SMART MIRROR

VEDLIoT
Very Efficient Deep Learning in IoT

AN INTUITIVE INTERACTION INTERFACE

OVERVIEW

The smart mirror was developed as an intuitive interface to assist the interaction in smart home environments with a strong focus on local processing for data privacy.
It shows a mirror image of the user and displays personalized information and the status of the smart home. Based on the virtual mirror image generated from depth imaging cameras, the user's face, basic objects, and simple hand gestures are recognized and used to control the mirror. Furthermore, a voice recognition system supported by natural language processing (NLP) allows for voice-assisted interaction. All computations are performed locally on the device using open-source software, ensuring maximum privacy as no data is transferred to the cloud or any third-party service providers.
A significant challenge in this project is combining multiple machine-learning techniques. The methodology developed within the VEDLIoT project helps to keep up with high-performance while, at the same time, maintaining a low energy consumption. The newest security techniques ensure meeting the high privacy requirements.

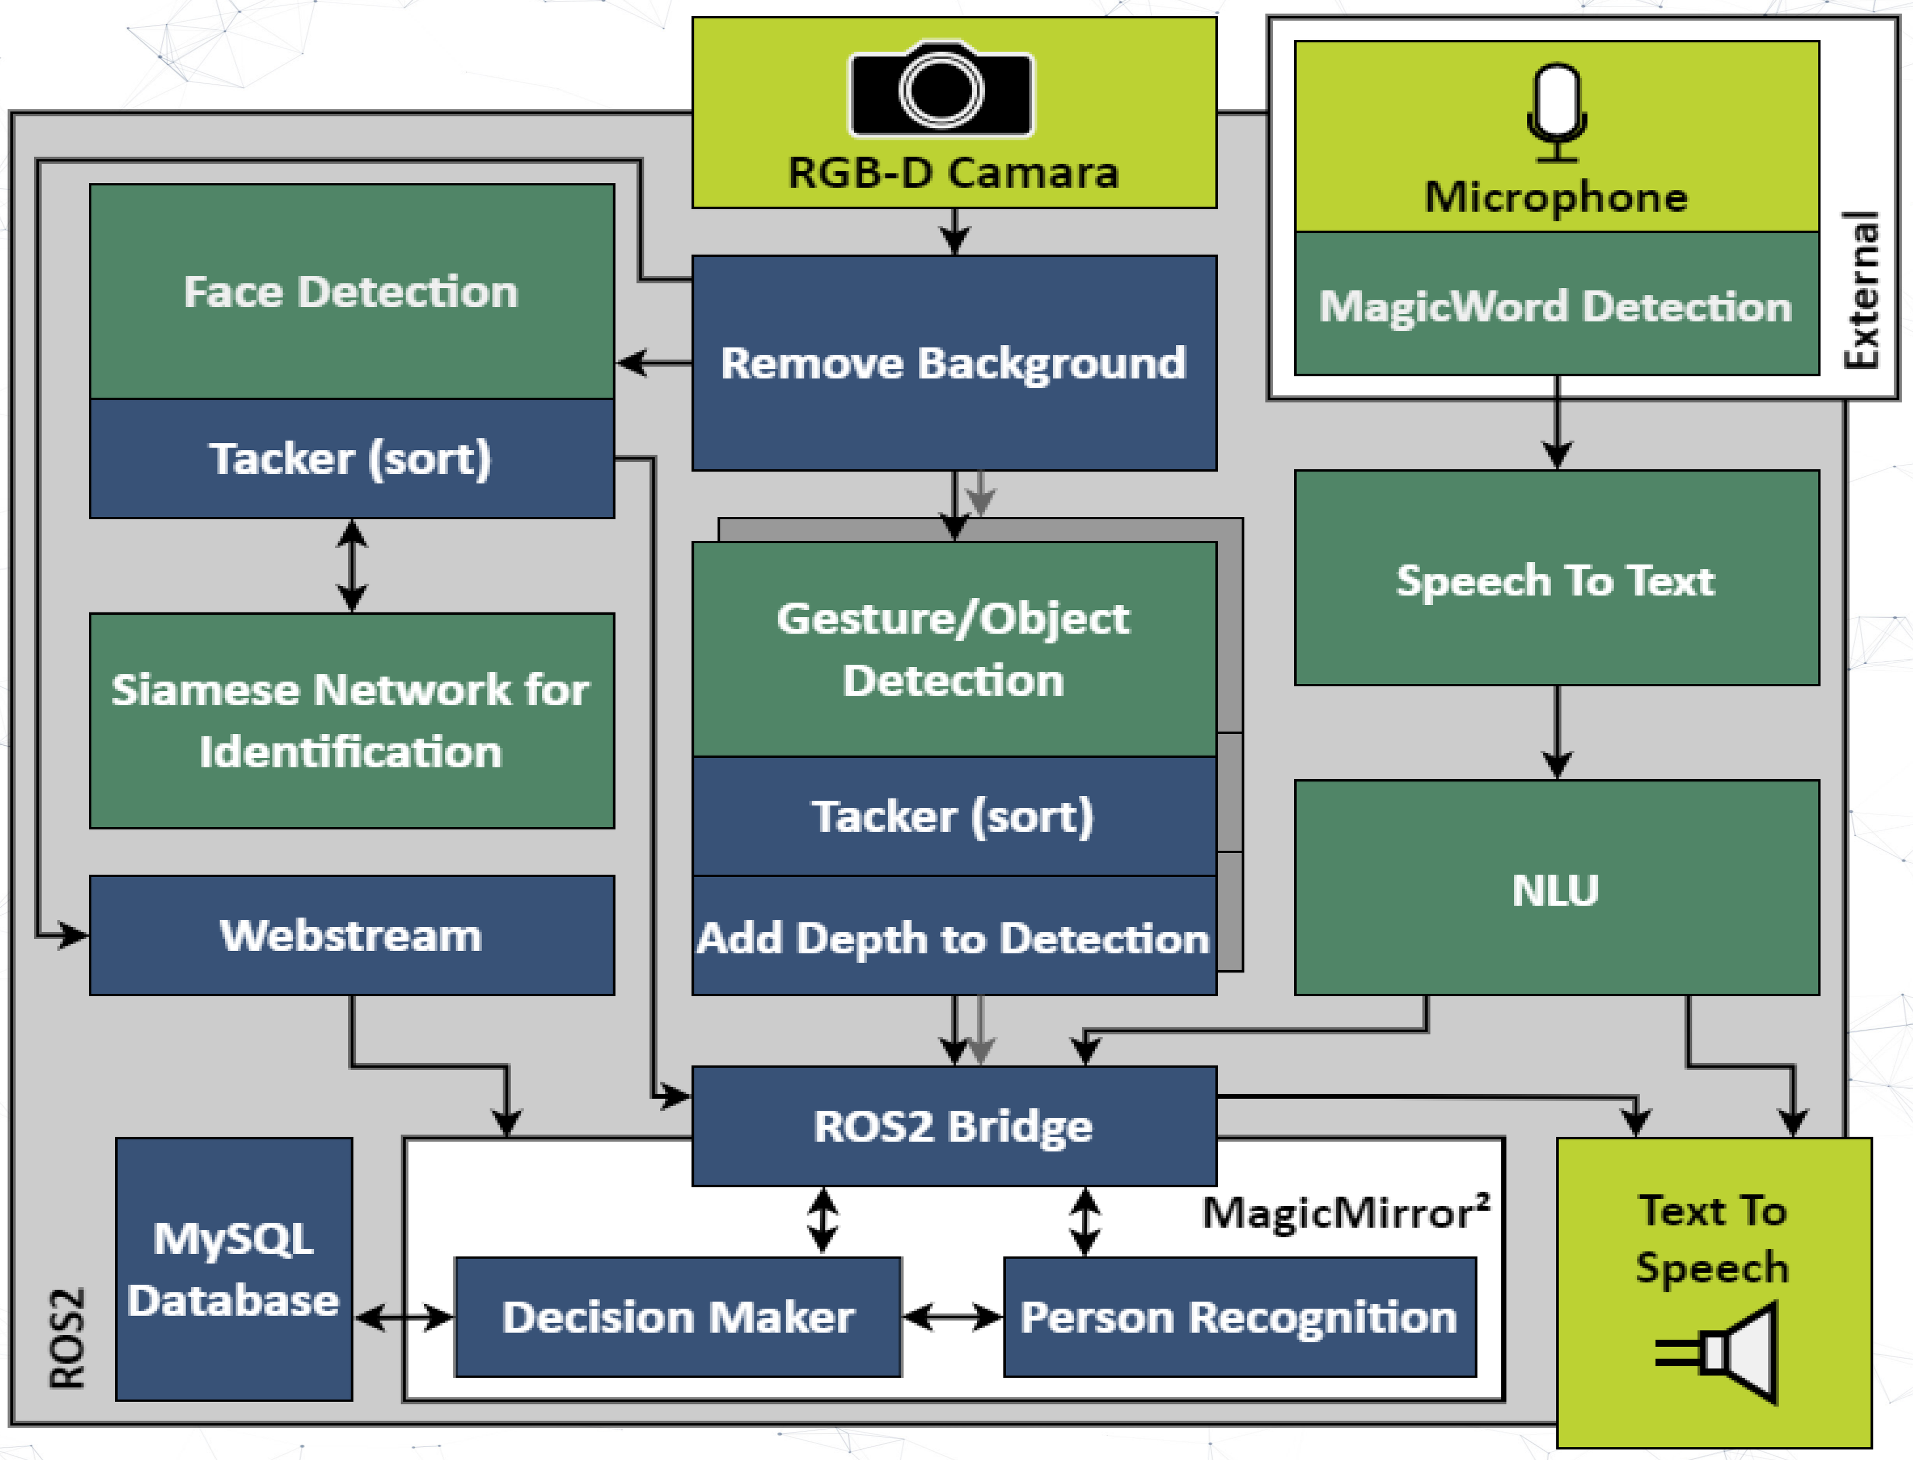

ACCELERATED MACHINE LEARNING

Multiple neural networks are combined via the common middleware ROS2 to perform all the necessary detections. Using ROS2 in this context allows for effortless interchangeability and varying the distribution of computation tasks. Every subsystem (e.g., for object- or gesture-detection) can be calculated in parallel on the heterogeneous hardware developed within this project. At the start of the project, the neuronal networks were calculated on two NVIDIA Jetson AGX Xavier microservers, interconnected on a t.RECS. This setup was moved to an NVIDIA Jetson AGX Orin in later stages, as it offers the same computational power while consuming less energy.
Additionally, the Orin was combined with a Hailo-8 AI accelerator to optimize energy efficiency further. In this setup, the Hailo-8 is responsible for object detection using YOLOv7. At the same time, the Orin calculates the remaining neural networks, e.g., YOLOv7 for gesture recognition, a feature extractor, and a Siamese network for face recognition. The current setup achieves a performance of 30 FPS with a power consumption of about 38 Watts, a significant improvement over the starting performance of 16 FPS and 150 Watt. The next target for this system is the u.RECS, which is combining an Orin NX and other accelerators. Therefore, the compute-intensive parts are optimized using the EmbeDL toolchain. The YOLOv7 models for object and gesture recognition are pruned, reducing the size needed GFLOPs in half. With this, a further reduction to 28 Watt is expected and partially verified.

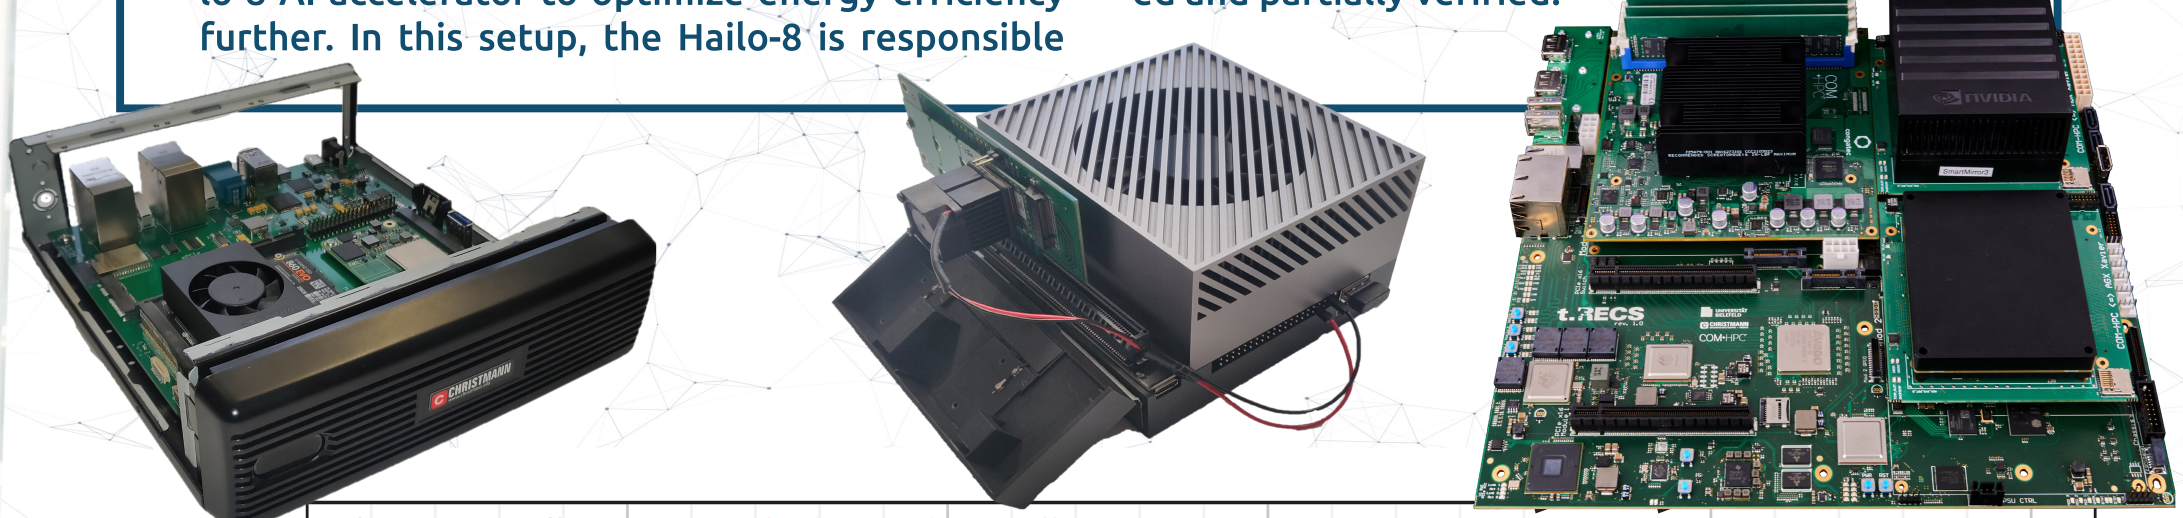

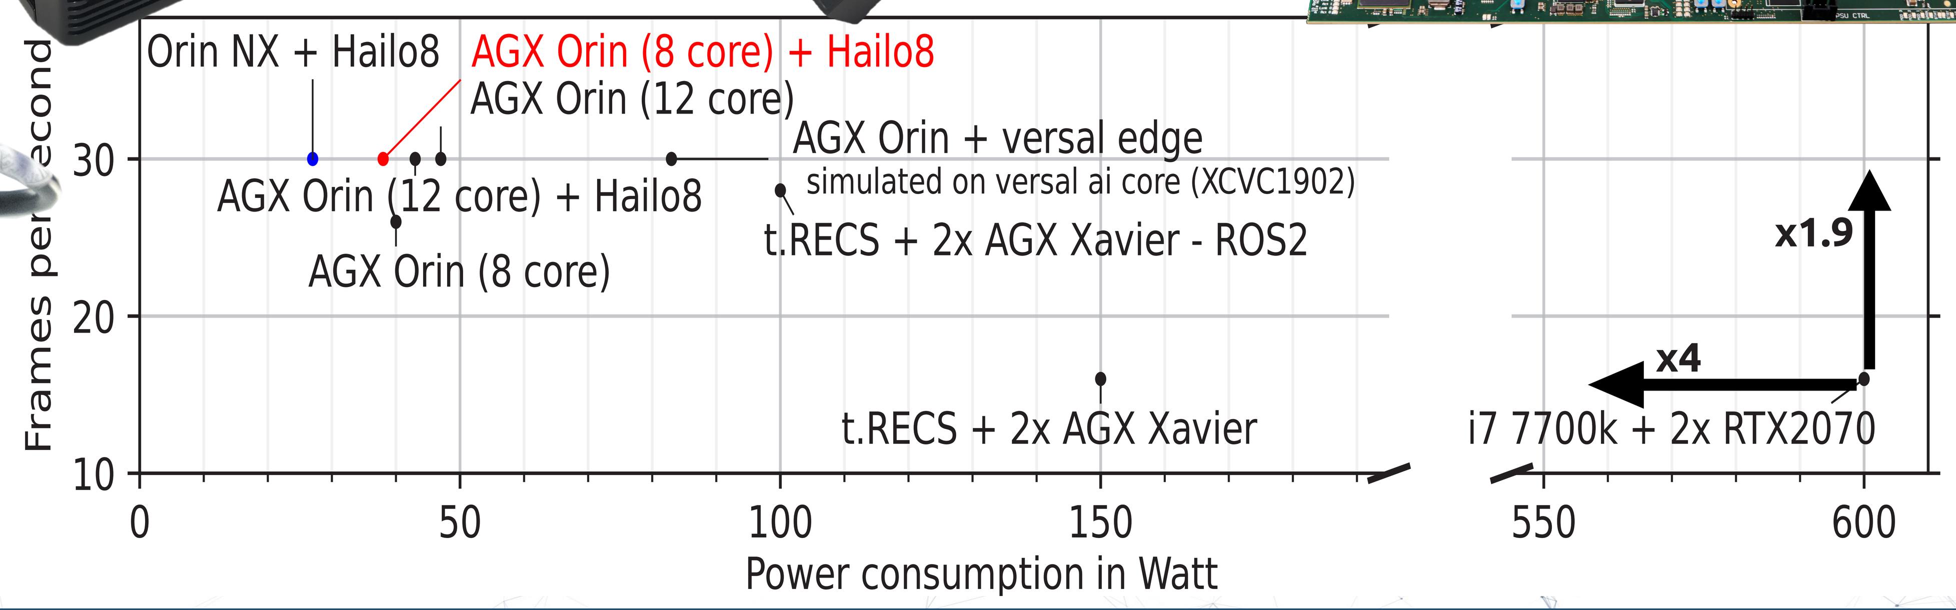

CONTACT

Nils Kucza, nkucza@techfak.uni-bielefeld.de
Bielefeld University, Germany
Cognitronics & Sensor Systems
https://cit-ec.de/en/

@VEDLIoT
VEDLIoT Project
VEDLIoT EU Project

PARTNERS

Universität Bielefeld, Christmann, veoneer, antmicro, SIEMENS, embedl, Ciências ULisboa, FCiências.ID, Universität Osnabrück, RISE Research Institutes of Sweden, unine Université de Neuchâtel, University of Gothenburg, Chalmers

Funded by the European Union
Grant No 957197

www.vedliot.eu