ercim-news.ercim.eu
Number 132 January 2023

ERCIM NEWS

Special theme:

Cognitive AI & Cobots

Also in this issue

Research and Innovation: The Interactive Classification System

Editorial Information

ERCIM News is the magazine of ERCIM. Published quarterly, it reports on joint actions of the ERCIM partners, and aims to reflect the contribution made by ERCIM to the European Community in Information Technology and Applied Mathematics. Through short articles and news items, it provides a forum for the exchange of information between the institutes and also with the wider scientific community. This issue has a circulation of about 2,000 printed copies and is also available online ,at https://ercim-news@ercim.eu.

ERCIM News is published by ERCIM EEIG
BP 93, F-06902 Sophia Antipolis Cedex, France
+33 4 9238 5010, contact@ercim.eu
Director: Philipp Hoschka, ISSN 0926-4981

Contributions
Contributions should be submitted to the local editor of your country

ERCIM News online edition: *https://ercim-news.ercim.eu/*

Next issue:
April 2023: Data infrastructures and management

Subscription
Subscribe to ERCIM News by sending an email to en-subscriptions@ercim.eu

Cover illustration drawn by Stefanos Poirazoglou.

JOINT ERCIM ACTIONS

SPECIAL THEME

This special theme Cognitive AI & Cobots was coordinated by the guest editors Theodore Patkos (ICS-FORTH) and Zsolt Viharos (SZTAKI)

RESARCH AND INNOVATION

ANNOUNCEMENTS / IN BRIEF

Forum on Digital Ethics in Research: Beyond Compliance

by Anaëlle Martin (CNPEN/CCNE)

The ERCIM Ethics Working Group organised jointly with Inria (the French National Institute for Research in Digital Science and Technology) and the CNPEN (the French National Pilot Committee for Digital Ethics) a forum targeted at scientists and Research Ethics Boards on 17-18 October 2022 in Paris.

This successful event, which brought together researchers interested in ethics of digital and connected technologies, consisted of keynotes, presentations and tutorials, and provided ample space for interactive sessions and open discussions. The participants from all over the world, and from different fields of research (computer science, robotics, philosophy of technology, research integrity, ethics and law) met in Paris and online on the 17th and 18th of October to discuss ethical challenges resulting from interactions between humans and technologies in fields related to research involving ICT. The aim of this workshop was to provide a forum for researchers in digital sciences who face tough ethical questions in their daily activity for which there are not yet consensual answers among the research community. Research in digital sciences raises many new ethical issues such as the use of "autonomous" robots (human responsibility, agency and oversight), the use of data (which raises privacy concerns), the negative impacts of AI systems in terms of discrimination, bias, lack of robustness and reliability, and also the transformative impact of disruptive technologies on society and work. In order to achieve scientific excellence and public trust, a responsible conduct of research in science and engineering is needed. Although some efforts have been made, the situation is far from satisfactory, and more progress needs to be made. On the one hand, it is important to train and raise awareness among researchers, as many are either ill-informed or reluctant to deal with potential ethical implications of their work. On the other hand, it is crucial to consider the current research framework that seems ill-suited to address all the risks and acute concerns about research ethics and research integrity.

Claude Kirchner (top left), Catherine Tessier (top right) and Andreas Rauber (bottom right) chairing the keynotes and sessions; Björn Levin (bottom left) addressing the audience.

Ethical issues arise from the intersections of business, technology, and information such that the whole digital ecosystem must be taken into account, including stakeholders such as entrepreneurs, users and scientists from fields other than digital sciences, applied mathematics and computer science. In addition, there must be appropriate trade-offs between innovation, ethics, regulation and compliance, since the current digital model is no longer sustainable. Last but not least, if we do not promote "digital sovereignty" in Europe and an operational agenda, there will be no alternatives to the GAFAM or Chinese models. Therefore, agenda for digital sovereignty and agenda for digital ethics should be deeply connected. To meet this ambitious program, the seminar consisted of six sessions devoted to research ethics from the perspective of cross-disciplinary and global setting, review, training, big data and AI, new challenges and opportunities, as well as a tutorial on research ethics for young researchers. Over two days, about 15 researchers spoke on issues relating to governance of AI, research ethics guidelines for the computer sciences, responsible innovation in a neuroinformatics project, ethics reviews in modern research, challenges of ethics oversight in digital science, the limits of human subjects review, research accountability in machine learning, ethics reflection as a precondition to research funding, teaching and research ethics training, open science, data trusts and the search for anonymous data. Each of the speakers made a high-quality contribution based on their academic work and personal experience in the field of research ethics or research integrity (as a researcher and/or teacher). Many suggestions were made and some outcomes may be directed towards policy makers. In addition to dealing with core issues such as human rights violation (but also environment well-being and animal welfare in the case of bee-robots) the workshop addressed epistemological and methodological issues. The tutorial addressed the distinction between morality, ethics, deontology and law. The exchanges with the audience allowed to deepen other conceptual distinctions, such as research ethics and research integrity, but also scientific soundness. Indeed, some highly publicised studies, often pointed out by the speakers, are questioning the boundary between the scientific quality of research and its ethical impact. One could almost argue that when we have some ethics research controversy, it is often based on "bad science", which is likely to undermine the public trust in science.

Despite its success, this seminar is not sufficient to fill the gaps in research ethics. There is a need in Europe for more events of this kind focusing not just on digital ethics but on research ethics for digital scientists. The richness of the contributions of that forum on digital ethics was due to their interdisciplinarity and to the international backgrounds of the speakers. We therefore need a series of other workshops of the same type dedicated to research ethics, especially of cross-disciplinary research like digital humanities research. To conclude, it is worth highlighting a message underpinning many of the speeches throughout the two days of the conference: Europe, and more specifically the European Union, has a duty to provide answers to the many questions raised by researchers and scientists. It needs to adopt a doctrine in line with its values and its vision of human rights. If not, European states will be condemned to follow the path decided by the United States or by other public or private actors.

More information:
https://www.ercim.eu/beyond-compliance

The recorded presentations are available on a Youtube channel at https://kwz.me/hwe

ERCIM Fellows Community Event 2022

by Emma Lière (ERCIM Office)

ERCIM organized an online community event for its fellows and guests for the second time on 4 November 2022.

Due to high demand and real interest, ERCIM decided to renew this event, which was originally organised during the pandemic in 2021. The main objective remained the same than for the first community event, namely to promote and improve the interaction between fellows. For this second edition, not only fellows but also their coordinators and representatives of ERCIM members were invited. A total of 59 participants met online to exchange ideas and get to know each other.

Several constraints were taken into account when designing this event. One was the diversity of the audience, which consisted of postdocs from all over the world who were not at the same stage of their fellowship. Another was the duration, which should not be too long for an online event, but allowed for the presentation of 36 posters. In addition, this event should be free of charge for all participants. By using the innovative online platform “Gather Town” all these parmeters could be taken into account.

Four virtual rooms were set up for the heart of the event, the poster sessions. These sessions allowed to showcase the remarkable scientific work of ERCIM postdocs and help them find common points of interest. Three host institutes (Fraunhofer-Gesellschaft, VTT and Inria) had a stand where they could also present their organisations. In addition, Professor Keith G Jeffery gave a very interesting talk on "Open Access".

The results of this community event far exceeded expectations. An anonymous post-event survey showed that for all participants who responded, the programme met their expectations, the event was useful and new contacts were established.

TESTIMONIALS

“The ERCIM Fellowship Programme Community Event exceeded my expectations. It was the second time I participated in this event. Not only did I meet new people, but I also received research assistance with new insights. Emma and Catherine did a great job in bringing all ERCIM community members together on the Gather Town platform. I would enthusiastically recommend this event in the future.”
Shipra Singh,
ERCIM Fellow at NTNU

“As a participant in the ERCIM Fellowship Programme Community Event 2022, I gained very useful insights from the keynote sessions, while meeting other ERCIM fellows was also very interesting, especially because it provided a platform for visibility of ERCIM fellows among fellows as well as a networking opportunity for further interactions and possible collaborations. Finally, the Gather Town app platform provided an eco-friendly virtual space and made the interaction during the poster sessions very intentional and engaging.”
Aderonke Favour-Bethy Thompson

“Many thanks for providing such an excellent and interactive online platform - 'Gather Town' - to meet other ERCIM fellows and exchange ideas. Overall, it was a fruitful and fun way to share research ideas with other ERCIM fellows. The ERCIM fellowship is a great opportunity to engage in cutting-edge research and build a network with internationally renowned researchers. My sincere thanks go to the ERCIM staff, especially Emma and Catherine, who are excellent administrators, unpretentious and always helpful. Once again, thank you very much!”
Nitin Muchhal

The event was organized by Emma Lière, coordinator of the ERCIM scholarship programme, with the support of Catherine Riou and moderated by Monica Divitini, NTNU, Chair of the ERCIM Human Capital Task Group.

http://fellowship.ercim.eu

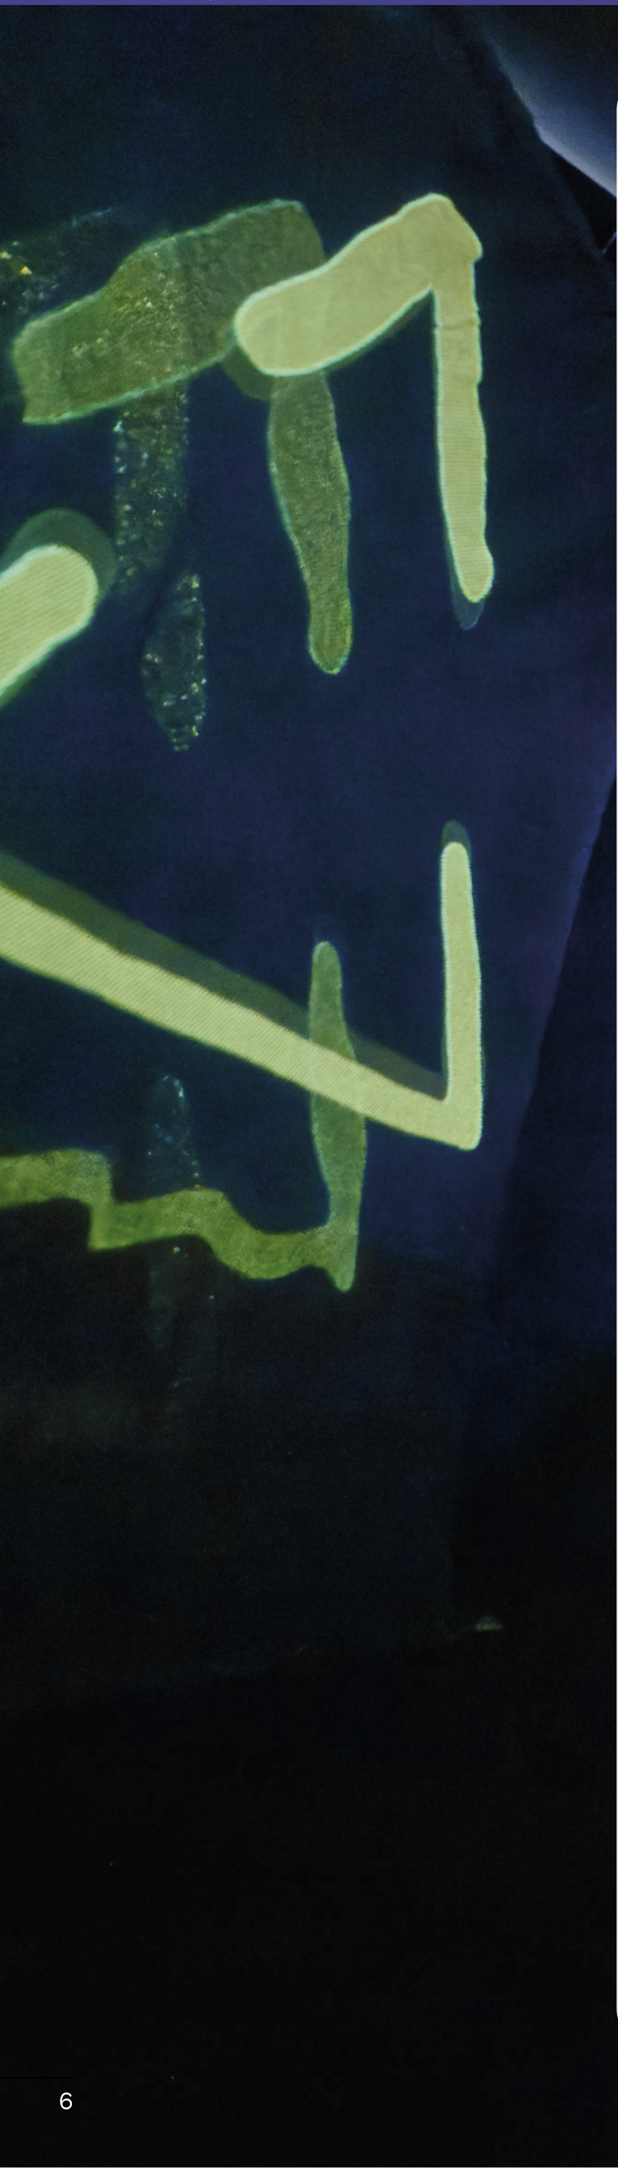

Introduction to the Special Theme

Cognitive AI & Cobots

by Theodore Patkos (ICS-FORTH) and Zsolt Viharos (SZTAKI)

The wave of popularity of modern Artificial Intelligence (AI) systems is creating well-justified expectations that its application to diverse domains will lead to even bigger advancements in the future. At the same time, there is an open debate in the research community regarding the limitations of existing methods, and whether the current success can scale up to a broader spectrum of problems than those current AI is focusing on. Especially considering that modern intelligent systems are affecting our everyday lives at an increasing pace, the request for future intelligent machines is to exhibit capabilities that are not only effective, but also closer to human intuition and intellect.

In order achieve a reliable human-machine symbiotic collaboration, AI needs to make progress on skills that humans excel in. These skills include: a general understanding of how the world works; the exploitation of common sense knowledge that is hidden, yet pervasive, in the majority of human-to-human interactions; the ability to understand other people's intentions and to comply with social norms and values, or the ability to explain with grounded justification their decisions. The explainability of AI models is one of the crucial challenges in today's machine learning solutions, in order to help designers and users understand those more-often-experienced cases, when AI systems significantly outperform human capabilities.

Cognitive AI intends to augment intelligent machines with human-like cognition. It builds on the advancements of modern, data-driven AI technologies, but also calls for progress in symbolic, knowledge-based methods, in order to enable machines to learn from how humans create rich cognitive models about the world they live in, and how they ascribe mental states to themselves and others, such as beliefs, intentions, emotions, perceptions, even thoughts. Cognitive AI will also help intelligent systems engage more smoothly in social interactions, accomplish collaborative tasks, and in general broaden their intelligence and the spectrum of problems they can tackle.

The field of collaborative robots (cobots) is a characteristic case, where Cognitive AI is expected to achieve broad impact. Cobots are receiving continuously increasing roles and benefits in social life and in the industrial sectors as well. This technical equipment shares the field of activities with human beings in space and/or time and, building on tight human-robot collaboration, it results in additional value-added processes beyond the summarised outputs of these two separate contributors. Coexistence, cooperation, sequential and responsive collaboration form the different levels of the common activities, in each of which their cross-understanding and harmonisation is an essential feature to be realised and exploited, giving rise to various novel challenges for science.

As such, Cognitive AI lies at the intersection of both Computer and Cognitive Sciences, crossing through a broad range of research fields, ranging from Robotics, and Human-Robot interaction and collaboration, to aspects of Computer Vision, such as scene understanding, or to the modelling and automation of Human-Machine social encounters and argumentative dialogues. The current ERCIM News special theme presents progress on all those matters, showcasing remarkable achievements in the industry, but also discussing highly interesting applications in everyday human life and in art.

Next, we provide a short overview on the articles included under the topic of Cognitive AI and Cobots, grouping them according to their thematic similarity.

AI in Manufacturing

Papoutsakis and Pateraki (page 8) describe a vision-based framework for real-time monitoring of 3D human motion and classification of ergonomic work postures and risk assessment during assembly tasks under real conditions, in the car manufacturing industry. Similarly, Anagnostopoulos et al. (page 10) present an AI-augmented multi-stereo camera system that collects ergonomic information of human operators in their working space, monitoring diverse anthropometric characteristics, in order to adjust the behaviour of robots working in the same space, accordingly. Vento et al. (page 12), in addition to visual information, also try to recognise speech commands, in order to achieve a more natural human-machine interaction in a cooperative, assembly line, industrial environment.

Models for Improving Human-Machine Collaboration in Industry

In order to achieve reliable human-machine symbiotic collaboration, Calabrò and Marchetti (page 14) discuss a desktop application for managing domain-specific protocols that can verify and validate a cobot's quality properties, such as safety and security. Behery et al. (page 15) propose a generalisation of "Behaviour Trees" as a modelling framework for industrial human-robot collaboration. This representation enhances safety, explainability, and verification, but also modification and exchange of production processes. Karvonen and Saariluoma (page 17) propose human digital twins as a conceptual design tool for a more human-centric approach to future industrial design. It is based on cognitive mimetics, a methodology for constructing technologies by mimicking human intelligent information processing. Dagioglou and Karkaletsis (page 18) investigate how humans perceive and experience the collaboration with AI-embodied agents, by means of a collaborative learning testbed that uses objective and subjective measures to rate the human-machine interaction.

Argument Exchange in Human-Machine Interaction

Loizos (page 20) proposes a structured, dialectical model for teaching an AI entity, through the exchange of arguments and counterarguments. Kakas (page 21) goes one step further and presents a system that implements Cognitive Machine Argumentation to investigate the effect of machine explanations on human reasoning.

AI in Arts

Thomay et al. (page 22) apply Cognitive AI in multimodal arts installations to create interactive systems that learn from their visitors, improving not only gesture recognition, but also usability and user experience. Cserteg et al. (page 24) combine Cognitive AI with Robotics to develop a portrait-drawing robot, which takes the picture of a person and, after a series of processing and execution steps, delivers a drawing with the person's facial features as a set of lines drawn by an average industrial robot.

Cognitive AI in Everyday Life

Hung (page 26) stresses the value of social encounters for humans and explores ways of measuring the quality of experience in social encounters through the interpretation of body language. Catricalà et al. (page 28) focus on the benefits that human-robot interaction can have on older adults, and discuss how serious games, based on and exploiting users' memories, can help them maintain their cognitive functional level. Hoang et al. (page 30) provide insights into a novel method for recognising various types of activities of humans performing everyday or, in some cases, professional tasks.

Please contact:
Theodore Patkos, ICS-FORTH, Greece
patkos@ics.forth.gr

Zsolt Viharos, SZTAKI, Hungary
viharos.zsolt@sztaki.hu

Automatic Vision-based Monitoring of Work Postures and Actions for Human-Robot Collaborative Assembly Tasks

by Konstantinos Papoutsakis (FORTH-ICS), Maria Pateraki (National Technical University of Athens & FORTH-ICS)

Visual monitoring of human behaviour during work activities is a key ingredient for task progress monitoring and fluency in human-robot collaboration as well as for supporting worker safety in industrial environments. The proposed framework employs low-cost, fixed sensors in a realistic manufacturing environment for the real-time observation of work postures and actions during car door assembly actions, as part of the FELICE project. This task aims to estimate workers' states and potential ergonomic risks in order to initiate robot collaboration with a specific worker, delegate tasks to the robot and ease the physical burden, while also contributing to the well-being of the worker.

The recent surge of interest in the manufacturing sector in technological advances related to collaborative robots is directly linked with the fast-growing demands for increasing production efficiency and flexibility in assembly processes and for lowering operating costs. Hence, apart from enhancing competitiveness, the introduction of collaborative robots in the workplace can also positively impact the human workforce and safety. Delegating ergonomically stressful and repetitive tasks to robots or recruiting cobots to aid workers during such strenuous tasks based on tool handover or other collaborative actions, eases physical burden on humans and contributes to

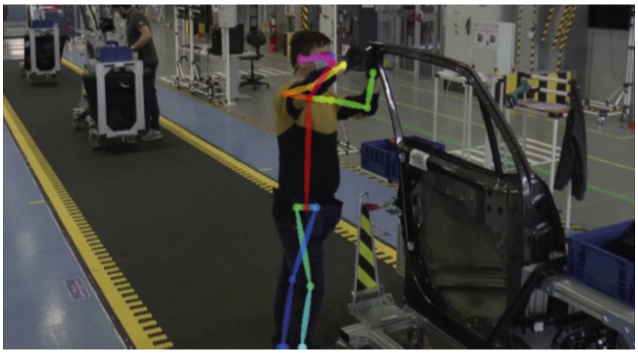

Figure 1: The car-door production line at the CRF workplace, which comprises three assembly workstations. A line worker is assigned to each workstation. The human body pose of a worker is estimated and overlaid (colour-coded skeletal body model). Image courtesy of Stellantis – Centro Ricerche FIAT (CRF) / SPW Research & Innovation department.

their occupational health (e.g., avoid muscle strains and injuries). FELICE project was initiated in 2021 [L1] with the aim to address one of the greatest challenges in robotics, i.e., coordinated interaction and combination of human and robot skills with the application priority area of agile production. It envisages adaptive workspaces and a cognitive robot collaborating with workers in assembly lines uniting multidisciplinary research in collaborative robotics, AI, computer vision, IoT, machine learning, data analytics, cyber-physical systems, process optimisation and ergonomics. Its overarching goal is to deliver a modular platform that integrates an array of autonomous and cognitive technologies in order to increase the agility and productivity of a manual assembly production system, ensure the safety, and improve the physical and mental well-being of workers.

As part of the FELICE system, a vision-based framework is introduced for real-time monitoring of 3D human motion and classification of ergonomic work postures during assembly tasks [1]. Contributions of the related developments and out-

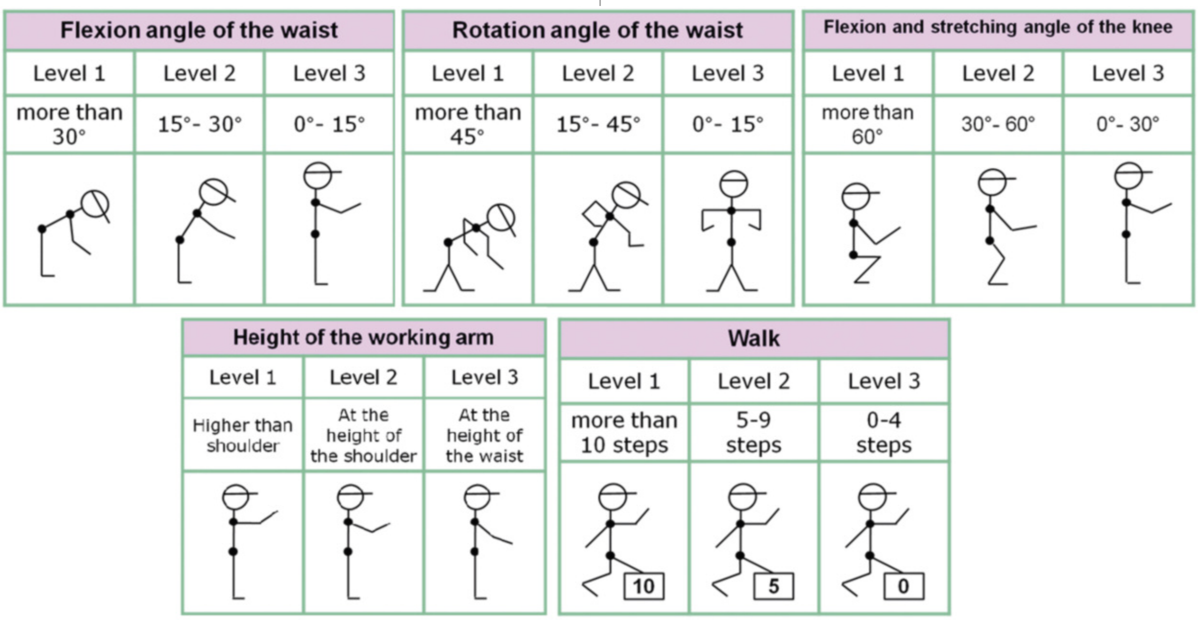

Figure 2: Five classes of ergonomic body postures during work tasks according to the MURI risk analysis approach. Image courtesy of Stellantis – Centro Ricerche FIAT (CRF) / SPW Research & Innovation department.

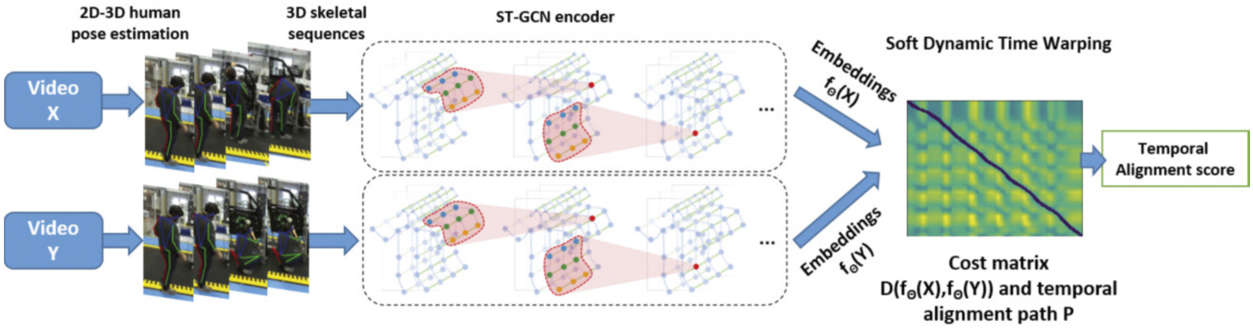

Figure 3: An overview of the proposed approach for vision-based classification of work postures given 3D skeleton-based body pose sequences captured during manufacturing activities [1].

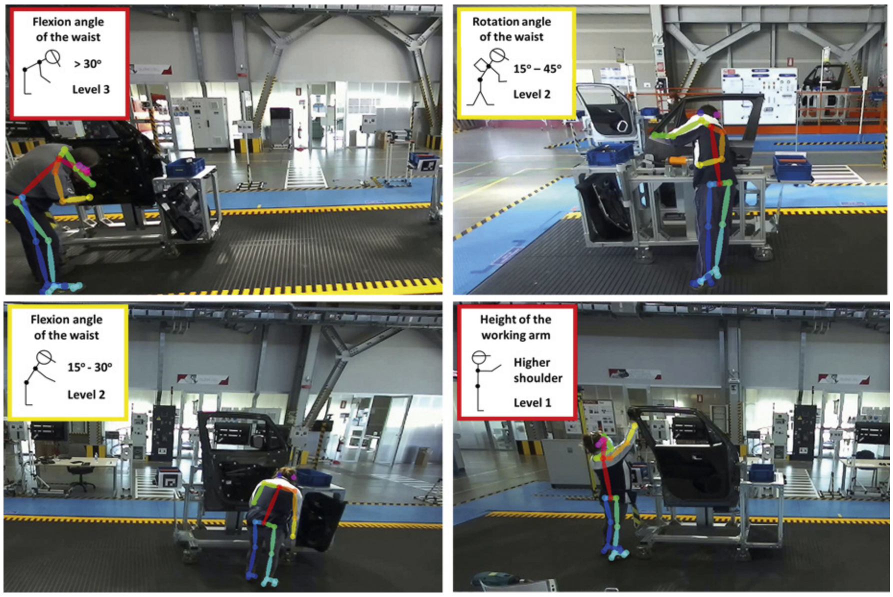

Figure 4: Snapshots of car-door assembly activities captured in a real manufacturing environment. Experimental results of the estimated 3D human body poses (overlaid as colour-coded skeletal body model) and of the classification of working postures that are associated with the ergonomic risk for increased physical strain are illustrated (text and sketch overlaid on the top left).

comes are: i) realising visual perception and cognitive capabilities, which will allow the system to assess ergonomic risks for physical strain of workers towards the prevention of work-related musculo-skeletal disorders (WMSD); ii) advancing human-robot collaboration by enabling the cobot to operate safely and to promptly intervene during assembly actions of high ergonomic risk to humans, sharing and reallocating tasks between them in an efficient and flexible manner and; iii) providing potential input to a manufacturing digital twin for the real-time worker motion, actions and physical state.

Our pilot case study addresses the car manufacturing industry in a real workplace located at the CRF-SPW Research & Innovation department of the Stellantis group in Melfi, Italy. We focus on workers that work in shifts, each assigned to a single workstation of a simulated car-door assembly line (Figure 1) to repeatedly execute a set of sequential actions, noted as a task cycle (total duration 4 to 5 minutes). In this context, we aim to estimate the 3D skeleton-based body motion in the presence of severe long-term occlusions for the automatic classification of five classes of ergonomically unsafe work postures during assembly actions (Figure 2). Posture classes are selected according to the postural grid of the MURI risk analysis method – a generic and widely used tool for efficiency evaluation and screening of physical ergonomics in workstations in different production contexts, according to the World Class Manufacturing strategy (WCM). Each class features three postural deviations of increasing risk for physical strain imposed to specific body parts that range from high ("Level 1") to low ("Level 3") risk level. It is known that Level 1 and Level 2 postural deviations lead to increased physical discomfort and stress in the context of particular work activities and serve as risk indicators for WMSD [2].

For each workstation, visual information (RGB image sequence) is acquired using two camera sensors installed in the actual workplace. Using the popular OpenPose method [2] the 2D body pose is estimated per frame. Subsequently, the MocapNet2 method [3] is employed to efficiently regress a view-invariant 3D skeleton-based pose based on ensembles of deep neural networks. Our non-invasive solution for estimating 3D human motion alleviates the need for the installation of special expensive equipment and wearable suits/reflectors (i.e., a motion-capture system). The 3D skeletal data sequence is forwarded to a Spatio-temporal Graph Convolutional Network model (ST-GCNs) for learning efficient discriminative representations of the spatio-temporal dependencies of the human joints (Figure 3). The ST-GCN model represents the locations of the human skeletal joints, as intra-graph edges of graph-based CNN per frame, whereas inter-graph edges connecting the same joints between consecutive graph models model the spatial temporal dynamics of the human motion.

The soft Dynamic Time Warping approach (softDTW) is used to estimate the minimal-cost temporal alignment between sequences as differentiable loss function and as classification metric.

A video dataset of assembly activities was compiled [L2] to facilitate the implementation and evaluation of the proposed method. Annotation data for assembly actions (310 instances) and work postures performed by two different workers during 12 task-cycle executions were provided by experts in manufacturing and ergonomics. The efficiency of the proposed deep-learning based classifier was measured based on the F1 scores per posture type using trimmed videos (Figure 4). The mean performance score is 68% of correctly classified postures, which is considered satisfactory considering the challenging conditions for acquiring high-quality estimation of the 3D human motion in a real-world manufacturing environment. Future work includes the development of a novel methodology for joint recognition of assembly actions and unsafe work postures and deployment of the FELICE system in the real workplace for evaluation using trained line workers.

The authors would like to acknowledge the consortium partners Stellantis – Centro Ricerche FIAT (CRF) / SPW Research & Innovation department in Melfi, Italy, for their valuable feedback in the design, implementation, visual-data acquisition and annotation involved in this study. This work has received funding from the European Union’s Horizon 2020 research and innovation program under grant agreement No. 101017151 (FELICE).

Links:
[L1] https://www.felice-project.eu/
[L2] https://zenodo.org/record/7043787

References:
[1] K. Papoutsakis et al., “Detection of Physical Strain and Fatigue in Industrial Environments Using Visual and Non-Visual Low-Cost Sensors”, in MDPI Technologies Journal, 2022. https://doi.org/10.3390/technologies10020042
[2] Z. Cao et al., “OpenPose: Real-time Multi-Person 2D Pose Estimation Using Part Affinity Fields”, in IEEE Transactions on PAMI, 2021. https://doi.org/10.48550/arXiv.1611.08050
[3] A. Qammaz and A.A. Argyros, “Occlusion-tolerant and personalized 3D human pose estimation in RGB images”, in IEEE ICPR 2020. https://doi.org/10.1109/ICPR48806.2021.9411956

Please contact:
Konstantinos Papoutsakis, ICS-FORTH
papoutsa@ics.forth.gr

A Multi-stereo Camera System for Improving Physical Ergonomics in Human-Robot Collaboration Scenarios

by Christos Anagnostopoulos (ISI, ATHENA R.C), Gerasimos Arvanitis (University of Patras), Nikos Piperigkos (ISI, ATHENA R.C), Aris S. Lalos (ISI, ATHENA R.C) and Konstantinos Moustakas (University of Patras)

Robot behaviour affects worker safety, health and comfort. Enhancing operator physical monitoring reduces the risk of accidents and musculoskeletal disorders. Within CPSoSAware, we propose a multi-stereo camera system and novel distributed fusion approaches that are executed on accelerated AI platforms, continuously monitoring the posture of the operators and assisting them to avoid uncomfortable and unsafe postures, based on their anthropometric characteristics and a real-time risk assessment standard methodology.

In manufacturing, the work of the robots is usually performed as a sequence of actions that are already known and pre-defined (it is not safe yet for the robots to improvise). A parameter that can affect the robot's movement is a characteristic of the operator's body (i.e., height). This parameter does not change the trajectory of the robot’s action but can be used to select the best adjustment of the robot's position and configuration in order to optimise the ergonomics assessment of the operator. The adjustment happens only one time at the beginning of the robot's operation since, for security reasons, it is not suitable to have a dynamic adjustment.

In the proposed architecture (Figure 1), 3 stereo cameras are utilised in order to cover the most visible area of the working space, so as to increase the possibility to have a good capture of the operator pose, with the most suitable direction, at least from one of the cameras, while the operator is moving in different directions and positions. Each stereo camera is used to extract the 3D pose landmarks of the operator for the height estimation and to calculate the current anthropometric state (based on the RULA score [L1]), in real-time. Each RGB camera of the stereo set is used to monitor the human's actions. A pose estimation algorithm is running to extract the posture 2D landmarks that are used, then for the 3D landmark, estimation is based on a Direct Linear Transform (DTL) triangulation approach [1], [L3].

Based on the ergonomic information, a decision-making module finds the right set of actions in order to minimise strain and ergonomic risk factors. Finally, the robot control receives the control information that the wanted set of actions is carried out. A total of three different classes are used based on the operator’s height, leading to three adaptable robot responses (adaptation scenarios) to allow the human to work in an optimal ergonomic state. The algorithmic framework identifies the operator’s class, based on their height, and the corresponding selected scenario configures the cooperating robot's movement

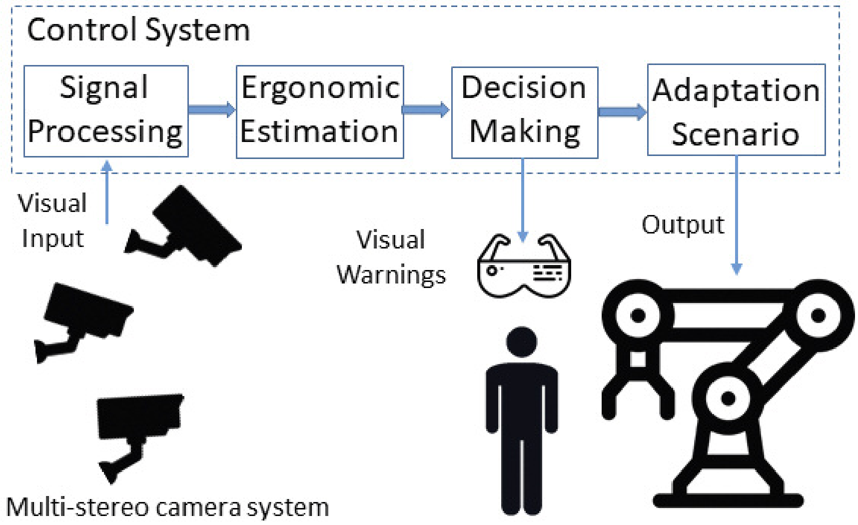

Figure 1: Proposed concept architecture.

parameters to adapt to the environment in a way that is ergonomically most comfortable for the interacting user. We assume the existence of three different classes based on the operator's height. More specifically, class 1 consists of operators with a height < 175 cm, in class 2 the operator's height is between 175 and 185 cm, and finally in class 3 the operators are taller than 185 cm. Additionally, a real-time joint angle estimation is performed to estimate the RULA score. Based on this value, the appropriate warning messages are sent to the operators, informing them if their working pose is ergonomically correct or not (Figure 2). Simulations in a virtual environment were also used for further ergonomic analysis in order to optimise the workstation (Figure 3), so as to obtain a trajectory of the robot, in a real environment, as much as possible adapted to the human and additionally to be used as a validation framework since the ground truth information of the 3D landmarks is known.

For the reliable height estimation of the operator, we have to select the best possible landmarks configuration from the involved cameras. Due to varying camera viewpoints and operator's movement, it is expected that landmarks, detected from different cameras, will not be quite the same, and nor will their accuracy. To this end, it is required to design a fusion approach that couples the multicamera system in order to provide as output the final group of anthropometric landmarks of the operator. For this task, we create an undirected graph, consisting of cameras and landmarks as nodes and cameras along with detected landmarks as edges. In order to take advantage of this graph topology formulation, we apply the Graph Laplacian Processing technique [3] using as anchor points the average landmark from each camera configuration, for re-estimating landmarks' positions in an optimal manner.

For the seamless integration of the components and the algorithms described above, we chose the Robotic Operating System (ROS) as our middleware. ROS is an open-source framework that contains a set of tools and libraries for developing distributed applications. Programs run on isolated nodes that can communicate using a publish-subscribe model. We have developed one ROS node for every camera used in the experiment. Each node captures a frame of the camera with a frequency of 10 Hz, extracts and processes the landmarks, and finally publishes them to a relevant ROS topic. The decentralised approach of ROS enables us to connect the cameras to

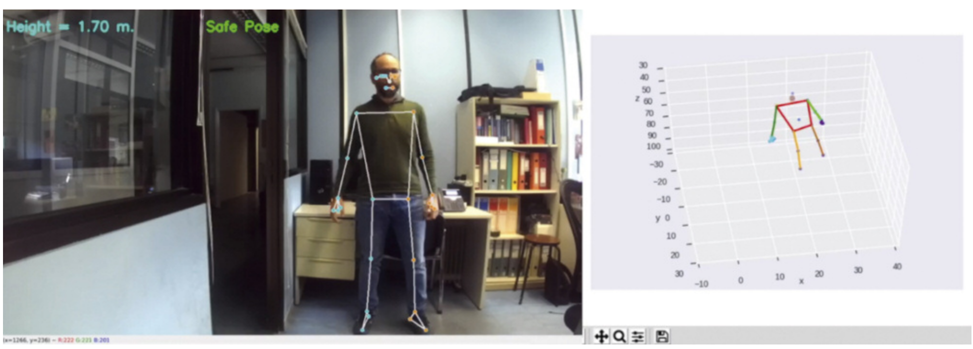

Figure 2: (a) Height estimation and safe pose identification, (b) Warning for wrong pose.

Figure 3: Examples of the simulator in different configurations (height of the operator), for ergonomic analysis in a safe virtual environment.

different physical machines, like the Jetson TX2 embedded device. A different node subscribes to the topics published by the cameras and extracts an optimal landmark set that will be used subsequently for the calculation of the posture of the operator.

CPSoSaware is an H2020-funded program, which targets shaping the future of Smart Mobility and Smart Manufacturing by developing autonomic, secure, cognitive cyber-physical system of systems that will support operators and users using artificial intelligence, dynamic reconfiguration and augmented reality technologies. The CPSoSaware solution aims to create self-aware cyber-physical systems of systems (CPSoSs) that detect and respond to physical and cyber environment changes, helping to avoid information overload, and reduce their management complexity by the human administrators in the automotive and manufacturing domains.

Link:
[L1] https://cpsosaware.eu/
[L2] https://ergo-plus.com/rula-assessment-tool-guide/
[L3] https://en.wikipedia.org/wiki/Direct_linear_transformation

Reference:
[1] N. Piperigkos, et al., "Cooperative Multi-Modal Localization in Connected and Autonomous Vehicles", 2020 IEEE 3rd Connected and Automated Vehicles Symposium (CAVS), 2020, pp. 1-5. https://doi.org/10.1109/CAVS51000.2020.9334558

Please contact:
Aris S. Lalos, Industrial Systems Institute, ATHENA Research Center, Greece
lalos@isi.gr

Speech and Gesture Command Recognition to Improve Human-Robot Interaction in Manual Assembly Lines

by Mario Vento, Antonio Greco and Vincenzo Carletti (University of Salerno)

UNISA is working on a more natural human-robot interaction and cooperation in manual assembly lines through speech and gesture commands. Where are we with the FELICE project?

FELICE [L1] is an ambitious project, started in the beginning of 2021, which has received funding from the European Union's Horizon 2020 Research and Innovation program. The project aims at delivering a modular platform that integrates and harmonises different technologies and methodologies, from robotics to computer vision and ergonomics, to increase the agility and productivity in manual assembly lines, ensure the safety of factory workers that interact with robots, and improve the workers' physical and mental well-being.

To pursue these goals FELICE involves different academic and research institutions (FHOOE, FORTH, ICCS, IFADO, IML, PRO, TUD, UNISA), one the largest automotive manufacturers in Europe, i.e., Centro Ricerche Fiat (CRF) belonging to Stellantis, three SMEs with significant commercial expertise in engineering and IT (ACCREA, AEGIS, CALTEK), and a legal consultancy service provider (EUN).

The Department of Information and Electrical Engineering and Applied Mathematics (DIEM) at the University of Salerno (UNISA) [L2] has a long-time experience in the fields of artificial intelligence, computer vision, cognitive robotics and autonomous vehicle navigation. DIEM is involved in the project to provide its competencies to improve the interaction between robots and humans in manual assembly lines. Such interaction usually comes through hardware interfaces like keyboards, buttons, and displays. Even though this is the traditional and commonly adopted way to interact with robots, it is not that natural for humans that are used to communicating with each other using speech and gestures. This is where the recent outcomes of artificial intelligence and cognitive robotics can make the way workers and robots to interact smoother, by providing the latter with the ability to understand speech and gesture commands and react coherently to them.

In the past two years, several steps have been taken towards realising a more natural human-machine interaction. The first step has been the formalisation of the tasks to be performed through artificial intelligence methodologies and the selection of gestures and speech commands of interest. The recognition of speech commands has been firstly analysed in the most general scenario possible, where there is a speaker who can produce complex non-predefined sentences that may or may not contain a command, namely the Conventional Spoken

Figure 1: The partners of the FELICE project in the test environment with the robot.

Language Understanding (SLU). In this case, it requires an acoustic model that translates speech to text and a linguistic model that extracts the intent of the speaker from the text; the intent is then associated with one of the commands provided. Successively, a set of short and predefined commands have been identified and the freedom of the speaker has been limited to short sentences and slight variants of the commands, in order to deal with practical problems, such as the limited quantity of speech-command data and the limited computational capabilities of the hardware embedded in the robot. Under the previous assumptions, the speech-recognition system does not have to extract the semantic information from sentences, so it does not require the acoustic model for the translation of speech into text, allowing the use of an end-to-end model where the audio input is directly associated with one of the defined commands of interest. The system has achieved an average accuracy of more the 93% during the tests performed in the real working environment at different noise levels.

Considering the gesture-recognition task, an experimental analysis highlighted the complexity of recognising gestures in real-time using state-of-the-art approaches due to their computational demands [1]; the solution has been to exclude dynamic gestures that require analysis across the duration of actions and focus on static gestures that can be recognised using convolutional neural networks (CNNs). Therefore, once the set of gestures to be recognised has been defined, the gesture-command-recognition system is realised by using a two-stage approach: firstly, the hand is detected in the image, then the gesture is classified analysing the pose of the hand. The effectiveness of this solution has also been tested in the wild where the system has achieved an accuracy of 95.4% while classifying among 17 different hand poses.

The two command-recognition systems have been installed in the device embedded on the robot, then a batch of tests in operative conditions have been recently performed and have confirmed the effectiveness of the proposed approaches. The next steps will be: (i) the extension of the dataset for both speech- and gesture-recognition systems by acquiring more data in real-world environments, (ii) exploring other approaches and improving those that have been adopted so far using the data collected, optimising the recognition systems for NVIDIA-equipped devices using frameworks like Tensor RT.

Links:
[L1] https://www.felice-project.eu
[L2] https://mivia.unisa.it

References:

[1] S. Bini, et al., "Benchmarking deep neural networks for gesture recognition on embedded devices", 31st IEEE Int. Conf. on Robot and Human Interactive Communication (RO-MAN), 2022, pp. 1285-1290. https://doi.org/10.1109/RO-MAN53752.2022.9900705

Please contact:

Vincenzo Carletti, University of Salerno, Italy
vcarletti@unisa.it

Antonio Greco, University of Salerno, Italy
agreco@unisa.it

Cobot Protocol Customisation Manager – PLUME

by Antonello Calabrò and Eda Marchetti (ISTI-CNR)

Making cobot safety protocols closer to a cookbook than to methodologies: Let users implement protocols without overhead in knowledge and understanding of the procedures.

The current intensive research activities around data-driven AI technologies, computer science and cognitive science still need to completely solve the stringent need for compliance and conformity with safety and security regulations. They provide essential documentation, requirements, and directives that must be satisfied, validated, and verified during the cobot's development. Indeed, to achieve a reliable human-machine symbiotic collaboration, regulations must be considered when applying data-driven AI technologies into cognitive robotics environments. However, the rules' complexity could impact the application of cognitive AI & cobots unless the required directives are not expressed as structured, easy-to-apply operational guidelines.

In this context, protocols are documents containing the methods, the collection of the procedures, and the prerequisites necessary to carry out safety experiments. These aim to verify and validate a cobot's quality (like safety and security) properties and the risk assessment [1].

Conceived inside the COVR project [L1], PLUME (ProtocoL cUstomization ManagEr) wants to provide an easy-to-use desktop application for writing/updating domain-specific protocols.PLUME is based on a model-driven approach that supports the protocol during its lifecycle. It guides users in creating, managing, and updating protocols to leverage the quality of the services they can offer to society.

As shown in Figure 1, through the PLUME user-friendly interface, a classification of the specific application domains and the relative device types is provided. Consequently, PLUME allows to:
- Gather the required directives representing the legislative acts and goals that all EU countries must achieve. In the example, the Machinery Directive can be selected.
- Gather the domain-specific standards and requirements directly connected to the selected directive. In the example, the ISO/TS 15066 standard is automatically associated, and requirements like "Limit energy transfer during collision" are shown to the users.
- Select or specify the abilities (skills) of the cobot. These will be verified and validated for risk reduction. The "Limit physical interaction energy" safety skill is selected in the example.
- Specify the working environment conditions in which the cobot device will operate. In the example, the "indoor - factory" can be identified.
- Select or specify the metrics, measures, and boundaries used during the validation. In particular, the metrics can be associated with values or Boolean when a threshold-based assessment is necessary. In the example, the "holding" with the "power on" is the condition in which the device will be tested.
- Identify possible test procedures or cases useful for the cobot's validation and assessment.

As shown in Figure 1, PLUME creates an interactive protocol skeleton reporting the collected information. It provides the users with a manageable document for cobot testing and validation. It allows them to:
- finalise the content with additional information,
- modify or refine the different fields according to specific needs or constraints,
- include existing protocols (or part of them) into a new one.

From a process point of view, the PLUME life-cycle includes two kinds of users: the generic Users (for instance, developers or testers) who are responsible for the protocol creation and the Experts (like a member of insurance, legislation government/organisations, or professionals) who are responsible for the approval of the protocols.

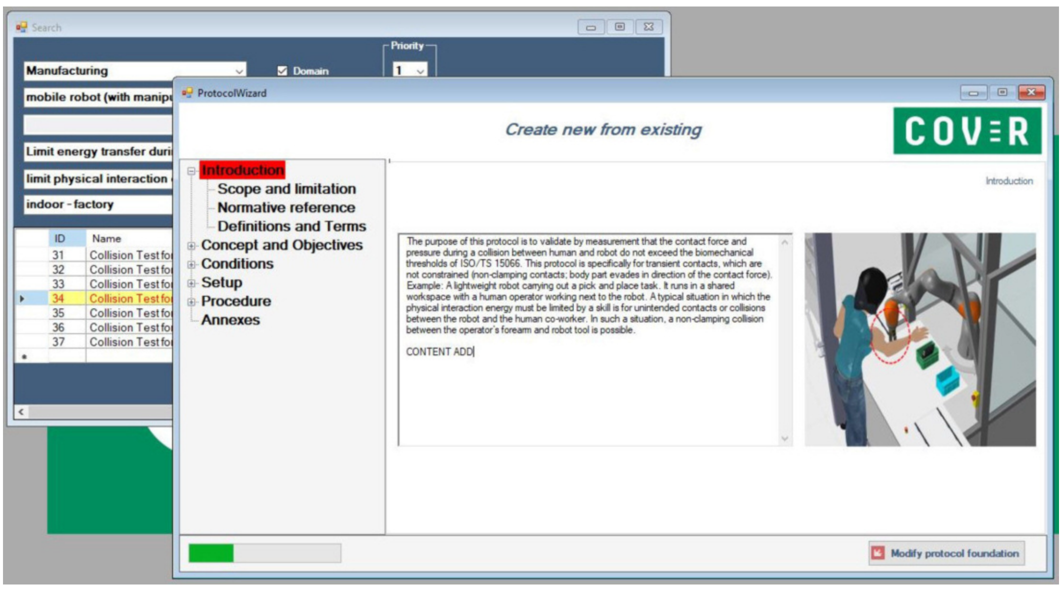

Figure 1: PLUME interactive protocol skeleton interface.

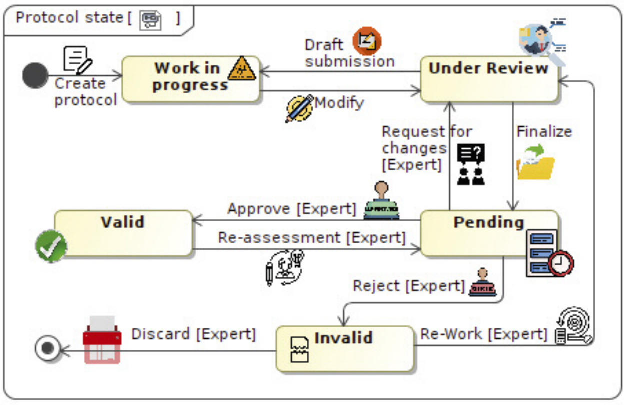

Figure 2: PLUME protocol life cycle.

As reported in Figure 2, during its lifecycle, the protocol passes through different stages:

- Work in progress is when the protocol skeleton fields need to be completed by the Users using the protocol wizard provided by PLUME.
- Under review is when the User provides a Draft submission or the protocol. In this stage, if the Users modify the protocol (Modify arrow in Figure 2), it returns to the Work in progress stage. Otherwise, it is classified as finalised and passed to the next Pending step (Finalize arrow in Figure 2).
- Pending is when the Experts evaluate the protocol to assess its correctness and completeness. In this case, three different situations can happen: Request for changes, i.e., the Experts describe needed modifications, and protocol returns to the Under review stage. Reject, i.e., the Experts do not approve the protocol, and it is moved to an "Invalid" stage. Approve, i.e., the Experts consider the protocol valid and move it into the Valid stage.
- Invalid is when a protocol needs to be discarded or integrated with Experts' suitable instruction to leverage it again to the Under review stage.
- Valid is when a protocol is available and ready to be used. If Experts decide to modify the protocol, it returns to the Pending stage for the second round of evaluation.

The current implementation of PLUME's standalone application relies on the MONO-Project framework [L2] and can be integrated with different datasets and repositories. Inside the COVR project, PLUME can interact with the Toolkit management system.

Links:
[L1] https://www.safearoundrobots.com/home
[L2] https://kwz.me/hvQ
[L3] https://www.safearoundrobots.com/toolkit/home

Reference:
[1] M. Askarpour, et al., "Formally-based Model-Driven Development of Collaborative Robotic Applications", in J. Intell, Robotic Syst. 102(3): 59 (2021). https://doi.org/10.1007/s10846-021-01386-2

Please contact:
Antonello Calabrò, ISTI-CNR, Italy
antonello.calabro@isti.cnr.it

Behaviour Trees for Representing Human-Robot Collaboration Processes in the World-Wide Lab

Mohamed Behery, Philipp Brauner, Martina Ziefle and Gerhard Lakemeyer (RWTH Aachen University)

Shorter product lifecycles, more product variants, individualised production, and the desire for sustainable production call for agile control frameworks that enable smarter robotic control and collaborating human-robot teams. We propose generalising and standardising "Behaviour Trees" that use human action nodes as a process model and task-execution-monitoring approach for human-robot collaborative assembly processes to increase the agility of human-robot teams while ensuring a safe and trusted human-robot interaction. Within the DFG (Deutsche Forschungsgemeinschaft) funded Cluster of Excellence "Internet of Production", we take a cross-disciplinary approach to conceptualisation and validation to ensure algorithmic soundness, technical viability, and social acceptance by the workers of increasingly agile human-robot teams.

Production processes that involve Human-Robot Collaboration (HRC) are required to meet the demands for more product variant requests, initialised production, and short product lifecycles, which calls for agile control frameworks that ensure a safe, trusted, and socially accepted integration of humans in the robots' workspace.

Moreover, the exchange and replication of production setups and pipelines is desirable in the World-Wide Lab (WWL) [1] and requires an abstract modular transferable process representation. The Cluster of Excellence Internet of Production (IoP) [L1] tackles these problems, among others, using process Digital Shadows (DSs) [2]. DSs are context- and task-specific process models that can be stored in online servers to be used in many ways including process control, verification, and decision support systems. Additionally, they can be used to efficiently couple the different entities of the WWL.

HRC processes have so far not been integrated into the WWL. The integration is challenging due to the lack of DSs for these processes. This can be attributed to the ad-hoc robot-specific programming involved in these use cases, which can be laborious and time consuming. They are inherently challenging because they cover a wide range of setups, all of which have a proximity to human workers.

HRC processes need DS to enable real-time introspection, code reuse, and integration into the WWL. A DS of an HRC process must conserve the safety requirements (and guarantees) of the process. For example, an HRC assembly must represent the requirement of a fully assembled product while maintaining the safety of the human co-worker.

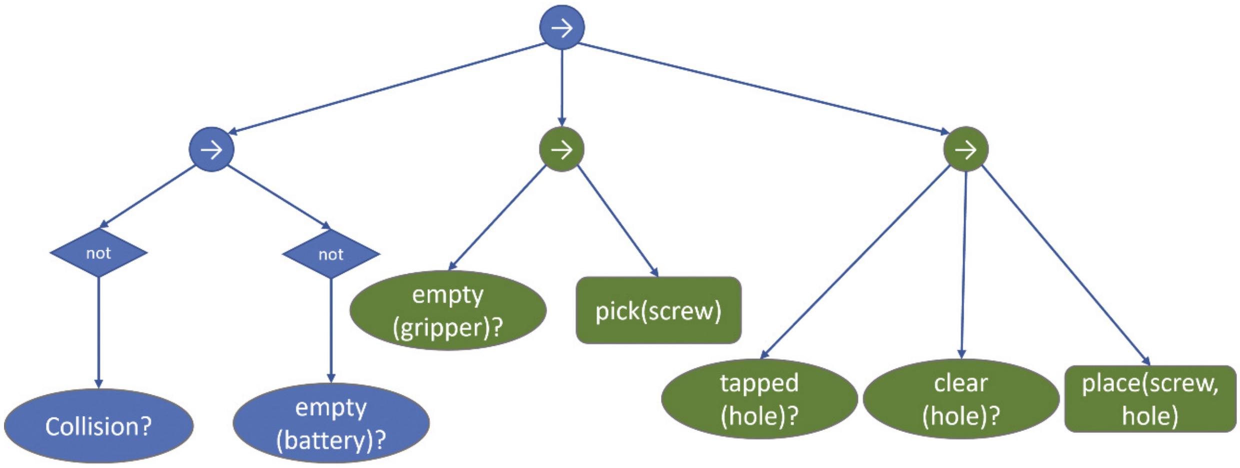

Figure 1: A sample behaviour tree for an assembly task. The leftmost branch of the tree (blue) ensures that the robot is not in a collision state and that the battery is not empty. The two remaining branches (green) pick a screw and place it in a clear hole. The latter branches contain guards that ensure that the preconditions of pick and place actions hold.

We propose generalising Behavior Trees (BTs) [3] as a modelling framework for industrial HRC applications and using BTs as a DS in the WWL. Recent work showed several extensions and applications of BTs to support HRC processes, heterogenous multi-agent systems, verification, as well as learning and synthesis of trees [3].

BTs are tick-based action representation- and execution-monitoring models. They represent tasks in a tree structure whose execution starts by ticking the root; internal control flow nodes propagate the tick to the leaves, which are responsible for the robot-specific execution. This structure splits processes into standalone modules with different responsibilities (e.g., a branch of the tree can focus on safety while others focus on navigation and manipulation). This can be seen in Figure 1, where the leftmost branch prevents the robot from moving in cases of collision and low battery and the core pick and place branches contain guards that ensure the respective action preconditions hold before execution. They represent the general execution flow and contain placeholders for robot-specific behaviour. They also represent priority of tasks (or subtrees) over others thanks to the tree’s ordering; they are executed from left to right and top to bottom. Their reactivity, modularity, and flexibility give them an edge as a task-level controller in many domains [3].

Given a BT representation, a process can be modified by replacing one subtree with another. This replacement can take place between nodes of different types (or a leaf and a branch) as long as the tree remains valid. Additionally, we can have a tree that ensures safety and use it as a subtree (with high priority) to guarantee the safety of the rest of the tasks in the new tree. This ensures that BT representations are flexible, modular, and reusable.

Our previous work proposed using a human-action node with a rule-based system allowing robots to switch tasks and react to the outcomes of human sub-tasks instead of waiting for them. Using BTs improves the “explainability” of the robots' actions and thus increases trust between the human workers and the robots. One of the demonstrators in the IoP shows how BTs can represent and execute processes that involve keeping safe distances between humans and robots during the execution. Additionally, ongoing work in our cluster explores the safety requirements involved in collaborative processes between humans and robots and how we can use different properties of BTs to satisfy them.

BT models of HRC processes are easy to share, find, and synthesise in the WWL. After the integration of BTs as the standard DS representation in the WWL, we plan to further extend BTs for increased robustness and reactivity through mixed initiative planning and supporting further use cases such as robot teleoperation.

This work is funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC-2023 Internet of Production – 390621612. It is partially supported by the EU ICT-48 2020 project TAILOR (No. 952215).

Links:
[L1] https://www.iop.rwth-aachen.de

References:
[1] P. Brauner, et al., “A computer science perspective on digital transformation in production”, in ACM Trans. Internet Things 3(2) (2 2022). https://doi.org/10.1145/3502265
[2] M. Liebenberg, M. Jarke, “Information systems engineering with digital shadows: concept and case studies” in Int. Conf. on Advanced Information Systems Engineering, pp. 70–84, Springer (2020).
[3] M. Iovino, et al., “A survey of behavior trees in robotics and ai”, Robotics and Autonomous Systems 154 (2022): 104096.

Please contact:
Mohamed Behery, Knowledge-Based Systems Group
RWTH Aachen University, Germany
behery@kbsg.rwth-aachen.de

Cognitive Mimetics and Human Digital Twins – Towards Holistic AI Design

by Antero Karvonen (VTT) and Pertti Saariluoma (Jyväskylä University)

AI is replacing and supporting people in many intelligence-requiring tasks. Therefore, it is essential to consider the conceptual grounds of designing future technical artefacts and technologies for practical use. We are developing two new practical design tools: cognitive mimetics and human digital twins for AI designers. Cognitive mimetics analyses human information processing to be mimicked by intelligent technologies. Human digital twins provide a tool for modelling what people do based on the results of cognitive mimetics. Together they provide a new way of designing intelligent technology in individual tasks and industrial contexts.

Today, the designers of industrial processes and technologies must find answers to novel problems of designing intelligent technical systems. The issues emerge at the intersection of major trends. First, AI is predicted to gradually enter further into aspects of industrial processes, in particular replacing and augmenting the work of operators. Second, more industrial work and design is moving into digital formats such as digital twins, augmented reality and the metaverse. Third, Industry 5.0 calls for a human-centric approach to industrial design.

A question facing a designer is how the human aspect should be included in this large-scale digitalisation of the industrial processes (in particular in the cluster of issues arising from AI) and even the design process itself. Human digital twins (HDTs) provide a conceptual design tool that can address some these questions, presenting a concretisation of a methodology called cognitive mimetics. HDTs and cognitive mimetics unify into single design scientific paradigm the basic questions of human work, action, and thought with intelligent technology – and their joint development [1].

The cognitively oriented analysis of human-technology interaction, embedded in cognitive mimetics, provides a set of questions and methods for the gradual development of human digital twins. Thought and action in relation to technical systems can be simultaneously analysed and modelled. As more aspects of the human information processing in a context are modelled in the HDT, implemented using a variety of extant AI or cognitive science, the closer it approaches intelligent technology. The value of such a process is potentially great, since it allows for the mutual development of human operations and technology as a joint cognitive system [2].

Cognitive mimetics is an empirical design scientific method. Like biomimetics mimics solutions of nature, cognitive mimetics constructs technologies by mimicking human intelligent information processing. Mapping between human information processing and computer models need not (and perhaps cannot) be slavishly similar, but similar to the way airplane wings follow the same laws of aerodynamics as bird wings. Intelligent computational processes can provide performance capacities that were previously only possible in people on the grounds of design knowledge collected by means of cognitive mimetics.

Cognitive mimetics does not deal purely with the mapping of abstract theories or models of AI in cognitive science. These are possible implementation methods for cognitive mimetics, which is focused methodologically on the analysis of empirical human information processes and contents. The main idea is to research and "open up" the thought and perception of, for example, human operators in a paper mill, to gain new ideas for the design of intelligent technology. The basic form of this activity is a mapping relation between a source and a target, here a human operator and an intelligent technology target.

Below are the main phases of cognitive mimetic design:
- Phase 1: Understand and define the intelligent technology design problem.
- Phase 2: Collect information about human information processes being able to carry out the task to be designed.
- Phase 3: Generate concepts for intelligent technology.
- Phase 4: Generate a solution structure.
- Phase 5: Finalise and deliver.

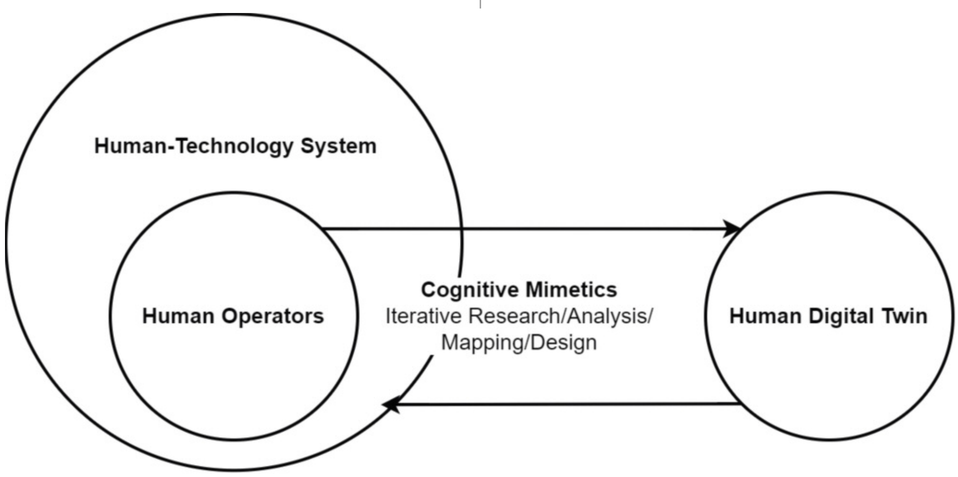

Figure 1: The basic idea of cognitive mimetics is iterative research and design, cycling between a source and a target. Here the source is specified as human operators (in an industrial context) and the target as human digital twins. Design research into human operators thinking progressively opens up the human-technology system as a whole. Ultimately the system is changed via the results of cognitive mimetic design, such as human digital twins.

There are naturally iterative loops between the phases. Each phase is completed through an empirical investigation of human action and thought using well-established methods such as observation, interviews, focus groups, think aloud and protocol analysis. Value can be extracted during all phases of the design process. The main point is to start from a general understanding of the context of human action and thought, advance through a concept-generation phase of how the joint cognitive system could be improved upon, and then focus in on a particular concept. At this point, cognitive and expertise research methods are employed to develop a rich understanding of the actual thought processes of operators. These are then mapped computationally to a human digital twin.

Human digital twins are computational models of people processing information (or perceiving, attending, and thinking) to carry out the critical tasks under design. Like digital twins of cyber-physical systems, human digital twins can be used to support design processes. They can be used to create alternative solutions, test the validity of solutions, and refine them.

The picture that emerges now addresses potential for intelligent technology, development of the human tasks and activities, places where human and machine intelligence can support each other, and the gradual emergence of intelligent and useful technologies. Using human digital twins makes the process digital and documented.

Human action can't be analysed separately from a context (and a space of possible action). Thus, human digital twins (constructed via cognitive mimetics) reflect and refer to the technical systems and processes in which industrial operators work. It is a short, logical step then to attach human digital twins to traditional digital twins and metaverse-type technologies. These yield concrete dataflow and control pipelines, as well as (machine and human) learning systems, for bringing together digitally the human and technical sides of industry. Achieving this basic vision is clearly bubbling under many researchers' thinking and visions of future industry. Arriving at this point via the cognitive mimetic route yields a deep and rich picture of human-technology co-agency in industrial processes.

Industry is just one example. AI will gradually change the way people live and work and how society operates. Therefore, it makes sense to develop AI design as a holistic process in which social and technical aspects of intelligent technologies can be simultaneously considered. Such new design practice can be called holistic AI design. Cognitive mimetics and HDTs provide key elements in such AI design.

References:

[1] P. Saariluoma, A. Karvonen, L. Sorsamäki, "Human Digital Twins in Acquiring Information About Human Mental Processes for Cognitive Mimetics", in Frontiers in Artificial Intelligence and Applications, 2022. https://doi.org/10.3233/faia210484.

[2] E. Hollnagel, D. A. Woods, "Joint cognitive systems: Foundations of cognitive systems engineering", 2005.

Please contact:
Antero Karvonen, VTT, Finland
antero.karvonen@vtt.fi

Understanding Partners' Behaviour for Transparent Collaboration

by Maria Dagioglou (NCSR 'Demokritos', Greece) and Vangelis Karkaletsis (NCSR 'Demokritos', Greece)

For this teaser we were asked to describe what is the article about, who is doing what and why. Fittingly enough, this is literally what we discuss in this article: in human-robot collaboration, who is doing what and why?

From carrying a heavy object to supporting a rehabilitation exercise, humans coordinate their bodies and minds to move together and achieve mutually agreed goals. During collaborative tasks, each partner needs to: know what the other perceives (or not), be able to predict his/her actions based on action observation and the requirements of the shared task, decide when and how to act for supporting team performance and efficiency [1]. In addition to all these, the spatio-temporal proximity of actions during collaboration affects the sense of self- and joint-agency. Who bears the authorship of a joint action's outcome? The experience of acting as a team, the nature of the task (competitive or complementary) and the fairness of resource distribution are only some of the factors that shape the perceptual distinctiveness of each partner's actions.

Now, what about human-robot collaboration (HRC)? What if a human were to collaborate with a robot to carry a sofa up the stairs (Figure 1)? In the context of this article, HRC is considered as the case of a human and a collaborative robot (cobot) that work in close proximity (within the intimate space of a human, which, based on Proxemics [L1], includes interactions from physical contact up to approximately 0.5 metres) towards a mutual goal that demands interdependent tasks. Similar to the processes described above for human-human collaboration (HHC), cobots must have the perceptive, cognitive and action capabilities that support joint attention, action observation, shared task representations and spatio-temporal action coordination.

In terms of the sense of agency, the (few) existing studies provide mixed evidence with respect to how this mechanism is manifested during human- (embodied) agent joint action and about the extent to which it follows similar patterns compared to HHC [2]. However, this is a crucial issue in HRC related to the delivery of ethical and trustworthy Artificial Intelligence (AI). Transparency in the authorship of each partner's actions is related to the ethical dimensions [L2] of: a) human agency, including over- or under-relying on the cobot's actions, as well as the social interaction instigated; b) accountability, in relation to the ascription of responsibilities if need be; and c) communicating and explaining the decisions of an AI system.

Thus, it becomes clear that HRC must be observed "in action", that is in real-world and in real-time, in order to not only evaluate the performance of AI methods, but also to study human behaviour towards cobots, and AI agents in general. Luckily, state-of-the-art AI methods now allow us to do so [3]. Towards this, in Roboskel, the robotics activity of SKEL-the AI Lab [L3], we have integrated an HRC testbed [L4] and have initi-

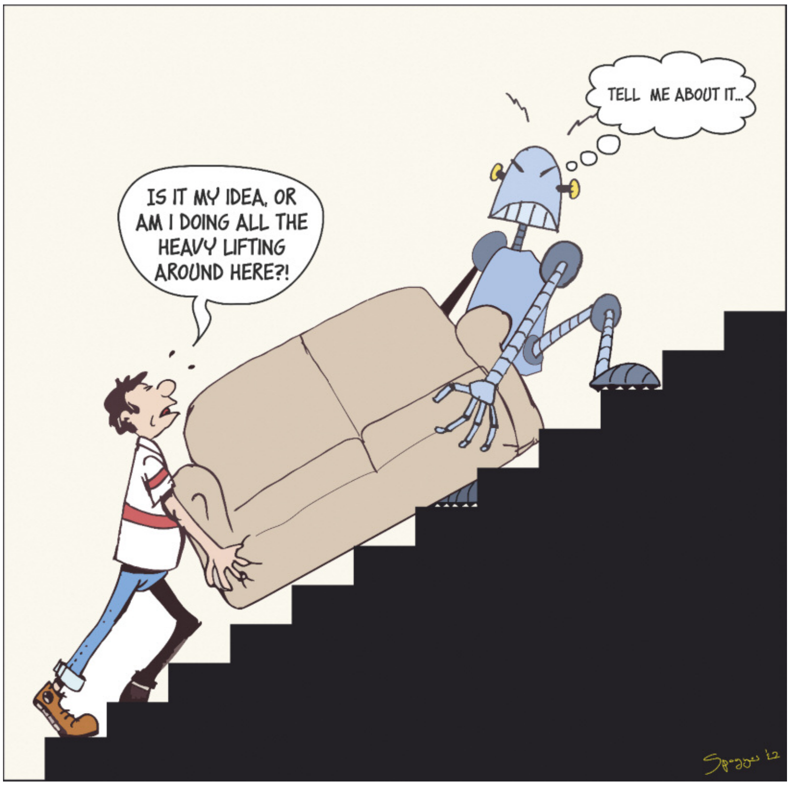

Figure 1: During collaboration it might be difficult to differentiate between self- and other-generated actions. Understanding our partner is crucial for fluent and transparent performance. Drawing by Stefanos Poirazoglou for the purposes of this article (CC BY 4.0).

ated a line of studies that aim at exploring team performance using objective measures, as well as subjective human partners' perceptions of the collaboration. As an example, we briefly describe below our most recent study [L5].

A human and an AI agent had to collaborate and drive the end-effector (EE) of a UR3 cobot, which was moving in a 2D plane, towards a goal position. Each partner was controlling the acceleration of the EE in one axis; the task was quite demanding and involved considerable collaborative learning. The human partner provided desired commands via a keyboard placed on the UR3's workspace. Seventeen participants were randomly assigned in one of two groups that involved collaboration with one of two different kinds of deep Reinforcement Learning agents: an agent that transferred knowledge (TKa) from an expert team and one that did not. The participants were unaware of the group they were assigned to.

As expected, the collaboration with the TKa significantly affected the time of the collaborative learning and the overall performance; the partners of these teams managed to reach the expert performance within the time provided. Moreover, the overall duration of the training was less than half compared to the other group (33.7 vs 73.1 minutes, respectively). In addition to the objective measures, perceived fluency was also considerably different in the two groups. Surprisingly, it was observed that participants that collaborated with the TKa tended to rate their own contribution and improvement quite high, without acknowledging at the same time the contribution of the robot.

This last result clearly demonstrates the need for further pursuing the research questions described earlier; there is an urgent need for systematically studying and better understanding how people perceive and experience the collaboration with AI (embodied) agents. Just transferring knowledge and results from human-human joint action studies does not seem appropriate; HRC triggers different behaviours. Naturally, this knowledge is necessary for developing ethical and trustworthy AI. On the one hand, AI agents must support fluent collaboration; in this context, some lack of self-agency might be desirable. At the same time, major issues of human agency, transparency and accountability are raised. Real-world studies that shed light on human behaviour during HRC are thus necessary for the alignment of the various ethical dimensions at stake and for adequately advancing the state-of-the-art AI methods.

Links:

[L1] https://en.wikipedia.org/wiki/Proxmics
[L2] https://altai.insight-centre.org/
[L3] https://www.skel.ai/, https://www.roboskel.iit.demokritos.gr
[L4] https://kwz.me/hvV
[L5] https://kwz.me/hvW

References:

[1] N. Sebanz, H. Bekkering, G. Knoblich, "Joint action: bodies and minds moving together", in Trends in Cognitive Sciences, 10(2), pp.70-76, 2006. https://doi.org/10.1016/j.tics.2005.12.009
[2] M. Dagioglou, V. Karkaletsis, "The sense of agency during Human-Agent Collaboration", HRI 2021 Workshop: Robo-Identity: Artificial Identity and Multi-Embodiment, March 8, Virtual, 2021. https://kwz.me/hvA
[3] F. Lygerakis, M. Dagioglou, V. Karkaletsis, "Accelerating Human-Agent Collaborative Reinforcement Learning", in The 14th PETRA Conference, 2021. https://doi.org/10.1145/3453892.3454004

Please contact:

Maria Dagioglou
National Centre for Scientific Research 'Demokritos', Greece
mdagiogl@iit.demokritos.gr

Coachable AI

by Loizos Michael (Open University of Cyprus
& CYENS Center of Excellence)

AI systems that are compatible with human cognition will not come as an afterthought by patching up opaque models with post-hoc explainability. This article reports on our research program to develop Machine Coaching, a computationally efficient and cognitively light paradigm that supports a form of interactive and developmental machine learning, through supervision that is not only functional, on "what" decisions should be made, but also structural, on "how" they should be made. We discuss the role that this paradigm can play in neural-symbolic systems, in operationalizing fast intuitive and slow deliberative thinking, and in supporting the explainability and contestability of AI systems by design.

A salient characteristic of data-driven neural-based approaches to machine learning is the autonomous identification of intermediate-level concepts, which then serve to scaffold higher-level decisions. These identified concepts do not necessarily align with familiar human concepts, which leads to a perception of opaqueness of the learned models. Post-hoc explainability methods can serve only in explicating the intermediate-level concepts in a learned model, but not in addressing the misalignment problem.

The problem is manifest in learning settings where the empirical generalizations leading to intermediate-level concepts are not directly pitted against a ground truth. Tackling misalignment requires, therefore, a non-trivial human-machine collaboration [1], during which a learner and a coach — who could be the learner's parent, teacher, or senior colleague in the case of a human learner, and a domain expert, or a casual user of the learned model in the case of a machine learner — engage in a dialectical and bilateral exchange of explanations of whether and how a decision should relate to intermediate-level concepts.

A coaching round begins when the learner faces a situation that calls for the making of a decision using the current learned model. Regardless of the decision's functional validity (i.e., whether the decision is correct), the coach can inquire for evidence of the decision's structural validity (i.e., whether the reasons for making that decision are correct). Upon receiving such a query, the learner returns arguments that decompose the decision as a simple function (e.g., a short conjunction) of intermediate-level concepts. Further inquiries can prompt the learner to decompose those concepts to even lower-level concepts. If the coach identifies an argument to be missing or unacceptable, the coach can offer a counterargument of how the decomposition could have been done better. The round concludes with the learner integrating the counterargument into the learned model, giving it preference over all existing conflicting arguments, and adopting it as part of the learner's pool of arguments that are available for use in subsequent rounds.

It can be formally shown that the process converges, in a computationally efficient manner, to a learned model that is, with high probability, structurally (and functionally) valid in a large fraction of all future situations [2]. The resulting structural validity entails the alignment of the learned model's intermediate-level concepts with those of the coach, while the rich interactive feedback enhances learnability beyond that of standard supervised learning. Concept alignment and enhanced learnability come with minimal extra burden for the coach, as suggested by evidence from Cognitive Psychology that argumentation is inherent in human reasoning, and that the offering of counterarguments is cognitively light. Relatedly, the coach's counterarguments are not required to be optimal, but merely satisficing in a given situation.

Depending on the nature of the application domain, the online engagement of the coach can be alleviated by alternating between a deployment phase, where the learned model makes multiple decisions, and a development phase, where the coach examines the functional and structural validity of a selected subset of those decisions. In domains where the functional validity (and perhaps the criticality) of decisions can be ascertained objectively, a certain number of incorrect (critical) decisions could trigger the transition to the development phase, where only those decisions will be further probed for their structural validity.

Nothing restricts the learned model to being purely symbolic, nor the irreducible concepts, reached by the iterated decomposition of intermediate-level concepts, to being directly observable. Instead, irreducible concepts can act as an interface between a neural and a coachable module in the learned model. Neural-symbolic techniques can be used to train the neural module so that its outputs seamlessly feed into the coachable module, and so that the latter module's arguments provide su-

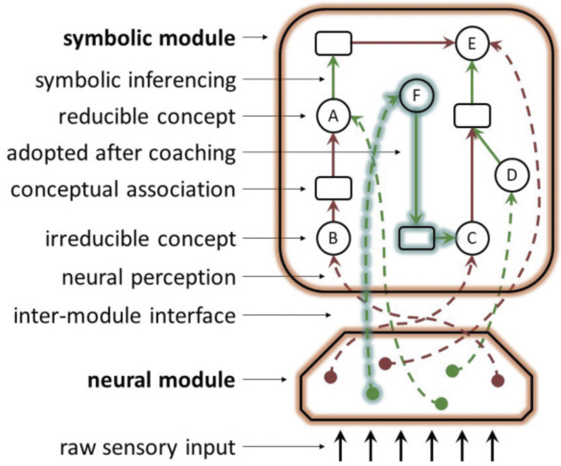

Figure 1: Depiction of a learned model with a neural-symbolic structure. The learned model receives raw sensory inputs through its neural module, maps them into a vector of output values, and maps each output value to one of the concepts in the symbolic module in a bijective manner. Reducible concepts in the symbolic module relate to other concepts through conceptual associations, which are viewed as expressing supporting or attacking arguments. Each of the two modules implements three methods: (a) deduction, for returning the output of a module when given a certain input, (b) abduction, for returning potential inputs of the module that give rise to a certain output, and (c) induction, for revising the module's part of the learned model. Induction for the neural module can utilise standard neural training techniques, while induction for the symbolic module can utilise machine coaching. The learned model as a whole makes decisions and is revised by combining the methods of its two constituent modules.

pervision signals for the former [3]. The development phase can be extended to include a neural training round after each coaching round. The structure of the learned model is shown in Figure 1.

According to this structured learned model, each potential decision with respect to some concept has an explanation that comes from direct neural perception. Decisions that structurally employ this first type of explanation resemble those coming from the fast and intuitive (System 1) thinking process of the human mind. For some concepts (namely, the reducible ones), there exist additional explanations that come from associative symbolic inferencing. Decisions that structurally employ this second type of explanation resemble those coming from the slow and deliberative (System 2) thinking process of the human mind. The argumentative nature of the learned model offers, then, an operationalization of the System 1 versus System 2 distinction as a difference in the effort put into determining the acceptability of arguments in the symbolic module of the learned model. With minimal effort, the argument "predict the neural module percept" is accepted uncontested, while extra effort allows the consideration of additional arguments that may defeat the former one, and lead to the making of possibly distinct decisions.

The argumentative nature of the symbolic module supports the offering of explanations of the learned model's decisions not only in the standard attributive fashion used in post-hoc explainability, but also in a contrastive fashion against other candidate decisions that rely on defeated arguments. Additionally, the structural validity of the symbolic module allows the learned model to demonstrate, by design, the compliance of its decisions against the coach's held policy, in case those are contested.

This work was supported by funding from the EU's Horizon 2020 Research and Innovation Programme under grant agreements no. 739578 and no. 823783, and from the Government of the Republic of Cyprus through the Deputy Ministry of Research, Innovation, and Digital Policy. Additional resources for Machine Coaching, including online services and demonstration scenarios, are available at [L1].

Link:
[L1] https://cognition.ouc.ac.cy/prudens

References:
[1] L. Michael, "Explainability and the Fourth AI Revolution" in Handbook of Research on Artificial Intelligence, Innovation, and Entrepreneurship, Edward Elgar Publishing, 2023. https://arxiv.org/abs/2111.06773
[2] L. Michael, "Machine Coaching", in Proc. of XAI @ IJCAI, p. 80–86, 2019. https://doi.org/10.5281/zenodo.3931266
[3] E. Tsamoura, T. Hospedales, L. Michael, "Neural-Symbolic Integration: A Compositional Perspective", in Proc. of AAAI, p. 5051–5060, 2021. https://doi.org/10.1609/aaai.v35i6.16639

Please contact:
Loizos Michael, Open University of Cyprus
& CYENS Center of Excellence, Cyprus
loizos@ouc.ac.cy

Cognitive Machine Argumentation

by Antonis Kakas (University of Cyprus)

How can machines think in a human-like way? How can they argue about and debate issues? Could machines argue on our behalf? Cognitive Machine Argumentation studies these questions based on a synthesis of Computational Argumentation in AI with studies on human reasoning in cognitive psychology, philosophy and other disciplines.

Cognitive AI systems are required to operate in a cognitively compatible way with their human users, naturally connecting to human thinking and behaviour. In particular, these systems need to be "explainable, contestable and debatable", thus able to adapt their knowledge and operation based on rich forms of interaction with humans.

One way to build such systems is to base their design on argumentation, exploiting its strong and natural link to human cognition. The relatively recent work of Mercier and Sperber [1] emphasises that human thinking is regulated by argumentation: humans (are motivated to) reason in order to justify their conclusions against counterarguments that come either from their external social environment or indeed from their own introspection. We can thus consider Computational Argumentation as a possible foundational calculus for Cognitive or Human-centric AI [2]. Nevertheless, in order to serve well the requirements of Cognitive AI it is necessary to inject within the formal framework of computational argumentation, cognitive principles, drawn from the many other disciplines in which human reasoning and in particular human argumentation is studied (Figure 1).

Cognitive Machine Argumentation is built through such a synthesis of argumentation in AI with elements from these other disciplines, such as cognitive psychology, linguistics, and philosophy. The task is to let these cognitive principles regulate the computation of argumentation in order to make it effective and cognitively viable within the dynamic and uncertain environment of application of AI systems. By "humanising" the form of machine argumentation, Cognitive Argumentation aims to facilitate an effective and naturally enhancing human-machine integration.

The framework of Cognitive (Machine) Argumentation has been validated by showing that it models well the empirical data from cognitive science in three classic human reasoning domains, that of Syllogistic Reasoning, the Suppression task and the Wason Selection task. COGNICA is a system that implements the framework of Cognitive Argumentation. Using argumentation it accommodates and extends the mental models theory on human conditional reasoning [3], from individual conditionals to sets of conditionals of different types, that together form a piece of knowledge on some subject of interest and human discourse. The existence of systems, such as COGNICA, opens up the opportunity for carrying out new systematic empirical studies of comparison between human and machine reasoning. In particular, such experiments would allow us to examine the nature of human-machine interaction and how we could beneficially use Machine Argumentation in the way that humans reason.

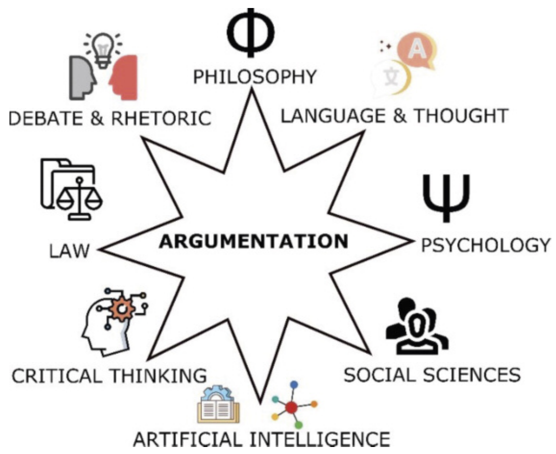

Figure 1: The interdisciplinary nature of Argumentation.

A program of experiments to examine the interaction of humans with COGNICA and its reasoning explanations has been setup. The aim of the experiments is to examine the effect that different types of machine explanations can have on human reasoning, ranging from no explanation to visual or verbal explanations, to summary explanations or extensive analytical explanations. A first pilot experiment has been completed in which participants were asked to answer questions based on some information that typically included three to five conditionals from everyday life. They were then asked to reconsider the questions after they were shown the answers of the system together with the explanations (verbal and/or visual) offered by the system. The results of this study show that in around 50% of the cases where the conclusion of the human participants differed from that of the machine, the participants changed their answer when they saw the explanations of the system. It was also observed that this kind of interaction with the system motivates participants to "drift" to more "careful reasoning" as they progress in the experiment, in accordance with the argumentation theory of Mercier and Sperber [1].

We are currently performing new experiments in order to investigate how this machine-human interaction varies across the population with different cognitive and personality characteristics. The interested reader can try this type of exercise, used in the COGNICA experiments, by visiting the links [L1, L2] and following the simple instructions given there.

Links:
[L1] http://kwz.me/hvR
[L2] http://kwz.me/hvU

References:

[1] H. Mercier, D. Sperber, "Why do humans reason? Arguments for an argumentative theory", in Behavioral and Brain Sciences, 34(2):57–74, 2011.
[2] E. Dietz, A. Kakas, L. Michael, "Argumentation: A calculus for Human-Centric AI", in Frontiers in Artificial Intelligence, 5, 2022. https://doi.org/10.3389/frai.2022.955579
[3] P.N. Johnson-Laird, R.M.J. Byrne, "Conditionals: a theory of meaning, pragmatics, and inference", in Psychological Review, 109(4):646–678, 2002.

Please contact:

Antonis Kakas , University of Cyprus, Cyprus
antonis@ucy.ac.cy

AI Meets Culture: Cognitive AI Enabling User Interaction in Multimodal Installation Art

by Christian Thomay and Benedikt Gollan (Research Studios Austria FG), Anna-Sophie Ofner (Universität Mozarteum Salzburg) and Thomas Scherndl (Paris Lodron Universität Salzburg)

We propose using Cognitive AI methods for use in multimodal arts installations to create interactive systems that learn from their visitors. Researchers and artists from Research Studios Austria FG and the Universität Mozarteum Salzburg have created a multimodal installation in which users can use hand gestures to interact with individual instruments playing a Mozart string quartet. In this article, we outline this installation and our plans to realise Cognitive AI methods that learn from their users, thereby improving the performance of gesture recognition as well as usability and user experience.

Cognitive AI – AI systems that apply human-like common sense and learn based on their interactions with humans [1] – is a field of AI that has plentiful applications in fields such as Internet of Things, cybersecurity, and content recommendation [2]. Especially in installation art, Cognitive AI offers the potential for a fundamental paradigm shift: allowing machines to learn to adapt interaction modalities to their visitors, instead of the visitors adapting to the installation. Approaching Cognitive AI in interactive systems, as a first step we explore natural, intuitive interaction in an artistic installation.

Computers and interactive art have a storied past, dating back to the 1950s, with early electronic art installations such as Nicolas Schöffer's CYSP-1 sculpture [L1] moving in reaction to the motions of the people around it. Since the rise of artificial intelligence, artists have begun to use AI algorithms to create visual, acoustic, or multi-sensory art, or to allow visitors to interact with installation art in different ways [3].

A foundation for many installation art projects is an intent to proactively involve their audience, allowing visitors to control and shape immersive experiences in the installation. Regarding how explicit this interaction is, there is a scale with two extreme points: "random" art systems, where art is created based on visitor behaviour using rules that are either opaque to the audience, or entirely non-deterministic; on the other end of the spectrum, there are explicit interactive systems, where the audience is meant to assume conscious control over the content of the installation. Modern AI art projects have fallen on various points on this scale, frequently encouraging the visitor to infer the rules without fully spelling them out. This leads to technical challenges for identifying user input, which may still be vague and fuzzy.

One such art installation is Mozart Contained! [L2], a project in the Spot on MozART [L3] program at the University Mozarteum Salzburg. The project was conceived in 2019 as an

Figure 1: Mozart Contained! Installation.

approach to create a new level of perception of the music of Mozart, with an interdisciplinary team of artists from Mozarteum and researchers from the Research Studio Austria creating this multimodal art installation (Figure 1). It was first presented in July 2021, and an enhanced version was opened to visitors in October 2022.

Visitors are able to use hand gestures to influence the individual voices of Mozart's dissonance string quartet (KV 465). In each corner of the installation there is a light sculpture representing one instrument (violin 1, violin 2, cello, viola). Visitors can interact with the light sculpture using hand gestures, controlling playback and volume of the individual tracks. The light sculpture gives feedback whether the activation was successful and also visually indicates the volume by using more or less pronounced lines and brightness. Visitors can activate each instrument separately: if they activate all four instruments, the quartet is played concurrently as intended; if instead only a subsample of the instruments is activated, visitors can focus solely on the respective instruments and thus experience Mozart's music in a new way, allowing both single and multiple visitors at the same time to create their own interpretation of the piece.

The technical setup consists of a sound system, four ultra short throw projectors creating the light sculptures, and two Microsoft Kinect skeleton tracking systems that observe the visitors' movements and thereby recognise input hand gestures. The interaction concept in Mozart Contained! is at this point entirely rule based: the system detects gestures performed by the audience and maps them to control inputs to the installation. However, here Cognitive AI offers the groundwork for a fundamentally novel approach.

In Mozart Contained!, Cognitive AI offers an inversion of the traditional interaction concept: instead of asking its users to

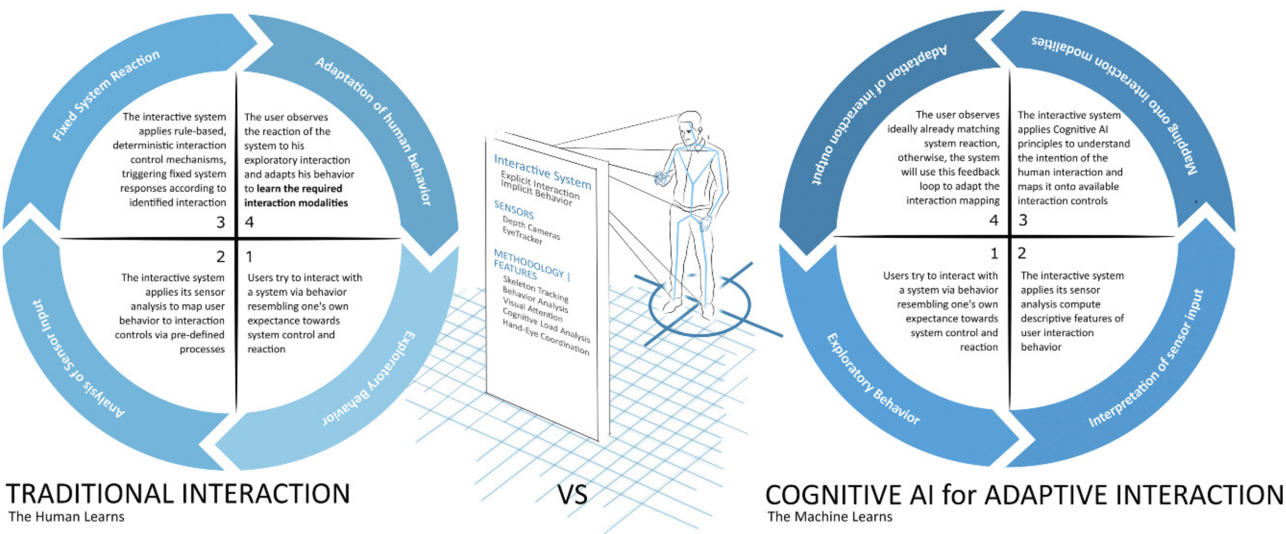

Figure 2: Schematic of traditional interaction vs interaction in a Cognitive AI system.

learn to use the system, the AI learning system infers its users' intent, rather than relying on exact execution of predefined input patterns. This has the potential to offer a much more intuitive experience true to the visitor's intent; Figure 2 illustrates this shift from traditional to Cognitive AI interaction.

The two active periods of Mozart Contained! so far have served as a fundamental validation step of the interaction concept, and as data gathering. Based on data collected so far and to be collected in the future, AI methods can be trained to use kinetic observables of visitors such as overall position in the installation, body/head/torso orientation, and arm or hand motion to infer which sculpture they wish to engage with, and with which specific intent. It is our aim to use the collected digital data, together with questionnaire data to serve as a baseline of intent and engagement, to create such a Cognitive AI model, and create a learning interactive system that learns together with its audience. Challenges such an AI model would face include the non-deterministic nature of human behaviour and the danger of feedback loops, where human participants adapt to the AI system before the AI can learn from their behaviour; this will be addressed via careful curation of algorithm input data, evaluating participant questionnaire data, and dedicated training sessions.

Consequently, realising such a Cognitive AI-based adaptive interaction would serve as a fundamental improvement in terms of human-centred technologies and interaction quality, and represent a fascinating use case example of how Cognitive AI can enhance and improve user experience.

Links:
[L1] http://dada.compart-bremen.de/item/artwork/670
[L2] https://www.spotonmozart.at/en/projekt/mozart-contained/
[L3] https://www.spotonmozart.at/en/

References:
[1] Y .Zhu, et al., "Dark, beyond deep: A paradigm shift to cognitive AI with humanlike common sense", in Engineering 6.3 (2020): 310-345. https://doi.org/10.1016/j.eng.2020.01.011
[2] K. C. Desouza, G. S. Dawson, and D. Chenok, "Designing, developing, and deploying artificial intelligence systems: Lessons from and for the public sector", in Business Horizons 63.2 (2020): 205-213. https://doi.org/10.1016/j.bushor.2019.11.004
[3] Jeon, Myounghoon, et al., "From rituals to magic: Interactive art and HCI of the past, present, and future", in Int. Journal of Human-Computer Studies 131 (2019): 108-119. https://doi.org/10.1016/j.ijhcs.2019.06.005

Please contact:
Christian Thomay
Research Studios Austria FG, Austria
christian.thomay@researchstudio.at

Piktor-O-Bot: The Robotic Face-Drawing Solution

by Tamás Cserteg, Anh Tuan Hoang, Krisztián Kis, János Csempesz (SZTAKI) and Zsolt János Viharos, (SZTAKI and John von Neumann University)

The combination of cognitive AI and robotics has already had a significant impact on the way we live. Novel developed solutions not only cover industrial applications, but are widely spread in the field of social robotics as well. To bring cognitive AI and robotics closer to people not exposed to these technologies on a day-to-day basis, a portrait-drawing demonstrator, called Piktor-O-Bot has been developed at SZTAKI. Over the past year and a half, portraits of more than a thousand people have been drawn, which shows great interest in the demonstrator and that people are genuinely interested in cognitive robotic applications.

As cognitive AI and robotics technology continues to advance and becomes more affordable, it is likely that we will see an even greater spread of cognitive robotics applications and use-cases. It is important to educate those who do not have frequent exposure to robots and AI applications to get a better understanding of these technologies, so they can welcome new generations of robotic applications.

The Piktor-O-Bot portrait-drawing demonstrator developed at SZTAKI provides a unique and entertaining experience for visitors, while demonstrating state-of-the-art AI and robotic technologies [L1]. It features machine-learning-based image processing for generating the to-be-drawn image, optimal path planning for quick drawing, and force-feedback robot control for precise execution. The camera takes a picture of a visitor, then our specialised image-processing algorithm replicates the person's facial features as a set of lines (Figure 1). The drawing time can be minimised by specifying the order of the lines. Ordering is carried out with a sequence-planning algorithm based on a VRP (vehicle routing problem) solver. The drawing itself takes place on a calibrated board, with a simple brush pen and no additional instruments. The cobot is programmed to draw each line individually while maintaining a constant force between the board and the pen. To draw as many portraits as possible even on busy days at a fair, the process overall is designed to be quick – an average drawing takes about two minutes to draw.

The demonstrator gained substantial attraction since the beginning of its development [L2]. It has been featured at many exhibitions including conferences and fairs across Hungary and even at the 2nd Stuttgart Science Festival in Germany [L3], reflecting high interest of people in cognitive robotics and insights into the fields of artificial intelligence and robotics (Figure 2).

The portrait-drawing process flow of Piktor-O-Bot is a much more complex process than it would be thought on first impression. It consists of seven image-processing steps (numbers of the steps reflect the numbering in Figure 3), but two steps are realised with the same artificial neural network model, so their index is the same.

Face detection: It is the first step of the portrait drawing process, which includes both face recognition and localisation. For this purpose, RTNet [1] was used, which is a pre-trained, convolution-based, multi-level, pyramidal neural network that can recognise human faces in the image in real time with high accuracy:
- Background removal: A pre-trained neural network called MODNet [2] removes irrelevant objects from the portrait point of view from the image, leaving only the human face and human body.
- Contour highlighting: A modern state-of-the-art, convolution-based neural network, DexiNed [3] was used to efficiently highlight facial contours.
- Face segmentation: Segmenting different parts of the face, e.g., eyes, mouths, and noses is also a difficult task, which was solved again by the RTNet.
- Iris detection: The eye is a relatively sensitive part of portrait drawing, thus iris and pupil detection were integrated into the process yielding to create an even more lifelike portrait.
- Edge thinning: The Zhang-Suen thinning method was used for edge thinning, which is used to reduce the thickness of edges.
- Vectorisation: The raster facial contour images obtained (as a result of the process) are transformed into a set of mathematically describable lines and curves by the AutoTrace program.

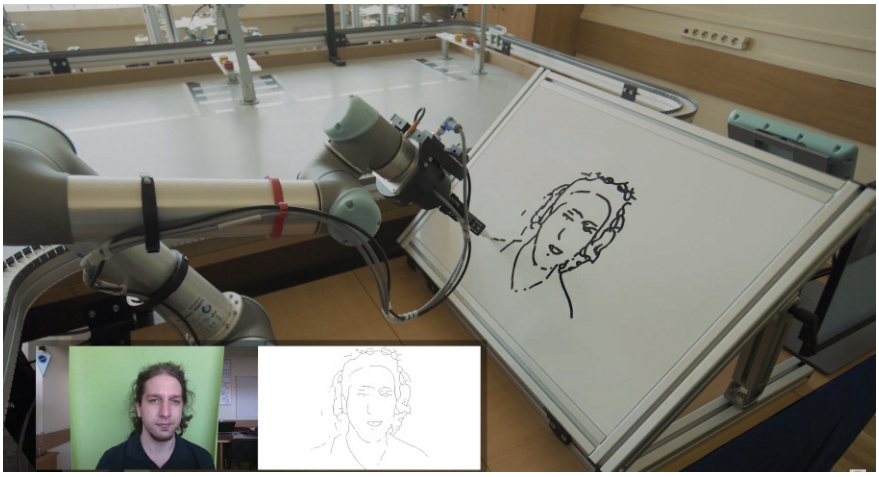

Figure 1: From the real face through the vector-graphic until the final drawing realised by the Piktor-O-Bot demonstrator of SZTAKI.

Figure 2: High interest of people for cognitive robotics presented in various public fairs and events.

Links:
[L1] https://www.youtube.com/watch?v=8ULIP_5nEJ0
[L2] https://www.youtube.com/watch?v=NhOibxWIpH0
[L3] https://www.facebook.com/watch/?v=562239958892035

References:

[1] Y. Lin, et al., “RoI Tanh-polar transformer network for face parsing in the wild”, in Image and Vision Computing, 2021.
[2] Z. Ke, et al., “MODNet: Real-Time Trimap-Free Portrait Matting via Objective Decomposition,” in Association for the Advancement of Artificial Intelligence, 2022.
[3] X. Soria, E. Riba, A. Sappa, “Dense Extreme Inception Network: Towards a Robust CNN Model for Edge Detection”, IEEE Winter Conference on Applications of Computer Vision (WACV), 2020.

Please contact:

Tamás Cserteg, SZTAKI, Hungary
cserteg.tamas@sztaki.hu

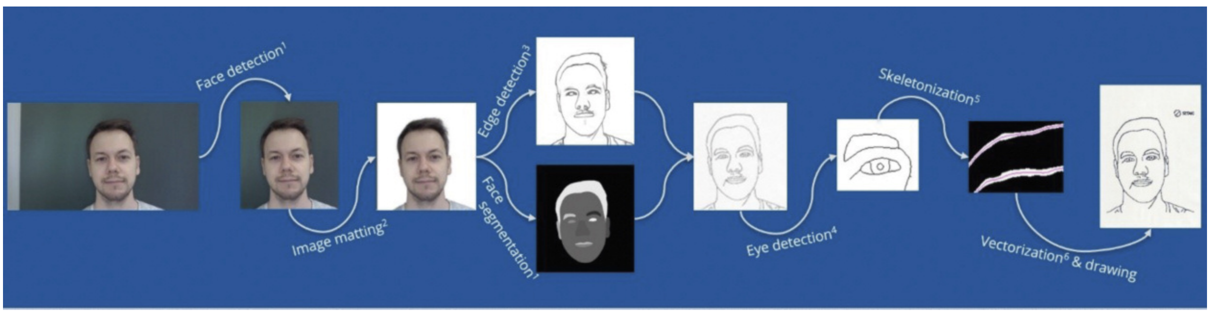

Figure 3: The portrait drawing process flow of Piktor-O-Bot consists of seven image-processing steps; two of them apply the same model.

How AI can Exploit Body Language to Understand Social Behaviour in the Wild

by Hayley Hung (Delft University of Technology)

Researchers at the Socially Perceptive Computing Lab at Delft University of Technology in the Netherlands investigate novel ways to measure our quality of experience in social encounters. Their aim is to understand the difference between good and bad social encounters, from speed-dates, to professional networking events, to long-term simulated space missions. This paves the way for the development of technologies to understand and ultimately improve our social encounters.

Have you ever been anxious about attending social networking events where it wasn't clear if you might meet the right people or the whole thing might have been a waste of time? We spend so many hours of our lives in face-to-face conversations and yet most of the time, we have no control over how well used that time is. The MINGLE project (Modelling Group Dynamics in Complex Conversational Scenes from Non-Verbal Behaviour) has been involved in developing new approaches to automatically analyse conversational interactions in large social gatherings such as professional-networking events. Research has shown that attending such events contributes greatly to career and personal success. However, it can often feel like there is too much serendipity involved or that how a conversation goes with a new acquaintance can determine whether a future relationship will be fruitful or not. Once machines can automatically interpret what is happening, they can start to help us to navigate our social experiences better. While much progress has been made in the development of automated analysis tools of small pre-arranged meetings, scaling up robustly to settings outside of meetings such as professional-networking events has remained an open and challenging research problem.

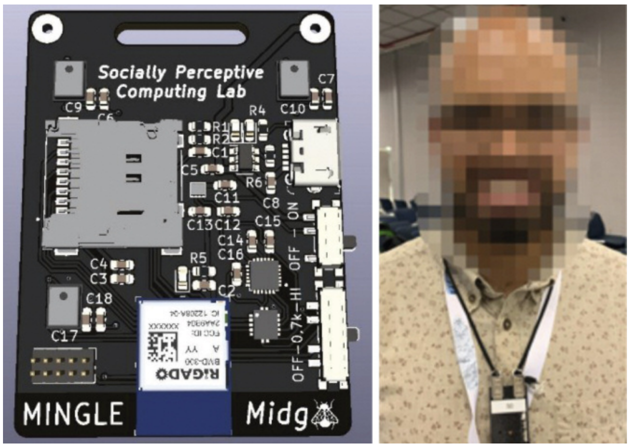

Figure 1: The MINGLE Midge Smart ID badge (left). Demonstration of how it is worn with a black lanyard around the neck (right). It contains a 9-degree-of-freedom inertial measurement unit, Bluetooth transmitter and receiver for measuring proximity to other sensors, and the possibility of recording high or low frequency audio for assisting in the training of body-language-based models. The design is open source. More details and access can be found at [1] and [L1].

Moving Out of the Lab and into the Wild

One of the problems was that scaling up looking at social behaviour outside of lab conditions greatly changes the nature of the machine inference problem. In lab-based settings, former smart meeting rooms (See Figure 2) have been equipped with a microphone with clean audio data per person in the meeting, and a camera capturing each person's face, gaze, and body behaviour (because they are mostly seated). The lighting can be easily controlled and we know beforehand that there is a single conversation occurring and who the participants are. We also have access to a rich body of social science literature that can help to inform and inspire the design of appropriate machine learning strategies to automatically estimate the role of participants, who is dominant, what someone's personality is, who is disagreeing with whom, or even how cohesive the team appears.

Scaling up to mingling settings and moving outside of lab conditions means that crowd density can increase, and multiple simultaneous conversations can occur in the same room (see Figure 3). Sensor data is more likely to capture cocktail party

Figure 2: Snapshot taken from the AMI meeting corpus. This shows example footage taken from an instrumented meeting room where more controlled lab experiments can occur.

Figure 3: Example snapshots of data we used for capturing and training machine learning models to detect laughter from a real-life professional-networking event [3]. Faces have been blurred to protect the identity of the participants.

noise in audio and high levels of occlusion, resulting in people being visually hidden behind other people in video. Determining who is talking with whom is more challenging because groups can split and merge at will. Finally, in such mingling settings, people are there for their own motivations so recording sensitive information in private conversations may not be as appropriate. Moreover, recording the spontaneous facial expressions can feel quite invasive. A solution to this would require an unconventional approach – what if a single wearable sensor capturing your body movements was enough?

Finding Unconventional Solutions Guided by Social Science

Despite a seemingly impossible task, there are some tantalising hints from social science that may enable us to reimagine how machines might perceive the world. Rather than using the human senses (sight, hearing) as the primary inspiration for artificial social intelligence, the MINGLE project exploited findings from social science related to the role of body language in all aspects of social communication. The key was to investigate how body movements could be captured, even in environments that would be considered classically extremely challenging for computer vision or speech processing problems. Fortunately, at conferences or professional-networking events, participants often wear an ID badge on a lanyard around their neck. Imagine if this ID badge was smart and was equipped with sensors that could capture your body language? As a result of this, a Smart ID badge was designed: the open-source Midge Badge [L1],[1] (see Figure 1).

Social Interpretation of Conversations using Wearable Accelerometers

One of the most foundational indicators of conversation involvement is the pattern of speaking turns that participants in a conversation have. The more frequent the turns, the more engaged the group is, the more equally spread the turns are, the more likely that everyone is able to express their opinions. Without access to the audio, how does one estimate if people are speaking? Fortunately, social scientists have long since cited a relationship between speaking and head and hand gesturing. The project took this as inspiration, showing that just a single accelerometer embedded inside the Midge Badge was enough to capture when someone was speaking [1]. Combining this with body language represented by skeletons extracted from body key points led to an improvement in the estimation of speaking [1]. Going one step further, we also built an approach to measure the quality of the conversation based on their joint movement patterns [2], and to detect different intensities of laughter [3].

The Remaining Open Challenges

So far, the speaking behaviours that we have trained our systems for have been labelled based on observing the videos only. We need to capture data with high-quality audio to train the machine learning systems with higher-quality labels. We are currently missing important information about shorter and more subtle speaking turns, such as backchannels, which convey important information such as that someone is actively listening to a speaker. To truly understand the ebb and flow of conversations, having systems that can accurately capture short and long speaking turns is crucial. Further research will be required to see if these backchannels can be detected better using multi-modal sensor data and then whether combining longer and shorter turns, laughter, and coordination behaviours together can improve estimates of conversation quality.

The work described in this article was carried out in collaboration with Stephanie Tan, Jose Vargas Quiros, Chirag Raman, Ekin Gedik, Ashraful Islam, Navin Raj Prabhu, Catharine Oertel, and Laura Cabrera Quiros. It was funded by the Netherlands Organization for Scientific Research (NWO) under project number 639.022.606 with associated Aspasia Grant.

Link:
[L1] https://kwz.me/hwT

References:

[1] C. Raman, et al., "ConfLab: A Data Collection Concept, Dataset, and Benchmark for Machine Analysis of Free - Standing Social Interactions in the Wild", in Proc. of the Neural Information Processing Conf., 2022.

[2] C. Raman, N. Raj Prabhu, and H. Hung, "Perceived Conversation Quality in Spontaneous Interactions", in arXiv preprint, 2022. https://arxiv.org/abs/2207.05791.

[3] J. Vargas-Quiros, et al., "Impact of annotation modality on label quality and model performance in the automatic assessment of laughter in-the-wild", in arXiv preprint, 2022. https://arxiv.org/abs/2211.00794

Please contact:

Hayley Hung
Delft University of Technology, The Netherlands
h.hung@tudelft.nl

Personalisation of Humanoid Robots: Serious Games for Older Adults Based on Biographical Memories

by Benedetta Catricalà, Marco Manca, Fabio Paternò, Carmen Santoro and Eleonora Zedda (ISTI-CNR)

We present an approach to novel digital cognitive training through serious games able to adapt to personally relevant material from the older adult's life. The games are based on memories associated with the older adult's biography, thus making interactions personalised, relevant, and more engaging. The serious games are accessed through humanoid robots, which can make the training exercise more engaging because of their human-like behaviour.

Currently, cognitive stimulation exercises conducted with mild cognitive impairment (MCI)-affected older adults are typically administered by care professionals through exercises, which require a high degree of involvement by professional therapists and are often perceived as tedious by older adults.

One way to overcome such issues is to offer older adults opportunities to engage in mentally stimulating activities. Though researchers have found that digital games may provide benefits for some cognitive domains, the game industry seems to have ignored this important part of the population as a special category of users. Tahmassebi [1] found that older adults complain about the complexity of both hardware and software, which can be problematic due to their age-related cognitive and physical limitations. So far there have been mixed results related to game-based interventions [2]. Unfortunately, despite the potential benefits to cognition described above, older adult populations are generally not targeted audiences when it comes to the design of digital games. However, we have found [3] that emerging humanoid robots may open up new possibilities in more effectively engaging MCI older adults during repetitive cognitive training. While the use of multiple modalities such as speech and visual interfaces has already been explored for older adults, social robots have recently shown to have the potential to support care and independent living because they are now capable of exhibiting natural-appearing social characteristics. Thus, they can help formal and informal caregivers to monitor and support older adults, in particular those with MCI, while also providing a mixture of empathy, motivation, encouragement and companionship.

For this purpose, we have designed a novel solution of serious games, based on users' memories and played through the social humanoid robot, which aims to help older adults to maintain their cognitive functional level, and prolong independent living. The serious games, installed on the humanoid robot, motivate older adults by engaging them in playful situations that draw on their personal memories with which they can interact. Since the games are based on elements associated with

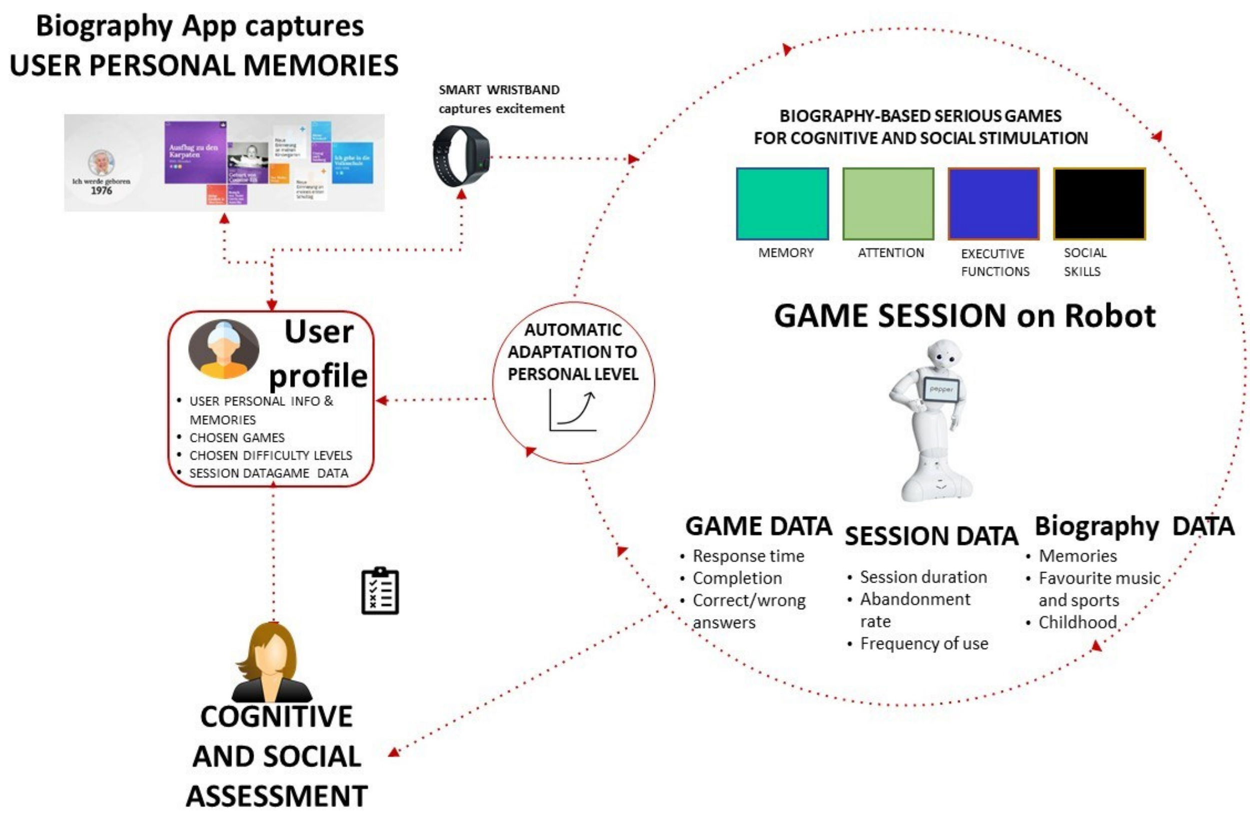

Figure 1: The SERENI Platform.

the biography of the older adults, they can make interactions more relevant and more likely to keep the older adults engaged while enhancing their well-being as well.

Design

For the motivations discussed in the previous section, in the context of the SERENI (SERious gamEs with humanoid robots in cogNItive training) CNR Project, we have designed a platform (see Figure 1) to deliver serious games able to exploit personally relevant material from an older adult's life through a humanoid robot. It aims to stimulate cognitive functions through play sessions, which should last about 20 minutes. The exercises should be useful for making the participants think and reason to provide the correct answers. The platform can be a solution for day-care centres where older adults with mild cognitive impairments can go to perform relevant exercises.

The platform is based on various components. The first one is the Remind App, which is a responsive multimodal web application to collect memories from the older adults and their relatives, which can be entered by both graphical and vocal interaction.

Six categories of memories have been identified:
- Music-related, concerning songs or singers of their youth
- Game-related, activities they liked to play
- Event-related, specific events that they still remember
- Locations, which are associated with particular memories of travels
- Food-related, with the possibility to indicate the corresponding recipes by listing the steps that should be performed
- Affects, memories related to family members and loved people.

The entered memories are stored in a database from which they can be automatically retrieved and exploited to provide contents for the games to play with the humanoid robot. We have used a Pepper robot for proposing the games. It is a 1.2-metre-tall robot with the ability to use various interaction modalities (speech, touch, gestures, graphics through a screen on the chest).

An initial set of four types of games have been identified:
- Memory completion, where a memory presented by the robot has a missing detail, which should be indicated by the user
- Activities ordering, which can be applied to food recipes, in which a list of relevant tasks should be put in the right order by the user
- Associate memories, where 3–4 memories are briefly described along with a set of details, and the user has to connect each memory with the corresponding detail (for example associate song titles with the corresponding singers)
- Memory-related question, where the user has to answer specific questions about an event drawn from the memories (for example: what happened in 1945?).

The platform is also able to store some data regarding the user performance, such as when and for how long the user played with a given game, the number of errors made in a session, and the type of games that have been played. Such information can be delivered to older adults and their caregivers both through the robot and the web app.

Conclusions

We have presented an approach to memory-based games delivered through a Pepper humanoid robot. We have identified a first set of memory and game types, which can be extended in the following. A first prototype of the platform for collecting the memories, and exploiting them in a set of games for cognitive stimulation has been delivered, and is under test with older adults with mild cognitive impairments in a Train the Brain programme to assess their user experience with such games.

Links:
[L1] https://hiis.isti.cnr.it/lab/home
[L2] https://hiis.isti.cnr.it/sereni/index.html

References:

[1] S. Tahmassebi, "Digital Game Design for Elderly People", in Digitala Vetenskapliga Arkivet, 2018.
[2] K. A. Blocker, T. J. Wright, W. R. Boot, "Gaming preferences of aging generations", in Gerontechnology: International Journal on the Fundamental Aspects of Technology to Serve the Ageing Society, 12(3), 174, 2014.
[3] M. Manca, et al., "The impact of serious games with humanoid robots on mild cognitive impairment older adults", in International Journal of Human-Computer Studies, 2021.

Please contact:

Fabio Paternò, CNR-ISTI, Italy
fabio.paterno@isti.cnr.it

Novel Deep Neural Architecture Search Algorithm for Human Activity Recognition

by Anh Tuan Hoang (SZTAKI) and Zsolt János Viharos (SZTAKI and John von Neumann University)

Human Activity Recognition (HAR) is a prevalent topic in the field of cognitive intelligence, closely related to Human-Centred Artificial Intelligence (HCAI). The emergence of deep learning has brought breakthroughs in many HAR problems. This research proposes a novel search algorithm to determine the optimal deep neural model structure that considers the internal relationships as well as the information content incorporated in the data, and as a result serving with superior capabilities in recognition of various human activity types.

Human-Centred Artificial Intelligence is a rapidly developing and increasingly popular topic that puts humans at the focus of attention when building AI systems. It seeks to preserve human control in a way that meets the needs of artificial intelligence. The applicability of HAR lies in the fact that during any type of movement, the human body generates a characteristic sensor signal, in other words, a pattern, and it can be effectively recognised with the help of combined digital signal processing and deep learning.

The idea of using deeper neural networks and new types of neural layers, such as convolutional layers, is gaining popularity among researchers. These new learning methods automatically form a higher-level representation from the raw sensory dataset. As a result, deep learning provides a more general solution, as it automates the extraction of features, as opposed to classical machine learning approaches, where the features are extracted statically.

An interesting concept is followed by Abid et al. [1] and Wu and Cheng [2] who developed an efficient feature-selection method based on autoencoder. Both studies minimise the cost of reconstruction error based on the selected features; however, the use of different types of deep models and the correlation of internal parameters are not taken into account. Moreover, their success has not yet been validated on sensory signal data. By the concept for internal-relationships exploration inside the data, we can achieve significant superiority.

To explore the internal relationships incorporated in the data, by deep neural architecture search, a modified adaptation of the classic autoencoder model was used, which can determine the appropriate input-output configuration and is generally used to learn efficient data coding in an unsupervised manner (the preliminary research for shallow network was published by Viharos and Monostori [3]). 1D convolutional and 1D

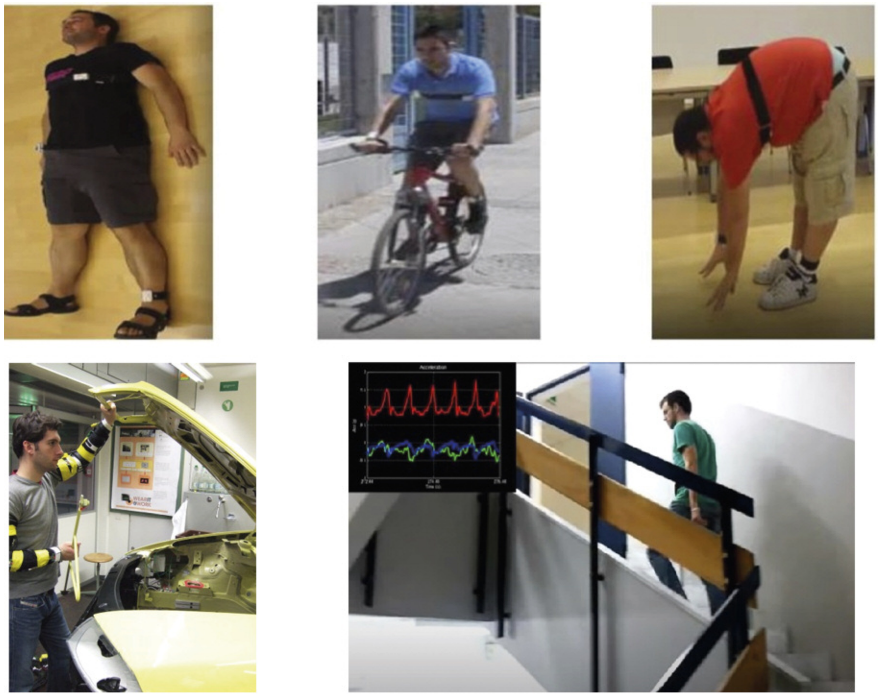

Figure 1: Example movements and measuring sensors of the human activity recognition datasets. The three pictures on the top represent the MHEALTH dataset, the bottom-left picture shows typical activity of the SKODA dataset (realised during reparation of SKODA cars), and the picture in the bottom-right shows a daily activity of people measured by their hand brought typically in their pockets.

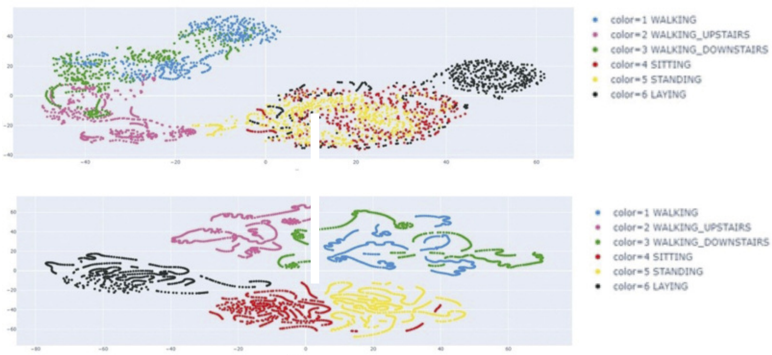

Figure 2: The T-SNE representation of the learned dependencies on the Smartphone dataset. The classical convolutional neural network (CNN) architecture resulted in the classification results shown in the top picture, while the proposed, novel deep neural architecture search (NAS) based solution results are represented in the bottom part of the picture trained on the same dataset. It can be seen clearly that the proposed algorithm realises a more distinctive, superior modelling solution.

transposed layers were used for efficiently processing sensory signal data.

Thanks to the widespread use of modern wearable sensors, more and more public datasets became available to researchers about various human activities. Nowadays, several different datasets can be freely used, including MHealth [L1], Skoda [L2], Smartphones [L3] on which the proposed algorithm was validated (Figure 1).

After the sensory data were pre-processed, training was carried out by the proposed novel method, which selects the most significant and descriptive features by exploring the internal dependencies. If the reconstructed data is added to the input, which brings the internal relationships of the data, significant improvement can be achieved. For the evaluation, a shallow convolutional neural network is used as a reference classifier.

Improvements for all the benchmarking datasets we have chosen were achieved, thereby validating the success of the proposed method. The correctness of the novel algorithm was also validated by visualisation on the Smartphone dataset. For the visualisation, T-SNE dimensionality reduction procedure was used, thereby achieving the possibility of 2D representation (Figure 2).

Links:
[L1] http://archive.ics.uci.edu/ml/datasets/mhealth+dataset
[L2] http://har-dataset.org/doku.php?id=wiki:dataset
[L3] https://archive.ics.uci.edu/ml/datasets/human+activity+recognition+using+smartphones

References:

[1] A. Abid, M. F. Balin, J. Y. Zou: "Concrete Autoencoders for Differentiable Feature Selection and Reconstruction", International Conference on Machine Learning (ICML), 2019.
[2] X. Wu, Q. Cheng: "Fractal Autoencoders for Feature Selection", AAAI Conference on Artificial Intelligence, 2020.
[3] Zs. J. Viharos, L. Monostori, "Automatic input-output configuration of ANN-based process models and its application in machining", in Lecture Notes of Artificial Intelligence - Multiple Approaches to Intelligent Systems, Springer, pp. 659-668., 1999.

Please contact:

Zsolt János Viharos, SZTAKI and John von Neumann University, Hungary
viharos.zsolt@sztaki.hu

László Hoang (Anh Tuan Hoang), SZTAKI, Hungary
hoang.laszlo@sztaki.hu

CALL FOR PAPERS

FMICS 2023: 28th International Conference on Formal Methods for Industrial Critical Systems

Antwerp, Belgium, 20-22 September 2023

FMICS is the annual conference organised by the homonymous ERCIM working group. The aim of the conference series is to provide a forum for researchers who are interested in the development and application of formal methods in industry. FMICS brings together scientists and engineers who are active in the area of formal methods and are interested in exchanging their experiences of the industrial usage of these methods. The FMICS conference series also strives to promote research and development for the improvement of formal methods and tools for industrial applications. As in previous years, FMICS 2023 is part of the CONFEST umbrella conference. In addition to FMICS, CONFEST also comprises CONCUR, FORMATS, and QEST, as well as pre- and post-conference workshops and tutorials. CONFEST will be held in Antwerp, Belgium on 18-23 September 2023.

Topics

Topics of interest include: Case studies and experience reports on industrial applications of formal methods; Methods, techniques, and tools to support automated analysis, certification, debugging, descriptions, learning, optimization, and transformation of complex, distributed, real-time, embedded, mobile and autonomous systems; Verification and validation methods that address shortcomings of existing methods with respect to their industrial applicability; Impact of the adoption of formal methods on the development process and associated costs; Application of formal methods in standardization and industrial forums.

Deadlines

The paper submission deadline is 15 May 2023. Accepted papers will be included in the Conference Proceedings published as part of Springer's Lecture Notes in Computer Science series. An award will be presented to the authors of the submission selected by the Program Committee as the FMICS 2023 Best Paper.

Programme Committee Chairs

- Alessandro Cimatti (FBK, Italy)
- Laura Titolo (NIA/NASA LaRC, USA).

More information:
https://www.uantwerpen.be/en/conferences/confest-2023/fmics/

Submission link:
https://easychair.org/conferences/?conf=fmics2023

AI-Enabled Services for Reconnaissance

by Refiz Duro, Rainer Simon (AIT Austrian Institute of Technology GmbH) and, Christoph Singewald (Syncpoint GmbH)

Climate change, pandemics and unstable geopolitical and economic circumstances on the global level are complex challenges necessitating approaches leveraging technological advances and human cross-domain expertise and experience. One piece of the puzzle addressing these challenges is to provide efficient services assisting decision and policy makers with insights extracted from the available data. Combining AI-based natural language processing and computer vision services with human-centred annotation and information enrichment service to build knowledge graphs for reconnaissance has a prospect to make a difference in the safety and security domain. But what does such an implementation look like, and shouldn't such services already be integrated into most decision-making processes?

We are directly witnessing the impacts that current world-changing events have on our daily lives. Pandemics have and still are affecting the movement of people and goods globally, climate change has brought more frequent and severe events and destruction, while the geopolitical situation has forced the EU and state governments to re-evaluate their current state of defence capabilities.

The commonality of the events mentioned is the necessity of acquiring insights for accurate and timely decision and policy making. Considering defence, the acquisition, processing, and analysis of data for creating intelligence are military operations' backbone. The dynamic nature of reconnaissance operations, however, require that the information be delivered as efficiently as possible, meaning with no or minimal delay and in a form that is easily consumable. To achieve this requirement and to leverage the increasing amount of available data and data sources, it is necessary to embrace technological advancements. They offer the potential to process, analyse, filter, and deliver information much more efficiently to ultimately accelerate the decision-making process.

The PIONEER project funded by the Austrian Defence Research Programme, has taken on the challenge of improving automation and digitalisation processes for information flow for reconnaissance and intelligence of Austrian Armed Forces (AAF). The initial overview of capabilities and practices has shown that the available ICT support for automated execution of the intelligence gathering and distribution can be significantly improved. For example, it can be challenging to quickly merge, structure and analyse collected sensor data (UAV footage, written reports) due to heterogeneity of data formats and multimodalities, limiting data fusion necessary for advanced analytics. Furthermore, although the experienced personnel involved in reconnaissance activities can quickly and accurately establish intelligence products using traditional methods (e.g., populating operational maps on walls with geotagged information), the ever-increasing amount of available

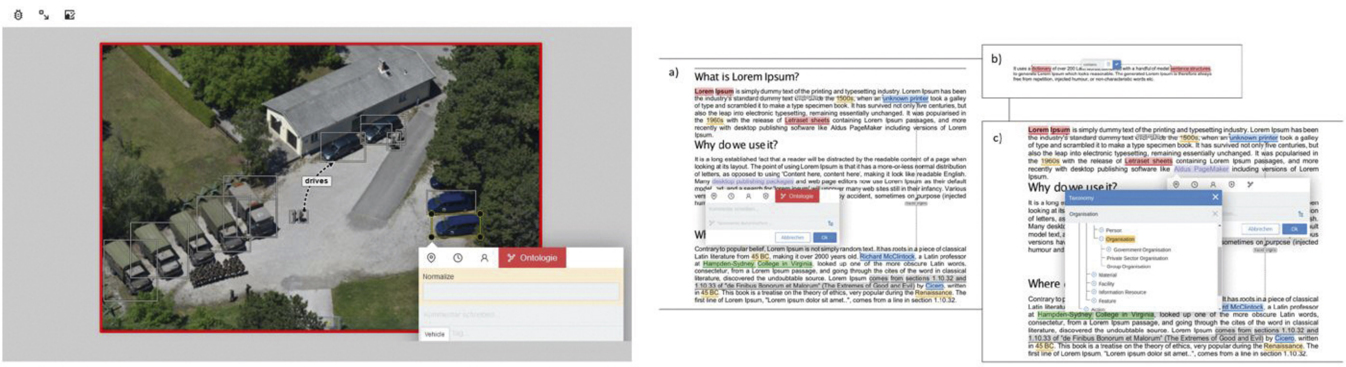

Figure 1: Annotation results are highlighted through boxes in photographs (left) or are highlighted with colours in text-based data (right; from [1]). Annotations can be set in relations ("person drives car") and enriched with additional information through taxonomy feature.

data and diversification of data sources are forcing the practices to be enhanced through integrating state-of-the-art information extraction technologies.

Due to the selected type of data (PDF reports and photographs), two technological fields heavily integrating Artificial Intelligence (AI) have seemed to be the most obvious for the experimental try-outs: computer vision and natural language processing (NLP). In the first one, the focus is on the object detection, leveraging an open-source machine learning model, which uses the YOLOv3 architecture for multi-class object recognition, and weights trained on the Microsoft Common Objects in Context dataset. The photographs provided by the project partner AAF are typical for reconnaissance operations. The applied model on these images shows promising results, i.e., vehicles (e.g., trucks) and persons, which are two categories of high intelligence value, are detected with high accuracies. The focus of NLP in the context of the project is on the extraction of entities through named entity recognition (NER). The basic idea is to automatically tag and classify words in a document according to predefined categories such as person or place. Similar to object detection, there are available open-source libraries for building suitable NER pipelines. In our case, the documents are reconnaissance reports integrating information collected by the field scouts and applying the model on these reports again shows promising results (Figure 1).

Further processing implies services exploiting the extracted information moving in the direction of building knowledge graphs (KG) through enhancing NER and entity and event linking. A KG enables semantic search and serves as a basis for different forms of representation – spatial, temporal and relational, supports analysts in recognising patterns and drawing conclusions – and for automated reasoning. Within the project, the automatically generated annotations are mapped to an ontology (schema.org) via classifiers, converted to RDF 1.1 N-Quads and integrated into the existing KG, thus allowing for automated spatial and temporal patterns and anomaly detection, detection of otherwise difficult-to-detect networks of groups of people and organisations, and in general knowledge engineering.

To conclude, integrating AI capabilities into the intelligence-gathering process looks promising. Recent advances have made it possible to achieve respectable gains with little effort, so that open-source models can be used relatively easily for tasks intended for reconnaissance systems. The suitability of these models is, however, not optimal, as misclassification can occur due to inaccuracies or training datasets not including, e.g., "tank" as a class, but "truck" class. This requires updating the existing model through "transfer learning" so that it is more appropriate for the context of reconnaissance and daily environments relevant for the AAF. It is critical that the objects and entities found in images and documents are linked and set in a context through additional information enrichment. This is not achieved by the object detection and NER components, but by open-source annotation libraries RecogitoJS and Annotorius [L1]. The tool necessitates a human user to, e.g., link the NER-detected entities through relations ("entity-1 is seen at the location entity-2"), and thus brings in the concept of human-centred in the process, making the system less prone to errors due to uncertainties and erroneous decisions. Implementing a KG in an RDF Quad Store is limiting, as the methods for network analyses resembles a property graph. In order not to lose the advantages of an RDF store, we consider using a hybrid model in future. We conclude that AI-based natural language processing and computer vision services with human-centred annotation and information enrichment service for reconnaissance has a prospect to make a difference in the safety and security domain.

Link:
[L1] https://recogito.github.io/annotorious/

Reference:
[1] G. Chroust, P. Doucek and, V. Oškrdal (editors), "IDIMT-2022 Digitalization of society, business and management in a pandemic : 30th Interdisciplinary Information Management Talk", in 30th Interdisciplinary Information Management Talks, Prague, 2022. https://doi.org/10.35011/IDIMT-2022

Please contact:
Refiz Duro,
AIT Austrian Institute of Technology GmbH, Austria
refiz.duro@ait.ac.at

The Interactive Classification System

by Andrea Esuli (ISTI-CNR)

ISTI-CNR released a new web application for the manual and automatic classification of documents. Human annotators collaboratively label documents with machine learning algorithms that learn from annotators' actions and support the activity with classification suggestions. The platform supports the early stages of document labelling, with the ability to change the classification scheme on the go and to reuse and adapt existing classifiers.

The Interactive Classification System (ICS) [1] is a novel open-source software [L1] from ISTI-CNR that is designed to support the activity of manual text classification. The application uses machine learning to continuously fit automatic classification models that are in turn used to actively support its users with classification suggestions, or to produce automatic classification of large datasets.

The aim of ICS is to support the activity of manual text classification, in collaborative scenarios, with a special attention to the early stages of labelling, i.e., when the dataset may be not yet complete, the classification schema is not yet consolidated. These scenarios receive less attention from the research on machine learning methods for text classification, which typically assume a more solid setup on which to operate.

ICS can be deployed as a web application (see Figure 1) and multiple users can collaboratively build datasets, define classification schemas, and label the documents in the datasets according to any of the classification schemas. A unique feature of ICS is that it gives its users total freedom of action: they can at any time modify any classification schema, any dataset, and any label assignment, possibly reusing any relevant information from previous activities. During these activities a machine-learning agent monitors those events and keeps a pool of automatic classifiers up to date.

Such freedom of action has an impact on how the machine learning algorithms used in the system must operate. The machine-learning approach used for ICS can be defined as "unobtrusive machine learning", as it never actively requests any action from the users, while it instead silently observes their actions, continuously adapting the automatic classification mod-

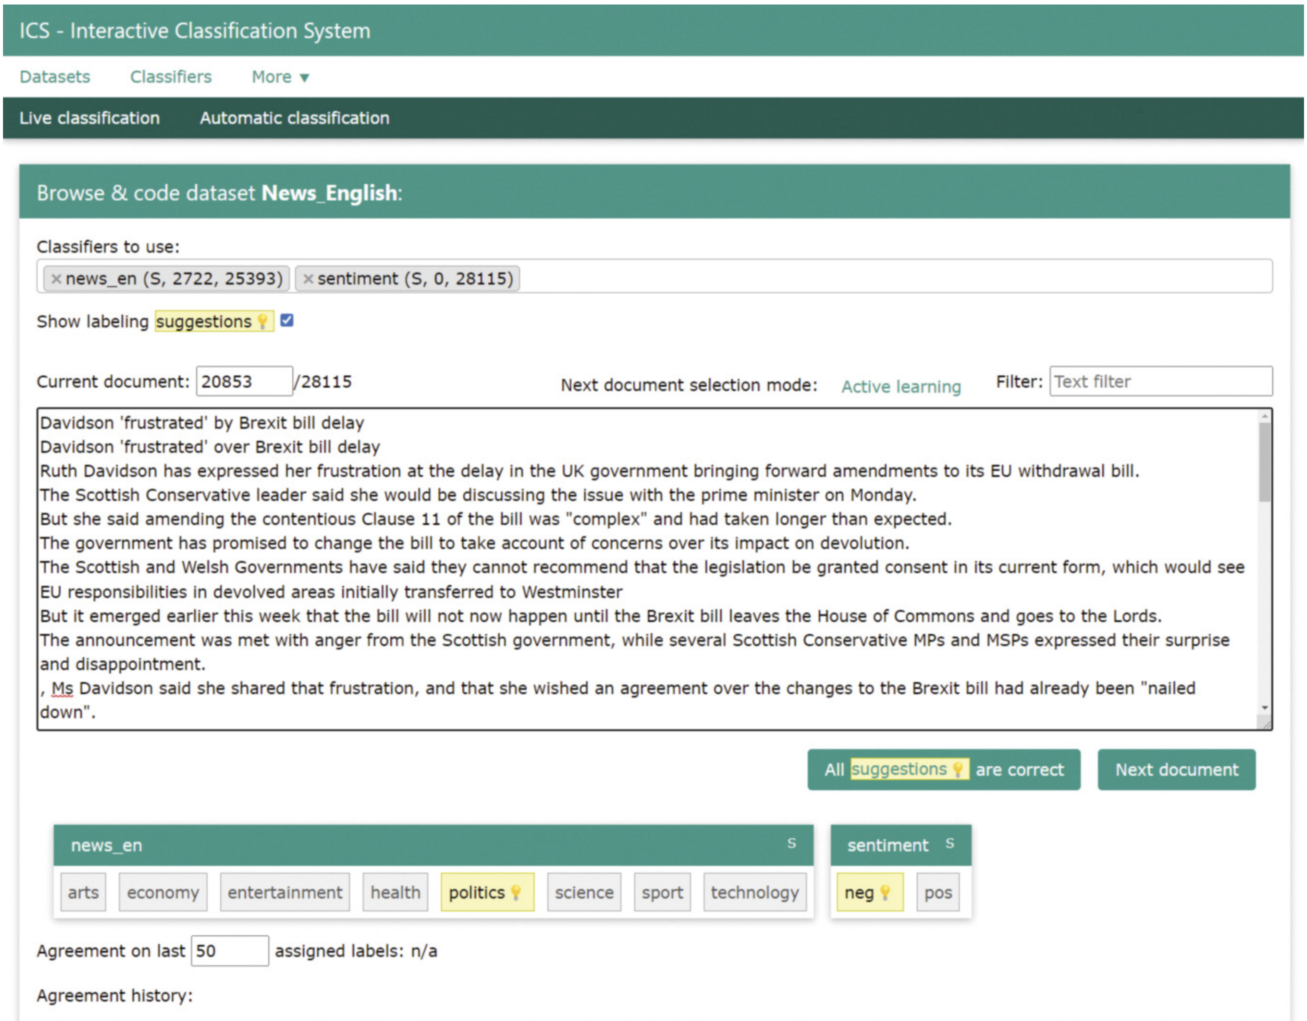

Figure 1: A screenshot of ICS. A document is shown, with two classifiers selected for labelling, i.e., by topic and by sentiment. Labels highlighted in yellow are the ones suggested by the machine learning model that ICS continuously updates according to users' actions.

els that are in turn used to provide the classification suggestions.

The unobtrusive approach challenges many of the assumptions made by the typical machine-learning approaches, i.e., batch-learning-based approaches, which require a complete training set before being able to be used, or active-learning-based approaches, which consider the user as a human-in-the-loop that acts on request from the algorithm, with very limited freedom of action.

The implementation of the unobtrusive machine learning required by ICS uses an online-learning algorithm (Passive-Aggressive, [2]). This algorithm is not yet sufficient to produce a usable implementation, as the algorithm can work in an online fashion only if the vector space in which the text documents are projected does not change. This is not the case for ICS, as its document datasets can be modified. A simple case of this is labelling a dataset of tweets that is continuously fed by an active streaming query. ICS solves this issue by using Lightweight Random Indexing [2], which defines a fixed vector space that is independent of any feature extracted from text, and it is efficient in reducing the number of dimensions in the vector space while not losing information from the resulting vectorial representation of text. This efficiency is a key aspect in having fast update time for any classifier, both due to lower computational complexity and due to reduced I/O from the DB that stores all the classification models. Many other theoretical and technological "tricks" are used in ICS to provide its users with responsive and accurate automatic classification, e.g., feature hashing, variable mini-batch learning, mini-sample-based active learning. The paper that presents ICS in details [1], also reports on some experimental evaluation that compares the performance of the model against traditional batch-learning approaches, including evaluating the possibility of performing transfer learning, reusing an existing classifier.

ICS is distributed under the BSD license and can be easily installed using the pip Python package installer (ics-pkg). More information, including installation and usage video tutorials are available on the GitHub repository [L1].

Link:
[L1] https://github.com/aesuli/ics

References:

[1] A. Esuli, "ICS: Total Freedom in Manual Text Classification Supported by Unobtrusive Machine Learning", in IEEE Access, vol. 10, pp. 64741-64760, 2022. https://doi.org/10.1109/access.2022.3184009
[2] K. Crammer, et al., "Online Passive-Aggressive Algorithms", in Journal of Machine Learning Research, 7, 551-585, 2006.
[3] A. Moreo, A. Esuli and F. Sebastiani, "Lightweight random indexing for polylingual text classification", in Journal of Artificial Intelligence Research 57, pp. 151-185, 2016. https://doi.org/10.1613/jair.5194

Please contact:

Andrea Esuli, ISTI-CNR, Italy
andrea.esuli@isti.cnr.it

How Quickly do Trees Grow?

by Refiz Duro (Austrian Institute of Technology), Hanns Kirchmeir (E.C.O. Institut für Ökologie Jungmeier), Anita Zolles (Bundesforschungs- und Ausbildungszentrum für Wald, Naturgefahren und Landschaft) and Günther Bronner (Umweltdata)

Changing climatic circumstances have a significant impact on forests: besides higher temperatures, more intense and frequent storms and drought spells affect forest growth. To see what the future is bringing, and to be able to deal with forest conservation and management, it is necessary to answer the question "how quickly do trees grow in an environment of climate change?" We take on a challenge to answer this question by integrating state-of-the-art data collection and AI-based methods.

Forests are the largest terrestrial sinks for carbon and some of the richest biological areas on Earth. Climate change is, however, affecting forests through increasing temperatures, changing precipitation patterns and the growing number of biotic and abiotic disturbances [1]. Saving forest ecosystems is thus one of the key measures to mitigate climate change and save biodiversity. To maintain and improve forest biodiversity and forest resilience to climate change, updated forest policies and forest management strategies are being developed and implemented in adaptive forest management. They all require up-to-date data of the forest conditions including the vitality and health of trees, as well as the ongoing tree growth (i.e., carbon sequestration).

Initially, tree and forest growth have been assessed mainly for economic reasons in order to build forest yield tables as simple "growth models" and a basis for improved forest management and taxation. Only within the past 50 years, improved tree and forest growth models have become available (e.g., [2]). They do not, however, allow for consideration of instantaneous changes in growth due to climatic extremes (e.g., drought, heat) and the changes of phenological patterns (i.e., length vegetation season).

Nowadays, new and more accurate measurement equipment has become available and new tree and forest characteristics have become the focus of forest and environmental science. The latest measurement equipment allows, for example, to assess intraday variation of tree radial growth (automatic dendrometers), to assess the water status of trees during a day (sap flow). This is used to characterise the tree and crown architecture including their habitat and microhabitat functions with terrestrial laser scanning in highest precision, or to assess tree health and vitality with airborne and satellite imaging / laser scanning. These new measurement techniques provide a huge amount of quantitative forest data, but their correct analysis and interpretation strongly requires advanced analytics to fully utilise the obtained data, achievable nowadays with the advanced probabilistic and machine learning approaches (e.g., neural networks).

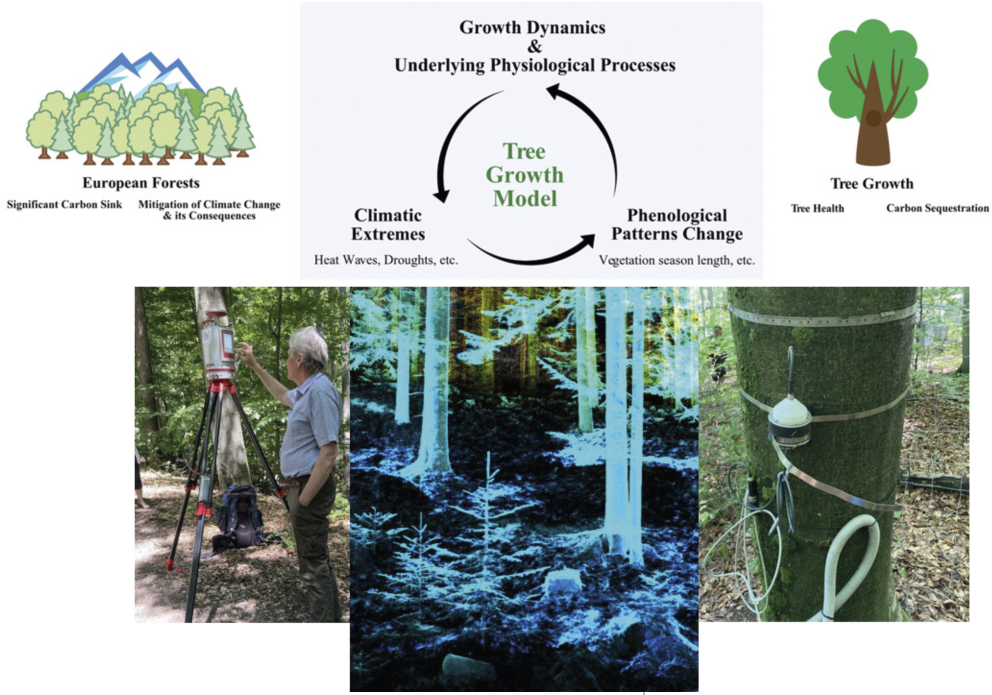

Figure 1: Top row: the basic concept of the AI4Trees project targeting development of a single tree growth model integrating climatic extremes and phenological patterns change. Bottom row: Data collection through TLS (left) providing a 3D image of the surroundings (middle), and dendrometer measuring tree growth (right).

In this sense, the AI4Tress project [L1] funded by the Austrian Federal Ministry for Climate Protection, Environment, Energy, Mobility, Innovation and Technology and the Austrian Research Promotion Agency's (FFG) "AI for Green" initiative, has taken on the challenge of developing a predictive AI-based climate-sensitive tree-growth model, exploiting the availability of the data, and applying the optimal machine learning strategies. Multidisciplinary experts from forestry to data science come from 6 Austrian project partners: AIT Austrian Institute of Technology GmbH (coordinator, AIT), Bundesforschungs- und Ausbildungszentrum für Wald, Naturgefahren und Landschaft (BWF), Umweltdata GmbH (UWD), Know-Center GmbH Research Center for Data-Driven Business & Big Data Analytics (KC), GeoVille Informationssysteme und Datenverarbeitung GmbH (GeoVille) and E.C.O. Institut für Ökologie Jungmeier GmbH (E.C.O.).

The project's first phase (started April 2022) has been focused on collecting high-quality, in-field data directly from forest sites selected from the Europe-wide forest monitoring program (ICP-Forests), which has been providing high-quality data on the vitality and adaptability of trees, nutrient cycles, critical load rates, and water balance [3]. Leveraging these high-quality datasets, statements can be made about climate change, air pollution, biodiversity, and the overall state of single tree stands. The project partner BFW maintains six ICP-Forests core measurement sites, where for each site, there are 10 trees equipped with automatic dendrometers, delivering tree diameter data at every 60 (or even 15) minute intervals. Dendrometers are basically highly sensitive measuring instruments that detect minuscule changes in a tree diameter (resolution of 1–5 micrometres), making them excellent to assess intraday variation of tree radial growth, which could be due to instantaneous changes in growth due to environmental extremes over short time spans (e.g., heat or cold waves). Furthermore, terrestrial laser scanning (TLS) technique delivers single tree point-clouds not only allowing extraction of traditional tree features like diameters at different heights, tree height and crown dimensions, but also providing the possibility of statistical approaches for calculation of various metrics, e.g., point-cloud percentiles along the tree height and competition between neighbouring trees. Repeating TLS measurements on a regular basis allows to estimate tree growth within a short period (1–3 years) in point-cloud data, which is still a challenge. Some measurement runs have already been performed and the collected TLS data has been processed, and together with the dendrometer data give an excellent basis for the forthcoming activities targeting development of an explainable and predictive AI-based climate-sensitive single tree growth model.

The successful outcome of the project has the potential to empower response to minimise potentially harmful climate change impacts on modern societies in line with the UN Sustainable Development Goals.

Links:
[L1] https://ai4trees-project.at/

References:
[1] R. Seidl, M.-J. Schelhaas, W. Rammer, and P. J. Verkerk, "Increasing forest disturbances in Europe and their impact on carbon storage," in Nat. Clim. Change, vol. 4, no. 9, pp. 806–810, Sep. 2014. https://doi.org/10.1038/nclimate2318
[2] H. Hasenauer, "Concepts Within Tree Growth Modeling," in Sustainable Forest Management: Growth Models for Europe, H. Hasenauer, Ed. Berlin, Heidelberg: Springer, 2006, pp. 3–17. https://doi.org/10.1007/3-540-31304-4_1
[3] M. Lorenz et al., "Forest Condition in Europe," Federal Research Centre for Forestry and Forest Products (BFH), United Nations Economic Commission for Europe Convention on Long-Range Transboundary Air Pollution, 2004.

Please contact:
Refiz Duro
AIT Austrian Institute of Technology GmbH, Austria
refiz.duro@ait.ac.at

Food Waste Reduction in Healthcare – Challenges in Integrating Usage Data with Scorings

by Johann Steszgal (Steszgal Informationstechnologie GmbH), Peter Kieseberg (St. Pölten UAS) and, Andreas Holzinger (University of Natural Resources and Life Sciences Vienna)

Reduction of food waste is an important target for reducing the human footprint and achieving better utilisation of natural resources. The healthcare sector especially offers a lot of potential for more sustainable handling of food. In this article we present major challenges for food waste reduction in real-world healthcare environments, as well as a solution approach.

The drastic reduction of food waste at the retail and consumer level has been identified by the United Nations as a key target (Target 12.3) to achieve the Global Sustainability Goals (SGDs) and is also closely related to SDG 2 ("Zero Hunger") [1]. A key challenge is the structured recording of food waste, its impact on the environment, as well as the adaptation of user behaviour to a more sustainable approach to food consumption.

Especially in mass catering, and in particular in the care sector, two special situations come together: on the one hand, food is distributed in a highly standardised way, especially with regards to portion sizes and the composition of the individual menus, which creates a great potential for waste; on the other hand, the actual consumption can be measured quite precisely in the case of inpatient admission. Therefore, this area in particular would be predestined for a pioneering role in food waste reduction, as first research results demonstrate: for example, a study of Swedish hospitals shows a very high level of waste, but this picture is very heterogeneous when looking at the details, as the actual impact in terms of sustainability was not measured, i.e., they made no distinction whether, e.g., meat or vegetables were left over. Furthermore, most studies do not take non-served food, e.g., due to patients being discharged after early rounds, into consideration. It should be noted that no intersection with sustainability scores was made in any of the studies known to us.

In the “Necta against Foodwaste” project [L2] we therefore focus on the utilisation of healthcare records related to food consumption for in-house patients and calculate the sustainability costs for wasted food as a basis for reducing the ecological footprint of catering in the health sector. This requires harnessing and integrating different data masters, especially intelligent unification of item master data to allow integration with external information like nutrition protocols that map the actual use of food in the hospital and care sector, and sustainability scores that allow ratings of individual foods in terms of sustainability, especially the one provided by Eaternity [2]. This leads to the following research challenges:

The “Necta against Foodwaste” project calculates the sustainability costs for wasted food as a basis for reducing the ecological footprint of catering in the health sector. Image from Pixabay.

- Applicability in real-world environments: In order to make the data usable and derive measures to reduce waste, it must be made analysable and evaluated. As real-world data is typically not flawless, this means that standard techniques for data cleansing, as well as their impact on the final result, need to be analysed in order to guarantee a certain accuracy [3]. In addition, the impact of privacy measures needs to be taken into account.
- Data normalisation and integration: Currently, there exists no standard for data collections regarding food consumption of residents in the health sector, more precisely, often the data is collected in free text form. In order to make this information usable, the data has to be transferred into a structured form and requires normalisation between the different environments.
- Establishing comparability: A major problem for meal scoring is the heterogeneity of the ingredient landscape and its comparability. In particular, establishing item master data comparability is a common problem, as equivalence is very difficult to define in general terms and is based on the application purpose, number of servings (package sizes), and many other side parameters that must be dynamically selected and integrated accordingly.
- Scoring: Just measuring the amount of wasted food is not good enough to estimate the ecological impact of food waste: wasting 1 kg of potatoes, while not good, does far less harm than wasting 1 kg of beef. Furthermore, for more fine-grained analysis, sustainability needs to be calculated with respect to different parameters, e.g., water consumption. Introduction of the Eaternity score is especially interesting here, as it provides a wide variety of different parameters and is still actively developed.

- Researching measures to reduce food waste: Depending on the operational scenario and organisation, different measures need to be applied – starting with the detour of unused food, to the adjustment of portion sizes, to the redesign of recipes with regard to unpopular parts.
- Sustainable recipes: By integrating sustainability scores, not only ingredients, but also the sustainability of recipes and portions can be calculated through the knowledge of recipes stored in necta. Thus, these can be made more sustainable, e.g., by replacing or removing problematic ingredients or searching for alternatives.

However, all data related to individual patients has high privacy and data security requirements, especially since anonymisation in dynamic data sets is extremely problematic from a data quality point of view. Thus, while anonymisation is applied where possible, intelligent and flexible pseudonymisation of data that cannot be anonymised due to subsequent evaluation options needs to be designed, including a high degree of flexibility in the implementation of access controls in de-pseudonymisation in order to be able to implement concepts such as the 4-eyes principle and the like.

Links:
[L1] https://www.necta-group.com/
[L2] https://kwz.me/hwC

References:
[1] A. Holzinger, et al., "Digital Transformation for Sustainable Development Goals (SDGs) – A Security, Safety and Privacy Perspective on AI", in Proc. of CD-MAKE 2021 (pp. 1-20), LNCS vol 12844, Springer, 2021. https://doi.org/10.1007/978-3-030-84060-0_1
[2] L. Eymann, M. Stucki and E. Hirsiger, "Nose to tail: how to allocate the environmental burden of livestock production systems to different meat cuts?", in 7th Int. Conf. on Life Cycle Management, 2015, 2015. https://doi.org/10.21256/zhaw-2739
[3] K. Stöger, et al., "Legal aspects of data cleansing in medical AI", Computer Law & Security Review, vol. 42, p.105587, 2021. https://doi.org/10.1016/j.clsr.2021.105587

Please contact:
Peter Kieseberg, St. Pölten UAS, Austria
peter.kieseberg@fhstp.ac.at

GLACIATION

Sponsored contribution

Launch of the New Horizon Europe Project GLACIATION

Fourteen leading European organisations join forces to provide a framework for green and privacy-preserving data operations in an edge-to-cloud architecture.

Existing network architectures offer a variety of options for storage and analysis of data. Currently, however, network architecture follows a cloud-based logic thereby hampering the exploitation of distributed architectures' potential. Therefore, organisations that operate on an edge-to-cloud architecture have to deal with suboptimal results with regards to energy efficiency, computation and resources allocation for data storage and management performance. This problem is most relevant in large organisations with many distributed entities, such as public administrations, or organisations that make an intensive use of sensors, such as the ones operating in Industry 4.0.

By contrast, edge applications are most suited to the increasing volume and heterogeneity of users, data, and devices. In that regard, 14 leading organisations join the Horizon Europe project GLACIATION (Green responsibLe privACy preservIng dAta operaTIONs) to develop a digital solution that harnesses the potential of the edge and combines it with the cloud. The project, coordinated by the Italian Ministry of Economics and Finance, kicked off its activities in October 2022, with an energising online meeting, where all partners took the floor to present their ambitious plans to make the project a great success. The GLACIATION consortium partners bring to the project expertise in several fields (universities and research centres, businesses, technology leaders, public organisations) and are among the leaders in the development of both cloud and edge solutions.

GLACIATION aims to improve the efficiency and the effectiveness of trustworthy digital technologies in addressing the requirements of citizens, companies and public organisations concerning privacy and commercial and administrative confidentiality, as well as regarding responsible, fair and environmentally friendly (e.g. in terms of energy/carbon/material footprint) data operations in data spaces, across the data life cycle.

Specifically, GLACIATION will develop a data management and analytic framework that prioritises energy-efficiency and privacy along an edge-core-cloud architecture, thereby advancing the state of the art in the fields of AI-based technologies applied to dynamic and predictive placement of data, as well as to dynamic merging and movement of distributed knowledge graphs. The project will demonstrate its solution in three real-life scenarios including a National government-wide platform, an Industry 4.0 setting, and a cross-company optimisation of processes. In this way, the project will contribute to the "European Green Deal" and to a Europe fit for the digital age" by developing a digital solution to improve the energy efficiency of the European data spaces.

ERCIM/W3C mainly contributes to the development of a metadata model for data centric architectures and leads the standardisation activities of the project.

Please contact:
Rigo Wenning, ERCIM/W3C
rigo@w3.org

SCHLOSS DAGSTUHL
Leibniz-Zentrum für Informatik

Call for Proposals

Dagstuhl Seminars and Perspectives Workshops

Schloss Dagstuhl – Leibniz-Zentrum für Informatik is accepting proposals for scientific seminars/workshops in all areas of computer science, in particular also in connection with other fields.

If accepted, the event will be hosted in the seclusion of Dagstuhl's well known, own, dedicated facilities in Wadern on the western fringe of Germany. Moreover, the Dagstuhl office will assume most of the organisational/ administrative work, and the Dagstuhl scientific staff will support the organizers in preparing, running, and documenting the event. Thanks to subsidies the costs are very low for participants.

Dagstuhl events are typically proposed by a group of three to four outstanding researchers of different affiliations. This organizer team should represent a range of research communities and reflect Dagstuhl's international orientation. More information, in particular details about event form and setup, as well as the proposal form and the proposing process, can be found on

https://www.dagstuhl.de/dsproposal

Schloss Dagstuhl – Leibniz-Zentrum für Informatik is funded by the German federal and state government. It pursues a mission of furthering world class research in computer science by facilitating communication and interaction between researchers.

Important Dates

- *Next submission period:* April 1 to April 15, 2023
- *Seminar dates:* Between March 2024 and February 2025 (tentative).

ERCIM "Alain Bensoussan" Fellowship Programme

The ERCIM postdoctoral Fellowship Programme has been established as one of the premier activities of ERCIM. The programme is open to young researchers from all over the world. It focuses on a broad range of fields in Computer Science and Applied Mathematics.

The fellowship scheme also helps young scientists to improve their knowledge of European research structures and networks and to gain more insight into the working conditions of leading European research institutions. The fellowships are of 12 months duration (with a possible extension), spent in one of the ERCIM member institutes. Fellows can apply for second year in a different institute.

Where are the fellows hosted?

Only ERCIM members can host fellows. When an ERCIM member is a consortium the hosting institute might be any of the consortium's members. When an ERCIM Member is a funding organisation, the hosting institute might be any of their affiliates. Fellowships are proposed according to the needs of the member institutes and the available funding.

Conditions

Candidates must:
- have obtained a PhD degree during the last eight years (prior to the application year deadline) or be in the last year of the thesis work with an outstanding academic record. Before starting the grant, a proof of the PhD degree will be requested;
- be fluent in English.

The fellows are appointed either by a stipend (an agreement for a research training programme) or a working contract. The type of contract and the monthly allowance/salary depends on the hosting institute.

ERCIM encourages both researchers from academic institutions and scientists working in industry to apply.

Equal Opportunities

ERCIM is committed to ensuring equal opportunities and promoting diversity. People seeking fellowship within the ERCIM consortium are not discriminated against because race, color, religion, gender, national origin, age, marital status or disability.

Application deadlines

Deadlines for applications are currently 30 April and 30 September each year.

Since its inception in 1991, over 790 fellows have passed through the programme. In 2022, 40 young scientists commenced an ERCIM PhD fellowship and 69 fellows have been hosted during the year. The Fellowship Programme is named in honour of Alain Bensoussan, former president of Inria, one of the three ERCIM founding institutes.

http://fellowship.ercim.eu

"ERCIM fellowship is one of the best things that have happened to me in my research career. I have forged new scientific contacts and I am more confident now to handle independent research work.

Nimisha GHOSH
Former ERCIM Fellow

ERCIM – the European Research Consortium for Informatics and Mathematics is an organisation dedicated to the advancement of European research and development in information technology and applied mathematics. Its member institutions aim to foster collaborative work within the European research community and to increase co-operation with European industry.

ERCIM is the European Host of the World Wide Web Consortium.

Consiglio Nazionale delle Ricerche
Area della Ricerca CNR di Pisa
Via G. Moruzzi 1, 56124 Pisa, Italy
www.iit.cnr.it

Norwegian University of Science and Technology
Faculty of Information Technology, Mathematics and Electrical Engineering, N 7491 Trondheim, Norway
http://www.ntnu.no/

Centrum Wiskunde & Informatica
Science Park 123,
NL-1098 XG Amsterdam, The Netherlands
www.cwi.nl

RISE SICS
Box 1263,
SE-164 29 Kista, Sweden
http://www.sics.se/

Fonds National de la Recherche
6, rue Antoine de Saint-Exupéry, B.P. 1777
L-1017 Luxembourg-Kirchberg
www.fnr.lu

SBA Research gGmbH
Floragasse 7, 1040 Wien, Austria
www.sba-research.org/

Foundation for Research and Technology – Hellas
Institute of Computer Science
P.O. Box 1385, GR-71110 Heraklion, Crete, Greece
www.ics.forth.gr

SIMULA
PO Box 134
1325 Lysaker, Norway
www.simula.no

Fraunhofer ICT Group
Anna-Louisa-Karsch-Str. 2
10178 Berlin, Germany
www.iuk.fraunhofer.de

Eötvös Loránd Research Network
Számítástechnikai és Automatizálási Kutató Intézet
P.O. Box 63, H-1518 Budapest, Hungary
www.sztaki.hu/

INESC
c/o INESC Porto, Campus da FEUP,
Rua Dr. Roberto Frias, nº 378,
4200-465 Porto, Portugal
www.inesc.pt

University of Cyprus
P.O. Box 20537
1678 Nicosia, Cyprus
www.cs.ucy.ac.cy/

Institut National de Recherche en Informatique
et en Automatique
B.P. 105, F-78153 Le Chesnay, France
www.inria.fr

Universty of Warsaw
Faculty of Mathematics, Informatics and Mechanics
Banacha 2, 02-097 Warsaw, Poland
www.mimuw.edu.pl/

I.S.I. – Industrial Systems Institute
Patras Science Park building
Platani, Patras, Greece, GR-26504
www.isi.gr

VTT Technical Research Centre of Finland Ltd
PO Box 1000
FIN-02044 VTT, Finland
www.vttresearch.com