NRP News Archive

27 November 2018

New NRP release 2.1

The Neurorobotics Platform is proud to annouce its new 2.1 release, featuring:

- Improved easy install of the NRP using docker
- Now supporting Windows, Linux and Mac
- Changelog informs user of new features and compatibility breaks
- New frontend design with layouts and new toolbars
- Code editors auto-persistence (no save button anymore)

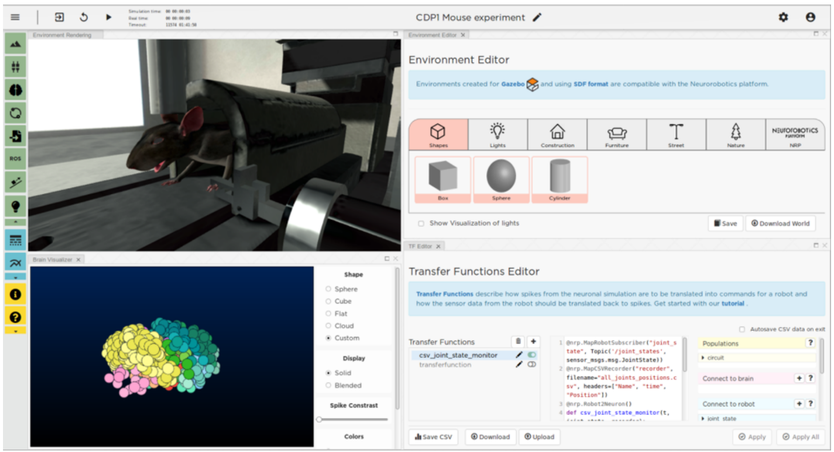

And coming very soon:

- Support for Nengo brains (coming in 2.1.1)
- Multiple robot support (coming in 2.1.2)

Available as usual as source or docker installation, and online from http://neurorobotics.net.

We hope you enjoy our new release and as always if you need support, please visit our forum.

02 November 2018

HTM Embodiment – A case study on how embodying a brain in a body unveils the full power of neural computation

As Neuroroboticists in the Human Brain Project we claim that a body interacting in the real world is necessary to understand the process learning and cognition in the brain.
Additionally, we claim that a body is not just an addon to let a brain do something in the world, but the body itself executes a part of computation (morphological computation, see for example the concept of a passive dynamic walker) and even further a neural network can only demonstrate its full capabilities if its linked within the sensory-motor space. These conceptual ideas gain more and more popularity, even though most AI applications in robotics still only see the body as a data generation engine, or in opposite direction as an execution engine of the network output.
We conducted a case study with the Hierarchical Temporal Memory[1] implementation to demonstrate how embodiment unveils new capabilities of a neural network in closed loop of a sensory-motor space. The neocortex inspired artificial neural network already demonstrates very good results for on data prediction, classification and find application in language processing scenarios. However, it still lacks breakthrough applications in robotic control scenarios.
The experiments where conducted with a Myororobotics[2] biomimetic robot arm: 3D printed ball and socket as well as a hinge joint imitate the skeleton of the upper human arm. Tendons, that are controlled by electric motors and routed via springs, imitate muscle activation to enable the compliant characteristics musculoskeletal bodies integrate naturally. The HTM is a very good fit for robotic applications as learning is done continuously and online.
Research was done in three stages, from an HTM application for prediction towards closing the sensory-motor loop for conditioned control:

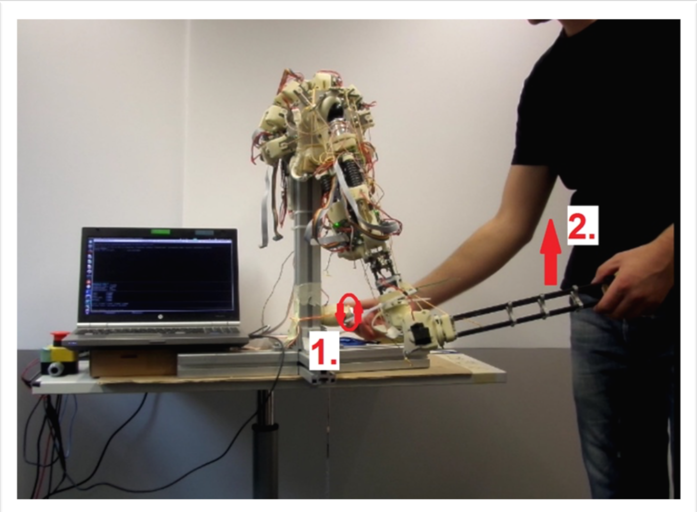

- **Proprioceptive motion prediction**

We demonstrate that the HTM can well predict motions of the robotic arm, that where generated by a human hand-shake like interaction. The HTM can deal with natural human motion variations and predicts arm positions well after only few iterations.

- **Neural motion storage and recall for robot control**

Once the HTM has learned a periodic motion, that means trajectory points are stored in neural synaptic connections, it can be recalled for motion control. In this base experiment a sinewave motion is learned, and afterwards the trajectory prediction output controls the motion. Hereby, the same motion that was learned can be continued for several iterations, and also parts of the motions a recalled occasionally.

- **Classical Conditioning of motion primitives**

The spatial and temporal associative capabilities of the HTM enable the HTM for Classical Conditioning. According to the figure above, we demonstrated subsequent stimuli and arm liftings to the robotic arm. The HTM learned the stimuli to behaviour relationship as well as the motion trajectory itself. After learning, same stimuli let the arm execute the same learned motion by neural recall. Also, variances of stimuli in contrast to the learned ones recall similar motions, e.g. a weaker stimulus triggers a weaker motion execution.
While the first experiment is a special application of HTM predictions, the second one is a first control example Jeff Hawkins already suggested that prediction outputs are similar to motor control commands. The third experiment combines both results. Here, the neural network is rather a client of the real-world situation. The HTM reveals a behaviour according to Classical Conditioning that would not have been observed just sticking with the principle of applying AI output to robotics but letting robotics "control" the AI neural network.
Ultimately, we see our results a first step towards agents that are able to learn from their environment and executing behaviours autonomously triggered by environmental stimuli. Due to the important role of Classical Conditioning in every mammal, we also hope to better understand the role of unsupervised learning for behaviour learning.
We appreciate to see our research results published in the "Bioinspiration and Biomimetics" journal and hope to foster the discussions on the importance of AI embodiment for truly intelligent robot development. Check out the paper results at https://doi.org/10.1088/1748-3190/aae1c2

- Hierarchical Temporal Memory: https://numenta.org/
- Myorobotics: https://myorobotics.readthedocs.io/en/master/

26 October 2018

NRP at the 28th Annual Conference of Japanese Neural Network Society

We presented our simulation platform at the 28th Annual Conference of Japanese Neural Network Society (JNNS) at Okinawa Institute of Science and Technology (OIST).

After a short introduction of Neurorobotics research we did a live demonstration and showcased our latest experiments. The platform gained a lot of interest and we could already talk about specific research experiments that we will integrate in the NRP in collaboration with institutes from Japan in the next weeks.
During the poster session we presented our two accepted papers "Multisensory Integration in the HBP Neurorobotics Platform" and "A Control Hierarchy Inpspired by the Spinal Cord to Exploit Self-Organizing Motion Primitives for Purposeful Trajectory Generation". Looking forward to the last day of interesting presentations about enhancements in Deep Learning, Spiking Neural Networks and Neurorobotics.

17 October 2018

Human Brain Project Summit 2018

The Neurorobotics group is currently attending the annual Human Brain Project Summit, this year taking place in Maastricht, and well there is quite a lot to report!
Starting on Monday we organized a joint booth with the division of Cognitive Architectures at the Open Day. People from all ages, research and industry as well as quite diverse disciplines had a look at the latest features of our Neurorobotics Platform, interacted with our biologically inspired robots and could even use Virtual Reality glasses to look and feel through a robots eyes. The Open Day was a great success as even more people attended compared to our last year booth representation.

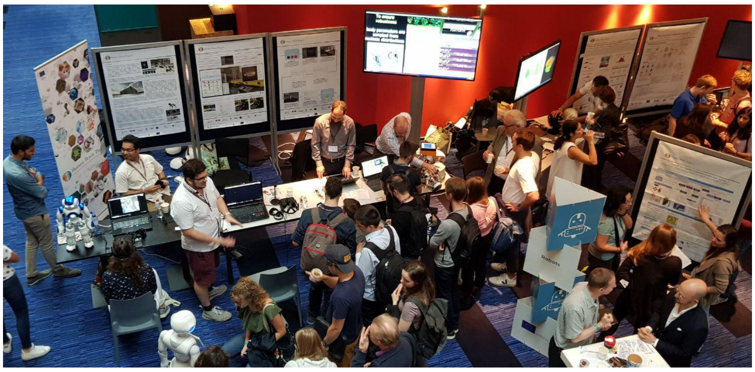

On Tuesday Mariya Gabriel, the European Commissioner in charge of Digital Economy and Society, opened the HBP Summit and we got the opportunity to present our latest research results at our robots booth. Apparently, she really liked the ongoing development as not only she stressed the Human Brain Project being one of the research priorities for the next EU budget, but also the robots booth got featured on her twitter account!

We now are gathering to exchange the latest results, discuss further collaboration opportunities within the Human Brain Project and outline the next research steps in talks, parallel sessions and workshpos. Afterwards, Thursday and Friday, we will conclude the week with an SP10 Neurorobotics meeting.

08 October 2018

Neurorobotics in the Human Brain Project at IROS Conference in Madrid

Last week we presented the Human Brain Project and in particular the Neurorobotics Subproject to the international robotics community at the International Conference on Intelligent Robots and Systems (IROS) in Madrid. During the exhibition our Human Brain Project booth presented facts and figures about the project and we answered any question from the robotics community. As a major, the Neurorobotics Platform was advertised on our posters but as well on our video screen and to be directly tested accessing our servers on a laptop. We could engage a lot of interested researchers and companies. Potential users of our platform but also multiple collaboration perspectives have been discussed and will be refined in the ongoing weeks.

Our booth was located right next to other European robotics projects. In particular we could also exchange and strengthen our collaboration with our booth neighbours such as the EDEN2020 to develop "An Enhanced Ecosystem for Neurosurgery in 2020". Also funded by the Horizon 2020 European Funding Programme, intersections in terms of brain health and robotics are key elements for exchange and support across the European project landscape.

25 June 2018

New Playlist on Youtube!

The Neurorobotics research team made 4 didactic videos to the NRP experiments they have been conducting in the last year. These videos give a nice overview of both the current features of the Neurorobotics Platform and the research being carried on in SP10.
Check our youtube channel at https://www.youtube.com/watch?v=Ya1VNq8X-g8&list=PLFfa5EHopIFrZRsJoNX4lNL8ZgZvXWhK3

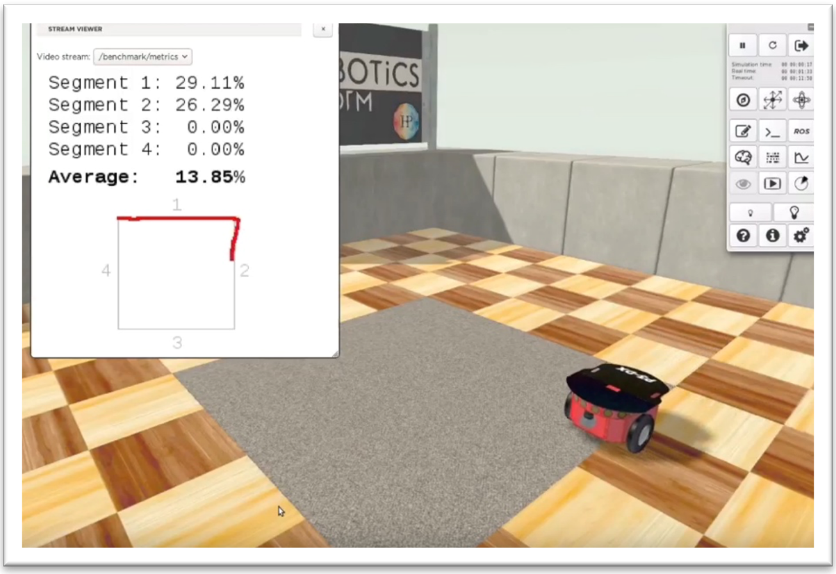

19 June 2018

The NRP supports Windows!

The Neurorobotics Platform docker installer now supports Windows!
Go here to learn more: http://neurorobotics.net/local_install.html

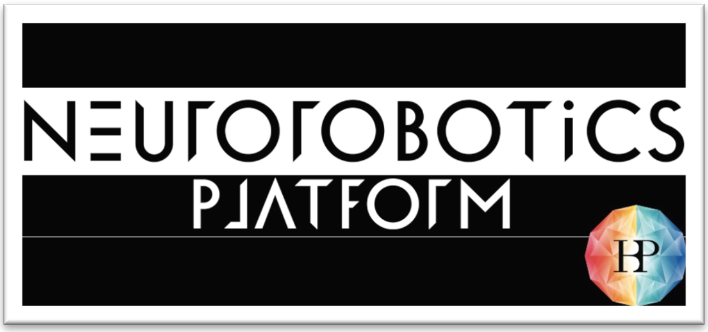

18 June 2018

Booth CeBit 2018

Date: 12.06.2018
Venue: Hannover
Duration: 4 Days
TUM presented the Neurorobotics Platform and our Robot Mice at the CeBIT.
For one week you could see our robots live and learn how they and the NRP are used within

the HBP research.
Thanks to the many people showing interest in the NRP and our Robots.

13 June 2018

Easy install using Docker available!

The Neurorobotics Platform is proud to announce the availibility of a binary and easy to install package. It is based on Docker and an installation script, supporting Linux only for now.
To install a ready-to-use NRP in minutes, follow the "Local Install" button at http://neurorobotics.net

11 June 2018

Come visit us at CEBIT!

The TUM is presenting the NRP and our robots at the booth of Bayern Innovativ.
Come see us in hall 27 at Booth F82 an see our Robot Mice in action!

11 June 2018

Bio-inspiration and modularity make robotic locomotion adaptable in the NRP

Current robotic control strategies are mainly based on trajectory plans that adjust the movements based on the next desired state. These control policies do not efficiently perform where the dimensionality of the control problem increases, or disturbances are perturbing the system from the external environment. These issues are critical in locomotion tasks and the need for different control methods arises. To make autonomous robots able to move in a real and dynamic environment, the research has focused on biologically inspired controller, such as neuro-controller.

The interaction among different bio-inspired motion controllers whose communication represents a simplified model of the neural locomotion control in vertebrates is possible in the Neurorobotics Platform.

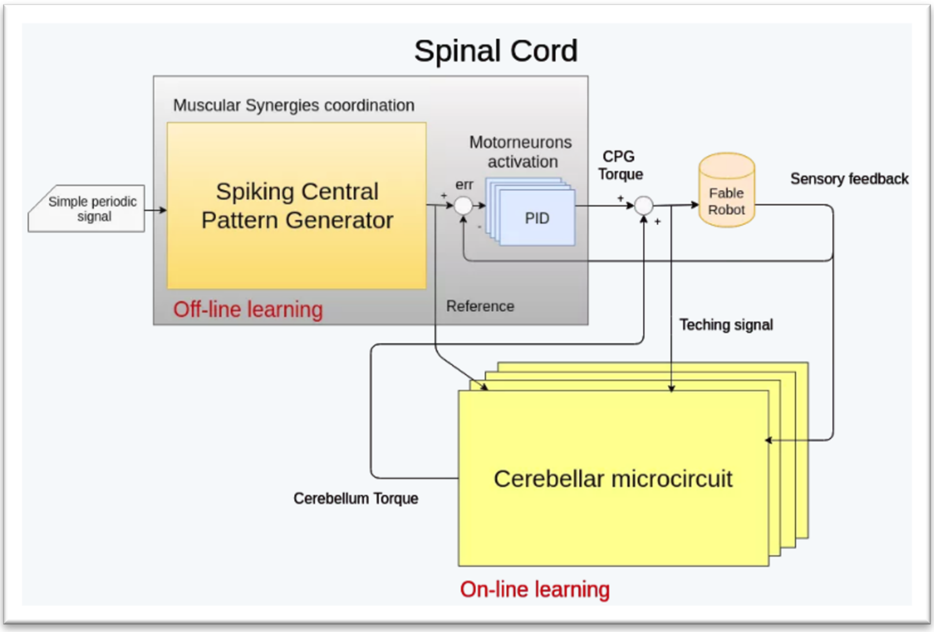

The presented solution combines classical control strategies with reservoir computing and spiking neural networks (*Reservoir computing with spiking populations* by Alex Vandesompele) to obtain a scalable and adaptable controller by taking advantage of the different learning properties of neural networks. To reflect the scalability of the controller, the experiments are performed on the simulated modular *Fable Robot.* In the experiment, it is built in a quadruped configuration so that the control architecture is composed of 4 cerebellar microcircuits, called Unit Learning Machines (ulm).
The use of a spiking neural network with reservoir computing as a trajectory planner (Central Pattern Generator, cpg) allows the learning of complex periodic trajectories for the movements of the robotic modules, whose frequency modulation is possible by just changing the frequency of the input signal to the network. Not optimally tuned PIDs give to the robot early stability during the first part of the simulation and provide a torque command for each module. Thus, a cerebellar network composed of 4 micro complexes computes and provides corrective effort contributions based on the inverse dynamics model of each robotic modules.
In the video below, it is possible to appreciate the locomotion improvements of the robot in an experimental simulation. The recording shows the simulation around second 100-130 when the position error is decreasing and stabilizing. The brain visualizer shown the spiking activity of the input population of the Central Pattern Generator (the higher groups) which are reflected in the blinking of one population of the reservoir (lower group). The spike train window (on the left) shows the periodicity of the spike trains which generate the trajectories for the modules (starting from the bottom, the activities of the input populations and one reservoir population are displayed).
The feed-forward cerebellar effort contribution decreases the mean of the position error of 0.3 radiant and its variance of 0.01 radiant compared to the case when just the effort command from the cpg is provided to the robot (the plot concerning the behavior of the second module is shown below). Moreover, the trend of the error is decreasing along the simulation time and the distance covered by the robot with the cerebellar-like neural

network contribution is 9.48 m while the cpg controller contributes to have the robot walk for 1.39 m.

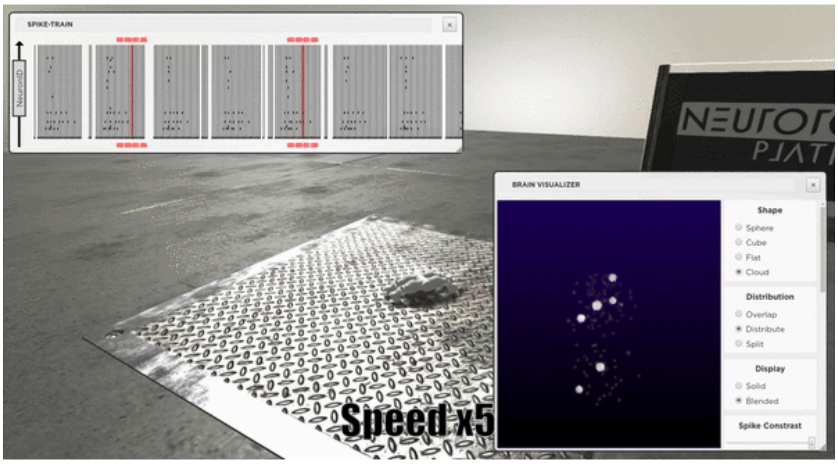

The modular configuration of the *Fable Robot* makes easier to test the control strategy for different configurations of the robot and patterns of locomotion, having the cerebellar-like neural network compensate the error after a short learning phase, since it has previously learned the internal model of the module.
This work was done in collaboration between DTU, Ghent and SSSA teams.

18 May 2018

A new integration experiment for the visual system of the NRP

We implemented a new experiment to demonstrate that the NRP is able to run many models together, as a single visual system. Here, a retina model, a deep neural network for saliency computation, a spiking cortical model for early-stage visual segmentation and an echo state network for saccade generation collaborate, despite the differences in their structures. The NRP provides a common framework where models can talk to each other easily.
In the video below [available at https://youtu.be/i4nFl9xzbuk] the robot has to keep track of a stimulus that moves on a screen.

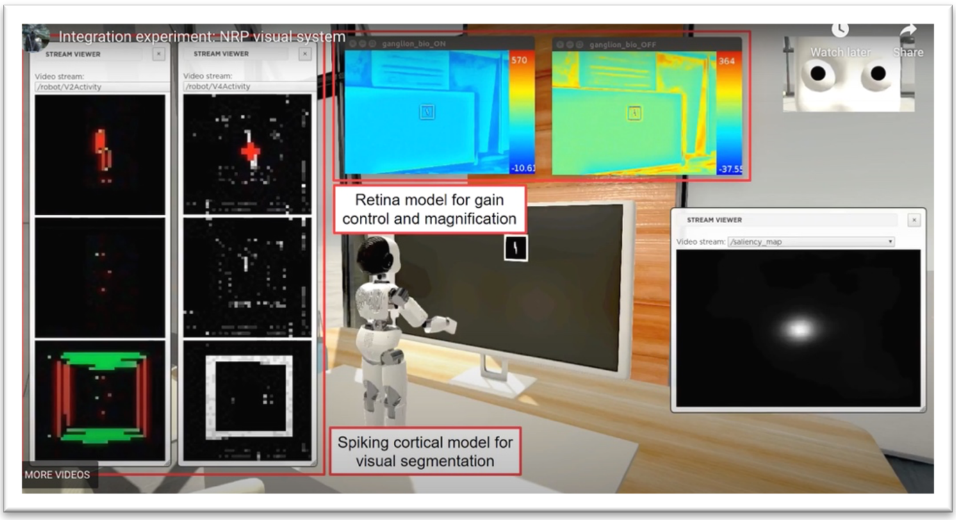

The saliency model computes where the stimulus lies (most salient region in the visual field), and whenever it is not in the fovea, it delivers a signal to the saccade model, so that an eye movement is generated towards the visual stimulus. When the stimulus reaches the fovea, the second task is to segregate the target (small tilted bars) from the flanking square (impairs target detection). For this purpose, when a saccade is generated, the segmentation model triggers saccade inhibition and then sends local signals that initiate a spreading segmentation process that lies in the networks dynamics, attempting to segregate the target from the flanker. These local signals are sent, using the saliency output as a 2D probability density distribution, so that segmentation is only triggered around regions that are worthy of attention. Only when the segmentation is successful, the target is detected. During the whole experiment, the retina model delivers adaptation to the lighting of the scene. Thanks to it, the cortical representation of the stimulus is stable and segmentation is possible even in low-lighting conditions.
The NRP demonstrates its ability to run large-scale, collaborative simulations. Being able to run a whole visual system gives many novel opportunities to vision research, especially to explain global effects in human vision.

16 April 2018

Reservoir computing with spiking populations

Passively compliant robots are robots with passive compliant parts, for instance springs or soft body parts. They can be cheaper, safer and more versatile than traditional stiff robotics. Since the compliance introduces non-linearities that are not easy to model analytically, we need to monitor the body with sensors and use machine learning to interpret those sensors.

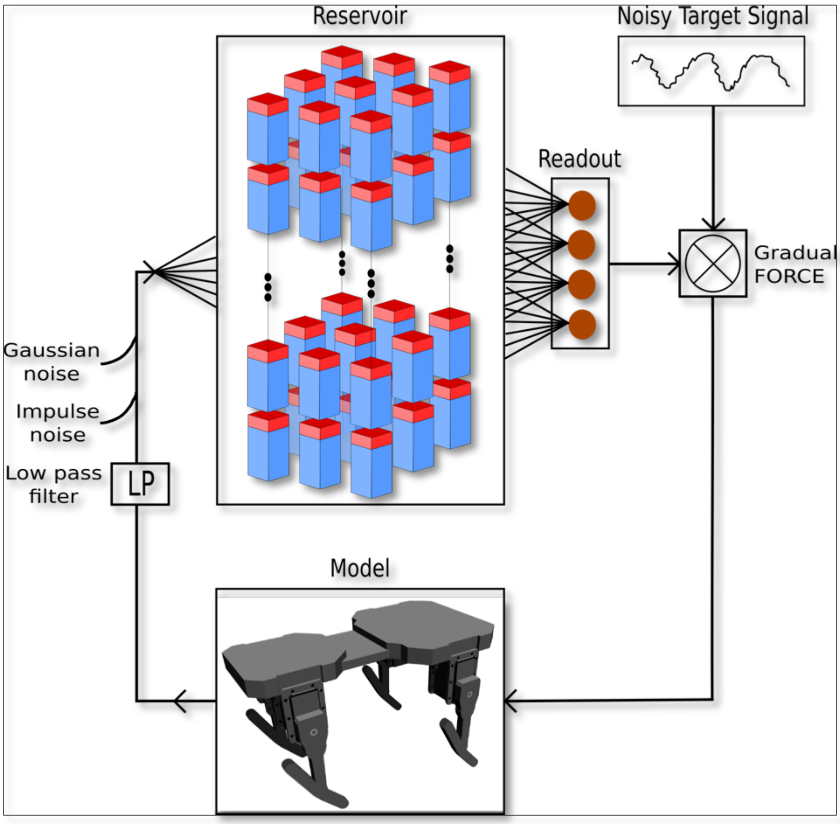

Reservoir computing allows to train non-linear dynamical systems using only simple machine learning techniques. The unit of our reservoir is a population of spiking neurons. During training the desired motor commands are gradually taught to the closed loop system (with gradual FORCE learning), as illustrated in the figure below. The only weights that need to be learned are those to the readouts.

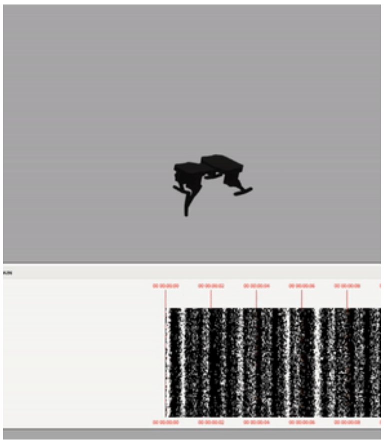

The neurorobotics platform provides a convenient interface between the robot model and a spiking 'brain'. After training, we have a closed loop system consisting of only the body and its 'brain'. The body sensors drive the 'brain' activity which in turn drives the actuators. The motor commands are 'embedded' into the dynamics of this system. The animation below shows the resulting trained closed loop gait controller. Spike trains are shown for all neurons in one population.

If desired, the system can be trained with an extra input to the reservoir (in addition to sensor inputs). This extra input can be coupled with different motor commands (for instance different gaits, or different frequency of the same gait). After training, this extra input can be set by an external actor (be it another 'brain' region or just a human) to control the system in real time:

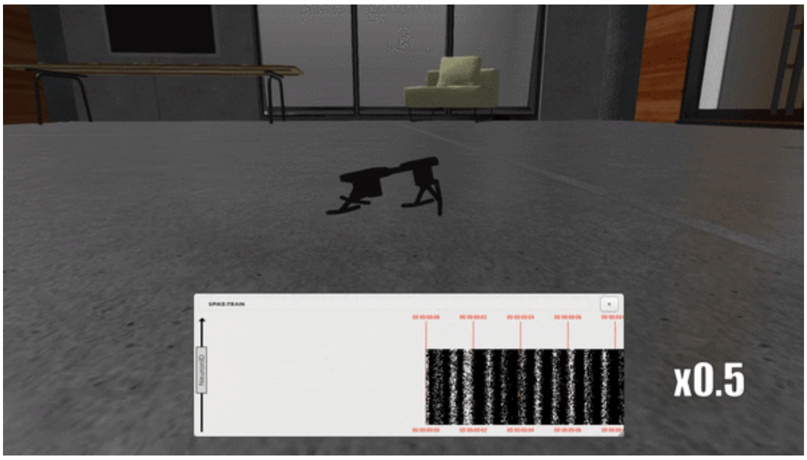

03 April 2018

New mechanisms in the Laminart model

The Laminart model is a spiking cortical network for early-stage visual segmentation. Using network dynamics simulated in NEST, it is able to parse its visual input into several perceptual groups. It is the largest simulated network available on the Neurorobotics Platform. It explains very well many behavioural results about visual crowding (a behavioural paradigm in which human observers try to identify a target being disrupted by nearby flanking shapes, the target being "crowded" when identification performance is impaired by the flankers – see fig. 1).

Fig. 1. a) Crowding in real life. If you look at the bull's eye, the kid on the left is easily identifiable. However, the one on the right is harder to identify, because the nearby elements have similar features (orange color, human shape). c) Crowding in behavioural experiments. The visual stimuli on the x-axis are presented in the periphery of the visual field of human observers. The task is to identify the direction of the offset of the target (small tilted bars). Flanking squares try to decrease performance. What is plotted on the y-axis is the target offset at which observers give 75% of correct answers (low values indicate good

performance). When the target is alone (dashed line), performance is very good. When only one square flanks the target, performance decreases dramatically. However, when more squares are added, the task becomes easier and easier. All classical models of crowding fail at explaining the latter condition (uncrowding), because they all predict that more flankers induce more interferences.

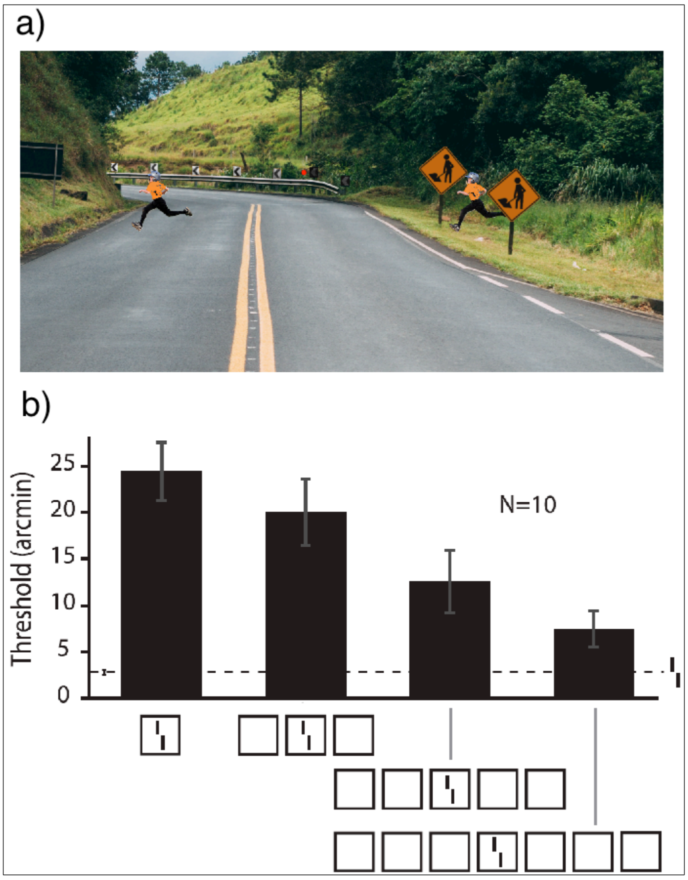

The model explains the results by stating that the target is crowded when it is considered as in the same perceptual group as the flanker, allowing interferences. When the flankers form a group on their own, they "frame" the target, increasing performance (less interference happens between elements that are in different perceptual groups). This idea explains the behavioural results very well. More than that, the model offers neural mechanisms for perceptual grouping and segmentation. However, these mechanisms work properly only with very simple flanking shapes (squares, rectangles, etc.). As soon as the shapes are slightly more complex (more orientations, various scales, etc.), the mechanisms that are necessary to produce the results do not work properly. This is a problem, because uncrowding happens for any flanking shape (not only squares). Here, we introduce new ways of implementing the key mechanisms of the model.

The first mechanism groups boundaries together (see fig. 2). This is linked to illusory contours computation. Activity spreads to connect edges that are well aligned. The new idea that was implemented here is that the illusory contours can now spread at different scales, according to the configuration of the stimulus. For example, if you are reading, you will tune your illusory contours mechanism, so that horizontal grouping happens between letters of the same word, but not between words, using the fact that letters of a same word are closer than letters from neighbouring words. We used the same idea for our crowding paradigms, where different grouping scales lead to different crowding results.

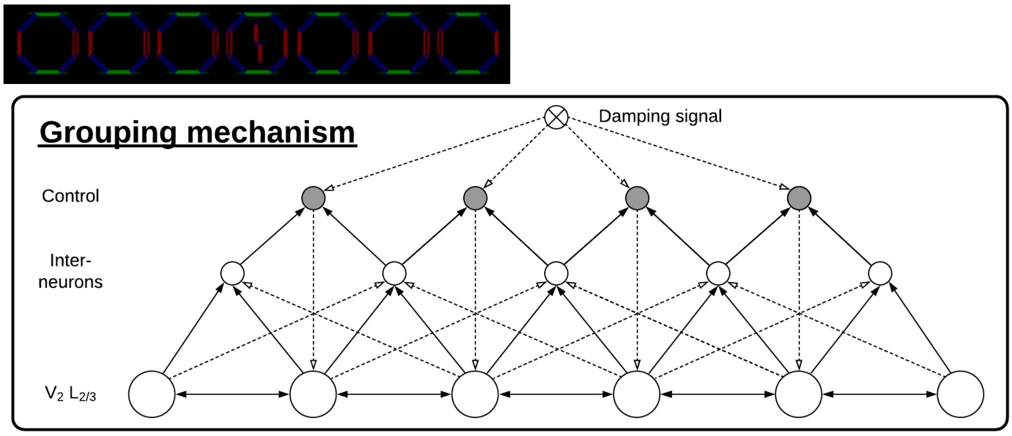

Fig. 2. Top: illustration of how the multi-scale grouping mechanism works (V2, layer 2/3 activity). Left: if tuned for low distances, no illusory contour will spread from the flankers to the target. Right: if tuned for long distances along the horizontal direction, illusory contours easily spreads through long distances, to link all the flanking octagons. Bottom: dynamics of the grouping mechanism. Activity naturally spreads in V2. However, spreading also activates interneurons that activate spreading control. In this state, nothing spreads. However, stimulus onset triggers a damping signal whose duration is specific for the orientation of spreading. It is highly dependent on the stimulus shape. For example, in the top-left stimulus, the closed surfaces impair spreading around the target. In the top-right stimulus, the horizontal alignement of the flankers makes that the damping signal is quite long along the horizontal direction, allowing horizontal boundaries to spread across flankers. When the damping signal stops, the illusory contours are already stabilised and do not spread back.

The second mechanism parses subsets of the image that are linked by boundary grouping to different segmentation layers (different regions of the network's activity – see fig. 3). In our case, this allows the model to block interferences between parts of the image that belong to different groups. The new mechanism is better, because the former one required tonic activity of many neurons in the segmentation network, which was biologically implausible. Now it is only driven by input-related activity and by a brief segmentation signal.

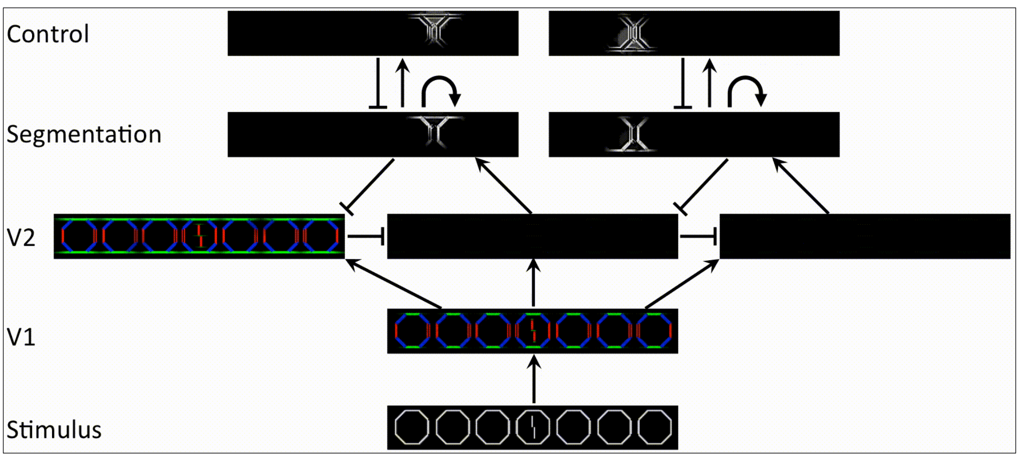

Fig. 3, top: Illustration of the mechanisms of segmentation. After the segmentation signal is sent, activity spreads from V2 segmentation layer 0 (SL0) to V2 SL1-2 (after competition, all the flanker-related activity ends up in SL2). Activity can spread throughout all flankers thanks to the illusory contours between them. This is only possible with the new grouping mechanism. Bottom: dynamics of the segmentation mechanism. The interneuron layer is the one that does all the spreading. It also triggers its own control by spreading. The control layer is being inhibited by activity in SL1, allowing the interneurons to shut down activity in SL0, disinhibiting activity in SL1. Because all the dishinibition relies on V2 activity, it will only spread along connected boundaries. You can see in the figure on the top that the control layer acts as a sheath to make sure that the interneuron activity only spreads along V2 activity.

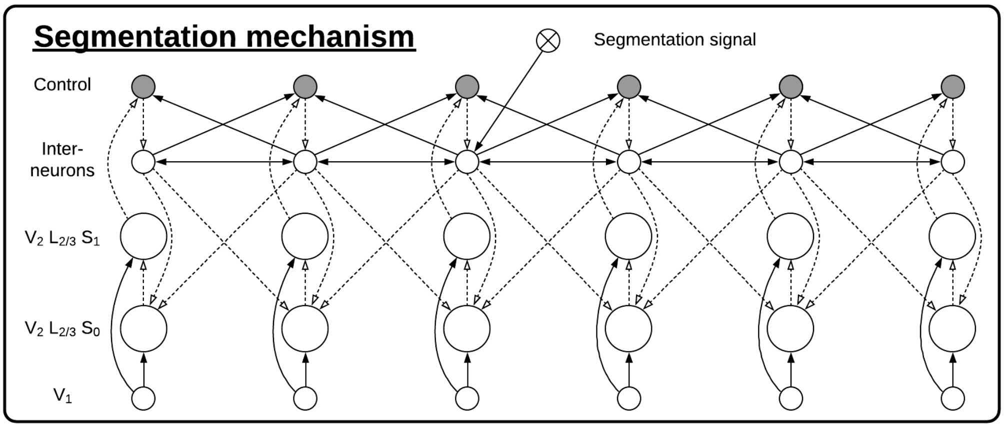

In the future, the new mechanisms will be generalised to more orientations (ongoing work), and the new version of the model is going to be integrated on the NRP. Also, top-down influence on how damping signals distribute across orientations will be investigated.

21 March 2018

We were mentioned in the press!

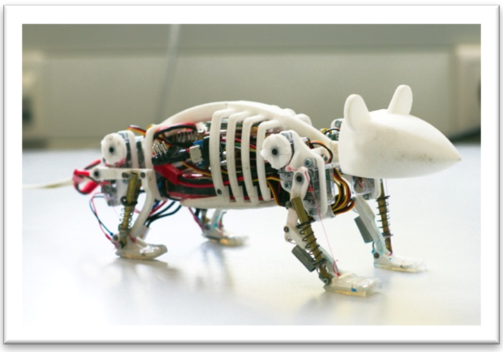

Robots with brain, an article explaining what we do in the Neurorobotics Platform. Thanks to fortiss GmbH for featuring our work in their website For the full article click here

19 March 2018

Going beyond conventional AI

European Robotics Forum 2018 in Tampere

The Neurorobotics Platform developed in the SP10 keeps improving its usability and reliability, and is looking to expand its user base. If the feedback obtained from the audience at the European Robotics Forum (900 registered guests, all roboticists from research and industry) is anything to go by, the NRP is in prime position to fill the need expressed by this community for an interdisciplinary simulation platform than connects neuroscience, AI and robotics.

Indeed, during our Workshop at the ERF and the various discussions that ensued, we were able to speak with a large number of researchers and company representatives from different backgrounds and activities. The overwhelming majority has clearly caught on the potential advantages of using the NRP, especially with standard AI tools such as TensorFlow. Furthermore, we found they were open to considering the ability of the NRP to establish brain-derived intelligent controllers that go beyond conventional AI. Finally, compliant robotics based on the up-and-coming technology of elastic elements that can make robots safe by design is an active area of research where ERF participants also saw potential for the NRP (OpenSim, custom robot designer, etc.).

We are thus looking forward to collaborating with our new contacts in the industry, and to improving the platform even further for their benefit.

Benedikt Feldotto (TUM) walking the audience through the NRP's many features)

19 March 2018

Using the NRP with motor actuated robots

The Neurorobotics Platform Development and Research teams are joining their forces to give the users the possibility to interface their own motor actuated robots with the NRP. This new intriguing feature will help all the roboticists to deploy their neural controllers on real robots, without the need of any external software.

A new working-mode for the NRP has been created by the SSSA team in Pisa, supporting NEST as the neural interface and ROS for the robotic interface.

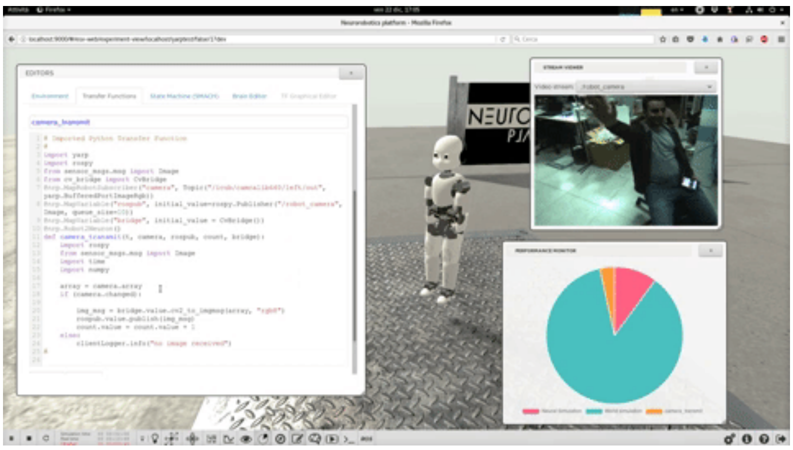

The iCub humanoid robot doesn't have a native support for ROS, but instead it uses YARP (Yet Another Robotic Platform) as a middleware.
This makes it necessary for the NRP to support YARP in order to interact with the hardware iCub platform. Such integration is also being currently developed by the SSSA team.
While not being able to send commands to the iCub yet, this extension allows already to read from its sensors, as for instance the eye cameras.
You can see in the picture below one of our team members waving at the iCub seen from the NRP web frontend.

A further mode is being developed by the KIT team in Karlsruhe, to support the SpiNNaker neuromorphic platform and thus allow whole experiments to go completely on hardware, leaving only the communication and synchronization to the NRP closed loop engine.

In the next months, the NRP will support the execution of transfer function at a fixed rate, in real time, thus allowing the implementation of controllers with the well-estabilished fixed rate approach, widely used in the robotic community.

19 March 2018

NRP Workshop at ERF2018

Date: 13.03.2018
Venue: Tampere
Duration: 3 Days
We organized a Workshop titled "Beyond conventional AI – the Neurorobotics Platform" that attracted significantly more people than last year during ERF2017 in Edinburgh.
You can learn more about the ERF here: http://erf2018.eu/

16 March 2018

Handling experiment-specific python packages in the NRP

In this blog post I, [*Jacques Kaiser, FZI, Karlsruhe*] share a method to handle experiment-specific python packages. Some of my experiments require TensorFlow v1.6, some others need Keras – which itself requires a prior version of TensorFlow – how to handle all this on your locally installed NRP?
My method relies on the package virtualenvwrapper, which allows you to keep your python virtualenv in a single place.

```
pip install virtualenvwrapper --user
```

Additionally, I have a custom config which adds a virtualenv to the $PYTHONPATH when I activate it. Copy the *postactivate and postdeactivate scripts* to $WORKON_HOME – the configuration folder of virtualenvwrapper.
Now, let's say you have an NRP experiment with custom python packages listed in a *requirements.txt.* Create a virtualenv for this experiment and install the experiment-specific packages:

```
mkvirtualenv --system-site-packages my_venv
pip install -r requirements.txt
```

To access your experiment-specific packages from within the NRP, simply start the NRP from the same terminal, where the virtualenv is activated:

```
workon my_venv
cle-start
```

That's it!

13 March 2018

SP10-Workshop in Barcelona

Date: 08.03.2018
Venue: Barcelona
Duration: 2 Days
A Neurorobotics Subproject Workshop titled "Does the Brain Need a Body to be a Brain?" was held in Barcelona, Spain, on the 8th-9th March, 2018.
Find out more here: https://hbpbrainbody.wordpress.com

13 February 2018

Real Robot Control with the Neurorobotics Platform

Goals
Thanks to its architecture, the NRP should be well suited for directly controlling a real robotic platform with spiking neural networks. Indeed the closed-loop mechanism of the NRP software but also the use of ROS as a middleware enables developments in this direction.
A first motivation for such a project is to outsource the heavy computation load of simulating spiking neural networks on embedded hardware to a fixed server, which can itself interface with neuromorphic hardware like SpiNNaker if required. Consequently, it helps to reach real-time performance on small and low-energy robotic platforms where neuronal computation would have been impossible otherwise. A second motivation is the possibility to partially train the neural network in the NRP, to avoid mechanical and electrical wear of the physical robot. This, however, requires the transferability of neural control from simulation to the real robot after the training; this challenging field is more known as transfer learning and requires a minimum level of accuracy in the simulation models and equations.

Methodology
Our work is focused on real-time locomotion of a compliant quadruped robot using CPGs. To outsource the controller to the NRP as discussed above, we have designed both a robot and its 3D clone in simulation. In this setup, they both have four actuators (one for each "body-to-leg" joint) and four sensors (one for each unactuated "knee" joint). The motor position follows a simple open-loop CPG signal with the same amplitude and phase for each leg, such that the robot will alternate between standing-up and sitting-down periodically for fifty seconds. During this experiment, the sensor values are merely recorded, and not used to regulate the control signal. Given the structure of the kinematic chain with springs and

dampers in the knee joints, the system can be explicitly described with a Lagrangian equation. The latter is a function of the physical parameters of the robot which can only be evaluated with a large range of uncertainty as we work with laser-cut or 3D-printed parts assembled with a non-quantified amount of slack. However, the parameters of the simulation model can be set roughly and then optimized to maximize the similarity between the sensors signals output from the hardware and the physics engine. In this setup, we use CMA-ES for this job.

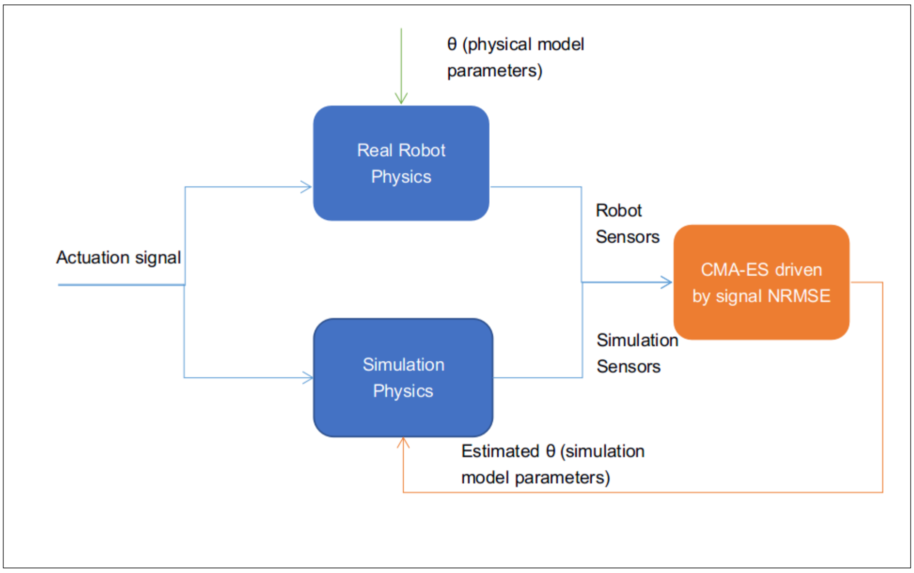

Results

After optimization, we can proceed to a qualitative visual validation using different controllers. To this goal, the NRP is installed locally on a machine connected to the same network as the robot. ROS is configured on the robot and in the NRP to enable streaming of actuation and sensing topics between the different machines. A proper calibration of the robot sensors and actuators is also needed to provide meaningful results. A NRP experiment embedding the new optimized model robot is created and used to pilot the robot and the simulated physics. In the current progress, the process seemed to give encouraging results regarding the accuracy and reliability of the simulation, but further tuning is still necessary. An illustration is presented in the following animated figure. The small delay observed between the image on the screen and the robot can be due to the NRP visual rendering in the browser or to the motor PID values.

The aim of pre-training is to exclude actuation patterns that lead to instability of the physical robot (stumbling, falling) and to tune the controller into a good regime. As an interesting fact, our first experiments indicate that there is a good correlation between failure in the simulator and failures in real observations.

13 February 2018

Structuring your local NRP experiment – some tips

Within the context of CDP4, we created a NRP experiment showcasing some functional models from SP1/4:

- A trained deep network to compute bottom-up saliency
- A saccade generation model

Since these models are generic, we want to package them so that they can easily be reused in other experiment, such as the WP10.2 strategic experiment. In this post, we quickly explain the structure of the CDP4 experiment on how modularity is achieved.

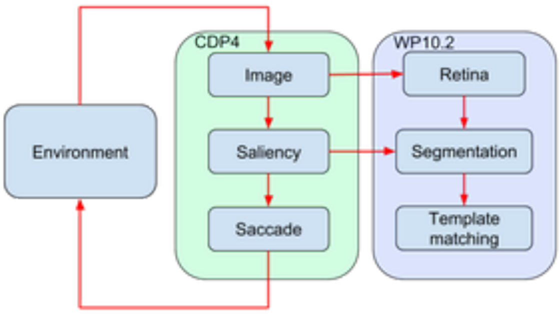

We decided to implement the functional modules from SP1/SP4 as ROS packages. Therefore, these modules can be used within the NRP (in the *GazeboRosPackages* folder), but also independently without the NRP, in any other catkin workspace. This has the advantage that the saliency model can be fed webcam images, and easily mounted on a real robot.

The main difference compared to implementing them as transfer function is synchronicity. When the user runs the saliency model on is CPU, processing a single camera image takes around 3 seconds. If the saliency model was implemented as a transfer function, the simulation would pause until the saliency output is ready. This causes the experiment to run slower but conserves reproducability. On the other hand, implemented as a ROS-node, the simulation does not wait for the saliency network to process an image, so the simulation runs faster.

The saliency model is a pre-trained deep network running on TensorFlow. The weights and topology of the network are saved in data files, loaded during the execution. Since these files are heavy and not interesting to version-control, we uploaded them on our owncloud, where they are automatically downloaded by the saliency model if not present. This also makes it simple for our collaborators in SP1/4 to provide us with new pre-trained weights/topology.

The CDP4 experiment itself has its own repo and is very lean, as it relies on these reusable modules. Additionally, an install script is provided to download the required modules in the *GazeboRosPackages*.

The topic of installing TensorFlow or other python libraries required by the CDP4 experiment, so that they do not collide with other experiment-specific libraries, will be covered in another blog post.

13 February 2018

10th Performance Show

Date: 18.01.2018
Venue: Geneva
Duration: 2 Days
The 10th HBP Neurorobotics Performance Show was held on **18 May – 19 May 2018** at the Biotech Campus in Geneva.

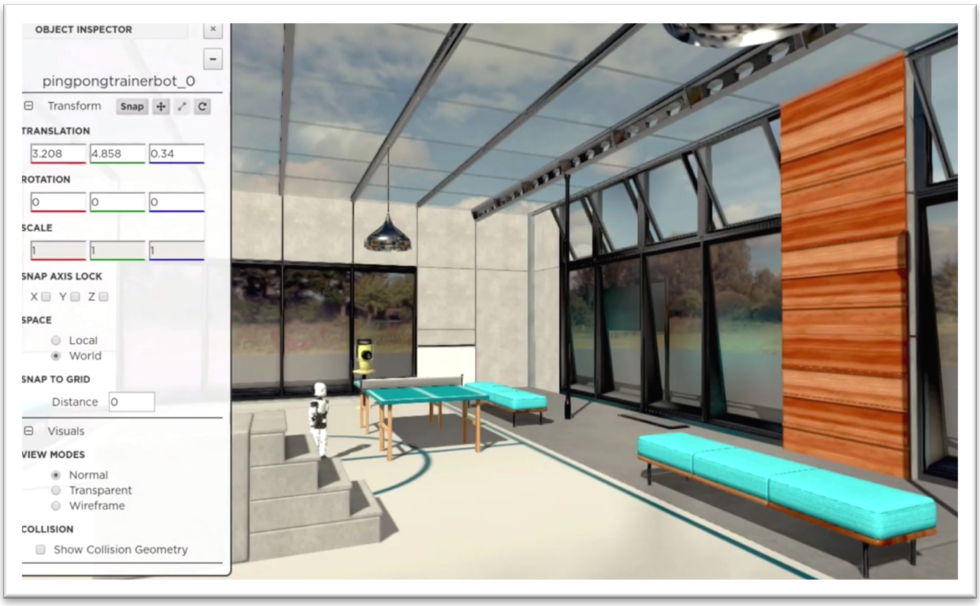

22 January 2018

Roboy Recruiting and Demo Night

When: Wednesday, 24.01 @ 5 PM
Where: TUM Entrepreneurship Center
What: Come by and get a chance to see the progress of Roboy Team and learn about the options for students to get involved with the project.

17 January 2018

24. Handelsblatt Jahrestagung Strategisches IT-Management 2018

Date: 16.01.2018
Venue: Sofitel Munich Bayerpost
Alois Knoll gave a keynote talk on neurorobotics and the Human Brain Project at the conference Handelsblatt Jahrestagung Strategisches IT-Management 2018. IT managers learned how HBP technology such as the Neurorobotics Platform can be used as tools for creating innovative brain-derived products and solutions. In the break, participants had the chance to meet our robot mouse.
http://veranstaltungen.handelsblatt.com/it-jahrestagung/

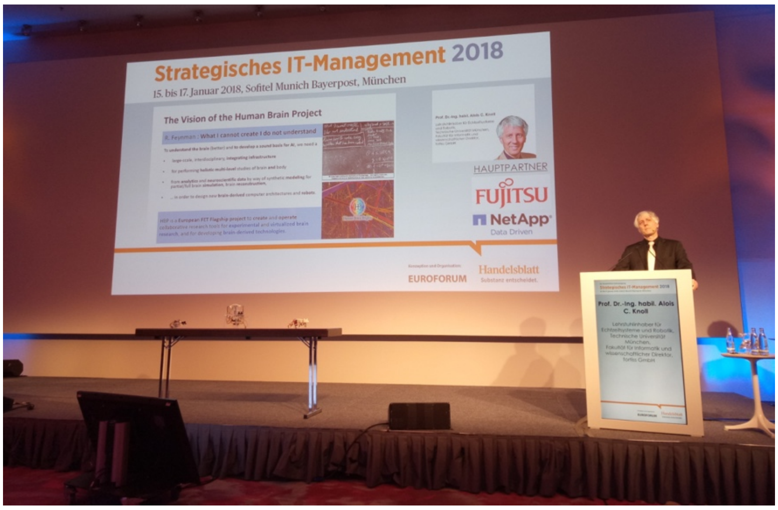

16 January 2018

SP10 at the interdisciplinary college

We present a poster and host an interactive demo session during this year's annual interdisciplinary college (IK), which takes place from 9-16th March.

Find out more here: https://www.interdisciplinary-college.de

16 January 2018

Implementing cerebellar learning rules for NEST simulator

The cerebellum is a relatively small center in the nervous system that accounts around half of the existing neurons. As we previously documented, research from the University of Granada are taking advantage of the NeuroRobotics Platform (NRP) in order to prove how cerebellar plasticity may contribute to vestibule-ocular reflex (VOR) adaptation.
Implementing neurorobotic experiments often requires some multidisciplinary efforts as:

1. Establishing a neuroscience-relevant working hypothesis.
2. Implementing an avatar or robot simulator to perform the task.
3. Developing the brain model with the indicated level of detail.
4. Transforming brain activity -spikes- into signals that can be used by the robot and viceversa.

The NRP provides useful tools in order to facilitate most of these steps. However, the definition of complex brain models might requires the implementation of neuron and synapsis models for the brain simulation platform (NEST in our particular case). The cerebellar models that we are including involves plasticity at two different synaptic sites: the parallel fibers (PF) and the mossy fibers (MF, targeting the vestibular nuclei neurons).

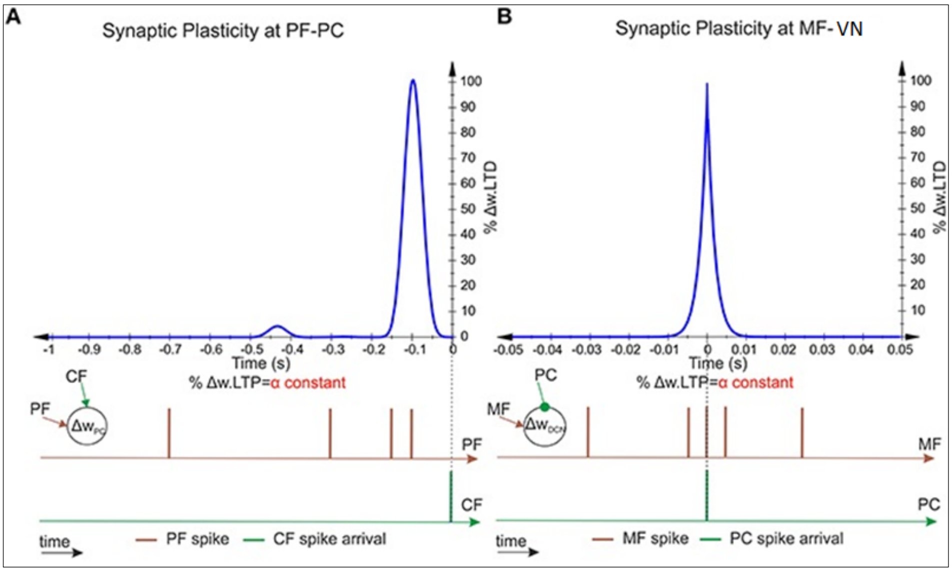

Although we will go deeper into the equations (see the reference above for further details) each parallel fiber synapsis will be depressed (LTD) when a presynaptic spike occurs closely to the occurrence of a complex spike of the target Purkinje cell (PC, see figure). Similarly, the plasticity at the mossy fiber/vestibular nuclei (VN) synapsis will be driven by the inhibitory activity coming from the Purkinje neurons.
These learning rules have been previously implemented for EDLUT simulator and used for complex manipulation tasks in [1]. The neuron and synapsis models have been released in GitHub and also as part of the NRP source code. This work in the framework of the HBP will allow researchers to demonstrate the role that plasticity at the parallel fibers and mossy fibers play in vestibule-occular reflex movements.

[1] Luque, N. R., Garrido, J. A., Naveros, F., Carrillo, R. R., D'Angelo, E., & Ros, E. (2016). Distributed cerebellar motor learning: a spike-timing-dependent plasticity model. Frontiers in computational neuroscience, 10.

16 January 2018

Upcoming Hack Roboy #4

The next Hack Roboy will take place from 13.04-15.04.
Find out more here: https://www.hackroboy.com/home

16 January 2018

Upcoming Nerd Nite in Munich

During this event Florian Röhrbein will give a talk on " Florian Röhrbein: Neurorobotik – Über künstliche Dummheit und natürliche Intelligenz", which translates to "Neurorobotics – of artificial stupidity and natural intelligence"
Find out more here: https://de-de.facebook.com/events/572210989783890/

16 January 2018

Using the NRP: a saliency computation model drives visual segmentation in the Laminart model

Recently, a cortical model for visual grouping and segmentation (the Laminart model) has been integrated to the NRP. From there, the goal was to build a whole visual system on the NRP, connecting many models for different functions of vision (retina processing, saliency computation, saccades generation, predictive coding, ...) in a single virtual experiment, including the Laminart as a model for early visual processing. While this process is on-the-go (see here), some scientifically relevant progress already arose from the premises of this implementation. This is what is going to be explained right now.

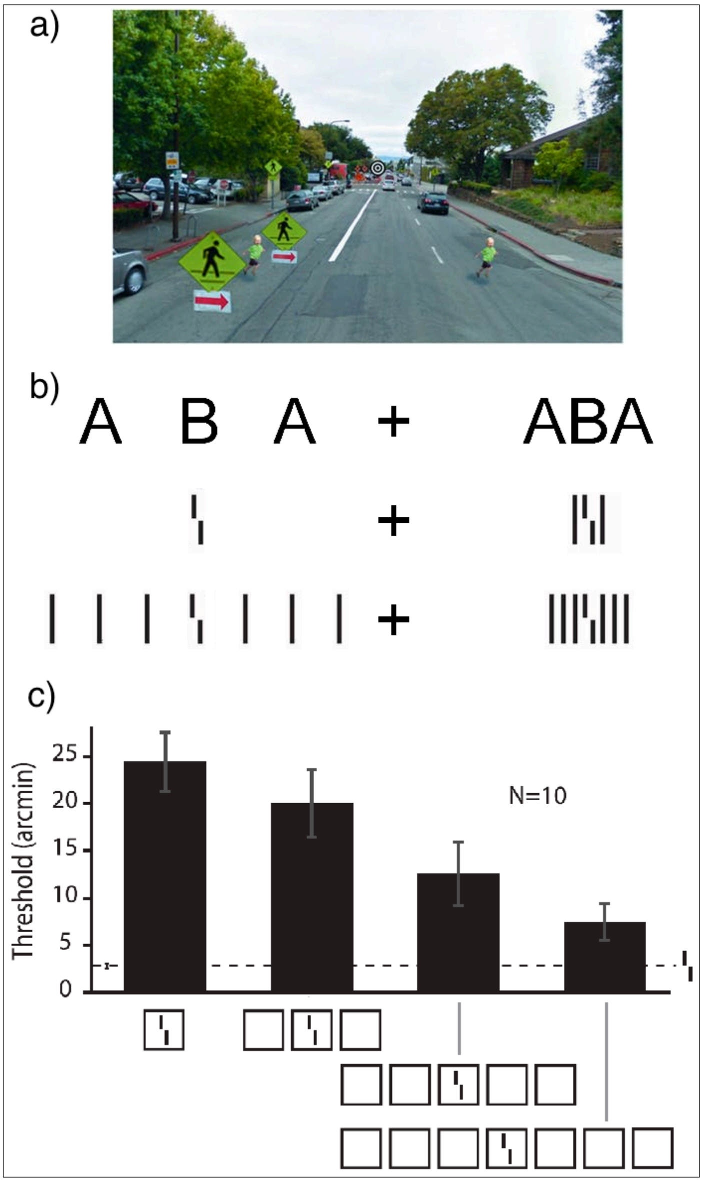

The Laminart is one of the only models being able to satisfactorily explain how *crowding* occurs in the visual system. Crowding is a visual phenomenon that happens when perception of a target deteriorates in the presence of nearby elements. Crowding occurs in real life (for example when driving in the street, see fig. 1a) and is widely studied in many psychophysical experiments (see fig. 1b). Crowding happens ubiquitously in the visual system and must thus be accounted for by any complete model of the visual system.

While crowding was for a long time believed to be driven by *local* interactions in the brain (e.g. decremental feed-forward pooling of different receptive fields along the hierarchy of the visual system, jumbling the target's visual features with the one from nearby elements), it recently appeared that adding remote contextual elements can still modulate crowding (see fig. 1c). The *entire* visual configuration is eligible to determine what happens at the very *tiny* scale of the target!

Fig. 1: a) Crowding in real life. If you look at the bull's eye, the kid on the right will be easily identifiable. However, the one on the left will be harder to identify, because the nearby elements have similar features (yellow color, human shape). b) Crowding in psychophysical experiments. Top: the goal is to identify the letter in the center, while looking at the fixation cross. Neighbouring letters make the task more difficult, especially if they are very close to the target. Center and bottom: the goal here is to identify the offset of the target (small tilted lines). Again, the neighbouring lines make the task more difficult. c) The task is the same as before (visual stimuli on the x-axis are presented in the periphery of the visual field and observer must report the offset of the target). This time, squares try to decrease performance. What is plotted on the y-axis is the target offset at which observers give 75% of correct answers (low values indicate good performance). When the target is alone (dashed line), performance is very good. When only one square flanks the target, performance decreases dramatically. However, when more squares are added, the task becomes easier and easier.

To account for this exciting phenomenon (named *uncrowding*), Francis et al. (2017) proposed a model that parses the visual stimulus in several groups, using low-level, cortical dynamics (arising from a biologically plausible and laminarly structured network of spiking neurons, with fixed connectivity). Crucially, the Laminart is a 2-stage model in which the input image is segmented in different groups *before* any decremental interaction can happen between the target and nearby elements. In other words: how elements are grouped in the visual field determines how crowding occurs, making the latter a simple and behaviourally measurable phenomenon that unambiguously describes a central feature of human vision (grouping). In fig. 1c (right), the 7 squares form a group that frames the target, instead of interfering with it, hence enhancing performance. In the Laminart model, the 7 squares are grouped together by illusory contours and are segmented out, leaving a privileged access to the target left alone. However, in order to work, the Laminart model needs to start the segmentation spreading process somewhere (see fig. 2).

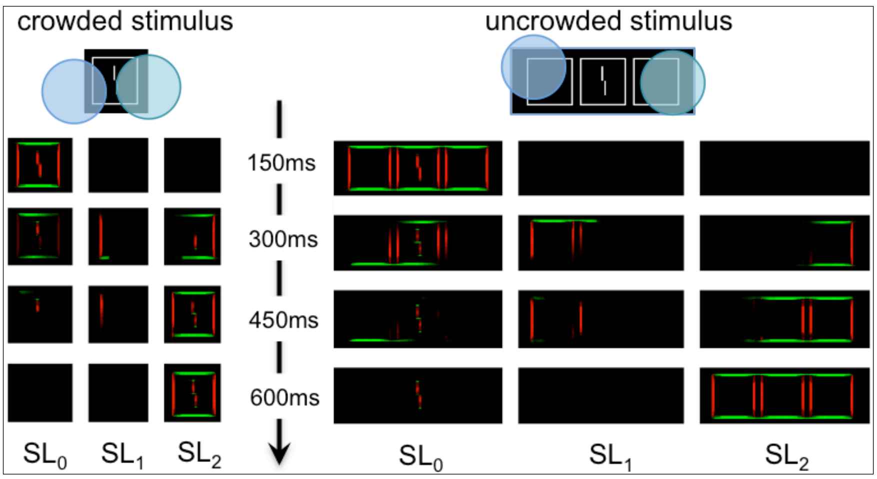

Fig. 2: Dynamics of the layer 2/3 of the area V2 of the Laminart model, for two different stimuli. The red/green lines correspond to the activity of the neurons that detect a

Co-funded by the European Union

vertical/horizontal contrast contour. The three different columns for each stimulus correspond to three segmentation layers, where the visual stimulus is parsed in different groups. The blue circles are spreading signals that start the segmentation process (one for each segmentation layer that is not SL0). Left: the flanker is close to the target. It is thus hard for the spreading signals to segment the flanking square from the target. Right: the flankers extend further and are linked by illusory contours. It is more easy for the the signals to segment them from the target. Thus, this condition produces less crowding than the other.

Up to now, the model was sending ad-hoc top-down signals, lacking an explicit process to generate them. Here, using the NRP, we could easily connect it to a model for saliency computation that was just integrated to the platform. Feeding the Laminart, the saliency computation model delivers its output as a bottom-up influence to where *segmentation signals* should arise. On the NRP, we created the stimuli appearing in the experimental results shown in fig. 1c, and presented them to the iCub robot. In this experiment, each time a segmentation signal is sent, its location is sampled from the saliency computation map, linking both models in an elegant manner. Notably, when only 1 squares flanks the target, the saliency map is merely a big blob around the target, whereas when 7 squares flank the target, the saliency map is more peaky around the outer squares (see fig. 3). Consequently, the more squares there are, the more probable it is that the segmentation signals succeed in creating 2 groups from the flankers and the target, releasing the target from crowding. This fits very well with the results of figure 1c. The next step for this project is to reproduce the results quantitatively on the NRP.

Fig. 3: Coexistence of the Laminart network and the saliency network. Top: crowded condition. Bottom: uncrowded condition. In both situations, the saliency computation model drives the location of the segmentation signals in the Laminart model and explains very well how crowding and uncrowding can occur. The windows on the left display the saliency model. The ones on the right display the output of the Laminart model (up: V2 contrast borders activity ; down: V4 surface activity).

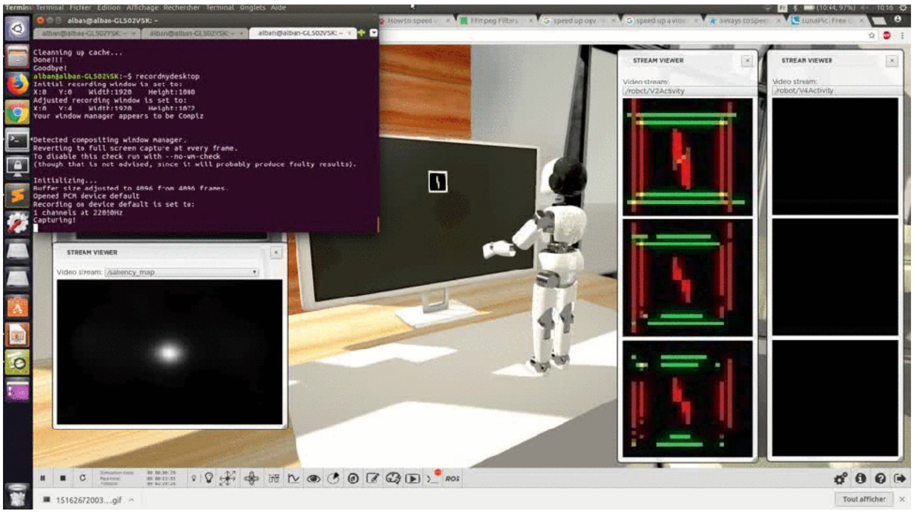

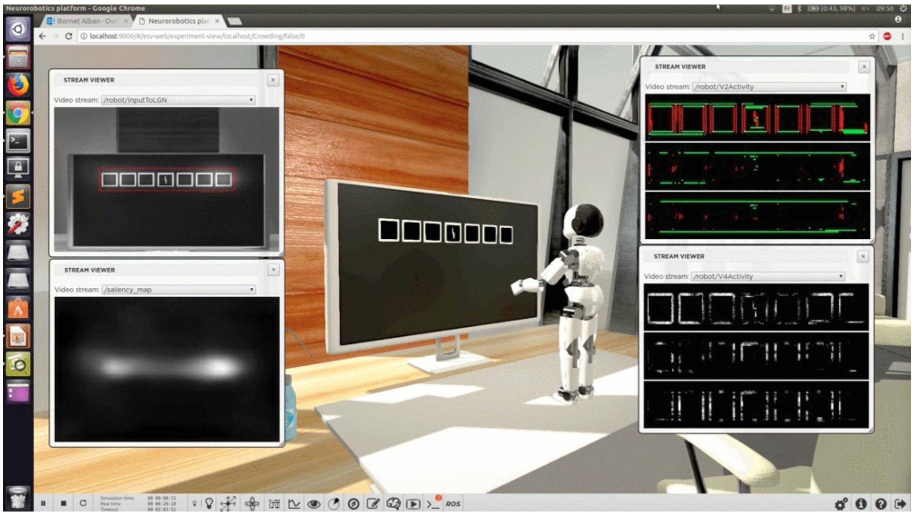

To sum up, building a visual system on the NRP, we could easily make the connection between a saliency computation model and our Laminart model. This connection greatly enhanced the latter model and gives it the opportunity to explain very well how uncrowding occurs in human vision and the low-level mechanisms of visual grouping. In the future, we will run psychophysical experiments in our lab, where it is possible to disentangle top-down from bottom-up influence on uncrowding, seeing whether a strong influence of saliency computation on visual grouping makes any sense.

15 January 2018

Fable robot simulator

Fable is a 2 DoF modular robot arm that is being used by the group of DTU in order to develop the task of "Self-Adaptation in Modular Robotics".

Thanks to the modularity provided by Fable, it is feasible to combine several modules together in order to create different robotic configurations increasing the complexity of the system. In this way, one is able to work on manipulation tasks as well as in locomotion tasks just by plugging a few modules together to form an arm, a worm, a spider,…

In the process to make the Fable robot as accessible as possible to the community, here at DTU we have been working on the implementation of the Fable v2.0 simulator.

We have created 3 different configurations:

A simple robotic arm, 2 DoF (1 Fable module)

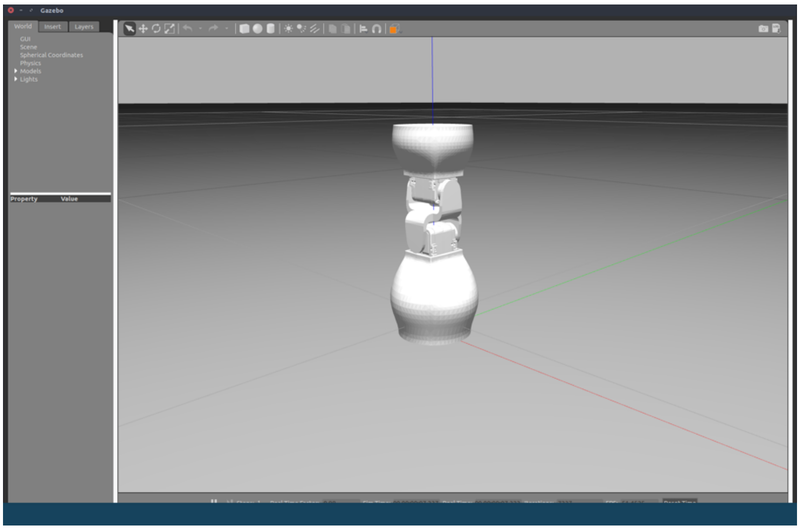

A worm-like robot, 4 DoF (2 Fable module)

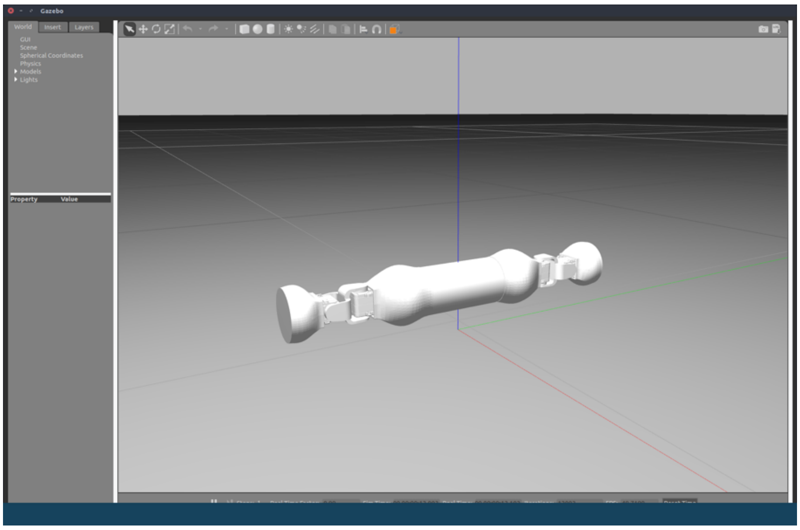

A quadruped-like robot, 8 DoF (4 Fable module)

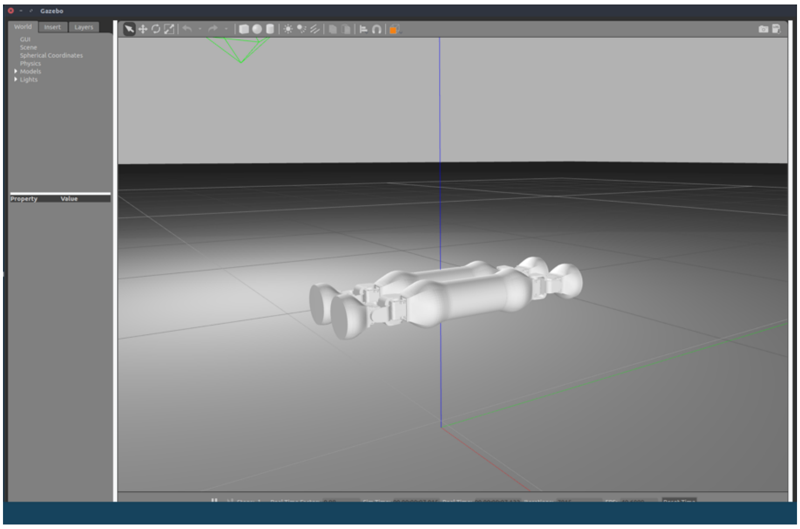

This robot model has not been included to the NRP yet, but soon will be available for users. We will keep you updated.

12 January 2018

SP10 mouse in german television!

Our robotic mouse (below) was featured in a german talkshow! Florian Röhrbein presented our mouse and talked about neurorobotics live on german television.
You can find part of this segment here:
http://www.ardmediathek.de/tv/Morgenmagazin/Künstliche-Intelligenz-Längst-keine-Zuk/Das-Erste/Video?bcastId=435054&documentId=49037372

08 January 2018

Upcoming Neurorobotics Lectures

Florian Röhrbein from TUM will give two lectures on Neurorobotics in general and in the HBP in the BCCN Lecture Series. Details can be found here: http://bccn-munich.de/teaching/computational-neuroscience

08 January 2018

Learning to walk

Favorit Bar + Kunstverein München, in collaboration with the TU München and the Human Brain Project, present:
Learning to Walk: Künstliche Intelligenz und Körper

Performance and grand finale at Kunstverein München: *Learning to Walk: 42nd Street of 2018* takes place on 26th January.

Find out more here: http://www.kunstverein-muenchen.de/en/program/events/2017/learning-to-walk
and here https://www.facebook.com/events/168991993875889/?active_tab=about

08 January 2018

Upcoming Performance Show in Geneva

The next HBP SP10 Performance Show will take place in Geneva on 18th and 19th January at the EPFL Biotech Campus.

08 January 2018

Plenary and Panel Discussion (Prof. Knoll)

Prof. Knoll (TUM, SP10) will participate in a plenary and a panel discussion on 16th January in Munich during a handelsblatt conference.
Learn more here: http://veranstaltungen.handelsblatt.com/it-jahrestagung/jahrestagung-2018/programm-2018/