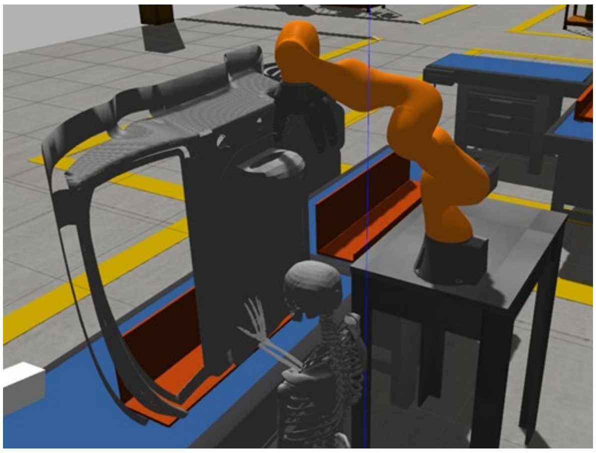

Figure 5 Human model affixing interior panel to a car door's frame.

3.2 Emerging Collaborations

In addition to the above, a number of productive collaborations have emerged from the work described. In particular, activities related to motor control constitute a natural extension of those on scaffold models in Task 3.2 (D3.4). A roadmap is being defined to provide Task 3.2 contributors with models developed for the present demonstrator, supporting embodiment (specifically, functional spinal cord and musculoskeletal models). The opportunity of extending this collaboration will be investigated, with the long-term perspective of developing functional models able to describe the mapping from skeletal movements to neural plasticity stimulation, with applications to the specialisation of post-stroke physical rehabilitation treatments. Another emerging area of collaboration is that with contributors to Showcase 4. In particular, specific aspects considered in the perception model discussed above (related to temporal patterns and decoding of internal, implicit information) are of relevance to the work conducted in WP2, where they are being investigated in a different setting. This illustrates the complementarity between work performed in both WPs, which approach comparable problems from different perspectives (with respective emphasis on biological plausibility in WP2, and functional performance in WP3). Finally, discussions are emerging on the possibility of merging (relevant aspects of) both showcases in WP3; for instance, including the anthropomorphic manipulator considered in Showcase 5 (and corresponding dextrous manipulation capabilities) in the present demonstrator to extend the range of tasks performed.

4. References

Abadía, Ignacio, Francisco Naveros, Eduardo Ros, Richard R. Carrillo, and Niceto R. Luque. 2021. "A cerebellar-based solution to the nondeterministic time delay problem in robotic control." Science Robotics. P2938

Abadía, Ignacio, Francisco Naveros, Jesús A. Garrido, Eduardo Ros, and Niceto R. Luque. 2019. "On robot compliance: a cerebellar control approach." IEEE transactions on cybernetics. P2362

Hogendoorn, Hinze, and Anthony N. Burkitt. 2018. "Predictive coding with neural transmission delays: A real-time temporal alignment hypothesi." Eneuro.

Kroner, Alexander, Mario Senden, Kurt Driessens, and Rainer Goebel. 2020. "Contextual encoder-decoder network for visual saliency prediction." Neural Networks. P1761

Long, Jonathan, Evan Shelhamer, and Trevor Darrell. 2015. "Fully convolutional networks for semantic segmentation." IEEE conference on computer vision and pattern recognition.

Lotter, William, Gabriel Kreiman, and David Cox. 2017. "Deep predictive coding networks for video prediction and unsupervised learning." International Conference on Learning Representations.

Medina, Javier F. 2019. "Teaching the cerebellum about reward." Nature neuroscience.

Naveros, Francisco, Jesus A. Garrido, Richard R. Carrillo, Eduardo Ros, and Niceto R. Luque. 2017. "Event-and time-driven techniques using parallel CPU-GPU co-processing for spiking neural networks." Frontiers in neuroinformatics.

Naveros, Francisco, Niceto R. Luque, Jesús A. Garrido, Richard R. Carrillo, Mancia Anguita, and Eduardo Ros. 2014. "A spiking neural simulator integrating event-driven and time-driven computation schemes using parallel CPU-GPU co-processing: a case study." IEEE transactions on neural networks and learning systems.

Ortenzi, Valerio, Akansel Cosgun, Tommaso Pardi, Wesley P. Chan, Elizabeth Croft, and Dana Kulić. 2021. "Object handovers: a review for robotics." IEEE Transactions on Robotics.

Pearson, Martin J., Shirin Dora, Oliver Struckmeier, Thomas C. Knowles, Ben Mitchinson, Kshitij Tiwari, Ville Kyrki, Sander Bohte, and Cyriel MA Pennartz. 2021. "Multimodal Representation Learning for Place Recognition using Deep Hebbian Predictive Coding." Frontiers in Robotics and AI. P3027

Pehle, Christian-Gernot, and Jens Egholm Pedersen. 2021. "Norse - A deep learning library for spiking neural networks." Zenodo. https://doi.org/10.5281/zenodo.4422025.

Plamondon, Réjean, Adel M. Alimi, Pierre Yergeau, and Franck Leclerc. 1993. "Modelling velocity profiles of rapid movements: a comparative study." Biological cybernetics.

Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. 2015. "U-net: Convolutional networks for biomedical image segmentation." International Conference on Medical image computing and computer-assisted intervention.

Schlegl, Thomas, Torsten Kröger, Andre Gaschler, Oussama Khatib, and Hubert Zangl. 2013. "Virtual whiskers—Highly responsive robot collision avoidance." IEEE/RSJ International Conference on Intelligent Robots and Systems.

Thelen, Darryl G. 2003. "Adjustment of muscle mechanics model parameters to simulate dynamic contractions in older adults." J. Biomech. Eng.

Tsianos, George A., Jared Goodner, and Gerald E. Loeb. 2014. "Useful properties of spinal circuits for learning and performing planar reaches." Journal of neural engineering.

Zenke, Friedemann, and Surya Ganguli. 2018. "SuperSpike: Supervised Learning in Multilayer Spiking Neural Networks." Neural Computation.